

# Spatio-temporal dynamics of human fMRI resting state

Katharina Glomb

---

TESI DOCTORAL UPF / 2016

Director de la tesi

Prof. Dr. Gustavo Deco,  
Department of Information and Communication Technologies





---

# Acknowledgements

---

First of all I would like to acknowledge the INDIREA Marie Curie Initial Training Network as the funding source of my PhD work (European Union FP7 Marie Curie ITN Grant N. 606901). Everyone in this network has made this PhD an amazing experience, especially the other “early stage researchers” with numbers ranging from 1 to 13 (me being number 10), and many of the PIs that were involved. And of course Eli Fulcini, I cannot even list all the things she has done for all of us and without her, all of our meetings would have been a mess instead of the fun and useful training events that they were.

Of course, I thank especially the one and only, Gustavo! Thank you for encouraging me, letting me work way more independently and giving me more trust than anyone could wish for - you (and Marie Curie) have spoiled me for the future.

Huge thanks to my collaborators and co-supervisors: Adrián Ponce-Alvarez, Matthieu Gilson, Petra Ritter - whose lab provided the beautiful fMRI data analyzed here - as well as Romain, Raphael, Selen, Adrià, Gorka, Rikkert, Mohit, and Mario, all of whom at some point have given me useful advice and support for the work presented on the following pages. Also, Jordi and David Varela, without whom I would have surely despaired in the face of cluster trouble; and the group as a whole has made this PhD an outstandingly pleasant experience that I will never forget. I would also like to mention all the admin staff who make sure that this place keeps running: Florencia, Silvia, Xavi and Cristina from the CBC and Lydia, Joana, Lourdes, Montse, Vanessa, and the whole team from the DTIC secretariat.

I had the privilege to take advantage of various research stays, trips, conferences, meetings and so on, one of which brought me to Japan and the research lab of Professor Andrzej Cichocki without whose help, and the help of his co-workers Dr Zhou Guoxu, Dr Qibin Zhao, and Dr Anh Huy Phan, I would have never gotten to the bottom of this tensor story. Another one brought me to Magdeburg and the lab of Professor Jochen Braun, and I thank João V Dornas and Professor Braun himself for all their help and support in those weeks; the latter deserves some more thanks for mentoring me and giving me insightful comments and feedback on my manuscripts at many INDIREA meetings and during my stay at his lab. I am looking forward to our future collaborations!

Last but definitely not least, I give the biggest thank you to my family back home who have supported me and kept me company during all the ups and downs - I cannot name you all, but my sister Anne, my parents Dorit and Michael, my grandparents Anneliese, Christa and Siegfried have to be here as well as Andi, Ivana, Lennart, Susi, Simon, and Evan. It's difficult to make a list of "most important friends" because there are so many meetings and events I remember that contributed to my life throughout this thesis work, but I also would feel bad not putting some of the names here. Laura, Marvin, Josephine, Katharina, Cornell, Jonas, the other Jonas, Marcus, Magdalena, Owen, Uli, Manuel, Takuya - I cannot not name you. My friends here were direct witnesses of what was going on, and many of them shared the same experience with me: Loreto, Pallabi, Kalinka, Giovanni, Camillia, Mathilde, Daria, Joanna, Stefanie, Nils, Lauren - thank you for being there and for being yourselves. Most of all I thank Cristina, without whom my life would be decidedly dull.

---

---

---

---

---

# Abstract

Spontaneous brain activity, measured under the absence of any overt task, has been investigated under the label of “resting state” for about 20 years with rising interest. While it was known since the beginnings of modern electrophysiology that the brain exhibits spontaneous fluctuations also during rest, the discovery, in 1995, that these fluctuations possess a robust spatio-temporal structure had a profound impact on how we understand and investigate brain activity. In this dissertation, we characterize the spatio-temporal dynamics of resting state on a macroscopic level using fMRI recordings from humans and combining novel data analysis tools with theoretical models on the level of the whole brain. We demonstrate the presence of common patterns of functional connectivity, known as resting state networks (RSNs), that evolve in time in both empirical and model data. We show that spontaneous fluctuations and their statistics are determined by the structure of the brain network and its dynamics.

---

## Resumen

La actividad cerebral espontánea, o actividad de reposo, es aquella que uno puede registrar cuando el cerebro no está involucrado en ninguna tarea impuesta del exterior (tal como sería la presentación de un estímulo). El estudio de la actividad de reposo ha conocido un interés creciente durante los últimos 20 años. Si bien las fluctuaciones en la actividad de reposo eran conocidas desde los inicios de la electrofisiología moderna, el descubrimiento, en 1995, de que estas fluctuaciones muestran patrones espaciotemporales robustos ha tenido un impacto profundo en la manera de entender e investigar la actividad del cerebro. En esta disertación caracterizamos la dinámica espaciotemporal de la actividad de reposo a nivel macroscópico usando registros de fMRI en humanos y combinando nuevas herramientas de análisis y modelos teóricos del cerebro a gran escala. Observamos patrones comunes de conectividad funcional evolucionando en el tiempo tanto en los datos empíricos como en las simulaciones. Demostramos que las fluctuaciones de reposo y su estadística son determinadas por la estructura de la red cerebral y su dinámica.

---

# Preface

The brain is the basis of behavior, personality, and conscious experience. One of the main challenges it faces is to balance flexibility and stability in order to enable us to survive in unknown situations, develop new solutions, and learn, in other words: to adapt, and at the same time draw from experience and remember. This is illustrated in an important concept of brain function: segregation versus integration. Segregation is necessary in order to be able to optimally perform a certain function in a specialized brain region, something that we have studied extensively over the last decades by identifying countless brain areas that support certain functions tested in experiments, like decision making, perception, problem solving, emotional processing, and many more. This endeavor has been immensely aided by imaging methods such as functional magnetic resonance imaging (fMRI). Nevertheless, we have always been interested in how brain regions communicate, recognizing that flexible integration between functions is necessary to perform complex cognitive functions.

Computational neuroscience is an approach that casts the idea of a “function” into the concept of computations and asks what kind of computation a brain area is suited for and why. Interplay between different kinds of computations over a staggeringly large range of temporal and spatial scales in a hierarchical fashion is the basis for performing tasks like navigating a car on a crowded highway, learning how to speak, or writing a PhD thesis.

We have long since learned that what fires together, wires together. A main topic in this thesis is the connectivity between brain areas that allows this teamwork to occur, and which structures a bunch of brain regions into a hierarchy inside of which flexible integration and segregation of functions

takes place. On a microscopic level, this entails synaptic plasticity, on a mesoscopic scale, neural assemblies, but the scale that is important here is the macroscopic one, the one that is measured with fMRI.

About 20 years ago, fMRI was instrumental to a discovery that radically changed the way we think about macroscopic connectivity: Even in the fMRI scans that were taken while subjects were not performing any tasks - until then referred to as “noise scans” - the correlation between fluctuations in the involved brain regions (in this case, motor regions) was very close to that observed during task. On the one hand, this meant that activations that are observed during task could no longer be considered as merely differing from a flat baseline consisting of uncorrelated noise. On the other, it allowed scientists to map distributed brain networks that are jointly involved in similar tasks - i.e. they are functionally connected - without having to design specific experiments. The emerging term “resting state” stands as much for an “idling” brain state as for an experimental paradigm. Immediately the question arises of what supports these functional relationships, and computational models have been invaluable in linking temporal and spatial scales and thus findings from different fields of neuroscience ranging from single cell recordings to fMRI. But it becomes more complicated: About a decade ago, scientists started realizing that this spatially very complex baseline also had a temporal structure, meaning that functional relationships between brain regions change all the time. Chapter 1 of this thesis will introduce “Resting state and its dynamics”.

Chapter 2, “Robust extraction of spatio-temporal patterns from resting state fMRI”, contains the first study that was conducted within this PhD. Its goal is to show that a certain dynamic model of resting state, the dynamic mean field model, does not only reproduce the average functional connectivity, as was already shown elsewhere, but also exhibits specific functional patterns over time. It shows how a decomposition technique that is routinely in use in other fields, but so far not in neuroscience, namely tensor factorization, can be used to obtain common sets of regions that tend to be functionally connected from empirical as well as simulated data; these “sets of regions”, or communities, are well-known in the literature as resting state networks. A discussion of which consequences our findings have for our understanding of the dynamics present in resting state is also included.

Chapter 3, “Temporal dynamics of human resting state fMRI”, continues where chapter 2 leaves off and characterizes the temporal changes in functional connectivity. We establish suitable measures to track these changes



and show that they are not explained by random fluctuations in global connectivity, but are a specific signature of resting state dynamics. This is tied to the temporal evolution of the resting state networks mentioned above. Furthermore, we discuss which properties a model that reproduces these dynamics should have.

Chapter 4 contains the “General discussion” and ties together all the findings presented in this thesis, including a discussion of future steps.

The work presented in chapters 2 and 3 could not have been done without collaboration, and two manuscripts resulting from it, carrying the same titles as the chapters of this dissertation, list the authors as follows: Katharina Glomb, Adrián Ponce-Alvarez, Matthieu Gilson, Petra Ritter, and Gustavo Deco.



---

# Contents

<b>Abstract</b>	<b>v</b>
<b>Resumen</b>	<b>vi</b>
<b>Preface</b>	<b>vii</b>
<b>List of Figures</b>	<b>xiii</b>
<b>List of Tables</b>	<b>xix</b>
<b>1 Introduction: Resting state and its dynamics</b>	<b>1</b>
1.1 Overview .....	1
1.2 Temporal dynamics of resting state .....	9
1.3 Tensor factorization for neural data.....	14
1.4 Studies in this thesis .....	21
<b>2 Robust extraction of spatio-temporal patterns from resting state fMRI</b>	<b>23</b>
2.1 Introduction.....	24
2.2 Methods .....	26
2.3 Results.....	38
2.4 Discussion .....	48
<b>3 Temporal dynamics of human resting state fMRI</b>	<b>55</b>
3.1 Introduction.....	56
3.2 Methods .....	58

3.3 Results.....	61
3.4 Discussion .....	75
<b>4 General Discussion</b>	<b>81</b>
4.1 Contributions .....	84
4.2 Conclusion .....	87
<b>A Appendix</b>	<b>89</b>
<b>Bibliography</b>	<b>91</b>

—

|

|

---

# List of Figures

1.1	Areas of the brain that decrease their activity in attention-demanding tasks and external stimulation, as measured with PET. <i>Image from Raichle et al. (2001)</i> .....	2
1.2	Illustration of anticorrelated networks in the human brain. The seed is located in the posterior cingulate cortex (PCC; yellow curve), i.e. in the task-negative network. Positive correlations are obtained with other voxels within this network (medial prefrontal cortex, MPF; orange curve), and negative ones with the “task-positive network” (intraparietal sulcus, IPS; blue curve). <i>Image from Fox et al. (2005)</i> .....	4
1.3	RSNs extracted with (spatial) ICA, <i>image from Beckmann et al. (2005)</i> . RSNs are labelled as follows: (a) medial visual cortical areas, (b) lateral visual cortical areas, (c) auditory system, (d) sensorymotor system, (e) visuo-spatial system, (f) executive control, (g, h) dorsal visual stream. ....	5
1.4	Example for how the brain is modelled as a graph. The color of each node stands for its degree, i.e. how many other nodes it is connected to, yellow = 1, green = 2, red = 3, blue = 4, other colours > 4, and the blue edges stand for functional connections in one hemisphere. <i>Image from Sporns et al. (2004)</i> .....	6

1.5 Example for the result of a time-frequency analysis via WTC for three subjects, using the PCC seed and another DMN node (medial prefrontal cortex). For each point in time (x-axes), the phase is plotted for each Fourier period (y-axes). Dark blue regions indicate no significant coherence, bright blue stands for a phase difference of  $0 \pm \pi/4$ , orange:  $\pi/2 \pm \pi/4$ , green:  $-\pi/2 \pm \pi/4$ , red:  $\pi \pm \pi/4$ . *Image from Chang and Glover (2010)* ..... 11

1.6 Time courses of windowed correlations, averaged over voxels belonging to “oculomotor” network (OCM, red), ventral pre-motor network (vPM, blue), and white matter (WM, cyan), from one monkey in the study of *Hutchison et al. (2012)*..... 12

1.7 Illustration of Tucker-N-decomposition when  $N = 3$ . *Image from Cichocki (2013)*..... 18

1.8 Illustration of CPD. The core tensor  $\mathbf{G}$  has ones on its superdiagonal and is zero everywhere else. One set of features is extracted for each mode of the tensor ( $\mathbf{A}$ ,  $\mathbf{B}$ ,  $\mathbf{C}$ ). Bottom: The tensor can then be approximated by rank-1 tensors made up of unique combinations of rows/columns from the feature matrices. *Image from Cichocki et al. (2009)* ..... 20

2.1	<p>Overview of analysis pipeline that leads to tensorized data. For simulated data, panels A to E apply, for empirical data, panels C to E. <b>A</b> Schematic view of dynamic mean field (DMF) model used to simulate synaptic activity. Each brain area is modelled by a pair of excitatory (E) and inhibitory (I) pools. The local connectivity is governed by the four weights <math>w_{EE}</math>, <math>w_{EI,i}</math>, <math>w_{IE}</math>, and <math>w_{II}</math>, whereby <math>w_{EI,i}</math> is adjusted for each population individually so as to keep the firing rates at a low level (3-10 Hz). Black lines with spheres signify GABA connections, black arrows, NMDA. Gray arrows are long range connections mediated by AMPA synapses whose weights are set by the SC or EC matrix. <b>B</b> Synaptic activity is simulated on a millisecond scale. <b>C</b> BOLD activity on the scale of seconds is obtained from simulated activity by applying the Balloon-Windkessel hemodynamic model. <b>D</b> Starting from these BOLD time courses, we compute time dependent dFC-matrices from overlapping time windows <math>w</math> of width 120s (60 frames) starting at time points <math>t = 1, 2, \dots, W</math>. Each dFC-matrix has dimensions <math>N \times N</math>, <math>N</math> being the number of ROIs. <b>E</b> This results in an <math>N \times N \times W</math> tensor. The entry at <math>dFC_{ij}(w)</math> is the functional connectivity value of ROIs <math>i</math> and <math>j</math> in window <math>w</math>. dFC-matrices are symmetric. ....</p>	29
2.2	<p>Illustration of tensor factorization with <math>F = 3</math> features. The original data tensor (top left) is decomposed into three components, i.e. three communities with their associated time courses are extracted. As shown in the bottom row, each set of vectors, consisting of two identical “community vectors” <math>a_f</math> and <math>b_f</math> and a time course, <math>c_f</math>, constitutes a rank-1 tensor, which can be explicitly computed by taking the outer product as in equation 1.8 (taking the outer product of <math>a_f</math> and <math>b_f</math>, one would obtain an FC matrix like the “layers” of the tensor). By adding these rank-1 tensors, we obtain an approximation of the data tensor, i.e. a reconstructed tensor (top right), and can evaluate the fit between the two by taking the distance. . .</p>	32
2.3	<p>Result of template extraction with optimal parameter settings. <b>A</b> Silhouette values for clustering all features extracted from MI tensors for the best threshold, i.e. 98th percentile. <b>B</b> Correlation between the same features, ordered by assignment to <math>K = 4</math> clusters. The clusters are in the same order as the templates in figure 2.4. ....</p>	40

2.4	Templates plotted in 3D. From top to bottom: default mode network, somatomotor network, control networks, visual network. Although the templates are obtained at a resolution of 66 cortical regions, for better visualization we use 998 centers of mass, each of which is clearly assigned to one of the 66 areas. The opacity of each sphere is proportional to the weight assigned to it in the spatial map of the template. ....	41
2.5	Overall template matches in comparison to correlation fits between FC matrices, results for simulations using SC (top) and EC (bottom). Black curves: template matches (left axes), thick black curve is real data, dashed line surrogate data. Shaded areas indicate the 95% confidence interval. Red curves: correlation of average FC matrices (right axes). ....	44
2.6	Spatial features extracted from 24 simulation runs using SC as underlying connectivity and the $G$ with the best match (figure 2.5: $G = 3.7$ ) with $\theta = 98$ th percentile and $F = 3$ , plotted next to the templates with which they exhibit the best match. Templates correspond to DMN - default mode network, SM - sensorimotor network, CTR - left and right control networks, VIS - visual network. For each set of features, the number of matched features and the average overlap for only this set is given.....	45
2.7	Spatial features extracted from 24 simulation runs using EC as underlying connectivity and the $G$ with the best match (figure 2.5: $G = 4.2$ ) with $\theta = 98$ th percentile and $F = 3$ , plotted next to the templates with which they exhibit the highest match. Templates correspond to DMN - default mode network, SM - sensorimotor network, CTR - left and right control networks, VIS - visual network. For each set of features, the number of matched features and the average overlap for only this set is given.....	46
2.8	Connectivity matrices. <b>A</b> Functional connectivity matrix averaged over 24 subjects. Each entry is the correlation over the entire time course of a pair of ROIs. <b>B</b> Average structural connectivity matrix from the same 24 subjects. <b>C</b> Effective connectivity matrix derived from the FC, using the method in section 2.2.9. <b>D</b> Degrees of the nodes. For SC (blue), in and out degree are identical. Here, nodes are ordered according to their degree. EC nodes (in-degree: orange, out-degree: yellow) are shown in the same order as for SC. ....	47



2.9	Time courses of the number of supra-threshold FC pairs per window for three subjects. The tensors were created using MI and $\theta = 98$ th percentile. ....	49
3.1	Illustration of three measures of RS dynamics. Top: We use sliding windows of width 120s (60 frames) and slide by 1 frame. Bottom left: Time courses in one window. The variance of each ROI is computed. Averaging over all ROIs yields the variance of the BOLD signal (“BOLD var”) for this window. Bottom right: The average FC (avFC) is obtained by correlating the full time courses of all pairs of ROIs (top). The same method is applied for each window, yielding the dynamic FC matrix, $dFC(w)$ (bottom). Averaging over all “unique” pairs yields the overall FC strength for this window, and correlating the unique pairs of the dFC with those of the avFC yields the similarity to avFC (“sim to avFC”). ....	63
3.2	Time courses of three measures BOLD variance, FC strength and similarity to average FC for the same subject as in figure 3.1. The measures are normalized to lie between 0 and 1 for illustration purposes. ....	64
3.3	Original data of one example subject together with a phase randomized version. The obvious similarity is due to the preservation of spectral properties and long-term correlations.	65
3.4	Illustration of the effect of phase randomization on the three measures of global RS dynamics. The time courses computed from an example subject are shown in blue, together with one phase shuffled version derived from the same subject. Top: BOLD variance, middle: FC strength, bottom: similarity to average FC. While specifics are clearly different, the ranges are the same. ....	66
3.5	Range of fluctuations for real data (black bars, marking minimum and maximum of the time course) and surrogate data (colored dots). The empty circles mark the medians of the real data. Each dot lying above the black bar stands for a surrogate time course reaching a higher value than the original time course from which it was derived; each orange dot lying below its bar stands for a time course reaching a lower value. Top: BOLD variance, middle: FC strength, bottom: similarity to average FC. ....	67

3.6 Correlations between all combinations of measures for all subjects..... 68

3.7 Distributions of all three measures averaged across all subjects, using 20 bins. The bins do not represent the same values for all subjects, since their ranges are different (see figure 3.4). Surrogate histograms use the same bins as their respective subject. .... 69

3.8 Time courses associated to communities identified as the indicated RSNs from two subjects. Thresholding the tensors before decomposing them leads to the activations being 0 in some windows because the baseline is effectively removed. DMN: Default mode network, SMN: somatomotor network, CTR: right and left control networks, VIS: visual network .... 71

3.9 Co-activation analysis of RSNs extracted via tensor factorization. Left: example time courses of “contributions” for the activation time courses shown in figure 3.8 (top). The sum over all values at each time point is 1. Right: Histogram of contribution values for each RSN, pooled across all subjects. DMN: Default mode network, SMN: somatomotor network, CTR: right and left control networks, VIS: visual network .... 73

3.10 Joint distribution of BOLD variance (x-axes) and similarity to avFC (y-axes) for empirical data (left) and surrogates (right). The marginal distributions are shown at the top and sides..... 74

---

## List of Tables

2.1	Parameters of the DMF model. ....	26
A.1	ROIs in the same order in which they appear in the vectors in figures 2.6 and 2.7, as well as in the matrices in figure 2.8. .	90

---



---

# Introduction: Resting state and its dynamics

## 1.1 Overview

---

Biswal et al. (1995) showed that low-frequency fMRI signals are correlated among anatomically separate regions even under the absence of any task. Specifically, they showed that the correlation pattern between voxels that were found to be activated during a finger tapping task was very much the same while this task was not being performed. This was true even for voxels located in the opposite hemisphere to the region used as seed. The authors make the point that the fluctuations observed in the band below 0.1 Hz are of neural origin, which has since then been corroborated by other authors (Shmuel and Leopold, 2008; Nir et al., 2008; Schölvinck et al., 2010; Britz et al., 2010; de Pasquale et al., 2010). Likewise, they state that their observations are not due to the subjects imagining themselves doing the same task because such imagery studies had yielded slightly different results. Thus, they showed that spontaneous brain activity is spatially and temporally structured.

Today, this publication is nearing 5000 citations, and many of the papers that cite it have thousands of citations themselves. Let us consider two reasons for the popularity of the authors' discovery.

On the one hand, it raised the question as to what exactly the brain is doing during rest. The idea of a “default mode of brain function” (Raichle et al., 2001; Greicius et al., 2003) states that during rest, the brain is busy with




Figure 1.1: Areas of the brain that decrease their activity in attention-demanding tasks and external stimulation, as measured with PET. *Image from Raichle et al. (2001)*

certain self-referential tasks which are put on hold as soon as an external task arrives (Fransson, 2006) (figure 1.1). While this has been connected to cognitive correlates like mind wandering (Mason et al., 2007), similar patterns have been observed under anaesthesia in macaques (Vincent et al., 2007), suggesting that spontaneous fluctuations are (also) a network effect. Thus, spontaneous activity is not only relevant under the absence of task (“resting state”, RS) but also for any study that has as its goal the localization of a function to a brain area, because we can no longer assume a “flat” baseline, i.e. uncorrelated fluctuations that can be averaged out.

On the other hand, it was realized that these correlations can be used to map the functional connectivity (FC) structure of the brain without having to apply a large number (or a small number, for that matter) of tasks, and thus, large-scale connectivity can be mapped for many functions (Lowe et al., 1998; Cordes et al., 2000; De Luca et al., 2005; Vincent et al., 2006) in any person that is able to lie still for five minutes (Van Dijk et al., 2010).

FC can be defined at the level appropriate for fMRI as any kind of statistical dependence between the measured activity of voxels or brain regions, for example Pearson correlation, covariance, or mutual information (Friston et al., 1993). This is taken to be indicative of coherence in the activities of the brain regions. Thereby we assume that brain regions that show coherent activity patterns exchange information and thus, integrate it in order to subserve a specific function. The coherent sets of regions that were revealed using correlations between time courses of spontaneous activity during RS were named “resting state networks” (RSNs). This term, however,

is slightly misleading, because these correlation patterns are not abolished during task, and specific tasks evoke co-activation patterns that are similar to RSNs (De Luca et al., 2005; Cole et al., 2014; Glasser et al., 2015). Therefore, some authors use the term “intrinsic connectivity networks” (ICNs) instead; see for example Calhoun and Adali (2012). On a systems level, Fox et al. (2005) observed sets of regions whose time courses were correlated with each other, but anticorrelated with other sets (figure 1.2). This lead to the notion that “the human brain is intrinsically organized into [...] anticorrelated functional networks” (Fox et al., 2005). This organisation is also apparent during rest, where the spontaneous activity of task-negative regions is anticorrelated with that of a number of other regions; thus the definition of the “default mode network” (DMN) - perhaps the best-known RSN - and the “task-positive network” (TPN) (Fox et al., 2005).

The DMN seems to play a particularly important role in mental disorders (Broyd et al., 2009), and it has been linked to consciousness (Gusnard et al., 2001) and recently, to many real-life behavioral measures (Smith et al., 2015) like IQ, life satisfaction, or cannabis use. Its level of activation is predictive of participants’ performance or behavior (Boly et al., 2007; Hesselmann et al., 2008; Sadaghiani et al., 2009; Coste et al., 2011). Specifically, higher levels of activity in the DMN are typically reported to yield weaker performance, while lower levels together with stronger activation of attention-related regions and networks linked to the task modality predict better performance. These observations have lead to the integration of the “default mode”-hypothesis into Bayesian ideas of predictive coding (Ringach, 2009; Sadaghiani et al., 2010; Carhart-Harris and Friston, 2010). In this framework, fluctuations represent the brain’s attempt to prepare for potential incoming stimuli and facilitate a swift reaction (Deco et al., 2013), as well as being indicative of the hierarchical organization of the brain, where higher areas (DMN, salience network) minimize prediction errors in lower ones (sensory cortices).

Put bluntly, information is believed to be integrated inside of RSNs/ICNs, and segregated between them. Instead of using seeds to unveil functional networks, this idea of integration/segregation motivates another, more data-driven formulation of the problem as blind source separation: the signals of a limited number of spatially independent sources (the RSNs/ICNs) mix to produce the data that we pick up in the MRI voxels. The goal is to find the unmixing matrix that allows us to extract the original components. Due to the segregation idea, it is reasonable to apply the constraint of spatial independence, which leads to spatial independent component analysis (sICA,

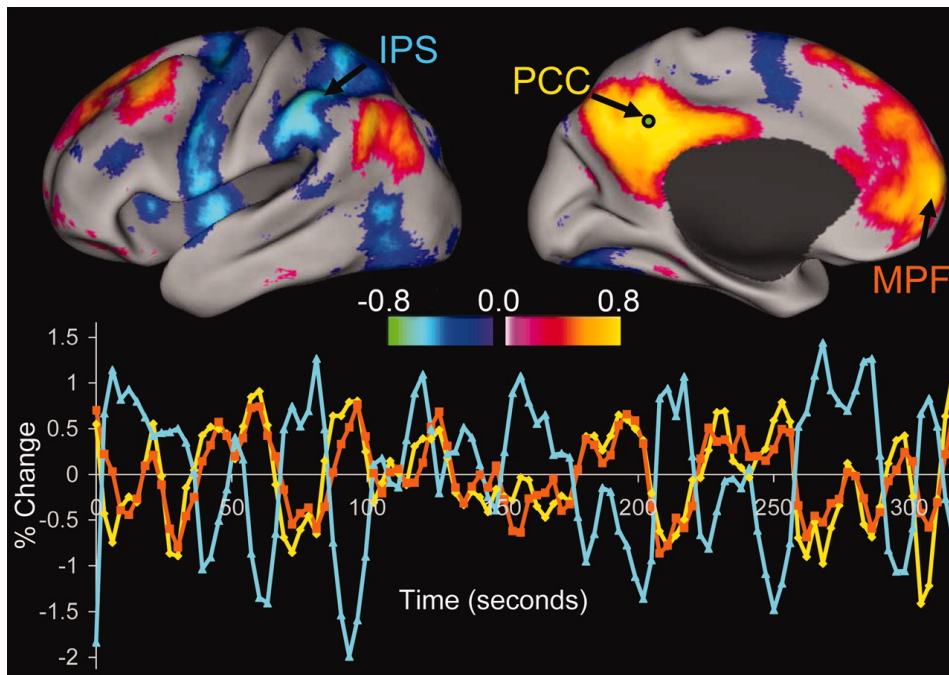


Figure 1.2: Illustration of anticorrelated networks in the human brain. The seed is located in the posterior cingulate cortex (PCC; yellow curve), i.e. in the task-negative network. Positive correlations are obtained with other voxels within this network (medial prefrontal cortex, MPF; orange curve), and negative ones with the “task-positive network” (intraparietal sulcus, IPS; blue curve). *Image from Fox et al. (2005)*

or just ICA) (Bell and Sejnowski, 1995; McKeown et al., 1998; Hyvärinen et al., 2004). Using this technique, all RSNs can be extracted simultaneously (Kiviniemi et al., 2003; Beckmann et al., 2005) (figure 1.3). It was shown that they are reproducible across subjects (Beckmann and Smith, 2005; Damoiseaux et al., 2006; Yeo et al., 2011) and related to functional networks found in activation studies (Smith et al., 2009). It is also possible to apply the independence constraint on the temporal domain (tICA), but this is rarely done (Brookes et al., 2011; Smith et al., 2012) because of the limited number of time points in a typical scanning session. In section 1.3 below, we present a more thorough introduction to this topic since it is of particular relevance to this thesis.

Given the connection between fundamental concepts like integration and segregation and RSNs, it is not surprising that the methods developed for






Figure 1.3: RSNs extracted with (spatial) ICA, *image from Beckmann et al. (2005)*. RSNs are labelled as follows: (a) medial visual cortical areas, (b) lateral visual cortical areas, (c) auditory system, (d) sensory-motor system, (e) visuo-spatial system, (f) executive control, (g, h) dorsal visual stream.

investigating FC in RS have been used to search for, and lead to the discovery of, changes of FC linked to ageing (Damoiseaux et al., 2008; Dosenbach et al., 2010; Chan et al., 2014; Betzel et al., 2014), different disorders (Greicius et al., 2004; Cherkassky et al., 2006; Yu-Feng et al., 2007; Castellanos et al., 2008; Greicius, 2008; Rotarska-Jagiela et al., 2010; Veer et al., 2010; Kühn and Gallinat, 2013), and brain states (Horovitz et al., 2008; Boveroux et al., 2010; Brewer et al., 2011; Carhart-Harris et al., 2012). Importantly, the described changes are often not merely markers, but they establish that RSNs are functionally significant and rooted in the underlying anatomical connections. For example, it was shown (Sorg et al., 2007; Seeley et al., 2009) that different neurodegenerative disorders spread through a particular RSN after initiating in a spatially delimited seed region. Recently, it was




Figure 1.4: Example for how the brain is modelled as a graph. The color of each node stands for its degree, i.e. how many other nodes it is connected to, yellow = 1, green = 2, red = 3, blue = 4, other colours > 4, and the blue edges stand for functional connections in one hemisphere. *Image from Sporns et al. (2004)*

---

found that gene expression patterns are correlated with RSNs (Richiardi et al., 2015), a finding that illustrates how different fields like neuroscience and molecular biology can link and produce new insights.

