

Complexity Changes in Neuroimaging Signals Point to Mechanisms of Neuropsychiatric Disease

David Sutherland Blair

TESI DOCTORAL UPF / 2022

Thesis supervisor

Prof. Dr. Gustavo Deco

DEPARTMENT OF INFORMATION AND COMMUNICATION
TECHNOLOGIES



*I couldn't do it. I couldn't reduce it to the freshman level.
That means we don't really understand it.*

Richard P. Feynman
(Goodstein 1989)

In Appreciation Of

If raising a child takes a village, writing a thesis takes a team. No modern thesis can or should be written alone.

With that in mind, I must first thank Prof. Deco for taking a chance on a young, inexperienced M.Sc. student who wanted to start modeling brain activity. Without his initial and continued support over the past five years, this thesis simply would not have been possible. At no point did he question my ability nor stop me from, to paraphrase Dr. Feynman, studying what interested me in the most undisciplined, irreverent, and original manner possible.

My co-authors, Dr. Joana Cabral and Dr. Carles Soriano-Mas, deserve much credit for the success of this thesis. Joana's explanation of the LEiDA framework and assistance with the ever-troublesome illustration thereof proved essential to the development and publication of the current work. Likewise, without Dr. Soriano-Mas, any attempt to interpret our findings for neurobiological meaning would have been entirely useless. I hope to continue working with them in future endeavors.

Thanks is due to my colleagues, particularly Gorka Zamora-López, Ane López-González and Manel Vida-Vidal, who offered much-needed advice at difficult points during this work. It is likely that without them, I would not have finished this project, nor would I have learned how to escape the traps into which I stumbled. For that I owe them gratitude.

More recent members of the CNS, including Elvira del Agua Banyeres, Eider, Lou Zonca, Sebastian Geli, Wiep Stikvoort, and Elvira García Guzman, deserve mentions as well. While the pandemic scattered many of us, those who returned or joined the group afterwards have created a vibrant and active group of researchers and friends. It is difficult to overstate the importance of such a community to a doctoral project, both in providing material aid through advice and code and in the often overlooked but utterly crucial role of maintaining morale. Even if unused, simply knowing where and when a helping hand or listening ear is available does much to keep one moving.

Outside the CNS, I should extend thanks to several friends who helped smooth the undeniably difficult transition into Barcelona. María Palazzi and Oscar Fajardo both deserve places on this list, as does Magdalena Matyjek. For the same reason, I must thank those from Zürich who remained in contact during these past five years: Marie Carré, Maud Galletti, Christopher Tcherenkov, Nina Stumpf, Harun Mustafa, Katherina Wellstein, and Rebecca Westphal; and my new old friends Travis Smith, Tommy La Voy, Allan La-Grenade Finch, Ivie Mariah, and Kirsie Mittini, all of whom helped maintain morale through dependable support.

Finally, of course, I must thank my parents and family above all. I can only imagine the thoughts which passed through their minds upon learning that their eldest son planned to complete his entire graduate education in countries foreign in language, culture, and continent. That they never attempted to dissuade me speaks much of their faith in their eldest. That they maintained it even at times when he did not says more. I could not ask for a more loving or supportive home. Thank you, from the bottom of my heart.

Abstract

The past decade has seen several attempts to employ the entropy of neuroimaging signals as a potential biomarker for cognitive decline or traumatic brain injury (C. Y. Liu et al. 2013; Adhikari et al. 2017; Li et al. 2018). Not all these studies properly account for the distributed nature of cognition, however, which raises the possibility of erroneous estimates of global entropy. This thesis proposes a novel means of estimating the complexity of fMRI signals and demonstrates its efficacy in detecting the effects of psychiatric disease on neuroimaging signals. The method determines the minimum number of orthogonal dimensions necessary to capture nonrandom signal dynamics, then projects the dynamic functional connectivity signal into the resultant low-dimensional space. In this space, the dynamic functional connectivity signal's entropy may be estimated along each dimension independently and summed to find the total entropy per subject, thus avoiding the need to estimate interregional effects. Tests on two independently collected datasets indicate that this pipeline can distinguish between healthy controls and psychiatric patients, and that a Hopf bifurcation-based effective connectivity model is able to recover meaningful differences between control and patient groups when trained in this space.

Keywords: whole-brain model; obsessive-compulsive disorder; schizophrenia; bipolar disorder (type I); attention-deficit hyperactive disorder; entropy; Shannon entropy; independent component analysis; eigendecomposition; Hopf bifurcation; effective connectivity; network-based statistic

Resumen

La última década ha sido testigo de varios intentos de emplear la entropía de las señales de neuroimagen como un biomarcador potencial para el deterioro cognitivo o la lesión cerebral traumática (C. Y. Liu et al. 2013; Adhikari et al. 2017; Li et al. 2018). Sin embargo, no todos estos estudios explican adecuadamente la naturaleza distribuida de la cognición, lo que plantea la posibilidad de estimaciones erróneas de la entropía global. Esta tesis propone un medio novedoso para estimar la complejidad de las señales de fMRI y demuestra su eficacia en la detección de los efectos de la enfermedad psiquiátrica en las señales de neuroimagen. El método determina el número mínimo de dimensiones ortogonales necesarias para capturar la dinámica de la señal no aleatoria, luego proyecta la señal de conectividad funcional dinámica en el espacio resultante de baja dimensión. En este espacio, la entropía de la señal de conectividad funcional dinámica puede estimarse a lo largo de cada dimensión de forma independiente y sumarse para encontrar la entropía total por sujeto, evitando así la necesidad de estimar los efectos interregionales. Las pruebas en dos conjuntos de datos recopilados de forma independiente indican que esta tubería puede distinguir entre controles sanos y pacientes psiquiátricos, y que un modelo de conectividad efectiva basado en la bifurcación de Hopf puede recuperar diferencias significativas entre los grupos de control y pacientes cuando se entrena en este espacio.

Palabras clave: modelo de cerebro completo; desorden obsesivo compulsivo; esquizofrenia; trastorno bipolar (tipo I); trastorno por déficit de atención con hiperactividad; entropía; entropía de Shannon; análisis de componentes independientes; descomposición propia; bifurcación de Hopf; conectividad efectiva; estadística basada en la red

Preface

To say that the field of neuroimaging has exploded in the past decades would be an understatement. Between methodological developments, technological improvements, and increased computing power, the field has ballooned at an almost inconceivable rate. So much information is now available that the challenge of the era seems to have shifted from locating relevant research to locating *which* research is relevant.

One of the newest and fastest-growing sections of neuroimaging has been in the field of MRI-based dynamic functional connectivity (dFC). Although the first publication of this methodology occurred scarcely a decade ago, it has quickly become the basis of a vast quantity of research into both healthy and disordered cognition. Of note has been the evolving consensus that the human brain's resting state activity is not random, but rather explores a highly organized phase space dominated by a finite number of recurrent connectivity patterns, or *networks*. These *resting-state networks* (RSNs) consist of both task-positive and task-negative networks, with task-positive networks becoming more active while engaged in a cognitive or behavioral task and task-negative networks increasing their activity in the awake, alert, but inactive rest state. The most famous of these RSNs, the *default-mode network*, falls firmly into the category of task-negative, as its activity appears to be suppressed during active tasks. The majority of its companion RSNs are task-positive networks, many of which are visually indistinguishable from networks known to support active thought and cognition. Examples include visual, sensorimotor, and auditory networks.

The discovery of these networks has prompted an immense amount of effort to be poured into cataloguing them. The sheer number of methodologies applied to this problem is difficult to account. To our knowledge, dynamic clustering algorithms, various types of Markov models, independent component analysis, principal component analysis, and a wide variety of other methods have been used to isolate these connectivity patterns, or *substates*, from dFC data and to track their activity over time. Many studies have been published on this subject, most of which attempt to describe the relevant dynamics in terms of transition probabilities or energy landscapes—concepts borrowed in no small part from dynamic

systems theory. Yet while these studies have provided valuable insights into transition dynamics, they can be difficult to summarize. Some method to directly compare the complexity or predictability of dFC dynamics between different subjects is needed. This thesis will attempt to provide such a method.

The role of effective connectivity must also be addressed. While not a part of the neuroimaging field *per se*, effective connectivity is undeniably connected to it. It is, after all, fundamentally an attempt to predict or explain the signals detected in neuroimaging data. But this raises the question, what *is* effective connectivity? Simply put it is an attempt to predict how various regions of the brain influence one another. In this case, the question of *how* is restricted to the questions, *which regions influence one another, how strong is this influence, and in which direction does this influence flow?* Note that this is distinct from functional connectivity, which cannot determine directionality and offers no explanation as to how such influence comes to be.

The concept of effective connectivity is not new; the question of which regions influence which is as old as the idea that long-range influence is possible at all. What is new is the wide availability of computational power necessary to make detailed estimates of such influence. That development, and the enormous quantity of research going into computational modeling and artificial intelligence, has permitted impressive advancements in the field of effective connectivity over the past decade. It remains, however, a developing art, as the question of how regions influence each other remains an ill-posed one. While useful information can be gleaned from such models, their results should be taken with a healthy dose of skepticism.

Despite this, this thesis attempts to determine interregional influence based on dynamic functional connectivity. Specifically, it searches for the connectivity which best predicts the predictability of the functional connectivity dynamics. If that sounds confusing, well, hopefully by the end of this you'll have an idea of how it works. Summarizing it is, unfortunately, not an obvious problem.

But why investigate the level of order in functional connectivity dynamics at all? Several reasons present themselves, but the most

pressing is the question of psychiatric disease. For all the advances in our knowledge of how the brain functions, advances in psychiatric knowledge and practice remain frustratingly slow. Two of the many unanswered questions which remain is the degree to which psychiatric disorders relate with one another, and whether psychiatric symptoms are fundamentally a problem of impaired or overactive brain dynamics. As it happens, these are questions which the methods developed in this thesis can begin to answer, as will be demonstrated in studies on obsessive-compulsive disorder (OCD), bipolar disorder (type I), schizophrenia patients, and attention-deficit hyperactivity disorder (ADHD). In the process, marked similarities between these disorders will in fact be observed, and potential effective connectivity misalignments identified. It is the author's hope that this will provide a step, however small, towards the effective treatment of these disorders.

Now. Let us begin.

Table of Contents

<i>In Appreciation Of</i>	v
<i>Abstract</i>	viii
<i>Resumen</i>	ix
<i>Preface</i>	xi
<i>List of figures</i>	xviii
<i>List of tables</i>	xix
1 CONNECTIVITY	2
1.1 Definitions	3
1.1.1 Connectivity in the Brain	3
1.1.2 The Connectivity Matrix	4
1.1.3 Node Strength	4
1.2 Structural Connectivity	5
1.2.1 Definitions	5
1.2.2 Tracing Tracts	5
1.2.3 Atlases	6
1.3 Functional Connectivity	7
1.3.1 Definitions	7
1.3.2 Communication Through Coherence	8
1.4 Effective Connectivity	11
1.4.1 Definitions	11
1.4.2 How is a neuron like a nitrogen atom?	11
1.5 Summary	18
2 The Resting State	20
2.1 Definitions: Brain Activity During Rest	20
2.2 The Resting State: A History	20
2.2.1 Early Days: Discovery and Validation	20
2.2.2 Dynamism in Data	23
2.2.3 Summary	26
2.3 Current Hypotheses	26
2.3.1 The Resting State as a Multistable System	26
2.3.2 The Resting State as a Metastable System	29
2.3.3 Metastable Multistability?	30
2.4 Summary	31

3	<i>Statistical Neurophysics: or, Measuring Entropy</i>	32
3.1	What is Entropy?	33
3.1.1	The Thermodynamic Entropy	33
3.1.2	Entropy in Statistical Mechanics	34
3.1.3	The Entropy in Information Theory	37
3.1.4	Tied Together	40
3.2	Why to Measure It?	40
4	<i>Datasets</i>	42
4.1	Motivation	42
4.2	Datasets	43
4.2.1	Obsessive-Compulsive Disorder	43
4.2.2	Cross-Disorder	44
4.3	Summary	46
5	<i>Extracting Entropy</i>	48
5.1	Reasoning	48
5.2	Pipeline	48
5.2.1	Functional Connectivity	48
5.2.2	Eigendecomposition	50
5.2.3	Component Detection	53
5.2.4	Component Extraction	53
5.2.5	Entropy Estimation	60
5.2.6	Comparisons	61
5.3	Results	61
5.3.1	Functional Connectivity	61
5.3.2	Functional Dimensions	62
5.3.3	Joint Entropy	63
5.3.4	Dimension-Specific Entropy	66
5.4	Conclusions	73
6	<i>Predicting Connectivity</i>	75
6.1	Motivations	75
6.2	Determining the Method	75
6.3	Pipeline	77
6.3.1	Brain Network Model	77
6.3.2	Optimization	78
6.3.3	Node Strength	79
6.3.4	Network-Based Statistic	79

6.4	Results	80
6.4.1	Goodness-of-Fit	80
6.4.2	Node Strength Analysis	84
6.4.3	Network-Based Statistic	89
6.5	Conclusions	97
7	<i>Discussion & Future Steps</i>	100
7.1	General Review	100
7.2	Review: Obsessive-Compulsive Disorder	101
7.2.1	The Phenomenology of Obsessive-Compulsive Disorder	101
7.2.2	Effective Connectivity in Obsessive-Compulsive Disorder	102
7.3	Review: Schizophrenia, Bipolar Disorder (Type I), and ADHD	103
7.3.1	Phenomenology	103
7.3.2	Connectivity	103
7.4	General Comments	104
7.5	Limitations	104
7.6	Future Steps	105
7.6.1	Biomarker Research	105
7.6.2	Consolidation of Resting-State Networks	107
7.6.3	Tracing Activity Propagation	108
7.7	Final Thoughts	108

List of figures

Figure 1: Behold the mighty duck-rabbit. _____	28
Figure 2: computing the phase of the BOLD signal. _____	50
Figure 3: Eigendecomposition. _____	52
Figure 4: the cocktail party problem _____	55
Figure 5: Kurtosis _____	58
Figure 6: joint entropy of controls vs. OCD patients _____	64
Figure 7: joint entropy of controls, bipolar, schizophrenia, and ADHD patients. _____	65
Figure 8: significant dimensions of controls and OCD patients. _____	67
Figure 9: significant dimensions of controls, bipolar (type I), schizophrenia, and ADHD patients. _____	69
Figure 10: three components display significant alterations between control and patient entropies (A). _____	70
Figure 11 distances between empirical and simulated entropies (OCD) _____	82
Figure 12 distances between empirical and simulated entropies (schizophrenia, bipolar type I, ADHD). _____	84
Figure 13 in- and out-strength analysis of controls vs. patient groups _____	86
Figure 14 in- and out-strength analysis of patient groups vs. patient groups. _____	88
Figure 15 NBS results for controls vs. obsessive-compulsive patients. _____	91
Figure 16 NBS results for controls vs. schizophrenia patients. _____	92
Figure 17 NBS results for controls vs. bipolar type I patients. _____	93
Figure 18 NBS results for controls vs. ADHD patients. _____	94
Figure 19 NBS results for bipolar type I vs. ADHD patients. _____	95
Figure 20 NBS results for bipolar type I vs. schizophrenia patients.. _____	96
Figure 21 NBS results for schizophrenia vs. ADHD patients. _____	97

List of tables

<i>Table 1</i> <i>p-values for pairwise joint entropy comparisons.</i>	65
<i>Table 2</i> <i>displays the regions of the first dimension with $z > 1.3$.</i>	68
<i>Table 4</i> <i>Significant LEICA Dimensions</i>	70
<i>Table 5</i> <i>LEICA Dimensional Entropies</i>	73
<i>Table 6</i> <i>Number of Significant Regions per Comparison</i>	89

1 CONNECTIVITY

Interconnectedness defines the nervous system, not just in humans, but in all animals known to possess one. The processing power of an individual neuron is not adequate for even the most basic cognition of a multicellular organism, and a pool of many isolated neurons offers no meaningful improvement. Only by forming organized networks does complex processing emerge. This pattern extends across spatial scales in the human brain, as neurons network into columns, columns into regions, and regions into the interregional networks thought to underly cognition in humans and similar vertebrate species. From cell to cerebrum, connectivity is key to complex behavior.

Given the almost self-evident importance of connectivity in supporting complex thought, it may surprise some to find that the study of large-scale connectivity in the human brain is scarcely thirty years old. Surely attempts to map connections in the brain started decades earlier. Unsurprisingly, they did; researchers have realized since first examining the nervous system that its various elements and parts must communicate with one another. Yet until the early 1990s, technology did not allow the mapping of neural systems larger than a few dozen cells. Indeed, by 1990, the only fully mapped neural connectome was that of *C. elegans*, consisting of 302 neurons—a full seven orders of magnitude smaller than a human brain.

In 1990, technology able to precisely locate regions of neural activation first became available in the form of the BOLD signal. Since its discovery, neuroscientists have raced to make up for lost time. A veritable explosion of research into brain connectivity has occurred in the past thirty years, an explosion which has birthed entire new fields of neuroscience and medicine. Perhaps the most famous of these fields is that of *connectomics*, the study of high-speed interregional communication between physically separated brain regions. While originally focused on the physical wiring which facilitates this communication, connectomics has now expanded into studying how frequently disparate regions communicate, and even in which direction such communication

travels. A brief overview of this field will prove necessary to understanding the methods and motivations of this thesis.

1.1 Definitions

1.1.1 Connectivity in the Brain

The past two decades have seen a dizzying array of innovations at every scale of neuroscience, from individual cells to the entire nervous system. One of the most consequential innovations has been the introduction of network analysis to neuroscience (Sporns 2014). Briefly put, network analysis is the study of networks; what may be considered a network, how one is formed, how it operates, what rules govern them and so on. As might be expected, this results in an enormously broad spectrum of research, covering everything from molecular interactions to global economics. Nonetheless, many networks have been found to follow consistent structural trends, hinting that even such diverse systems may be

At the scale of the human brain, network analysis is most often applied to three major types of connectivity. The first is structural connectivity, which generally refers to the anatomical connections which carry neural signals between regions of the brain. The second is functional connectivity, which describes the similarities (or lack thereof) of the time-resolved activity between different regions of the brain (Friston 2011). Finally, there is effective connectivity, which predicts how neural signals travel between regions in the brain. This requires a fundamentally different approach; whereas structural and functional connectivity primarily catalog phenomena without modeling its generative mechanisms, effective connectivity must employ such models and ensure that their predictions align with empirical data (Gilson et al. 2016). The challenges of model selection and fitting do not fall within the realm of network analysis *per se*, being more closely aligned with neuromodeling and machine learning respectively; nonetheless, they play a vital role and so must be included in any discussion of effective connectivity.

1.1.2 The Connectivity Matrix

Central to almost all procedures in network analysis is the *adjacency* or *connectivity matrix*, which provides a convenient and mathematically tractable means to record interregional connections (Newman 2014). Given a network with N nodes, its *adjacency matrix* is an $N \times N$ matrix A in which

- Element $A_{ij} = 1$ if and only if there exists a link pointing from node j to node i
- Element $A_{ij} = 0$ if there does not exist a link pointing from node j to node i

Note that the adjacency matrix records no information on the relative strength of these connections, only that these connections exist. Obviously, not all connections are equal; some regions are more tightly linked than others. Including this information produces the network's *connectivity matrix* C , in which

- Element $C_{ij} = w_{ij}$ if and only if there exists a link pointing from node j to node i , where w_{ij} is the weight (strength) of the link between nodes j and i .
- Element $C_{ij} = 0$ if there does not exist a link pointing from node j to node i .

1.1.3 Node Strength

A common question to arise when examining networks is the relative influence of an individual region. Network science has developed many tools to study this question, but this thesis will primarily employ the node strength. The *strength* of a node j in a weighted network is simply the sum of all connections coming into (*in-strength*) or out of (*out-strength*) that node (Barabási 2016):

$$s_j^{in} = \sum_{i=1}^N C_{ij}$$
$$s_j^{out} = \sum_{j=1}^N C_{ij}$$

1.2 Structural Connectivity

1.2.1 Definitions

Structural connectivity is, simply put, the network of physical fibers which transmit neural activity or information between different regions of the brain (Sporns 2010). This definition holds at all scales of brain architecture, from the white matter fiber bundles which connect regions to the individual synapses and dendrites which allow intercellular communication. The spatial scale of these connections reflects their temporal stability; at the synaptic, cellular, and even columnal levels, the mechanisms of neural plasticity ensure that connections change on an order of milliseconds, while interconnections visible to the naked eye develop over years. This thesis will remain at the whole-brain level, so its use of the term “structural connectivity” will refer exclusively to the network of white matter fibers which carry action potentials from one anatomically defined brain region to another. These long, heavily myelinated axons tend to aggregate in dense bundles to efficiently carry messages from one region to another, before once again branching in to separate smaller tendrils to connect with local processing units. Detailed maps of these fiber bundles have been produced over the past decade as interest in network neuroscience has increased and the brain’s capacity for distributed cognition have gained recognition, and many atlases of such connectivity are now available.

1.2.2 Tracing Tracts

How are such networks charted? This is a question with a variety of answers depending on the study in question. In animal studies, such as rats and macaques, physical tract-tracing techniques may be employed (Lanciego and Wouterlood 2020). These techniques may involve a straightforward dissection of the postmortem brain or may involve injecting a labeled tracer molecule into a living specimen and observing the diffusion pattern of this tracer. In these cases, the animal must be sacrificed to observe the pattern in the tissue, which renders such procedures ethically unsuitable for human studies.

To bypass this problem, human studies make use of the heavy myelination of the white-matter tracts of interest. A family of MRI

techniques collectively referred to as *diffusion-weighted imaging* allow researchers to track the average direction of the diffusion of water molecules in biological tissue (Basser et al. 2000). As biological tissue contains substantial numbers of water-impermeable structures, water molecules do not diffuse evenly throughout it; instead, they tend to follow the path of least resistance. Myelin is predominantly made up of fat, so it is not easily water-permeable; thus, water molecules tend to flow along a sheet of myelin rather than across it. Thus, when embedded in a bundle of myelinated fibers, such as white matter tracts, water molecules tend to travel along the lengths of the fibers rather than crossing them. This differential flow can be measured using a *diffusion tensor imaging* process (Basser, Mattiello, and LeBihan 1994; Basser et al. 2000), and from this the approximate direction and size of the fiber bundle estimated. The process remains imperfect; it cannot easily distinguish between tracts when fiber bundles cross, it cannot determine the direction of neural signals in the bundles it traces, and its limited resolution means it cannot trace individual fibers or small bundles with meaningful precision. Nonetheless, it has proven an exceptionally valuable tool for human neuroimaging.

1.2.3 Atlases

The two studies described in this thesis utilize the Automated Anatomical Labels (AAL) atlas (Tzourio-Mazoyer et al. 2002) and the Desikan- Killiany atlas (Desikan et al. 2006a), respectively. These atlases represent early and relatively coarse parcellations of the human cortical surface, with the AAL atlas containing 90 cortical regions (45 per hemisphere) and the Desikan-Killiany atlas 68 cortical regions (34 per hemisphere). Both atlases are well-established in the neuroimaging communities and include structural connectivity matrices as well as anatomical labels and spatial coordinates. For the purposes of this thesis, the anatomical labels and structural connectivity matrices provide all required information.

1.3 Functional Connectivity

1.3.1 Definitions

Functional connectivity, in the neuroimaging context, refers to the temporal coincidence of spatially separated neurophysiological events (Friston 1994). This is generally interpreted as the degree of statistical dependence between the time series of separate regions in the brain. Such statistical dependence may be estimated in a variety of ways, such as the Pearson correlation coefficient, the covariance, the mutual information, or, as shall be elaborated on momentarily, the phase coherence of neural activity oscillations. The Pearson correlation has been the favored method in the fMRI literature due to its conceptual and computational simplicity, although recent work has called its use into question due to its inability to capture nonlinear interactions (Zhang et al. 2018; Sayed Hussein Jomaa et al. 2019).

Why choose to define the functional connectivity in this manner? This definition rests on the intuition that events which consistently co-occur in time are likely to be related to one another. Simply put, the researcher chooses not to believe in coincidence. This decision has risks of its own, of course—coincidences may indeed occur, or observed coincidences be mediated by an unobserved factor—but it is an intuitive and generally accepted means of mapping links between parts of a network.

One can argue that the great strength of functional connectivity is its relative agnosticism to how coincident events occur. Whereas effective connectivity attempts to model how regions communicate and is thus susceptible to all the simplifications and assumptions that such models require, functional connectivity simply reports which regions coactivate and which do not. This minimizes experimental assumptions. On the other hand, this agnosticism means that functional connectivity reveals relatively little about how such communication occurs. This is no small detail, as the means for interregional communication in the human brain remains an unsolved problem. Such communication is almost certainly mediated by white matter fiber bundles—a presumption reinforced by the convergence of structural and functional connectivity

(Bettinardi et al. 2017)—but information on the “neural code”, so to speak, is absent.

1.3.2 Communication Through Coherence

Rhythmic oscillations across multiple frequency bands have proven a consistent theme for neuronal activity across spatial scales, from single neurons to whole-brain imaging techniques (Gyorgy and Andreas 2004; Engel, Fries, and Singer 2001). It has been proposed (Fries 2005) that these oscillations modulate interregional communication by coordinating periods of high and low excitability. Neural excitability is known to vary in time, largely as a function of the time since generating an action potential(s). This, coupled with the tendency for neural groups to predictably oscillate in activity level, implies that neural groups experience rhythmic variations in receptibility to novel inputs. Input signals which arrive out of sync with a neural group’s inherent oscillations are unlikely to activate the group, which implies that such signals will not be transmitted further. Signals which arrive in phase with that group’s oscillations, on the other hand, will almost certainly be passed on. Prof. Fries thus hypothesizes that oscillatory coherence modulates interregional communication, such that regions which communicate with one another must necessarily display coherent activity. This *communication through coherence* (CTC) hypothesis neatly solves a pressing question in neuroscience, namely how different parts of the brain can flexibly communicate on a fixed anatomical structure.

If one presumes the CTC hypothesis to be correct, this raises the question of how to detect which regions are communicating. Synchronization between brain areas appears the obvious solution, but transmission and activation delays may render perfect, zero-phase synchrony suboptimal. Regions must lock phases to communicate, with the phase difference dependent on the oscillatory frequency and transmission delay. It follows that, to measure functional connectivity, one must determine the pairwise phase coupling between all pairs of regions in the brain. But if this is the case, how to go about it?

1.3.2.1 The Analytic Signal

Before we continue, we should discuss the analytic signal and its foundation. Given a real-valued signal $x(t)$ with a Fourier transform $\mathcal{F}[x(t)]$, there exists a complex signal $x_a(t)$ such that $\mathcal{F}[x_a(t)] = \{\mathcal{F}[x(t)] | \mathcal{F}[x(t)] > 0\}$. Put another way, $x_a(t)$ has the same positive frequency components as $x(t)$, but its negative frequency components are not defined. This analytic representation can be assembled by means of the Hilbert transform:

$$x_a(t) = x(t) + iH[x(t)]$$

where $H[x(t)]$ is the Hilbert transform and i is the imaginary unit. The mathematically minded may observe that this relation resembles Euler's formula,

$$e^{i\varphi(t)} = \cos \varphi(t) + i \sin \varphi(t)$$

where *sin* is the sinusoid function and *cos* is the cosine function. This similarity is, in fact, crucial, for under certain conditions, the analytic signal may be written in terms of Euler's formula:

$$x_a(t) = a(t)e^{i\varphi(t)}$$

$$x(t) = a(t) \cos \varphi(t)$$

$$H[x(t)] = a(t) \sin \varphi(t)$$

with $a(t)$ being the instantaneous amplitude and $\varphi(t)$ being the instantaneous phase. Note that this relationship only produces practicable results under the terms of Bedrosian's theorem, i.e. if the Fourier transforms of $a(t)$ and $\cos \varphi(t)$ are separable. A narrow-band signal is essential to ensure such separability. Fortunately, bandpass filters of 0.01-0.1 Hz have already been adopted as standard practice (Biswal et al. 1995; Buckner et al. 2009), so this does not necessitate major changes to analysis. Having found the equivalence between the analytic representation of a signal and Euler's formula, extracting the signal's time-resolved phase $\varphi(t)$ is relatively straightforward.

1.3.2.2 Phase Synchronization

We now return to the challenge of measuring functional connectivity under the CTC hypothesis. A method from physics introduces itself here, namely phase synchronization. Introduced to study the behavior of weakly coupled oscillators, it was first applied to fMRI in 2012 (Glerean et al. 2012). The process compares two

signals m and n by decomposing them into phase and amplitude components. This is achieved by converting their real-valued timeseries $x(m, t)$ into its analytic representation $x_a(m, t)$, as described above. Upon obtaining the time-resolved phases of these regions, the real-valued phase coherence coefficient at time t is

$$\begin{aligned} \text{dFC}(m, n, t) &= \text{Re}[e^{i(\varphi(m, t) - \varphi(n, t))}] \\ \text{dFC}(m, n, t) &= \cos(\varphi(m, t) - \varphi(n, t)) \end{aligned}$$

The cosine function has the notable advantage of continuously spanning the range $[-1, 1]$, with $\cos(\varphi(m, t) - \varphi(n, t)) = 1$ indicating that m and n are perfectly in phase ($\theta(m, t) - \theta(n, t) = 0, \pm 2\pi$), and $\cos(\varphi(m, t) - \varphi(n, t)) = -1$ indicating that they are perfectly out of phase ($\theta(m, t) - \theta(n, t) = \pm\pi$). Repeating this process for all pairs of regions in a brain atlas produces an $N \times N$ matrix of phase coherence coefficients for each time point t .

Why select this method for computing the functional connectivity, rather than the more commonly utilized Pearson correlation? Two arguments present themselves. First, there is the practical reality that the phase coherence method provides a far superior time resolution. Whereas Pearson correlation requires several samples and thus necessitates a time window several TRs long, the phase coherence method can estimate connectivity on the order of a single time-to-repetition. This order of magnitude improvement in time resolution allows far more precise estimates of functional connectivity dynamics and likely hews far more closely to the actual timescale of brain activity (Deco, Cruzat Grand, and Kringelbach 2019).

The second reason may be presented in either a positive or a negative light. Simply put, the communication through coherence hypothesis offers a clear, mechanistic argument for how different regions in the brain communicate. On one hand, this discards the agnostic nature of correlation-based functional connectivity; in logical terms, it introduces a new assumption to the analysis, a new assumption which is by no means guaranteed to be correct. On the other hand, this very fact renders the CTC a testable hypothesis; if CTC-based functional connectivity dramatically differs from that found by correlation-based analysis, it will be clear that one of these approaches is flawed. As results from these two approaches are, in fact, consistent (Glerean et al. 2012), it appears that the CTC

hypothesis does hold; as such, phase-based functional connectivity is a viable means to improve temporal resolution in functional connectivity analysis.

1.4 Effective Connectivity

1.4.1 Definitions

Effective connectivity can be defined as the influence one neural system exerts on another (Friston 1994). Such influence can be described at any spatial or temporal scale, from the synaptic to the cortical level. At the synaptic level, the effective connectivity can often be determined directly by observing the efficacy of synaptic transmission between two neurons. Such direct measurements are seldom possible at the cortical level, and entirely impossible in non-invasive neuroimaging studies. Human neuroscientists must instead rely on mathematical models and computational fitting algorithms to estimate effective connectivity at the whole-brain level. This necessarily introduces several problems: which model should one use to simulate regional activity? How should one estimate the connectivity strength between regions? Which algorithm is best suited to solving the optimization? Which cost function will provide the best fit? These are only a few of the questions which such optimization requires researchers to answer.

1.4.2 How is a neuron like a nitrogen atom?

The naïve approach to simulating the activity of a cortical region is to simulate the interaction of every neuron in said region. Efforts in this direction are in fact being made in the Blue Brain Project, specifically in the mouse brain. The fact that such efforts have thus far been ongoing for fifteen years, receive the undivided attention of multiple research groups, require multiple computing clusters, and are still orders of magnitude short of a complete simulation should illustrate the difficulties in achieving such goals. Suffice it to say, it is not a realistic approach for most studies. Thus, connectivity researchers must grapple with the question of how to extract metrics

which can describe system-level behavior from an immensely complicated system composed of billions of interacting parts.

Fortunately, such a dilemma has been at least partially solved in another scientific discipline: statistical physics. At sea level, a single cubic meter of air contains some 10^{25} molecules. It should be obvious that attempting to track the positions, velocities, and interactions of all these molecules is impossible with today's technology. Fortunately, it is also unnecessary, for the simple reason that when dealing with such enormous numbers, the activity of a single molecule, or even many thousands of molecules, has no meaningful effect on the behavior of the system. It is not necessary to track the behavior of individual molecules; it is only necessary to track their average behavior, which turns out to be a far more tractable problem. In the same way, it is not necessary for modern computational neuroscientists working at the whole-brain level to simulate the behavior of individual neurons, but only to simulate their average behavior over a specified region. So the question then arises: how to simulate such behavior?

1.4.2.1 Neural Masses

The approach followed in this thesis is to consider individual cortical and subcortical grey matter regions to be fully spatially insulated from one another. Each region (node) possesses independent intrinsic dynamics which are generally assumed to be distinct from one another. All interregional interaction takes place via coupling described by a *connectivity matrix* and *global coupling parameter*, in which the element C_{ij} describes the influence node j has on node i . This formulation considers the spatial orientation of regions relative to one another irrelevant; each region's behavior is spatially continuous, and interactions occur strictly through the coupling matrix. Those with a background in physics may observe a parallel between such a simplification and that of the *point mass*; both simplify an object's interactions with the environment to a single parameter and discards information on spatial extent or boundaries. With this parallel in mind, each self-contained region is described as a *neural mass*, with a network of such masses comprising a *neural mass model*.

