Stereotypical modulations in dynamic functional connectivity explained by changes in BOLD variance

Katharina Glomb[1,⋆], Adrián Ponce-Alvarez[1], Matthieu Gilson[1], Petra Ritter[2,3,4,5], and Gustavo Deco[1,6]

1 Center for Brain and Cognition, Dept. of Technology and Information, Universitat Pompeu Fabra, Carrer Tànger, 122-140, 08018 Barcelona, Spain
2 Minerva Research Group Brain Modes, Max Planck Institute for Human Cognitive and Brain Sciences, Stephanstrasse 1, 04103 Leipzig, Germany
3 Dept. of Neurology, Charité - University Medicine, Charitéplatz 1, 10117 Berlin, Germany
4 Bernstein Focus State Dependencies of Learning & Bernstein Center for Computational Neuroscience, Philippstrasse 12, 10115 Berlin, Germany
5 Berlin School of Mind and Brain & Mind and Brain Institute, Humboldt University, Luisenstrasse 56, 10117 Berlin, Germany
6 Institució Catalana de la Recerca i Estudis Avançats, Universitat Barcelona, Passeig Lluís Companys 23, 08010 Barcelona, Spain
⋆ katharina.glomb@upf.edu

Keywords: fMRI, human, functional connectivity, dynamic functional connectivity, tensor decomposition, feature extraction, mean field models, whole-brain models

Conflict of interest: The authors declare no competing financial interests.

Funding (authors' initials given after grant numbers):
This work was supported by the European Union, FP7 Marie Curie ITN "INDIREA" (Grant N. 606901; KG), FP7 FET ICT Flagship Human Brain Project (Grant N. 604102; MG), ERC Advanced Human Brain Project (Grant N. 604102; GD), Horizon2020 ERC Consolidator (Grant N. ...; PR);
the Spanish Ministry for Economy, Industry and Competitiveness (MINECO) project "PIRE-PICCS" (Grant N. PCIN-2015-079), SEMAINE ERA-Net NEURON Project (Grant N. PCIN2013-026; APA), and ICoBAM (Grant N. PSI2013-42091-P; GD);
the James S. McDonnell Foundation (Brain Network Recovery Group, Grant N. JSMF22002082; PR);
the German Ministry of Education and Research (Grant N. 01GQ1504A and 01GQ0971-5; PR);
the Max-Planck Society (Minerva Program; PR)

Abstract

Spontaneous activity measured in human subject under the absence of any task exhibits complex patterns of correlation that largely correspond to large-scale functional topographies obtained with a wide variety of cognitive and perceptual tasks. These "resting state networks" (RSNs) fluctuate over time, forming and dissolving on the scale of seconds to minutes. While these fluctuations, most prominently those of the default mode network, have been linked to cognitive function, it remains unclear whether they result from random noise or whether they index a non-stationary process which could be described as state switching.

In this study, we use a sliding windows approach to relate temporal dynamics of RSNs to global modulations in correlation and BOLD variance. We compare empirical data, phase-randomized surrogate data, and data simulated with a stationary model. We find that RSN time courses exhibit a large amount of coactivation in all three cases, and that the modulations in their activity are closely linked to global dynamics of the underlying BOLD signal.

We find that many properties of the observed fluctuations in FC and BOLD, including their ranges and their correlations amongst each other, are explained by fluctuations around the average FC structure. However, we also encounter interesting characteristics that are not explained in this way. In particular, we find that the brain spends more time in the peaks and troughs of modulations than can be expected from stationary dynamics.

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" (RSNs), enables us to probe integration (within networks) and segregation (between networks) in the brain.

More recent research has shown that the FC between brain regions forming the basis of RSNs are not constant over time. Dynamic functional connectivity (dFC) is defined in contrast to average or "static" functional connectivity (avFC), and dFC often appear to greatly deviate from it. Changes in the global FC structure apparent in dFC have been taken to be signs of time-dependent integration and segregation between functional networks (see Hutchison et al. (2013); Calhoun et al. (2014); Preti et al. (2016) for reviews). A growing number of studies are proposing and studying new and innovative approaches to characterizing dFC (Chang and Glover, 2010; Britz et al., 2010; Kiviniemi et al., 2011; Hutchison et al., 2012; Allen et al., 2012; Leonardi and Van de Ville, 2013a; Liu et al., 2013; Zalesky et al., 2014; Damaraju et al., 2014; Karahanoğlu and Van De Ville, 2015; Betzel et al., 2016a). Many of these studies have focused on characterizing recurring "states" of brain connectivity. 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.

Since the beginning of RS research, it was asked whether the observed fluctuations are "meaningful" or not. This question is tightly related to the question of what the mechanism behind the fluctuations is. It has already been recognized that several different sources likely contribute to the variability observed in dynamic FC (dFC) (Messé et al., 2014; Barttfeld et al., 2015): First, the observed fluctuations can be shown to be relevant to behavior (Fox et al., 2005; Bassett et al., 2011; Kucyi and Davis, 2014; Yang et al., 2014; Davison et al., 2015) which has recently lead to the idea of "controllability", i.e. that the large-scale architecture of the human brain is organized in such a way that frequently visited states are reachable with minimal effort and maximal energy efficiency (Gu et al., 2015). Thus, observed dynamics are partially linked to ongoing cognition and are implied to arise from temporal dynamics on a neural level (He and Raichle, 2009; Kopell et al., 2014). Second, a certain part of the fluctuations is related to changes in brain state, like arousal (Chang et al., 2013), sleep state (Larson-Prior et al., 2009; Horovitz et al., 2008), or aneasthesia (Vincent et al., 2007; Bettinardi et al., 2015; Hutchison et al., 2013). Third, recent work shows a link between ongoing fluctuations in the brain and rhythms generated in internal organs like the heart and stomach, sug-