In the present context, the FC we are interested in is on a macroscopic scale, incorporating the entire brain; we are interested in how distributed, large-scale networks integrate information in order to subservise complex cognitive functions. Therefore, it makes sense to conceptualize the brain as a graph, i.e. a network consisting of (essentially equal) nodes and edges (Sporns et al., 2004; Rubinov and Sporns, 2010) (figure 1.4). On the one hand, this allows us to assess the brain in terms of its network properties (Watts and Strogatz, 1998), providing interesting insights into how the brain network is optimized for processing information, cost and efficiency, and resisting damage (Sporns et al., 2004; Kaiser and Hilgetag, 2004; Bullmore and Sporns, 2012). The focus is on the modular structure and the hierarchy of regions and groups of regions (Meunier et al., 2010).

Furthermore, using graphs provides an entry point for modelling studies which strive to uncover the relationship between structure and function (Sporns et al., 2004, 2005; Honey et al., 2009; Deco et al., 2010). This

topic received a boost not only from whole-brain FC obtained from RS scans, but also from the development of non-invasive imaging techniques for structural connectivity, diffusion-weighted MRI (dwMRI) or diffusion tensor imaging (DTI) (Basser et al., 1994). Combining these techniques, it is possible to map both the functional and structural connectome (Sporns et al., 2005; Van Dijk et al., 2010) of the human brain, one of the main efforts of neuroscience today (Van Essen et al., 2013).

In terms of the graph formalism, the edges can stand for different kinds of relationships between the nodes (Sporns et al., 2004). In the case of FC, they stand for the strength of the statistical dependence found between the time courses of brain regions; in the case of structural connectivity, they stand for fiber densities, anisotropy, or some other measure derived from DTI describing the physical, anatomical connections that exist in the brain; a third possibility is to use effective connectivity (Friston, 1994), which, via a model, uses causal influences that brain regions exert upon each other. Also the nodes can be defined in different ways, depending on the parcellation. We have to consider not only the resolution (i.e., using more or fewer nodes), but also what criterion we use to define a brain region. This can happen purely anatomically, based on landmarks or cell types (Tzourio-Mazoyer et al., 2002; Fischl et al., 2004), or functionally, using findings from activation studies (Glasser et al., 2015), or even uniformity of RS activity (Shen et al., 2013).

Even though this way of modelling the brain has been extremely useful, it is important to understand that it has its limitations. Actually, the data do not warrant a clear-cut modular structure because interactions over time are influenced by the hierarchy of brain regions between and across RSNs/ICNs (Sadaghiani et al., 2010). When taking into account actual anatomy (Kopell et al., 2014), we realize that the brain is a continuous sheet of tissue which is not easily divided into “nodes” and “edges”; brain signals like local field potentials or travelling waves do not stop at the border of a “node”, nor are they transmitted solely through “edges”. Furthermore, the hierarchical organization of the brain is heavily influenced by the laminar structure of feedforward and feedback projections (Markov et al., 2014; Wang and Kennedy, 2016), something that is rather hard to include in a graph representation.

While the initial focus of RS research was on delineating RSNs/ICNs, i.e. entities that are thought to be more or less separated in both time and space, the fluctuations themselves have received more and more interest over the

last years. Indeed, it has been appreciated already two decades ago that the brain state directly preceding a stimulus linearly superimposes with the responses of neurons to a repeated visual stimulus (Arieli et al., 1996), explaining a large part of the inter-trial variability in neuronal responses. This was shown in fMRI by Fox et al. (2006). Thus, neuronal responses to incoming stimuli are state-dependent (the stimulus-evoked activity interacts with the ongoing activity). Horowitz et al. (2008) corroborated this on a global level by comparing the magnitude of spontaneous fluctuations between RS, task, and sleep. Likewise, Garrett et al. (2010) demonstrated changes over the life span.

A number of recent studies have focused on the temporal integration of RSNs/ICNs in rest and task (de Pasquale et al., 2012; Cole et al., 2013; Mitra et al., 2014). The rich dynamics present in the fMRI time series are currently being explored, and a number of methods is being employed in this endeavor (Chang and Glover, 2010; Allen et al., 2012; Jones et al., 2012; Liu and Duyn, 2013; Betzel et al., 2016). This adds another layer to the question of how the brain integrates and segregates information because we need to take into account that areas can be integrated at one point and segregated at another (Hutchison et al., 2013; Calhoun et al., 2014; Kopell et al., 2014). Likewise, there are differences depending on age, brain state, or due to disorders, with studies on these topics already emerging (Jafri et al., 2008; Garrett et al., 2010; Jones et al., 2012; Barttfeld et al., 2015). We go more into detail below since this topic is crucial to the present thesis (section 1.2).

In general, the question arises as to how a relatively unchanging anatomical connectivity can at the same time support extremely versatile and adaptive cognitive function. Computational models are invaluable because they provide simulated time courses which can be tested against the data (Haimovici et al., 2013; Hansen et al., 2014; Messé et al., 2014). Importantly, this can help answer a questions that is overarching to the mere characterization, description and reproduction of time courses and their properties: What kind of dynamical system is the brain? In this framework, co-activation patterns that occur repeatedly and reliably are formalized as *attractors*, and the patterns themselves as well as the dynamics which govern their activation and the transition between them are abstracted as a “dynamical repertoire” of the brain. This “repertoire” subserves the flexibility and adaptiveness that is necessary for cognitive function.

How is the observed “dynamical repertoire” generated? How does switching

occur between different states? How is dynamical behavior organized across spatial and temporal scales, spanning from the firing of single neurons to complex behaviors? In multistability (Schoner and Kelso, 1988; Freyer et al., 2011), several stable states co-exist and the system switches between them due to noise. In metastability (Kelso, 1995; Friston, 1997; Tognoli and Kelso, 2014), the switching behavior itself follows a stable pattern and does not require any additional drive, be it by external inputs or noise. This can be explained in terms of “ghost attractors” (Deco and Jirsa, 2012; Kelso, 2012), thereby establishing a link between multi- and metastability. A related approach is that of self-organized criticality (Chialvo, 2010; Tagliazucchi et al., 2012; Haimovici et al., 2013), where changing attractor landscapes are explained by small events “avalanching” to larger scales. Common themes are the gradual transition between integration and segregation via coupling of the elements of the system, and generally, phase transitions and bifurcations, close to which attractors emerge from the system’s constituent elements - i.e., single neurons and small circuits. Importantly, this theory of dynamical systems strives to explain brain dynamics as emergent from a relatively low-dimensional system, a point of view that is particularly interesting in the age of the human brain project ([www.humanbrainproject.eu](http://www.humanbrainproject.eu)) which aspires to ever-growing detail in theoretical models of the brain.

If the discovery of RS FC structure triggered a paradigm shift from assuming a “flat” baseline to taking into account a spatially and temporally complex baseline, the characterization of temporal dynamics should trigger a further shift towards understanding task as a deviation from RS. In the following, we will go into more detail with two topics of special relevance to this thesis, and then describe the studies presented in the next chapters.

## 1.2 Temporal dynamics of resting state

### 1.2.1 Methodological approaches

One of the first, and an influential, study on the temporal dynamics of human RS as measured with fMRI is Chang and Glover (2010). The authors analyze how the connectivity of the posterior cingulate cortex (PCC) fluctuates over time; the PCC is a node in the DMN, and connectivity to other nodes inside and outside of the DMN is quantified. They apply two techniques that subsequently were adopted in other publications: time-frequency analysis using wavelet transform coherence (WTC), and sliding window FC. The former measures the power of interactions between pairs

of brain regions over time depending on both frequency and phase. This is valuable because the FC between two regions can fluctuate on several time scales. Thus, for each point in time, for each pair of regions, we obtain the power and the phase lag of the interaction for each frequency (see figure 1.5 for an example), as was previously done by Sato et al. (2006). The latter simply works by sliding a window (rectangular or with tapered edges) of fixed width along the time courses of a pair of regions and computing the correlation inside this window (see figure 1.6 for an example from a subsequent study by Hutchison et al. (2012)). The advantages and disadvantages of both techniques become obvious: while using WTC is more specific and allows us to examine all temporal scales involved, it also produces a large amount of data that can be difficult to interpret, in part due to the large variability across subjects. The sliding window approach allows only one width at a time and thereby fixes the time scale, but it is much simpler and easier to do. By combining both, the authors established that 1) FC between the examined nodes (brain regions) varies over time, 2) the variations are scale-dependent.

Sliding windows have been used in different ways. Valuable insights were provided by studies that combined them with ICA, investigating the changes of relationships between RSNs/ICNs at different spatial scales (Calhoun et al., 2008; Sakoğlu et al., 2010; Kiviniemi et al., 2011; Jones et al., 2012; Allen et al., 2012). This way, it was established that changes occur on the level of large-scale configurations of inter-network relationships, where some regions belonging to attention networks and the DMN act as “hubs” not only in terms of their static connectedness, but also in terms of the variability of their connections. This argues against the rigid assignment of brain regions to a functional network, and stresses that at certain points in time, the FC patterns can diverge greatly from the average. Furthermore, the idea of “connectivity states” was introduced: Along with relatively fixed FC patterns that recur over time, the way in which they interact also is recurrent in time.

When analyzing “states” occurring over time, it is also possible to directly consider co-activation of voxels, as was done by Tagliazucchi et al. (2012); Liu and Duyn (2013). They show that maps defined this way generally resemble those found using FC (although there are also some differences), suggesting that a big part of the occurrence of RSNs can be explained based on discrete events confined to a small number of time points.

A way to quantify global changes in network configurations is to use graph

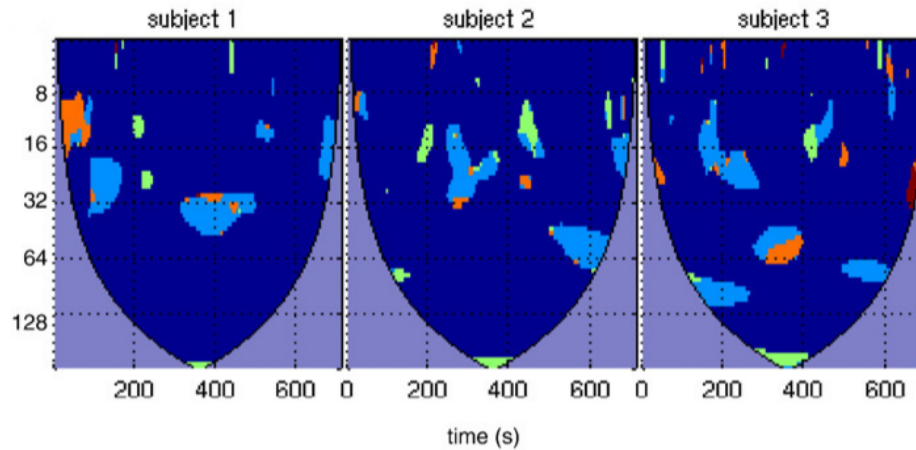


Figure 1.5: Example for the result of a time-frequency analysis via WTC for three subjects, using the PCC seed and another DMN node (medial prefrontal cortex). For each point in time (x-axes), the phase is plotted for each Fourier period (y-axes). Dark blue regions indicate no significant coherence, bright blue stands for a phase difference of  $0 \pm \pi/4$ , orange:  $\pi/2 \pm \pi/4$ , green:  $-\pi/2 \pm \pi/4$ , red:  $\pi \pm \pi/4$ . *Image from Chang and Glover (2010)*

measures as done for example by Bassett et al. (2011); Jones et al. (2012); Zalesky et al. (2014); Betzel et al. (2016). Here, the modular structure plays the most important role because it is a fingerprint of integration and segregation (Sporns and Betzel, 2016) and allows the examination of network organization on different spatial and temporal scales.

### 1.2.2 Bridging scales

The question of how observations from different imaging modalities fit together has always been of importance, emphasizing the relevance of different scales. Mantini et al. (2007) were the first to directly compare RSNs found in fMRI to patterns of power variations in different bands of EEG. Since EEG (as well as MEG, and electrophysiological recordings) possesses a much better temporal resolution than fMRI, it is no surprise that it has been known for a long time that measurements exhibit transiently stable configurations, termed “microstates”. A number of papers (Britz et al., 2010; Musso et al., 2010; Van de Ville et al., 2010) established a link between RSNs and EEG microstates, thus relating different temporal and spatial scales. Importantly, this speaks to a neural origin of RS fluctuations because EEG




Figure 1.6: Time courses of windowed correlations, averaged over voxels belonging to “oculomotor” network (OCM, red), ventral premotor network (vPM, blue), and white matter (WM, cyan), from one monkey in the study of *Hutchison et al. (2012)*

recordings are not influenced by the neurovascular coupling.

In MEG, it was found that maps similar to RSNs can be obtained by correlating the power envelopes of signals filtered in different frequency bands (band limited power), meaning that in MEG, RSNs are based on the modulations of the power in these bands (de Pasquale et al., 2010). Interestingly, coherence between nodes was found to exhibit a local maximum in the same band that was reported for fMRI, suggesting that slow fluctuations (below 0.1 Hz) are of neural origin (de Pasquale et al., 2010; Brookes et al., 2011; Hipp et al., 2012). In concordance with findings in EEG, correlations within (de Pasquale et al., 2010) and between (de Pasquale et al., 2012) networks depend on the frequency band, and occur transiently. Brookes et al. (2011) reported that fMRI-based RSNs share temporal patterns of fluctuations obtained via temporal ICA. This approach was also applied to fMRI data by Smith et al. (2012), who stressed that temporal modes differ in important points from spatial maps because RSNs interact over time.

Even though the temporal resolution of M/EEG is much higher than that of fMRI, the spatial resolution is in both cases very limited and requires complicated source reconstruction procedures. Although fMRI is superior in that regard, it is still not enough to resolve single cell populations, despite progress in decoding cortical columns (Yacoub et al., 2008). The only method that can provide both is that of electrophysiological recordings. The problem here, however, is the limited coverage and difficulties with *in vivo* recordings, especially in humans, and we can assume that this is the



reason why for a long time, the activity of the brain during rest was considered noise: when recording from a limited area, the coherent activity of macroscopic populations due to long-range connectivity is not evident.

Findings from simultaneous fMRI and electrophysiological recordings in animals (Shmuel and Leopold, 2008; Schölvinck et al., 2010; Pan et al., 2013) and humans (Nir et al., 2008) show that slow BOLD oscillations are correlated with neural activity, including LFP modulations, but also directly with slow components of the LFP (He and Raichle, 2009; Pan et al., 2013). This is consistent with abovementioned findings in M/EEG.

Another prominent feature in animal recordings are propagating waves during rest (Ferezou et al., 2007; Luczak et al., 2007; Matsui et al., 2016); they are believed to be directly related to the firing of single neurons. Travelling waves have also been reported with fMRI in rats (Majeed et al., 2011) and have recently been proposed as a mechanism for RSN generation for humans (Mitra et al., 2014).

### 1.2.3 Origin and function of spontaneous slow fluctuations

While it seems fairly clear that fluctuation in RS observed with fMRI BOLD are at least in part due to underlying neural activity, this does not answer the question as to how these fluctuations are generated. A first hint comes from the fact that the spatial patterns are greatly shaped by the underlying connectivity, both on the micro- (Bosking et al., 1997; Tsodyks et al., 1999; Kenet et al., 2003) and macroscopic level (Vincent et al., 2007; Hagmann et al., 2008; Greicius et al., 2009). Computational models provide valuable insights into how different kinds of dynamics on the microscopic level together with the anatomical connectivity can lead to the correlated slow fluctuations observed in fMRI (Ghosh et al., 2008; Honey et al., 2009; Deco et al., 2009) and MEG (Cabral et al., 2011, 2014). Specifically, this way it can be shown how slow whole-brain oscillations can emerge from faster (30-80 Hz) oscillations on a smaller scale. This direction argues that despite the low-pass filtering properties of the hemodynamic response and neurovascular coupling, correlations can be transmitted through spatial and temporal scales due to the dynamical system sitting close to a bifurcation.

As mentioned above, a number of studies has reported that slow fluctuations are not only the result of filtering from smaller to larger scales, but are present in the LFP or M/EEG measurements themselves (de Pasquale et al., 2010; Pan et al., 2013). This speaks to the fact that investigations of

long range connectivity in low frequency bands did not just arise as a side-effect of the properties of fMRI, but that communication between neural assemblies at long distances might involve lower frequencies due to longer delays (Varela et al., 2001; Monto et al., 2008). The discovery of resting state slow fluctuations at lower levels has led to the hypothesis that they are related to consciousness (He and Raichle, 2009) by virtue of integrating information across the entire brain. Also on the level of higher cognition, the notion has been put forward that fast and slow fluctuations are part of a hierarchy between fast, lower-order (sensory, motor, emotions) functions and slow, higher-order ones (deliberation, forming of beliefs) (Kringelbach et al., 2015). This is in line with the predictive coding-framework in which spontaneous activity has been placed (Ringach, 2009; Sadaghiani et al., 2010; Carhart-Harris and Friston, 2010). In this view, ongoing activity is modulated by top-down expectations, thereby influencing the processing of incoming stimuli.

While it is tempting to speculate about the functional significance of spontaneous fluctuations, on the level of BOLD, it has been shown that surprisingly big fluctuations can be explained by null models that do not assume any non-stationarity of interregional FC (Handwerker et al., 2012; Lindquist et al., 2014; Hindriks et al., 2015; Laumann et al., 2016). This is illustrated in Betzel et al. (2016) by the fact that the fluctuations are bounded by the average correlation. Several approaches have been employed to identify “real” non-stationary deviations (Sakoğlu et al., 2010; Zalesky et al., 2014; Hindriks et al., 2015), and the interpretation of the results is controversial.

### 1.3 Tensor factorization for neural data

We will now turn our attention to a more technical topic, which is nonetheless of immense importance in the field of RS fMRI. The extraction of “features” is a very important element when identifying modules and large-scale networks, and in this thesis, some unconventional methods are introduced, which require clarification beforehand.

#### 1.3.1 Formulation of the problem

Blind source separation is a very general problem, perhaps best-known as the “cocktail party problem”, where a number of microphones placed in different locations in a room full of people having different conversations picks up different mixtures of the audio signal emitted by the people. The

goal is to isolate the original signals, i.e. single conversations or people. Mathematically, the problem can be stated in matrix form as follows:

$$\mathbf{Y} = \mathbf{A}\mathbf{X} + \mathbf{E} \quad (1.1)$$

Here,  $\mathbf{Y}$  is an  $I \times T$  matrix containing the recordings of  $T$  time points from  $I$  microphones (or generally, sensors), which are assumed to be the result of a mixing process encoded in the matrix  $\mathbf{A}$  which is applied to the original conversations, encoded in  $\mathbf{X}$ .  $\mathbf{X}$  is of dimensionality  $J \times T$ , i.e. the number of sources can be different from the number of signals recorded by the sensors, and is usually unknown; this leaves the mixing matrix  $\mathbf{A}$  to be of dimensionality  $I \times J$ .  $\mathbf{E}$  is the noise matrix. In element-wise form, this equation becomes

$$y_i(t) = \sum_{j=1}^J a_{ij}x_j(t) + e_i(t) \quad (1.2)$$

Here, we have made the time dependence explicit. Translated to fMRI, or neural recordings in general,  $\mathbf{X}$  represents the neural signals that we want to estimate, which undergo some unknown transformation and filtering encoded in  $\mathbf{A}$  and are corrupted by noise  $\mathbf{E}$  before arriving at our sensors, which record the data in  $\mathbf{Y}$ . Estimating  $\mathbf{X}$  from  $\mathbf{Y}$  is also known as the inverse problem. This leads us to another way of formulating the problem, which is connected to the general linear model (GLM):

$$\mathbf{y}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{e}(t) \quad (1.3)$$

Now,  $\mathbf{y}(t)$  is a vector containing observations from the sensors at a certain point in time. The fact that  $I$  and  $J$  are not necessarily equal means that we have an ill-conditioned set of equations, i.e. the system is under- or overdetermined. Indeed, most of the time, the number of neural sources is unknown and large, and the number of signals is not big enough to accurately estimate them, for example the number of surface electrodes in EEG.

There are a number of well-known methods to deal with the inverse problem

$$\hat{\mathbf{X}} = \mathbf{W}\mathbf{Y}, \quad (1.4)$$

where  $\mathbf{W} = \hat{\mathbf{A}}^\dagger$  is the estimated unmixing matrix with  $(\cdot)^\dagger$  denoting the Moore-Penrose inverse. Thus, the task is to estimate  $\mathbf{W}$  such that it is optimal, i.e. minimizes the error between the estimated  $\hat{\mathbf{X}}$  and the real  $\mathbf{X}$ . In an under-determined system, this is achieved by imposing additional constraints. For example, in principal component analysis (PCA) as well as in singular value decomposition (SVD), it is assumed that the vectors in  $\mathbf{X}$  are orthogonal and that the original data are low-dimensional. Independent component analysis (ICA) assumes independence between the vectors in  $\mathbf{X}$ . In non-negative matrix factorization (NMF), the constraint is non-negativity, which can be sensible depending on the interpretation of the results. For example, if the vectors in  $\mathbf{X}$  are assumed to be membership weights, it makes sense to assume that they should be non-negative.

### 1.3.2 Generalization to tensors

The methods sketched above can be extended to higher dimensional arrays. In order to not confuse the dimensionality of a matrix, for example  $\mathbf{Y}$  in equation 1.1 being of dimensionality  $I \times T$ , with the fact that  $\mathbf{Y}$  is a matrix and therefore a two-dimensional array, we will use the term “2-way array” and refer to the “ways” of the arrays as “modes”. We will denote tensors in the following way:  $\underline{\mathbf{Y}}$  to distinguish them from matrices, although a tensor could also be 2-way and thus, a matrix (or even a vector).

Mostly, multiway arrays are known in neuroscience from group analysis, where additionally to the temporal and the sensor (spatial) dimension, we have a subject- and/or session dimension, or from time-frequency analysis, where frequency is an additional dimension. For the former, and most relevant to fMRI, a generalization of ICA has been developed and finds wide application (Beckmann and Smith, 2005).

The most general way of formulating this problem is known as the Tucker-N-model:

$$\underline{\mathbf{Y}} = \underline{\mathbf{G}} \times \{\mathbf{U}\} + \underline{\mathbf{E}}, \quad (1.5)$$

Figure 1.7 shows an illustration.  $\times$  denotes the n-mode product (the meaning of which will become clear in equation 1.6). For ease of description, we will limit our explanation to three modes, but it is easy to see that this can be generalized to any number of modes.  $\underline{\mathbf{Y}}$  is the data tensor and  $\{\mathbf{U}\}$  consists of three matrices  $\mathbf{A}$ ,  $\mathbf{B}$  and  $\mathbf{C}$  which contain components or features,

one set for each of the three modes of the tensor. The data is of dimensionality  $I \times T \times Q$ , and the matrices  $\mathbf{A}$ ,  $\mathbf{B}$  and  $\mathbf{C}$  are  $I \times J$ ,  $T \times R$ , and  $Q \times P$ , respectively.  $\mathbf{Y}$  and  $\{\mathbf{U}\}$  are linked by the core tensor,  $\mathbf{G}$ , of dimensionality  $J \times R \times P$ , which describes interactions between the components (Cichocki et al., 2009; Cichocki, 2013).

This can be written in element-wise form as for two dimensions in equation 1.2:

$$y_{itq} = \sum_{j=1}^J \sum_{r=1}^R \sum_{p=1}^P g_{jrp} a_{ij} b_{tr} c_{qp} + e_{itq} \quad (1.6)$$

It is important to note a few properties of this decomposition which are quite counter-intuitive and are due to the generality of this approach:

1. The number of components can differ from mode to mode, i.e.  $J \neq R \neq P$ . For example, we could have  $J$  modules/communities (spatial dimension),  $R$  temporal patterns, and  $P$  patterns of variability across subjects.
2. These components can interact whichever way possible, both between and across modes, as encoded in the core tensor. This means that any combination of components can contribute to the data.
3. Thus, components do not have any meaning *per se*.

Taken together, this obviously poses problems in the interpretation and visualization of decomposition results, which makes this formulation too general for neuroscience. Furthermore, there is usually a large number of possible solutions, as is the case for underdetermined linear systems, and consequently, any solution cannot be guaranteed to be optimal or unique. Regardless of these problems, it is interesting to consider the Tucker-N-model, because it reveals connections between approaches and makes it possible to understand virtually any decomposition, from PCA to tensor pICA, as a special case of the Tucker-N-model.

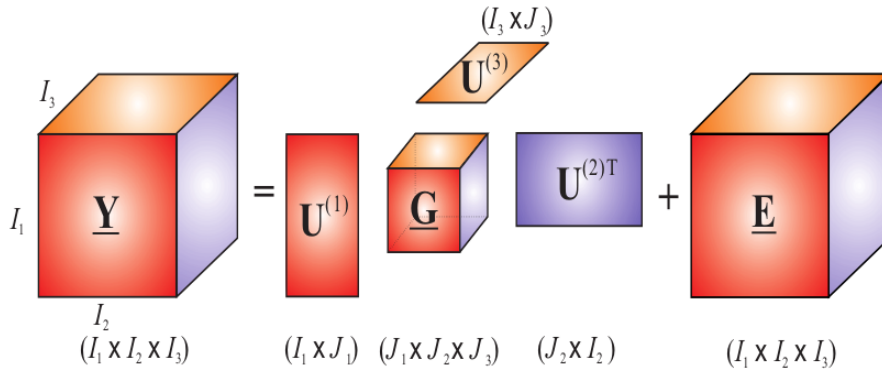


Figure 1.7: Illustration of Tucker-N-decomposition when  $N = 3$ . Image from Cichocki (2013)

### 1.3.3 Canonical Polyadic Decomposition and Non-negative Tensor Factorization

The solution to the lack of optimality and uniqueness is to impose further constraints, as is done in PCA, SVD, ICA, NMF, and other decomposition methods that are in use. The choice of these constraints will have a huge impact on the results and requires *a priori* knowledge of the problem. As mentioned above, a non-negativity constraint can make intuitive sense, and for the extraction of resting state networks, an independence constraint is derived from notions about integration and segregation of brain function.

A particularly intuitive and useful approach is Canonical polyadic decomposition (CPD), also known as CANDECOMP or parallel factor analysis (PARAFAC). Describing this as a special case of the Tucker-N-model, CPD can be denoted in three dimensions as

$$\mathbf{Y} = \sum_{j=1}^J \lambda_j \mathbf{a}_j \circ \mathbf{b}_j \circ \mathbf{c}_j + \mathbf{E} \quad (1.7)$$

Here,  $\lambda_j$  is a scaling factor of the components  $\mathbf{a}_j$ ,  $\mathbf{b}_j$ ,  $\mathbf{c}_j$ , meaning in turn that each component is of unit length, and  $\circ$  is the *outer product*, defined as

$$\mathbf{Y} = \sum_{j=1}^J a_j \circ b_j + \mathbf{E} = \sum_{j=1}^J a_j b_j^T + \mathbf{E}, \quad (1.8)$$

i.e. it contains all possible pairwise combinations of the elements of its constituent vectors. A 3-way-tensor that can be represented by three vectors is a rank-1 tensor and any tensor can be approximated by a sum of rank-1 tensors (see figure 1.8 for an illustration). Using this concept, the problems with uniqueness and optimality can be explained in terms of the impossibility to determine the true rank of a tensor, a problem which is readily solved for a matrix. In CPD, the additional constraints come from a requirement on the core tensor - which indeed does not even occur anymore in the above formulation because it is assumed to be zero everywhere except on its superdiagonal, where it is 1 everywhere (figure 1.8). As a consequence, the number of components is the same for all modes, and there is a clear matching between one component of each mode. For example, we can obtain one spatial map with an associated time course which varies across subjects.

Decompositions rely on numerical solutions which gradually optimize an objective function, subject to constraints. This function usually contains the distance between the tensor which is reconstructed from the extracted components and the original data tensor:

$$D_F(\underline{\mathbf{Y}} \parallel \llbracket \mathbf{A}, \mathbf{B}, \mathbf{C} \rrbracket) = \|\underline{\mathbf{Y}} - \llbracket \mathbf{A}, \mathbf{B}, \mathbf{C} \rrbracket\|_F^2 \quad (1.9)$$

$(\cdot)_F$  denotes the Frobenius norm. The equation is complemented with regularization terms according to the applied constraints.

Non-negative tensor factorization (NTF), then, is derived from CPD, adding a non-negativity constraint. Adding constraints shrinks the space in which the solution can reside, which typically leads to a poorer fit between the original data and the approximation made by the extracted components. However, the results are, if the constraint is appropriate, very interpretable and meaningful.