Despite the intuitive nature of neural mass models, their practical implementation did not begin until the late 2000s. This delay can be attributed to a lack of computational power, a lack of detailed structural connectivity atlases, and the inability to accurately map anatomical connections in living patients. The publication of detailed structural connectivity atlases (Kötter 2004) and the development of efficient structural MRI algorithms allowed realistically wired neural mass models to make an appearance in the late 2000s (Honey et al. 2007; 2009; Ghosh et al. 2008a; 2008b; Deco et al. 2009; “Correction for Deco et al., Key Role of Coupling, Delay, and Noise in Resting Brain Fluctuations” 2009). Since these publications, however, it’s been off to the races.

Neural mass models are generally derived from a preexisting models of neural membrane potential dynamics, of which several have proven useful. One such model is the Morris-Lecar model (Morris and Lecar 1981; Larter, Speelman, and Worth 1999), which predicts pyramidal cell membrane potential based on voltage-dependent ion conductance, specifically the conductance of sodium (Na^+), potassium (K^+), and calcium (Ca^{+2}) ions. Conversion of this neuron-level model to a neural mass model took place in the early 2000s (Breakspear, R. Terry, and J. Friston 2003; Breakspear 2004) and relied on the assumption of purely excitatory long-range coupling. Such a model has proven capable of replicating some features of BOLD signals, such as phase synchrony and self-organizing activation patterns (Cabral 2012).

The classic Fitzhugh-Nagumo model (FitzHugh 1961; Nagumo, Arimoto, and Yoshizawa 1962) has also been converted to a neural mass formulation and the dynamics of both isolated and networked masses examined as a function of connectivity strength and transmission velocity (Ghosh et al. 2008b; 2008a). Isolated neural masses acted as damped oscillators; however, when coupled with biologically plausible values of connectivity strength and transmission velocity, the network displayed oscillatory behavior near the frequency of 10 Hz. Such behavior was shown to replicate the slow BOLD signal fluctuations of the resting state, which demonstrates the crucial role that connectivity and transmission delays play in shaping brain function.

The final neural mass model of interest to this thesis (although far from the last one developed) is the Wilson-Cowan model (Hugh R.

Wilson and Cowan 1972; H. R. Wilson and Cowan 1973). This represents one of the earliest attempts to estimate the behavior of a neural ensemble rather than individual neurons—no small feat in a time before the widely available and powerful computers of modern research. The model treats each region of interest as composed of two coupled populations of neurons, one excitatory and one inhibitory, which each serves to regulate the other's output. Notably, this model is known to undergo a Hopf bifurcation in some regions of parameter space, transitioning from damped oscillatory to limit cycle behavior. When models at the edge of such a bifurcation are coupled according to the macaque connectivity matrix with biologically realistic transmission delays, two distinct anticorrelated functional modules emerge. These modules compete with one another for synchrony across the network, with fluctuations in the same frequency range as those observed in resting-state BOLD signals. It thus seems that regular, oscillation-like fluctuations in activity naturally emerge when coupling neural populations across large distances.

1.4.2.2 The Observation of Oscillations

Despite the disparate bases of these neural mass models, some similarities emerge between all of them. Of note is the fact that, for most biologically plausible parameter ranges, all three models report oscillatory dynamics which approximate the frequency and structure of those observed in resting-state fMRI. This leads to the question: if biologically inspired coupled neural masses naturally settle into oscillatory stable states, could a network of coupled oscillators reproduce those same dynamics?

This is no minor question. The neural mass models described above have clear bases in biological neural interactions. A network of canonical coupled oscillators, such as the Kuramoto model (Kuramoto 1983), has no such basis. If such a system imitates real brain dynamics, even imperfectly, it implies that researchers can ignore neural mechanics at the whole-brain level. *The macroscopic behavior is independent of the microscopic interactions!*

To test this hypothesis, Cabral and colleagues began testing whether a network of Kuramoto oscillators coupled according to known white matter circuitry could replicate BOLD signals in the human

cortex (Cabral et al. 2011). Conceptually, the Kuramoto model is quite straightforward. An individual Kuramoto oscillator's dynamics can be fully described by its initial phase and its angular frequency:

$$\frac{d\theta_n}{dt} = \omega_n$$

with ω_n being the natural frequency of node n when uncoupled. When dealing with two or more coupled oscillators, researchers must hypothesize how this coupling affects each one. One common coupling model is the phase difference coupling, which presumes oscillators tend to synchronize. Assuming a network of N phase oscillators connected according to a connectivity matrix \mathbf{C} , the behavior of any node n in the network can be described by

$$\frac{d\theta_n}{dt} = \omega_n + G \sum_{p=1}^N C_{np} \sin(\theta_p(t) - \theta_n(t))$$

with ω_n being the natural frequency of node n when uncoupled and G indicating the global coupling efficiency. This model predicts that any node n will attempt to oscillate at its natural frequency ω_n , while other nodes in the network will attempt to synchronize node n with their own oscillations. The relative influence of internal versus external dynamics depends on the global coupling efficiency G and the in-strength for each node n .

When applied to the brain, the Kuramoto model's parameters should be conceptualized as, if not biologically derived, then at least biologically inspired. The variables C_{np} and G may be conceptualized as the axonal and synaptic transmission efficacy, respectively. When modeling brain activity, the structural connectivity matrix is generally used to approximate the Kuramoto connectivity matrix \mathbf{C} based on the assumption that interregional coupling strength scales linearly with the number of white matter tracts connecting two regions. Implicit in this is the assumption that all brain regions are of comparable size; if this is not the case, the connectivity strength C_{np} must be normalized by the relative size of the target region n to balance the relationship between nodes' internal dynamics ω_n and the external influences mediated by \mathbf{C} . Finally, to account for the presence of random noise in brain activity, Cabral and colleagues added a Gaussian noise term $\eta_n(t)$

with mean zero ($\langle \eta_n(t) \rangle = 0$) and variance σ_n^2 , both measured in radians for a final model of

$$\frac{d\theta_n}{dt} = \omega_n + G \sum_{p=1}^N C_{np} \sin(\theta_p(t) - \theta_n(t)) + \eta_n(t)$$

Having defined the model, one can now turn to the question: does the Kuramoto model, given biologically plausible parameters, replicate the dynamics viewed in fMRI BOLD signals? Cabral and colleagues demonstrated one decade ago an affirmative answer one decade ago (Cabral et al. 2011; 2012). This suggests that one can model and, to a limited degree, predict human brain activity and its dysfunction in disease without reference to neural dynamics at all! A canonical oscillatory model is all that is required—a fact which vastly simplifies the problem of simulating brain dynamics and estimating their controlling parameters. More importantly, it suggests that much human brain dynamics is completely independent of microstructure.

1.4.2.3 The Hopf Bifurcation

While the Kuramoto oscillator has proven capable of replicating certain elements of human BOLD signals, it remains an imperfect model. One of its greatest drawbacks is the fact that, as the name implies, it is restricted to oscillatory dynamics. While strong oscillations have been observed in regional local field potentials, individual regions may also display noise-driven deviations from a relatively stable quiescent state, or switch between quiescent and oscillatory dynamics. In addition, fMRI BOLD signals do not maintain a constant amplitude; indeed, the BOLD signal's analytic representation explicitly includes a time-dependent amplitude term. As the Kuramoto model only models the phase of a given region and assumes that each node has an inherent frequency ω_n , it is entirely unable to represent these behaviors. A more flexible regional model is necessary to better capture global dynamics.

In 2011, Freyer and colleagues proposed that the normal form of a Hopf bifurcation could fill this need in the context of EEG recordings (Freyer et al. 2011; 2012). The Hopf bifurcation, or Stuart-Landau oscillator, is a standard model for systems which can transition between stable and oscillatory behavior—precisely the

kind of dynamics observed in regional local field potentials. Deco and colleagues applied such a model to both MEG and BOLD data in 2017 (Deco, Cabral, et al. 2017; Deco, Kringelbach, et al. 2017) and found that that it can predict the slow (0.01-0.1 Hz) fluctuations observed in fMRI data well.

How does the Stuart-Landau oscillator differ from the Kuramoto oscillator? One can begin to answer this question by examining the governing equations:

$$z_j = \rho_j e^{-i\theta_j} = x_j + iy_j$$

$$\frac{dz_j}{dt} = z_j[\alpha_j + i\omega_j - |z_j|^2]$$

This set of governing equations makes clear that the Stuart-Landau oscillators is indeed an oscillator; the phase θ_j changes in time dependent on an angular frequency $\omega_j = 2\pi f_j$. However, the Hopf oscillator explicitly includes an amplitude term ρ_j , whereas the Kuramoto model ignores amplitude altogether. Further, this amplitude term can be controlled. The chain rule of derivatives implies that

$$\frac{d\rho_j}{dt} = \alpha_j - |z_j|^2$$

which brings us to an interesting conclusion. Remember that since $|z_j|^2 = \rho_j^2 (e^{-i\theta_j} \cdot e^{i\theta_j}) = \rho_j^2$, $\frac{d\rho_j}{dt} = 0$ only at $\rho_j = \sqrt{\alpha_j}$. For values of $\alpha_j < 0$, $\frac{d\rho_j}{dt} < 0$ for any value of ρ_j , so the amplitude decays monotonically to zero. For values of $\rho_j > \sqrt{\alpha_j}$, the amplitude will decay until $\rho_j = \sqrt{\alpha_j}$. However, for values of $\rho_j < \sqrt{\alpha_j}$, the amplitude will *grow* until $\rho_j = \sqrt{\alpha_j}$. Thus, at $\alpha_j = 0$, the system's behavior changes from a stable fixed point at zero to a stable oscillation with amplitude $\rho_j = \sqrt{\alpha_j}$. This provides considerably more flexibility to model behavior than the Kuramoto oscillator. The literature on Hopf models is split between polar and Cartesian coordinates. As this group tends to prefer Cartesian coordinates, further examples will appear in that system.

Having established that a Hopf bifurcation may better suit the needs of a brain modeling, the steps of adding interregional effects and noise largely followed the logic of the Kuramoto model. Translating into Cartesian coordinates, this produces the final model

$$\frac{dx_j}{dt} = x_j(\alpha_j - x_j^2 - y_j^2) - \omega_j y_j + G \sum_i C_{ij} (x_i - x_j) + \beta \eta_j(t)$$

$$\frac{dy_j}{dt} = y_j(\alpha_j - x_j^2 - y_j^2) - \omega_j x_j + G \sum_i C_{ij} (y_i - y_j) + \beta \eta_j(t)$$

where, as with the Kuramoto model, ω_j is the characteristic frequency of node j , C_{ij} is the connection strength from j to i , and G represents global coupling efficiency. This leaves the modeler with four parameters to set for each node: the bifurcation/amplitude parameter α_j , the characteristic frequency ω_j , the connection strength C_{ij} , and the noise η_j . For the present thesis, ω_j is estimated directly from the BOLD time series by extracting the dominant frequency of node j within the band of 0.01 to 0.08 Hz. α and G are set to the initial values of $\alpha = 0$ and $G = 0.2$, in line with previous work (Deco and Kringelbach 2016; Deco, Kringelbach, et al. 2017) while $\eta_j(t)$ is assumed to follow a standard normal distribution, i.e. zero mean and standard deviation of unity.

1.5 Summary

This section has attempted to summarize some thirty years of connectivity research. No doubt some details—perhaps some significant ones—have been omitted.

- Structural connectivity refers to the physical connections between network nodes, regardless of spatial scale.
- Functional connectivity refers to the temporal coincidence of spatially separated neurophysiological events—in the present case, regions which display similar phases in their time-resolved analytic signals. Functional connectivity is always symmetric.
- Effective connectivity describes the level of influence one region of a network has on another. This is not necessarily a symmetric measure; region k may influence region j without receiving a reciprocal influence.
- Effective connectivity, by necessity, must model the effect each node has on every other node. This necessitates the

development of neural mass models capable of reproducing regional time series dynamics.

- The current analyses utilize a normal form of a Hopf bifurcation, which has been proven capable of replicating the slow dynamics captured by blood oxygen-level dependent (BOLD) signals in fMRI.

2 The Resting State

Since the discovery of ordered resting state activity in the human brain, considerable effort has been placed into mapping its functional and structural characteristics. The development of dynamic functional connectivity, and with it the ability to characterize functional dynamics on the scale of seconds rather than minutes, has only added fuel to that fire. Parallel to these efforts at characterization have come efforts at modeling, which have the goal of predicting not just how but why the resting state networks display the structure and functional features which are currently known. Such efforts are ongoing, something which is unsurprising given the relatively novel state of the field; nonetheless, some trends have begun to emerge.

2.1 Definitions: Brain Activity During Rest

The resting state is generally understood as the state of consciousness which exists while sitting or lying still, fully conscious, in the absence of physical or mental tasks. The modifier of *fully conscious* is crucial in the definition of this state, as sleep, drowsiness, coma, and anesthesia are considered entirely distinct states which require independent study. Long considered a mere baseline state against which to compare task-based activity, the discovery that the resting state in fact displays rich and spatiotemporally ordered activity may be the single most revolutionary—and, in hindsight, obvious—discovery of the past two decades of neuroimaging research.

2.2 The Resting State: A History

2.2.1 Early Days: Discovery and Validation

It is perhaps fitting that a finding as revolutionary as ordered resting state activity was discovered quite by accident. Indeed, by the researchers' own confession (Biswal 2012), the discovery of temporally correlated activity between left and right sensorimotor

cortices (Biswal et al. 1995) was so unexpected that many researchers considered it an existential threat. As well they might have, for almost all fMRI studies up until that point had treated the resting state as simple background noise, to be sampled only as a means of filtering unwanted interference from task data. The discovery of order in such supposed noise raised serious questions that results from prior task-based studies might be skewed, thus rendering their conclusions invalid. Despite this fear—or perhaps because of it—research continued, demonstrating that the observed BOLD fluctuations were clearly related to low-frequency fluctuations observed in EEG studies (Biswal et al. 1997; Biswal, van Kylen, and Hyde 1997), and that some link existed between resting-state and task-based activity (Biswal et al. 1998).

With results trickling in and the existence of organized resting-state activity confirmed, a clear need for a defined baseline state emerged (Gusnard and Raichle 2001). Research in this direction soon revealed that a consistent set of regions displayed greater activity during rest than during task (Raichle et al. 2001), and that these areas consistently co-activated with one another (Greicius et al. 2003). At about the same time, magnetoencephalographic (MEG) studies and electroencephalographic (EEG) studies (Nikouline et al. 2001; Z. Liu et al. 2010; Laufs et al. 2003) found that beta-band power fluctuations (17-23 Hz) appear to underlie this specific resting-state hemodynamic activity. Thus the *default mode network* (DMN), so called because it appeared to operate as the standard, or “default”, activity pattern of a human brain not engaged in an active task, reached a form recognizable today.

It did not take long for other resting-state networks to be identified (Beckmann et al. 2005; Damoiseaux et al. 2006). A veritable explosion of studies soon revealed roughly ten distinguishable resting-state networks, with the exact number varying slightly from study to study. Of these networks, the majority were found to be task-positive; that is, cerebral blood flow to these networks increases during a task and decreases in the resting condition. This is perhaps unsurprising, as these task-positive networks appear to in fact *be* task-based networks; studies have identified visual, auditory, salience, sensorimotor, attention, and executive control networks in the resting state (Fox et al. 2005; Beckmann et al. 2005; Damoiseaux et al. 2006; de Luca et al. 2006; Mantini et al. 2007;

Brookes et al. 2011; Larson-Prior et al. 2011; Moussa et al. 2012). By contrast, the default-mode network sees its activity suppressed during cognitive tasks. A somewhat odd picture thus emerges: while the default-mode network dominates in the resting state, task-based networks do not fully deactivate. Instead, the brain appears to intermittently activate task networks even while resting, as a boxer might shift his position during a lull in the match. It is, perhaps, a telling parallel. A boxer must never set his weight while sparring, for once set, he cannot easily move it again, and is thus left unable to quickly respond to his opponent. Perhaps, in a similar fashion, the brain must avoid settling into a single state lest unexpected sensory inputs demand rapid action.

Simultaneous with the fMRI studies listed above, efforts were made to further characterize the electrophysiological signals of the resting state. These studies began to bear fruit in 2007, when Mantini *et al.* demonstrated that resting-state networks appear to correlate with electrophysiological signatures (Mantini et al. 2007). Intercranial studies soon confirmed that slow neural fluctuations did indeed underlie the observed MRI signals, while EEG studies showed that gamma-band local field potentials (LFPs) and firing rate modulations control task-positive RSNs (Miller, Weaver, and Ojemann 2009; Shmuel and Leopold 2008). Note that this contrasts with the default-mode network, which most strongly maps to the alpha- and beta-bands electrophysiological signal. Further studies suggested that each RSN maps to a specific electrophysiological signature, each spanning a range of frequency bands (Mantini et al. 2007; de Pasquale et al. 2010). This led to suggestions that the various frequency bands may control specific aspects of neural processing (Womelsdorf et al. 2007; Palva and Palva 2012), a suggestion strengthened by the discovery of frequency-specific spatial correlations in spontaneous neural oscillations (Hipp et al. 2012). These correlations appear to link functional hubs of the brain to specific frequency bands; for instance, theta-band frequencies most strongly correlate in the medial temporal lobes, while lateral parietal areas resonate most strongly with alpha- and beta-band frequencies. All in all, then, task-negative networks appear to associate with alpha- and beta-band signals, while task networks associate with the higher frequencies.

For all this apparent progress, one question so far remains unanswered: are the observed resting state networks consistent across subjects? Damoiseaux and colleagues demonstrated that this is indeed the case; not only do the RSNs display remarkable consistency across subjects (Damoiseaux et al. 2006), but what variation does exist can be explained almost entirely by changes in age and sex (Biswal et al. 2010; Andrews-Hanna et al. 2007; H. Liu et al. 2009). These findings were only further reinforced in 2011, when Brookes *et. al.* mapped the resting-state networks within magnetoencephalography (MEG), a far more direct measure than the blood-oxygen level-dependent (BOLD) signal of fMRI (Brookes et al. 2011).

So, come 2012, what image had emerged? The human resting state is organized into several separable functional networks. These networks may be separated into the categories of task-positive and task-negative. Each network produces unique spatial and frequency signatures. The primary task-negative network, namely the default-mode network, operates primarily in the alpha- and beta-band frequencies, while task-positive networks are controlled by gamma-band frequencies. Task-positive resting-state networks appear to in fact be task networks, which are not deactivated in the resting state but simply reduced to a low-power, or standby, state. The implication appears to be that the resting-state does not remain in a single functional network over the course of a scan, but rather alternates between many, with the default mode network only being the strongest amongst many attractive states. Without further temporal resolution, it would be difficult to explore the realm of resting-state functional connectivity further. Fortunately, this dilemma solved itself virtually before being noticed, for in 2009, dynamic functional MRI entered the stage.

2.2.2 Dynamism in Data

The development of dynamic functional MRI (dfMRI) has been yet another shock to a scientific community which is still attempting to metabolize the discovery and popularization of the resting state. While EEG and MEG studies have long demonstrated that brain activity changes on the scale of milliseconds, neither method yet possesses the spatial resolution of fMRI. Thus, for several decades,

researchers faced the dilemma of choosing between high spatial and poor temporal resolution, or poor spatial and high temporal resolution. Attempts to bridge this gap via joint fMRI-EEG studies yielded some progress, but such experiments add considerable complexity to analysis.

The introduction of an MRI method which could resolve hemodynamic activity on the scale of seconds, rather than minutes (Sakoğlu and Calhoun 2009; Sakoğlu et al. 2010), was little short of a revolution. While far from the millisecond resolution of EEG and MEG, it still represented a full order of magnitude's improvement. The original version, presented in 2009 and 2010, consists in essence of multiple consecutive correlation-based functional connectivity estimates. Rather than correlating regional activity over an entire scan, the scan is segmented into overlapping windows, usually several tens of seconds in length. The correlation coefficients of the time courses in each window are then estimated to produce a correlation matrix. This process is repeated for each window to produce a series of time-resolved correlation matrices, each representing the average functional connectivity of a separate window. Conceptually, Calhoun's approach was a natural evolution from the previous "static" functional connectivity; whereas static functional connectivity (sFC) correlates regional time series over an entire scan, *sliding-window* functional connectivity (so-called for the windowing of the time series) correlates regional time series over short windows which slide along the temporal dimension of the scan. Such an approach is not without disadvantages: the use of correlation- or covariance-based metrics requires that windows be several times-to-repetition (TRs) long, with predictable effects on temporal resolution. Nonetheless, its development virtually immediately confirmed the transience of resting-state networks first observed in MEG and allowed for an explosion of studies into the dynamics of functional connectivity.

Although introduced in 2010, dynamic functional MRI did not fully break into the neuroimaging world until the end of 2014 (Zalesky et al. 2014). Given this short timeframe, the field remains rather unconsolidated, although efforts are being made to rectify this (Preti, Bolton, and van de Ville 2017; Vohryzek et al. 2022; Kringelbach and Deco 2020). Despite this, certain trends have

clearly emerged, one of which is the existence of separable, recurrent connectivity patterns. This finding is not entirely unexpected, as EEG and MEG studies have long demonstrated recurrent microstates (Khanna et al. 2015). Nonetheless, the proliferation of MRI-based dFC has allowed previously unattainable spatial resolution and quantity of data to be collected, which in turn has allowed researchers to confirm the existence of repeating connectivity states in the resting brain.

How does one isolate those states? This is a question without a consistent answer. Several methods have been employed to this end, including principal component analysis (Leonardi et al. 2013), eigendecomposition and k -means clustering (Cabral et al. 2017; Lord et al. 2019; Figueroa et al. 2019), eigendecomposition and independent component analysis (Blair et al. 2022), hidden or semi-hidden Markov models (Shappell et al. 2019), and non-negative matrix factorization (Du, He, and Calhoun 2021; Glomb et al. 2017). Efforts have also been made to predict these functional substates from known anatomical and dynamic constraints (Ashourvan et al. 2021; Deng et al. 2022; Gu et al. 2018; Atasoy et al. 2018). Unfortunately, efforts to compare or harmonize the results from these various studies remain in their infancy, so little information is available on these methods' reliability in either spatial or temporal dimensions. Yet the essential finding—namely that recurrent connectivity substates exist and that the resting brain transitions between them in the absence of external stimuli—remains consistent. Increasingly, it appears that most, if not all, such functional substates correspond to identified resting state networks (Vohryzek et al. 2020).

In addition to the spatial structure of these dynamic resting state networks (dRSNs), considerable effort has been placed into describing their dynamics. Such efforts have employed both predictive (Gu et al. 2018; Atasoy et al. 2018) and empirical (Deco et al. 2019) strategies, and most have utilized some version of a transition matrix. While an intuitive and powerful method for describing the likelihood of each dRSNs' appearance as a function of the brain's current state, the transition matrix does imply that each dRSN exists in only one of two states, *on* or *off*, and that no two dRSNs may be active simultaneously. This is likely an

oversimplification of actual dRSN interactions, which may well display co-activation and large variations in activity levels. A method to capture and visualize dRSN dynamics on a continuous scale is needed to bridge this gap.

2.2.3 Summary

The human brain's resting state has proven to contain remarkably complex behavior. Laudable efforts have been and are being made to chart its organization, but these efforts remain in their infancy. It is nonetheless clear that resting-state activity contains multiple recurrent states, which dynamically shift activity levels. This raises the question: what form of dynamic system does this activity take?

2.3 Current Hypotheses

2.3.1 The Resting State as a Multistable System

A dynamic system may have points in its phase space from which the system will not move unless disturbed. For example, a pendulum at rest in a gravitational field will not spontaneously begin to oscillate. Similarly, a sled on a leveled area above a hill will not spontaneously begin to slide; it must enter the sloped area of the hill before gravity can pull the sled downwards. Although their behavior is very different once in motion—the pendulum will oscillate around its rest position, while the sled will accelerate away from that rest position—both examples share a common feature; absent some external force or a change in environmental conditions, they will remain in their rest position. A position in a dynamical system's phase space which satisfies this condition is called a *fixed point*, as a system which starts at that point is “fixed” there absent some external driving force.

It should be obvious that the two scenarios described above—the pendulum and the sled—have very different behaviors once in motion, however. Once in motion, a pendulum will oscillate around its rest position. A real pendulum's oscillations will decay over time due to the dampening effect of friction, while an ideal pendulum will continue to oscillate around its rest position forever.

In both cases, the pendulum remains in the vicinity of its rest position. The sled, on the other hand, will travel away from its rest position once in motion, and will continue to do so until it reaches another leveled area. These distinctions—a tendency to return to the resting position, a tendency to oscillate around the resting position, and the tendency to travel away from the resting position—define three of the basic trajectories of dynamical systems: a *stable fixed point*, a *cycle*, and an *unstable fixed point*. A system which is moved a small distance from a stable fixed point will eventually return to that fixed point, much like a real pendulum will eventually come to rest. A system which enters a cycle will trace a closed trajectory through its phase space, which is to say it will continue tracing the same path unless disturbed, just as a pendulum would in the absence of friction. Finally, a system leaving an unstable fixed point will accelerate away from that fixed point, just as a sled will accelerate away from the peak of a hill.

What if a system has more than one fixed point? Actually, all three systems described above do. The pendulum has an obvious fixed point at its rest position, i.e. with the main mass directly below the pivot point, but it also has a fixed point in the inverse position, i.e. with the mass directly above the pivot! This second fixed point is obviously unstable, as even a miniscule deviation from perfect balance will cause the pendulum to swing away from the vertical, but nonetheless, this second fixed point does exist. In the case of the sled, an unstable fixed point exists at the hill's peak, but a stable one also exists at its base. Any level sections of the hill's slope will form a *semi-stable* fixed point, i.e. a fixed point which attracts the sled from one direction and repels it from the other. In both cases, then, we find that the system contains several fixed points. However, it should be obvious that the system tends to enter only one of its fixed points, that being the stable one. For this reason, we can call these fixed points *attractors*, because they attract the system to enter a permanent state. An attractor's *basin of attraction* refers to the region of phase space under the attractor's influence; that is, if the system enters a basin of attraction, it will tend to move towards the attractor, just as a sled, once on the slope of a hill, tends to move towards the hill's base.

Suppose now that our sled is not at the top of one hill but is rather sitting at the top of a low ridge between two hills. Also suppose

that in the two pockets between this ridge and either hill, there are children who, if the sled falls towards them, will climb partway up the ridge to sled down. Each child will climb an arbitrary distance before growing tired and jumping on the sled. If one child carries the sled to the top of this ridge, he will have the option of either sledding to his original group or sledding to the other. Thus the sled may alternate between the two sides of the ridge, according to the (somewhat) random whims of whoever catches it on the downwards journey. It is important to note that both sides of this ridge are attractors, as absent the children's interference, the sled would slide to one side and remain there. It is only due to this unpredictable input of energy that the sled may move between the two basins.

A system such as this—one with two or more stable attractors which it may alternate between—is called a *multistable* system. The terminology is (hopefully) somewhat self-explanatory. Such systems are familiar from the realm of optical illusions, such as Rubin's vase or Ludwig Wittgenstein's duck-rabbit:




Figure 1: Behold the mighty duck-rabbit. Credit to the 23 October 1892 issue of *Blätter*.

Both these illusions are examples of bistability: they may be perceived as one of two potential objects, and the observer may alternate between which object he (or she) perceives. For the Rubin's vase, most observers will alternate between seeing an ornate vase or two facial silhouettes; for the duck-rabbit, an observer may see a duck, a rabbit, or may alternate between the two. Most observers will not see both at once, however. This is one of the characteristics of a multistable system: while it may switch between stable states, it does not exist in both simultaneously. Instead, external interference may cause the system to alternate between the two (or more) attractors.

Well and good, but what does this have to do with resting state functional connectivity? Recall that it has (very) recently been established that the human brain, while at rest, alternates between a finite and countable number of functional connectivity substates. It is not too great a leap of logic to suggest that these substates are attractors; stable fixed points which whole-brain connectivity tends to enter in the absence of external perturbation. The tendency to switch between these states is explained by the simple fact that the living brain is never inactive; intrinsic activity and neuronal noise are features which could easily provide the energy necessary to transition between connectivity substates. The suggestion that the human brain may exhibit a multistable regime is not unique to the realm of neuroimaging; it has been postulated to exist at virtually every level of cognition (Scott Kelso 2012). The idea has gained substantial traction in the past decade, particularly with the explosive growth of dynamic functional connectivity providing ever-greater evidence of recurring connectivity substates in the resting brain. However, this framework began to fall out of favor after several studies suggested that a multistable system is unable to reproduce both the static FC and the dynamics of dFC networks (Hansen et al. 2015; Deco et al. 2013). Some new framework was needed to explain these dynamics.

2.3.2 The Resting State as a Metastable System

The resting state may also be described as a metastable system. The difference between these two formulations lies in the existence of stable fixed points; whereas a multistable system has two or more

such points, a metastable system does not. Instead, metastable systems are characterized by excursions into recurrent, but temporary, non-equilibrium states. Herein lies the key difference between a multistable and a metastable system; even in the absence of noise, a metastable state will not persist indefinitely. Instead, a metastable system may wander through a landscape of intermittently stable states, spending substantial periods of time in each, without ever settling into a final configuration (Afraimovich, Zhigulin, and Rabinovich 2004). Such a system bears some resemblance to both the needs of adaptive cognitive functions (Rabinovich et al. 2008) and the subjective experience of mind-wandering (Christoff et al. 2016; Kucyi 2018). More importantly, metastable network models have proven capable of replicating both static and dynamic FC in the human resting state (Honey et al. 2007; Cabral et al. 2014; Vidaurre et al. 2018), and these results have plausible bases in known physiological phenomena (Beckmann et al. 2005; Smith and Nichols 2009; Brookes et al. 2011).

2.3.3 Metastable Multistability?

A recurrent debate in the neuroimaging literature is the question of whether the human brain is a multistable or metastable system. It has recently been proposed that it may satisfy both criteria via self-modulation of the dynamic connectivity energy landscape (Vohryzek et al. 2022). Under this hypothesis, certain fixed points of the brain's connectivity landscape may either emerge or disappear under a change in brain state.

Why consider metastability as a potential model for the resting state? To answer this, we refer to the fact that metastable states frequently occur in the vicinity of a fixed point which has disappeared due to a system parameter change. Recent research suggests that the resting brain spends a plurality of its time in a globally synchronized state, with brief excursions into task-relevant network configurations (Vohryzek et al. 2020). The transience of these excursions suggests multiple metastable states in the resting state energy landscape, with the most stable of these corresponding to the globally synchronized state. The resemblance of these metastable states to known task networks provides an interesting

hypothesis; that the brain's energy landscape alters upon engaging in a task, converting the relevant task network to a stable attractor while removing that of global synchrony. It is plausible that the resting state balances system parameters to keep these task-relevant states just below stability, yet close enough to stability that they may still affect global dynamics. Upon entering a task, some mechanism alters the system parameters to convert the relevant task network to a stable attractor while the global state is rendered unstable. Of yet, this hypothesis remains unvalidated; perhaps future research will shed additional light.

2.4 Summary

Research into the resting state of the human brain remains a dizzyingly active field of research. The number of methods and findings in the field increases virtually weekly, to the point that discerning useful signal from spurious noise is arguably the greatest challenge that today's neuroimaging researcher face. Nonetheless, clear themes have emerged over the past twenty years. These themes include the existence of distinct resting-state networks (RSNs) which appear to consistently jockey with one another for dominance of brain activity. The precise number of these networks remains a subject of active research, as does the nature of the mechanisms underlying them; indeed, the Organization for Human Brain Mapping (OHBM) has recently announced a collaboration dedicated to the task of identifying canonical resting-state and dynamic resting-state networks (Uddin et al. 2022). Given this ongoing and rapid evolution, it is difficult to make confident assertions on what conclusions the field will reach. Following its continued evolution will likely prove interesting.

3 Statistical Neurophysics: or, Measuring Entropy

Since the inception of the field, one of the key goals of neuroimaging has been to discover robust, precise biomarkers for psychiatric disease. Such a biomarker would provide something that the psychiatric field of medicine has sought for decades: a measurable diagnostic test for disease, which might inform the clinician which treatment a patient requires. No such test yet exists. Instead, psychiatric diagnoses rely upon self-reported tests or clinical interviews which seek to capture the symptoms of the afflicted. Such diagnoses provide little insight into the causal mechanisms underlying the symptoms, and as such, prescribing treatment remains something of an exercise in trial and error. Each advance in the field of neuroimaging has brought hope for the discovery of such a test, some feature visible on a scan which could clearly delineate the well from the unwell. Enormous progress has been made in this direction, with a dizzying quantity of studies demonstrating differences between healthy controls and psychiatric patients in any number of neuroimaging features. Nonetheless, a clear delineating line between both healthy controls and between patients of different disorders remains frustratingly difficult to locate.

Fundamentally, the search for a meaningful biomarker comes down to the question: *what brain features are affected by psychiatric disease?* As it stands, this is an impossibly broad question: as such, the first step is to shrink it to a manageable size. With this in mind, the specific question to be answered in this thesis is, *do psychiatric disorders affect the complexity of functional processing in the human brain?* While still a broad question, its phrasing allows us to begin parsing it. The key term is *complexity*, as this is a well-defined concept in several scientific disciplines. In information theory, a signal's complexity may be considered inversely proportional to its predictability, as one intuitively considers highly predictable signals to be easily interpretable (and vice versa). Each sample of a given signal contains information on the signal's state, with the amount of information per sample being inversely proportional to the signal's predictability. It is thus possible to draw a direct link between the information contained in a message and

the complexity of that message's source signal: a link made explicit in Claude Shannon's *entropy*.

3.1 What is Entropy?

Despite its ubiquitous roles in statistical mechanics and information theory, entropy has proven a remarkably difficult concept to define. It has been called a measure of information, but information in what sense? It has been described as measuring the disorder in a system, yet ordered systems can contain high levels of entropy. One can measure the richness of a system's dynamics with entropy, but again we come to the question, what does dynamical richness mean? Unfortunately, many articles and even some textbooks which use the entropy fail to clarify these questions or answer them tautologically. Perhaps this is what John von Neumann meant when he told Claude Shannon "You should call it [uncertainty measure] entropy, for two reasons. In the first place, your uncertainty function has been used in statistical mechanics under that name, so it already has a name. In the second place, and much more important, no one really knows what entropy really is, so in a debate you will always have the advantage."

It is, of course, entirely possible that von Neumann was speaking tongue in cheek in this instance. Physicists have senses of humor as well. However, a student first entering information theory or its applications could be forgiven for thinking that von Neumann had a point, i.e. that most of those using the entropy do not fully understand in what sense it quantifies information, disorder, surprise, or any of the other concepts it is meant to measure. So, the first step of this section shall be an attempt to answer that question, from the perspective of thermodynamics, statistical mechanics, and information theory.

3.1.1 The Thermodynamic Entropy

To begin this discussion, we must define a *system* in the thermodynamical or statistical mechanical sense. In these fields, a *system* is defined simply as "whatever part of the universe is

selected for study” (Blundell and Blundell 2010). This is a rather vague answer, but it does have the advantage of generalizability; a system may consist of a single cylinder of an internal combustion engine, a section of a turbine, or a biological organ, e.g. the human brain. All of these are systems which may (hypothetically) be studied in isolation from the rest of the universe.

One may also discuss the *state* of a system. A system’s state may be defined by the measurable—or, perhaps more precisely, the *controllable*—variables. Taking the example of an engine’s cylinder, one can either measure or control its temperature, pressure, volume, and the amount of energy injected into or removed from it. One cannot (nor generally does one want to) control the locations or trajectories of individual molecules within this cylinder. Thus, the pressure, volume, temperature, and energy can define a cylinder’s state; a change in any of these variables causes a *change in state*. There are conceptual parallels here to the concept of a *brain state*, which can be (crudely) controlled via inputs to the relevant brain’s sensory organs and instructions to the brain’s owner (assuming an *in vivo* study).

What does all this have to do with entropy? Thus far, the discussion of entropy has been based on its role in thermodynamics. Thermodynamics approaches entropy from the perspective of temperature and heat flow—unsurprising given the field’s foundation in the operation of steam engines during the Industrial Revolution. There is no need to go into the exact equation and derivation in this thesis, but the thermodynamic entropy may be summarized as a *measure of the reversibility of the system’s change in state*. If a change in state does not generate entropy, that change may be reversed with no net energy loss. A change in state which does generate entropy, on the other hand, that change cannot be perfectly reversed, nor the energy expended perfectly recovered.

3.1.2 Entropy in Statistical Mechanics

The definition of entropy given in statistical mechanics relates somewhat more directly to neuroscience. Roughly speaking, in statistical mechanics, one can interpret the entropy of a given state as a measure of *the probability of that state being observed given*

known constraints. Understanding this requires unpacking the definition of entropy used in statistical mechanics, which is

$$S = k_B \ln \Omega$$

with S being the specified state's entropy and k_B the Boltzmann constant. Understanding Ω requires some additional background. Statistical mechanics maintains a distinction between two types of states: *macrostates* and *microstates*. Roughly speaking, a *microstate* describes the condition of every particle in a system *individually*, while a macrostate considers their aggregate behavior. Given a box containing 100 coins, one could say that the system's microstate is the position of each individual coin: coin one is heads, two is heads, three is tails, etc. The macrostate, on the other hand, is simply the number of heads and tails visible; which coin displays which face is irrelevant.

The laws of statistics imply that certain macrostates are more probable than others due to the number of microstates that they contain. It should be clear, for instance, that our hypothetical box of coins has many, many microstates which produce the macrostate of fifty heads and fifty tails (on the order of 10^{29}), but only one which produces the macrostate of 100 heads. Assuming that every microstate is equally probable (the so-called *ergodic hypothesis*), the macrostate of fifty heads and fifty tails is some 10^{29} times more probable than that of 100 heads.

What does this have to do with the entropy? Return to Boltzmann's definition of entropy:

$$S = k_B \ln \Omega$$

As previously mentioned, S is the entropy of the macrostate under investigation, and k_B is the Boltzmann constant. Ω is the number of microstates which can produce the given macrostate. Evidently a macrostate containing many microstates will have a higher entropy than one containing few.

It is worth noting that despite common assertions, entropy is *not* a measurement of a state's "disorder" or "randomness". While disordered or random states are often more likely than ordered ones, this is not always the case. Imagine, for instance, that the system is a suitcase packed with some number of travel articles. If the size of the items to be packed is comparable to the size of the suitcase,

simply cramming everything into the case at random is unlikely to work; while there is no doubt some random arrangement which will make everything fit, there are many more which do not. Neatly packing the articles is much more likely to let the case close, because minimizing the space each article takes means that more arrangements which meet the constraints, i.e. fit in the suitcase, are available. Translating this into the terminology of statistical mechanics, the final state of the suitcase (packed and closed) is the macrostate, and the possible internal arrangements of articles (which shirt is where) the microstates. Intuitively, one perceives that packing each article neatly produces more microstates than random shoving would. This implies that, for this example system, a more ordered arrangement has *more* entropy than a random one—quite the contradiction to common assertion!

While theoretically coherent, the definition of entropy given above has one glaring practical problem. Microstates are difficult to measure directly. Indeed, in most cases they are *impossible* to measure directly; one cannot, even in theory, perfectly measure both a molecule's position and its velocity. This raises an obvious question: if microstates cannot be measured, how are they to be counted? If they cannot be counted, how is Boltzmann's definition of entropy to be used?

These questions reveal the need for a definition of macrostate entropy which does not require counting microstates. Dr. Josiah Willard Gibbs recognized this need and devoted some of his considerable intellect to it. His solution involves converting the entropy from a measure based on quantities to a measure based on probabilities—specifically the probabilities of the relatively easily measured and countable macrostates. To begin this process, define the total number of microstates as N and the number of microstates in the macrostate i to be n_i . Assuming that each microstate is equally likely, the probability of the system being in macrostate i is then

$$P_i = \frac{n_i}{N}$$

It should be observed that the system's total entropy does not only consist of the entropy of the macrostates, but that of the microstates as well. Until now, this has been neglected because microstates are (as mentioned) extremely difficult, in most cases impossible, to

measure or count. Nonetheless, using the terminology defined above, the entropy of the individual microstates in macrostate i is

$$S_{i_{micro}} = k_B \ln n_i$$

so the entropy of microstates across all available macrostates is

$$S_{micro} = \langle S_{i_{micro}} \rangle = \sum_i P_i S_{i_{micro}} = k_b \sum_i P_i \ln n_i$$

The system's total entropy is simply the entropy estimated across all possible microstates N :

$$S_{total} = k_B \ln N$$

and is the sum of the entropy of microstates and macrostates:

$$S_{total} = S + S_{micro}$$

This leads rather directly to the conclusion that the entropy associated with the system's macrostates—which one *can* measure and count—is

$$\begin{aligned} S &= S_{total} - S_{micro} \\ S &= k_B \ln N - k_b \sum_i P_i \ln n_i \\ S &= k_B \left(\sum_i P_i \ln N - \sum_i P_i \ln n_i \right) \\ S &= k_B \sum_i P_i (\ln N - \ln n_i) \\ S &= k_B \sum_i P_i \left(-\ln \frac{n_i}{N} \right) \\ S &= -k_B \sum_i P_i \ln P_i \end{aligned}$$

which is a far more practically useful definition than Boltzmann's.

3.1.3 The Entropy in Information Theory

Now we can turn to the question of how information theory approaches entropy. When Claude Shannon turned to the question of how to measure the information in a signal, he was not aware of the entropy's role in statistical physics. Thus, he approached the problem from quite a different angle. His challenge was to find a means of measuring the amount of information which a message carries. This process requires two steps: defining a measure for information and defining a model for messages.

How does one define the amount of information carried in a message? It's hardly a trivial question. One can start by asking, *how informative is an answer to a question?* To take a banal example, suppose I wish to know whether it will rain this afternoon. If, upon stepping outside, I find a bright, clear, dry day, I will expect that it will remain clear, as such days normally do. Intuitively, I expect a high probability of good weather in the afternoon, and thus a low probability of bad weather. If, upon reviewing the weather forecast, I see that it is indeed predicted to remain pleasant all day, this only confirms a previous belief. Little new information has been gained. On the other hand, if the weather report predicts a thunderstorm, I would be quite surprised and would have to reevaluate my prior belief in light of this substantial new information. Following this example, one can intuit that an unlikely message conveys more information than a likely one, which implies that information is an inverse function of probability:

$$I(a) \propto \frac{1}{P(a)}$$

What of the case of two independent messages? How would the joint information of these combine? If the messages are independent, e.g. a weather report and a reminder to pick up groceries, one intuitively expects the information they contain to combine linearly (add):

$$I(a\&b) = I(a) + I(b)$$

Yet the probabilities of two independent events occurring combine multiplicatively:

$$P(a\&b) = P(a)P(b)$$

and it is already established that information is an inverse function of probability. Fortunately, there is a class of functions which bridges this divide: the logarithm.

$$\begin{aligned} \log P(a\&b) &= \log P(a)P(b) = \log P(a) + \log P(b) \\ \log \frac{1}{P(a\&b)} &= \log \frac{1}{P(a)P(b)} = \log \frac{1}{P(a)} + \log \frac{1}{P(b)} \end{aligned}$$

Remembering that the logarithm of an inverse is the negative logarithm,

$$\log \frac{1}{P(a)} + \log \frac{1}{P(b)} = -\log P(a) - \log P(b)$$

Putting all this together, it seems that the information contained in two independent messages is proportional to the negative log probability of each message:

$$I(a) \propto -\log P(a) - \log P(b)$$

Extending this line of logic to an arbitrary number of independent messages N produces

$$I(M) \propto -\sum_{i=1}^N \log P(m_i)$$

with M being a combination of N independent messages and m_i being the i^{th} message.

Any measure of information must also account for dependencies between messages. Suppose that upon seeing the weather report, I retrieved an umbrella. At work, a colleague notices this and asks why I brought an umbrella to work given the good weather. To this, I reply that while the weather is good now, it will not remain so. My colleague has just received two messages:

1. It is predicted to rain this afternoon
2. I am carrying an umbrella *because* of the previous fact

It is patently obvious that these messages are not independent; the second follows from the first. It is also clear that these messages conflict with prior beliefs, as my colleague presumably did not expect rain and did not expect me to carry an umbrella. Thus, any measure of information must account for both pieces of information and their dependency. In his seminal article, Dr. Shannon demonstrates that the only expression which satisfies such a condition is (Shannon 1949)

$$S = -K \sum_i P_i \ln P_i$$

which, aside from the constant K , is identical to the Gibbs entropy described in Section 3.1.2.

Where does this leave the entropy? One can say that in this case, the entropy of a single sample (symbol) is a measurement of the *amount of flexibility possessed by the transmitter*. One could also say that it is the *amount of uncertainty which the message resolves when received*. These two statements are equivalent; a transmitter sending a highly stereotyped message will have very little flexibility in the message's contents, and the recipient will be able to predict

those contents with considerable certainty. Conversely, the sender of a highly unpredictable message—stream-of-consciousness poetry, for instance—will have a great deal of choice (flexibility) in the message's contents, and the recipient will find the contents very difficult to predict.

3.1.4 Tied Together

So in the final analysis, what *is* entropy? Combining the definitions of thermodynamics, statistical mechanics, and information theory, we can roughly say that it is a measure of the *flexibility of the system's dynamics*. A system or message with high entropy will always have more options for organizing its internal structure than a system with low entropy. This explains that classic thermodynamic finding that entropy always increases: a system with low entropy can contain only a few states, whereas a system with high entropy can contain many more. Assuming that all these states are equally likely (the so-called *ergodic hypothesis*), a system is more likely to enter a state which contains many substates than a state which contains only a few.

3.2 Why to Measure It?

As described in Section Two, the discovery of temporal variability in dynamic functional connectivity raised led to questions about these dynamics' organization. Do the dynamics of the human functional connectome lie in a multistable regime, a metastable regime, or some combination of the two?

Entropy provides a potential means of resolving this question. Intuitively, one expects a system composed of stable attractors to be more constrained in its dynamics than a system without such attractors. Thus, assuming equal numbers of attractors, a metastable system should have more display higher entropy than a multistable one. Such a finding could be further investigated by fitting a generative model to empirical data and examining the optimal working point of each region. Alternatively, novel methods have been proposed to directly measure the broken detailed balance and

turbulent dynamics of neuroimaging signals (Martín 2019; Deco and Kringelbach 2020; Deco et al. 2021), all of which are closely related to Shannon entropy.

Beyond the scientific question of whether the brain exists in a multistable or a metastable regime, entropy-based analyses allow direct comparisons of stability between healthy and patient groups. Several psychiatric disorders are hypothesized to display aberrant connectivity dynamics. For instance, major depressive disorder (MDD) is thought to be partially caused by hyperstability of the default mode network (Vohryzek et al. 2022), which may cause the ruminative thoughts known to characterize MDD. Such hyperstability should alter the entropy of depressive patients' network dynamics.

Alterations in entropy have the potential to act as more than a biomarker for disease, however. Recent work suggests that it may be possible to alter human network dynamics from a diseased to a healthy state (Deco et al. 2019), and the advent of noninvasive stimulation methods such as trans-cranial magnetic stimulation (TCM) make such interventions clinically feasible. Such interventions will need precise targets, however, so the ability to detect which functional networks display aberrant dynamics is essential. An analysis method which can identify both the functional networks of interest and directly compare the entropies of these networks between populations may provide a powerful tool for identifying such targets in future clinical practice.

4 Datasets

4.1 Motivation

We elected to test this pipeline on a dataset (Moreira et al. 2017) consisting of obsessive-compulsive disorder (OCD) patients and number of age-, gender-, and education-matched controls ($N_{\text{OCD}} = 40$, $N_{\text{control}} = 39$). The wide prevalence and severe effects of OCD factored into this choice of dataset; with some 2.1% of the population affected each year (DuPont et al. 1995), it is a widespread, yet poorly understood disorder that causes its victims great distress. Obsessive thoughts and compulsive behaviors often hinder victims' ability to concentrate, with predictable effects on learning and productivity (Piacentini et al. 2003; Weidle et al. 2014). These factors contribute to a high societal cost of illness (DuPont et al. 1995; Lenhard et al. 2021a) and reduced quality of life for patients. Despite its prevalence, the disorder's functional dynamics remain poorly understood; in particular, we have been unable to find any attempts to examine the functional complexity of OCD patients. In this study, we demonstrate that the obsessive-compulsive group displays elevated joint entropies compared to healthy controls. Indeed, not only can we identify which group has higher joint entropy, but also along which dimension the entropy changes.

Though primarily a methodological work, this thesis also intends to explore the functional signatures and potential network-based causes of psychiatric disease. I could cite any number of studies detailing the sociological or economic costs of psychiatric disease (Piacentini et al. 2003; Weidle et al. 2014), but the primary motive is far simpler. Multiple members of my family have suffered and continue to suffer the effects of anxiety and compulsive disorders, with mood disorders happy to play their typical supporting roles. The experience is not a pleasant one.

It should come as little surprise, then, that I have long had an interest in psychiatry and the neural mechanisms underlying it. More recently, the problem of screening for disease has caught my attention as well. It is widely believed that psychiatric disease is underreported and underdiagnosed due to a combination of social stigma and simple ignorance (Takayanagi et al. 2014). While

public awareness campaigns have begun to mitigate these problems, neither has been eliminated. Additionally, efforts to screen for vulnerability psychiatric distress have raised the hope that some, perhaps many, cases of psychiatric distress could be preempted via targeted intervention before symptoms reach the clinical stage. A methodologically sound means of screening for such disorders or vulnerability to them could prevent enormous suffering, to say nothing of the monetary savings and economic benefit that such prevention implies (DuPont et al. 1995; Lenhard et al. 2021b).

The above information adequately explains my interest in obsessive-compulsive disorder, which is thought to lie in or adjacent to the anxiety spectrum. What of the second dataset, however? The three disorders contained within—schizophrenia, bipolar type I, and attention-deficit hyperactivity disorder (ADHD)—have no known relationship to the anxiety or compulsive spectra, although many ADHD patients are known to be impulsive. Nonetheless, the general motives remain similar; a desire to mitigate emotional distress, a desire to reduce the monetary burden on patients and their relations, and frustration at the economic waste that mitigation, care, and treatment represent. Further, from a methodological perspective, there is the need to confirm the robustness of the LEiDA framework. Thus, I sought to test the pipeline on unrelated disorders to confirm that LEICA could recover functional alterations across a wide behavioral and cognitive spectrum.

4.2 Datasets

4.2.1 Obsessive-Compulsive Disorder

This study uses a dataset from a previous study at the Universidade do Minho, Portugal (Moreira et al. 2017). A detailed description may be found in that paper, but a summary is included here for completeness.

Eighty right-handed subjects (40 patients with OCD, 40 controls) participated in this study. Recruitment ensured that controls matched patients in age, sex, education, and ethnic origin. Both patients and controls were screened to remove subjects with

comorbid mental health, neurological or major medical disorders (except nicotine use or dependence). Patients were all confirmed to have been using stable doses of medication for three months prior to the study. Specifically, 72.2% used selective serotonin reuptake inhibitors (SSRIs), 11.1% tricyclic antidepressants (TCA), and 16.7% a combination of these medications.

Image acquisition occurred in a 1.5 T Siemens Magnetom Avanto MRI scanner (Siemens, Erlangen, Germany) with a standard 12-channel receiver coil. Images were visually examined for artifacts and the functional data preprocessed using FSL. Slice-timing correction used the first slice as a reference, a rigid-body spatial transformation aligned the volumes of each subject with the mean volume, and motion scrubbing identified time points contaminated by significant motion. Participants with more than 20 such time points were removed from analysis. Images were then non-linearly normalized to MNI standard space and linear regression used to remove motion-related variance and signals from white matter and cerebrospinal fluid. Acquisitions were filtered with a Gaussian spatial smoothing kernel (8 mm FWHM) and a temporal band-pass filter (0.01 to 0.08 Hz). This frequency band has demonstrated greater reliability and functional relevance in fMRI compared to others (Glerean et al. 2012; Biswal et al. 1995; Buckner et al. 2009; Achard et al. 2006). This low frequency band has the additional advantage of averaging out physiological noise and hemodynamic response functions, as these signals have frequencies above 0.08 Hz and thus fall outside the passband of this filter. Finally, following the preprocessing, Moreira *et al.* extracted the mean BOLD time series of the 116 cortical, subcortical, and cerebellar regions of the Anatomical Automatic Labeling atlas (Tzourio-Mazoyer et al. 2002). As our study focuses on cortical and subcortical regions, the 26 cerebellar regions of the Anatomical Automatic Labeling (AAL) atlas were removed.

4.2.2 Cross-Disorder

This article utilizes a publicly available dataset from the Consortium for Neuropsychiatric Phenomics (Poldrack et al. 2016). We include a summary of the dataset for completeness; details may be found in the accompanying article.

The dataset includes four groups of subjects: healthy controls, schizophrenia patients, bipolar disorder (Type I) patients, and persons with attention deficit hyperactivity disorder (ADHD). Subjects were recruited from the wider Los Angeles area through a combination of community advertisements and patient-oriented clinical outreach. Final sample demographics span both men and women, ages 21 to 50, with no less than eight years of formal education and who primarily speak English or Spanish. Applicants with substantial medical illness or which failed a urinary test for drugs of abuse (cocaine, methamphetamines, morphine, THC, or benzodiazepines) were excluded. Previous diagnosis for schizophrenia or other psychotic disorders, bipolar disorder, or substance abuse and dependence (nicotine and caffeine excepted) were grounds for exclusion from the control group, as was current diagnosis for major depressive disorder, attention deficit hyperactivity disorder (ADHD), suicidality, or an anxiety disorder (obsessive compulsive disorder, panic disorder, generalized anxiety disorder, post-traumatic stress disorder). Control applicants displaying sub-threshold ADHD or who had undergone medical treatment for ADHD in the preceding twelve months were also excluded.

Patients were diagnosed following a Structured Clinical Interview for the Diagnostic and Statistical Manual of Mental Disorders, Fourth Edition-Text Revision (DSM-IV) to ensure that diagnoses followed DSM-IV categories. The Adult ADHD Interview was also employed to fully characterize the lifetime history of ADHD in adults. Interviewers were trained according to , with minimum acceptable symptom agreement consisting of overall $\kappa \geq 0.75$, specificity of $\kappa \geq 0.75$, sensitivity of $\kappa \geq 0.75$, and diagnostic accuracy of $\kappa \geq 0.85$. Interviewer skill was measured according to the SCID Checklist of Interviewer Behaviors and the Symptom Checklist of Interviewer Behaviors , and no interviewer displayed an annual $\kappa < 0.75$ during the study. Patient groups were mutually exclusive, i.e. subjects which suffered from two or more of the disorders under study were removed. Stable medication was permitted. All participants gave written informed consent according to procedures approved by the institutional review boards of the University of California, Los Angeles and the Los Angeles County Department of Mental Health.

The resting-state fMRI scan lasted 304 seconds, with participants asked to remain relaxed with eyes open. No stimuli were presented during this scan. MRI data were acquired on two 3T Siemens Trio scanners, one at the Ahmanson-Lovelace Brain Mapping Center (Siemens version syngo MR B15) and one at the Staglin Center for Cognitive Neuroscience (Siemens version syngo MR B17) at the UCLA. Functional MRI used a T2*-weighted echoplanar imaging sequence with the parameters: slice thickness = 4mm, 34 slices, TR = 2s. TE = 30ms, flip angle = 90°, 64 × 64 matrix, 192mm field of view, with an oblique slice orientation. Echo-planar imaging was again used to acquire diffusion-weighted imaging data with parameters: slice thickness = 2mm, 64 directions, TR = 9000ms. TE = 93ms, flip angle = 90°, 96 × 96 matrix, axial slice orientation, $b = 1000 \frac{\text{s}}{\text{mm}^2}$, one average. An MPRAGE and a T2-weighted matched-bandwidth high-resolution anatomical scan were also collected using the same slice prescription as the fMRI scan. High-resolution scan parameters were: slice thickness = 4mm, TR = 5000ms. TE = 34ms, flip angle = 90°, 128 × 128 matrix, four averages. MPRAGE parameters were: slice thickness = 1mm, 176 slices, TR = 1.9s. TE = 2.26ms, 256 × 256 matrix, 250mm field of view, slices in the sagittal plane. No MRI data were collected for left-handed subjects, subjects who might be pregnant, had metal in the body, or were otherwise unable to safely enter an MRI machine.

Of 272 initial participants, we received resting-state fMRI data for 260, of which 121 were healthy controls, 50 schizophrenia patients, 49 bipolar disorder patients, and 40 in treatment for ADHD. The data for six subjects were found to be unusable, reducing the final numbers to 118 healthy controls, 48 schizophrenia patients, 49 bipolar disorder patients, and 39 ADHD patients. fMRI data were segmented according to the Desikan-Killiany atlas integrated with subcortical areas (Desikan et al. 2006b), and the standard Desikan-Killiany structural connectivity template used as a connectivity mask.

4.3 Summary

The datasets employed in this thesis cover multiple psychiatric disorders with a wide range of symptoms. So wide a distribution

offers a thorough test for any analysis pipeline, as it much be able to both detect and distinguish multiple symptoms and causal mechanisms. The application of LEICA to this task, and the results of said applications, will be the topic of the following chapters.

5 Extracting Entropy

5.1 Reasoning

It has been hypothesized that the dynamic resting-state networks may operate somewhat like the proposed independent channels which Shannon imagined when envisioning his measure of entropy, with each acting largely independently of the others. Under this hypothesis, each substate may be considered a separate channel (or, equivalently, processing unit) of the resting state signal. This line of logic results in two immediate corollaries:

- It should be possible to separate the signals from each channel (source reconstruction)
- The signals from each channel should be (partially) independent

From these corollaries spring two immediate consequences:

- It should be possible to estimate the complexity of each processing unit's activity
- The complexity of the entire fMRI signal may be estimated as a function of these (semi-)independent signals

Given the mathematical equivalence of complexity to information, it should be possible to use metrics from information theory to quantify the complexity of biological signals. The most general such metric, the Shannon entropy, also has the advantage of being mathematically tractable for a multichannel signal. Assuming statistically independent channels, the Shannon entropy of a multichannel signal is simply the sum of the entropies of each channel.

5.2 Pipeline

5.2.1 Functional Connectivity

We elected to calculate the dynamic functional connectivity (dFC) using Coherence Connectivity Dynamics (Glerean et al. 2012; Ponce-Alvarez et al. 2015; Deco and Kringelbach 2016). Unlike sliding window correlation, Coherence Connectivity Dynamics can achieve temporal resolution of a single time to repetition and has no free parameters to set, making it both simpler and more precise than

sliding window correlation. It achieves this by passing each region's time series through the Hilbert transform, which outputs the analytic representation of the BOLD signal. As the analytic representation is a complex signal, it can be described by a time-varying amplitude multiplied by a time-varying phase angle:

$$s_a(t) = s_m(t)e^{i\varphi(t)}$$

where $s_m(t) \equiv |s_a(t)|$ and $\varphi(t) \equiv \text{arg}(s_a(t))$, with the argument being defined by Euler's formula

$$e^{i\varphi} = \cos \varphi + i \sin \varphi$$

Instantaneous interregional coherence is estimated as the difference between regional time-dependent phases. More specifically, phase coherence between regions m and n at time t ($\text{dFC}(m, n, t)$) is

$$\mathbf{dFC}(m, n, t) = \cos(\theta(m, t) - \theta(n, t))$$

with \cos (the cosine function) normalizing the value into the range $[-1, 1]$. The result is an $N \times N \times T$ dFC array, with N being the number of ROIs and T being the number of time points. $\cos(\theta)$ being an even function, each $N \times N$ matrix $\mathbf{dFC}(t)$ is symmetric.

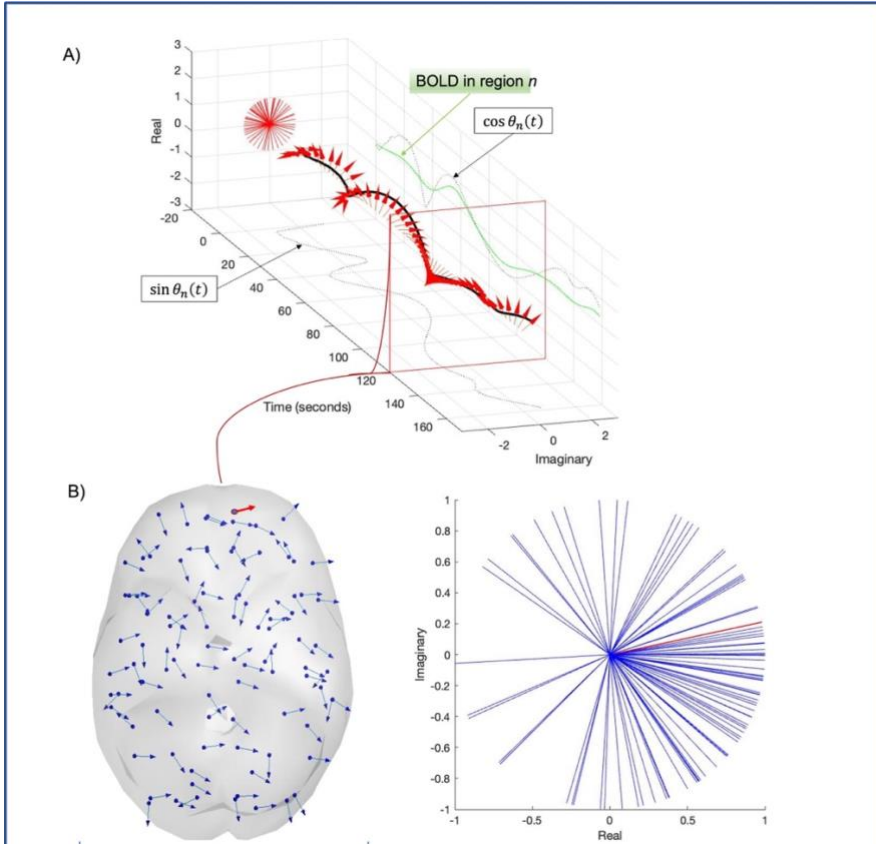


Figure 2 | computing the phase of the BOLD signal. Each regional time series (green trace) is converted into an analytic signal using the Hilbert transform. Euler's formula converts this analytic signal into a time-resolved phase signal (A) with both real and imaginary parts (dashed black traces). This process is repeated for all brain regions (B) at across time points.

5.2.2 Eigendecomposition

The fundamental goal of the LEiDA process is to project the dominant spatial connectivity pattern dynamics into a lower-dimensional space for ease of analysis. This dominant pattern is easily identifiable at each time point because functional connectivity matrices are always both real and symmetric. Such matrices are always diagonalizable, so the dFC at any time point t can be decomposed into

$$dFC(t) = VDV^{-1}$$

with V being the eigenvectors of $dFC(t)$ and D the diagonal matrix of eigenvalues. As the eigenvectors of a symmetric matrix must be orthogonal, $V^{-1} = V^T$; thus,

$$dFC(t) = VDV^T$$

which may be equivalently written as

$$dFC(t) = VDV^T = \sum_n \lambda_n v_n v_n^T$$

where v_n is the n^{th} eigenvector and λ_n the n^{th} eigenvalue of $dFC(t)$. At each time point, the instantaneous FC matrix may be decomposed into a weighted sum of eigenvector outer products $v_n v_n^T$ weighted according to the respective eigenvalue λ_n . Thus, finding the dominant spatial pattern at any time point simply involves finding the eigenvector with the largest eigenvalue at that time point. In addition, one may easily compute the proportion of variance which this pattern captures simply by dividing the leading eigenvalue by the sum of all eigenvalues:

$$\rho = \frac{\lambda_l}{\sum_n \lambda_n}$$

Previous work demonstrates that the leading eigenvector consistently represents more than 50% of data variance (Cabral et al. 2017; Lord et al. 2019), a finding confirmed in the present thesis. Further, experiments with the use of additional eigenvectors demonstrated no improvement in performance or clinical interpretability. The author thus believe that a single eigenvector is sufficient to represent functional connectivity dynamics.

This compression has three distinct advantages for signal analysis. First, compressing each $N \times N$ $dFC(t)$ matrix to an $N \times 1$ vector p_1 reduces sample dimensionality from $\frac{N(N-1)}{2}$ to N . Second, the primary connectivity pattern should contain no noise, as noise components generally appear in trailing eigenvectors. Finally, previous work in spectral community detection (Newman 2006; Leicht and Newman 2008) demonstrated that the leading eigenvector $p_1(t)$ separates network nodes into communities based on the sign of each node $r \in p_1(t)$, with the magnitude of r indicating that assignment's "strength". Thus, transforming the $dFC(t)$ matrix to $p_1(t)$ converts interregional phase-locking values into regional community assignment values. The leading eigenvector of an FC matrix naturally separates networks into two mutually opposing communities.

The Leading Eigenvector Dynamics Analysis framework (LEiDA) (Cabral et al. 2017; Figueroa et al. 2019; Lord et al. 2019) examines only the leading eigenvector $\mathbf{v}_1(t)$ of each $N \times N$ $\mathbf{dFC}(t)$ matrix. At each time point, the leading eigenvector of the $N \times N$ $\mathbf{dFC}(t)$ is extracted; once the leading eigenvectors of all time points have been extracted, they are concatenated horizontally to form a space-time matrix \mathbf{E} . Each row r of \mathbf{E} represents one brain region r , and each column t contains the leading eigenvector $\mathbf{v}_1(t)$ for time t . The laws of linear algebra render $\mathbf{v}_1(t)$ and $-\mathbf{v}_1(t)$ equivalent, so we follow the convention that most elements in each eigenvector should be negative (Figueroa et al. 2019).

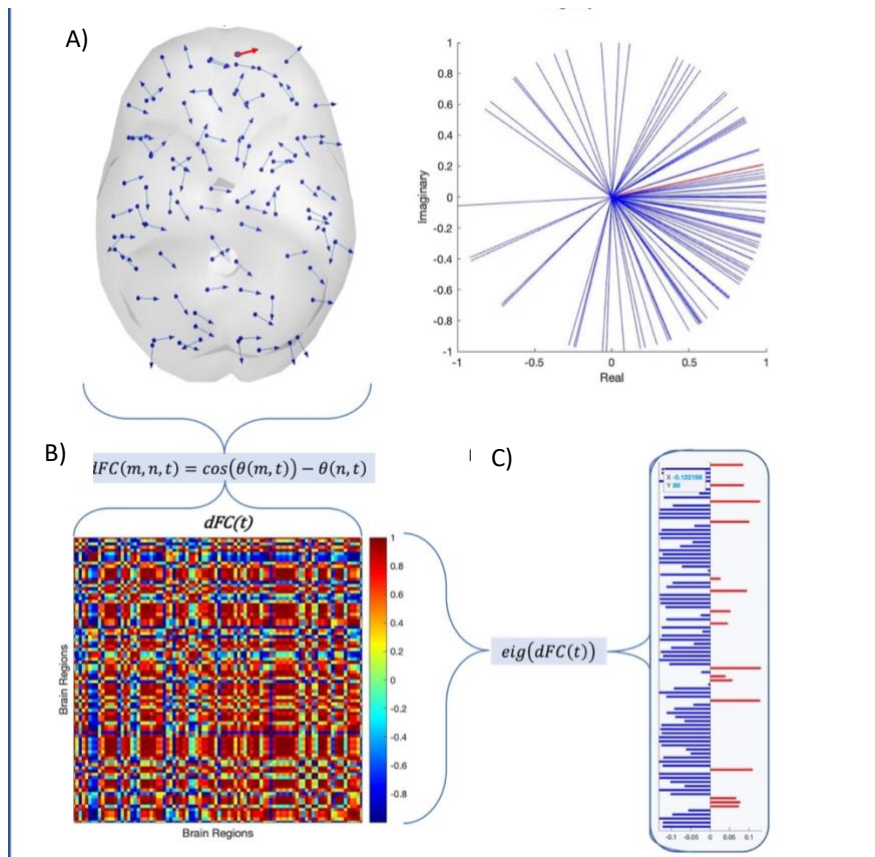


Figure 3 | Eigendecomposition. After computing the phases of each region at time t (A), the cosine distance between each pair of regions is computed to produce an instantaneous functional connectivity matrix (B). The leading eigenvector \mathbf{V}_1 of this functional connectivity matrix is then isolated (C). Repeating this process across time points results in a 2-D array of eigenvectors, each representing a distinct time point in the data.