#### 1.3.4 CPD and NTF for community extraction and relationship to ICA

ICA relies on the assumption of independence in one of the modes. In fMRI, this is virtually always the spatial mode, with a few exceptions (Brookes




Figure 1.8: Illustration of CPD. The core tensor  $\mathbf{G}$  has ones on its superdiagonal and is zero everywhere else. One set of features is extracted for each mode of the tensor ( $\mathbf{A}$ ,  $\mathbf{B}$ ,  $\mathbf{C}$ ). Bottom: The tensor can then be approximated by rank-1 tensors made up of unique combinations of rows/columns from the feature matrices. *Image from Cichocki et al. (2009)*

et al., 2011; Smith et al., 2012). The reason for this is that we have many more voxels (usually tens of thousands) than time points (usually between 150 and 600, corresponding to scans of 5 to 20 minutes with a resolution of 1 frame per 2 seconds). Thus, ICA can be seen as a special case of a two-way Tucker decomposition, with an additional independence constraint. In order to do group analysis, which means handling an additional dimension, the data is usually “flattened”, i.e. concatenated in the temporal dimension (Calhoun et al., 2001a; Beckmann et al., 2009), thus making it possible to apply matrix methods. The method by Beckmann and Smith (2005) is an exception because it explicitly takes advantage of the multiway structure of the data. Cichocki (2013) directly proposed tensor methods for the analysis of neural data since tensors are a natural representation of these data in many cases.

However, neither of the authors applied tensor methods directly for community detection, i.e. starting from adjacency matrices. Community detection, or detection of modular structure, has become a major field in data science



because the representation of relationships between elements of a system as a network is extremely intuitive and useful (Fortunato, 2010). As mentioned above, this is also a major topic in neuroscience (Sporns et al., 2004; Rubinov and Sporns, 2010). Gauvin et al. (2014) used tensor techniques to detect *time-evolving* community structure in a social context; this was applied to neural data by Leonardi and Van de Ville (2013); Ponce-Alvarez et al. (2015). When using adjacency matrices, which are usually symmetrical, the first two modes become identical, i.e. the tensors have two spatial modes and one temporal mode:

$$\underline{\mathbf{Y}} = \sum_{j=1}^J a_j \circ a_j \circ b_j + \underline{\mathbf{E}} \quad (1.10)$$

While the contribution of ICA to science in general and neuroscience in particular cannot be stressed enough, independence is a rather strong assumption. CPD does not require such a constraint at all, and non-negativity, which can be considered a weak constraint (Cichocki, 2013), typically leads to unique solutions in NTF decompositions. Thus, it is possible to consider the spatial and the temporal dimension truly simultaneously, which promises a different point of view and new insights into the problem of community detection in neural data.

## 1.4 Studies in this thesis

In the following chapters, we will present two studies that make specific contributions to the field of fMRI RS. In chapter 2, we show explicitly that a dynamic mean field model of resting state (Deco et al., 2014), which relies on simple node dynamics and realistic anatomical connectivity, can reproduce spatial patterns over time (not only on average). These patterns resemble well-known RSNs and generalize across 24 healthy subjects. To this end, we apply CPD and NTF to low-resolution fMRI RS data and show under which circumstances this leads to a satisfying solution. Our results confirm that underlying (structural or effective) connectivity shapes functional connectivity to a large degree, although there are also important differences to the FC. It is shown that simple dynamics are sufficient to explain large-scale functional patterns, without the need for explicit non-stationary or oscillatory elements. Furthermore, only a strikingly small portion of the FC values over time is necessary to obtain these results, in line with previous findings (Tagliazucchi et al., 2012; Liu and Duyn, 2013).

In the second study, the analysis is expanded to include the temporal dimension. We show that FC fluctuates strongly on a global level by introducing three measures for assessing global changes in RS. We confirm that these fluctuations can be largely explained by the presence of a long-term correlation structure, but go on to identify certain elements that exhibit further structure. The relationship to RSN activations is investigated, and suggests a great degree of temporally specific integration across networks. Thus, we propose an interpretation of RSNs that stresses their overlap on multiple levels instead of their segregation. This thesis hopes to make its own contribution to understanding the spatio-temporal dynamics of human resting state, and triggering new directions in the approach of this topic.

---

---

---

---

---

## Robust extraction of spatio-temporal patterns from resting state fMRI

It is well-established that the patterns of correlation - or more generally, functional connectivity (FC) - between pairs of voxels or regions found in the human brain in many different tasks, are preserved in spontaneous activity recorded during rest. Furthermore, these patterns are not static, but fluctuate over time, exhibiting complex spatio-temporal dynamics. In this study, we use a whole-brain approach combining data analysis and modelling of FC dynamics between 66 ROIs covering the entire cortex. We simultaneously utilize temporal and spatial information by creating tensors that are subsequently decomposed into sets of brain regions (“communities”) that share similar temporal dynamics, and their associated time courses. The tensors contain pairwise FC computed inside of overlapping sliding windows. Communities are discovered by clustering features obtained from 24 healthy subjects, thereby ensuring that they generalize across subjects. First, we determine that at this resolution, four communities that resemble known RSNs can be clearly discerned in the empirical data: DMN, visual network, control networks, and sensorimotor network. Second, we use a noise-driven stationary mean field model which possesses simple node dynamics and realistic anatomical connectivity derived from DTI and fiber tracking. It has been shown to explain resting state FC as averaged over time and multiple subjects, however, this average FC summarizes the spatial distribution of correlations while hiding their temporal dynamics. Thus,

it is unclear whether the same type of model can reproduce FC at different points in time. We find that this is the case for all four networks using the spatio-temporal information revealed by tensor decomposition if nodes in the simulation are connected according to model-based effective connectivity. Furthermore, we find that these results require only a small part of the FC values, namely the highest values that occur across time and ROI pair. Our findings suggest that in resting state fMRI, FC patterns that occur over time are mostly derived from the average FC, are shaped by underlying structural connectivity, and that the activation of these patterns is limited to brief periods in time. We provide an innovative method that does not make strong assumptions about the underlying data and is generally applicable to resting state or task data from different subject populations.

## 2.1 Introduction

The question of how large-scale cortical function arises from underlying anatomical connectivity has been the object of much investigation since the advent of non-invasive imaging techniques (Vincent et al., 2007; Matsui et al., 2011; Wang et al., 2013), in particular since it was discovered that inter-areal functional relationships found under task conditions are maintained during rest (Biswal et al., 1995; Cordes et al., 2000; Beckmann and Smith, 2004; Fox et al., 2005). With magnetic resonance imaging (MRI) it is possible to obtain both functional and structural connectivities (FC and SC, respectively). Although there is large variability across subjects and sessions, both in SC (Heiervang et al., 2006) and FC measures (Mueller et al., 2013; Finn et al., 2015), studies using group averages have revealed general principles of information processing in the brain (Raichle et al., 2001; Doucet et al., 2011; Van den Heuvel and Sporns, 2011; Deco and Jirsa, 2012; Haimovici et al., 2013).

In order to connect SC and FC, computational models are an important tool for understanding how activity propagates from one node to another to produce the observed data (Honey et al., 2009; Cabral et al., 2012; Deco et al., 2014). Most models optimize their parameters by fitting the average FC. Only recently, the question of whether and how relevant information can be extracted from the fluctuations in pairwise FC strength, and how to describe the richness of the temporal dynamics, has received increasing attention in data analysis (Chang and Glover, 2010; Hutchison et al., 2012; Allen et al., 2012; Liu et al., 2013) and modelling (Hansen et al., 2014; Ponce-Alvarez et al., 2015). This has led to the notion of dynamic functional connectivity

(dFC).

Here, we use a dynamic mean field model (Wong and Wang, 2006) of the human cortex which has been shown to reproduce average resting state (RS) FC (Deco et al., 2014). It is our goal to determine whether simulated data exhibit FC patterns over time that resemble those of empirical data. To this end, we analyze RS data from 24 healthy subjects (Schirner et al., 2015) and compare to simulated data. The cortex is modelled by 66 nodes corresponding to 66 brain areas also used to parcellate the empirical data. The nodes are connected according to empirical SC derived from the same subjects.

We opt for tensor decomposition for extracting relevant and general features of the spatio-temporal dynamics. This method has been shown to work well for community detection (Gauvin et al., 2014) and has been applied to brain data (Cichocki, 2013; Leonardi and Van de Ville, 2013; Ponce-Alvarez et al., 2015). Unlike ICA, which has become the standard method for extracting RSNs (McKeown et al., 1998; Beckmann et al., 2005; Mantini et al., 2007), tensor factorization does not assume spatial independence of the underlying components, which is a strong constraint not directly motivated by the data. Here, such a constraint is not required and hence, the space of possible solutions is not unnecessarily restricted. Furthermore, it has the advantage that it can readily be used at our level of coarse spatial resolution.

The modelling approach aims at linking FC and SC. One conceptual problem of SC is that it provides neither directionality information nor the weights of the connections. These two points are addressed by the concept of effective connectivity (EC). SC can be viewed as an approximation to EC, and it is the latter that is genuinely related to the dynamics in network models (Friston, 1994). Reversely, underlying connectivity (SC or EC) can be inferred from FC, or more generally, from the dynamics found in the data, through the same kinds of models. Gilson et al. (2016) developed a method to extract EC from RS fMRI data using a noise diffusion model. They show that the EC that accounts best for empirical FC significantly differs from the SC in a number of points. We use both SC and EC as underlying connectivity in our model and explore how their properties are linked to the spatio-temporal patterns found in empirical and simulated data.

Table 2.1: Parameters of the DMF model.

Excitatory	Inhibitory
$a_E = 310(\text{nC}^{-1})$	$a_I = 615(\text{nC}^{-1})$
$b_E = 125(\text{Hz})$	$b_I = 177(\text{Hz})$
$d_E = 0.16(\text{s})$	$d_I = 0.087(\text{s})$
$\tau_E = \tau_{\text{NMDA}} = 100(\text{ms})$	$\tau_I = \tau_{\text{GABA}} = 10(\text{ms})$
$I_{0,E} = 0.382$	$I_{0,I} = 0.267$
$w_{EE} = 0.21$	$w_{II} = 1, w_{IE} = 0.15$

## 2.2 Methods

### 2.2.1 Empirical data

RS fMRI as well as corresponding diffusion weighted (dw) MRI data were collected from 24 healthy participants (11 female) at the Charité Berlin, Germany, by Petra Ritter and co-workers. The original dataset consisted of 49 subjects, but we chose only those aged 18 to 35 years (mean 25.7 years) since it is known that FC changes with age (Meunier et al., 2009). Each dataset amounts to 661 time points recorded at TR=2s, i.e. about 21 minutes. In the same session, EEG was also recorded, but we do not use the data here. RS BOLD was recorded while subjects were asked to stay awake with their eyes closed, using a 3T Siemens Trim Trio scanner and a 12 channel Siemens head coil. Voxel time courses are averaged inside ROIs defined by the Desikan-Killiany atlas (Desikan et al., 2006) as implemented in FreeSurfer. We removed the areas labeled as corpus callosum on both sides since they only contain white matter, amounting to 33 cortical ROIs for each hemisphere. The dwMRI data were subjected to fiber tracking to obtain structural connectivity (SC) matrices for each subject. Details are available in Schirner et al. (2015).

### 2.2.2 Model data

A mean field approximation of a network of populations of leaky integrate and fire neurons (Wong and Wang, 2006) was used to simulate RS activity as described in Deco et al. (2014). The excitatory populations are connected using a) an average over the SC matrices derived from dwMRI data (section 2.2.1), and b) an effective connectivity (EC) matrix derived from these SC matrices as described in section 2.2.9.

The activity of the populations is computed using a set of coupled nonlinear stochastic differential equations:

$$\begin{aligned}
I_i^{(E)} &= I_{0,E} + w_{EE}S_i^{(E)} + G \sum_j C_{ij}S_j^{(E)} - w_{EI,i}S_i^{(I)} \\
I_i^{(I)} &= I_{0,I} + w_{IE}S_i^{(E)} - w_{II}S_i^{(I)} \\
r_i^{(E)} &= H^{(E)}(I_i^{(E)}) = \frac{a_E I_i^{(E)} - b_E}{1 - \exp(-d_E(a_E I_i^{(E)} - b_E))} \\
r_i^{(I)} &= H^{(I)}(I_i^{(I)}) = \frac{a_I I_i^{(I)} - b_I}{1 - \exp(-d_I(a_I I_i^{(I)} - b_I))} \\
\frac{dS_i^{(E)}(t)}{dt} &= -\frac{S_i^{(E)}}{\tau_E} + (1 - S_i^{(E)})\gamma r_i^{(E)} + \sigma\nu_i(t) \\
\frac{dS_i^{(I)}(t)}{dt} &= -\frac{S_i^{(I)}}{\tau_I} + r_i^{(I)} + \sigma\nu_i(t)
\end{aligned} \tag{2.1}$$

Constants are listed in table 2.1; see figure 2.1A for an illustration. Super-/subscripts  $E$  and  $I$  denote the excitatory and inhibitory pools of population  $i$ , respectively.  $I_i$  denotes synaptic input currents, which are turned into population firing rates  $r_i$  via sigmoid transfer functions  $H(\cdot)$ .  $S_i$  denotes the average synaptic gating variable, or activity, and  $\nu(t)$  is Gaussian noise with amplitude  $\sigma = 0.01$ . The kinetic parameter  $\gamma = 0.641$ .

Synaptic currents  $I_i$  are a result of of inputs from the local network, i.e.  $S_i^{(E)}$  and  $S_i^{(I)}$ , and inputs from other network nodes  $j$ , i.e.  $S_j^{(E)}$ . Local inputs are governed by four weights,  $w_{EE}$ ,  $w_{EI,i}$ ,  $w_{IE}$ , and  $w_{II}$ . Additionally, there are constant inputs to each population, denoted by  $I_{0,E}$  and  $I_{0,I}$ . Inputs from other parts of the network are provided by the excitatory populations and weighted by the entries  $C_{ij}$  for the connectivity from region  $i$  to region  $j$ , noted in the SC or EC. The diagonal of  $\mathbf{C}$  is set to 0. Weights are scaled by the global coupling parameter  $G$ .

The feedback inhibition,  $w_{EI,i}$ , is adjusted before the simulation to ensure that the network is in its asynchronous state where we only have one stable

fixed point with firing rates between 3 and 10 Hz for all regions (Deco and Jirsa, 2012). We can determine the stability of the system by taking advantage of a semi-analytical solution (Deco et al., 2014). We calculate the Jacobian matrix and confirm that all eigenvalues are negative with zero imaginary parts. Simulations were only performed for values of  $G$  that warranted stability of the system.

Number and length of simulations are matched to the empirical data. BOLD time courses are obtained from the synaptic activities of the excitatory pools via the Balloon-Windkessel model (Friston et al., 2003; Deco et al., 2013). Time courses are downsampled to match the TR of the experimental data (figure 2.1B,C).

---

Figure 2.1: (see next page) Overview of analysis pipeline that leads to tensorized data. For simulated data, panels A to E apply, for empirical data, panels C to E. **A** Schematic view of dynamic mean field (DMF) model used to simulate synaptic activity. Each brain area is modelled by a pair of excitatory (E) and inhibitory (I) pools. The local connectivity is governed by the four weights  $w_{EE}$ ,  $w_{EI,i}$ ,  $w_{IE}$ , and  $w_{II}$ , whereby  $w_{EI,i}$  is adjusted for each population individually so as to keep the firing rates at a low level (3-10 Hz). Black lines with spheres signify GABA connections, black arrows, NMDA. Gray arrows are long range connections mediated by AMPA synapses and whose weights are set by the SC or EC matrix. **B** Synaptic activity is simulated on a millisecond scale. **C** BOLD activity on the scale of seconds is obtained from simulated activity by applying the Balloon-Windkessel hemodynamic model. **D** Starting from these BOLD time courses, we compute time dependent dFC(w)-matrices from overlapping time windows of width 120s (60 frames) starting at time points  $t = 1, 2, \dots, W$ . Each dFC(w)-matrix has dimensions  $N \times N$ ,  $N$  being the total number of ROIs. **E** This results in an  $N \times N \times W$  tensor. The entry at  $dFC_{ij}(w)$  is the functional connectivity value of ROIs  $i$  and  $j$  in window  $t$ . dFC(w)-matrices are symmetric.

---



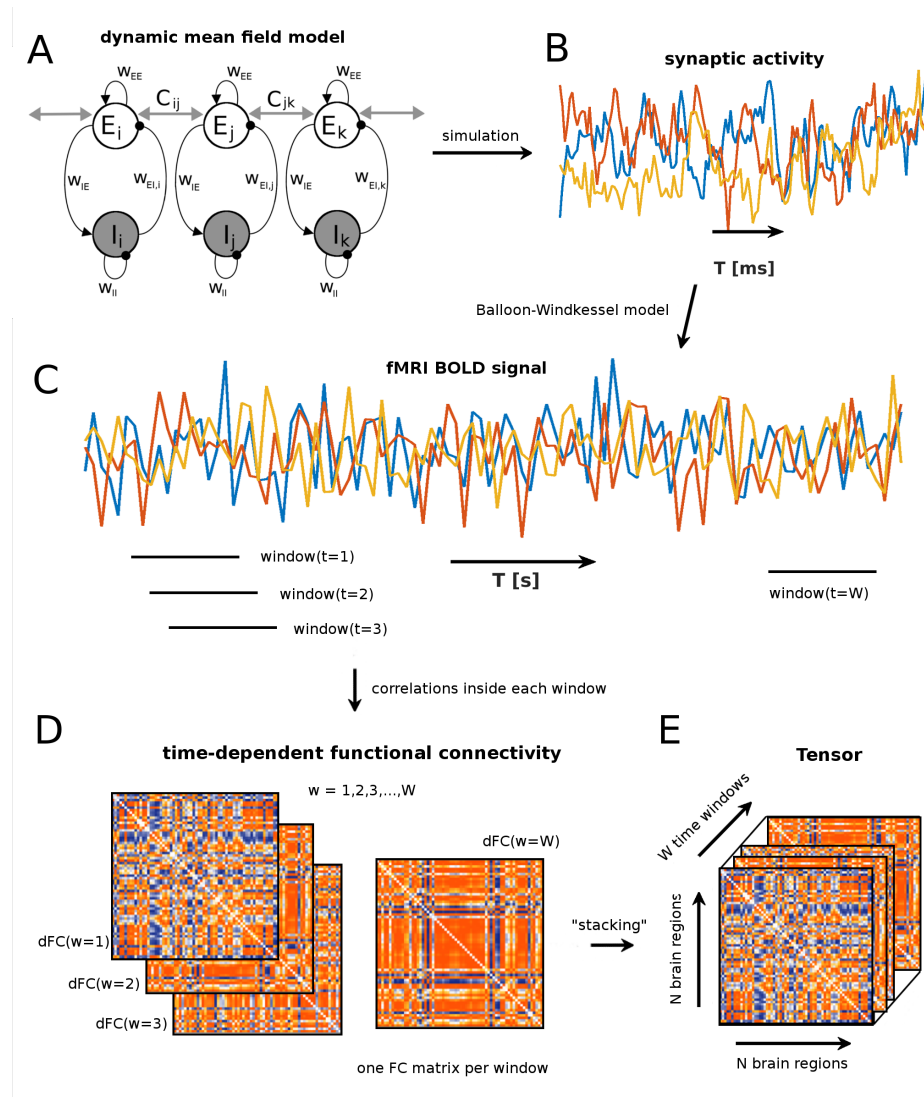


Figure 2.1: Description on previous page

### 2.2.3 Tensorization of the data

In order to take advantage of the temporal information, we adopt the widely used approach of sliding windows to compute time-dependent dynamic FC (dFC). We use overlapping windows  $w$  of width 120s (60 data points, TR=2s) that we advance along the time course in increments of 2s (1 frame), which results in  $W = T - 60$  windows for each dataset (subject or simulation),  $T$  being the number of frames. For each window, we compute an  $N \times N$  matrix of pairwise connectivity values,  $dFC(w)$ . By concatenating these matrices along the temporal dimension, we represent each dataset as a tensor of dimensions  $N \times N \times W$  (see figures 2.1D and E).

We use two measures of FC: on the one hand the most widely used one, i.e. Pearson correlation, on the other, mutual information (MI), a non-negative and nonlinear measure that allows us to constrain the tensor decomposition. The calculation of MI follows Kraskov et al. (2004) and is based on nearest neighbor distances, thus being adaptive and continuous.

In short, when estimating the MI between the 60 data points inside our window, we determine the distance between any two points  $[x_i, y_i]$  and  $[x_j, y_j]$  with  $i \neq j$  and take the maximum norm, i.e.

$$d_{ij} = \| z_i - z_j \| = \max\{\| x_i - x_j \|, \| y_i - y_j \|\} \quad (2.2)$$

The nearest neighbor to each point  $[x_i, y_i]$  is the point with the minimum  $d_{ij}$ , and we term this distance  $\epsilon_i$ . For each point  $[x_i, y_i]$ , we count how many points are within this minimum distance  $\epsilon_i$ , separately for the x- and y-directions, resulting in two numbers  $n_x(i)$  and  $n_y(i)$ . We estimate MI as

$$I(X, Y) = \psi(k) - \langle \psi(n_x + 1) + \psi(n_y + 1) \rangle + \psi(N) \quad (2.3)$$

where  $X$  and  $Y$  are the two time series, and  $\psi(\cdot)$  denotes the digamma function

$$\psi(x) = \frac{d}{dx} \ln(\Gamma(x)) \quad (2.4)$$

$k = 1$  because we consider the nearest neighbor;  $N$  is the number of data points, i.e. 60. Since this is a continuous measure, it is possible for  $I(\cdot, \cdot)$  to become negative whenever the following inequality is satisfied:

$$\psi(k) + \psi(N) < \langle \psi(n_x + 1) + \psi(n_y + 1) \rangle$$

This happens when there is very little MI and many points are closer than the nearest neighbor. In these cases, we simply set MI to zero.

#### 2.2.4 Extracting communities from low resolution fMRI data

We apply tensor factorization to both the empirical and the model data (figure 2.2). The problem for the three-dimensional case treated here can be formulated as follows (Cichocki et al., 2009):

$$\underline{\mathbf{Y}} = \sum_{f=1}^F a_f \circ b_f \circ c_f + \underline{\mathbf{E}} = \hat{\underline{\mathbf{Y}}} + \underline{\mathbf{E}} \quad (2.5)$$

where  $\underline{\mathbf{Y}}$  is the data tensor of dimensions  $N \times N \times W$ , defined as in section 2.2.3.  $F$  is the number of features we wish to extract.  $\mathbf{A} = [a_1, a_2, \dots, a_f]$ ,  $\mathbf{B} = [b_1, b_2, \dots, b_f]$  and  $\mathbf{C} = [c_1, c_2, \dots, c_f]$  are the factor matrices that contain the features of each of the three dimensions, respectively. Here,  $\mathbf{A} = \mathbf{B}$  due to the symmetry of the  $dFC(w)$  matrices. They contain  $F$  vectors of length  $N$  with weights that are interpreted as membership values for a community. The third set  $\mathbf{C}$  contains their associated time courses and is of dimensions  $F \times W$ .

$\hat{\underline{\mathbf{Y}}}$  is an approximation of the data based on the features, and  $\underline{\mathbf{E}}$  is the error/noise not described by the features. Hence, the distance between  $\underline{\mathbf{Y}}$  and  $\hat{\underline{\mathbf{Y}}}$  can be used to assess how well the extracted features approximate the data. We use the Frobenius norm in the case of continuous, non-thresholded tensors and the Hamming distance for thresholded, binarized tensors (section 2.2.5). Note that in the latter case, the result of the decomposition is continuous although the input is binary, so we threshold and binarize the reconstructed tensor such that the number of ones is preserved.

This technique is based on the very general Tucker model which can be viewed as a generalization of SVD to tensors. Unlike for SVD, though, convergence to a unique and optimal solution is not guaranteed. Consequently, it is impossible to determine the true rank of the tensors and thus, the appropriate number of features  $F$ . This problem is mitigated by the inclusion of further constraints, in this case, non-negativity when using absolute value of correlation or MI to construct the tensors (section 2.2.3).

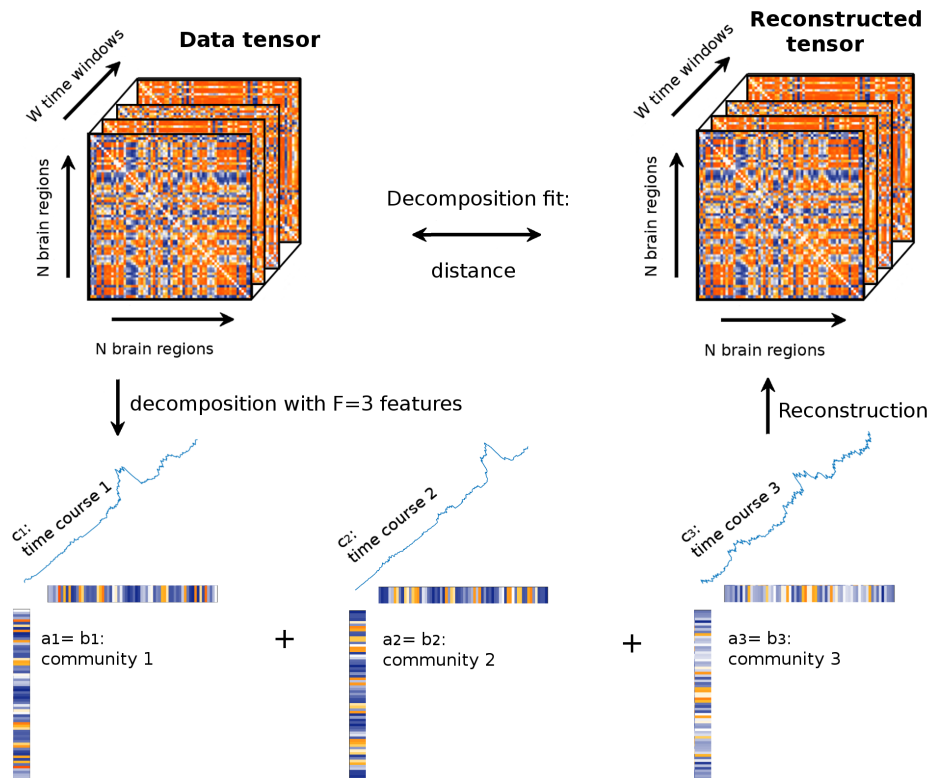


Figure 2.2: Illustration of tensor factorization with  $F = 3$  features. The original data tensor (top left) is decomposed into three components, i.e. three communities with their associated time courses are extracted. As shown in the bottom row, each set of vectors, consisting of two identical “community vectors”  $a_f$  and  $b_f$  and a time course,  $c_f$ , constitutes a rank-1 tensor, which can be explicitly computed by taking the outer product as in equation 1.8 (taking the outer product of  $a_f$  and  $b_f$ , one would obtain an FC matrix like the “layers” of the tensor). By adding these rank-1 tensors, we obtain an approximation of the data tensor, i.e. a reconstructed tensor (top right), and can evaluate the fit between the two by taking the distance.

To decompose tensors constructed using correlation, we use the algorithm described in Phan et al. (2013). For the non-negative measures (absolute value of correlation, mutual information), we apply non-negativity of the resulting features as an additional constraint and use the algorithm described in Kim and Park (2012). Both algorithms are implemented in Matlab, requiring the tensor toolbox (Bader et al., 2015; Acar et al., 2011), and available on-line.

### 2.2.5 Thresholding

We reduce noise by thresholding and binarizing the tensors. This approach was also chosen in Ponce-Alvarez et al. (2015). It would perhaps be most desirable to use only the dFC values that are significant. However, it is too time demanding to generate the appropriate number of surrogate datasets to achieve the desired significance level to account for the high number of ROI pairs, windows and subjects/simulations. Hence, we simply use the  $x$ -th percentile as a threshold  $\theta$ , where  $x = \{0, 75, 80, 90, 91, \dots, 99\}$  and, for  $x > 0$ , set all elements  $Y_{ijt}$  of tensor  $\hat{\mathbf{Y}}$  which are bigger than or equal to that percentile to one and everything else to zero. Note that we do not make any claims about non-stationarity of the FC time courses.

### 2.2.6 Surrogate data

To further validate our results, we conduct analyses with surrogate data alongside those for real data. Surrogate data are constructed by removing the pairwise correlation structure of the original time series while keeping the Fourier spectrum constant. More specifically, we Fourier transform our original time courses  $x_i(t)$  of region  $i$  for time points  $t = 1, 2, \dots, T$  using Fast Fourier Transform and add random phases  $\varphi_r$  to each frequency bin before transforming back.

$$\begin{aligned}\tilde{x}_i(k) &= \sum_{t=1}^T x_i(t) e^{-j \frac{2\pi kt}{T}} \\ x_i^{\text{surr}}(t) &= \frac{1}{T} \sum_{k=1}^T |\tilde{x}_i(k)| e^{-j \frac{2\pi kt}{T} + \varphi_r}\end{aligned}\tag{2.6}$$

$k = 1, 2, \dots, T$  and  $\varphi_r$  is uniformly distributed between  $-\pi$  and  $\pi$ . This results in time courses that have the same spectral properties and autocorrelations as the original data, but are uncorrelated.

### 2.2.7 Constructing templates from empirical data

We construct spatial templates directly from the empirical data. First, we decompose each subject’s data tensor as described above, using a range of numbers  $F$  of features. The true value of  $F$  is impossible to determine due to the aforementioned problems with uniqueness and optimality of the solution. In our case, the goal is to extract spatial maps that are common across subjects, so we apply a simple clustering algorithm (K-means clustering) to the set of all  $F \cdot S$  features (i.e. pooled from all subjects) with different numbers  $K$  of clusters. We test the quality of the clustering by evaluating the mean silhouette value (de Amorim and Hennig, 2015). The resulting cluster centers of the parameter combination with the maximum silhouette value are used as templates.

For each point  $i$  (here, an  $N$ -dimensional feature) the average distance (in terms of correlation) to points assigned to the same cluster is evaluated, denoted by  $a_i$ , as well as the smallest average distance to all points assigned to a different cluster (i.e., the closest cluster), denoted by  $b_i$ . Then the silhouette value is calculated as

$$s(i) = \frac{b_i - a_i}{\max\{a_i, b_i\}} \quad (2.7)$$

Obviously,  $s(i)$  will lie between -1 if the point is entirely in the wrong cluster, and 1 if the assignment is perfect. Taking the mean over all features gives an estimate of how well the data points are clustered.