5.2.3 Component Detection

Finding the communities that recur above chance requires determining a significance threshold for regional co-activation. Although surrogate methods, e.g. a permutation test, can establish such a threshold, they are slow and computationally intensive. A far cheaper and more elegant method is available by leveraging based on autocorrelation matrix eigenvalues (Peyrache et al. 2009; 2010). It has been established for several decades that if an $m \times n$ matrix \mathbf{M} has statistically independent rows (as would be expected for uncoupled noisy oscillators), the eigenvalues of its autocorrelation matrix follow the Marčenko-Pastur distribution (Marčenko and Pastur 1967). Crucially, this distribution has analytically tractable limits

$$\lambda_{\min}^{\max} = \sigma^2 \left(1 \pm \sqrt{\frac{1}{q}} \right)^2$$

where σ is the standard deviation of \mathbf{M} and $q \equiv \frac{n}{m} \geq 1$. Thus, if communities do not recur over time, the eigenvalues of \mathbf{E} 's correlation matrix should lie within the limits imposed by λ_{\min}^{\max} . Conversely, should any communities of \mathbf{E} recur at a rate significantly above chance, a corresponding number of eigenvalues of the correlation matrix of \mathbf{E} should exceed the upper limit λ_{\max} . This method has been validated in the spike activity context (Lopes-dos-Santos et al. 2011; Lopes-dos-Santos, Ribeiro, and Tort 2013) and in a previously published fMRI study (Deco, Cruzat Grand, and Kringelbach 2019).

5.2.4 Component Extraction

Upon finding the total number of recurrent communities with the Marčenko-Pastur distribution, we utilize the fastICA algorithm (Laubach, Shuler, and Nicolelis 1999; Hyvärinen and Oja 2000) to extract these communities and their activity time courses from the matrix \mathbf{E} . Since the fastICA algorithm requires the user to manually specify the number of independent components, the Marčenko-Pastur distribution threshold is crucial to providing an objective, data-driven metric for the number of components.

After computing \mathbf{E} 's covariance matrix, twelve (12) eigenvalues were found to surpass the Marčenko-Pastur upper bound. ICA was then run to extract these twelve distinct and temporally independent components. As fastICA can only extract the magnitude of an independent component, not its sign, the spatial map's positive and negative signs should be understood to represent relative orientations rather than absolute weights.

Why select independent component analysis to extract functional connectivity states? ICA minimizes the statistical dependencies between its output components, a fact which has been confirmed in the dynamic functional imaging context (Calhoun et al. 2013). Thus, running ICA along the time dimension of dynamic functional imaging data may be understood as a search for connectivity states which display no temporal relationships between one another. This should completely—or at least almost completely—prevent the temporal dependencies between components, which greatly simplifies calculation of the entropy.

But how, precisely, does independent component analysis function? It is a relatively new method which seeks to solve an old problem: the problem of blind source separation. We can turn to the classic example of the cocktail party problem: given a group of perhaps twenty persons in a room, all speaking at once, how can one identify and extract a single conversation of interest from the overall noise?

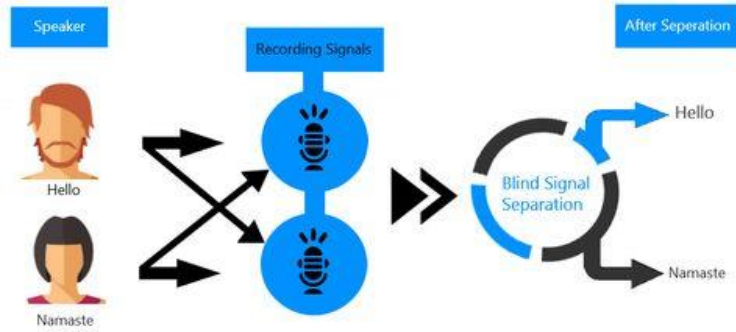


Figure 4: the cocktail party problem: given two or more signals which have been mixed, how can one isolate the original signals?

As is often the case in mathematical analysis, one must start with assumptions. The first assumption is that the signals are mixed linearly, i.e. they combine via simple addition:

$$\mathbf{x} = \mathbf{A}\mathbf{s}^T$$

where \mathbf{x} are the received signals and \mathbf{s} are the original signals (which we hope to recover). The second assumption is that the signals \mathbf{s} are statistically independent, i.e. knowledge of any one signal provides no information on the content of the others. Given an absence of additional knowledge, independent component analysis (ICA) must satisfy the following conditions:

1. it must be a linear model
2. the signals it isolates must be statistically independent
3. it must require no inputs aside from the mixed signals \mathbf{x} and the number of original signals \mathbf{s}

These constraints enforce a linear mixture model

$$\mathbf{x} = \mathbf{A}\mathbf{s}^T = \sum_{j=1}^n \mathbf{a}_j s_j$$

with variables

$\mathbf{s} = [\mathbf{s}_1, \dots, \mathbf{s}_n]$	n original, independent components (S_n)
--	---

$\mathbf{A} = \begin{bmatrix} \mathbf{a}_{11} & \cdots & \mathbf{a}_{1n} \\ \vdots & \ddots & \vdots \\ \mathbf{a}_{m1} & \cdots & \mathbf{a}_{mn} \end{bmatrix}$	mixed according to the coefficients (a_{ij})
$\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_m]$	to produce m observed linear mixtures (x_m)

Of course, ICA does not want to find the mixing model; its goal is to find the *unmixing*, or *inverse*, model, which should satisfy

$$\mathbf{s} = \mathbf{W}\mathbf{x}^T = \sum_{j=1}^n \mathbf{w}_j x_j$$

Solving this problem is equivalent to finding the inverse matrix \mathbf{W} , which might seem obvious; setting $\mathbf{W} = \mathbf{A}^{-1}$ would suffice:

$$\begin{aligned} \mathbf{x} &= \mathbf{A}\mathbf{s} \\ \mathbf{A}^{-1}\mathbf{x} &= \mathbf{A}^{-1}\mathbf{A}\mathbf{s} = \mathbf{s} \end{aligned}$$

The problem is that we do not know \mathbf{A} , so inverting it is not a simple application of linear algebra. Instead, one must estimate it based on the available constraints and a suitable cost function.

It is here that one encounters—again—the Central Limit Theorem. Briefly, the Central Limit Theorem states that any linear combination of independent variables will tend to form a normal distribution, *regardless of the distribution of the original variables*. Thus, for any original signal s_k , unless s_k is Gaussian white noise—that is unless it follows a normal distribution—then the combined signal \mathbf{x} will more strongly resemble a normal distribution than s_k does. More importantly, any reconstructed signal \mathbf{y} will be closer to a normal distribution than s_k unless $\mathbf{y} = s_k$ (or unless s_k follows a normal distribution itself!).

Can this cost function help us estimate the inverse mixing matrix \mathbf{W} ? To answer that question, we can engage in some slight-of-hand by introducing a new variable \mathbf{z} :

$$\mathbf{z} = \mathbf{A}^T \mathbf{W}$$

$$y = \mathbf{W}^T \mathbf{x} = \mathbf{W}^T \mathbf{A} \mathbf{s} = \mathbf{z}^T \mathbf{s} = \sum_j z_j s_j$$

where y is the estimate of the k^{th} signal s_k . By the Central Limit Theorem, y is more Gaussian than $\mathbf{z}^T \mathbf{s}$ unless $y = s_k$. But $y = s_k$ is only true if $z_k = 1$ and $z_{j \neq k} = 0$, which is only possible if $\mathbf{W}^T \mathbf{A} = \mathbf{I}$, i.e. if $\mathbf{W} = \mathbf{A}^{-1}$. Thus, by minimizing the Gaussianity of $\mathbf{z}^T \mathbf{s}$, we can in principle recover the original signal s_k . Better yet, there's no need to introduce the variable \mathbf{z} at all; since $\mathbf{W}^T \mathbf{x} = \mathbf{z}^T \mathbf{s}$, minimizing the Gaussianity of $\mathbf{W}^T \mathbf{x}$ implies minimizing the Gaussianity of $\mathbf{z}^T \mathbf{s}$, so we simply need to find a \mathbf{W} such that $\mathbf{W}^T \mathbf{x}$ is maximally non-Gaussian.

Thus, one approach to solving the source separation problem is to establish a linear model

$$y = \mathbf{W}^T \mathbf{x}$$

and minimize the Gaussianity of y . This leads to another question: how should we determine how Gaussian or non-Gaussian our variable has become? In the language of machine learning, what cost function should we employ for our algorithm?

Two primary metrics for Gaussianity exist: the *excess kurtosis* and the *negentropy*. Given a probability distribution y , excess kurtosis is defined as the expected value of y 's fourth moment minus three times the expected value of its second moment:

$$kurt(y) = E\{y^4\} - 3E\{y^2\}$$

This formulation is computationally straightforward and has the interesting property of equalling zero if and only if y follows a Gaussian distribution. If y is flatter than a Gaussian distribution, the excess kurtosis falls below zero; if y is more strongly peaked than a Gaussian, the excess kurtosis rises above zero. Thus, the kurtosis fulfills two necessary features for a cost function: it is easily estimated, and it is easily interpretable.




Figure 5: Kurtosis

Unfortunately, kurtosis has one highly undesirable property for a cost function; it is not robust to outliers. Given the fact that real data seldom confirms perfectly to a specified probability distribution, this renders it unsuitable for independent component analysis.

The second potential Gaussianity metric is known as the *negentropy*, which is defined as the difference in entropy between y and a Gaussian variable y_{gauss} with a covariance matrix identical to y :

$$J(\mathbf{y}) = H(\mathbf{y}_{\text{Gauss}}) - H(\mathbf{y})$$

where \mathbf{y} is the variable of interest, $\mathbf{y}_{\text{Gauss}}$ is a Gaussian variable with same covariance matrix as \mathbf{y} , and $H(\mathbf{y})$ is the entropy of \mathbf{y} . In principle, this is an ideal cost function; highly robust and zero if and only if y follows a Gaussian distribution exactly. Unfortunately, estimating y 's entropy requires estimating y 's probability distribution, which is seldom practical. Thus analysts use approximations. Hyvärinen and Oja—the authors of the fastICA algorithm—chose a somewhat non-obvious formulation:

$$J(\mathbf{y}) \approx \sum_{j=1}^p k_j [E\{G_j(\mathbf{y})\} - E\{G_j(\mathbf{v})\}]^2$$

with k_j being arbitrary constants ($k_j \geq 0$ for all j), G_j being a non-quadratic function, and \mathbf{v} being a Gaussian variable of zero mean and unitary variance (Hyvärinen 1999; Hyvärinen and Oja 2000).

The authors explain that while this is not the classical estimate of the negentropy, it does have advantages of consistency and robustness given appropriate choices for G_j :

$$G_1(u) = \frac{1}{a} \ln(\cosh(v)) \text{ for } 1 \leq a \leq 2$$

$$G_2(u) = -e^{-\frac{u^2}{2}}$$

Hyvärinen and Oja found that, after whitening \mathbf{y} , these selections resulted in a $J(\mathbf{y})$ which was highly robust to outliers, universally nonnegative ($J(\mathbf{y}) \geq 0$ for all \mathbf{y}), and zero if and only if \mathbf{y} follows a Gaussian distribution—which is to say, an ideal cost function.

Having identified a suitable cost function, the next challenge is to find a suitable optimization algorithm. After all, a traveler's ability to measure how far he is from his goal is of little use unless he can also determine which direction he must go. This process can be simplified by centering and whitening our observed data \mathbf{x} : that is, by removing its mean $\mathbf{m} = E\{\mathbf{x}\}$ from \mathbf{x} , by removing correlations between the component variables, and by rescaling component variables to achieve a unitary variance:

$$E\{\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T\} = \mathbf{I}$$

$\tilde{\mathbf{x}}$ being the whitened, demeaned version of \mathbf{x} .

These alterations are well-defined in linear algebraic literature and are easily reversed. Thus, any alteration in the estimated source signals \mathbf{y} are easily compensated at the end of optimization.

Having prepared the data, our next step is to find in which direction to take it. The goal is to find a transformation \mathbf{w} of \mathbf{x} which minimizes the Gaussianity of $\mathbf{w}^T \mathbf{x}$, i.e. which maximizes $J(\mathbf{w}^T \mathbf{x})$. The maxima of any function will be a fixed point, i.e. a point where the derivative $\frac{J(\mathbf{w}^T \mathbf{x})}{\mathbf{w}} = 0$; thus, the goal is to find a \mathbf{w}' which satisfies $\frac{J(\mathbf{w}'^T \mathbf{x})}{\mathbf{w}} = 0$. To do this, Hyvärinen and Oja propose the following algorithm:

- Choose initial (random) weight vector \mathbf{w}_0
- Set $g \equiv \vec{\nabla}(G)$

- Set the update rule:
 - At each iteration n , define

$$\mathbf{w}^+ = E\{\mathbf{x}g(\mathbf{w}_{n-1}^T\mathbf{x})\} - E\{g'(\mathbf{w}_{n-1}^T\mathbf{x})\}\mathbf{w}_{n-1}$$
 - Set $\mathbf{w}_k = \frac{\mathbf{w}^+}{\|\mathbf{w}^+\|}$
- Repeat until convergence, i.e. $\mathbf{w}_{k-1} \cdot \mathbf{w}_k \approx 1$.

The proof for this algorithm is provided in (Hyvärinen and Oja 2000).

Well and good, provided we only want to find a single original component $y = \mathbf{w}^T\mathbf{x}$. In general, however—and certainly in the human brain—the goal is to find multiple components. This raises a problem: re-running the iteration algorithm described above may produce duplicate results. Clearly, it is of little use to extract the same \mathbf{x} several times; one must find a means to prevent this. The solution presented by Hyvärinen and Oja is to decorrelate each output after each iteration. This can be accomplished via another iterative algorithm (Hyvärinen 1999):

- Given n desired outputs $\mathbf{w}_1 \dots \mathbf{w}_n$, define $\mathbf{W} \equiv (\mathbf{w}_1, \dots, \mathbf{w}_n)^T$
- Set $\mathbf{W} = \frac{\mathbf{W}}{\sqrt{\|\mathbf{W}\mathbf{W}^T\|}}$
- Iterate Step 2 until convergence: $\mathbf{W} = \frac{3}{2}\mathbf{W} - \frac{1}{2}\mathbf{W}\mathbf{W}^T\mathbf{W}$

This step completes the optimization algorithm to find \mathbf{w} such that $y = \mathbf{w}^T\mathbf{x} \approx \mathbf{s}$.

5.2.5 Entropy Estimation

Since independent component analysis forbids statistical dependencies between its component outputs, the joint entropy over all components is simply the sum of the individual components' Shannon entropies (Cover and Thomas 2005):

$$H(C_1, \dots, C_N) = \sum_{j=1}^N H(C_j)$$

It is possible to compute the joint entropy of each subject by computing the Shannon entropy of each component's activation time series and summing them. This allows the construction of a distribution of subject joint entropies, which can then be analyzed for group-level differences.

There is, of course, the problem of estimating the entropy of each individual component. Dr. Shannon, in his seminal article, provides a definition of entropy for continuous variables:

$$H(x) = - \int p(x) \ln p(x)$$

where $p(x)$ is the probability density function of the continuous variable x . However, this definition is seldom practical as, in most cases, researchers cannot access $p(x)$. This led to considerable research into the problem of estimating $p(x)$ based on a finite sample of x (Dudewicz and van der Meulen 1981; Gorla et al. 2005; Delattre and Fournier 2017). This thesis employs a variant of the Kozachenko and Leonenko estimator (Singh et al. 2003) as implemented by Zoltan Szabó in the Information Theoretical Estimators toolbox (Szabó 2014).

5.2.6 Comparisons

Group-level differences are examined via a difference-of-means permutation test (Krol 2021) with 10,000 permutations, and provide multiple-comparison correction via the false discovery rate (Benjamini and Hochberg 1995). The Bonferroni (Bonferroni 1935) and Šidák (Šidak 1967) thresholds verify these results.

5.3 Results

5.3.1 Functional Connectivity

5.3.1.1 Study I: Obsessive-Compulsive Disorder

Both control and patient time series are parcellated according to the AAL atlas (Tzourio-Mazoyer et al. 2002). Each subject's dynamic functional connectivity is computed using Coherence Connectivity Dynamics (Deco, Cabral, et al. 2017). Analysis is restricted to the cortical and subcortical regions; as such, the 26 cerebellar regions of the AAL atlas are discarded. The resultant three-dimensional

array must be converted into two dimensions for further analysis. Three methods are tested. In the first method, we extract the leading eigenvector (90×1) of each time point's connectivity matrix. The eigenvectors of all time point are then concatenated to form a subject-level 90×175 eigenvector time series **E**. In the second method, each time point's connectivity matrix is averaged horizontally, and the resulting average coherence vectors (90×1) are concatenated to form a subject-level 90×175 mean coherence time series **M**. Finally, each time point's connectivity matrix is vectorized to form a 4005×1 connectivity vector, and these vectors are again concatenated to form a subject-level 4005×175 **dFC** time series (as each connectivity matrix is symmetric and the main diagonal neglected, only the upper triangle is vectorized).

5.3.1.2 Study II: Multiple Disorders

We elect to define the basis based on the control group. In order to find the number of necessary basis vectors, a “global” time series is constructed by concatenating all controls along the time dimension. The autocorrelation matrix of this global signal undergoes eigendecomposition; the number of eigenvalues which surpass the upper bound of the Marčenko-Pastur distribution (Marčenko and Pastur 1967) is the number of dimensions necessary to describe the nonrandom activity of the control time series. In this study, fourteen (14) dimensions were found to be sufficient.

Upon determining the number of dimensions necessary, ICA converts the 68-dimensional control time series T_C into its fourteen-dimensional representation A_C (Lopes-dos-Santos et al. 2011; Lopes-dos-Santos, Ribeiro, and Tort 2013). The time series of the three patient groups are mapped to this fourteen-dimensional space via the mixing matrix **W**. No additional formats are utilized: all data are analyzed using the LEiDA framework.

5.3.2 Functional Dimensions

5.3.2.1 Study I: Obsessive-Compulsive Disorder

To determine the number of dimensions necessary, all subjects' time series are concatenated and the autocorrelation matrix of this global time series array calculated. The number of significant

dimensions is then the number of eigenvalues in the autocorrelation matrix which exceed the upper bound of the Marčenko-Pastur distribution (Marčenko and Pastur 1967). Applying this method to the eigenvector time series \mathbf{E} identifies twelve independent dimensions across the resting state of all subjects. ICA can then convert the 90-dimensional eigenvector time series \mathbf{E} into its twelve-dimensional representation \mathbf{T}_E (Lopes-dos-Santos et al. 2011; Lopes-dos-Santos, Ribeiro, and Tort 2013). Repeating this process for the vectorized \mathbf{dFC} produces the 347-dimensional representation \mathbf{T}_F , and the spatially averaged \mathbf{M} produces the eleven-dimensional representation \mathbf{T}_M .

5.3.2.2 Study II: Multiple Disorders

We elect to define the basis based on the control group. In order to find the number of necessary basis vectors, a “global” time series is constructed by concatenating all controls along the time dimension. The autocorrelation matrix of this global signal undergoes eigendecomposition; the number of eigenvalues which surpass the upper bound of the Marčenko-Pastur distribution (Marčenko and Pastur 1967) is the number of dimensions necessary to describe the nonrandom activity of the control time series. In this study, fourteen (14) dimensions were found to be sufficient.

Upon determining the number of dimensions necessary, ICA converts the 68-dimensional control time series T_C into its fourteen-dimensional representation A_C (Lopes-dos-Santos et al. 2011; Lopes-dos-Santos, Ribeiro, and Tort 2013). The time series of the three patient groups are mapped to this fourteen-dimensional space via the mixing matrix \mathbf{W} .

5.3.3 Joint Entropy

5.3.3.1 Study I: Obsessive-Compulsive Disorder

Since each time series in the low-dimensional space is statistically independent, each dimension’s Shannon entropy may be calculated (Singh et al. 2003; Delattre and Fournier 2017) independent of the others’. Computing the subject-level Shannon entropy of each substate’s time series results in a $D \times S$ array of entropy values for patients and controls, where S is the number of dimensions and S

the number of subjects per group. This format means that computing the subject-level joint entropy simply requires summing along this array’s first dimension. This produces two $1 \times S$ joint entropy distributions, which can be compared with any standard statistical test. Applying this process to eigenvector-based entropy scores again shows elevated entropy in patients relative to controls ($p = 0.0119$, Hodges’ $G = -0.5833$); controls display a joint entropy of 14.5695 ± 1.2473 , while patients display a mean joint entropy of 15.2214 ± 1.1535 . However, the joint entropy distributions of \mathbf{T}_F and of \mathbf{T}_M display no significant group-level differences. Eigenvector-based analysis thus appears to preserve the information of the full signal while reducing dimensionality almost thirtyfold—a crucial consideration, as the curse of dimensionality states that patterns become exponentially harder to detect as dimensionality increases.

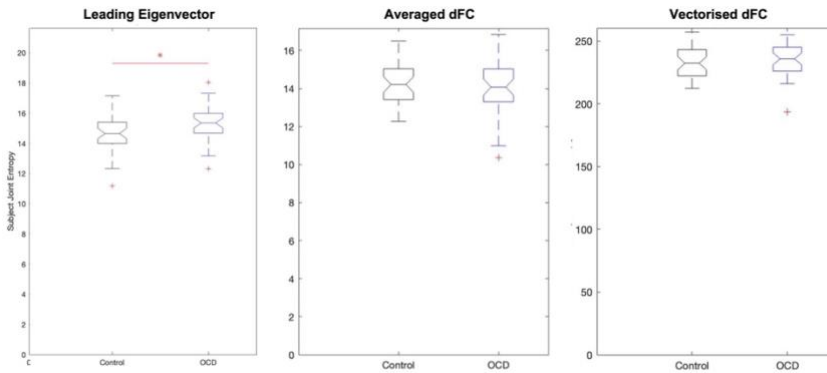


Figure 6: Analysis of eigenvector-based component time series (T_E) shows that obsessive-compulsive patients display significantly higher joint entropy than age-, gender-, and education-matched controls. Notably, only the

5.3.3.2 Study II: Multiple Disorders

In the control basis space, the statistical independence of each basis vector means that each subject’s joint entropy is simply the sum of its dimensional entropies. Thus, by independently estimating Shannon entropy of subject activities along each dimension (Singh et al. 2003; Delattre and Fournier 2017) and summing over the rows of the resulting $D \times S$ matrix (D being the number of control-based basis vectors and S being the number of subjects), we obtain a $1 \times S$ joint entropy distribution which can be examined by standard

statistical tests. Applying this procedure to the current dataset reveals that the control group displays significantly lower joint entropy than any of the patient groups (Figure 7 and Table 1), but no patient group displays changes in joint entropy relative to another patient group. Elevated levels of randomness (variability) in connectivity time series thus appear to be a core characteristic of several psychiatric disorders.




Figure 7: Group-level comparisons of the joint entropy. Controls display significantly lower joint entropy than any patient group, while no such difference is apparent between patient groups.

On average, controls display a joint entropy of 18.2676 ± 0.9845 . By contrast, schizophrenia patients display a mean joint entropy of 19.0099 ± 1.0344 , bipolar patients a mean of 18.8994 ± 1.1336 , and ADHD patients a mean of 18.7945 ± 1.0140 . After multiple comparison correction, patient groups display no significant differences in joint entropy between one another.

Table 1: *p*-values for pairwise joint entropy comparisons. Controls display significantly lower entropy than any patient group. No such difference is visible between patient groups.

Comparison	<i>p</i> -value	Hodges' G	Significant?
<i>Controls vs. schizophrenia</i>	9.999×10^{-5}	-0.7430	True
<i>Controls vs. bipolar</i>	2.9997×10^{-4}	-0.6134	True
<i>Controls vs. ADHD</i>	0.0059	-0.5313	True
<i>Schizophrenia vs. bipolar</i>	0.6224	0.1017	False
<i>Schizophrenia vs. ADHD</i>	0.3322	0.2101	False
<i>Bipolar vs. ADHD</i>	0.6586	0.0970	False

5.3.4 Dimension-Specific Entropy

5.3.4.1 Study I: Obsessive-Compulsive Disorder

To determine whether alterations in entropy concentrate in specific dimensions, start with the same $D \times S$ patient and control arrays of entropy values as the previous section. Each row of these arrays was compared and corrected with the false discovery rate. As above, this analysis was run for all three compression methods: \mathbf{T}_E (eigenvectors), \mathbf{T}_F (uncompressed), and \mathbf{T}_M (spatial average). Only the eigenvector-based representation (\mathbf{T}_E) detects a significant alteration along any dimension, specifically the first (ordered according to mean activity). This dimension consists of paired anticorrelated communities and both display significantly higher entropy in patients than in age-, gender-, and education-matched controls.

In component space, one LEICA component displays higher entropy in patients than in controls (1.1818 ± 0.1401 , 1.3075 ± 0.2276 , $p = 0.0020$, Hedges' $g = -0.6634$). This substate consists of two opposing communities, with the sign of each brain region denoting to which community that region belongs and the magnitude of that region's weight denoting the strength of its association with that community. Concentrating on regions with absolute z -scores above 1.3 ($|z| > 1.3$) (Figure 8). Under these constraints, the first community contains the left precentral gyrus, left and right frontal superior cortex (orbital), left middle frontal gyrus (orbital), the left inferior frontal gyrus (opercular), left cuneus, right olfactory bulb, and right inferior parietal gyrus. Its opposite number includes the right lingual gyrus, right occipital medial gyrus, right putamen, right pallidum, left amygdala, right middle temporal gyrus, and right temporal pole of the middle temporal gyrus (Figure 8). This result survives both the false discovery rate and the Sidak multiple comparison correction.

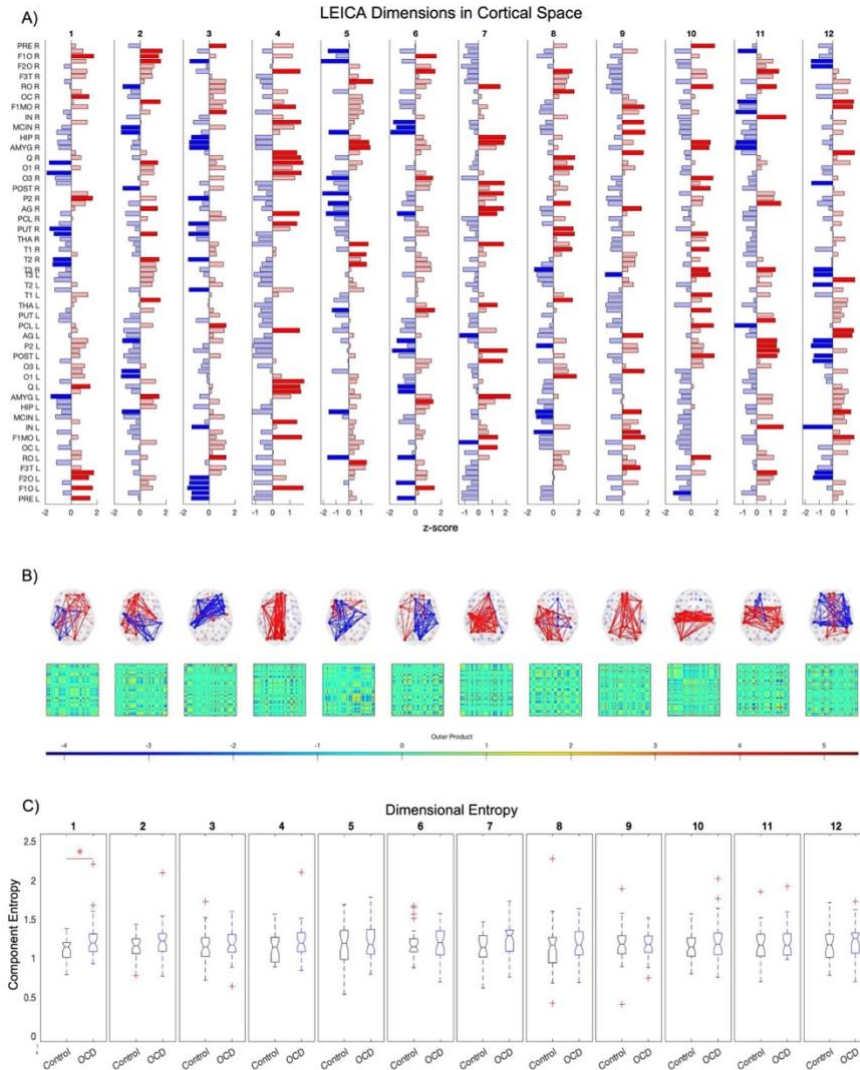


Figure 8 / Twelve of the eigenvalues of \mathbf{E} 's autocorrelation matrix exceed the upper limit of the Marčenko-Pastur distribution, suggesting that twelve dimensions are necessary to capture \mathbf{E} 's activity. Independent component analysis reveals how these dimensions map to brain regions (A). Map weights have been converted into z-scores for this figure and regions with a weight $z < 1.3$ are depicted in faded color. Plotting these mapping vectors in the brain and as connectivity (B) reveals that the trailing dimensions (nine, ten, eleven, and twelve) display notable homotopic symmetry, while leading dimensions are strongly asymmetric. Finally, group-level entropy analysis shows that the first dimension displays significantly higher entropy in obsessive-compulsive patients than in controls (C). Note that dimensions are ordered according to average activity level across the dataset.

Table 2 | High-Weighted Regions in OCD

Component 1 ($z > 1.3$)	
Positive	Negative
L Precentral Gyrus	L Amygdala
L Superior Frontal Gyrus, Orbital Part	R Temporal Pole: Middle Temporal Gyrus
L Middle Frontal Gyrus, Orbital Part	R Middle Temporal Gyrus
L Inferior Frontal Gyrus, Opercular Part	R Lenticular Nucleus, Pallidum
L Cuneus	R Lenticular Nucleus, Putamen
R Inferior Parietal Gyri	R Middle Occipital Gyrus
R Olfactory Cortex	R Lingual Gyrus
R Superior Frontal Gyrus, Orbital Part	

Table 2 | displays the regions of the first dimension with absolute z -scores exceeding 1.3 ($|z| > 1.3$). The sign of each regional weight indicates to which of two communities it belongs, with the magnitude of its weight indicating its centrality to that community. Regions with absolute z -scores exceeding 1.3 ($|z| > 1.3$) can be considered core nodes in a more distributed network which covers the entirety of the brain space.

5.3.4.2 Study II: Multiple Disorders

The independence of each dimension in the control basis space also allows researchers to determine whether entropy alterations concentrate in individual dimensions or spread evenly across the space. This may be done via statistical tests on the individual rows of the entropy array $D \times S$ (with appropriate multiple comparison correction). Three dimensions are found to display substantial inter-group alterations, specifically the sixth, seventh, and twelfth (ordered from highest to lowest mean activity) (Figure 8).