### 2.2.8 Calculating the overlap between templates and features

In order to determine how well, overall, communities extracted from simulated data match those from empirical data, we compute the average overlap between all simulated spatial features with any of the templates, comparing connectivities (EC vs. SC) and using the full range of global coupling parameter values ( $G$ ), i.e.  $G = 0.5, 0.6, \dots, 4$  for simulations with the SC and  $G = 1, 1.1, \dots, 6$  for simulations with the EC.

Due to the high thresholds, the features are somewhat sparse. Therefore, correlation is not a good choice to measure the distance between features and templates. Instead, we use confusion matrices and Cohen’s kappa. Briefly, we quantize the features on three levels and compute the overlap

between any two vectors (one template, one feature extracted from simulated data) by determining the overlap for each level, creating a confusion matrix. Cohen's kappa is a summary of the confusion matrix:

$$\kappa = \frac{P_a - P_e}{1 - P_e} \quad (2.8)$$

$P_a$  is the overlap,  $P_e$  is the expected overlap.

As the overall match between a feature extracted from simulated data and the templates of the empirical data, we consider the mean over the maximum  $\kappa$  values for each feature in each simulation. In other words, we assume that each simulated feature corresponds to only one template. We compute an overall match for each value of  $G$  by averaging this value over the features and simulations.

### 2.2.9 Effective connectivity

In this study, we compare simulation results from two different underlying connectivities. On the one hand, we have dwMRI-derived SC containing estimates of fiber densities from the same 24 subjects whose BOLD signals are analyzed, on the other, model-based EC. In the following, we briefly describe the method developed by Gilson et al. (2016) for constructing EC matrices by combining SC and FC. The key is to extract information about cortical interactions from BOLD covariances with non-zero time shifts:

$$Q_{ij}^\tau := \langle (x_i^t - \bar{x}_i)(x_j^{t+\tau} - \bar{x}_j) \rangle \quad (2.9)$$

$Q_{ij}^\tau$  is the covariance between BOLD time courses  $x$  of ROIs  $i$  and  $j$ , with  $x_j$  shifted by  $\tau$  against  $x_i$ . The angular brackets denote averaging over randomness due to noise in the model, such that the mean BOLD for node  $i$  is  $\bar{x}_i = \langle x_i^t \rangle$ . For  $\tau \neq 0$ , this matrix is non-symmetric. The goal is to estimate the underlying connectivity such that the model minimizes the error between model covariances ( $Q^0$  and  $Q^\tau$ ) and their empirical counterparts ( $\hat{Q}^0$  and  $\hat{Q}^\tau$ ), for a given  $\tau$  equal to 1 TR.

EC is model based, meaning that there is an underlying assumption of how activity propagates through the brain using the present connections to activate the nodes. We use a noise diffusion model with a static nonlinearity,  $\Phi$ :

$$dx_i^t = \left[ -\frac{x_i^t}{\tau_x} + \Phi \left( \sum_{k \neq i} C_{ik} x_k^t + e \right) \right] dt + dB_i^t \quad (2.10)$$

Time course  $x_i$  of region  $i$  is subject to an exponential decay with time constant  $\tau_x$  at each time point  $t$ .  $C$  is the connectivity matrix that contains weights linking regions  $i$  and  $j$ , and the sum is over all regions  $k$  from which  $i$  receives input. This means that activation is only provided by the input from other nodes, weights for which are defined in  $C$ . The background input  $e$  is shared by all  $i$ . Fluctuations of the activities are driven by Gaussian noise  $dB_i^t$ . The model directly simulates BOLD activity, hence the Jacobian of the system is simply:

$$J_{ij} = -\frac{\delta_{ij}}{\tau_x} + C_{ij} \Phi' \left( \sum_{k \neq i} C_{ik} \bar{x}_k + e \right), \quad (2.11)$$

where  $\delta_{ij}$  is the Kronecker delta and  $\Phi'$  denotes the first time derivative of  $\Phi$ . Therefore, the model is solely constrained by  $C$ .

We want to estimate  $J$  and therefore  $C$  such that it satisfies the steady state of the second order fluctuations:

$$\begin{aligned} JQ^0 + Q^0 J^\dagger + \Sigma &= 0 \\ Q^\tau &= Q^0 \mathbf{exp}(J^\dagger \tau) \end{aligned} \quad (2.12)$$

$\Sigma$  is the noise matrix with diagonal terms  $(\sigma_i)^2 = \langle dB_i^t dB_j^t \rangle$ ;  $\dagger$  denotes the matrix transpose and  $\mathbf{exp}$  the matrix exponential. Since we are using empirical covariances estimated from fMRI data, the objective covariances  $\hat{Q}^0$  and  $\hat{Q}^\tau$  are very noisy and make a direct estimation via an analytical approach unfeasible. Therefore, we use the iterative Lyapunov optimization (LO) procedure described in Gilson et al. (2016).

The update works by simulating the BOLD activity using the model defined in equation 2.10 without noise and the current connectivity  $C$ , so as to evaluate the mean activity  $\bar{x}_i$  for all regions, yielding the Jacobian  $J$  in equation 2.11. Then, the model covariances  $Q^0$  and  $Q^\tau$  are given by equation 2.12, using the Bartels-Stewart algorithm for the first line and using the current values for  $J$  and  $\Sigma$ . The model covariance matrices are compared to the objective covariances  $\hat{Q}^0$  and  $\hat{Q}^\tau$ . Then, the Jacobian update is evaluated according to:



$$\Delta J = \frac{1}{\tau} [(Q^0)^{-1} (\Delta Q^0 + \Delta Q^\tau \mathbf{exp}(-J^\dagger \tau))]^\dagger \quad (2.13)$$

Finally, we obtain the connectivity update  $\Delta C_{ij} = \Delta J_{ij} / \Phi(\sum_k C_{ik} \bar{x}_k + e)$ .

We apply a mask when updating  $C$ , thus only tuning connections that are present in the SC as well and are above a certain threshold. The only exception are the elements on the secondary diagonal, which are added regardless of whether they are present in the SC or not, in order to account for homotopic connections.

In addition, the noise matrix  $\Sigma$  (see equation 2.12) is optimized at the same time as  $C$ . We assume that each node receives independent noise, meaning that  $\Sigma$  is diagonal. The update is performed such that the model variances coincide with the empirical values:

$$\Delta \Sigma_{ii} = \epsilon_\Sigma (\hat{Q}_{ii}^0 - Q_{ii}^0) \quad (2.14)$$

It was shown that a time shift  $\tau$  equal to 1 or 2 TR gives a good estimation performance (Gilson et al., 2016). In fact,  $\tau$  has to roughly match the time scale on which the neural activity decays, i.e.  $\tau_x$  in equation 2.11. The latter is estimated from the slope of the autocorrelation of each region (the slope is close to  $1/\tau_x$ ), and results in  $\tau = 5.3s$  which leads us to set the time shift to 1 TR=2s.

### 2.2.10 Analysis pipeline

To summarize the methodology of the paper, we give a brief overview of the steps involved in obtaining the results described in the next section. Illustrations are shown in figures 2.1 and 2.2. We have  $S = 24$  sets of empirical resting state data (from 24 subjects) of length  $T = 661$  frames (TR=2s). Voxel time courses are averaged inside of  $N = 66$  ROIs that cover the entire cortex. To match this number, we simulate 24 times for each one of evenly spaced values of the global coupling parameter  $G$  from a suitable range, using the DMF model. We obtain 2 sets of simulations, one using the SC matrix and one using the EC matrix to set the underlying connectivity between the 66 regions. Each dataset is tensorized by computing dFC matrices inside of rectangular sliding windows  $w$  of width 2 minutes (60 frames) which are moved by one frame, resulting in  $W$   $dFC(w)$  matrices for each subject/simulation. These matrices are concatenated into

tensors of dimensions  $N \times N \times W$ . Entries of  $dFC(w)$  are calculated using correlation (positive and negative values), the absolute value of correlation (non-negative values) and mutual information (non-negative values). Each tensor is then decomposed into sets of spatial (communities) and temporal (time courses) features.

In a first step, in order to obtain community templates, we decompose the empirical tensors using different thresholds  $\theta$ , binarizing the tensors in the case of non-zero thresholds. Since the rank of the tensors is unknown, we use different numbers  $F$  of features, where  $F$  ranges from 3 to 9. Hence, we obtain  $F \cdot S$  spatial (and temporal) features for each  $F$ . We do not expect the temporal features to have anything in common except some general dynamic properties, so we continue with only the spatial features and cluster them, calculating the silhouette value as a quality measure for each instance of clustering. We choose the combination of  $F$ ,  $K$  and  $\theta$  that yields the most well-defined clusters, i.e. the highest silhouette value. The means of those clusters are the templates (i.e.,  $K$  is the number of templates). In a second step we extract features from simulated data ( $S = 24$  runs, different values of  $G$ ) using the same  $F$  and  $\theta$ , and compare each feature to the templates. We calculate an overall match for each value of  $G$  by taking the mean over the maximum match of each feature with any of the templates.

## 2.3 Results

### 2.3.1 A method for extracting RSNs from single subject data

Our general goal is to understand the spatio-temporal dynamics of human resting state (RS) fMRI, using time-dependent, or dynamic, FC (Chang and Glover, 2010; Hutchison et al., 2012; Allen et al., 2012; Liu et al., 2013). To this end, we combine data analysis of long-range connectivity and a whole-brain modelling approach (Deco et al., 2014) to investigate whether the dynamics of the model can reproduce the empirical data. We apply tensor decomposition (Cichocki et al., 2014), a method that allows us to simultaneously consider spatial and temporal structure of FC. We compute pair-wise dynamic FC (dFC) using mutual information (MI) (Kraskov et al., 2004) inside of overlapping sliding time windows (Chang and Glover, 2010; Kiviniemi et al., 2011; Hutchison et al., 2012; Allen et al., 2012; Leonardi and Van de Ville, 2013), obtaining a dFC matrix,  $dFC(w)$ , for each window  $w$ . These matrices are then concatenated along the temporal dimension

into a 3-way-tensor (figure 2.1C-E). Additionally, we apply a binarization threshold in order to reduce noise. This is done for both empirical as well as simulated data.

Tensor factorization allows us to decompose dFC of a single subject (or simulation run) into spatial patterns, so-called “communities”, and associated time courses. The spatial patterns are expected to be common among subjects, while the temporal evolution of each community is specific to a subject. We extract  $F = 3, 4, \dots, 9$  communities/time courses (features) (Beckmann and Smith, 2004; Mantini et al., 2007) from each of  $S = 24$  subjects which gives a pool of  $F \cdot S$  communities. These are then grouped into  $K = F + 1, F + 2, \dots, 10$  clusters using K-means clustering in order to find common patterns. Goodness of clustering is assessed using the silhouette value (section 2.2.7); a high value indicates that the clusters represent the data well, i.e. cluster centers can be seen as “prototype communities”, or templates, that can be used on the whole group of subjects. In order to validate our results, we compare them against surrogate data and consider the difference in clustering performance. The surrogate data do not exhibit any long-term correlations, so any clustering structure is due to their Fourier spectrum, autocorrelations, and due to the method itself. We find that  $F = 3$  and  $K = 4$  results in the maximum value of 0.54. Phase randomized surrogate data (section 2.2.6), on the other hand, only reach 0.08. This indicates that, while there are clearly still a lot of inter-individual differences, assuming a cluster structure is supported by the data.

Figure 2.3A shows the difference between silhouette values for real and surrogate data for all combinations of  $F$  and  $K$ . Panel B visualizes the clusters by showing the correlations between communities. The clusters are quite well separated and correspond to four common patterns that are similar to previously described resting state networks, namely the default mode network, somatomotor network, right and left control networks, and visual network (figure 2.4).

### 2.3.2 Mutual information and a high threshold produce communities that generalize well

We use MI to compute dFC (Kraskov et al., 2004) because we find that MI is a better choice than correlation, based on clustering performance as well as reconstruction fits. Additionally, we apply a binarization threshold to the MI values, keeping only the highest ones, in order to reduce noise,

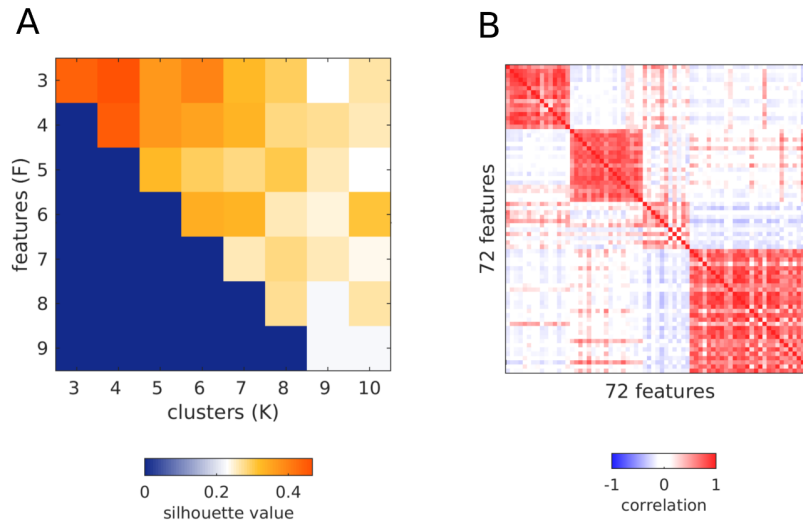


Figure 2.3: Result of template extraction with optimal parameter settings. **A** Silhouette values for clustering all features extracted from MI tensors for the best threshold, i.e. 98th percentile. **B** Correlation between the same features, ordered by assignment to  $K = 4$  clusters. The clusters are in the same order as the templates in figure 2.4.

which is in the following demonstrated to be necessary in order to obtain the desired generalized communities.

We start out without applying a binarization threshold, computing pairwise dFC values with Pearson correlation. Decomposing these tensors and clustering the resulting communities, we obtain the best silhouette value at  $F = K = 3$  with a value of 0.31 for real data, and 0.19 for surrogates. This means that most of the cluster structure is due to spectral properties and autocorrelation of the data rather than real dFC. In comparison, when using MI without a threshold, we obtain the best value at  $F = K = 3$  with 0.45 for real and 0.10 for surrogate data.

One advantage of using MI is that it allows us to apply an additional constraint when decomposing the data, i.e. non-negativity. In order to exclude that the poor result obtained with correlation is due to this, we repeat the procedure replacing the negative correlations with their absolute values. Again,  $F = K = 3$ , and we obtain 0.44 for real data, i.e. just as high as for MI, but again, 0.19 for surrogate data.




Figure 2.4: Templates plotted in 3D. From top to bottom: default mode network, somatomotor network, control networks, visual network. Although the templates are obtained at a resolution of 66 cortical regions, for better visualization we use 998 centers of mass, each of which is clearly assigned to one of the 66 areas. The opacity of each sphere is proportional to the weight assigned to it in the spatial map of the template.

Taken together, clustering can be improved by applying the non-negativity constraint for both correlation and MI, but when using correlation, also the surrogate data exhibit more of a cluster structure, which makes MI the better choice in this application.

Another way to improve clustering is to reduce noise and thereby mitigate the effect of inter-individual differences. In principle, it would be best to apply significance testing and keep only the significant dFC pairs/windows, however due to the large number of windows and pairs, this is computationally not feasible since we would have to generate thousands of surrogate tensors to reach a satisfying significance level, which would prove prohibitive in terms of computation time and storage space. Therefore, we just apply different thresholds and decompose the resulting binary tensors. We use different percentiles as thresholds,  $\theta = 0, 75, 80, 90, 91, 92, \dots, 99$ , and compare the results, using the silhouette value.

For correlation, the absolute value of correlation, and MI, the thresholds are determined to be the 98th, the 97th and the 98th percentile, respectively. For all measures,  $F = 3$  and  $K = 4$ . Correlation reaches a maximum mean silhouette value of 0.54 with a corresponding surrogate value of 0.13. The two non-negative measures produce better and equally good results: using the absolute value of correlation, we obtain silhouette values of 0.59 and 0.13 for real and surrogate data, respectively; for MI, the values are 0.54 and 0.08 (these templates are the ones shown in figure 2.4).

Apart from the clustering performance, we consider the reconstruction fits that quantify how well the extracted features describe the original tensor. At their best thresholds, MI and absolute value of correlation reach average fits of 0.39 and 0.48, respectively. The corresponding surrogate tensors yielded 0.09 and 0.34, confirming that while using the absolute value of correlation results in a good decomposition performance, the surrogate tensors constructed in this way also exhibit a lot of structure. Taken together, MI is the better choice.

### 2.3.3 Dynamic mean field model reproduces RSNs

In the next step, we use the templates shown in figure 2.4 to determine whether the dynamic mean field (DMF) model can produce data that exhibit spatio-temporal patterns similar to those found in the empirical data. We run  $S = 24$  simulations of the same length as the empirical data and use the parameters previously determined with the empirical data to the resulting tensors, i.e.  $F = 3$  and  $\theta = 98$ . The only free parameter is the

global coupling,  $G$  (figure 2.1A) which is a factor by which the connectivity matrix is multiplied and is hence related to the overall amount of activation in the system. We determine for each of the  $F \cdot S = 72$  extracted features the maximum correspondence to any of the  $K = 4$  empirical templates and use the mean across all features and all runs as a measure for the match between features (simulated) and templates (empirical) and thus of how well the model reproduces the empirical data. As before, we consider the difference to matches obtained using surrogates, i.e. of features extracted from tensors that are calculated from phase randomized simulated data to empirical templates.

We use the parameters obtained from the empirical data instead of running the same parameter selection procedure on the simulated data. This is because the simulations run on an average connectivity matrix and therefore, clustering across simulation runs does not make sense; the variance in FC structure is small across runs for a given value of  $G$ .

Figure 2.5A shows the overall overlap between simulated features and templates depending on the global coupling parameter. The 95% confidence intervals of real and surrogate data overlap for most values of  $G$ , except for  $G = 1.9, 2.5 - 2.7, 2.9$ , and  $3.1 - 3.7$ , the latter values matching the region in which the fit of the average FC is best. For these values, the simulated data can be shown to contain FC patterns that match empirical data to a higher degree than surrogate features. However, even the best match at  $G = 3.7$  is moderate. Figure 2.6 shows the simulated features next to the empirical template it best agrees with. We display the communities in vector form, and the ROIs are ordered as indicated in table A.1 in the appendix, i.e. symmetrically. Translated to the communities on the cortex, it is obvious that the main problem is that the features are not symmetrical, but rather, lateralized. For example, all features that match with the somato-motor template have their members in the right hemisphere. Surprisingly, the visual network, which has proven to be the most clearly pronounced one in the empirical data, is not found at all.

#### 2.3.4 Effective connectivity is crucial for modelling realistic communities

SC is derived by applying fiber-tracking algorithms to diffusion tensors obtained via dwMRI. We use the method developed by Gilson et al. (2016), briefly described in section 2.2.9, to obtain from this SC an effective connectivity (EC) for our dataset. This EC contains meaningful weights as




Figure 2.5: Overall template matches in comparison to correlation fits between FC matrices, results for simulations using SC (top) and EC (bottom). Black curves: template matches (left axes), thick black curve is real data, dashed line surrogate data. Shaded areas indicate the 95% confidence interval. Red curves: correlation of average FC matrices (right axes).

well as directionality information. Importantly, only the weights that are also present in the SC are tuned by the procedure, plus the weights on the secondary diagonal to account for homotopic connections that are not represented well by fiber tracking.

Figure 2.5B shows the overall overlap between simulated features and templates when EC is used (note that the scale on the y-axis is different). The non-overlapping region of the 95% confidence intervals of real and surrogate data covers a wider range of  $G$  (2.9 to 6.0), the maximum being at  $G = 4.2$ . Figure 2.7 shows the features extracted at this point next to the templates in vector form. Clearly, they are now symmetrical and have much more in common with the templates although the overlap in figure 2.5B may still



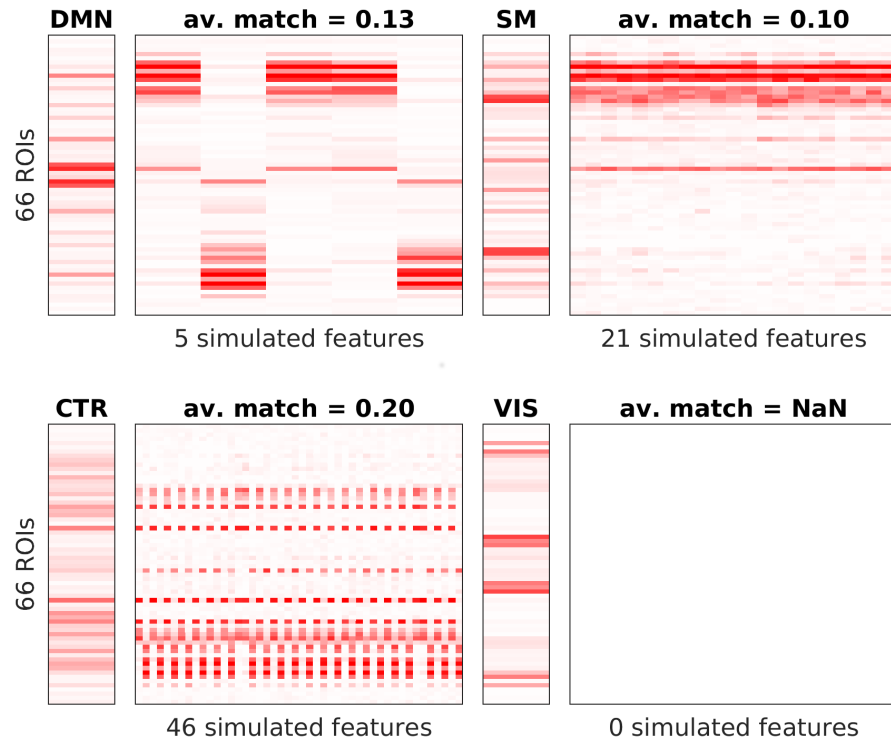


Figure 2.6: Spatial features extracted from 24 simulation runs using SC as underlying connectivity and the  $G$  with the best match (figure 2.5:  $G = 3.7$ ) with  $\theta = 98$ th percentile and  $F = 3$ , plotted next to the templates with which they exhibit the best match. Templates correspond to DMN - default mode network, SM - sensorimotor network, CTR - left and right control networks, VIS - visual network. For each set of features, the number of matched features and the average overlap for only this set is given.

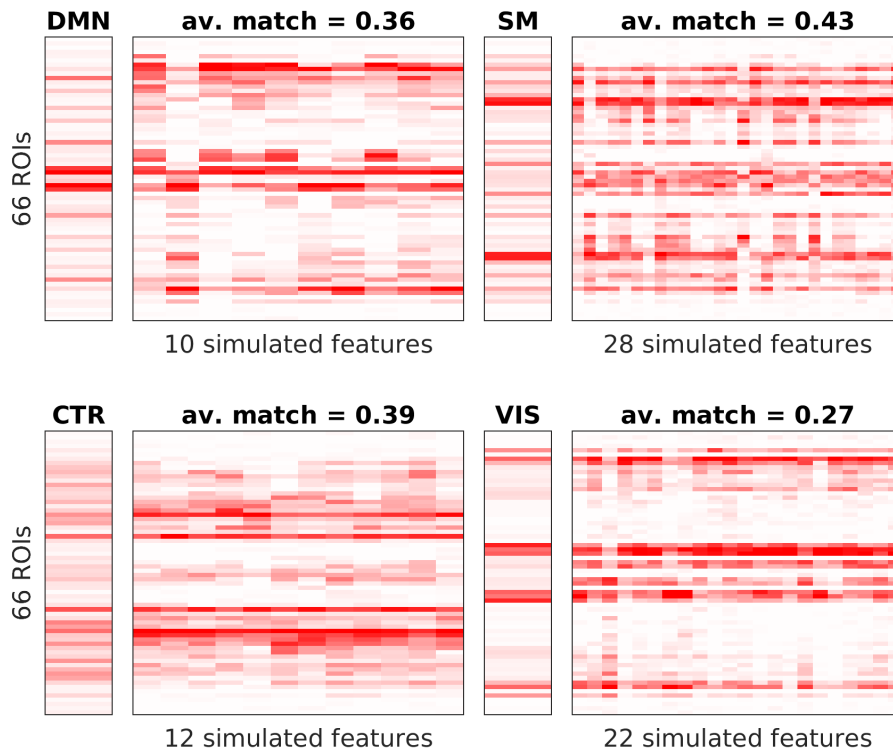


Figure 2.7: Spatial features extracted from 24 simulation runs using EC as underlying connectivity and the  $G$  with the best match (figure 2.5:  $G = 4.2$ ) with  $\theta = 98$ th percentile and  $F = 3$ , plotted next to the templates with which they exhibit the highest match. Templates correspond to DMN - default mode network, SM - sensorimotor network, CTR - left and right control networks, VIS - visual network. For each set of features, the number of matched features and the average overlap for only this set is given.

seem somewhat modest.

We explain this better match by considering panels B and C of figure 2.8. Several differences between SC and EC are immediately obvious, although we first note that both of them differ greatly from the FC shown in panel A. First of all, homotopic connections that are missing in SC are prominent in EC. This can explain the symmetry in the features obtained from EC-based simulations. Furthermore, the weights are more uniformly distributed in EC, making it appear more dense. Panel D of figure 2.8 reveals that the node degrees are largely equalized, such that most of them range in the

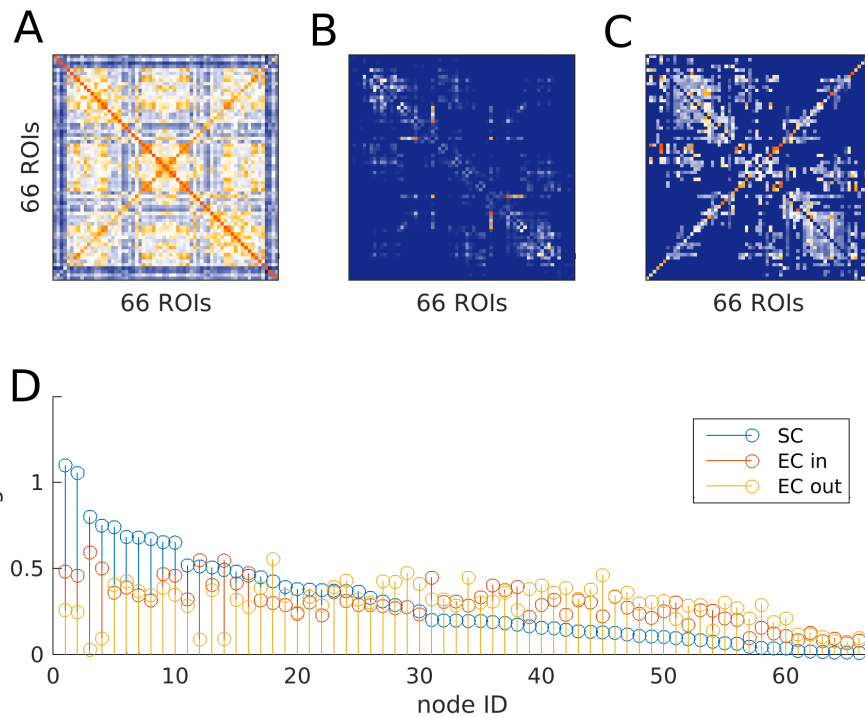


Figure 2.8: Connectivity matrices. **A** Functional connectivity matrix averaged over 24 subjects. Each entry is the correlation over the entire time course of a pair of ROIs. **B** Average structural connectivity matrix from the same 24 subjects. **C** Effective connectivity matrix derived from the FC, using the method in section 2.2.9. **D** Degrees of the nodes. For SC (blue), in and out degree are identical. Here, nodes are ordered according to their degree. EC nodes (in-degree: orange, out-degree: yellow) are shown in the same order as for SC.

middle of the distribution. This explains the different ranges of  $G$  that can be used in the simulations. For SC, the nodes with the largest degree cause the firing rates in the corresponding excitatory pools to rise to the point where the inhibitory pools cannot compensate for them any longer, and the asynchronous, low-activity regime of the system becomes unstable. The more uniform node degree distribution in the EC matrix allows for higher values of  $G$  and thus, the communities become more pronounced due to improved signal to noise-ratio. The EC matrix is also non-symmetric. This property likely further contributes to the stability of the simulations and more generally to the more realistic shape of extracted communities, both on the spatial and temporal level because it allows for a more diverse

propagation of the activity through the entire network.

## 2.4 Discussion

Our goal was to characterize spatio-temporal features of human resting state (RS) fMRI and to quantify to what extent a noise-driven stationary mean field model (Wong and Wang, 2006; Deco et al., 2014) can reproduce them. Using tensor decomposition (Cichocki et al., 2009), we identify four communities that generalize across subjects and resemble known RSNs (Fox et al., 2005; Beckmann et al., 2005; Mantini et al., 2007). We utilize temporal information by computing pair-wise dynamic FC (dFC) in sliding windows of 2 minutes to build our tensors. We compare three dFC measures: correlation, absolute value of correlation, and mutual information (MI). We determine the dFC measure, number of extracted features  $F$ , number of templates  $K$ , and binarization threshold  $\theta$ , that yield the best clustering performance, and take cluster centers as templates. We find that using a low number of features ( $F = 3$ ) and of clusters ( $K = 4$ ) combined with a high threshold ( $\theta = 98$ th percentile) applied to dFC calculated from MI works best (figures 2.3, 2.4).