Differences between controls and ADHD patients concentrate in the sixth dimension ($p = 8.991 \times 10^{-4}$, $Hodges' g = -0.6167$) while the seventh appears to most influence schizophrenia patients ($p = 0.0011$, $Hodges' g = -0.5821$). Bipolar patients, unlike ADHD or schizophrenia, display elevated entropy in both the seventh ($p = 9.999 \times 10^{-5}$, $Hodges' g = -0.7020$) and twelfth ($p = 6.9993 \times 10^{-4}$, $Hodges' g = -0.6120$) dimensions. Patients display elevated entropy relative to controls in all pairings. The most influential regions in each of these dimensions may be viewed in Table 3, and a list of dimensional entropies per condition

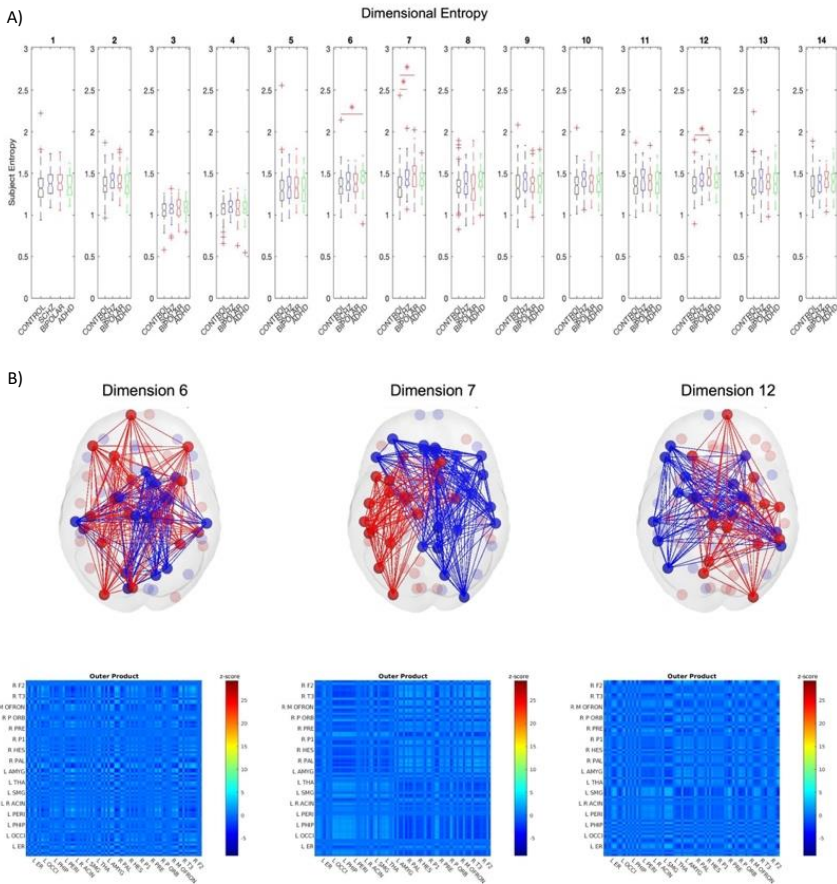
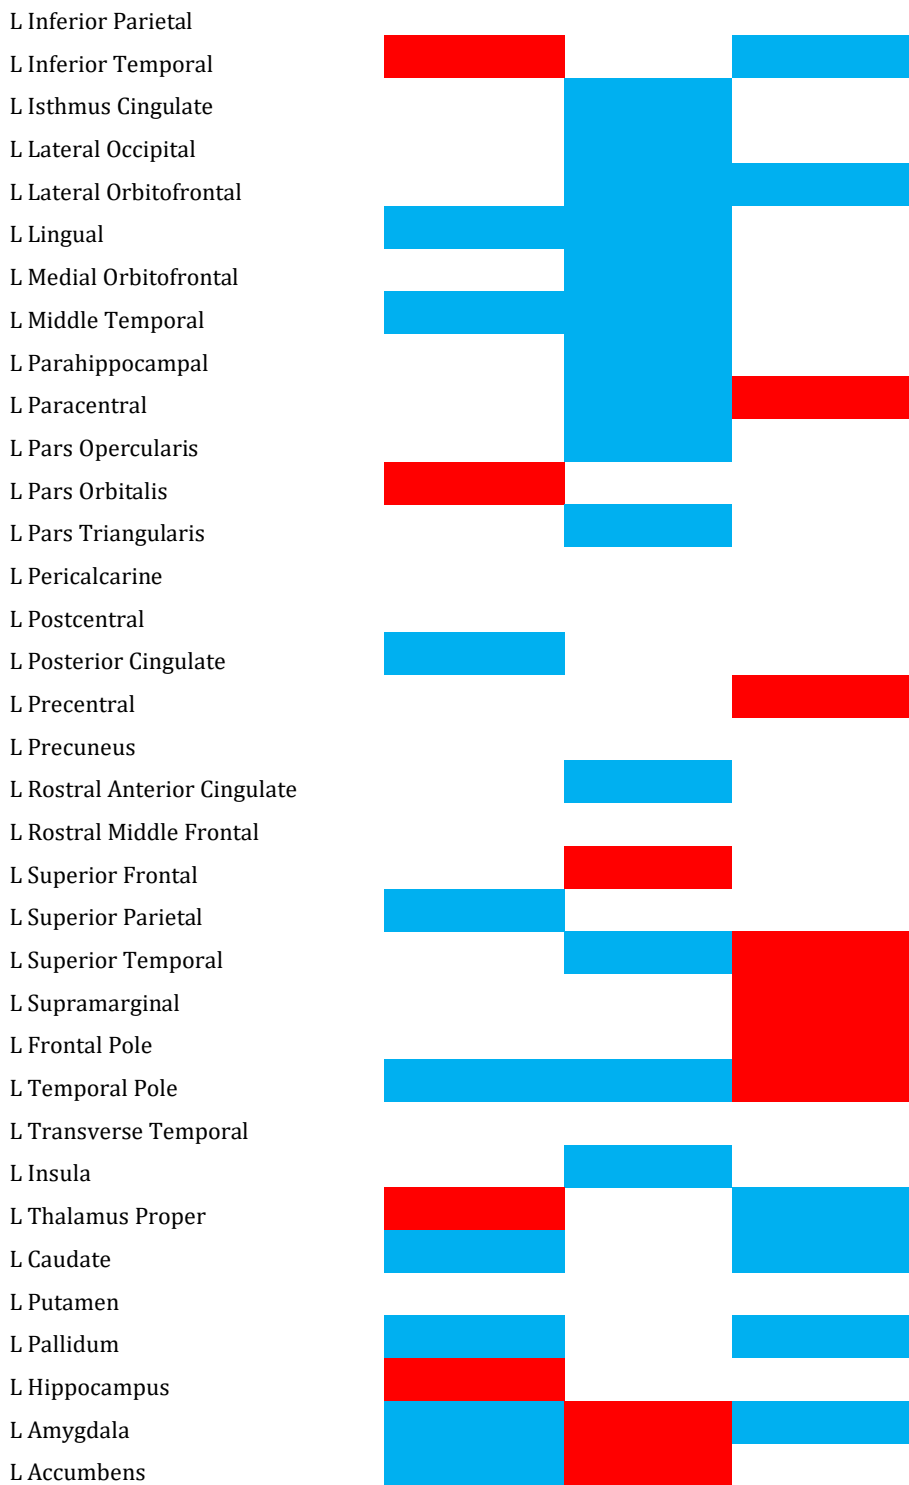
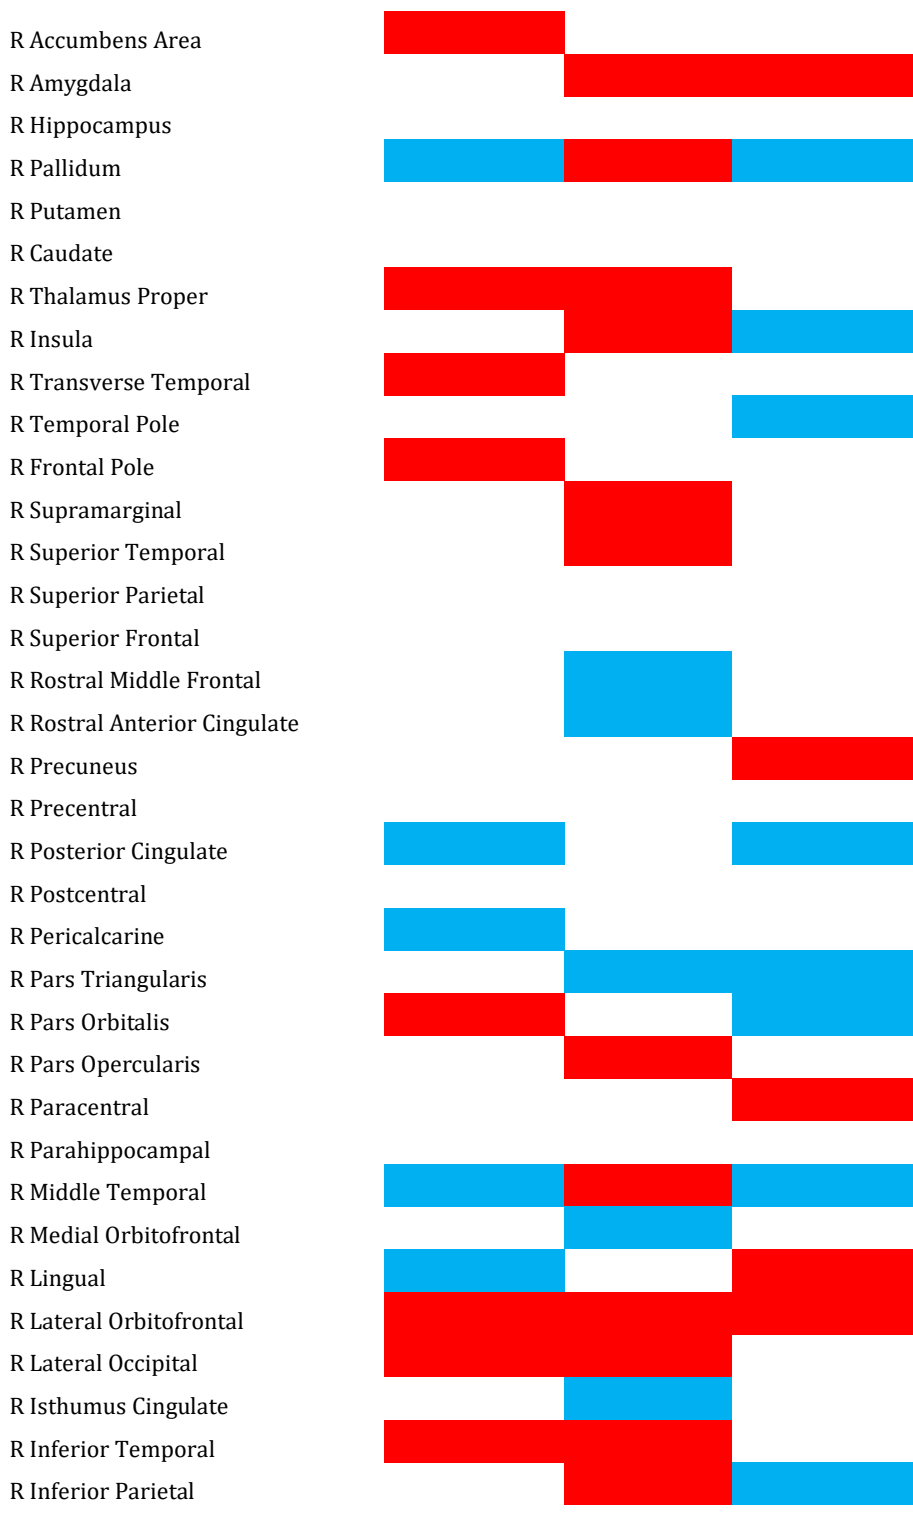


Figure 10 | Group-level entropy analysis shows that three components—six, seven, and twelve—display significant alterations between control and patient entropies (A). These alterations appear disorder-specific, with ADHD displaying elevated entropy only along the sixth component and schizophrenia only along the seventh. Bipolar disorder (type I) displays alterations along both components seven and twelve. The spatial maps of these components are displayed in greater detail in panel (B).

Table 3 | Significant LEICA Dimensions

	Dimension 6	Dimension 7	Dimension 12
L Bank's Superior Temporal Sulcus			Blue
L Caudal Anterior Cingulate	Blue	Blue	
L Caudal Middle Frontal	Red		
L Cuneus			
L Entorhinal	Blue		Red
L Fusiform	Red		Red





R Fusiform
R EntoRinal
R Cuneus
R Caudal Middle Frontal
R Caudal Anterior Cingulate
R Bank's Superior Temporal Sulcus




Table 4 | LEICA Dimensional Entropies

	Control	Schizophrenia	Bipolar	ADHD
Comp. 1	1.3314 ± 0.18354	1.3817 ± 0.17252	1.392 ± 0.15928	1.3476 ± 0.14804
Comp. 2	1.3631 ± 0.15676	1.417 ± 0.13407	1.4133 ± 0.14623	1.3654 ± 0.15778
Comp. 3	1.0527 ± 0.10891	1.0712 ± 0.11667	1.0808 ± 0.13456	1.087 ± 0.1068
Comp. 4	1.0678 ± 0.11119	1.0959 ± 0.09644	1.0774 ± 0.13626	1.0662 ± 0.131
Comp. 5	1.3068 ± 0.20813	1.3259 ± 0.17096	1.3431 ± 0.16761	1.3052 ± 0.16197
Comp. 6	1.3509 ± 0.1608	1.4065 ± 0.13561	1.3872 ± 0.14497	1.4481 ± 0.14761
Comp. 7	1.339 ± 0.19451	1.4498 ± 0.17964	1.4759 ± 0.19629	1.4251 ± 0.13253
Comp. 8	1.3518 ± 0.17209	1.3756 ± 0.19924	1.3438 ± 0.22108	1.4191 ± 0.17177
Comp. 9	1.3356 ± 0.18963	1.4173 ± 0.16614	1.3672 ± 0.15345	1.3842 ± 0.15147
Comp. 10	1.359 ± 0.16396	1.4222 ± 0.13863	1.3724 ± 0.14273	1.3831 ± 0.1449
Comp. 11	1.3482 ± 0.15999	1.4234 ± 0.15426	1.396 ± 0.16183	1.3835 ± 0.16683
Comp. 12	1.3608 ± 0.15143	1.4212 ± 0.12474	1.4522 ± 0.1445	1.3993 ± 0.12304
Comp. 13	1.3493 ± 0.17731	1.4133 ± 0.1802	1.3844 ± 0.14786	1.3936 ± 0.18206
Comp. 14	1.3513 ± 0.18844	1.3888 ± 0.13996	1.4137 ± 0.1372	1.3869 ± 0.14364

5.4 Conclusions

The LEICA framework has proven remarkably capable of separating psychiatric patients from controls, regardless of the condition with which it is presented. This is a substantial achievement, as many studies focus on a single disorder rather than several simultaneously. Further, the LEICA framework identifies which resting state network(s) are affected by each disorder, which provides researchers with clear targets for future investigations.

Our initial study demonstrated that the use of the functional connectivity's leading eigenvector provides a demonstrable improvement in discriminability compared to the use of the full functional connectivity matrix, thus validating the use of this format in our second study. Thus, one can state that the LEICA framework, using Shannon entropy as a functional feature, provides a powerful tool in the continued search for explanatory biomarkers in neurological and psychiatric disease.

6 Predicting Connectivity

6.1 Motivations

Having established that psychiatric disorders affect the functional complexity of resting-state brain activity, the next logical question is how these effects are implemented in the brain. This is, as most questions in neuroscience are, an exceedingly complex question. It should surprise no one that neuroscience and complexity science substantially overlap. Thus, as with the metrics used before, one must approach this question by first establishing which aspect of it to address. The present thesis addresses it through the filter of whole-brain effective connectivity, which seeks to estimate the influence that each region of the brain has on other regions. Such an approach has gained increasing popularity in the past decade, as an abundance of neuroimaging data, the increased availability of powerful computational resources, and an explosive growth in machine learning techniques have made implementing such estimates ever more feasible.

Why estimate the influence that a node has over any other? As mentioned previously, there is the scientific impulse to learn the causal mechanisms underlying observed functional alterations. One must remember, however, that neuroscience is motivated not only by scientific curiosity, but also by the need for more effective interventions for neurological and psychiatric disease. Such interventions must also be targeted to avoid affecting healthy function. As such, knowledge of where aberrant activity originates—whether for surgical intervention, as is the case for epilepsy, or for pharmaceutical intervention, if a neurotransmitter network is suspected—is vital both for planning immediate interventions and for researching new ones.

6.2 Determining the Method

While the approach of whole-brain effective connectivity represents only a small part of the methods for evaluating how psychiatric disease affects brain function, even this small part contains a vast number of potential implementations. Each study must select

between a model-based or model-free approach, the brain parcellation to use, the machine learning algorithm to employ, and the features on which to fit the estimate, along with other variables. Thus, as with the question of how psychiatric disease affects functional complexity, one must choose which models to implement and which to ignore.

The studies composing this thesis employ a model-based approach, with a Hopf bifurcation (Section 2.4.b.iii) serving to simulate the slow dynamics of each region in the brain. This approach has become standard in the Computational Neuroscience (CNS) Group following confirmation of its ability to replicate the slow dynamics detected in fMRI signals (Deco, Cabral, et al. 2017; Deco, Kringelbach, et al. 2017). One of the principal advantages of the Hopf model is its simplicity. The dynamics of each individual node j , depend upon only three parameters: the intrinsic frequency ω , the oscillation amplitude parameter α , and the incoming connections $\sum_i C_{ij}$. This simplicity imposes a high degree of abstraction from neurobiologically interpretable parameters, but it also simplifies the both the fitting procedure and the mathematical interpretation of the resulting model. Of note, such a model demonstrating that brain regions in the resting state appear to operate near a critical point (Deco, Kringelbach, et al. 2017), which maximizes the metastability and functional flexibility of connectivity dynamics (Hansen et al. 2015).

The parcellations employed were largely determined by the datasets available for study. The first dataset arrived in the 116-region Automated Anatomical Labeling parcellation (Tzourio-Mazoyer et al. 2002), with the second using the 68-region Desikan-Killany atlas (Desikan et al. 2006a). Both parcellations include standardized structural connectivity atlases, which are used as masks¹ for the estimation of effective connectivity. More fine-grained parcellations have since become available, but computational requirements scale with the square of number of regions; as such, studies must always accommodate available resources in their pursuit of detail.

¹ Only connections corresponding to the atlas' nonzero elements are estimated; all others are set to zero.

The choice of machine learning algorithm is perhaps the least constrained of those detailed in this thesis. While the Hopf model has become the standard model in the CNS group and parcellations were determined externally, the field of machine learning algorithms has grown to a dizzying size and shows no signs of stopping. Whichever method was selected, however, would need the ability to fit a high-dimensional space with relatively few constraints. Ideally, it would not be constrained to strictly linear models, as brain dynamics tend to be nonlinear. Finally, demonstrable efficacy in fitting biological data was necessary. These three considerations led the author to, on his advisor's suggestion, employ the particle swarm optimization method (Kennedy and Eberhart 1995), which has proven to fit the above criteria.

6.3 Pipeline

6.3.1 Brain Network Model

For each study, the employed parcellation defines the brain network model. Specifically, each network model consists of N nodes, with each node representing a single region in the parcellation of interest. The parcellation's standard connectivity template \mathbf{C} defines the model's interregional coupling (network links). Thus, for the OCD-patient dataset, the network model consists of 90 nodes representing the 90 cortical and subcortical regions of the AAL parcellation, coupled according to the standard AAL connectivity matrix. For the second, four-group dataset, the network model contains the 68 nodes of the Desikan-Killany parcellation and is coupled according to that parcellation's standard connectivity template.

In both studies, internal node dynamics are modeled as the normal form of a Hopf oscillator set to criticality (Deco, Kringelbach, et al. 2017). This produces

$$\frac{dx_j}{dt} = x_j(\alpha_j - x_j^2 - y_j^2) - \omega_j y_j + G \sum_i C_{ij} (x_i - x_j) \beta \eta_j(t)$$

$$\frac{dy_j}{dt} = y_j(\alpha_j - x_j^2 - y_j^2) - \omega_j x_j + G \sum_i C_{ij} (y_i - y_j) \beta \eta_j(t)$$

where C_{ij} is the connection strength from j to i and G represents global coupling efficiency. ω_j is estimated directly from the BOLD time series by extracting the dominant frequency of node j within the band of 0.01 to 0.08 Hz. α and G are set to the initial values of $\alpha = 0$ and $G = 0.2$, in line with previous work (Deco and Kringelbach 2016; Deco, Kringelbach, et al. 2017).

6.3.2 Optimization

The nonzero connection strengths C_{ij} are optimized using the population swarm algorithm (Kennedy and Eberhart 1995; Erik, Pedersen, and Pedersen 2010; Mezura-Montes et al. 2011), which simulates a population of individuals (particles) randomly moving within an N -dimensional space. N is the number of free parameters in the model. At each optimization step, each particle can continue exploring the space, move to its optimal prior position, or move to the global optimal prior position. The model is then tested using the new positions of each particle as parameters, and the individual and global optima are updated as necessary. This method has proven robust and sensitive in biological optimization problems, particularly problems with a high dimensionality to data ratio; as such, it was a natural choice for the current model.

As with every optimization algorithm, the particle swarm algorithm needs a cost function to quantify the difference between simulated and empirical data distributions. For the present studies, we found that the Euclidean distance between entropy distributions,

$$d(S, E) = \sqrt{\sum_{j=1}^N (S(j) - E(j))^2}$$

provided a robust and easily interpretable measure of distance between simulation and observation. In the above formula, $S(j)$ indicates the simulated entropy of component j , while $E(j)$ represents the empirical entropy of the same component. The overall goal of the fitting procedure was to find the location in parameter space which minimized this distance.

With a parameter space composed of nonzero elements of each connectivity matrix, the entropy cannot be estimated directly. Instead, at each optimization step, a temporary network model is built for each particle with particle location values forming connection weights. The regional BOLD signals of this temporary network are then estimated, these simulated signals mapped to component time series via the mixing matrix W , and the Shannon entropy of each component is computed. The Euclidean distance between the simulated entropy distribution and its empirical counterpart is used as the optimization cost function, which guides the particle swarm algorithm's estimates for optimal model parameters.

It should be emphasized that only the nonzero elements of the structural connectivity matrix \mathbf{SC} are optimized; all elements satisfying $C_{ij} = 0$ was not altered. This choice was necessary as the large number of parameters to be optimized raised the danger of overfitting. Empirical tests demonstrated that optimizing the full matrix did not improve performance.

6.3.3 Node Strength

In order to find network-level connectivity changes in obsessive compulsive disorder patients, two group-level analyses are applied to the connectivity estimates obtained in the previous section. The first of these is group-level strength of the individual nodes in the network. A node's strength is simply the sum of the weights of all connections leading into (in-strength) or out of (out-strength) the node in question and provides some insight into the importance of that node to the network dynamics. As a rule, the higher a node's strength, the more crucial a role it plays in the network's function (although exceptions have been found in nature and human designs).

6.3.4 Network-Based Statistic

The network-based statistic (Zalesky, Fornito, and Bullmore 2010) represents an attempt to circumvent the multiple comparison problem in network analysis. As the number of possible links l

scales according to $l \propto N^2$, link-based analyses of even small networks will require extreme family-wise error correction. This means that most link-level effects will be ignored. However, while the number of links scales quadratically with network size, the number and size of components may not. Equally important, a null model for the number of expected components may be simulated by permuting network link locations while keeping the number of links constant. This allows researchers to estimate a null distribution for the number and size of components in the network and determine which, if any, extant communities fall outside this distribution.

Based on this intuition, the network-based statistic (NBS) employs the following procedure for detecting significant components. First, a group-level t -test establishes linkwise effect sizes in the network(s) of interest. A threshold is then applied to prune the network, and a breadth-first search establishes the number and size of surviving communities. The null distribution of suprathreshold community sizes is then determined via permutation and this null distribution used to estimate the p -value of each empirical community. In most networks, this test provides substantially greater power than link-based family-wise error control at the cost of being unable to identify the significance of individual links.

6.4 Results

6.4.1 Goodness-of-Fit

Subject-level effective connectivity was estimated by fitting a network model of homogenous Hopf oscillators (Deco, Kringelbach, et al. 2017) to each subject's entropy profile. We selected the particle swarm algorithm for model optimization (Kennedy and Eberhart 1995) as this algorithm has demonstrated a good mixture of accuracy, robustness, and computational efficiency in biological models (Mezura-Montes et al. 2011; Poli, Kennedy, and Blackwell 2007; Erik, Pedersen, and Pedersen 2010). The component-level Euclidean distance between simulated and empirical entropy serves as cost function for model optimization.

Only the eigenvector-based decomposition produced a generative model which displays significant group-level alterations in network connectivity; the spatially averaged and uncompressed decompositions failed to find meaningful results. As such, only eigenvector-based estimates are displayed in the following sections.

6.4.1.1 Study I: Obsessive-Compulsive Disorder

Modeling each subject separately provides a complete distribution of distances for each group and each component, illustrated in Figure 11. Comparisons of pre- and post-fit distances demonstrate that the particle swarm algorithm performs admirably on the control group, as all components' estimated entropy distributions are brought into alignment with those from empirical data.

Optimization on the patient group is markedly less successful. While patients' simulated distribution means grow closer to empirical means, the patients' simulated distribution variances grow enormously. Repeated efforts at model fitting demonstrate that this explosion in estimated variance appears to be a consistent theme, but one restricted to patients alone. As both groups undergo identical procedures and the empirical variances of the two groups are similar, the cause of this discrepancy is not clear.

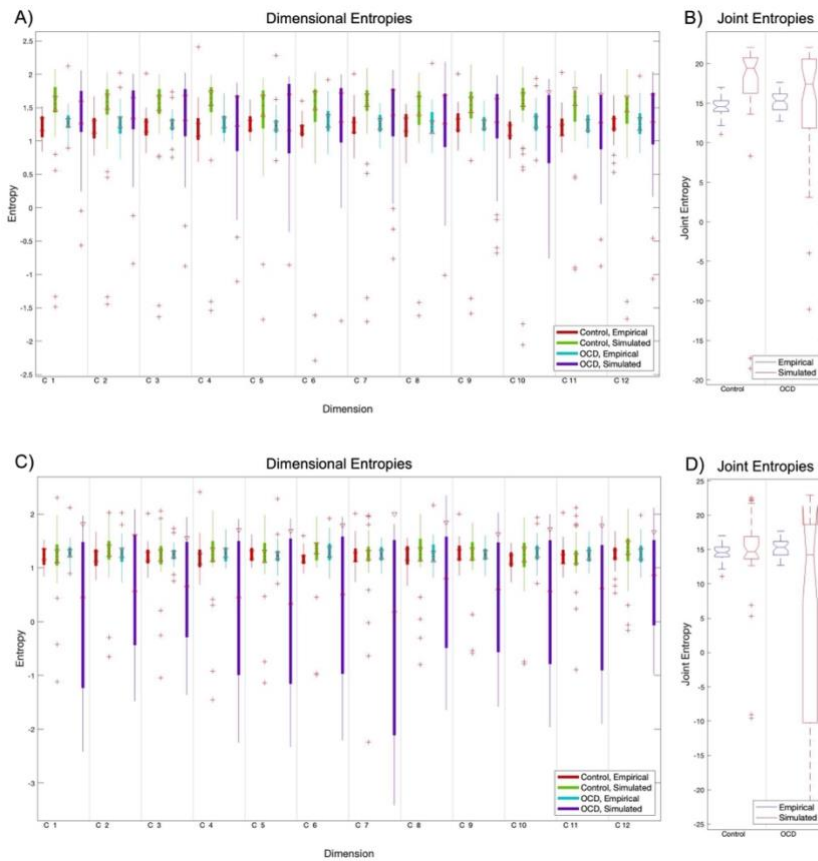


Figure 11 | The particle swarm fitting algorithm, like most optimization algorithms, minimizes a cost function which (in principle) measures how well the predicted solution aligns with measurable data. We chose the Euclidean distance between empirical and simulated entropy vectors as a cost function due to its conceptual simplicity and confirmed its superiority versus absolute maximum distance. Comparisons of component-level entropy distributions pre-fit (A) and post-fit (B) demonstrate that this method does improve the model for controls, a finding confirmed in measurements of the joint entropy. While optimization brings the mean entropies of patient models closer to those of empirical subjects, its performance is quite inconsistent in this group. This is reflected in the extremely high variance in post-optimization dimensional and joint entropies (B).

6.4.1.2 Study II: Multiple Disorders

The optimization procedure for the second dataset is identical to that employed in the obsessive-compulsive data, with the resultant distance distributions illustrated in Figure 12. Unlike the original dataset, comparisons of pre- and post-fit distances do not demonstrate a dramatic condition-dependent effect on optimization

performance. Dramatic component-dependent effects are evident, however, as trained models substantially underestimate the entropy of components six (6) and seven (7) (Figure 12). Component eight (8) is less dramatically but still consistently underestimated, while models overestimate the entropy of component three (3). While these components represent fewer than a third of the total signal, their poor fit suggests that connectivity alone may be unable to replicate the functional signals of the brain.

Despite these errors, optimization generally improves the match between modeled and empirical data (Figure 12). The general tendency to underestimate empirical data can likely be ascribed to components six and seven depressing the estimate of joint entropy. It is notable that the control group is substantially better fit than patient groups, suggesting that the decision to define components on control data alone has affected the model's sensitivity to other groups. This implies that psychiatric disorders may affect the components' structure as well as their functional signatures, a possibility which may warrant investigation in a future study.

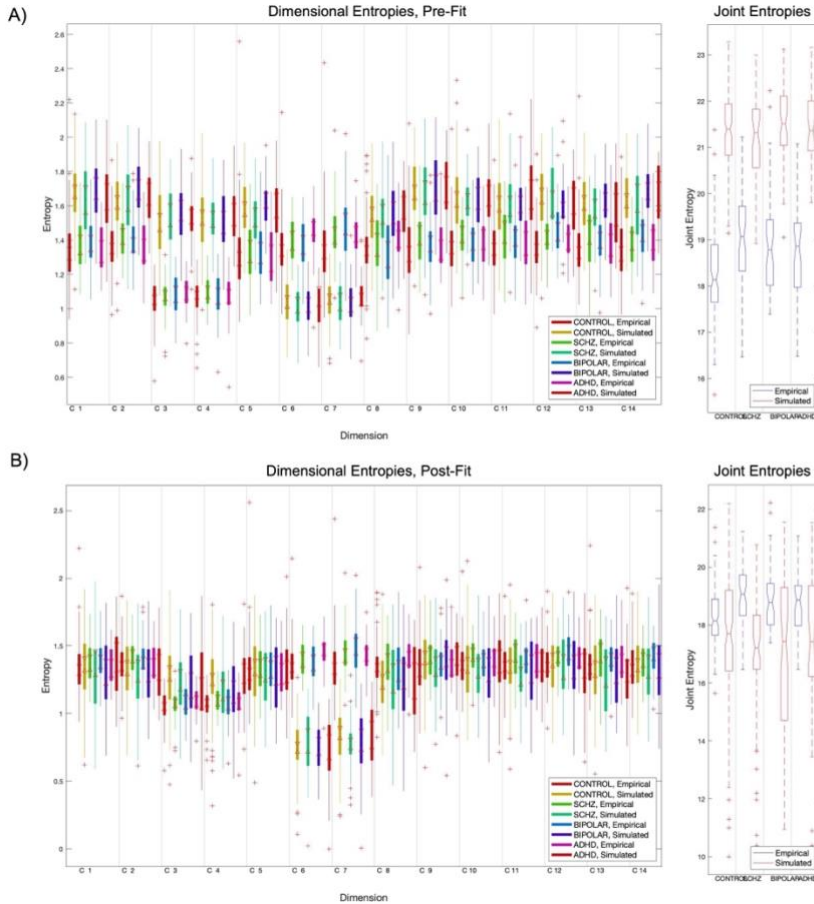


Figure 12 | The particle swarm fitting algorithm, like most optimization algorithms, minimizes a cost function which (in principle) measures how well the predicted solution aligns with measurable data. We chose the Euclidean distance between empirical and simulated entropy vectors as a cost function due to its conceptual simplicity and confirmed its superiority versus absolute maximum distance. Comparisons of component-level entropy distributions pre-fit (A) and post-fit (B) demonstrate that this approach substantially improves the model's fidelity to real data across multiple groups.

6.4.2 Node Strength Analysis

Upon obtaining effective connectivity estimates, node strength distributions were compared across groups to determine which, if any, conditions express over- or under-connectivity. Identical procedures were applied to both datasets, with appropriate multiple comparison corrections applied to the four-group data.

6.4.2.1 Study I: Obsessive-Compulsive Disorder

No meaningful differences in node strength were apparent between obsessive-compulsive patients and healthy controls. This may be due to the relatively poor fit between patients and controls in this study, as the wide distribution of patient entropies suggests a lack of consistency in estimated connectivity.

6.4.2.2 Study II: Multiple Disorders

In the four-group data, results were surprisingly robust: all three patient groups express an increase in both in- and out-strength over the control group (Figure 14). This result is both widespread and monotonic, with most regions displaying increased connectivity in both directions and in all patient groups (Table 5). This suggests that schizophrenia, ADHD, and bipolar disorder may be characterized by overconnectivity across a large part of the cortex and subcortical regions. This, in turn, may contribute to the elevated levels of entropy found in patients (Figure 12).

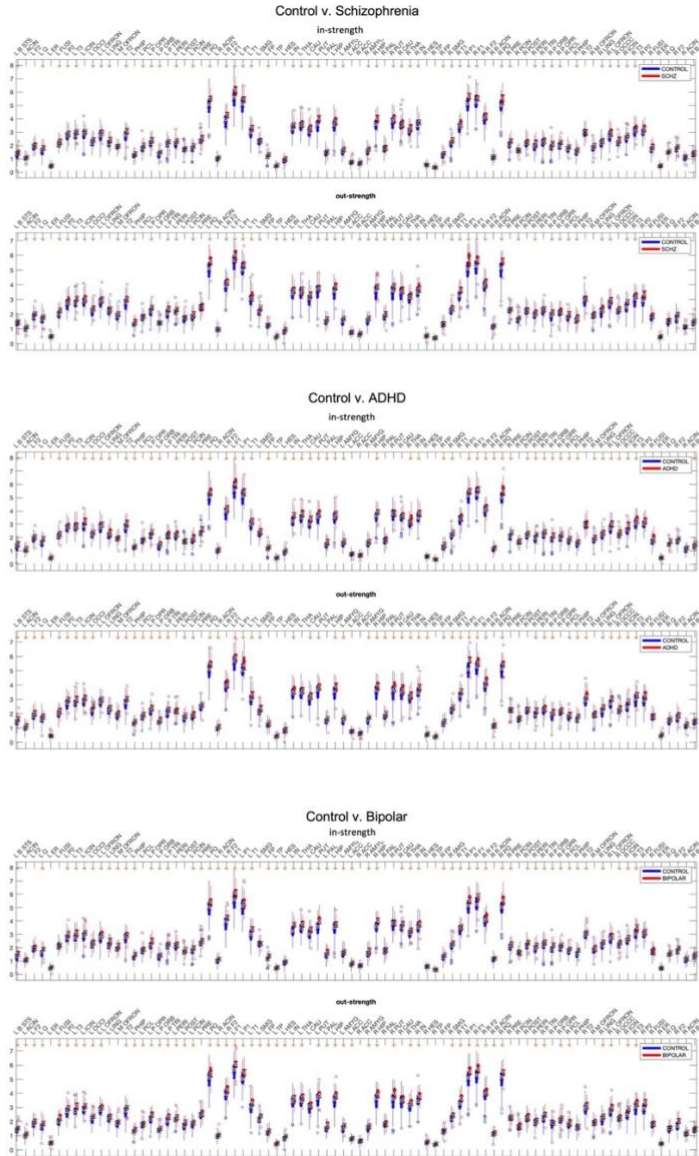


Figure 13 | Results from in- and out-strength analysis of controls vs. patient groups. After familywise error correction, all six pairwise combinations display significant node-level differences in strength in both directions. Control strengths are lower than patient strengths across conditions, and this difference remains significant for most regions even after multiple comparison correction. It appears that schizophrenia, bipolar disorder, and ADHD may be characterized by overconnectivity across much of the brain, which may be a contributing factor to the elevated disorder in patients.