We determine the range of global coupling for which the match of communities extracted from the model data to the templates is maximal (figure 2.5). We compare two underlying connectivities of the model: dwMRI derived structural connectivity (SC) and model-based effective connectivity (EC), which estimates directionality and weights of connections using time-shifted covariances (Gilson et al., 2016). EC produces more realistic FC patterns than SC alone.

### 2.4.1 Few data points are sufficient to recover RSNs

We find that applying a threshold and binarizing the tensors much improves the clustering performance. For MI, with  $\theta = 0$ , the maximum silhouette value is 0.45 for real data and 0.10 for surrogates, finding only three templates that do not match known functional networks. For the best threshold of  $\theta = 98$ , we find a silhouette value of 0.54 (surrogates: 0.09) and we succeed in extracting FC patterns that resemble RSNs (figures 2.3, 2.4).  $\theta = 98$  translates to using only the 2% biggest dFC values in the decomposition.

Tagliazucchi et al. (2012) transformed RS fMRI time courses into a point process, reducing the data by 94% and keeping only extreme events. Even so, the authors were able to recover RSNs, suggesting that “avalanching

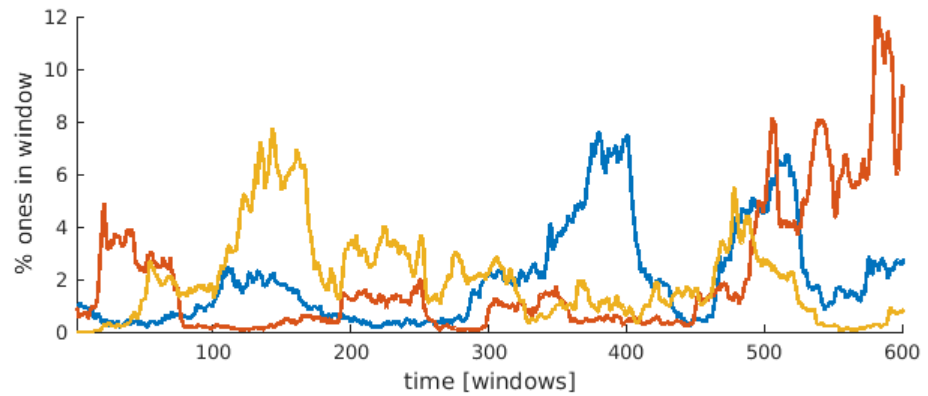


Figure 2.9: Time courses of the number of supra-threshold FC pairs per window for three subjects. The tensors were created using MI and  $\theta = 98$ th percentile.

events which involve short and long range cortical co-activations” explained their results. Here, we find that only the very largest dFC values are necessary to recover RSNs, suggesting that periods of highly structured FC are equivalent to periods of high variance in the BOLD signal. Indeed, figure 2.9 shows a strongly fluctuating time course of information content (as measured by supra-threshold MI pairs). We conclude that the peaks of these fluctuations represent periods of high modularity which are detected by the factorization algorithm. Modularity has been shown to dynamically fluctuate by Betzel et al. (2016).

Considering these large fluctuations, our findings are compatible with results by Mitra et al. (2014). The authors explain observed data in terms of waves of activity that propagate from regions acting as sources to others acting as sinks, involving the entire cortex. This results in stereotypical dFC patterns derived from average FC. Since so few data points are necessary to obtain our findings, and since they are concentrated around a very small number of peaks in each subject, it seems plausible to assume that in any given peak, *all* encountered communities are present. This concurs with the notion of a stereotypical global event. Furthermore, Messé et al. (2014) showed that a stationary model assuming fluctuations around a single fixed point, i.e., the average FC, predicts spatial structure of FC more accurately than more complex models, even when dFC is included. This is in parallel to what we show here: even a model that possesses only one attractor can produce time-varying FC patterns that are picked up by a decomposition algorithm.

### 2.4.2 Why we use tensor decomposition

In this study, we use a method that is relatively unknown in neuroscience, i.e. tensor decomposition (Cichocki, 2013; Leonardi and Van de Ville, 2013; Gauvin et al., 2014; Ponce-Alvarez et al., 2015). There are several reasons for our choice. First of all, with a resolution of 661 time frames and 66 ROIs, ICA cannot be readily applied. We would have to use temporal ICA (Calhoun et al., 2001b) instead of the more widely used spatial ICA (McKeown et al., 1998; Beckmann and Smith, 2004; Calhoun et al., 2009), making it difficult to compare to studies that identify communities (Beckmann et al., 2005; Damoiseaux et al., 2006; De Luca et al., 2006; Mantini et al., 2007; Van den Heuvel and Hulshoff Pol, 2010). However, the low spatial resolution was necessary in order to assess whether our specific model - the dynamic mean field model (Deco et al., 2014) - does indeed have the capacity to reproduce RSNs. Apart from that, investigating specifically long-range connections in a whole-brain approach has its own merit. In any case, tensor decomposition is nothing more but a generalization of PCA or SVD and is thus a fairly generic data analysis technique.

Our second reason is that this method allows us to use temporal fluctuations of FC explicitly, but without making any assumptions. Indeed, non-negativity is the only constraint we apply, and it is completely natural to do so when using a non-negative FC measure like MI. In other words, while ICA starts with the BOLD time courses, we use pair-wise FC, investigating time-evolving community structure explicitly. The result is that we get spatial and temporal features at the same time, without having to assume independence in either dimension.

Third of all, and importantly, this method allows us to extract features from single subjects/simulations, making it unnecessary to map group results back to the individual like with dual regression in group ICA. It is straightforward to match extracted features to a group template and estimate how well a subject agrees with the group average, which is summarized in the silhouette value used here to validate clustering results and tune parameters.

### 2.4.3 Correlation versus mutual information

Another perhaps slightly unusual choice is to favor MI over correlation as a measure of FC. It has, however, several advantages. First, it is a non-negative measure which allows us to further constrain the tensor decompo-

sition, leading to more reliable results. Second, spurious dFC structure in correlation-based tensors is reduced, as consistently shown by our results.

The reason for this difference has to lie in the spectral properties of the data, because they are preserved in the surrogate data. We found many time windows that contain outliers, i.e. a handful of time points with much higher activity than the rest. MI as computed here (Kraskov et al., 2004) is robust against these outliers, but correlation overestimates dFC in such windows. We hypothesize that these outliers are a result of global slow fluctuations in the signal that are preserved in the surrogates. MI as a non-parametric measure is robust against these outliers while correlation assumes normality and therefore is sensitive to them.

This observation calls into question the usefulness of Pearson correlation for investigating dFC despite its popularity (Allen et al., 2012; Hutchison et al., 2013; Hansen et al., 2014). Concerns have been addressed in some publications (Lindquist et al., 2014; Hindriks et al., 2015), and cross-validation with appropriate surrogate data is highly commendable.

#### 2.4.4 Using state-of-the-art connectivity matrices

In SC, each connection is symmetric and its weight is determined by the number of fibers detected by the tracking algorithm. However, it is known that many long-range connections are missed by these algorithms because of crossing fibers, notably in the region of the corpus callosum but also connections between frontal and occipital regions. Therefore, a lot of the interhemispheric connections between homotopic areas are missing. Furthermore, the results of fiber tracking do *not* allow direct inference of the weight of a connection. Lastly, fiber tracking results exhibit a high variability across subjects and sessions (Jones et al., 2013; Jeurissen et al., 2013).

On top of these shortcomings, we know that, *in vivo*, the asymmetry of the underlying connectivity shapes the dynamics of the system. To estimate weights and directionality, we need suitable observables and a dynamical model. We use the conceptual framework of EC (Friston, 1994), which describes the influence that one cortical area has over another. It depends, in principle, on synaptic plasticity, previous brain activity, and neuromodulation. Therefore, *in vivo*, EC changes dynamically depending on the experimental condition, individual differences, etc.

The method developed in Gilson et al. (2016) allows us to extract a single (group level) EC from our data that remains unchanged for all sim-

ulations. Note that we constrain the optimization to direct connections that are present in the SC, plus homotopic connections. Directionality and weights of connections are estimated using time-shifted, i.e., asymmetric, covariances as observables, as opposed to correlations, which allows us to also estimate the noise matrix, i.e. the amount of variability in each node. We apply a simple noise diffusion model, the type of model which has been found to reproduce RS FC quite well (Messé et al., 2014).

The use of EC leads to more realistic communities that are impossible to obtain with SC alone. One of the main benefits of using EC is that homotopic connections are strengthened, enabling realistic symmetric communities. Messé et al. (2014) found that just adding homotopic connections by setting them to a fixed value greatly improved predictive power. It would be interesting to see how model performances play out when using EC, because the shortcomings of SC are likely to impair modelling results very strongly. Perhaps, once this limitation is mitigated, we would favor a more complex model, as was the authors' initial hypothesis.

We emphasize that the communities themselves do not contribute in any way to the EC optimization procedure. They are extracted using temporal information in the shape of  $dFC(w)$  matrices that are obtained on time scales that are far greater than those used in the optimization procedure. Thus, the two methodologies make completely different use of dFC. Furthermore, the template extraction which drives the adjustment of parameter setting (choice of  $F$ ,  $K$ , dFC measure, and  $\theta$ ) is done on the single subject basis and is independent of any simulations.

### 2.4.5 Conclusion

We have shown that a dynamic mean field model with a single attractor that is explored through noise is sufficient to explain a lot of the spatio-temporal structure found in large-scale resting state fMRI. Noise-driven fluctuations around the average functional connectivity structure are shaped by the underlying connectivity and the simple dynamics of the model in such a way that over time, known functional networks are expressed. We emphasize that we added to previous findings according to which the average FC is reproduced by the model by demonstrating that dFC patterns occurring *over time* are contained in the simulated data. We achieved this by decomposing tensors, or “stacks” of time-dependent dFC matrices. We also showed that applying a high binarization threshold to these tensors, keeping only



the overall largest pair-wise dFC values, is necessary to obtain communities that generalize well across subjects.

Future studies should try to further characterize the temporal structure, for example determining how separate in time different networks are. This would enable us to pinpoint differences between experimental groups (different ages, gender, patient groups) and regarding tasks (e.g. attention, decision-making, motor) by describing and modelling how large-scale networks interact and how these characteristics are related to performance or clinical markers.



---

## Temporal dynamics of human resting state fMRI

It is well known that the pattern of pairwise correlations of BOLD signals is preserved during rest. These correlations are more generally referred to as functional connectivity (FC) because they are thought to reflect functional integration between brain regions. Recently, it has been established that FC is not static, but that it changes over time, however, the dynamics that determine of these changes remain elusive. In this study, we use a sliding windows approach to track global fluctuations in FC and BOLD variance. We find that both go hand in hand, suggesting that one cannot be interpreted without taking the other into account. Additionally, we analyze activation time courses of resting state networks (RSNs) obtained via tensor factorization, a method which makes very few assumptions about the underlying structure of the data. RSN time courses exhibit a large amount of co-activation which varies with global fluctuations in FC and BOLD. We find that many properties of the observed fluctuations in FC and BOLD, including their ranges and their correlations amongst each other, are predicted by the presence of long-term correlations. However, we also encounter interesting characteristics that are not explained in this way. In particular, we find that fluctuations in the BOLD signal and fluctuations of FC exhibit a non-linear relationship which suggests that the brain transitions between states of high synchronization and states of low synchronization in a non-trivial manner.

### 3.1 Introduction

In their seminal paper, Biswal et al. (1995) report that correlation patterns in motor cortex during rest are remarkably similar to those found during a finger tapping task. Following this discovery, a number of large-scale functional networks was unveiled using the experimental paradigm of “resting state” (RS) and the concept of functional connectivity (FC) (Lowe et al., 1998; Cordes et al., 2000; Kiviniemi et al., 2003; Fox et al., 2005; Beckmann et al., 2005; De Luca et al., 2006; Damoiseaux et al., 2006; Mantini et al., 2007; Smith et al., 2009; Yeo et al., 2011). Investigating these functional networks, termed “resting state networks”, enables us to probe integration (within networks) and segregation (between networks) in the brain.

Since the beginning of RS research, it was asked whether the observed fluctuations are “meaningful” or not, i.e. if they are the signature of ongoing computations. This question has been approached from three different angles. First, there are several brain regions that are more active during rest than during task, and their activity has been linked to mind-wandering and consciousness (Raichle et al., 2001; Gusnard et al., 2001; Mason et al., 2007). In this line, the observed fluctuations are interpreted as an expression of ongoing self-referential tasks. Second, it has been discovered that the statistical dependencies between brain regions’ activity persist at least partially during sleep (Larson-Prior et al., 2009) and anaesthesia (Vincent et al., 2007; Bettinardi et al., 2015), and that even the magnitude of fluctuations is maintained in states with different levels of consciousness (Horovitz et al., 2008; Hutchison et al., 2013). This suggests that the fluctuations are (also) a result of a “network effect”, i.e. random noise reverberating through a structured network, which nonetheless could be functionally relevant to the maintenance of synaptic connections. Third, the statistical properties of these fluctuations have been analyzed without assuming or rejecting any functionality, simply investigating whether observations can be explained by random noise or not. The focus has been mostly on the magnitude of the fluctuations in FC, investigating whether correlations are stationary in a strictly statistical sense, or not (Handwerker et al., 2012; Lindquist et al., 2014; Hindriks et al., 2015; Betzel et al., 2016; Laumann et al., 2016). Of course, these three approaches are by no means mutually exclusive, but rather different aspects of the same question.

Only recently, studies have begun to investigate the rich temporal dynamics of RS (Chang and Glover, 2010; Kiviniemi et al., 2011; Allen et al., 2012; Hutchison et al., 2012; Liu et al., 2013; Hutchison et al., 2013; Zalesky et

al., 2014). This led to the idea of the “chronnectome” (Calhoun et al., 2014) or “dynome” (Kopell et al., 2014), where the notion is put forward that the brain cannot be understood without considering temporal dynamics of interactions. From a modelling point of view, the observed variability is thought to constitute a “dynamic repertoire” of the brain. Analyzing this repertoire is crucial for understanding the brain in terms of general principles of information processing, pattern formation, integration across temporal and spatial scales, etc. Computational models are an indispensable tool in this endeavor (Ghosh et al., 2008; Honey et al., 2009; Deco and Jirsa, 2012; Hansen et al., 2014; Messé et al., 2014).

Fluctuations occur on many levels, for example, spatial configurations of functional networks (Kiviniemi et al., 2011), correlations between spatially fixed networks (Allen et al., 2012), network measures like modularity (Betzel et al., 2016), amount of co-activation in general (Tagliazucchi et al., 2012), and co-activation patterns (Liu et al., 2013). Precisely characterizing the ongoing dynamics will help connect them to task (Fox et al., 2006), development (Grady and Garrett, 2014), disorders (Fröhlich et al., 2010), and different states (Hudetz et al., 2015).

---

The temporal dynamics of spontaneous fluctuations present in BOLD time courses measured with fMRI are the focus of the present chapter. We start by defining three measures of global RS dynamics and apply them to scans of 24 healthy participants (the same data as used in the previous chapter). We use these measures to characterize the dynamics and compare our findings against those from surrogate data. These surrogates are constructed under the null hypothesis that fluctuations can be fully explained by the presence of long-term correlations, or in other words, by random fluctuations around the average FC. We then connect our results to the activation time courses of resting state networks, showing that they strongly co-activate during periods of increased global FC and strong BOLD fluctuations. Our findings suggest that the brain transitions in a non-trivial way between rather stereotypical states of high and low synchronization, and that these dynamics influence the time courses of RSNs.

## 3.2 Methods

### 3.2.1 Data and extraction of RSN time courses

We use the same data as in the previous chapter (section 2.2.1), i.e. 21-minute ( $T = 661$  frames,  $TR=2s$ ) RS scans from 24 healthy subjects aged 18 to 35 years. We apply non-negative tensor factorization and select the parameters as summarized in section 2.2.10. This results in decomposing each subject's tensor into  $F = 3$  rank-1 tensors, as illustrated in figure 2.2 in section 2.2.4. Previously, we focused only on the communities ( $a_f$  and  $b_f$  in the figure), whereas now, we work with the time courses that are associated with them ( $c_f$  in the figure). In order to match them to RSNs, we use the same clustering as before, resulting in four common RSNs being discovered: default mode network, somatomotor network, right and left control networks, and visual network (figure 2.4) in section 2.3.1 of chapter 2).

### 3.2.2 Measures of RS dynamics

We use three measures to quantify global dynamics of RS, illustrated in figure 3.1. Two measures are related to changes in functional connectivity (FC), and one to changes in the BOLD signal itself. For each of them, one time course is created for each subject  $s$  (i.e., we end up with three time courses for each subject) using sliding windows of width 120s, i.e. 60 frames, which are moved along the time course in increments of one frame ( $TR=2s$ ). This gives us a total of  $W = T - 60 = 601$  windows  $w$  per measure per subject.

The FC-related measures rely on pairwise correlations, i.e. for each window, the correlation is computed for each pair of  $N$  ROIs using the 60 time points falling inside this window. This results in an  $N \times N$  dynamic FC matrix,  $dFC(w)$  with  $w = 1, 2, \dots, W$  which is symmetrical and whose diagonal is filled with ones. The **FC strength** is just the average over all unique pairs, i.e. using the upper or lower triangle of the  $dFC(w)$  matrix, excluding the diagonal:

$$\text{FC str}_s(w) = \tanh \frac{1}{((N \cdot N) - N)/2} \sum_{i=1}^N \sum_{j=(i+1)}^N \tanh^{-1}(C_{ijw}) \quad (3.1)$$

$C_{ijw}$  is the entry in the  $dFC(w)$  matrix at pair  $ij$  and window  $w$ ,  $s$  is the subject index,  $\tanh^{-1}$  is the arcus tangens hyperbolicus, needed to z-transform the correlation values before averaging them, and likewise,  $\tanh$

is the tangens hyperbolicus, needed to back-transform the result into a correlation value.

We evaluate how similar each  $dFC(w)$  matrix is to the average FC (avFC) matrix, i.e. the FC matrix obtained from correlating the full time courses of the ROIs. This is done by simply correlating the upper or lower triangle of the matrices with each other after flattening them into a vector.

$$\text{sim to avFC}_s(w) = \text{corr} \left\{ \text{vec} \left[ \mathbf{U}(\tanh(dFC(w))) \right], \text{vec} \left[ \mathbf{U}(\tanh(avFC)) \right] \right\} \quad (3.2)$$

$\text{vec}(\cdot)$  stands for the vectorization of a matrix,  $\mathbf{U}$  is the upper triangle of a matrix, and  $\text{corr}$  stands for Pearson correlation.

Additionally, we use a measure that tracks changes in the BOLD signal itself, namely of its variance. For this, a normalization step is necessary because the variability of variances across subjects does not carry any meaning, only the relative differences between ROIs within a subject do. Therefore, for each subject, we normalize such that the variance of its most variable ROI is 1. We then compute the variance of the signal  $\mathbf{x}_i$  for each ROI  $i$  inside each window ranging from  $t$  to  $t + 60$  (in frames) and average over ROIs, resulting in a measure for the mean variance of the signal, which we refer to as “BOLD var”.

$$\text{BOLD var}_s(w) = \frac{1}{N} \sum_{i=1}^N \text{var}_{t:t+60}(\mathbf{x}_i) \quad (3.3)$$

Note that this is not the variance of the mean signal.

### 3.2.3 Surrogates

In the previous chapter, we used surrogates that preserve the Fourier spectrum of the data, but destroy any long-term correlations. This corresponded to the null hypothesis that any cluster structure in extracted communities was the result of just random fluctuations and not of the actual correlation structure of the data. Here, we have a different null hypothesis, which corresponds to the assumption that observed dynamics are a result of the long-term correlations, but that no additional structure is present. In parallel to what is described in section 2.2.6 of chapter 2 (equation 2.6), we use

the Fast Fourier Transform (FFT) to obtain a representation of the signal in frequency space and then add random phases to each bin.

In particular, we have  $x_1, x_2, \dots, x_T$  and  $y_1, y_2, \dots, y_T$ , where  $\mathbf{x}$  and  $\mathbf{y}$  are two ROI's signals in time whose long-term correlation we want to preserve. After applying the FFT and obtaining signals  $X_1, X_2, \dots, X_N$  and  $Y_1, Y_2, \dots, Y_N$  in frequency space, a random phase vector  $\varphi_{r,1}, \varphi_{r,2}, \dots, \varphi_{r,N}$ , values of which are drawn from a uniform distribution in the interval  $[-\pi, \pi]$ , is used to randomize phases in each frequency bin. The long-term correlations are preserved by adding the same phases to both signals before transforming back.

### 3.2.4 Statistical test for histograms

We frequently compute distributions using histograms. In order to test whether two histograms are significantly different, we use the G-test, which is similar to the  $\chi^2$ -test. The G-test is generally applicable for contingency tables, i.e. the bins used in a histogram in any number of dimensions are interpreted as categories  $i$ . The test statistic  $G$  is computed as:

$$G = 2 \cdot \sum_i (O_i \ln \frac{O_i}{E_i}) \quad (3.4)$$

$O_i$  is the number of observed samples in category  $i$ ,  $E_i$  is the expected number of samples in that category, derived from a model distribution or surrogates. That is, the total number of samples across all categories has to be the same in the observed and expected case, and the category boundaries have to be identical. The sum is over all non-empty categories. The factor 2 leads to the distribution of  $G$  approximately fitting the  $\chi^2$ -distribution for large enough samples. Then, the p-value is computed as

$$p = 1 - \chi^2(G, \text{df}) \quad (3.5)$$

Here,  $\chi^2(G, \text{df})$  is the value of the cumulative  $\chi^2$ -function with  $\text{df}$  degrees of freedom, at point  $G$ .



### 3.3 Results

#### 3.3.1 Three measures to track resting state dynamics

We start with the simple question whether spontaneous fluctuations measured during resting state (RS) are random or not. We use 21 minute long RS scans from  $S = 24$  healthy participants. We explore this temporal structure by looking at its changes in short (2 minutes, 60 frames) overlapping (shifted by 1 frame) time windows (figure 3.1, top). Since we are interested in global changes and want to compare to surrogate data later, we do not take the approach to compare ROI by ROI or pair by pair. Instead, we use three global measures of RS fluctuations, all of which are illustrated in figure 3.1. This results in time courses of the measures, one for each measure and each subject (figure 3.2). The measures are:

1. Average BOLD variance (BOLD var): Inside each window, the variance over time is computed for each region and averaged over all regions (figure 3.1 bottom left). Note that this is not the same as taking the variance of the mean signal since in this case, strong fluctuations in separate regions average each other out if they are uncorrelated. This does not happen in our measure.
2. FC strength: Inside each window, the correlation between all pairs of regions is computed, just like when assessing the average FC (avFC) by computing correlations over the entire time course (figure 3.1 bottom right). This results in a time-dependent dynamic FC matrix,  $dFC(w)$  with  $w = 1, 2, \dots, W$ ,  $W$  being the number of time windows. The FC strength is just the average over all correlation values in the upper (or lower) triangle part of this matrix, ignoring the diagonal. It can be interpreted as a proxy for overall “synchronization” in the system, although we do not explicitly examine oscillations.
3. Similarity to avFC (sim to avFC): The  $dFC(w)$  matrix is correlated with the avFC matrix (again using only the upper or lower triangle; figure 3.1 bottom right).

In figure 3.1, top, we can clearly discern the “bands” of orange and blue showing that almost all regions are participating in large fluctuations of the BOLD signal. This observation raises the question whether and how much these fluctuations in the BOLD activity contribute to fluctuations in FC.

In the previous chapter, we hypothesized that the observed fluctuations in suprathreshold FC values are related to the large BOLD fluctuations described by Tagliazucchi et al. (2012). In order to pursue this, we include measures for both, BOLD variance and FC variance, in this study. Figure 3.2 shows an example of the time courses of all three measures for one subject. Strong infraslow ( $< 0.01$  Hz) fluctuations are evident in all of them.

### 3.3.2 Specifying the null hypothesis

In order to determine whether the observed fluctuations are random or not, we need to define very carefully what we mean by “random”, i.e. what our null model is. For example, if one creates two time series by drawing random numbers from a bivariate Gaussian distribution with a fixed covariance, the degree to which they correlate fluctuates over time. Even if the two variables are independent, there will be time windows in which they reach correlations that would suggest otherwise. In other words, the true covariances between the time series explain a lot of the variance in correlations over time. In our BOLD data, the “true” correlation is estimated by taking into account the entire time course.

However, BOLD signals cannot be modelled by Gaussian processes due to their autoregressive properties. These properties result in autocorrelations that are sustained over several seconds, and in a particular power spectrum. Clearly, our null model must include these properties because they contribute to the statistics of the fluctuations. We use a method that randomizes the phases while preserving the pairwise correlations (Handwerker et al., 2012; Hindriks et al., 2015). This works by taking the Fourier transform and adding to each frequency bin a random phase; the correlations are preserved by using the same phase for both signals. This results in the avFC being the same, as well as the autocorrelations and the spectra, but the specific alignment of the fluctuations across regions is destroyed. Using this kind of surrogates corresponds to testing the null hypothesis that observed properties can be fully explained by the presence of long-term correlations. Importantly, the model we used in the previous chapter also possesses only one attractor, the avFC, and therefore, we have already shown that spatial patterns can be explained by this null hypothesis.

Figure 3.3 shows the same signal as we used before together with a phase-randomized version - the time courses are strikingly similar and the presence of the “bands” mentioned above are preserved, although their timing is of course different. This goes to illustrate just how much of the interesting

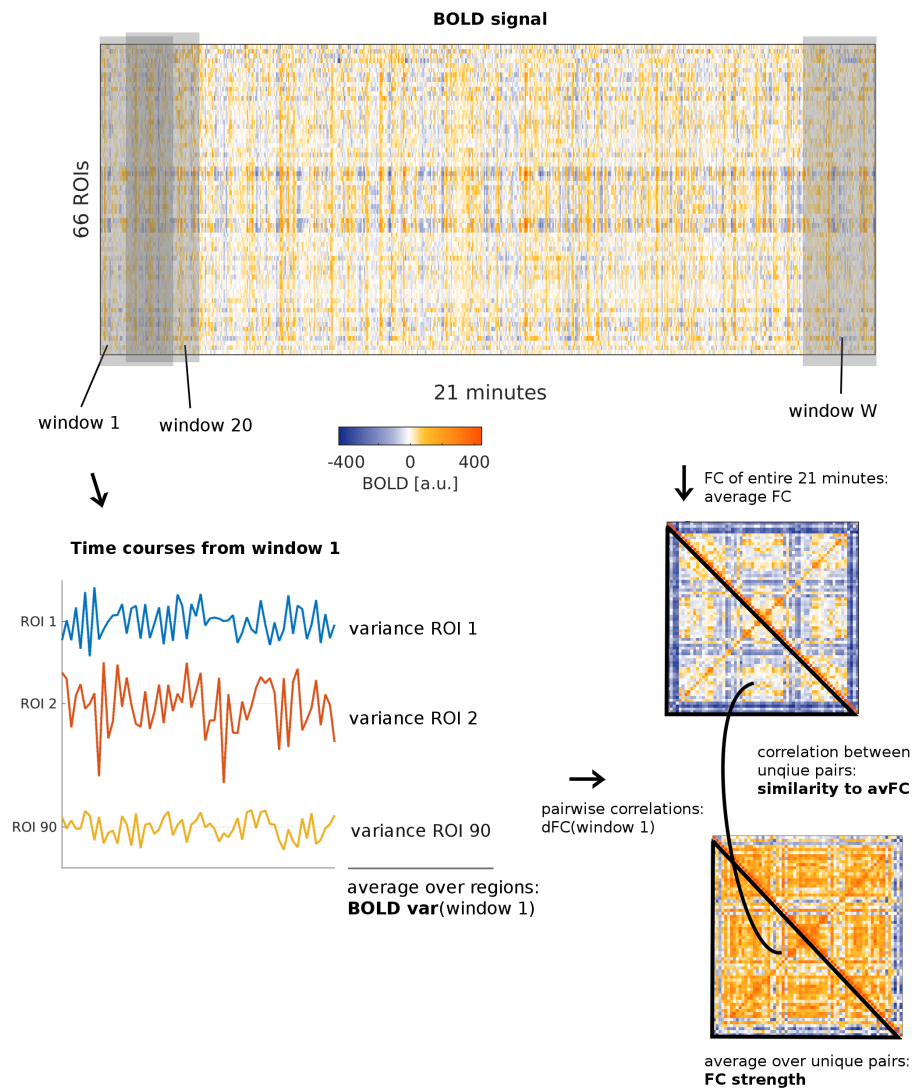


Figure 3.1: Illustration of three measures of RS dynamics. Top: We use sliding windows of width 120s (60 frames) and slide by 1 frame. Bottom left: Time courses in one window. The variance of each ROI is computed. Averaging over all ROIs yields the variance of the BOLD signal (“BOLD var”) for this window. Bottom right: The average FC (avFC) is obtained by correlating the full time courses of all pairs of ROIs (top). The same method is applied for each window, yielding the dynamic FC matrix,  $dFC(w)$  (bottom). Averaging over all “unique” pairs yields the overall FC strength for this window, and correlating the unique pairs of the  $dFC$  with those of the avFC yields the similarity to avFC (“sim to avFC”).

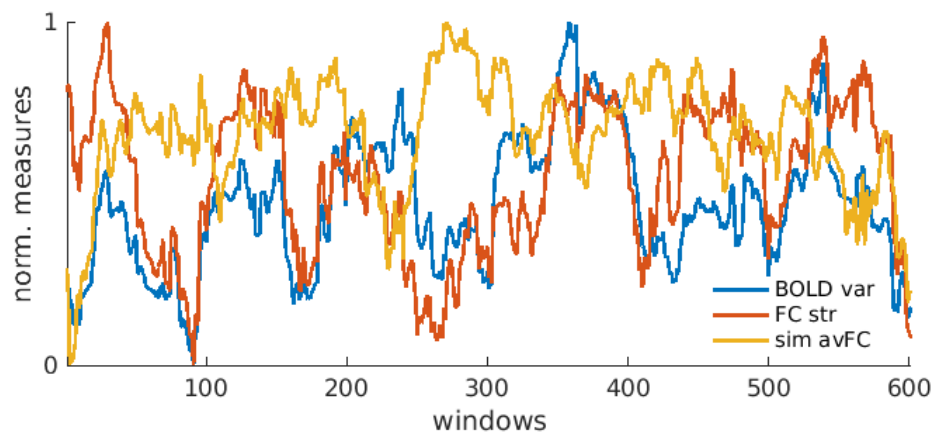


Figure 3.2: Time courses of three measures BOLD variance, FC strength and similarity to average FC for the same subject as in figure 3.1. The measures are normalized to lie between 0 and 1 for illustration purposes.

properties that we find in RS fMRI time series are a result of the general properties of the BOLD signal (Fourier spectrum, autocorrelation) which in turn result from the properties of the hemodynamic function, and can be explained by the existence of long-term correlations.

### 3.3.3 Fluctuations are explained by surrogates

Figure 3.4 shows time courses of the three measures taken from the same subject as used before, together with one phase randomized version. It is evident that there are considerable fluctuations in both cases, and that the surrogates exhibit a range that is comparable to that of the real data. We compute 100 surrogate data sets for each subject and count how many times the surrogates reach values that are smaller than the minimum or bigger than the maximum of the real data, respectively. Figure 3.5 shows the results for each subject.

The first thing we notice, independent of the surrogates, is that there is a large inter-subject variability in where the ranges lie and how broad they are. Likewise, while for some subjects all dots are within the range of the bar, for others, there are many surrogates that reach higher and lower values than the original data. A general observation is that BOLD variance seems to be the measure that is best preserved in the surrogates, with a total of 403 (16.8%) out of 2400 surrogates reaching a higher and 165 (6.9%) reaching a lower value than the original time courses they are derived from. For

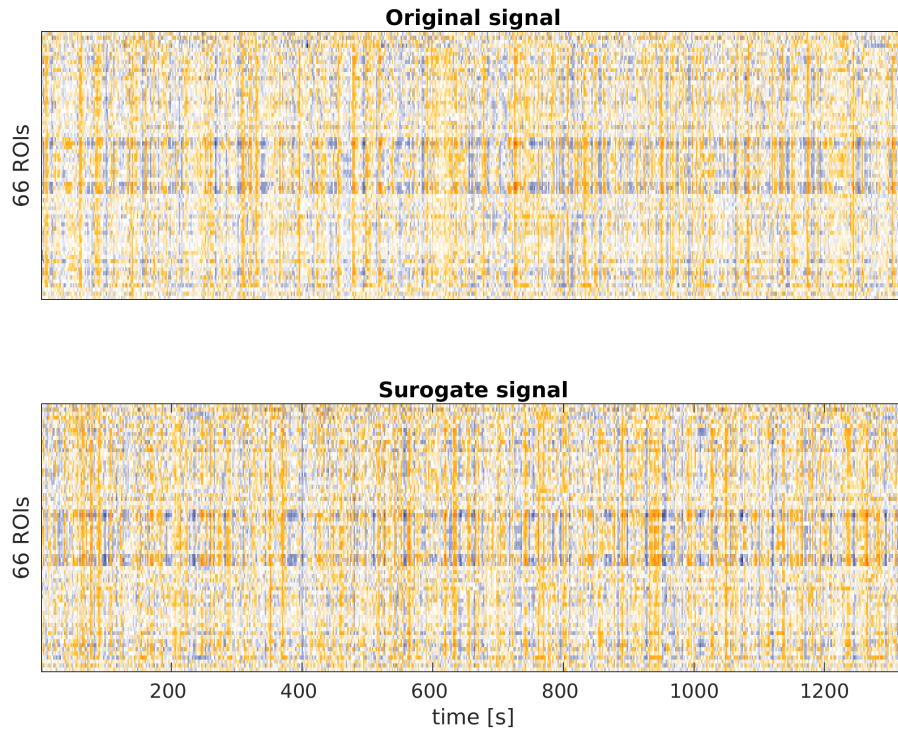


Figure 3.3: Original data of one example subject together with a phase randomized version. The obvious similarity is due to the preservation of spectral properties and long-term correlations.

FC strength, 21.8% of the surrogates have a higher maximum and 52.4% a lower minimum; for similarity to avFC, these numbers are 81.0% and 11.6%. Taken together, while on a single subject level, sometimes the surrogates seem to fail to reproduce the ranges of the original time courses, on a group level, the evidence is not sufficient to reject our null hypothesis. This is in line with findings by Handwerker et al. (2012); Hindriks et al. (2015) who used the same surrogates on different measures for RS dynamics.

### 3.3.4 Measures are closely related

Furthermore, we observe that the three measures seem to be closely related, especially BOLD variance and FC strength (figure 3.2). We compute the correlations between all three possible pairs (BOLD var-FC strength, BOLD var-sim to avFC, FC strength-sim to avFC) for each subject (figure 3.6).




Figure 3.4: Illustration of the effect of phase randomization on the three measures of global RS dynamics. The time courses computed from an example subject are shown in blue, together with one phase shuffled version derived from the same subject. Top: BOLD variance, middle: FC strength, bottom: similarity to average FC. While specifics are clearly different, the ranges are the same.

The correlations are highly significant for all combinations and all subjects, with five exceptions (empty circles). Again, there is a large variability across the subjects, but the correlation between BOLD variance and FC strength is the most reliable (average over subjects: 0.81). Correlations of these measures with the similarity to avFC is high in many, but not all subjects, suggesting a more complex relationship.

We repeat the analysis for surrogate data and find that also these correlations are explained by the long-term correlation structure, with many surrogates displaying a higher correlation than the original data. In fact, in some cases, *all* 100 surrogates do so, suggesting that the correlations are actually *lower* than expected. However, since this is again not consistent over subjects, we do not reject the null hypothesis.

However simple the relationship between BOLD variance and FC strength may be, it confirms that changes in BOLD activity and changes in FC contribute to each other. Fluctuations observed in the BOLD signal (fig-

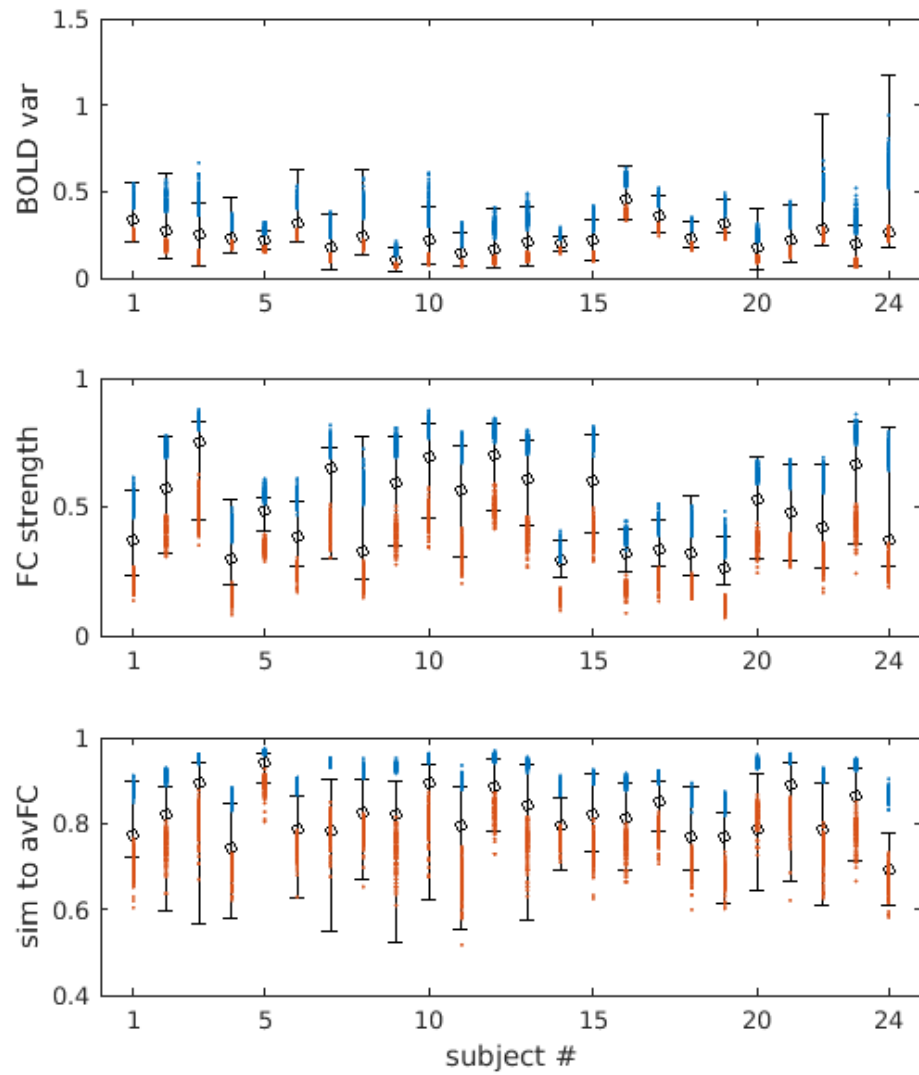


Figure 3.5: Range of fluctuations for real data (black bars, marking minimum and maximum of the time course) and surrogate data (colored dots). The empty circles mark the medians of the real data. Each dot lying above the black bar stands for a surrogate time course reaching a higher value than the original time course from which it was derived; each orange dot lying below its bar stands for a time course reaching a lower value. Top: BOLD variance, middle: FC strength, bottom: similarity to average FC.




Figure 3.6: Correlations between all combinations of measures for all subjects.

ure 3.3) go hand in hand with fluctuations in FC. In general, the system is most synchronized, and closest to its highly structured average FC, when fluctuations in the BOLD signal are strongest. This means that these fluctuations increase the signal to noise-ratio, rather than acting as perturbations, and emphasizes that changes in FC cannot be interpreted without taking changes in BOLD itself into account.

Our observation explains why similar observations have been reported by so many studies using a diversity of analysis approaches which include co-activations as well as correlations (Chang and Glover, 2010; Kiviniemi et al., 2011; Allen et al., 2012; Hutchison et al., 2012; Liu et al., 2013; Zalesky et al., 2014; Betzel et al., 2016). Although our finding may seem “trivial”, it is not explicitly taken into account in the cited studies.

### 3.3.5 Significant differences between data and surrogates are found when using full distributions of measures

So far, we have confirmed that neither the presence of large fluctuations in BOLD variance, FC strength and similarity of pairwise FC structure to the average structure, nor the correlation of these measures amongst each other ought to be a surprise. They alone are not evidence that would allow us to reject the null hypothesis that our observations can be explained by the noisy fluctuation of the average FC structure. However, not only the



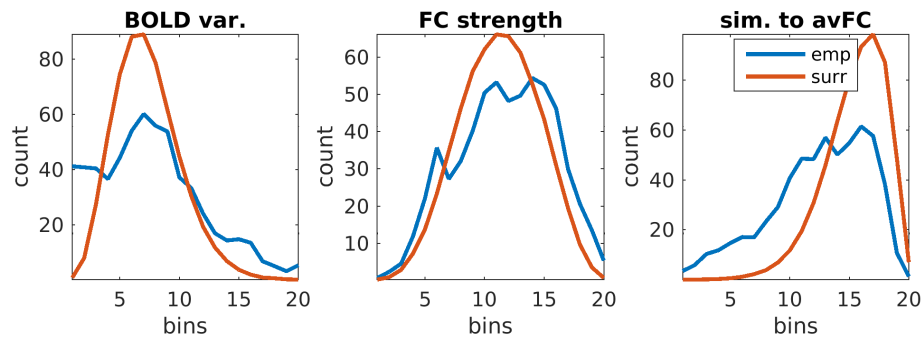


Figure 3.7: Distributions of all three measures averaged across all subjects, using 20 bins. The bins do not represent the same values for all subjects, since their ranges are different (see figure 3.4). Surrogate histograms use the same bins as their respective subject.

range of fluctuations is important, but also the trajectory upon which the system evolves from troughs to peaks. Therefore, our next step is to take into account the entire distribution of each measure. For each measure and each subject, we have  $W = 601$  samples, which we distribute into 20 equally spaced bins. We apply a G-test, using the average histogram of 100 surrogates as a prediction against which the counts from the real data are tested. We find that the distributions of all three measures of all 24 subjects are different from those predicted by the surrogates on a level of  $\alpha = 0.05$ , Bonferroni-corrected for  $S$  comparisons. This means that the time courses exhibit dynamical properties that cannot be explained by the average correlation structure. Figure 3.7 shows the distributions averaged across subjects. Also on the group level, the distributions are significantly different.

Thus, we have established that the shapes of the distributions differ from the predictions made by the null hypothesis. Clearly, the dynamics are not explained by the avFC alone even though the size of the fluctuations and their relationships to each other are.

### 3.3.6 Extracting RSN time courses

We have introduced three new measures to evaluate dynamics of RS FC and relate them to BOLD activity changes. In order to show that these measures are relevant to existing research, we relate them to dynamics of resting state networks (RSNs). We use tensor factorization to extract recur-

ring FC patterns and associated time courses, as explained in chapter 2 and illustrated in figure 2.2. We select the number of features,  $F$ , the number of clusters,  $K$ , and the binarization threshold,  $\theta$ , by considering how well the communities (spatial features) cluster, i.e. generalize across subjects. Each cluster is identified with an RSN known from the literature: default mode network (DMN), somatomotor network (SMN), right and left control networks (CTR), and visual network (VIS; figure 2.4). While in the previous chapter, we focussed only on the communities, we now analyze their associated time courses,  $c_f$  in the illustration in figure 2.2.

Figure 3.8 shows two example subjects' time courses. Note that not each of the four RSNs occurs in each subject. This is due to the procedure used to select the parameters of the algorithm, which results in  $F = 3$  communities being extracted from each subject, but  $K = 4$  communities being detected on the group level. Furthermore, the activations are zero in some windows because we use only the highest values of pairwise FC in the factorization (see section 2.2.5). This leads to some windows possessing no or almost no supra-threshold pairwise FC (see figure 2.9 in chapter 2 for an example).

We compute the correlation between the sum of the activations recorded in the RSN time courses and our three measures from above. The correlations are significant for all subjects and all measures except for two instances for similarity to avFC. Thus, RSNs are most strongly activated during windows of high BOLD variance (mean correlation with activation: 0.81), high FC strength (0.79), and strong similarity to avFC (0.49). Again, there is a large variability across subjects, with one subject even displaying a significant *negative* correlation between RSN activations and similarity to avFC. Still, the group result is quite strong. Since we threshold the tensors, so that only the highest FC values enter into the extraction of FC patterns and the time courses analyzed here, it is perhaps not surprising to find a strong correlation with FC strength. However, since we consider single pairs and windows when thresholding, the fact that these values cluster in certain windows - i.e. the windows that have the highest FC strength - is not an *a priori* assumption. Although we have established that FC strength and BOLD variance are highly correlated, it should be duly noted that the variance of the BOLD signal itself does not contribute in any way to the factorization. Thus, our finding establishes an interesting perspective on the data where RSNs are most strongly expressed during periods of high BOLD variance.

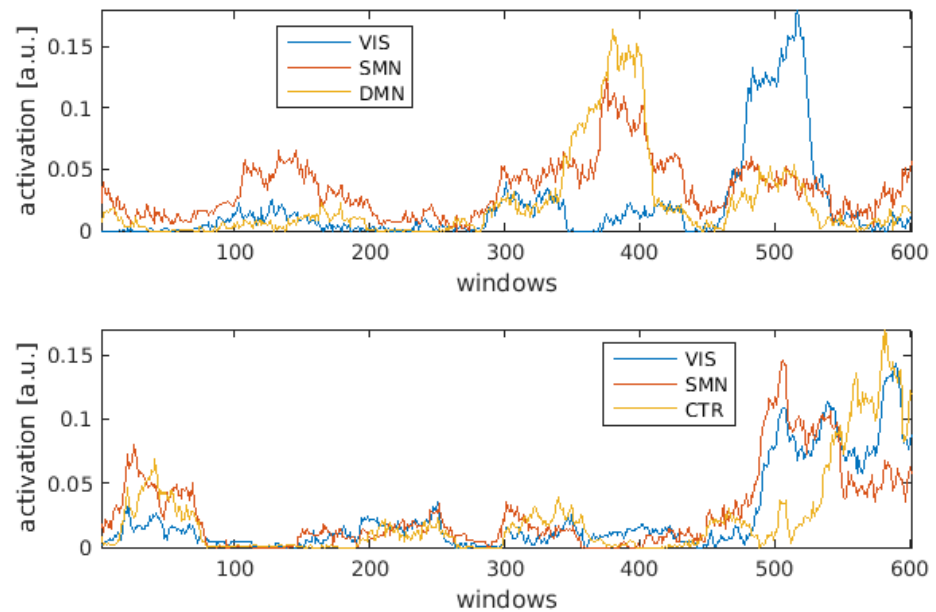


Figure 3.8: Time courses associated to communities identified as the indicated RSNs from two subjects. Thresholding the tensors before decomposing them leads to the activations being 0 in some windows because the baseline is effectively removed. DMN: Default mode network, SMN: somatomotor network, CTR: right and left control networks, VIS: visual network

### 3.3.7 Co-activation of RSNs depends on dynamical properties

From the time courses (figure 3.8), it is clear that the RSNs overlap in time to a certain degree. In order to quantify the degree of co-activation, we consider the “contributions” of each time course to the overall activation, which is just the ratio between an activation and the sum of all activations in this window. Figure 3.9 (left), shows these contributions for the activation time courses in figure 3.8 (top). If this value is close to 1, it means that this network dominates. If the value is close to  $1/F$ , i.e. one third in this case, it means that there is a large degree of co-activation; and if the value is very small, it means that the network is not active at all.

The avFC is a summary of all FC patterns that occur over time. Therefore, if RSNs are more strongly activated in windows with high similarity to FC, we would expect that during those peaks, RSNs co-activate more strongly

than away from the peaks. We use the maximum contribution in each window to obtain a global index of co-activation. Indeed, we find that this index is negatively correlated with our three measures. For BOLD variance, FC strength, and similarity to avFC, 23, 24 and 20 out of 24 subject exhibit a significant correlation, respectively. The mean correlations, averaging over the significant values, are - 0.41, - 0.47, and - 0.48, respectively. Thus, there is a tendency of the RSNs to co-activate more strongly during windows with strong BOLD fluctuations, high FC strength and high similarity to avFC, while dominance of a single RSN tends to occur during windows when these measures are lower.

Figure 3.9 (right) shows the contribution values for single RSNs (DMN, somatomotor, control, visual). The distributions show a peak around  $1/3$ , a value which clearly indicates co-activation. There are also many zeros, indicating windows in which the network was not active (windows in which all three were zero were excluded from the analysis since “contribution” cannot be defined in this case), but the number of values falling in a higher range does not counter-balance this peak, suggesting that there is a bias towards RSNs being co-activated or not active at all and away from single RSNs being dominant. There seems to be a difference between the networks, with the DMN showing the strongest tendency towards co-activation, which is in line with its role as “hub” (Hagmann et al., 2008).

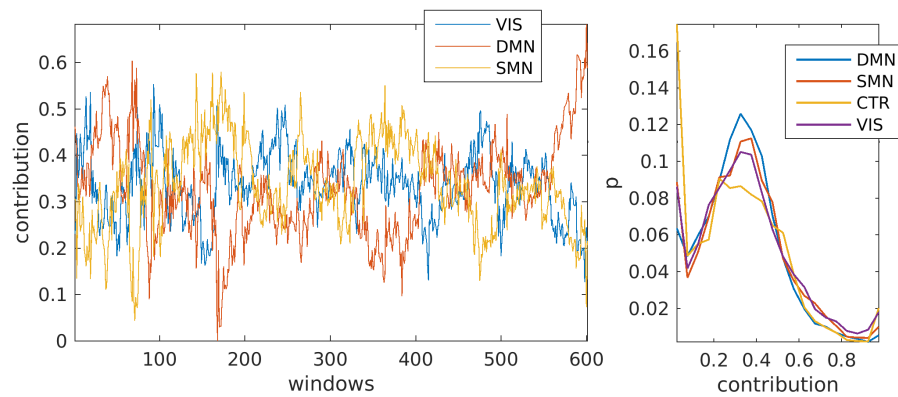


Figure 3.9: Co-activation analysis of RSNs extracted via tensor factorization. Left: example time courses of “contributions” for the activation time courses shown in figure 3.8 (top). The sum over all values at each time point is 1. Right: Histogram of contribution values for each RSN, pooled across all subjects. DMN: Default mode network, SMN: somatomotor network, CTR: right and left control networks, VIS: visual network

### 3.3.8 Evidence for non-linear relationship between BOLD variance and FC

We have established that the dynamics of RS, tracked via our three measures BOLD variance, FC strength, and similarity to avFC, cannot be fully explained by surrogate data which exhibit only a noisy fluctuation of the avFC. However, we do not know what exactly the difference is. As mentioned above, the correlation between FC strength and BOLD variance is very high (0.81 across all subjects), but the relationship of both measures to the similarity to avFC is less clear. Likewise, RSNs are most strongly activated during windows of high FC strength and high BOLD variance, but the relationship to similarity to avFC is not as strong, the correlation only an average of 0.49 across subjects, with two of them exhibiting no significant correlation, and one even a slightly negative one.

Following these discrepancies, we consider the joint distribution of BOLD variance and similarity to avFC. We use equally spaced bins in order to overcome the problems posed by the large variability across subjects as before when constructing the univariate distributions (figure 3.7). In other words, we are not interested in the actual values of these measures, but only in the temporal fluctuations that they exhibit. Figure 3.10 (left) shows the joint distribution of the two measures, averaged across all subjects. Although we

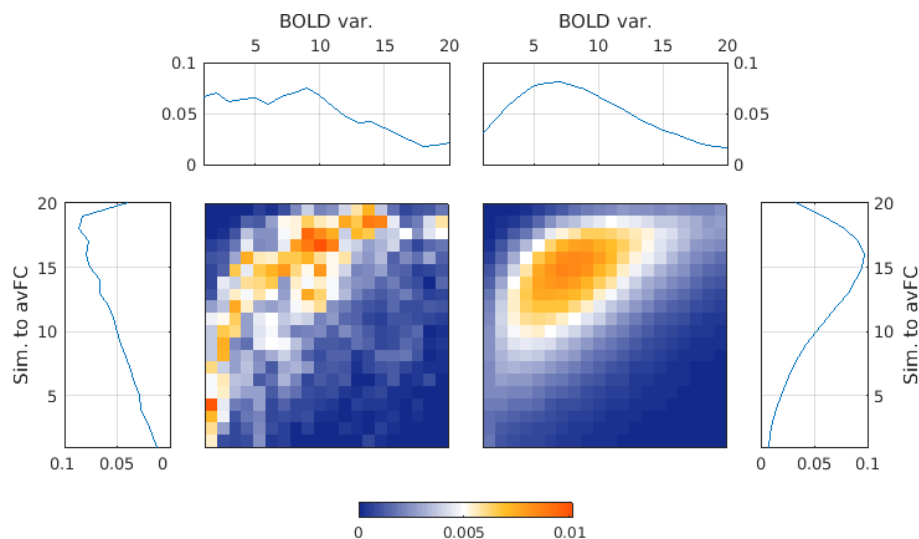


Figure 3.10: Joint distribution of BOLD variance (x-axes) and similarity to avFC (y-axes) for empirical data (left) and surrogates (right). The marginal distributions are shown at the top and sides.

assessed the relationship between the two measures using correlation, it is now evident that we are looking at a non-linear relationship, which explains why the correlation values were not as high, averaging 0.53 (figure 3.6).

The shape of the joint distribution hints at BOLD variance and similarity to avFC being related over time in a non-trivial way. This is confirmed by plotting the same distribution using the surrogates (figure 3.10, right). The marginal distributions differ significantly between empirical and surrogate data ( $\alpha < 0.001$ ), and it is obvious that the joint distributions are very different.

We conclude that BOLD variance and fluctuations in FC structure are tied together in a way that is specific to the temporal dynamics of resting state and not just an effect of the average spatial FC structure. Specifically, the non-linear structure in the left panel of figure 3.10 shows that during periods where BOLD variance increases (going from left to right on the x-axis), the “expression” of the FC structure - i.e. the degree to which the  $dFC(w)$  matrix resembles the avFC - increases more quickly than expected. This suggests the presence of a mechanism that actively “boosts” this structure.

Furthermore, there are many windows located at the bottom left of the joint distribution of the empirical data that do not occur in the surrogates at all. The presence of these windows is in line with the distribution in figure 3.7 (right) being shifted towards smaller values. Likewise, the distribution is slightly shifted towards higher values as well, as is more easily visible in the marginal distributions of the similarity to avFC (to the left and right). This hints at the system exhibiting longer dwell times in states of high and low values of the two measures than in states of intermediate values, suggesting some degree of bistability in the system.

### 3.4 Discussion

With the help of three measures of resting state (RS) dynamics and surrogate data constructed under the appropriate null hypothesis, we have shown that many statistical properties of spontaneous fluctuations in fMRI BOLD are explained by the presence of long-term correlations. Both the range of fluctuations in BOLD variance and FC, and the correlations between the two are reproduced by the surrogates corresponding to the null hypothesis of a fluctuating average FC (avFC). However, when considering the full distributions of the three measures we establish, it is revealed that dynamics of the empirical data exhibit additional structure that is not explained by the surrogates, and the null hypothesis is rejected.

We also find that the dynamic behavior of resting state networks (RSNs) is intimately related to our measures in such a way that they activate most strongly during periods of high BOLD variance, high FC strength (synchronization), and strong expression of the average FC structure (high similarity to avFC). They also tend to co-activate more during these windows.

By considering the joint distribution of two of our measures, BOLD variance and FC strength, we show a possible origin of the additional structure present in the empirical data. Our results suggest that there is a mechanism that leads to the average FC structure being expressed more strongly and more quickly than predicted by the surrogates.

#### 3.4.1 Are fluctuations during rest meaningful?

On the one hand, it has been asserted that ongoing spontaneous activity reflects self-referential tasks or mind-wandering (Raichle et al., 2001; Greicius et al., 2003; Fransson, 2006; Mason et al., 2007), on the other, it has been contested that the observed fluctuations are an actual sign of

non-stationarities, and proposed that observed “states” are explained by a fluctuating average FC structure (Handwerker et al., 2012; Lindquist et al., 2014; Hindriks et al., 2015; Laumann et al., 2016). If the latter case is true, it is unlikely that observed fluctuations reflect any kind of switching between states. We show that the relationship between fluctuations in activity and fluctuations in FC is non-linear, and that RSN temporal dynamics are tied to these fluctuations. This lends support to the hypothesis that “states” are not fully determined by the average FC, but that excursions from it occur at certain points in time, even though these excursions may be relatively scarce.

However, it should be noted that there is a large inter-subject variability: for some subjects and some measures, the surrogates do clearly not reproduce the data (figure 3.4), in contradiction to what we find on the group level. Determining the sources of this variability should be a goal in future studies to make sure that findings accurately represent single subject data. Here we tackled this problem by reporting results on both the single subject and the group level, and making sure that differences existed on both. Even so, there is no guarantee that these differences have the same direction in all subjects. The three measures introduced here are an attempt to find a low-dimensional representation of the temporal evolution of RS. Further research is necessary to further explore and refine them.

### 3.4.2 Relationship between changes in activity and in FC

We have shown very clearly that there is a strong positive correlation between changes in BOLD activity and changes in FC, in parallel to findings on a finer spatial and temporal scale (Chawla et al., 1999). While it may be clear that this relationship is “trivial” in the sense that it is explained by the long-term correlations, it is not necessarily intuitive or without alternative. Such an alternative hypothesis could be that fluctuations in BOLD act as noise which would lead to a negative correlation between BOLD variance and similarity to avFC. Instead, we find the opposite, which implies that, when reporting changes in FC, for example due to task, disease, anaesthesia, etc., it should be carefully assessed which role the changes in activity play. In other words, changes in FC could just be “apparent” changes and not really indicate any difference in the communication between the areas in question. In general, while it is easy to define FC in terms of statistical dependency, it seems that to what exactly it refers in terms of the real brain is not as clear.



To quantify the contribution of BOLD fluctuations to observed changes in FC, the surrogates used here could be useful because no real changes in FC occur. A suitable metric to quantify the distance between distributions of measures of dynamics, for example based on entropy, could capture the “residual” differences between surrogates and real data caused by real changes in FC.

### 3.4.3 Time scales and possible origins of fluctuations

Since the Fourier spectra of the time courses are preserved in the surrogate data, the fluctuations exhibit the same time scales as per construction (figure 3.4). Therefore, one limitation of our methods here is that we cannot make clear predictions about the origin of the observed time scales. It has been shown that slow ( $<0.1$  Hz) fluctuations observed on the level of BOLD activity can arise from gamma band oscillations on the neuronal level (Cabral et al., 2011, 2014), and it has been suggested (Deco and Kringelbach, 2016) that they are a signature of long-range communication in the communication through coherence (CTC)-framework (Fries, 2005). Furthermore, it was recently found by Matsui et al. (2016) in a mouse study that global waves of activity occur in both Calcium signal and in the hemodynamic signal, and that the latter follows the former, implying a causal relationship. Therefore it seems that the observed fluctuations are at least partly of neural origin. It remains to be seen what the exact relationship between the slow fluctuations and the travelling waves is, but they might originate from the same mechanisms.