Inter-disorder strength comparisons are sparser, but still notable. Schizophrenia patients displays significantly lower in- and out-strength than either bipolar or ADHD patients in most regions, with the exceptions of the right fusiform gyrus (in-strength, bipolar) and the right pars orbitalis (out-strength, bipolar). This suggests that lowered connectivity may differentiate schizophrenia from bipolar or ADHD patients, although the sparsity of these results should be considered before treating such findings as definitive. As for bipolar disorder, the left parahippocampus, left putamen, and right parahippocampus display elevated in-strength versus ADHD patients, but this is reversed in the right fusiform gyrus. The general trend reverses in the outgoing direction, with ADHD expressing higher out-strength in the right pars orbitalis and right middle temporal gyrus compared to bipolar disorder. Only the right caudal middle frontal gyrus displays higher out-strength in bipolar disorder than in ADHD. Strength analysis thus suggests that all three psychiatric conditions examined in this study display dramatic hyperconnectivity compared to healthy controls, while schizophrenia appears to have the least elevated connectivity amongst patient groups. Complete lists of significant regions and their mean strength per condition may be found in Supplementary Table I and II.

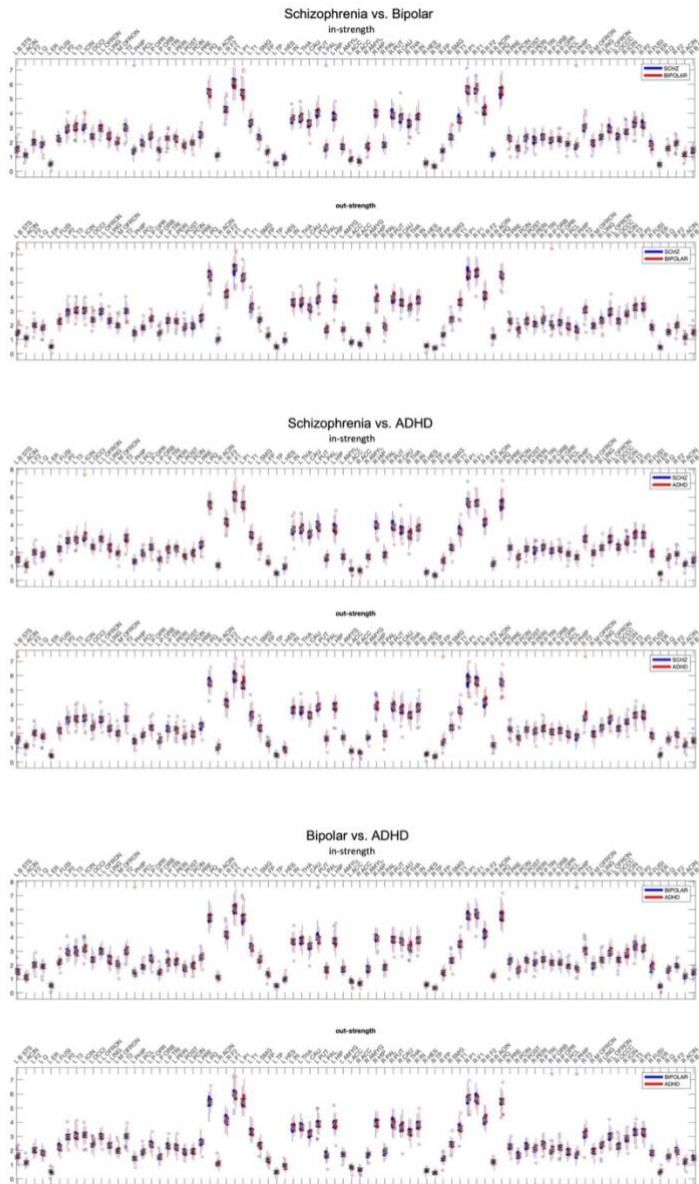


Figure 14 | Results from in- and out-strength analysis of patient groups vs. other patient groups. Results are much sparser between patient groups than between patients and controls, but trends remain apparent. For instance, all but one result show schizophrenia with reduced strength values compared to ADHD or bipolar disorder. Comparing bipolar disorder to ADHD, on the other hand, shows that bipolar disorder generally has elevated in-strengths but lower out-strengths. A rough hierarchy is thus apparent, in which bipolar disorder and ADHD are highly overconnected and schizophrenia slightly overconnected compared to controls. It is intriguing to note that bipolar disorder is also the only disorder to display elevated entropy in two LEICA dimensions rather than only one.

Table 5 / Number of Significant Regions per Comparison

	in-strength	out-strength
<i>Controls vs. schizophrenia</i>	52	52
<i>Controls vs. bipolar</i>	61	61
<i>Controls vs. ADHD</i>	59	59
<i>Schizophrenia vs. bipolar</i>	4	4
<i>Schizophrenia vs. ADHD</i>	1	1
<i>Bipolar vs. ADHD</i>	4	4

6.4.3 Network-Based Statistic

In order to hypothesize on causes for the observed alterations in Shannon entropy, a networked Hopf model (Deco, Kringelbach, et al. 2017) was fit to each subject’s entropy profile. While previous studies have examined the working point α of networked Hopf oscillators, these two studies instead sought to find which connectivity changes might cause the observed functional alterations. After estimating subject-level effective connectivity profiles from these models, we applied the network-based statistic (NBS) (Zalesky, Fornito, and Bullmore 2010) to determine which, if any, connections display significant group-level alterations.

Results from the network-based statistic depend upon the t -statistic chosen at the thresholding step. Unfortunately, no data-driven method for determining an optimal threshold has yet been developed, nor has such a threshold been established experimentally. As such, it must be treated as a free parameter.

6.4.3.1 Study I: Obsessive-Compulsive Disorder

Obsessive-compulsive connectivity estimates underwent NBS analysis at t -thresholds of 4.0, 4.5, 5.0, 5.5, and 6.0. A t -threshold of 4.5 reveals a single large hyperconnected component and eleven small hypoconnected components in the patient population (Figure 15). These components consist of

1. Left superior frontal gyrus (orbital), left superior frontal gyrus (medial orbital), and left lenticular nucleus (putamen)

2. Left middle frontal gyrus and left caudate nucleus
3. Left Rolandic operculum, left insula, left supramarginal gyrus, left superior temporal gyrus, left middle temporal gyrus, and left temporal pole (middle temporal gyrus)
4. Left middle occipital gyrus, left superior frontal gyrus (medial), left middle occipital gyrus, left inferior occipital gyrus, left and right precuneus, left superior parietal gyrus, and right superior occipital gyrus
5. Left calcarine fissure, left fusiform gyrus, left cuneus, left and right posterior cingulate gyrus, and left superior occipital gyrus
6. Right temporal pole (superior temporal gyrus), right inferior temporal gyrus, and right middle frontal gyrus (orbital)
7. Right middle temporal gyrus and right inferior frontal gyrus (orbital)
8. Right supplementary motor area and right paracentral lobule
9. Right amygdala and right fusiform gyrus
10. Right inferior frontal gyrus (triangular) and right middle frontal gyrus
11. Right inferior frontal gyrus (opercular) and right precentral gyrus

Other settings of the t -statistic threshold produce slightly different results. For example, raising the threshold to $t = 5$ causes the hyperconnected network to fragment into a single large component and two small ones. Similarly, one could expect that hypoconnected components would consolidate into fewer, larger networks at lower thresholds.

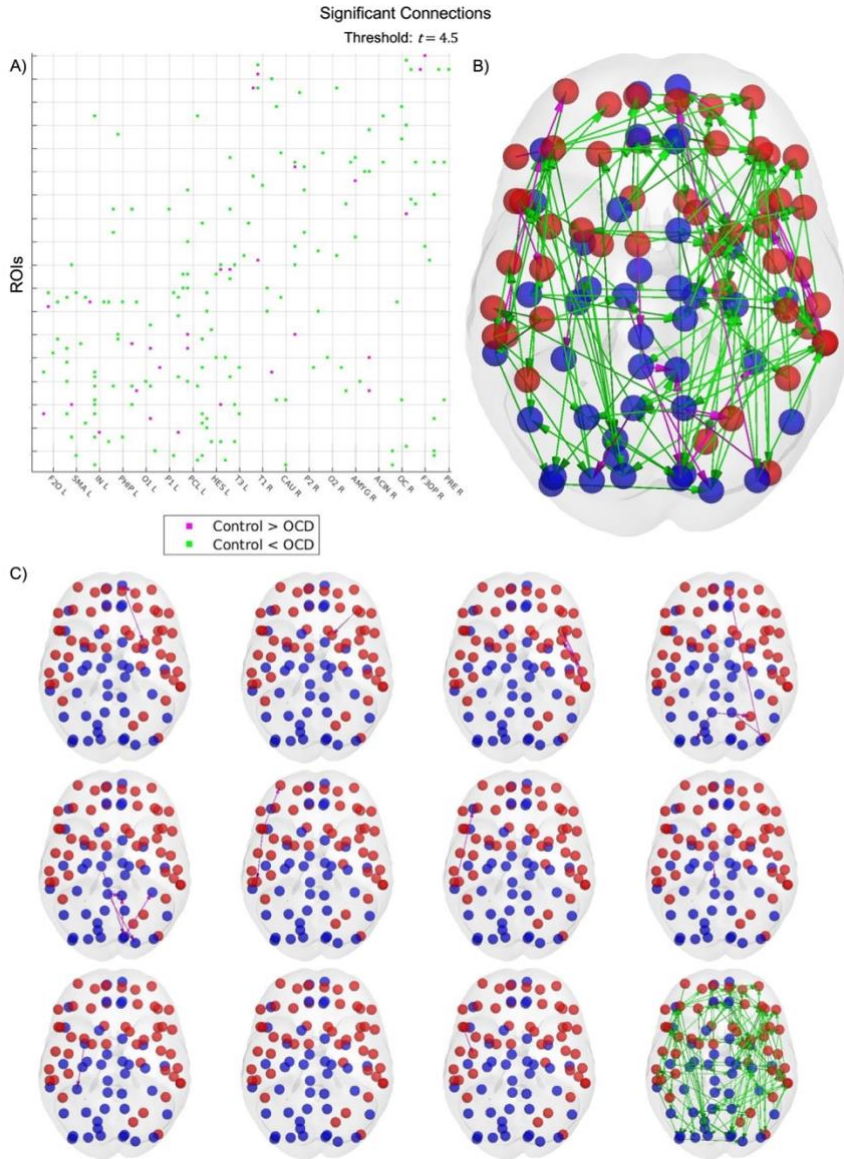


Figure 15 | Results from the network-based statistic. A t -statistic threshold of 4.5 returns twelve connected components (C), visualized as a connectivity matrix (A) and in cortical space (B). Cyan links indicate a connection which is stronger in OCD patients than in controls, while magenta links indicate the converse. One connected component including 87 of the 90 cortical nodes displays increased strength in OCD, suggesting widespread cortical hyperconnectivity. The eleven control-biased components consist of between one to six links, with large components clustering in small areas. Many regions displaying depressed connectivity in patients are involved in top-down control and impulse inhibition. OCD may thus be characterized by localized disruptions in top-down inhibitory activity, which may explain the widespread hyperconnectivity observed in patients.

6.4.3.2 Study II: Multiple Disorders

As with the obsessive-compulsive dataset, the t -statistic threshold is treated as a free parameter with the thresholds of 3.0, 3.5, 4.0, 4.5, and 5.0. being tested. Significant alterations were found in all three patient groups relative to controls. Notably, most of these alterations suggest hyperconnectivity in patients. Of the three conditions examined, only ADHD displays any hypoconnected pathways, and these are substantially outnumbered by hyperconnected tracts. Hyperconnectivity thus seems to be a characteristic trait of all the disorders examined here. The results also suggest a weak hierarchy amongst the disorders, with ADHD displaying relatively weak connections and bipolar disorder relatively strong ones.

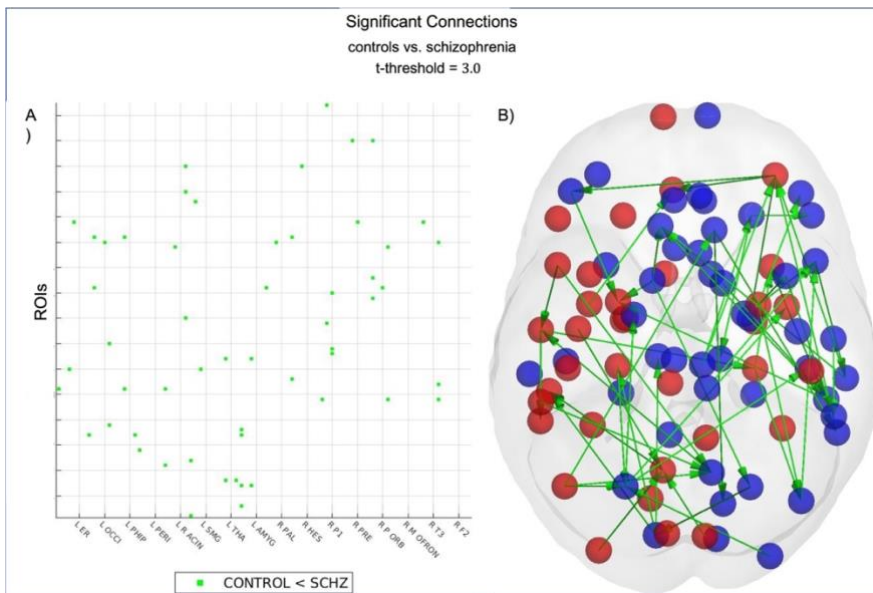


Figure 16 | Results from the network-based statistic comparing controls vs. schizophrenia patients. A t -statistic threshold of $t = 3.0$ reveals a single, large, fully connected network of heightened connectivity in patients. Raising the threshold to $t = 3.5$ splits this large component into three smaller subcomponents and raising it to $t = 4.0$ returns four significant interregional links. This figure visualizes the results as $t = 3.0$ via a connectivity matrix (A) and in cortical space (B). Patients display hyperconnectivity at all examined thresholds.

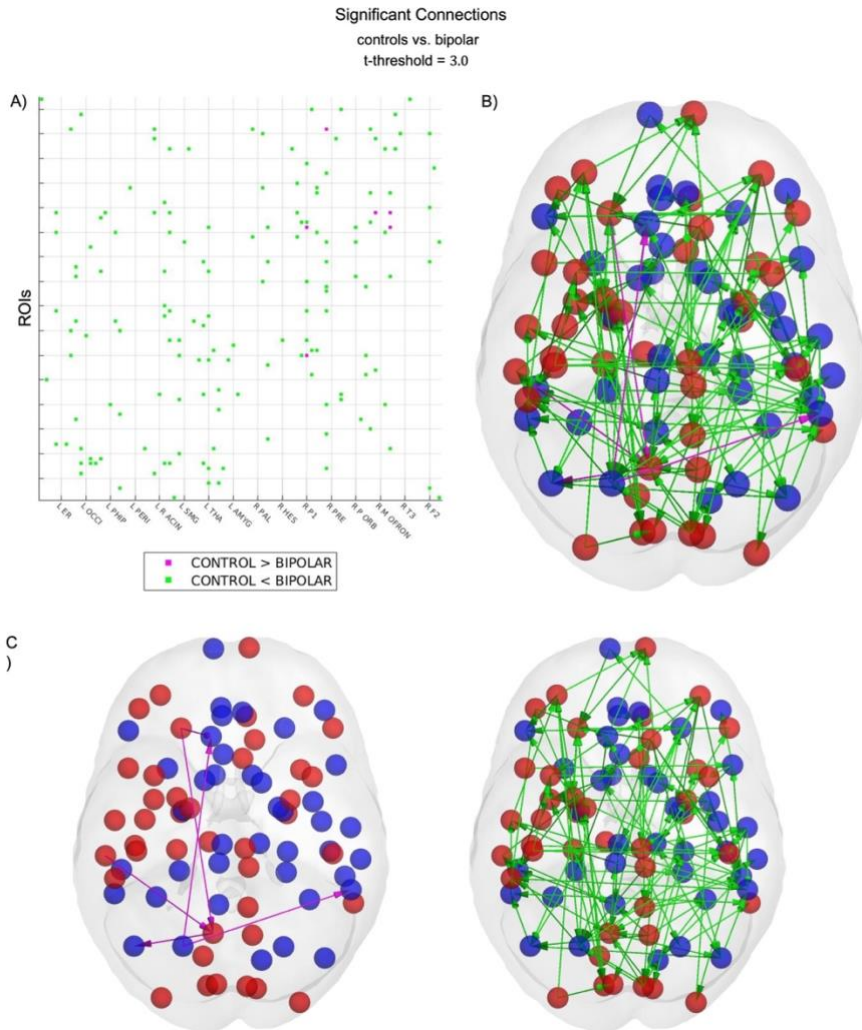


Figure 17 | Results from the network-based statistic comparing controls vs. bipolar patients. A t -statistic threshold of $t = 3.0$ reveals two connected networks. The first of these components displays lower connectivity in patients compared to controls. The other, substantially larger component displays heightened connectivity in patients. This hyperconnected component encompasses the majority of the regions examined in this study. Raising the threshold to $t = 3.5$ splits this large hyperconnected component into two smaller subcomponents. while raising it to $t = 4.0$ produces fourteen small hyperconnected components. Of these fourteen components, ten consist of only a single link; the remainder vary in size from three to five links. This figure visualizes the results of $t = 3.0$ as a connectivity matrix (A), in cortical space (B), and as individual components (C).

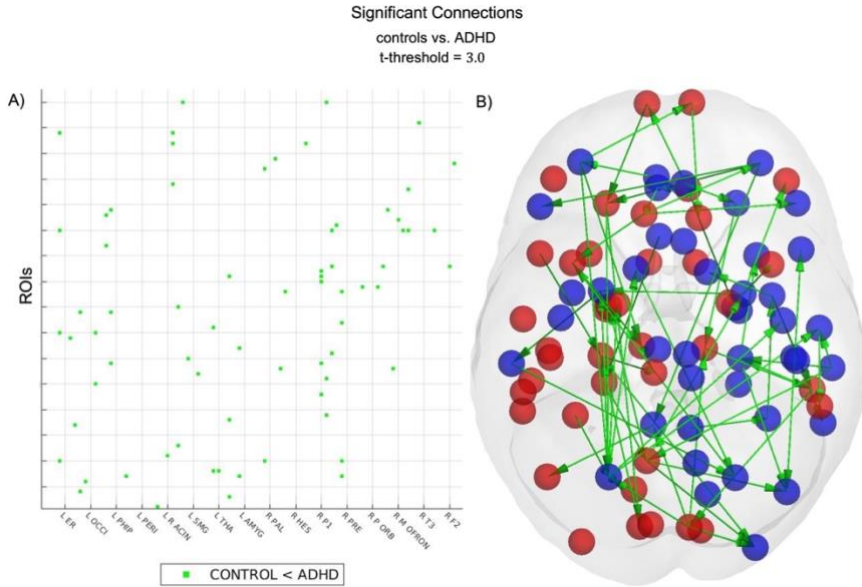


Figure 18 | Results from the network-based statistic comparing controls vs. ADHD patients. A t -statistic threshold of $t = 3.0$ reveals a single, large, fully connected network of heightened connectivity in patients. Raising the threshold to $t = 3.5$ splits this large component into six smaller subcomponents. Intriguingly, a threshold of $t = 4.0$ returns ten significant components, of which only nine are stronger in patients than in controls; the final link displays the opposite trend, i.e. a decrease in patient connectivity relative to controls. This hypoconnected component only appears at a threshold of $t = 4.0$ and consists of only a single link. Other thresholds reveal only hyperconnected components, with all components at thresholds of $t = 3.0$ and $t = 3.5$ consisting of multiple links.

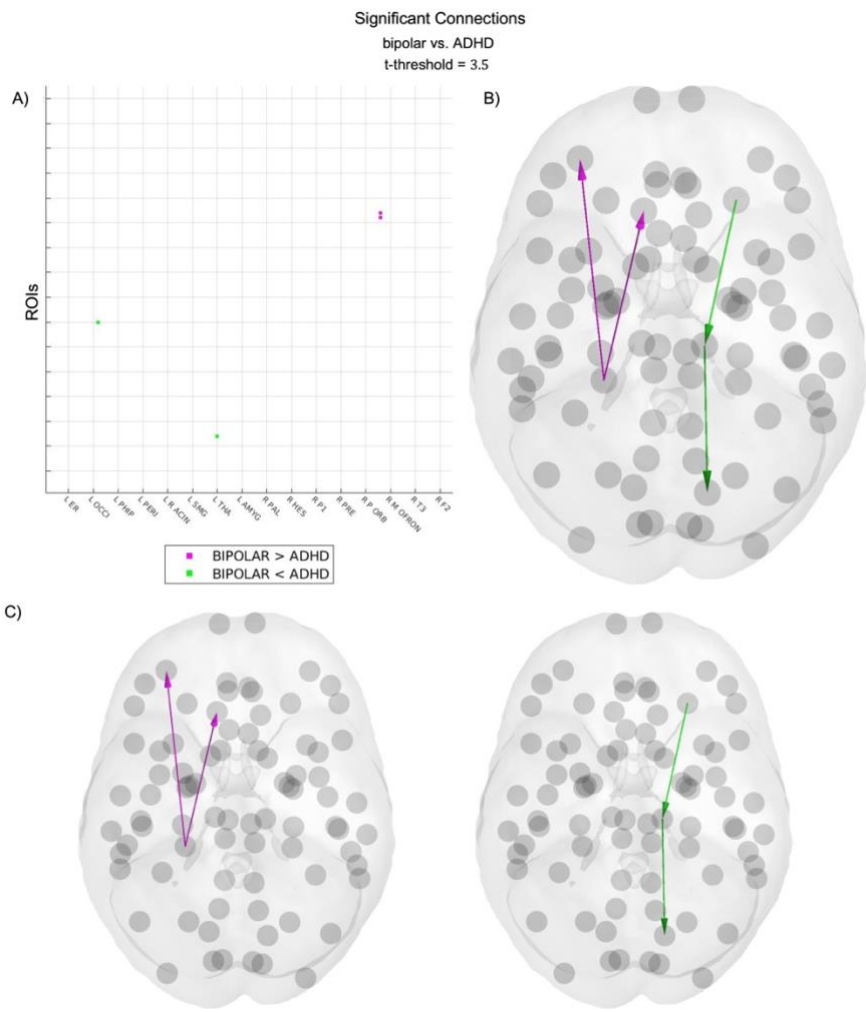


Figure 19 | Results from the network-based statistic comparing bipolar vs. ADHD patients. A t -statistic threshold of $t = 3.5$ returns two significant components (C), visualized together as a connectivity matrix (A) and in cortical space (B). Of these two components, one is stronger in bipolar patients than in ADHD patients, with the other being stronger in ADHD than in bipolar patients. All components consist of two links, none of which are interhemispheric.

Raising the threshold to $t = 4.0$ substantially changes the resulting connectivity map. This threshold reveals three links, two of which are stronger in bipolar patients than in ADHD patients. Surprisingly, only one of these three links is shared with the components revealed at $t = 3.5$; the other two are novel. This discrepancy is presumably an artifact of the alterations in probability distributions that these two thresholds generate.

increased entropy reported in the previous section; increased sensitivity to inputs may cause more erratic network behavior. In addition, some of the hypoconnected links appear to affect regulatory regions which normally mediate cortical activity. For example, the pallidum is known to inhibit output from the striatum to the thalamus and subthalamus; altered inhibitory output from the external globus pallidum to the subthalamus may cause the thalamic and cortical hyperexcitability reported in some OCD patients. Other affected regions include the Rolandic operculum, which has been linked to the impulsivity of autogenous obsessions and compulsions in some patients (Subirà et al. 2013; Alonso et al. 2013) and the premonitory obsessions of some tics and sensory phenomena (Wang et al. 2011) that many OCD patients display (Rosario et al. 2009).

The recovery of hyperconnectivity in all patient groups is somewhat surprising. Schizophrenia, in particular, has long been believed to be characterized by long-range hypoconnectivity (Friston et al. 2016). Recent studies have begun to contest this hypothesis (Rolls et al. 2020), suggesting that schizophrenia patients may display a mixture of hyper- and hypoconnectivity in long-range connections. Nonetheless, the fact that all significant links in schizophrenia patients are significantly stronger than in controls is, to this author's knowledge, unusual. Lowering the t -threshold of the network-based statistic may reveal weakened connections in schizophrenia and other disorders; however, the universal elevation of in- and out-strengths in schizophrenia patients makes this author suspect otherwise. Unfortunately, the author does not feel qualified to speculate on reasons for this apparent contradiction.

The relatively poor fit of patient data in the obsessive-compulsive dataset is worthy of further examination. The Hopf model's performance on the second dataset suggests that model suitability is not the problem; rather, some other factor prevents an adequate capture of the data's dynamics. The larger feature space of the AAL parcellation—90 regions compared to the Desikan-Killany's 68—may make finding an optimal solution more difficult. Alternatively, the decision to run independent component analysis across both patients and controls may have limited the specificity of the resultant components; however, it is unclear why this would affect patients more than controls, as both groups are very nearly

the same size. The causes for this poor performance and the discrepancy between datasets may warrant investigation in future studies.

7 Discussion & Future Steps

7.1 General Review

Overall, during this thesis, the author has demonstrated that a neuroimaging signal's entropy is a feasible and sensitive biomarker for several common psychiatric disorders. All disorders examined thus far display elevated entropy compared to controls, although it remains to be seen whether this trend holds for others. That four disorders spanning multiple behavioral families all display a similar global trend, albeit within different dimensions, is somewhat surprising. It is hardly beyond the realm of possibility, but verification via alternative datasets and methodologies is necessary before drawing firm conclusions.

The use of a critical Hopf bifurcation as a basis for model-based effective connectivity is, to the author's knowledge, entirely novel. While Hopf bifurcations have been used to explore the criticality and timescale of the human brain (Deco, Cruzat Grand, and Kringelbach 2019), attempting to predict the direction and magnitude of activity flow in the brain is a novel application. This is perhaps unsurprising; the Hopf bifurcation is itself a relatively new addition to the field, seeing its full debut only five years before the submission of this work. Nonetheless, the fact that a relatively simple and abstract model can recover functionally plausible network connectivity further suggests that the dependence of macroscale dynamics on microscale activity is not straightforward.

That all four disorders display elevated entropy and hyperconnectivity indicates that these may be common causes for a wide range of psychiatric symptoms. Previous studies on obsessive-compulsive disorder have produced contradictory findings on connectivity and functional complexity, while schizophrenia has long been considered an underperformance of long-range connectivity and global integration. Reports are contradictory, however, and none utilize the methodology or parcellations applied in these studies. Further investigation is required to comment on these findings.

7.2 Review: Obsessive-Compulsive Disorder

7.2.1 The Phenomenology of Obsessive-Compulsive Disorder

The discovery of elevated entropy across all four investigated disorders raises intriguing questions regarding dynamics and information flow in psychiatric disorders. This finding is of particular note in obsessive-compulsive disorder, as previous studies report decreased entropy in obsessive-compulsive patients (Aydin et al. 2015). Analysis of networks of cortical and subcortical nodes in the cortico-striatal-thalamo-cortical (CSTC) circuit, on the other hand, have found increased entropy in adolescents with OCD (Sen et al. 2020), and more general desynchronization has been demonstrated elsewhere (Pujol et al. 2019). How this apparent increase in complexity—which, in information-theoretic terms, is equivalent to randomness—maps to the well-established tendency of obsessive-compulsive patients to become “stuck” in stereotyped, repetitive patterns will be an interesting topic for future research. It may be that these stereotyped patterns represent a coping mechanism, intended to reduce the randomness of brain activity by imposing control of inputs and responses.

Such a hypothesis receives some support in the fact that the dimension found to increase patient entropy maps to two anticorrelated networks which roughly separate prefrontal-parietal regions vs. subcortical-temporal nodes. Prefrontal and parietal regions exert a top-down inhibitory control on striatal and limbic regions, which has been related to emotion regulation and cognitive control capacities (Ochsner, Silvers, and Buhle 2012; Etkin, Büchel, and Gross 2015). Alterations in such interregional interactions have been associated with mood and anxiety disorders, including OCD (Etkin and Wager 2007; Picó-Pérez et al. 2017). Decreased order within this network may disrupt of top-down inhibition and thus affect emotion regulation and cognitive control, both of which are affected in the context of CSTC dysfunction in OCD (van den Heuvel et al. 2016). Stereotyped, repetitive behaviors—i.e. compulsions—may thus act as a compensatory mechanism by which the brain attempts to impose order on its surroundings.

Interestingly, the affected dimension also contains several occipital nodes. Although the occipital cortex have not typically been considered a core part of neurobiological models of OCD, previous research has shown that such regions and their projections to limbic cortices may play an important role in the induction of increased anxiety levels in patients with contamination obsessions induced by actual or mental images (i.e., intrusive thoughts) of dirt (Göttlich et al. 2014; Moreira et al. 2017). In future research, it may be worth examining whether the patient's entropic alterations along this dimension correlates with anxiety or compulsive behavior, which could be as measured by e.g. the Hamilton Anxiety Rating Scale (HAM-A) (HAMILTON 1959) or the Yale-Brown Obsessive-Compulsive Scale (Y-BOCS) (Goodman et al. 1989).

7.2.2 Effective Connectivity in Obsessive-Compulsive Disorder

The networked Hopf model appears to have successfully recovered meaningful effective connectivity alterations in all examined conditions. To this author's knowledge, this is the first time the Hopf model has been employed to estimate connectivity directly; previous studies from the CNS group have optimized alternative parameters, most notably the bifurcation parameter α , to determine the whether the brain's working point operated in a critical, subcritical, or supercritical space. The success of these studies may encourage further use of the Hopf model in connectivity-based research, as its relative simplicity makes implementation straightforward.

Hyperexcitability has been reported in recent studies of obsessive-compulsive disorder (Cano et al. 2018), so the finding of a broad network of hyperconnectivity in this disorder is not overly surprising. Indeed, such hyperexcitability may underly the increased entropy reported in the previous section; increased sensitivity to inputs may cause more erratic network behavior. In addition, some of the hypoconnected links appear to affect regulatory regions which normally mediate cortical activity. For example, the palladium is known to inhibit output from the striatum to the thalamus and subthalamus; altered inhibitory output from the external globus palladium to the subthalamus may cause the thalamic and cortical hyperexcitability reported in some OCD patients. Other affected regions include the Rolandic operculum,

which has been linked to the impulsivity of autogenous obsessions and compulsions in some patients (Subirà et al. 2013; Alonso et al. 2013) and the premonitory obsessions of some tics and sensory phenomena (Wang et al. 2011) that many OCD patients display (Rosario et al. 2009).

The relatively poor fit of patient data in the obsessive-compulsive dataset is worthy of further examination. The Hopf model's performance on the second dataset suggests that model suitability is not the problem; rather, some other factor prevents an adequate capture of the data's dynamics. The larger feature space of the AAL parcellation—90 regions compared to the Desikan-Killany's 68—may make finding an optimal solution more difficult. Alternatively, the decision to run independent component analysis across both patients and controls may have limited the specificity of the resultant components; however, it is unclear why this would affect patients more than controls, as both groups are very nearly the same size. The causes for this poor performance and the discrepancy between datasets may warrant investigation in future studies.

7.3 Review: Schizophrenia, Bipolar Disorder (Type I), and ADHD

7.3.1 Phenomenology

The author is unsure what to make of the elevated entropy across patient groups. Given the alterations concentration in separate components per condition, it is plausible, but nonetheless surprises him. He does not feel qualified to comment on the phenomenological plausibility of such a finding. That LEICA is able to separate all disorders from healthy controls, and to a limited degree from one another, supports the framework's sensitivity to functional alterations.

7.3.2 Connectivity

The recovery of hyperconnectivity in all patient groups is somewhat surprising. Schizophrenia, in particular, has long been believed to be characterized by long-range hypoconnectivity (Friston et al.

2016). Recent studies have begun to contest this hypothesis (Rolls et al. 2020), suggesting that schizophrenia patients may display a mixture of hyper- and hypoconnectivity in long-range connections. Nonetheless, the fact that all significant links in schizophrenia patients are significantly stronger than in controls is, to this author's knowledge, unusual. Lowering the t -threshold of the network-based statistic may reveal weakened connections in schizophrenia and other disorders; however, the universal elevation of in- and out-strengths in schizophrenia patients makes this author suspect otherwise. Unfortunately, the author does not feel qualified to speculate on reasons for this apparent contradiction, and he is still awaiting input from a collaborator with more experience in neuroanatomy and regional function.

7.4 General Comments

The discovery of elevated entropy and connectivity across all four investigated disorders raises intriguing questions regarding dynamics and information flow in psychiatric disorders. While the author is not an academic expert in any single psychiatric disease, much less all four examined in this thesis, the universality of these findings should be confirmed on other datasets. The effect of medication must also be considered, as all patients in this study had been on stable medications for multiple weeks at time of scans.