Fluctuations of our three global measures occur on an even slower time scale ( $\approx 0.01$  Hz, figure 3.4). The time scale reported here is consistent with previous findings (Hutchison et al., 2012; Hansen et al., 2014; Ponce-Alvarez et al., 2015), and in particular, modelling studies (Hansen et al., 2014; Ponce-Alvarez et al., 2015) have shown that they arise from spontaneous fluctuations in the band around 0.1 Hz and are slow modulations in the level of synchronization, manifesting as global state transitions.

### 3.4.4 RSN activation

Previous studies have suggested that RSN activation occurs as a sequence of transients (Baker et al., 2014; Ponce-Alvarez et al., 2015), putting the focus on their separateness in time. Here, we find that they co-activate a lot of the time, which was also reported by de Pasquale et al. (2012); Smith et al. (2012). We find that RSNs are most strongly activated and co-activated

during periods of strong synchronization (i.e., high FC strength). This is at odds with the idea that each peak in the measures represents a distinct “state”, and suggests that, rather, we are looking at stereotypical events. This is in line with the interpretation given in Mitra et al. (2014): the activation of RSNs is embedded in stereotypical events triggered by global waves of BOLD activity. It is conceivable that the sequences observed by Baker et al. (2014) are due to the faster time scale in that study (MEG). Matsui et al. (2016) found that in mice, typical co-activations between brain areas are embedded in global waves of activity, and that this is true both on the level of Calcium signals and the hemodynamic signal.

If RSNs are most strongly activated and co-activated during periods of high FC strength, high similarity to avFC, and high BOLD variance, the activation of single RSNs should occur during periods of intermediate values of these measures, as is indeed suggested by the negative correlation between contribution index and all three measures (section 3.3.7).

### 3.4.5 Conclusion: Modelling of RS dynamics

We have shown in the previous chapter (chapter 2) that no additional assumptions about dynamics are necessary to model time-evolving community structure. There, we found that RSNs can be extracted from brief intervals which correspond to the peaks in the global measures used here. However, here, we do find interesting properties in the dynamics not explained by a fluctuating avFC.

The fluctuating global BOLD variance could represent an “echo” of global waves of activity occurring on a finer spatial and temporal scale (Matsui et al., 2016). Each wave can be understood as a stereotypical event (Mitra et al., 2014) which evolves from a state of low BOLD variance, low FC strength and low “expression” of the avFC structure, i.e. low similarity of FC structure to avFC, to a state where these measures are high. The transition from “low” to “high” does not happen linearly, as is evident from 1) the fact that the expression of the avFC grows faster than BOLD variance, and 2) the fact that there are many more “low” windows than expected, and also slightly more “high” windows. Both facts in conjunction lead to a strongly non-linear relationship between the two measures. The underlying mechanism could be the particular structure of the connectome or a boosting mechanism, perhaps resonance (Atasoy et al., 2016).

The model used in the previous chapter (Deco et al., 2014) is unlikely to reproduce these properties because it only possess one attractor. Possibly a

model which has previously been linked to multistability at the level of neural fluctuations (Freyer et al., 2011), and which has recently been shown to reproduce features of fMRI RS temporal dynamics (Deco and Kringelbach, 2016), could be of interest. As before, the key is that the system is close to a bifurcation, in this case, a subcritical Hopf bifurcation. Cast in these terms, the system would be in a regime of damped oscillations when the global measures are low, and in an oscillatory regime when they are high; switching occurs due to noise.

In terms of RSNs, the transition from “low” to “high” means a transition from none of them being activated to all of them being co-activated. Here, the non-linear relationship is evident from the histograms of the contributions (figure 3.9, right), where both strong co-activation (“high” windows) and absence of activation (“low” windows) are over-represented. The recruitment of single RSNs occurs at intermediate values, and could be seen as excursions from the stereotypical path from “low” to “high” and thus from the mere fluctuation of avFC. Therefore, more data will be necessary to obtain a sufficient sampling of different kinds of excursions, i.e. to tie the activation of certain RSNs to certain configurations of our three measures. This could provide a promising entry point to pinpoint differences between rest and task, in that specific recruitment of task-related networks should happen more strongly and reliably in task than in rest.



---

## General Discussion

The goal of this thesis was to elucidate the temporal dynamics of resting state (RS) as measured with fMRI in humans. In chapter 1, we gave an introduction to the field which developed over the last 20 years and whose perhaps best-known discovery is that of coherent sets of brain regions that form distributed functional networks, so-called resting state networks (RSNs). On the one hand, RSNs are stable, reproducible patterns of co-activated and correlated brain regions - which we have called communities in this thesis - that are functionally related. On the other hand, these RSNs have been shown to interact over time, exhibiting rich spatio-temporal dynamics. This places spontaneous fluctuations at the heart of one of the main topics in neuroscience: integration and segregation; both specialization and cooperation are necessary for brain function, cognition and behavior.

We have reviewed some evidence for the neural basis of the fluctuations observed with fMRI, stemming from electrophysiological recordings in humans and animals, as well as MEG and EEG. Computational models have demonstrated that gamma band oscillations on the neuronal level can lead to the slow ( $< 0.1$  Hz) fluctuations on the macroscopic scale. Dynamical systems theory explains the transmission of these phenomena through many levels of spatial and temporal resolution by placing the brain at the edge of a bifurcation, or phase transition. This can lead from local seemingly random fluctuations to well-ordered macroscopic relationships.

Another important point is the concept of graph models, which also underlies many of the computational models that are in use. In this view, activity propagates from one brain region (node) to another via a connection (edge) between them. Using such concepts, we can explore the properties of the

brain network which are at the base of its functional organization, like hierarchy, modularity, efficiency of communication, and resilience against damage, and, using simulations, their effect on the dynamics.

On the methodological side, an overview was given over decomposition techniques and the very general ideas behind them like inverse problem solving, dimensionality reduction, and feature extraction. We hope to have convinced the reader that, while ICA is one of the most widely-applied decomposition technique in fMRI today, there is neither a mathematical nor a biological reason why it should be the best, let alone only, possible approach. Especially for community detection, tensor factorization is a method whose results should be considered alongside those from established techniques because they can provide a different, and complementary, perspective on the data. It is important to understand that the data can only be approximated by a strictly modular structure, so any method attempting this has its limitations.

In the main part of the thesis, which contains the studies done in this PhD, chapters 2 and 3, we presented data from 21-minute fMRI resting state scans (TR=2 min) of 24 healthy subjects, coarsely parcellated into 66 regions covering the cortex.

In chapter 2, we used tensor factorization to extract communities that are representative for all subjects from the BOLD time courses and showed that they are also present in simulated data. In order to do so, communities were extracted from each subject individually and then clustered. Clustering performance was used to tune the parameters of the analysis pipeline, namely number of extracted features  $F$ , number of clusters  $K$ , and binarization threshold  $\theta$ . This approach was chosen in order to make sure that the factorization extracts features that generalize across subjects. We succeeded in doing so and identified four RSNs: the DMN, the visual network, the sensorimotor network, and the control networks.

We revealed a number of important points that should to be taken into account in future studies. First, the use of correlation-based measures for estimating functional connectivity was demonstrated to have its limitations in that it picked up on “residual” structure, meaning that a slight cluster structure could also be found in surrogate data that do not have any long-term correlations between the ROIs. This resulted in the real data only exhibiting slightly more clustering than the surrogates. We suggested mutual information as a measure that was superior in this application because it was shown not to be sensitive to this residual structure. Second, using

a non-negativity constraint lead to better clustering, and resulting features were more easily interpreted as membership weights. This is a constraint that is justified by the data and does not limit the space of solutions unnecessarily as ICA does. Third, we used a high threshold to binarize the tensors before decomposing them, the best clustering result being achieved when keeping only 2% of the pairwise FC values over the entire time course.

Subsequently, we employed a dynamic mean field model of the cortex to obtain simulated data, and showed that also in tensors created from these time courses, the abovementioned networks could be extracted. Using effective connectivity (EC) to couple the nodes of the network in our simulations proved essential, not only because it added missing interhemispheric connections, but also because it included asymmetrical weights of the connections. As a result, the node degrees were shown to be more or less uniformly distributed (as opposed to the SC, where there are a few nodes with much higher degrees than the rest). This, and the asymmetry, allows the activity to propagate throughout the entire network in a more diverse way, leading to more patterns that were more clearly pronounced.

---

In chapter 3, we developed a simple but effective approach which allowed us to directly analyze temporal dynamics. We used both the BOLD signals themselves as well as the time courses obtained from the tensor factorization which quantify RSN activation. We introduced three global measures of RS dynamics which were computed inside of sliding windows for each subject. Because we wanted to connect fluctuations in activity with fluctuations in FC, we chose one measure that tracks the former - namely the BOLD variance averaged over all regions - and two that track the latter - namely FC strength, i.e. correlation averaged over all pairs of ROIs, as well as similarity to average FC. This last measure was motivated by the null hypothesis applied in this chapter, stating that observations can be explained by a single fluctuating FC structure instead of transitions between (more or less) discrete states. Therefore, we measured, for each window, how similar the correlation structure was to the average one. To compare our results against the null hypothesis, we used surrogates, as before, but this time, the long-term correlations were preserved.

While we confirmed that the presence of fluctuations of the magnitude that we observed can be considered an effect of the long-term correlations, we also showed that the full distributions are significantly different from what is predicted by the surrogates. In particular, we found that there is a slightly bimodal structure in the three measures as well as the RSN time courses,

which lead us to postulate a “high” and a “low” state. The “high” state is characterized by high FC strength, high average BOLD variance, and high similarity to the average FC structure, as well as strong RSN activation and, importantly, co-activation; the opposite is the case for the “low” state. Furthermore, the transitions between “high” and “low” states were found to be non-linear, as was demonstrated by the joint distribution of BOLD variance and similarity to average FC. RSNs are largely co-activated during “high” periods, leading us to hypothesize that activation of single RSNs happens during periods of transition between highly and less synchronized states.

We take our results to support the idea that temporal dynamics of RS are due to underlying global waves of activity which spread over the entire cortex. Our findings contest the notion that the slow fluctuations themselves are the signatures of the brain transitioning between distinct states. Rather, we suggest that these states are embedded in the global fluctuations, and that the fluctuations themselves are the signatures of stereotypical events whose origin remains to be determined.

Future steps include the application of the methods developed here to task data. The goal is to define task dynamics in terms of a deviation from resting dynamics, or in other words, to determine how the brain prepares for task during rest, a notion that is firmly rooted in the predictive coding framework. One problem is that performing a task causes fluctuations in the brain that are related to the specifics of the experimental design, which would make it difficult to compare to RS. Furthermore, it is a major goal to model the temporal dynamics using the propositions made here and to test whether they produce meaningful predictions. Perhaps these predictions could even help with formulating hypotheses for the task-related analysis. For the modelling, the large inter-subject variability is definitely an obstacle, and clear criteria are needed in order to decide whether a model “reproduces” a certain property of the data or not. In this study, we used surrogate data to accomplish this goal, a practice that we will continue.

## 4.1 Contributions

In this dissertation, we make some unique contributions to the investigation of spatio-temporal dynamics of human fMRI RS. We employ a whole-brain perspective, using coarse parcellations, long-range connectivity, and mean field modelling. This way, we add to the evidence suggesting that impor-



tant insights can be had by considering the macroscopic level of description instead of striving for more and more detail. Indeed, especially in this field, the most remarkable observations can only be made when considering anatomically far away brain regions; considering small regions gives the appearance of randomness and leaves us blind for the beautiful regularities in brain dynamics.

To begin with, we describe in detail a methodology based on non-negative tensor factorization that can be applied for community detection in empirical and model data. Importantly, it works on the spatial resolution that is common for whole-brain modelling and graph analysis. It is based on generic models of n-way-decomposition which can be seen as a generalization of common methods such as PCA and SVD to higher-dimensional arrays. The main point about this method is that, in a natural way, it allows us to analyze space and time in a truly simultaneous fashion, without having to make any assumptions about independence in either dimension. A crucial part of our methodology is the way we validate our results. Since the factorization is applied to each subject's data separately, we use the clustering of the extracted communities to tune the parameters of the decomposition. This means we do not have to assume anything about the communities, except that they are common across subjects. Furthermore, we use surrogates which enable us to quantify which amount of clustering is obtained by chance, without any true cluster structure being present in the data. The pipeline we have developed can easily be applied to datasets of similar resolutions and will be made available on-line.

With this approach, we explicitly show that RSNs are present in data simulated by a dynamic mean field model of RS. Previously, it had been established that this model reproduces the average FC (Deco et al., 2014), but not that RSNs themselves can be isolated from the time courses. This lends support to the idea that a single attractor is sufficient to explain the presence of RSNs, and that they are largely derived from the average FC. Of course, this does not mean that this kind of model can reproduce the temporal dynamics, which indeed we contest with our results in chapter 3.

By using two different connectivities in the simulations of the model data, we also confirm that an underlying complex network and simple dynamics work together to produce FC patterns. Specifically, we apply a very recent method that allows us to derive an EC on the group level for our dataset, and we demonstrate that 1) interhemispheric connections, 2) asymmetry of in- and out-degree, and 3) relatively uniform distribution of node degrees

contribute to producing realistic communities. Further studies are necessary to determine how exactly the dynamics on the EC play out based on these properties.

Another notable result presented in this thesis is that brief instances of maximum FC are sufficient to obtain RSNs. It is remarkable that only the highest 2% of pairwise FC values, over all windows, are necessary for the tensor factorization algorithm to characterize the spatial dynamics. This speaks to its infraslow evolution and its stereotypical nature, and suggests that, while RSNs are undoubtedly of special interest, they may represent the most basic components of brain dynamics, much like the first principal component in PCA. Much work remains to be done in characterizing and explaining inter-individual differences, something that also whole-brain computational models are to date not able to account for.

We introduce three measures of resting state dynamics, which are very simple and capture global changes in FC and BOLD activity. As mentioned above, the whole-brain approach lies at the heart of this thesis, and using these measures works around statistical problems when considering measures pair by pair. We show that they exhibit ultraslow fluctuations ( $\approx 0.01$  Hz) as reported in other studies using different measures of global RS dynamics. Despite their simplicity, they are shown to carry meaning by connecting them to RSN activation time courses which are obtained in a completely different way (using tensor factorization). Instances of high overall FC are equivalent to periods of maximum variance. We suggest this in chapter 2 by demonstrating that the amount of FC pairs which exceed a high (98th percentile) threshold fluctuates strongly over time, and confirm it in chapter 3 by unambiguously connecting global measures of FC with one of activity, such that it becomes clear that they are two sides of the same coin. While this finding may be “trivial” in the sense that surrogates readily reproduce it, this perspective has not been thoroughly investigated and calls for a careful evaluation of the meaning of “FC” because measured changes between experimental conditions could be due to temporary changes in activity instead of any changes in the connections themselves. Furthermore, it is shown that the measures and the RSN time courses exhibit a strong positive correlation, suggesting that their temporal evolution is governed by these global properties.

We make our contribution to a growing body of literature that emphasizes that RSNs are not clearly separated in time. ICA, which is widely applied and has unquestionable merits, has had a big influence on how we under-

stand and interpret RSNs. By using an alternative approach which does not assume independence (tensor factorization), and deriving from the obtained time courses a very simple index of co-activation, we show that they have a preference for co-activating strongly. This co-activation occurs mostly during periods of high FC strength, high similarity to average FC, and strong BOLD variance. We conclude, therefore, that the ultraslow fluctuations do not represent different FC patterns, but instead, the transitions of the system from a “low” to a “high” synchronization state, which happens in a fairly stereotypical manner.

Concerning these “states”, we furthermore establish that the joint distributions of two of the three measures, BOLD variance and similarity to avFC, as well as that of RSN (co-)activations, shows a non-linear and bimodal distribution, i.e. the “high” and “low” states actually exhibit longer dwell times than intermediate states. Furthermore, the transition between the states does not seem to happen at a constant speed, but the similarity to average FC grows faster than is predicted by the surrogates. This suggests an additional mechanism, perhaps resonance, caused by the underlying connectivity, perhaps a synchronization effect caused by a Hopf bifurcation. Future modelling studies should try to explain this bimodal behavior, and we suggest to use the Hopf model which can exhibit bi- and multistable behavior, with transitions between states being driven by noise. We hypothesize here that using effective connectivity to connect nodes could lead to the desired dynamics in FC due to asymmetries.

## 4.2 Conclusion

One of the points that this thesis hopes to make is that it is not clear what kind of dynamics leads to the emergence of RSNs. It could be that we have a multistable or metastable system which transitions between different states identified with the dominance of a certain RSN; it could be that we only have one attractor and RSNs are nothing but a result of the modular structure of the underlying anatomical connections; or we could consider travelling waves which propagate through the connectome in a specific way, leading to the expression of subnetworks at different points in time. Computational models will be the key to further understanding these scenarios, making predictions for empirical data, and refining the ideas we have about the relationship between micro- and macroscale, and between structure and function. While it is clear that we have not fully reached the ambitious goal stated above - i.e. “to elucidate the temporal dynamics of resting state”

- we have developed numerous methods to investigate this question, have clarified many details, and identified promising future directions based on hypotheses derived from the results presented here.

APPENDIX **A**

---

Table A.1: ROIs in the same order in which they appear in the vectors in figures 2.6 and 2.7, as well as in the matrices in figure 2.8.

	name		name
1	R entorhinal cortex	34	L posterior-cing cortex
2	R parahippocampal gyrus	35	L precuneus cortex
3	R temp pole	36	L isthmuscing cortex
4	R front pole	37	L paracentral lobule
5	R fusiform gyrus	38	L cuneus cortex
6	R transverse temp cortex	39	L pericalcarine cortex
7	R lateral occipital cortex	40	L lingual gyrus
8	R sup parietal cortex	41	L medial orbital front cortex
9	R inf temp gyrus	42	L sup front gyrus
10	R inf parietal cortex	43	L rostral ant cing cortex
11	R supramarginal gyrus	44	L caudal ant-cing cortex
12	R banks sup temp sulcus	45	L lateral orbital front cortex
13	R middle temp gyrus	46	L pars orbitalis
14	R sup temp gyrus	47	L rostral middle front gyrus
15	R postcentral gyrus	48	L pars triangularis
16	R precentral gyrus	49	L pars opercularis
17	R caudal middle front gyrus	50	L caudal middle front gyrus
18	R pars opercularis	51	L precentral gyrus
19	R pars triangularis	52	L postcentral gyrus
20	R rostral middle front gyrus	53	L sup temp gyrus
21	R pars orbitalis	54	L middle temp gyrus
22	R lateral orbital front cortex	55	L banks sup temp sulcus
23	R caudal ant-cing cortex	56	L supramarginal gyrus
24	R rostral ant cing cortex	57	L inf parietal cortex
25	R sup front gyrus	58	L inf temp gyrus
26	R medial orbital front cortex	59	L sup parietal cortex
27	R lingual gyrus	60	L lateral occipital cortex
28	R pericalcarine cortex	61	L transverse temp cortex
29	R cuneus cortex	62	L fusiform gyrus
30	R paracentral lobule	63	L front pole
31	R isthmuscing cortex	64	L temp pole
32	R precuneus cortex	65	L parahippocampal gyrus
33	R posterior-cing cortex	66	L entorhinal cortex

---

## Bibliography

Each reference indicates the pages where it appears.

- Acar E, Dunlavy DM, Kolda TG (2011) A scalable optimization approach for fitting canonical tensor decompositions. *Journal of Chemometrics* 25:67–86. 33
- Allen EA, Damaraju E, Plis SM, Erhardt EB, Eichele T, Calhoun VD (2012) Tracking whole-brain connectivity dynamics in the resting state. *Cerebral Cortex* p. bhs352. 8, 10, 24, 38, 51, 56, 57, 68
- Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273:1868. 8
- Atasoy S, Donnelly I, Pearson J (2016) Human brain networks function in connectome-specific harmonic waves. *Nature communications* 7. 78
- Bader BW, Kolda TG et al. (2015) Matlab tensor toolbox version 2.6, Available online. 33
- Baker AP, Brookes MJ, Rezek IA, Smith SM, Behrens T, Smith PJP, Woolrich M (2014) Fast transient networks in spontaneous human brain activity. *Elife* 3:e01867. 77, 78
- Barttfeld P, Uhrig L, Sitt JD, Sigman M, Jarraya B, Dehaene S (2015) Signature of consciousness in the dynamics of resting-state brain activity. *Proceedings of the National Academy of Sciences* 112:887–892. 8
- Basser PJ, Mattiello J, LeBihan D (1994) MR diffusion tensor spectroscopy and imaging. *Biophysical journal* 66:259. 7
- Bassett DS, Wymbs NF, Porter MA, Mucha PJ, Carlson JM, Grafton ST

- (2011) Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences* 108:7641–7646. 11
- Beckmann CF, Smith SM (2005) Tensorial extensions of independent component analysis for multisubject fMRI analysis. *NeuroImage* 25:294–311. 4, 16, 20
- Beckmann CF, DeLuca M, Devlin JT, Smith SM (2005) Investigations into resting-state connectivity using independent component analysis. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* 360:1001–13. xiii, 4, 5, 25, 48, 50, 56
- Beckmann CF, Mackay CE, Filippini N, Smith SM (2009) Group comparison of resting-state fMRI data using multi-subject ICA and dual regression. *NeuroImage* 47:S148. 20
- Beckmann CF, Smith SM (2004) Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE transactions on medical imaging* 23:137–52. 24, 39, 50
- Bell AJ, Sejnowski TJ (1995) An information-maximization approach to blind separation and blind deconvolution. *Neural computation* 7:1129–1159. 4
- Bettinardi RG, Tort-Colet N, Ruiz-Mejias M, Sanchez-Vives MV, Deco G (2015) Gradual emergence of spontaneous correlated brain activity during fading of general anesthesia in rats: evidences from fMRI and local field potentials. *NeuroImage* 114:185–198. 56
- Betzel RF, Byrge L, He Y, Goñi J, Zuo XN, Sporns O (2014) Changes in structural and functional connectivity among resting-state networks across the human lifespan. *NeuroImage* 102:345–357. 5
- Betzel RF, Fukushima M, He Y, Zuo XN, Sporns O (2016) Dynamic fluctuations coincide with periods of high and low modularity in resting-state functional brain networks. *NeuroImage* 127:287–297. 8, 11, 14, 49, 56, 57, 68
- Biswal B, Zerrin Yetkin F, Haughton V, Hyde J (1995) Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic resonance in medicine* 34:537–541. 1, 24, 56
- Boly M, Balteau E, Schnakers C, Degueldre C, Moonen G, Luxen A, Phillips C, Peigneux P, Maquet P, Laureys S (2007) Baseline brain activity fluctuations predict somatosensory perception in humans. *Proceedings of the National Academy of Sciences* 104:12187–12192. 3
- Bosking WH, Zhang Y, Schofield B, Fitzpatrick D (1997) Orientation selectivity and the arrangement of horizontal connections in tree shrew striate



- cortex. *The Journal of neuroscience* 17:2112–2127. 13
- Boveroux P, Vanhauzenhuysse A, Bruno MA, Noirhomme Q, Lauwick S, Luxen A, Degueldre C, Plenevaux A, Schnakers C, Phillips C et al. (2010) Breakdown of within-and between-network resting state functional magnetic resonance imaging connectivity during propofol-induced loss of consciousness. *The Journal of the American Society of Anesthesiologists* 113:1038–1053. 5
- Brewer JA, Worhunsky PD, Gray JR, Tang YY, Weber J, Kober H (2011) Meditation experience is associated with differences in default mode network activity and connectivity. *Proceedings of the National Academy of Sciences* 108:20254–20259. 5
- Britz J, Van De Ville D, Michel CM (2010) BOLD correlates of EEG topography reveal rapid resting-state network dynamics. *NeuroImage* 52:1162–1170. 1, 11
- Brookes MJ, Woolrich M, Luckhoo H, Price D, Hale JR, Stephenson MC, Barnes GR, Smith SM, Morris PG (2011) Investigating the electrophysiological basis of resting state networks using magnetoencephalography. *Proceedings of the National Academy of Sciences* 108:16783–16788. 4, 12, 19
- Broyd SJ, Demanuele C, Debener S, Helps SK, James CJ, Sonuga-Barke EJ (2009) Default-mode brain dysfunction in mental disorders: a systematic review. *Neuroscience & biobehavioral reviews* 33:279–296. 3
- Bullmore E, Sporns O (2012) The economy of brain network organization. *Nature Reviews Neuroscience* 13:336–349. 6
- Cabral J, Hugues E, Kringelbach ML, Deco G (2012) Modeling the outcome of structural disconnection on resting-state functional connectivity. *NeuroImage* 62:1342–1353. 24
- Cabral J, Hugues E, Sporns O, Deco G (2011) Role of local network oscillations in resting-state functional connectivity. *NeuroImage* 57:130–139. 13, 77
- Cabral J, Luckhoo H, Woolrich M, Joensson M, Mohseni H, Baker A, Kringelbach ML, Deco G (2014) Exploring mechanisms of spontaneous functional connectivity in MEG: how delayed network interactions lead to structured amplitude envelopes of band-pass filtered oscillations. *NeuroImage* 90:423–435. 13, 77
- Calhoun V, Adali T, Pearlson G, Pekar J (2001a) A method for making group inferences from functional mri data using independent component analysis. *Human brain mapping* 14:140–151. 20

- Calhoun V, Adali T, Pearlson G, Pekar J (2001b) Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms. *Human Brain Mapping* 13:43–53. 50
- Calhoun VD, Adali T (2012) Multisubject independent component analysis of fMRI: a decade of intrinsic networks, default mode, and neurodiagnostic discovery. *IEEE reviews in biomedical engineering* 5:60–73. 3
- Calhoun VD, Kiehl KA, Pearlson GD (2008) Modulation of temporally coherent brain networks estimated using ICA at rest and during cognitive tasks. *Human brain mapping* 29:828–838. 10
- Calhoun VD, Liu J, Adali T (2009) A review of group ICA for fMRI data and ICA for joint inference of imaging, genetic, and ERP data. *NeuroImage* 45:S163–72. 50
- Calhoun VD, Miller R, Pearlson G, Adali T (2014) The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery. *Neuron* 84:262–274. 8, 57
- Carhart-Harris RL, Erritzoe D, Williams T, Stone JM, Reed LJ, Colasanti A, Tyacke RJ, Leech R, Malizia AL, Murphy K et al. (2012) Neural correlates of the psychedelic state as determined by fMRI studies with psilocybin. *Proceedings of the National Academy of Sciences* 109:2138–2143. 5
- Carhart-Harris RL, Friston KJ (2010) The default-mode, ego-functions and free-energy: a neurobiological account of freudian ideas. *Brain* 133:1265–1283. 3, 14
- Castellanos FX, Margulies DS, Kelly C, Uddin LQ, Ghaffari M, Kirsch A, Shaw D, Shehzad Z, Di Martino A, Biswal B et al. (2008) Cingulate-precuneus interactions: a new locus of dysfunction in adult attention-deficit/hyperactivity disorder. *Biological psychiatry* 63:332–337. 5
- Chan MY, Park DC, Savalia NK, Petersen SE, Wig GS (2014) Decreased segregation of brain systems across the healthy adult lifespan. *Proceedings of the National Academy of Sciences* 111:E4997–E5006. 5
- Chang C, Glover GH (2010) Time–frequency dynamics of resting-state brain connectivity measured with fMRI. *NeuroImage* 50:81–98. xiv, 8, 9, 11, 24, 38, 56, 68
- Chawla D, Lumer ED, Friston KJ (1999) The relationship between synchronization among neuronal populations and their mean activity levels. *Neural Computation* 11:1389–1411. 76
- Cherkassky VL, Kana RK, Keller TA, Just MA (2006) Functional connectivity in a baseline resting-state network in autism. *Neurore-*

- port* 17:1687–1690. 5
- Chialvo DR (2010) Emergent complex neural dynamics. *Nature physics* 6:744–750. 9
- Cichocki A, Mandic D, Phan A, Caiafa C, Zhou G, Zhao Q, De Lathauwer L (2014) Tensor decompositions for signal processing applications from two-way to multiway component analysis. *arXiv preprint* . 38
- Cichocki A (2013) Tensor decompositions: A new concept in brain data analysis? *arXiv* pp. 507–517. xiv, 17, 18, 20, 21, 25, 50
- Cichocki A, Zdunek R, Phan AH, Amari S (2009) *Nonnegative matrix and tensor factorizations: applications to exploratory multi-way data analysis and blind source separation* John Wiley & Sons. xiv, 17, 20, 31, 48
- Cole MW, Bassett DS, Power JD, Braver TS, Petersen SE (2014) Intrinsic and task-evoked network architectures of the human brain. *Neuron* 83:238–251. 3
- Cole MW, Reynolds JR, Power JD, Repovs G, Anticevic A, Braver TS (2013) Multi-task connectivity reveals flexible hubs for adaptive task control. *Nature neuroscience* 16:1348–1355. 8
- Cordes D, Haughton VM, Arfanakis K, Wendt GJ, Turski PA, Moritz CH, Quigley MA, Meyerand ME (2000) Mapping functionally related regions of brain with functional connectivity MR imaging. *American Journal of Neuroradiology* 21:1636–1644. 2, 24, 56
- Coste CP, Sadaghiani S, Friston KJ, Kleinschmidt A (2011) Ongoing brain activity fluctuations directly account for intertrial and indirectly for intersubject variability in stroop task performance. *Cerebral cortex* 21:2612–2619. 3
- Damoiseaux J, Beckmann C, Arigita ES, Barkhof F, Scheltens P, Stam C, Smith S, Rombouts S (2008) Reduced resting-state brain activity in the default network in normal aging. *Cerebral cortex* 18:1856–1864. 5
- Damoiseaux J, Rombouts S, Barkhof F, Scheltens P, Stam C, Smith S, Beckmann C (2006) Consistent resting-state networks across healthy subjects. *Proceedings of the national academy of sciences* 103:13848–13853. 4, 50, 56
- de Amorim RC, Hennig C (2015) Recovering the number of clusters in data sets with noise features using feature rescaling factors. *Information Sciences* 324:126–145. 34
- De Luca M, Beckmann C, De Stefano N, Matthews P, Smith S et al. (2006) FMRI resting state networks define distinct modes of long-distance interactions in the human brain. *NeuroImage* 29:1359–1367. 50, 56

- De Luca M, Smith S, De Stefano N, Federico A, Matthews PM (2005) Blood oxygenation level dependent contrast resting state networks are relevant to functional activity in the neocortical sensorimotor system. *Experimental brain research* 167:587–594. 2, 3
- de Pasquale F, Della Penna S, Snyder AZ, Lewis C, Mantini D, Marzetti L, Belardinelli P, Ciancetta L, Pizzella V, Romani GL et al. (2010) Temporal dynamics of spontaneous MEG activity in brain networks. *Proceedings of the National Academy of Sciences* 107:6040–6045. 1, 12, 13
- de Pasquale F, Della Penna S, Snyder AZ, Marzetti L, Pizzella V, Romani GL, Corbetta M (2012) A cortical core for dynamic integration of functional networks in the resting human brain. *Neuron* 74:753–764. 8, 12, 77
- Deco G, Ponce-Alvarez A, Hagmann P, Romani G, Mantini D, Corbetta M (2014) How local excitation-inhibition ratio impacts the whole brain dynamics. *The Journal of Neuroscience* 34:7886–7898. 21, 25, 26, 28, 38, 48, 50, 78, 85
- Deco G, Jirsa V, McIntosh A, Sporns O, Kötter R (2009) Key role of coupling, delay, and noise in resting brain fluctuations. *Proceedings of the National Academy of Sciences* 106:10302–10307. 13
- Deco G, Jirsa VK (2012) Ongoing cortical activity at rest: criticality, multistability, and ghost attractors. *The Journal of Neuroscience* 32:3366–3375. 9, 24, 28, 57
- Deco G, Jirsa VK, McIntosh AR (2010) Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nature Reviews Neuroscience* 12:43–56. 6
- Deco G, Jirsa VK, McIntosh AR (2013) Resting brains never rest: computational insights into potential cognitive architectures. *Trends in neurosciences* 36:268–274. 3
- Deco G, Kringelbach ML (2016) Metastability and coherence: extending the communication through coherence hypothesis using a whole-brain computational perspective. *Trends in neurosciences* 39:125–135. 77, 79
- Deco G, McIntosh AR, Shen K, Hutchison RM, Menon RS, Everling S, Hagmann P, Jirsa VK (2014) Identification of optimal structural connectivity using functional connectivity and neural modeling. *The Journal of Neuroscience* 34:7910–6. 24
- Deco G, Ponce-Alvarez A, Mantini D, Romani GL, Hagmann P, Corbetta M (2013) Resting-state functional connectivity emerges from structurally and dynamically shaped slow linear fluctuations. *The Journal of Neuro-*

- science* 33:11239–11252. 28
- Desikan RS, Segonne F, Fischl B, Quinn BT, Dickerson BC, Blacker D, Buckner RL, Dale AM, Maguire RP, Hyman BT, Albert MS, Killiany RJ (2006) An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage* 31:968–980. 26
- Dosenbach NU, Nardos B, Cohen AL, Fair DA, Power JD, Church JA, Nelson SM, Wig GS, Vogel AC, Lessov-Schlaggar CN et al. (2010) Prediction of individual brain maturity using fMRI. *Science* 329:1358–1361. 5
- Doucet G, Naveau M, Petit L, Delcroix N, Zago L, Crivello F, Jobard G, Tzourio-Mazoyer N, Mazoyer B, Mellet E, Joliot M (2011) Brain activity at rest: a multiscale hierarchical functional organization. *Journal of Neurophysiology* 105:2753–2763. 24
- Ferezou I, Haiss F, Gentet LJ, Aronoff R, Weber B, Petersen CC (2007) Spatiotemporal dynamics of cortical sensorimotor integration in behaving mice. *Neuron* 56:907–923. 13
- Finn ES, Shen X, Scheinost D, Rosenberg MD, Huang J, Chun MM, Papademetris X, Constable RT (2015) Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nature Neuroscience* 18:1664–1671. 24
- Fischl B, Van Der Kouwe A, Destrieux C, Halgren E, Ségonne F, Salat DH, Busa E, Seidman LJ, Goldstein J, Kennedy D et al. (2004) Automatically parcellating the human cerebral cortex. *Cerebral Cortex* 14:11–22. 7
- Fortunato S (2010) Community detection in graphs. *Physics reports* 486:75–174. 21
- Fox M, Snyder A, Vincent J, Corbetta M, Van Essen D, Raichle M (2005) The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences* 102:9673–9678. xiii, 3, 4, 24, 48, 56
- Fox MD, Snyder AZ, Zacks JM, Raichle ME (2006) Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses. *Nature neuroscience* 9:23–25. 8, 57
- Fransson P (2006) How default is the default mode of brain function?: Further evidence from intrinsic bold signal fluctuations. *Neuropsychologia* 44:2836–2845. 2, 75
- Freyer F, Roberts JA, Becker R, Robinson PA, Ritter P, Breakspear M (2011) Biophysical mechanisms of multistability in resting-state cortical rhythms. *The Journal of Neuroscience* 31:6353–6361. 9, 79

- Fries P (2005) A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in cognitive sciences* 9:474–480. 77
- Friston KJ (1994) Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping* 2:56–78. 7, 25, 51
- Friston KJ (1997) Transients, metastability, and neuronal dynamics. *NeuroImage* 5:164–171. 9
- Friston K, Frith C, Liddle P, Frackowiak R (1993) Functional connectivity: the principal-component analysis of large (PET) data sets. *Journal of Cerebral Blood Flow & Metabolism* 13:5–14. 2
- Friston K, Harrison L, Penny W (2003) Dynamic causal modelling. *NeuroImage* 19:1273–1302. 28
- Fröhlich F, Sejnowski TJ, Bazhenov M (2010) Network bistability mediates spontaneous transitions between normal and pathological brain states. *The Journal of Neuroscience* 30:10734–10743. 57
- Garrett DD, Kovacevic N, McIntosh AR, Grady CL (2010) Blood oxygen level-dependent signal variability is more than just noise. *The Journal of Neuroscience* 30:4914–4921. 8
- Gauvin L, Panisson A, Cattuto C (2014) Detecting the community structure and activity patterns of temporal networks: a non-negative tensor factorization approach. *PLoS ONE* 9:e86028. 21, 25, 50
- Ghosh A, Rho Y, McIntosh AR, Kötter R, Jirsa VK (2008) Noise during rest enables the exploration of the brain’s dynamic repertoire. *PLoS Comput Biol* 4:e1000196. 13, 57
- Gilson M, Moreno-Bote R, Ponce-Alvarez A, Ritter P, Deco G (2016) Estimation of directed effective connectivity from fMRI functional connectivity hints at asymmetries in cortical connectome. *PLoS Computational Biology* . 25, 35, 36, 37, 43, 48, 51
- Glasser M, Coalson T, Robinson E, Hacker C, Harwell J, Yacoub E, Ugurbil K, Anderson J, Beckmann C, Jenkinson M et al. (2015) A multi-modal parcellation of human cerebral cortex. *Nature* . 3, 7
- Grady CL, Garrett DD (2014) Understanding variability in the BOLD signal and why it matters for aging. *Brain imaging and behavior* 8:274–83. 57
- Greicius M, Krasnow B, Reiss A, Menon V (2003) Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences* 100:253–258. 1, 75
- Greicius M (2008) Resting-state functional connectivity in neuropsychiatric

- disorders. *Current opinion in neurology* 21:424–430. 5
- Greicius MD, Srivastava G, Reiss AL, Menon V (2004) Default-mode network activity distinguishes Alzheimer’s disease from healthy aging: evidence from functional MRI. *Proceedings of the National Academy of Sciences* 101:4637–4642. 5
- Greicius MD, Supekar K, Menon V, Dougherty RF (2009) Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cerebral cortex* 19:72–78. 13
- Gusnard DA, Akbudak E, Shulman GL, Raichle ME (2001) Medial prefrontal cortex and self-referential mental activity: relation to a default mode of brain function. *Proceedings of the National Academy of Sciences* 98:4259–4264. 3, 56
- Hagmann P, Cammoun L, Gigandet X, Meuli R, Honey CJ, Wedeen VJ, Sporns O (2008) Mapping the structural core of human cerebral cortex. *PLoS Biology* 6:e159. 13, 72
- Haimovici A, Tagliazucchi E, Balenzuela P, Chialvo DR (2013) Brain organization into resting state networks emerges at criticality on a model of the human connectome. *Physical Review Letters* 110:178101. 8, 9, 24
- Handwerker DA, Roopchansingh V, Gonzalez-Castillo J, Bandettini PA (2012) Periodic changes in fMRI connectivity. *NeuroImage* 63:1712–1719. 14, 56, 62, 65, 76
- Hansen EC, Battaglia D, Spiegler A, Deco G, Jirsa VK (2014) Functional Connectivity Dynamics: Modeling the switching behavior of the resting state. *NeuroImage* 105:525–535. 8, 24, 51, 57, 77
- He BJ, Raichle ME (2009) The fMRI signal, slow cortical potential and consciousness. *Trends in cognitive sciences* 13:302–309. 13, 14
- Heiervang E, Behrens TEJ, Mackay CE, Robson MD, Johansen-Berg H (2006) Between session reproducibility and between subject variability of diffusion MR and tractography measures. *NeuroImage* 33:867–77. 24
- Hesselmann G, Kell CA, Eger E, Kleinschmidt A (2008) Spontaneous local variations in ongoing neural activity bias perceptual decisions. *Proceedings of the National Academy of Sciences* 105:10984–10989. 3
- Hindriks R, Adhikari M, Murayama Y, Ganzetti M, Mantini D, Logothetis N, Deco G (2015) Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI? *NeuroImage* 127:242–256. 14, 51, 56, 62, 65, 76
- Hipp JF, Hawellek DJ, Corbetta M, Siegel M, Engel AK (2012) Large-scale cortical correlation structure of spontaneous oscillatory activity. *Nature*

- neuroscience* 15:884–890. 12
- Honey CJ, Sporns O, Cammoun L, Gigandet X, Thiran JP, Meuli R, Hagmann P (2009) Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences* 106:2035–40. 6, 13, 24, 57
- Horowitz SG, Fukunaga M, de Zwart JA, van Gelderen P, Fulton SC, Balkin TJ, Duyn JH (2008) Low frequency BOLD fluctuations during resting wakefulness and light sleep: A simultaneous EEG-fMRI study. *Human brain mapping* 29:671–682. 5, 8, 56
- Hudetz AG, Liu X, Pillay S (2015) Dynamic repertoire of intrinsic brain states is reduced in propofol-induced unconsciousness. *Brain connectivity* 5:10–22. 57
- Hutchison RM, Womelsdorf T, Allen EA, Bandettini PA, Calhoun VD, Corbetta M, Della Penna S, Duyn JH, Glover GH, Gonzalez-Castillo J, Handwerker DA, Keilholz S, Kiviniemi V, Leopold DA, de Pasquale F, Sporns O, Walter M, Chang C (2013) Dynamic functional connectivity: promise, issues, and interpretations. *NeuroImage* 80:360–78. 8, 51, 56
- Hutchison RM, Womelsdorf T, Gati JS, Everling S, Menon RS (2012) Resting-state networks show dynamic functional connectivity in awake humans and anesthetized macaques. *Human Brain Mapping* 34:2154–2177. xiv, 10, 12, 24, 38, 56, 68, 77
- Hyvärinen A, Karhunen J, Oja E (2004) *Independent component analysis*, Vol. 46 John Wiley & Sons. 4
- Jafri MJ, Pearlson GD, Stevens M, Calhoun VD (2008) A method for functional network connectivity among spatially independent resting-state components in schizophrenia. *NeuroImage* 39:1666–1681. 8
- Jeurissen B, Leemans A, Tournier JD, Jones DK, Sijbers J (2013) Investigating the prevalence of complex fiber configurations in white matter tissue with diffusion magnetic resonance imaging. *Human Brain Mapping* 34:2747–66. 51
- Jones DT, Vemuri P, Murphy MC, Gunter JL, Senjem ML, Machulda MM, Przybelski SA, Gregg BE, Kantarci K, Knopman DS et al. (2012) Non-stationarity in the resting brains modular architecture. *PloS one* 7:e39731. 8, 10, 11
- Jones DK, Knösche TR, Turner R (2013) White matter integrity, fiber count, and other fallacies: the do's and don'ts of diffusion MRI. *NeuroImage* 73:239–54. 51
- Kaiser M, Hilgetag CC (2004) Edge vulnerability in neural and metabolic



- networks. *Biological cybernetics* 90:311–317. 6
- Kelso JS (2012) Multistability and metastability: understanding dynamic coordination in the brain. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 367:906–918. 9
- Kelso JSea (1995) *The self-organization of brain and behavior* Cambridge, MA: MIT Press. 9
- Kenet T, Bibitchkov D, Tsodyks M, Grinvald A, Arieli A (2003) Spontaneously emerging cortical representations of visual attributes. *Nature* 425:954–956. 13
- Kim J, Park H (2012) Fast nonnegative tensor factorization with an active-set-like method In *High-Performance Scientific Computing*, pp. 311–326. Springer. 33
- Kiviniemi V, Kantola JH, Jauhiainen J, Hyvärinen A, Tervonen O (2003) Independent component analysis of nondeterministic fMRI signal sources. *NeuroImage* 19:253–260. 4, 56
- Kiviniemi V, Vire T, Remes J, Elseoud AA, Starck T, Tervonen O, Nikkinen J (2011) A sliding time-window ICA reveals spatial variability of the default mode network in time. *Brain connectivity* 1:339–347. 10, 38, 56, 57, 68
- Kopell NJ, Gritton HJ, Whittington MA, Kramer MA (2014) Beyond the connectome: the dynamo. *Neuron* 83:1319–1328. 7, 8, 57
- Kraskov A, Stögbauer H, Grassberger P (2004) Estimating mutual information. *Physical Review E* 69:066138. 30, 38, 39, 51
- Kringelbach ML, McIntosh AR, Ritter P, Jirsa VK, Deco G (2015) The rediscovery of slowness: exploring the timing of cognition. *Trends in cognitive sciences* 19:616–628. 14
- Kühn S, Gallinat J (2013) Resting-state brain activity in schizophrenia and major depression: a quantitative meta-analysis. *Schizophrenia bulletin* 39:358–365. 5
- Larson-Prior LJ, Zempel JM, Nolan TS, Prior FW, Snyder AZ, Raichle ME (2009) Cortical network functional connectivity in the descent to sleep. *Proceedings of the National Academy of Sciences* 106:4489–4494. 56
- Laumann TO, Snyder AZ, Mitra A, Gordon EM, Gratton C, Adeyemo B, Gilmore AW, Nelson SM, Berg JJ, Greene DJ et al. (2016) On the stability of BOLD fMRI correlations. *Cerebral Cortex* . 14, 56, 76
- Leonardi N, Van de Ville D (2013) Identifying network correlates of brain states using tensor decompositions of whole-brain dynamic functional

- connectivity In *2013 International Workshop on Pattern Recognition in Neuroimaging*, pp. 74–77. IEEE. 21, 25, 38, 50
- Lindquist MA, Xu Y, Nebel MB, Caffo BS (2014) Evaluating dynamic bivariate correlations in resting-state fMRI: a comparison study and a new approach. *NeuroImage* 101:531–46. 14, 51, 56, 76
- Liu X, Chang C, Duyn JH (2013) Decomposition of spontaneous brain activity into distinct fMRI co-activation patterns. *Frontiers in Systems Neuroscience* 7:101. 24, 38, 56, 57, 68
- Liu X, Duyn JH (2013) Time-varying functional network information extracted from brief instances of spontaneous brain activity. *Proceedings of the National Academy of Sciences* 110:4392–4397. 8, 10, 21
- Lowe M, Mock B, Sorenson J (1998) Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. *NeuroImage* 7:119–132. 2, 56
- Luczak A, Barthó P, Marguet SL, Buzsáki G, Harris KD (2007) Sequential structure of neocortical spontaneous activity in vivo. *Proceedings of the National Academy of Sciences* 104:347–352. 13
- Majeed W, Magnuson M, Hasenkamp W, Schwarb H, Schumacher EH, Barsalou L, Keilholz SD (2011) Spatiotemporal dynamics of low frequency BOLD fluctuations in rats and humans. *NeuroImage* 54:1140–1150. 13
- Mantini D, Perrucci M, Del Gratta C, Romani G, Corbetta M (2007) Electrophysiological signatures of resting state networks in the human brain. *Proceedings of the National Academy of Sciences* 104:13170. 11, 25, 39, 48, 50, 56
- Markov NT, Vezoli J, Chameau P, Falchier A, Quilodran R, Huissoud C, Lamy C, Misery P, Giroud P, Ullman S et al. (2014) Anatomy of hierarchy: feedforward and feedback pathways in macaque visual cortex. *Journal of Comparative Neurology* 522:225–259. 7
- Mason MF, Norton MI, Van Horn JD, Wegner DM, Grafton ST, Macrae CN (2007) Wandering minds: the default network and stimulus-independent thought. *Science* 315:393–395. 2, 56, 75
- Matsui T, Murakami T, Ohki K (2016) Transient neuronal coactivations embedded in globally propagating waves underlie resting-state functional connectivity. *Proceedings of the National Academy of Sciences* 113:6556–6561. 13, 77, 78
- Matsui T, Tamura K, Koyano KW, Takeuchi D, Adachi Y, Osada T, Miyashita Y, Tamura K (2011) Direct comparison of spontaneous functional connectivity and effective connectivity measured by intracortical

- microstimulation: An fMRI study in macaque monkeys. *Cerebral Cortex* 21:2348–2356. 24
- McKeown MJ, Makeig S, Brown GG, Jung TP, Kindermann SS, Bell AJ, Sejnowski TJ (1998) Analysis of fMRI data by blind separation into independent spatial components. *Human Brain Mapping* 6:160–88. 4, 25, 50
- Messé A, Rudrauf D, Benali H, Marrelec G (2014) Relating structure and function in the human brain: relative contributions of anatomy, stationary dynamics, and non-stationarities. *PLoS Computational Biology* 10:e1003530. 8, 49, 52, 57
- Meunier D, Achard S, Morcom A, Bullmore E (2009) Age-related changes in modular organization of human brain functional networks. *NeuroImage* 44:715–723. 26
- Meunier D, Lambiotte R, Bullmore ET (2010) Modular and hierarchically modular organization of brain networks. *Frontiers in neuroscience* 4:200. 6
- Mitra A, Snyder AZ, Hacker CD, Raichle ME (2014) Lag structure in resting-state fMRI. *Journal of Neurophysiology* 111:2374–91. 8, 13, 49, 78
- Monto S, Palva S, Voipio J, Palva JM (2008) Very slow eeg fluctuations predict the dynamics of stimulus detection and oscillation amplitudes in humans. *The Journal of neuroscience* 28:8268–8272. 14
- Mueller S, Wang D, Fox MD, Yeo BTT, Sepulcre J, Sabuncu MR, Shafee R, Lu J, Liu H (2013) Individual variability in functional connectivity architecture of the human brain. *Neuron* 77:586–95. 24
- Musso F, Brinkmeyer J, Mobascher A, Warbrick T, Winterer G (2010) Spontaneous brain activity and EEG microstates. a novel EEG/fMRI analysis approach to explore resting-state networks. *NeuroImage* 52:1149–1161. 11
- Nir Y, Mukamel R, Dinstein I, Privman E, Harel M, Fisch L, Gelbard-Sagiv H, Kipervasser S, Andelman F, Neufeld MY et al. (2008) Interhemispheric correlations of slow spontaneous neuronal fluctuations revealed in human sensory cortex. *Nature neuroscience* 11:1100–1108. 1, 13
- Pan WJ, Thompson GJ, Magnuson ME, Jaeger D, Keilholz S (2013) Infralow LFP correlates to resting-state fMRI BOLD signals. *NeuroImage* 74:288–297. 13
- Phan AH, Tichavský P, Cichocki A (2013) Low complexity damped Gauss–Newton algorithms for CANDECOMP/PARAFAC. *SIAM Journal on*

- Matrix Analysis and Applications* 34:126–147. 33
- Ponce-Alvarez A, Deco G, Hagmann P, Romani GL, Mantini D, Corbetta M (2015) Resting-state temporal synchronization networks emerge from connectivity topology and heterogeneity. *PLoS Computational Biology* 11:e1004100. 21, 24, 25, 33, 50, 77
- Raichle M, MacLeod A, Snyder A, Powers W, Gusnard D, Shulman G (2001) A default mode of brain function. *Proceedings of the National Academy of Sciences* 98:676–682. xiii, 1, 2, 24, 56, 75
- Richiardi J, Altmann A, Milazzo AC, Chang C, Chakravarty MM, Banaschewski T, Barker GJ, Bokde AL, Bromberg U, Büchel C et al. (2015) Correlated gene expression supports synchronous activity in brain networks. *Science* 348:1241–1244. 6
- Ringach DL (2009) Spontaneous and driven cortical activity: implications for computation. *Current opinion in neurobiology* 19:439–444. 3, 14
- Rotarska-Jagiela A, van de Ven V, Oertel-Knöchel V, Uhlhaas PJ, Vogele K, Linden DE (2010) Resting-state functional network correlates of psychotic symptoms in schizophrenia. *Schizophrenia research* 117:21–30. 5
- Rubinov M, Sporns O (2010) Complex network measures of brain connectivity: uses and interpretations. *NeuroImage* 52:1059–1069. 6, 21
- Sadaghiani S, Hesselmann G, Friston KJ, Kleinschmidt A (2010) The relation of ongoing brain activity, evoked neural responses, and cognition. *Frontiers in systems neuroscience* 4:20. 3, 7, 14
- Sadaghiani S, Hesselmann G, Kleinschmidt A (2009) Distributed and antagonistic contributions of ongoing activity fluctuations to auditory stimulus detection. *The Journal of neuroscience* 29:13410–13417. 3
- Sakoğlu Ü, Pearlson GD, Kiehl KA, Wang YM, Michael AM, Calhoun VD (2010) A method for evaluating dynamic functional network connectivity and task-modulation: application to schizophrenia. *Magnetic Resonance Materials in Physics, Biology and Medicine* 23:351–366. 10, 14
- Sato JR, Junior EA, Takahashi DY, de Maria Felix M, Brammer MJ, Moret-tin PA (2006) A method to produce evolving functional connectivity maps during the course of an fmri experiment using wavelet-based time-varying granger causality. *NeuroImage* 31:187–196. 10
- Schirner M, Rothmeier S, Jirsa VK, McIntosh AR, Ritter P (2015) An automated pipeline for constructing personalized virtual brains from multimodal neuroimaging data. *NeuroImage* 117:343–57. 25, 26
- Schölvinck ML, Maier A, Frank QY, Duyn JH, Leopold DA (2010) Neural

- basis of global resting-state fMRI activity. *Proceedings of the National Academy of Sciences* 107:10238–10243. 1, 13
- Schoner G, Kelso J (1988) Dynamic pattern generation in behavioral and neural systems. *Science* 239:1513–1520. 9
- Seeley WW, Crawford RK, Zhou J, Miller BL, Greicius MD (2009) Neurodegenerative diseases target large-scale human brain networks. *Neuron* 62:42–52. 5
- Shen X, Tokoglu F, Papademetris X, Constable RT (2013) Groupwise whole-brain parcellation from resting-state fMRI data for network node identification. *NeuroImage* 82:403–15. 7
- Shmuel A, Leopold DA (2008) Neuronal correlates of spontaneous fluctuations in fMRI signals in monkey visual cortex: implications for functional connectivity at rest. *Human brain mapping* 29:751–761. 1, 13
- Smith SM, Fox PT, Miller KL, Glahn DC, Fox PM, Mackay CE, Filippini N, Watkins KE, Toro R, Laird AR et al. (2009) Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences* 106:13040–13045. 4, 56
- Smith SM, Miller KL, Moeller S, Xu J, Auerbach EJ, Woolrich MW, Beckmann CF, Jenkinson M, Andersson J, Glasser MF et al. (2012) Temporally-independent functional modes of spontaneous brain activity. *Proceedings of the National Academy of Sciences* 109:3131–3136. 4, 12, 20, 77
- Smith SM, Nichols TE, Vidaurre D, Winkler AM, Behrens TE, Glasser MF, Ugurbil K, Barch DM, Van Essen DC, Miller KL (2015) A positive-negative mode of population covariation links brain connectivity, demographics and behavior. *Nature neuroscience* 18:1565–1567. 3
- Sorg C, Riedl V, Mühlau M, Calhoun VD, Eichele T, Läer L, Drzezga A, Förstl H, Kurz A, Zimmer C et al. (2007) Selective changes of resting-state networks in individuals at risk for Alzheimer's disease. *Proceedings of the National Academy of Sciences* 104:18760–18765. 5
- Sporns O, Betzel RF (2016) Modular brain networks. *Annual review of psychology* 67:613. 11
- Sporns O, Chialvo DR, Kaiser M, Hilgetag CC (2004) Organization, development and function of complex brain networks. *Trends in cognitive sciences* 8:418–425. xiii, 6, 7, 21
- Sporns O, Tononi G, Kötter R (2005) The human connectome: a structural description of the human brain. *PLoS computational biology* 1:e42. 6, 7
- Tagliazucchi E, Balenzuela P, Fraiman D, Chialvo DR (2012) Criticality

- in large-scale brain fMRI dynamics unveiled by a novel point process analysis. *Frontiers in Physiology* 3:15. 9, 10, 21, 48, 57, 62
- Tognoli E, Kelso JS (2014) The metastable brain. *Neuron* 81:35–48. 9
- Tsodyks M, Kenet T, Grinvald A, Arieli A (1999) Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 286:1943–1946. 13
- Tzourio-Mazoyer N, Landeau B, Papathanassiou D, Crivello F, Etard O, Delcroix N, Mazoyer B, Joliot M et al. (2002) Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage* 15:273–289. 7
- Van de Ville D, Britz J, Michel CM (2010) EEG microstate sequences in healthy humans at rest reveal scale-free dynamics. *Proceedings of the National Academy of Sciences* 107:18179–18184. 11
- Van den Heuvel MP, Hulshoff Pol HE (2010) Exploring the brain network: a review on resting-state fMRI functional connectivity. *European neuropsychopharmacology* 20:519–34. 50
- Van den Heuvel MP, Sporns O (2011) Rich-club organization of the human connectome. *Journal of Neuroscience* 31:15775–15786. 24
- Van Dijk KR, Hedden T, Venkataraman A, Evans KC, Lazar SW, Buckner RL (2010) Intrinsic functional connectivity as a tool for human connectomics: theory, properties, and optimization. *Journal of neurophysiology* 103:297–321. 2, 7
- Van Essen DC, Smith SM, Barch DM, Behrens TE, Yacoub E, Ugurbil K, Consortium WMH et al. (2013) The WU-Minn human connectome project: an overview. *NeuroImage* 80:62–79. 7
- Varela F, Lachaux JP, Rodriguez E, Martinerie J (2001) The brainweb: phase synchronization and large-scale integration. *Nature reviews. Neuroscience* 2:229–39. 14
- Veer IM, Beckmann C, Van Tol MJ, Ferrarini L, Milles J, Veltman D, Aleman A, Van Buchem MA, Van Der Wee NJ, Rombouts SA (2010) Whole brain resting-state analysis reveals decreased functional connectivity in major depression. *Frontiers in systems neuroscience* 4:41. 5
- Vincent J, Patel G, Fox M, Snyder A, Baker J, Van Essen D, Zempel J, Snyder L, Corbetta M, Raichle M (2007) Intrinsic functional architecture in the anaesthetized monkey brain. *Nature* 447:83–86. 2, 13, 24, 56
- Vincent JL, Snyder AZ, Fox MD, Shannon BJ, Andrews JR, Raichle ME, Buckner RL (2006) Coherent spontaneous activity identifies a hippocampal-parietal memory network. *Journal of neurophysiol-*

- ogy* 96:3517–3531. 2
- Wang XJ, Kennedy H (2016) Brain structure and dynamics across scales: in search of rules. *Current opinion in neurobiology* 37:92–98. 7
- Wang Z, Chen LM, Négyessy L, Friedman RM, Mishra A, Gore JC, Roe AW (2013) The relationship of anatomical and functional connectivity to resting-state connectivity in primate somatosensory cortex. *Neuron* 78:1116–1126. 24
- Watts D, Strogatz S (1998) Collective dynamics of ‘small-world’ networks. *Nature* 393:440–442. 6
- Wong KF, Wang XJ (2006) A recurrent network mechanism of time integration in perceptual decisions. *The Journal of Neuroscience* 26:1314–1328. 25, 26, 48
- Yacoub E, Harel N, Uğurbil K (2008) High-field fMRI unveils orientation columns in humans. *Proceedings of the National Academy of Sciences* 105:10607–10612. 12
- Yeo BT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zöllei L, Polimeni JR et al. (2011) The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of neurophysiology* 106:1125–1165. 4, 56
- Yu-Feng Z, Yong H, Chao-Zhe Z, Qing-Jiu C, Man-Qiu S, Meng L, Li-Xia T, Tian-Zi J, Yu-Feng W (2007) Altered baseline brain activity in children with ADHD revealed by resting-state functional MRI. *Brain and Development* 29:83–91. 5
- Zalesky A, Fornito A, Cocchi L, Gollo LL, Breakspear M (2014) Time-resolved resting-state brain networks. *Proceedings of the National Academy of Sciences* 111:10341–10346. 11, 14, 56, 68