7.5 Limitations

Three cautionary notes must be added. First, the LEICA method is, by necessity, agnostic as to the true orientation of its communities. Since eigendecomposition and ICA can determine only the orientation of communities relative to each other, not relative to the data itself, LEICA cannot determine which community is “positive” or “negative” in any absolute sense. This may be established by a parallel analysis observing which community is more or less active at any given time; such an analysis is unnecessary for the present purposes.

Second, the Hopf model should not be considered predictive at the level of individual links. It has been able to replicate known phenomena and mechanisms in past studies (Jobst et al. 2017) and brain-level results, e.g. the widespread cortical hyperconnectivity in obsessive-compulsive patients, appear robust. However, a Hopf bifurcation remains an idealized simplification of neural dynamics. To predict neurobiological mechanisms would require both more detailed data and a more sophisticated model, e.g. a model incorporating transmission delays and neuromodulation. Link-level model results in this paper should thus be considered starting points for future research rather than forming hard conclusions themselves.

Finally, while the network-based statistic (NBS) is a well-established method, its results remain dependent on the choice of t -statistic threshold employed. This does not affect the power of the results, only the effect size of the results reported. Unfortunately, no data-driven method has yet been established for determining an appropriate threshold. However, studying which connections survive the different thresholds allows us to partially quantify the group-level effect size.

7.6 Future Steps

7.6.1 Biomarker Research

Beyond the societal and medical costs of such anxiety and compulsive disorders, I am personally aware of the high cost these disorders place upon the families and friends of the afflicted. They do not make for pleasant memories. Better treatment and screening are desperately needed to alleviate the effects of such disorders, and to detect them in the general population. Psychiatric disorders are notoriously underreported and underdiagnosed due to social stigma or simple ignorance. The detection of robust, reliable biomarkers for such diseases would go far towards reducing that underdiagnosis.

To that end, I hope to continue researching biomarkers and possible treatment targets for such disorders in the future. Through Prof. Carles Soriano-Mas, I have discovered several efforts in the psychiatric community to thoroughly map the functional and

structural features of anxiety and compulsion-related disorders, particularly obsessive-compulsive disorder and anorexia nerviosa. Regrettably, I did not discover these efforts in time to incorporate them into my doctoral studies. However, I do hope to become involved in such work in the future. A potential start to such work might involve a translational role, in which LEICA or other methods in the Deco group are applied to the pooled datasets which Prof. Soriano-Mas and collaborators have been compiling.

The explosion of the neuroimaging field has largely extinguished hopes that a single biomarker will be adequate to define a single disorder (Fullana et al. 2020). Instead, research appears to be moving towards the concept of a biomarker space which represents individual patients as vectors of functional and structural features. Such a space, with the exponential increase in information that high dimensionality provides, may allow the delineation of patients into behaviorally and medically distinct groups. The search for this space has been a topic of research for a considerable period of time, but the past decade has demonstrated that some methods can detect behaviorally distinct groups within individual DSM-IV diagnoses (Brodersen et al. 2011; 2014). While isolated for the moment, this finding indicates that such spaces do exist, although to find them requires prudent feature selection.

It would be intriguing to explore which, if any, functional and structural features correlate or otherwise predict the behavioral dimensions hypothesized to describe psychiatric symptoms (Robbins et al. 2012). Some evidence already suggests that this is the case (Xia et al. 2018). While not as immediately clinically useful as a simple test, such a space does have the potential to meaningfully link observed structural and functional features to behavior, a link which has historically been murky at best. Such a map might prove more informative than a simple test, as would directly point to the features which cause behavioral and emotional distress. From such a position, estimating the functional alterations which cause these symptoms is (conceptually) a simple step, and predicting treatment is a natural next step from such estimates. With this in mind, the author hopes to begin surveying such a map between structural, functional, and behavioral features in disorders characterized by anxiety, obsessions, and compulsions.

7.6.2 Consolidation of Resting-State Networks

LEICA is in further need of validation. While it has proven able to find group-level, all studies in this thesis have assumed that patients and controls possess identical components. This assumption has not yet been validated. It is plausible that psychiatric disease so alters the patients' dynamic functional connectivity that patients and controls possess structurally distinct components. The implications of such a finding are beyond the scope of this work (the author has, albeit late, learned to take research one step at a time), yet might provide useful insights towards screening or treatment of the disorders in question.

Additional methods for choosing the number of components are also worthy of consideration. A recent study exploring the efficacy of various clustering validity indices indicates that not all such indices are made equal (Vergara et al. 2020). While the Marčenko-Pasteur method is theoretically sound, its efficacy has not been directly compared to other metrics. Such a study might further validate the LEiDA framework, LEICA included, as a useful tool in the search for functional biomarkers in psychiatric disease. These are projects to pursue in the immediate future and will require minimal adjustments to the current pipeline.

In the less immediate future, the author believes that the field of resting-state networks needs consolidation. To his knowledge, few studies have attempted to compare methods for extracting these resting-state networks directly. Given the sheer number of such methods which are now available, this presents the new researcher with a difficult choice of which method to employ for a given study. It further confuses the search for reliable biomarkers, as the plethora of methods obscures which structural or dynamic functional differences are robust and which are particular to a specific methodology. This problem appears to be gaining attention in the scientific community, as OHBM has recently announced a project aimed at producing a consistent taxonomy of static and dynamic resting-state networks (Uddin et al. 2022). Such a consistent taxonomy, along with the documentation and ratings of methods that the effort will necessitate, may prove a great boon to both

functional neuroimaging in general and the mapping of behavioral dimensions in particular.

7.6.3 Tracing Activity Propagation

The widespread alterations in cortical connectivity likely affect activity propagation and organization. While such alterations were outside the scope of this study, they are of great interest to the understanding of OCD's functionality. Leveraging established network analyses frameworks, such as community detection, node centrality measures, or mappings of node to function, may provide further insights into cortical activity adaptations OCD, and potentially in related disorders such as anxiety and depression. Several frameworks for such studies have been proposed in the Deco group (Klimm et al. 2014; Gilson et al. 2019; 2018); I hope to eventually begin applying them to psychiatric disease directly.

7.7 Final Thoughts

Initial results from this thesis suggest that the LEiDA framework, particularly its LEICA subform, is highly sensitive to the functional alterations representative of psychiatric disease. The ability to distinguish which components display altered functional complexity surpassed initial expectations for the framework, particularly given the high number of components and the resultant need for multiple comparison correction. In addition, the discovery that a simple Hopf bifurcation model is able to recover substantial alterations in effective connectivity adds further evidence that global dynamics can be abstracted from microscale mechanisms, an approach already common in statistical physics. Nonetheless, the imperfect matches between simulated and empirical dynamics suggest that further improvement is possible, and the entire framework should be further tested for robustness and sensitivity on both the same and other datasets. Comparisons with alternative methods should also be considered in order to confirm these results.

Bibliography

- Achard, Sophie, Raymond Salvador, Brandon Whitcer, John Suckling, and Edward T. Bullmore. 2006. "A Resilient, Low-Frequency, Small-World Human Brain Functional Network with Highly Connected Association Cortical Hubs." *Journal of Neuroscience* 26 (1): 63–72.
<https://doi.org/10.1523/JNEUROSCI.3874-05.2006>.
- Adhikari, Mohit H., Carl D. Hacker, Joshua S. Siegel, Alessandra Griffa, Patric Hagmann, Gustavo Deco, and Maurizio Corbetta. 2017. "Decreased Integration and Information Capacity in Stroke Measured by Whole Brain Models of Resting State Activity." *Brain* 140 (4): 1068–85.
<https://doi.org/10.1093/brain/awx021>.
- Afraimovich, V. S., V. P. Zhigulin, and M. I. Rabinovich. 2004. "On the Origin of Reproducible Sequential Activity in Neural Circuits." *Chaos: An Interdisciplinary Journal of Nonlinear Science* 14 (4): 1123. <https://doi.org/10.1063/1.1819625>.
- Alonso, Pino, Arantxa Orbeagozo, Jesús Pujol, Clara López-Solà, Miquel Àngel Fullana, Cinto Segalàs, Eva Real, Marta Subirà, Ignacio Martínez-Zalacaín, José M Menchón, Ben J Harrison, Narcís Cardoner, and Carles Soriano-Mas. 2013. "Neural Correlates of Obsessive-Compulsive Related Dysfunctional Beliefs." *Progress in Neuro-Psychopharmacology and Biological Psychiatry* 47 (December): 25–32.
<https://doi.org/10.1016/j.pnpbp.2013.07.016>.
- Andrews-Hanna, Jessica R., Abraham Z. Snyder, Justin L. Vincent, Cindy Lustig, Denise Head, Marcus E. Raichle, and Randy L. Buckner. 2007. "Disruption of Large-Scale Brain Systems in Advanced Aging." *Neuron* 56 (5).
<https://doi.org/10.1016/j.neuron.2007.10.038>.
- Ashourvan, Arian, Preya Shah, Adam Pines, Shi Gu, Christopher W. Lynn, Danielle S. Bassett, Kathryn A. Davis, and Brian Litt. 2021. "Pairwise Maximum Entropy Model Explains the Role of White Matter Structure in Shaping Emergent Co-Activation States." *Communications Biology* 4 (1).
<https://doi.org/10.1038/s42003-021-01700-6>.
- Atasoy, Selen, Gustavo Deco, Morten L. Kringelbach, and Joel Pearson. 2018. "Harmonic Brain Modes: A Unifying Framework for Linking Space and Time in Brain Dynamics."

- The Neuroscientist* 24 (3): 277–93.
<https://doi.org/10.1177/1073858417728032>.
- Aydin, Serap, Nafiz Arica, Emrah Ergul, and Oğuz Tan. 2015. “Classification of Obsessive Compulsive Disorder by EEG Complexity and Hemispheric Dependency Measurements.” *International Journal of Neural Systems* 25 (03): 1550010. <https://doi.org/10.1142/S0129065715500100>.
- Barabási, Albert-László. 2016. *Network Science. Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*. 1st ed. Cambridge University Press. <http://barabasi.com/book/network-science>.
- Basser, Peter J., J. Mattiello, and D. LeBihan. 1994. “MR Diffusion Tensor Spectroscopy and Imaging.” *Biophysical Journal* 66 (1). [https://doi.org/10.1016/S0006-3495\(94\)80775-1](https://doi.org/10.1016/S0006-3495(94)80775-1).
- Basser, Peter J., Sinisa Pajevic, Carlo Pierpaoli, Jeffrey Duda, and Akram Aldroubi. 2000. “In Vivo Fiber Tractography Using DT-MRI Data.” *Magnetic Resonance in Medicine* 44 (4): 625–32. [https://doi.org/10.1002/1522-2594\(200010\)44:4<625::AID-MRM17>3.0.CO;2-O](https://doi.org/10.1002/1522-2594(200010)44:4<625::AID-MRM17>3.0.CO;2-O).
- Beckmann, Christian F., Marilena DeLuca, Joseph T. Devlin, and Stephen M. Smith. 2005. “Investigations into Resting-State Connectivity Using Independent Component Analysis.” *Philosophical Transactions of the Royal Society B: Biological Sciences* 360 (1457). <https://doi.org/10.1098/rstb.2005.1634>.
- Benjamini, Yoav, and Yosef Hochberg. 1995. “Controlling the False Discovery Rate : A Practical and Powerful Approach to Multiple Testing.” *Journal of the Royal Statistical Society. Series B (Methodological)*. 57 (1): 289–300.
- Bettinardi, R. G., Gustavo Deco, V. M. Karlaftis, Tim J. van Hartevelt, H. M. Fernandes, Z. Kourtzi, Morten L. Kringelbach, and Gorka Zamora-López. 2017. “How Structure Sculpt Function: Unveiling the Contribution of Anatomical Connectivity to the Brain’s Spontaneous Correlation Structure.” *Chaos (Woodbury, N.Y.)* 27 (4): 047409. <https://doi.org/10.1063/1.4980099>.
- Biswal, Bharat B. 2012. “Resting State FMRI: A Personal History.” *NeuroImage* 62 (2): 938–44. <https://doi.org/10.1016/j.neuroimage.2012.01.090>.
- Biswal, Bharat B., Dana A. Eldreth, Michael A. Motes, and Bart Rypma. 2010. “Task-Dependent Individual Differences in

- Prefrontal Connectivity.” *Cerebral Cortex* 20 (9): 2188–97.
<https://doi.org/10.1093/cercor/bhp284>.
- Biswal, Bharat B., A. G. Hudetz, F. Zerrin Yetkin, Victor M. Haughton, and James S. Hyde. 1997. “Hypercapnia Reversibly Suppresses Low-Frequency Fluctuations in the Human Motor Cortex during Rest Using Echo-Planar MRI.” *Journal of Cerebral Blood Flow and Metabolism* 17 (3): 301–8.
<https://doi.org/10.1097/00004647-199703000-00007>.
- Biswal, Bharat B., Joel van Kylen, and James S. Hyde. 1997. “Simultaneous Assessment of Flow and BOLD Signals in Resting-State Functional Connectivity Maps.” *NMR in Biomedicine* 10 (4–5): 165–70.
[https://doi.org/10.1002/\(SICI\)1099-1492\(199706/08\)10:4/5<165::AID-NBM454>3.0.CO;2-7](https://doi.org/10.1002/(SICI)1099-1492(199706/08)10:4/5<165::AID-NBM454>3.0.CO;2-7).
- Biswal, Bharat B., John L. Ulmer, Robert L. Krippendorf, Harold H. Harsch, David L. Daniels, James S. Hyde, and Victor M. Haughton. 1998. “Abnormal Cerebral Activation Associated with a Motor Task in Tourette Syndrome.” *American Journal of Neuroradiology* 19 (8): 1509–12.
[/pmc/articles/PMC8338689/?report=abstract](https://pubmed.ncbi.nlm.nih.gov/9780199562/).
- Biswal, Bharat B., F Z Yetkin, Victor M. Haughton, James S. Hyde, F. Zerrin Yetkin, Victor M. Haughton, and James S. Hyde. 1995. “Functional Connectivity in the Motor Cortex of Resting Human Brain Using Echo-Planar Mri.” *Magnetic Resonance in Medicine* 34 (4): 537–41.
<https://doi.org/10.1002/mrm.1910340409>.
- Blair, David Sutherland, Carles Soriano-Mas, Joana R. B. Cabral, Pedro Moreira, Pedro Morgado, and Gustavo Deco. 2022. “Complexity Changes in Functional State Dynamics Suggest Focal Connectivity Reductions.” *Frontiers in Human Neuroscience* 16 (September): 1–18.
<https://doi.org/10.3389/fnhum.2022.958706>.
- Blundell, Stephen J., and Katherine M. Blundell. 2010. *Concepts in Thermal Physics*. *Concepts in Thermal Physics*. 2nd Editio. Vol. 9780199562. Oxford: Oxford University Press.
<https://doi.org/10.1093/acprof:oso/9780199562091.001.0001>.
- Bonferroni, C. E. 1935. “Il Calcolo Delle Assicurazioni Su Gruppi Di Teste.” In *Studi in Onore Del Professore Salvatore Ortu Carboni*.

- Breakspear, Michael. 2004. "'Dynamic' Connectivity in Neural Systems." *Neuroinformatics* 2 (2): 205–26.
<https://doi.org/10.1385/NI:2:2:205>.
- Breakspear, Michael, John R. Terry, and Karl J. Friston. 2003. "Modulation of Excitatory Synaptic Coupling Facilitates Synchronization and Complex Dynamics in a Nonlinear Model of Neuronal Dynamics." *Neurocomputing* 52–54: 151–58.
[https://doi.org/10.1016/S0925-2312\(02\)00740-3](https://doi.org/10.1016/S0925-2312(02)00740-3).
- Brodersen, Kay Henning, Lorenz Deserno, Florian Schlagenhaut, Zhihao Lin, Will D. Penny, Joachim M. Buhmann, and Klaas Enno Stephan. 2014. "Dissecting Psychiatric Spectrum Disorders by Generative Embedding." *NeuroImage: Clinical* 4: 98–111. <https://doi.org/10.1016/j.nicl.2013.11.002>.
- Brodersen, Kay Henning, Thomas M. Schofield, Alexander P. Leff, Cheng Soon Ong, Ekaterina I. Lomakina, Joachim M. Buhmann, and Klaas Enno Stephan. 2011. "Generative Embedding for Model-Based Classification of fMRI Data." Edited by Olaf Sporns. *PLoS Computational Biology* 7 (6): e1002079. <https://doi.org/10.1371/journal.pcbi.1002079>.
- Brookes, Matthew J., Mark Woolrich, Henry Luckhoo, Darren Price, Joanne R. Hale, Mary C. Stephenson, Gareth R. Barnes, Stephen M. Smith, and Peter G. Morris. 2011. "Investigating the Electrophysiological Basis of Resting State Networks Using Magnetoencephalography." *Proceedings of the National Academy of Sciences of the United States of America* 108 (40): 16783–88. <https://doi.org/10.1073/pnas.1112685108>.
- Buckner, Randy L., Jorge Sepulcre, Tanveer Talukdar, Fenna M. Krienen, Hesheng Liu, Trey Hedden, Jessica R. Andrews-Hanna, Reisa A. Sperling, and Keith A. Johnson. 2009. "Cortical Hubs Revealed by Intrinsic Functional Connectivity: Mapping, Assessment of Stability, and Relation to Alzheimer's Disease." *Journal of Neuroscience* 29 (6): 1860–73.
<https://doi.org/10.1523/JNEUROSCI.5062-08.2009>.
- Cabral, Joana R. B. 2012. "Brain Activity during Rest: A Signature of the Underlying Network Dynamics." *TDx (Tesis Doctorals En Xarxa)*. <http://www.tdx.cat/handle/10803/85414>.
- Cabral, Joana R. B., Etienne Hugues, Morten L. Kringelbach, and Gustavo Deco. 2012. "Modeling the Outcome of Structural Disconnection on Resting-State Functional Connectivity." *NeuroImage* 62 (3): 1342–53.
<https://doi.org/10.1016/j.neuroimage.2012.06.007>.

- Cabral, Joana R. B., Etienne Hugues, Olaf Sporns, and Gustavo Deco. 2011. "Role of Local Network Oscillations in Resting-State Functional Connectivity." *NeuroImage* 57 (1): 130–39. <https://doi.org/10.1016/j.neuroimage.2011.04.010>.
- Cabral, Joana R. B., Henry Luckhoo, Mark Woolrich, Morten Joensson, Hamid Mohseni, Adam Baker, Morten L. Kringelbach, and Gustavo Deco. 2014. "Exploring Mechanisms of Spontaneous Functional Connectivity in MEG: How Delayed Network Interactions Lead to Structured Amplitude Envelopes of Band-Pass Filtered Oscillations." *NeuroImage* 90 (April): 423–35. <https://doi.org/10.1016/j.neuroimage.2013.11.047>.
- Cabral, Joana R. B., Diego Vidaurre, Paulo Marques, Ricardo Magalhães, Pedro Silva Moreira, José Miguel Soares, Gustavo Deco, Nuno Sousa, and Morten L. Kringelbach. 2017. "Cognitive Performance in Healthy Older Adults Relates to Spontaneous Switching between States of Functional Connectivity during Rest." *Scientific Reports* 7 (1): 5135. <https://doi.org/10.1038/s41598-017-05425-7>.
- Calhoun, Vince D., Vamsi K Potluru, Ronald Phlypo, Rogers F Silva, Barak A Pearlmuter, Arvind Caprihan, Sergey M. Plis, and Tülay Adali. 2013. "Independent Component Analysis for Brain FMRI Does Indeed Select for Maximal Independence." Edited by Dante R. Chialvo. *PLoS ONE* 8 (8): e73309. <https://doi.org/10.1371/journal.pone.0073309>.
- Cano, Mónica, Pino Alonso, I. Martínez-Zalacaín, M. Subirà, E. Real, Cinto Segalàs, Jesús Pujol, Narcís Cardoner, J. M. Menchón, and Carles Soriano-Mas. 2018. "Altered Functional Connectivity of the Subthalamus and the Bed Nucleus of the Stria Terminalis in Obsessive-Compulsive Disorder." *Psychological Medicine* 48 (6): 919–28. <https://doi.org/10.1017/S0033291717002288>.
- Christoff, Kalina, Zachary C. Irving, Kieran C.R. Fox, R. Nathan Spreng, and Jessica R. Andrews-Hanna. 2016. "Mind-Wandering as Spontaneous Thought: A Dynamic Framework." *Nature Reviews Neuroscience* 17 (11): 718–31. <https://doi.org/10.1038/nrn.2016.113>.
- "Correction for Deco et al., Key Role of Coupling, Delay, and Noise in Resting Brain Fluctuations." 2009. *Proceedings of the National Academy of Sciences* 106 (29). <https://doi.org/10.1073/pnas.0906701106>.

- Cover, Thomas M., and Joy A. Thomas. 2005. *Elements of Information Theory*. *Elements of Information Theory*. <https://doi.org/10.1002/047174882X>.
- Damoiseaux, Jessica S., S. A. R. B. Rombouts, F. Barkhof, P. Scheltens, Cornelis J. Stam, Stephen M. Smith, and Christian F. Beckmann. 2006. "Consistent Resting-State Networks across Healthy Subjects." *Proceedings of the National Academy of Sciences* 103 (37): 13848–53. <https://doi.org/10.1073/pnas.0601417103>.
- Deco, Gustavo, Joana R. B. Cabral, Mark W. Woolrich, Angus B.A. Stevner, Tim J. van Hartevelt, and Morten L. Kringelbach. 2017. "Single or Multiple Frequency Generators in On-Going Brain Activity: A Mechanistic Whole-Brain Model of Empirical MEG Data." *NeuroImage* 152 (February): 538–50. <https://doi.org/10.1016/j.neuroimage.2017.03.023>.
- Deco, Gustavo, Josefina Cruzat Grand, Joana R. B. Cabral, Enzo Tagliazucchi, Helmut Laufs, Nikos K. Logothetis, and Morten L. Kringelbach. 2019. "Awakening: Predicting External Stimulation to Force Transitions between Different Brain States." *Proceedings of the National Academy of Sciences*, 201905534. <https://doi.org/10.1073/pnas.1905534116>.
- Deco, Gustavo, Josefina Cruzat Grand, and Morten L. Kringelbach. 2019. "Brain Songs Framework Used for Discovering the Relevant Timescale of the Human Brain." *Nature Communications* 10 (1): 1–13. <https://doi.org/10.1038/s41467-018-08186-7>.
- Deco, Gustavo, Viktor Jirs, A. R. McIntosh, Olaf Sporns, and Rolf Kötter. 2009. "Key Role of Coupling, Delay, and Noise in Resting Brain Fluctuations." *Proceedings of the National Academy of Sciences of the United States of America* 106 (25). <https://doi.org/10.1073/pnas.0901831106>.
- Deco, Gustavo, and Morten L. Kringelbach. 2016. "Metastability and Coherence: Extending the Communication through Coherence Hypothesis Using A Whole-Brain Computational Perspective." *Trends in Neurosciences* 39 (3): 125–35. <https://doi.org/10.1016/j.tins.2016.01.001>.
- . 2020. "Turbulent-like Dynamics in the Human Brain." *Cell Reports* 33 (10): 108471. <https://doi.org/10.1016/j.celrep.2020.108471>.
- Deco, Gustavo, Morten L. Kringelbach, Viktor K. Jirsa, and Petra Ritter. 2017. "The Dynamics of Resting Fluctuations in the

- Brain: Metastability and Its Dynamical Cortical Core.” *Scientific Reports* 7 (1): 3095. <https://doi.org/10.1038/s41598-017-03073-5>.
- Deco, Gustavo, Yonatan Sanz Perl, Jacobo D. Sitt, Enzo Tagliazucchi, and Morten L. Kringelbach. 2021. “Deep Learning the Arrow of Time in Brain Activity: Characterising Brain-Environment Behavioural Interactions in Health and Disease.” *BioRxiv*, 2021.07.02.450899. <https://www.biorxiv.org/content/10.1101/2021.07.02.450899v1%0Ahttps://www.biorxiv.org/content/10.1101/2021.07.02.450899v1.abstract>.
- Deco, Gustavo, Adrián Ponce-Alvarez, Dante Mantini, Gian Luca Romani, Patric Hagmann, and Maurizio Corbetta. 2013. “Resting-State Functional Connectivity Emerges from Structurally and Dynamically Shaped Slow Linear Fluctuations.” *Journal of Neuroscience* 33 (27): 11239–52. <https://doi.org/10.1523/JNEUROSCI.1091-13.2013>.
- Delattre, Sylvain, and Nicolas Fournier. 2017. “On the Kozachenko–Leonenko Entropy Estimator.” *Journal of Statistical Planning and Inference* 185 (June): 69–93. <https://doi.org/10.1016/j.jspi.2017.01.004>.
- Deng, Shikuang, Jingwei Li, B. T. Thomas Yeo, and Shi Gu. 2022. “Control Theory Illustrates the Energy Efficiency in the Dynamic Reconfiguration of Functional Connectivity.” *Communications Biology* 5 (1). <https://doi.org/10.1038/s42003-022-03196-0>.
- Desikan, Rahul S., Florent Ségonne, Bruce Fischl, Brian T. Quinn, Bradford C. Dickerson, Deborah Blacker, Randy L. Buckner, Anders M. Dale, R. Paul Maguire, Bradley T. Hyman, Marilyn S. Albert, and Ronald J. Killiany. 2006a. “An Automated Labeling System for Subdividing the Human Cerebral Cortex on MRI Scans into Gyral Based Regions of Interest.” *NeuroImage* 31 (3): 968–80. <https://doi.org/10.1016/j.neuroimage.2006.01.021>.
- . 2006b. “An Automated Labeling System for Subdividing the Human Cerebral Cortex on MRI Scans into Gyral Based Regions of Interest.” *NeuroImage* 31 (3): 968–80. <https://doi.org/10.1016/j.neuroimage.2006.01.021>.
- Du, Yuhui, Xingyu He, and Vince D. Calhoun. 2021. “A New Semi-Supervised Non-Negative Matrix Factorization Method For Brain Dynamic Functional Connectivity Analysis.” In

- 2021 *IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, 2021-April:1591–94. IEEE.
<https://doi.org/10.1109/ISBI48211.2021.9433988>.
- Dudewicz, Edward J., and Edward C. van der Meulen. 1981.
 “Entropy-Based Tests of Uniformity.” *Journal of the American Statistical Association* 76 (376): 967–74.
<https://doi.org/10.1080/01621459.1981.10477750>.
- DuPont, Robert L., Dorothy P. Rice, S Shiraki, and Clayton R. Rowland. 1995. “Economic Costs of Obsessive-Compulsive Disorder.” *Medical Interface* 8 (4): 102–9.
- Engel, Andreas K., Pascal Fries, and Wolf Singer. 2001. “Dynamic Predictions: Oscillations and Synchrony in Top-down Processing.” *Nature Reviews Neuroscience* 2 (10).
<https://doi.org/10.1038/35094565>.
- Erik, Magnus, Hvass Pedersen, and Magnus Erik Hvass Pedersen. 2010. “Good Parameters for Particle Swarm Optimization.” *Technical Report HL1001, Hvass Laboratories HL1001*: 1–12.
- Etkin, Amit, Christian Büchel, and James J. Gross. 2015. “The Neural Bases of Emotion Regulation.” *Nature Reviews Neuroscience* 16 (11): 693–700.
<https://doi.org/10.1038/nrn4044>.
- Etkin, Amit, and Tor D. Wager. 2007. “Functional Neuroimaging of Anxiety: A Meta-Analysis of Emotional Processing in PTSD, Social Anxiety Disorder, and Specific Phobia.” *American Journal of Psychiatry* 164 (10): 1476–88.
<https://doi.org/10.1176/appi.ajp.2007.07030504>.
- Figueroa, Caroline A., Joana R. B. Cabral, Roel J.T. Mocking, Kristina M. Rapuano, Tim J. van Hartevelt, Gustavo Deco, Paul Expert, Aart H. Schene, Morten L. Kringelbach, and Henricus G. Ruhé. 2019. “Altered Ability to Access a Clinically Relevant Control Network in Patients Remitted from Major Depressive Disorder.” *Human Brain Mapping* 40 (9): 2771–86. <https://doi.org/10.1002/hbm.24559>.
- FitzHugh, Richard. 1961. “Impulses and Physiological States in Theoretical Models of Nerve Membrane.” *Biophysical Journal* 1 (6): 445–66. [https://doi.org/10.1016/S0006-3495\(61\)86902-6](https://doi.org/10.1016/S0006-3495(61)86902-6).
- Fox, Michael D., Abraham Z. Snyder, Justin L. Vincent, Maurizio Corbetta, David C van Essen, and Marcus E. Raichle. 2005. “The Human Brain Is Intrinsically Organized into Dynamic, Anticorrelated Functional Networks.” *Proceedings of the*

- National Academy of Sciences of the United States of America* 102 (27): 9673–78. <https://doi.org/10.1073/pnas.0504136102>.
- Freyer, Frank, James A. Roberts, Robert Becker, Peter A. Robinson, Petra Ritter, and Michael Breakspear. 2011. “Biophysical Mechanisms of Multistability in Resting-State Cortical Rhythms.” *Journal of Neuroscience* 31 (17): 6353–61. <https://doi.org/10.1523/JNEUROSCI.6693-10.2011>.
- Freyer, Frank, James A. Roberts, Petra Ritter, and Michael Breakspear. 2012. “A Canonical Model of Multistability and Scale-Invariance in Biological Systems.” Edited by Tim Behrens. *PLoS Computational Biology* 8 (8): e1002634. <https://doi.org/10.1371/journal.pcbi.1002634>.
- Fries, Pascal. 2005. “A Mechanism for Cognitive Dynamics: Neuronal Communication through Neuronal Coherence.” *Trends in Cognitive Sciences* 9 (10): 474–80. <https://doi.org/10.1016/j.tics.2005.08.011>.
- Friston, Karl J. 1994. “Functional and Effective Connectivity in Neuroimaging: A Synthesis.” *Human Brain Mapping* 2 (1–2): 56–78. <https://doi.org/10.1002/hbm.460020107>.
- . 2011. “Functional and Effective Connectivity: A Review.” *Brain Connectivity* 1 (1): 13–36. <https://doi.org/10.1089/brain.2011.0008>.
- Friston, Karl J., Harriet R. Brown, Jakob Siemerikus, and Klaas Enno Stephan. 2016. “The Dysconnection Hypothesis (2016).” *Schizophrenia Research* 176 (2–3): 83–94. <https://doi.org/10.1016/j.schres.2016.07.014>.
- Fullana, Miquel A., Amitai Abramovitch, Esther Via, Clara López-Sola, Ximena Goldberg, Nuria Reina, Lydia Fortea, Aleix Solanes, Matthew J. Buckley, Valentina Ramella-Cravaro, André F. Carvalho, Miquel Tortella-Feliu, Eduard Vieta, Carles Soriano-Mas, Luisa Lázaro, Dan J. Stein, Lorena Fernández de la Cruz, David Mataix-Cols, and Joaquim Radua. 2020. “Diagnostic Biomarkers for Obsessive-Compulsive Disorder: A Reasonable Quest or Ignis Fatuus?” *Neuroscience & Biobehavioral Reviews* 118 (November): 504–13. <https://doi.org/10.1016/J.NEUBIOREV.2020.08.008>.
- Ghosh, Anandamohan, Y. Rho, A. R. McIntosh, R. Kötter, and Viktor K. Jirsa. 2008a. “Cortical Network Dynamics with Time Delays Reveals Functional Connectivity in the Resting Brain.” *Cognitive Neurodynamics* 2 (2): 115–20. <https://doi.org/10.1007/s11571-008-9044-2>.

- . 2008b. “Noise during Rest Enables the Exploration of the Brain’s Dynamic Repertoire.” Edited by Karl J. Friston. *PLoS Computational Biology* 4 (10): e1000196.
<https://doi.org/10.1371/journal.pcbi.1000196>.
- Gilson, Matthieu, Nikos E. Kouvaris, Gustavo Deco, and Gorka Zamora-López. 2018. “Framework Based on Communicability and Flow to Analyze Complex Network Dynamics.” *Physical Review E* 97 (5): 052301.
<https://doi.org/10.1103/PhysRevE.97.052301>.
- Gilson, Matthieu, Rubén Moreno-Bote, Adrián Ponce-Alvarez, Petra Ritter, and Gustavo Deco. 2016. “Estimation of Directed Effective Connectivity from fMRI Functional Connectivity Hints at Asymmetries of Cortical Connectome.” *PLoS Computational Biology*.
<https://doi.org/10.1371/journal.pcbi.1004762>.
- Gilson, Matthieu, Gorka Zamora-López, Vicente Pallarés, Mohit H. Adhikari, Mario Senden, Adrià Tauste Campo, Dante Mantini, Maurizio Corbetta, Gustavo Deco, and Andrea Insabato. 2019. “Model-Based Whole-Brain Effective Connectivity to Study Distributed Cognition in Health and Disease.” *Network Neuroscience* 4 (2): 338–73.
https://doi.org/10.1162/netn_a_00117.
- Glerean, Enrico, Juha Salmi, Juha M. Lahnakoski, Iiro P. Jääskeläinen, and Mikko Sams. 2012. “Functional Magnetic Resonance Imaging Phase Synchronization as a Measure of Dynamic Functional Connectivity.” *Brain Connectivity* 2 (2): 91–101. <https://doi.org/10.1089/brain.2011.0068>.
- Glomb, Katharina, Adrián Ponce-Alvarez, Matthieu Gilson, Petra Ritter, and Gustavo Deco. 2017. “Resting State Networks in Empirical and Simulated Dynamic Functional Connectivity.” *NeuroImage* 159 (November 2016): 388–402.
<https://doi.org/10.1016/j.neuroimage.2017.07.065>.
- Goodman, Wayne K., Lawrence H. Price, Steven A. Rasmussen, Carolyn Mazure, Roberta L. Fleischmann, Candy L. Hill, George R. Heninger, and Dennis S. Charney. 1989. “The Yale-Brown Obsessive Compulsive Scale: I. Development, Use, and Reliability.” *Archives of General Psychiatry*.
<https://doi.org/10.1001/archpsyc.1989.01810110048007>.
- Goodstein, David L. 1989. “Richard P. Feynman Teacher.” *Physics Today* 42 (2). <https://doi.org/10.1063/1.881195>.

- Goria, M. N., N. N. Leonenko, V. v. Mergel, and P. L. Novi Inverardi. 2005. "A New Class of Random Vector Entropy Estimators and Its Applications in Testing Statistical Hypotheses." *Journal of Nonparametric Statistics* 17 (3): 277–97. <https://doi.org/10.1080/104852504200026815>.
- Göttlich, Martin, Ulrike M. Krämer, Andreas Kordon, Fritz Hohagen, and Bartosz Zurowski. 2014. "Decreased Limbic and Increased Fronto-Parietal Connectivity in Unmedicated Patients with Obsessive-Compulsive Disorder." *Human Brain Mapping* 35 (11): 5617–32. <https://doi.org/10.1002/hbm.22574>.
- Greicius, Michael D., Ben Krasnow, Allan L. Reiss, and Vinod Menon. 2003. "Functional Connectivity in the Resting Brain: A Network Analysis of the Default Mode Hypothesis." *Proceedings of the National Academy of Sciences of the United States of America* 100 (1): 253–58. <https://doi.org/10.1073/pnas.0135058100>.
- Gu, Shi, Matthew Cieslak, Benjamin Baird, Sarah Feldt Muldoon, Scott T. Grafton, Fabio Pasqualetti, and Danielle S. Bassett. 2018. "The Energy Landscape of Neurophysiological Activity Implicit in Brain Network Structure." *Scientific Reports* 8 (1): 2507. <https://doi.org/10.1038/s41598-018-20123-8>.
- Gusnard, Debra A., and Marcus E. Raichle. 2001. "Searching for a Baseline: Functional Imaging and the Resting Human Brain." *Nature Reviews Neuroscience* 2 (10): 685–94. <https://doi.org/10.1038/35094500>.
- Gyorgy, Buzsaki, and Draguhn Andreas. 2004. "Neuronal Oscillations in Cortical Networks." *Science* 304 (June).
- HAMILTON, M. 1959. "The Assessment of Anxiety States by Rating." *The British Journal of Medical Psychology* 32 (1): 50–55. <https://doi.org/10.1111/j.2044-8341.1959.tb00467.x>.
- Hansen, Enrique C.A., Demian Battaglia, Andreas Spiegler, Gustavo Deco, and Viktor K. Jirsa. 2015. "Functional Connectivity Dynamics: Modeling the Switching Behavior of the Resting State." *NeuroImage* 105. <https://doi.org/10.1016/j.neuroimage.2014.11.001>.
- Heuvel, Odile A. van den, Guido van Wingen, Carles Soriano-Mas, Pino Alonso, Samuel R. Chamberlain, Takashi Nakamae, Damiaan Denys, Anna E. Goudriaan, and Dick J. Veltman. 2016. "Brain Circuitry of Compulsivity." *European*

- Neuropsychopharmacology* 26 (5): 810–27.
<https://doi.org/10.1016/j.euroneuro.2015.12.005>.
- Hipp, Joerg F., David J. Hawellek, Maurizio Corbetta, Markus Siegel, and Andreas K. Engel. 2012. “Large-Scale Cortical Correlation Structure of Spontaneous Oscillatory Activity.” *Nature Neuroscience* 15 (6). <https://doi.org/10.1038/nn.3101>.
- Honey, Christopher J., Rolf Kötter, Michael Breakspear, and Olaf Sporns. 2007. “Network Structure of Cerebral Cortex Shapes Functional Connectivity on Multiple Time Scales.” *Proceedings of the National Academy of Sciences of the United States of America* 104 (24).
<https://doi.org/10.1073/pnas.0701519104>.
- Honey, Christopher J., Olaf Sporns, Leila Cammoun, Xavier Gigandet, Jean Philippe Thiran, Reto Meuli, and Patric Hagmann. 2009. “Predicting Human Resting-State Functional Connectivity from Structural Connectivity.” *Proceedings of the National Academy of Sciences of the United States of America* 106 (6): 2035–40.
<https://doi.org/10.1073/pnas.0811168106>.
- Hyvärinen, Aapo. 1999. “Fast and Robust Fixed-Point Algorithms for Independent Component Analysis.” *IEEE Transactions on Neural Networks* 10 (3): 626–34.
<https://doi.org/10.1109/72.761722>.
- Hyvärinen, Aapo, and Erkki Oja. 2000. “Independent Component Analysis: Algorithms and Applications.” *Neural Networks*. Vol. 13. [https://doi.org/10.1016/S0893-6080\(00\)00026-5](https://doi.org/10.1016/S0893-6080(00)00026-5).
- Jobst, Beatrice M., Rikkert Hindriks, Helmut Laufs, Enzo Tagliazucchi, Gerald Hahn, Adrián Ponce-Alvarez, Angus B. A. Stevner, Morten L. Kringelbach, and Gustavo Deco. 2017. “Increased Stability and Breakdown of Brain Effective Connectivity During Slow-Wave Sleep: Mechanistic Insights from Whole-Brain Computational Modelling.” *Scientific Reports* 7 (1): 4634. <https://doi.org/10.1038/s41598-017-04522-x>.
- Kennedy, James, and Russell Eberhart. 1995. “Particle Swarm Optimization.” In *Proceedings of ICNN’95 - International Conference on Neural Networks*, 4:1942–48. IEEE.
<https://doi.org/10.1109/ICNN.1995.488968>.
- Khanna, Arjun, Alvaro Pascual-Leone, Christoph M. Michel, and Faranak Farzan. 2015. “Microstates in Resting-State EEG: Current Status and Future Directions.” *Neuroscience &*

- Biobehavioral Reviews* 49 (February): 105–13.
<https://doi.org/10.1016/j.neubiorev.2014.12.010>.
- Klimm, Florian, Javier Borge-Holthoefer, Niels Wessel, Jürgen Kurths, and Gorka Zamora-López. 2014. “Individual Nodes Contribution to the Mesoscale of Complex Networks.” *New Journal of Physics* 16 (125006): 1–25.
<https://doi.org/10.1088/1367-2630/16/12/125006>.
- Kötter, Rolf. 2004. “Online Retrieval, Processing, and Visualization of Primate Connectivity Data from the CoCoMac Database.” *Neuroinformatics* 2 (2). <https://doi.org/10.1385/NI:2:2:127>.
- Kringelbach, Morten L., and Gustavo Deco. 2020. “Brain States and Transitions: Insights from Computational Neuroscience.” *Cell Reports* 32 (10): 108128.
<https://doi.org/10.1016/j.celrep.2020.108128>.
- Krol, Laurens R. 2021. “Permutation Test.” Berlin: GitHub.
- Kucyi, Aaron. 2018. “Just a Thought: How Mind-Wandering Is Represented in Dynamic Brain Connectivity.” *NeuroImage* 180 (October): 505–14.
<https://doi.org/10.1016/J.NEUROIMAGE.2017.07.001>.
- Kuramoto, Y. 1983. *Chemical Oscillations, Waves and Turbulence*. Springer, Berlin.
- Lanciego, Jose L., and Floris G. Wouterlood. 2020. “Neuroanatomical Tract-Tracing Techniques That Did Go Viral.” *Brain Structure and Function* 225 (4): 1193–1224.
<https://doi.org/10.1007/s00429-020-02041-6>.
- Larson-Prior, Linda J., Jonathan D. Power, Justin L. Vincent, Tracy S. Nolan, Rebecca S. Coalson, John Zempel, Abraham Z. Snyder, Bradley L. Schlaggar, Marcus E. Raichle, and Steven E. Petersen. 2011. “Modulation of the Brain’s Functional Network Architecture in the Transition from Wake to Sleep.” In *Progress in Brain Research*. Vol. 193.
<https://doi.org/10.1016/B978-0-444-53839-0.00018-1>.
- Larter, Raima, Brent Speelman, and Robert M. Worth. 1999. “A Coupled Ordinary Differential Equation Lattice Model for the Simulation of Epileptic Seizures.” *Chaos* 9 (3).
<https://doi.org/10.1063/1.166453>.
- Laubach, Mark, Marshall Shuler, and Miguel A.L. Nicolelis. 1999. “Independent Component Analyses for Quantifying Neuronal Ensemble Interactions.” *Journal of Neuroscience Methods* 94 (1): 141–54. [https://doi.org/10.1016/S0165-0270\(99\)00131-4](https://doi.org/10.1016/S0165-0270(99)00131-4).

- Laufs, H., K. Krakow, P. Sterzer, E. Eger, A. Beyerle, A. Salek-Haddadi, and Andreas Kleinschmidt. 2003. "Electroencephalographic Signatures of Attentional and Cognitive Default Modes in Spontaneous Brain Activity Fluctuations at Rest." *Proceedings of the National Academy of Sciences of the United States of America* 100 (19). <https://doi.org/10.1073/pnas.1831638100>.
- Leicht, E. A., and Mark E. J. Newman. 2008. "Community Structure in Directed Networks." *Physical Review Letters* 100 (11): 1–4. <https://doi.org/10.1103/PhysRevLett.100.118703>.
- Lenhard, Fabian, Kristina Aspvall, Erik Andersson, Johan Ahlen, Eva Serlachius, Malin Lavner, Anna Brodin, and David Mataix-Cols. 2021a. "The Cost of Obsessive–Compulsive Disorder in Swedish Youth." *Child Psychiatry and Human Development*, September. <https://doi.org/10.1007/s10578-021-01261-z>.
- . 2021b. "The Cost of Obsessive–Compulsive Disorder in Swedish Youth." *Child Psychiatry & Human Development*, September. <https://doi.org/10.1007/s10578-021-01261-z>.
- Leonardi, Nora, Jonas Richiardi, Markus Gschwind, Samanta Simioni, Jean Marie Annoni, Myriam Schlupe, Patrik Vuilleumier, and Dimitri van de Ville. 2013. "Principal Components of Functional Connectivity: A New Approach to Study Dynamic Brain Connectivity during Rest." *NeuroImage* 83 (December): 937–50. <https://doi.org/10.1016/j.neuroimage.2013.07.019>.
- Li, Xuanyu, Zhaojun Zhu, Weina Zhao, Yu Sun, Dong Wen, Yunyan Xie, Xiangyu Liu, Haijing Niu, and Ying Han. 2018. "Decreased Resting-State Brain Signal Complexity in Patients with Mild Cognitive Impairment and Alzheimer's Disease: A Multi-Scale Entropy Analysis." *Biomedical Optics Express* 9 (4): 1916. <https://doi.org/10.1364/BOE.9.001916>.
- Liu, Collin Y., Anitha P. Krishnan, Lirong Yan, Robert X. Smith, Emily Kilroy, Jeffery R. Alger, John M. Ringman, and Danny J.J. Wang. 2013. "Complexity and Synchronicity of Resting State Blood Oxygenation Level-Dependent (BOLD) Functional MRI in Normal Aging and Cognitive Decline." *Journal of Magnetic Resonance Imaging* 38 (1): 36–45. <https://doi.org/10.1002/jmri.23961>.
- Liu, Hesheng, Steven M. Stufflebeam, Jorge Sepulcre, Trey Hedden, and Randy L. Buckner. 2009. "Evidence from

- Intrinsic Activity That Asymmetry of the Human Brain Is Controlled by Multiple Factors.” *Proceedings of the National Academy of Sciences of the United States of America* 106 (48). <https://doi.org/10.1073/pnas.0908073106>.
- Liu, Zhongming, Masaki Fukunaga, Jacco A. de Zwart, and Jeff H. Duyn. 2010. “Large-Scale Spontaneous Fluctuations and Correlations in Brain Electrical Activity Observed with Magnetoencephalography.” *NeuroImage* 51 (1). <https://doi.org/10.1016/j.neuroimage.2010.01.092>.
- Lopes-dos-Santos, Vítor, Sergio Conde-Ocazionez, Miguel A.L. Nicolelis, Sidarta T. Ribeiro, and Adriano B.L. Tort. 2011. “Neuronal Assembly Detection and Cell Membership Specification by Principal Component Analysis.” *PLoS ONE* 6 (6). <https://doi.org/10.1371/journal.pone.0020996>.
- Lopes-dos-Santos, Vítor, Sidarta T. Ribeiro, and Adriano B.L. Tort. 2013. “Detecting Cell Assemblies in Large Neuronal Populations.” *Journal of Neuroscience Methods* 220 (2): 149–66. <https://doi.org/10.1016/j.jneumeth.2013.04.010>.
- Lord, Louis-David, Paul Expert, Selen Atasoy, Leor Roseman, Kristina Rapuano, Renaud Lambiotte, David J. Nutt, Gustavo Deco, Robin L. Carhart-Harris, Morten L. Kringelbach, and Joana R. B. Cabral. 2019. “Dynamical Exploration of the Repertoire of Brain Networks at Rest Is Modulated by Psilocybin.” *NeuroImage* 199 (April): 127–42. <https://doi.org/10.1016/j.neuroimage.2019.05.060>.
- Luca, Marilena de, Christian F. Beckmann, N. de Stefano, P. M. Matthews, and Stephen M. Smith. 2006. “fMRI Resting State Networks Define Distinct Modes of Long-Distance Interactions in the Human Brain.” *NeuroImage* 29 (4): 1359–67. <https://doi.org/10.1016/j.neuroimage.2005.08.035>.
- Mantini, Dante, M. G. Perrucci, C. del Gratta, G. L. Romani, and Maurizio Corbetta. 2007. “Electrophysiological Signatures of Resting State Networks in the Human Brain.” *Proceedings of the National Academy of Sciences of the United States of America* 104 (32): 13170–75. <https://doi.org/10.1073/pnas.0700668104>.
- Marčenko, V A, and L A Pastur. 1967. “DISTRIBUTION OF EIGENVALUES FOR SOME SETS OF RANDOM MATRICES.” *Mathematics of the USSR-Sbornik* 1 (4): 457–83. <https://doi.org/10.1070/SM1967v001n04ABEH001994>.

- Martín, Alberto Vela. 2019. “Entropy, Chaos and Irreversibility in the Turbulence Energy Cascade.”
- Mezura-Montes, Efrén, Carlos Coello, Efrén Mezura, and Carlos Coello. 2011. “Constraint-Handling in Nature-Inspired Numerical Optimization: Past, Present and Future.” *Swarm and Evolutionary Computation* 1 (4): 173–94.
<https://doi.org/10.1016/j.swevo.2011.10.001>.
- Miller, Kai J., Kurt E. Weaver, and Jeffrey G. Ojemann. 2009. “Direct Electrophysiological Measurement of Human Default Network Areas.” *Proceedings of the National Academy of Sciences of the United States of America* 106 (29).
<https://doi.org/10.1073/pnas.0902071106>.
- Moreira, Pedro, Paulo Marques, Carles Soriano-Mas, Ricardo Magalhães, Nuno Sousa, José Soares, and Pedro Morgado. 2017. “The Neural Correlates of Obsessive-Compulsive Disorder: A Multimodal Perspective.” *Translational Psychiatry* 7 (8): e1224. <https://doi.org/10.1038/tp.2017.189>.
- Morris, C., and H. Lecar. 1981. “Voltage Oscillations in the Barnacle Giant Muscle Fiber.” *Biophysical Journal* 35 (1).
[https://doi.org/10.1016/S0006-3495\(81\)84782-0](https://doi.org/10.1016/S0006-3495(81)84782-0).
- Moussa, Malaak N., Matthew R. Steen, Paul J. Laurienti, and Satoru Hayasaka. 2012. “Consistency of Network Modules in Resting-State fMRI Connectome Data.” Edited by Yu-Feng Zang. *PLoS ONE* 7 (8): e44428.
<https://doi.org/10.1371/journal.pone.0044428>.
- Nagumo, J., S. Arimoto, and S. Yoshizawa. 1962. “An Active Pulse Transmission Line Simulating Nerve Axon*.” *Proceedings of the IRE* 50 (10).
<https://doi.org/10.1109/JRPROC.1962.288235>.
- Newman, Mark E. J. 2006. “Modularity and Community Structure in Networks.” *Proceedings of the National Academy of Sciences of the United States of America* 103 (23): 8577–82.
<https://doi.org/10.1073/pnas.0601602103>.
- . 2014. *Networks: An Introduction*. Oxford University.
<https://doi.org/10.1007/978-3-319-03518-5-8>.
- Nikouline, Vadim v., Klaus Linkenkaer-Hansen, Juha Huttunen, and Risto J. Ilmoniemi. 2001. “Interhemispheric Phase Synchrony and Amplitude Correlation of Spontaneous Beta Oscillations in Human Subjects: A Magnetoencephalographic Study.” *NeuroReport* 12 (11).
<https://doi.org/10.1097/00001756-200108080-00040>.

- Ochsner, Kevin N., Jennifer A. Silvers, and Jason T. Buhle. 2012. "Functional Imaging Studies of Emotion Regulation: A Synthetic Review and Evolving Model of the Cognitive Control of Emotion." *Annals of the New York Academy of Sciences* 1251 (1): E1–24. <https://doi.org/10.1111/j.1749-6632.2012.06751.x>.
- Palva, Satu, and J. Matias Palva. 2012. "Discovering Oscillatory Interaction Networks with M/EEG: Challenges and Breakthroughs." *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2012.02.004>.
- Pasquale, Francesco de, Stefania della Penna, Abraham Z. Snyder, Christopher Lewis, Dante Mantini, Laura Marzetti, Paolo Belardinelli, Luca Ciancetta, Vittorio Pizzella, Gian Luca Romani, and Maurizio Corbetta. 2010. "Temporal Dynamics of Spontaneous MEG Activity in Brain Networks." *Proceedings of the National Academy of Sciences of the United States of America* 107 (13). <https://doi.org/10.1073/pnas.0913863107>.
- Peyrache, Adrien, Karim Benchenane, Mehdi Khamassi, Sidney I. Wiener, and Francesco P. Battaglia. 2010. "Principal Component Analysis of Ensemble Recordings Reveals Cell Assemblies at High Temporal Resolution." *Journal of Computational Neuroscience* 29 (1–2): 309–25. <https://doi.org/10.1007/s10827-009-0154-6>.
- Peyrache, Adrien, Mehdi Khamassi, Karim Benchenane, Sidney I. Wiener, and Francesco P. Battaglia. 2009. "Replay of Rule-Learning Related Neural Patterns in the Prefrontal Cortex during Sleep." *Nature Neuroscience* 12 (7): 919–26. <https://doi.org/10.1038/nn.2337>.
- Piacentini, John, R Lindsey Bergman, Melody Keller, and James McCracken. 2003. "Functional Impairment in Children and Adolescents with Obsessive-Compulsive Disorder." *Journal of Child and Adolescent Psychopharmacology* 13 (supplement 1): 61–69. <https://doi.org/10.1089/104454603322126359>.
- Picó-Pérez, Maria, Joaquim Radua, Trevor Steward, José M. Menchón, and Carles Soriano-Mas. 2017. "Emotion Regulation in Mood and Anxiety Disorders: A Meta-Analysis of fMRI Cognitive Reappraisal Studies." *Progress in Neuro-Psychopharmacology and Biological Psychiatry* 79 (May): 96–104. <https://doi.org/10.1016/j.pnpbp.2017.06.001>.

- Poldrack, Russell A., E. Congdon, W. Triplett, K J Gorgolewski, K H Karlsgodt, J A Mumford, F W Sabb, N B Freimer, E D London, T D Cannon, and R M Bilder. 2016. "A Phenome-Wide Examination of Neural and Cognitive Function." *Scientific Data* 3 (1): 160110. <https://doi.org/10.1038/sdata.2016.110>.
- Poli, Riccardo, James Kennedy, and Tim Blackwell. 2007. "Particle Swarm Optimization: An Overview." *Swarm Intelligence* 1 (1): 33–57. <https://doi.org/10.1007/s11721-007-0002-0>.
- Ponce-Alvarez, Adrián, Gustavo Deco, Patric Hagmann, Gian Luca Romani, Dante Mantini, and Maurizio Corbetta. 2015. "Resting-State Temporal Synchronization Networks Emerge from Connectivity Topology and Heterogeneity." *PLoS Computational Biology* 11 (2): 1–23. <https://doi.org/10.1371/journal.pcbi.1004100>.
- Preti, Maria Giulia, Thomas A.W. Bolton, and Dimitri van de Ville. 2017. "The Dynamic Functional Connectome: State-of-the-Art and Perspectives." *NeuroImage* 160 (October): 41–54. <https://doi.org/10.1016/j.neuroimage.2016.12.061>.
- Pujol, Jesús, Laura Blanco-Hinojo, Dídac MacIá, Pino Alonso, Ben J. Harrison, Gerard Martínez-Vilavella, Joan Deus, José M. Menchón, Narcís Cardoner, and Carles Soriano-Mas. 2019. "Mapping Alterations of the Functional Structure of the Cerebral Cortex in Obsessive–Compulsive Disorder." *Cerebral Cortex* 29 (11): 4753–62. <https://doi.org/10.1093/CERCOR/BHZ008>.
- Rabinovich, Mikhail I., Ramón Huerta, Pablo Varona, and Valentin S. Afraimovich. 2008. "Transient Cognitive Dynamics, Metastability, and Decision Making." *PLOS Computational Biology* 4 (5): e1000072. <https://doi.org/10.1371/JOURNAL.PCBI.1000072>.
- Raichle, Marcus E., Ann Mary MacLeod, Abraham Z. Snyder, William J. Powers, Debra A. Gusnard, and Gordon L. Shulman. 2001. "A Default Mode of Brain Function." *Proceedings of the National Academy of Sciences of the United States of America* 98 (2): 676–82. <https://doi.org/10.1073/pnas.98.2.676>.
- Robbins, Trevor W., Claire M. Gillan, Dana G. Smith, Sanne de Wit, and Karen D. Ersche. 2012. "Neurocognitive Endophenotypes of Impulsivity and Compulsivity: Towards

- Dimensional Psychiatry.” *Trends in Cognitive Sciences* 16 (1): 81–91. <https://doi.org/10.1016/j.tics.2011.11.009>.
- Rolls, Edmund T., Wei Cheng, Matthieu Gilson, Weikang Gong, Gustavo Deco, Chun Yi Zac Lo, Albert C. Yang, Shih Jen Tsai, Mu En Liu, Ching Po Lin, and Jianfeng Feng. 2020. “Beyond the Disconnectivity Hypothesis of Schizophrenia.” *Cerebral Cortex* 30 (3): 1213–33. <https://doi.org/10.1093/cercor/bhz161>.
- Rosario, Maria Conceição, Helena Silva Prado, Sonia Borcato, Juliana Belo Diniz, Roseli Gedanki Shavitt, Ana Gabriela Hounie, Maria Eugênia Mathis, Rosana Savio Mastroso, Patricia Velloso, Eduardo Allende Perin, Victor Fossaluzza, Carlos Alberto Pereira, Daniel Geller, James Leckman, and Euripedes Miguel. 2009. “Validation of the University of São Paulo Sensory Phenomena Scale: Initial Psychometric Properties.” *CNS Spectrums* 14 (6): 315–23. <https://doi.org/10.1017/S1092852900020319>.
- Sakoğlu, Ünal, and Vince D. Calhoun. 2009. “Dynamic Windowing Reveals Task-Modulation of Functional Connectivity in Schizophrenia Patients vs Healthy Controls.” In *Proceedings 17th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Honolulu*.
- Sakoğlu, Ünal, Godfrey D. Pearlson, Kent A Kiehl, Y. Michelle Wang, Andrew M. Michael, and Vince D. Calhoun. 2010. “A Method for Evaluating Dynamic Functional Network Connectivity and Task-Modulation: Application to Schizophrenia.” *Magnetic Resonance Materials in Physics, Biology and Medicine* 23 (5–6): 351–66. <https://doi.org/10.1007/s10334-010-0197-8>.
- Sayed Hussein Jomaa, Mohamad el, Marcelo A. Colominas, Nisrine Jrad, Patrick van Bogaert, and Anne Humeau-Heurtier. 2019. “A New Mutual Information Measure to Estimate Functional Connectivity: Preliminary Study.” *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, 640–43. <https://doi.org/10.1109/EMBC.2019.8856659>.
- Scott Kelso, J. A. 2012. “Multistability and Metastability: Understanding Dynamic Coordination in the Brain.” *Philosophical Transactions of the Royal Society B: Biological Sciences* 367 (1591): 906–18. <https://doi.org/10.1098/rstb.2011.0351>.

- Sen, Bhaskar, Gail A. Bernstein, Bryon A. Mueller, Kathryn R. Cullen, and Keshab K. Parhi. 2020. "Sub-Graph Entropy Based Network Approaches for Classifying Adolescent Obsessive-Compulsive Disorder from Resting-State Functional MRI." *NeuroImage: Clinical* 26 (October 2019): 102208. <https://doi.org/10.1016/j.nicl.2020.102208>.
- Shannon, C. E. 1949. "A Mathematical Theory of Communication." *The Mathematical Theory of Communication*, no. April 1924: 1–54.
- Shappell, Heather M., Brian S. Caffo, James J. Pekar, and Martin A. Lindquist. 2019. "Improved State Change Estimation in Dynamic Functional Connectivity Using Hidden Semi-Markov Models." *NeuroImage* 191 (October 2018): 243–57. <https://doi.org/10.1016/j.neuroimage.2019.02.013>.
- Shmuel, Amir, and David A. Leopold. 2008. "Neuronal Correlates of Spontaneous Fluctuations in fMRI Signals in Monkey Visual Cortex: Implications for Functional Connectivity at Rest." *Human Brain Mapping* 29 (7). <https://doi.org/10.1002/hbm.20580>.
- Sidak, Zbynek. 1967. "Rectangular Confidence Regions for the Means of Multivariate Normal Distributions." *Journal of the American Statistical Association*. <https://doi.org/10.2307/2283989>.
- Singh, Harshinder, Neeraj Misra, Vladimir Hnizdo, Adam Fedorowicz, and Eugene Demchuk. 2003. "Nearest Neighbor Estimates of Entropy." *American Journal of Mathematical and Management Sciences* 23 (3–4): 301–21. <https://doi.org/10.1080/01966324.2003.10737616>.
- Smith, Stephen M., and Thomas E. Nichols. 2009. "Threshold-Free Cluster Enhancement: Addressing Problems of Smoothing, Threshold Dependence and Localisation in Cluster Inference." *NeuroImage* 44 (1): 83–98. <https://doi.org/10.1016/j.neuroimage.2008.03.061>.
- Sporns, Olaf. 2010. *Networks of the Brain*. Networks of the Brain. The MIT Press. <https://doi.org/10.7551/mitpress/8476.001.0001>.
- . 2014. "Contributions and Challenges for Network Models in Cognitive Neuroscience." *Nature Neuroscience* 17 (5). <https://doi.org/10.1038/nn.3690>.
- Subirà, Marta, Pino Alonso, Cinto Segalàs, Eva Real, Clara López-Solà, Jesús Pujol, Ignacio Martínez-Zalacaín, Benjamin J

- Harrison, José M. Menchón, Narcís Cardoner, and Carles Soriano-Mas. 2013. “Brain Structural Alterations in Obsessive-Compulsive Disorder Patients with Autogenous and Reactive Obsessions.” Edited by Noam Harel. *PLoS ONE* 8 (9): e75273. <https://doi.org/10.1371/journal.pone.0075273>.
- Szabó, Zoltán. 2014. “Information Theoretical Estimators Toolbox.” *Journal of Machine Learning Research* 15: 283–87.
- Takayanagi, Yoichiro, Adam P. Spira, Kimberly B. Roth, Joseph J. Gallo, William W. Eaton, and Ramin Mojtabai. 2014. “Accuracy of Reports of Lifetime Mental and Physical Disorders: Results from the Baltimore Epidemiological Catchment Area Study.” *JAMA Psychiatry* 71 (3). <https://doi.org/10.1001/jamapsychiatry.2013.3579>.
- Tzourio-Mazoyer, N., B. Landeau, D. Papathanassiou, F. Crivello, O. Etard, N. Delcroix, B. Mazoyer, and M. Joliot. 2002. “Automated Anatomical Labeling of Activations in SPM Using a Macroscopic Anatomical Parcellation of the MNI MRI Single-Subject Brain.” *NeuroImage* 15 (1): 273–89. <https://doi.org/10.1006/nimg.2001.0978>.
- Uddin, Lucina Q., Richard F. Betzel, Jessica R. Cohen, Jessica S. Damoiseaux, Felipe de Brigard, Simon B. Eickhoff, Alex Fornito, Caterina Gratton, Evan M. Gordon, Angie Laird, Linda J. Larson-Prior, Anthony Randal McIntosh, Lisa D. Nickerson, Luiz Pessoa, Ana Luísa Pinho, Russell A. Poldrack, Adeel Razi, Sepideh Sadaghiani, James M. Shine, et al. 2022. “Controversies and Current Progress on Large-Scale Brain Network Nomenclature from OHBM WHATNET: Workgroup for HARmonized Taxonomy of NETworks.” *Open Science Framework Preprints*. <https://doi.org/10.31219/osf.io/25za6>.
- Vergara, Victor M., Mustafa Salman, Anees Abrol, Flor A. Espinoza, and Vince D. Calhoun. 2020. “Determining the Number of States in Dynamic Functional Connectivity Using Cluster Validity Indexes.” *Journal of Neuroscience Methods* 337 (February): 108651. <https://doi.org/10.1016/j.jneumeth.2020.108651>.
- Vidaurre, Diego, Laurence T. Hunt, Andrew J. Quinn, Benjamin A.E. E. Hunt, Matthew J. Brookes, Anna C. Nobre, and Mark W. Woolrich. 2018. “Spontaneous Cortical Activity Transiently Organises into Frequency Specific Phase-Coupling Networks.” *Nature Communications* 9 (1): 2987. <https://doi.org/10.1038/s41467-018-05316-z>.

- Vohryzek, Jakub, Joana R. B. Cabral, Peter Vuust, Gustavo Deco, and Morten L. Kringelbach. 2022. "Understanding Brain States across Spacetime Informed by Whole-Brain Modelling." *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 380 (2227). <https://doi.org/10.1098/rsta.2021.0247>.
- Vohryzek, Jakub, Gustavo Deco, Bruno Cessac, Morten L. Kringelbach, and Joana R. B. Cabral. 2020. "Ghost Attractors in Spontaneous Brain Activity: Recurrent Excursions Into Functionally-Relevant BOLD Phase-Locking States." *Frontiers in Systems Neuroscience* 14 (April): 1–15. <https://doi.org/10.3389/fnsys.2020.00020>.
- Wang, Zhishun, Tiago v. Maia, Rachel Marsh, Tiziano Colibazzi, Andrew Gerber, and Bradley S. Peterson. 2011. "The Neural Circuits That Generate Tics in Tourette's Syndrome." *American Journal of Psychiatry* 168 (12): 1326–37. <https://doi.org/10.1176/appi.ajp.2011.09111692>.
- Weidle, Bernhard, Thomas Jozefiak, Tord Ivarsson, and Per Hove Thomsen. 2014. "Quality of Life in Children with OCD with and without Comorbidity." *Health and Quality of Life Outcomes* 12 (1): 1–12. <https://doi.org/10.1186/s12955-014-0152-x>.
- Wilson, H. R., and J. D. Cowan. 1973. "A Mathematical Theory of the Functional Dynamics of Cortical and Thalamic Nervous Tissue." *Kybernetik* 13 (2). <https://doi.org/10.1007/BF00288786>.
- Wilson, Hugh R., and Jack D. Cowan. 1972. "Excitatory and Inhibitory Interactions in Localized Populations of Model Neurons." *Biophysical Journal* 12 (1). [https://doi.org/10.1016/S0006-3495\(72\)86068-5](https://doi.org/10.1016/S0006-3495(72)86068-5).
- Womelsdorf, Thilo, Jan Mathijs Schoffelen, Robert Oostenveld, Wolf Singer, Robert Desimone, Andreas K. Engel, and Pascal Fries. 2007. "Modulation of Neuronal Interactions through Neuronal Synchronization." *Science* 316 (5831). <https://doi.org/10.1126/science.1139597>.
- Xia, Cedric Huchuan, Zongming Ma, Rastko Ciric, Shi Gu, Richard F. Betzel, Antonia N. Kaczurkin, Monica E. Calkins, Philip A. Cook, Angel García de la Garza, Simon N. Vandekar, Zaixu Cui, Tyler M. Moore, David R. Roalf, Kosha Ruparel, Daniel H. Wolf, Christos Davatzikos, Raquel E. Ruben C. Gur, Raquel E. Ruben C. Gur, Russell T. Shinohara, et al. 2018.

- “Linked Dimensions of Psychopathology and Connectivity in Functional Brain Networks.” *Nature Communications* 9 (1): 3003. <https://doi.org/10.1038/s41467-018-05317-y>.
- Zalesky, Andrew, Alex Fornito, and Edward T. Bullmore. 2010. “Network-Based Statistic: Identifying Differences in Brain Networks.” *NeuroImage* 53 (4): 1197–1207. <https://doi.org/10.1016/j.neuroimage.2010.06.041>.
- Zalesky, Andrew, Alex Fornito, Luca Cocchi, Leonardo L. Gollo, and Michael Breakspear. 2014. “Time-Resolved Resting-State Brain Networks.” *Proceedings of the National Academy of Sciences* 111 (28): 10341–46. <https://doi.org/10.1073/pnas.1400181111>.
- Zhang, Wei, Viktoria Muravina, Robert Azencott, Zili D. Chu, and Michael J. Paldino. 2018. “Mutual Information Better Quantifies Brain Network Architecture in Children with Epilepsy.” *Computational and Mathematical Methods in Medicine* 2018. <https://doi.org/10.1155/2018/6142898>.