gesting that self-referential behavior related to the default mode network includes the body as well as the mind (Babo-Rebelo et al., 2016; Richter et al., 2016; Babo-Rebelo et al., 2016). Fourth, a part of the fluctuations arises from random noise reverberating through a structured network (Deco and Jirsa, 2012; Messé et al., 2014; Ponce-Alvarez et al., 2015; Glomb et al., 2016). Nonetheless, this last source of variability could be functionally relevant to the maintenance of synaptic connections. Furthermore, these reverberations constitue a large part of the broadband-signal (He, 2014; Chaudhuri et al., 2016), and only recently are receiving increased attention.

Evidently, the mechanisms behind dFC are manifold and require careful interpretation of data as well as integration of results from many lines of inquiry and different methods (Ritter et al., 2013). Two tools that whose importance cannot be stressed enough in this context are proper statistical assessments on the one hand (Handwerker et al., 2012; Bassett et al., 2013; Lindquist et al., 2014; Hindriks et al., 2015; Betzel et al., 2016a; Laumann et al., 2016), and modelling approaches on the other (Messé et al., 2014; Hansen et al., 2014; Glomb et al., 2016). While the former is necessary to answer the question whether observed fluctuations contain hallmarks of non-stationarity, the latter is crucial to our understanding of underlying mechanisms and provides the interpretation of the observed variability as a "dynamic repertoire" of the brain (Deco et al., 2013). 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.

The temporal dynamics of modulations present in time courses of spontaneous brain activity measured with fMRI BOLD are the focus of the present study. We will first characterize these dynamics on the level of RSNs and the underlying BOLD signal, and then investigate in how far our findings are explained by stationary dynamics. To this end, we will use both stationary surrogate data and a stationary model.

2 Methods

2.1 Data

We use data previously published in Schirner et al. (2015), recorded at the Charite University Medicine in Berlin, Germany. Three datasets are recorded from each of 24 healthy subject aged 18 to 35 years: 1) 22-minute ($T = 661$ frames, TR=2s) RS scans (eyes closed), 2) diffusion-weighted (dw) MRI, 3) anatomical scans (T1-weighted). The latter are used for parcellation and for creating masks for fiber tracking. We use the Desikan-Killiany atlas (Desikan et al., 2006) to parcellate our data into large-scale brain regions; discarding the Corpus Callosum on both sides, we are left with $N = 66$ cortical areas.

DwMRI was performed with 64 directions. For each voxel, the fiber orientation distribution function was computed, which was then used by the probabilistic fiber tracking algorithm implemented in MRTrix (Tournier et al., 2004, 2007) to initiate tracts in the grey matter-white matter interface and prolong them until they reach a termination region. When constructing the structural connectivity matrices from the tracking output, assumptions about limitations on connection strength were used to improve upon several shortcomings of fiber tracking algorithms. For example, connection strengths were assumed to be bounded by the area available at the grey matter-white matter interface (see Schirner et al. (2015) for details). The number of streamlines detected per connection was then used as input to an algorithm which derives effective connectivity from structural connectivity using a noise diffusion model (see section 2.5).

2.2 RSN time courses

We use tensor decomposition, a general dimensionality reduction approach which has previously been shown to be useful in neuroscience (Beckmann and Smith, 2005; Cichocki, 2013; Leonardi and Van de Ville, 2013a; Leonardi and Van De Ville, 2013b; Ponce-Alvarez et al., 2015) and which we apply to community detection (Gauvin et al., 2014); in our case, the communities we wish to detect are recurring FC patterns. For each subject (or simulation), we compute pair-wise FC inside of sliding windows of width 120 s (60 frames) which are moved in increments of 1 frame (figure 2A), resulting in $W = 601$ time windows. This results in an $N \times N$ dynamic FC matrix, $dFC(w)$ with $w = 1, 2, ..., W$ which is symmetrical. The W dFC-matrices are then stacked along the temporal dimension into 3-way data structures of dimensions $N \times N \times W$, called tensors (figure 2B), denoted by $\underline{\mathbf{Y}}$. We use a special case of tensor decomposition, non-negative canonical polyadic decomposition (NCP). In order to do so, we use a non-negative measure of dFC, i.e. mutual information as described in Kraskov et al. (2004). We obtain a number F of features for each of these tensors (figure 2C, for $F = 3$). Each feature i is a rank-1-tensor, i.e. it can be represented as the outer product, denoted as $\circ$, between three vectors, each vector corresponding to one dimension of the tensor:

$$\begin{aligned} \underline{\mathbf{Y}} &= \sum_{i=1}^{F} u^i \circ v^i \circ w^i + \underline{\mathbf{E}} = \underline{\hat{\mathbf{Y}}} + \underline{\mathbf{E}}, \\ u^i &= v^i \end{aligned} \tag{1}$$

Since the dFC matrices are symmetric, the first two vectors (of length N) are identical ($u^i = v^i$) and are interpreted as containing membership weights. These weights describe how much each of the 66 brain regions is a member of community i. The third vector w^i (of length W) is the time course which describes how much this community as a whole contributes to the dFC in each window. Of course, the decomposition which we obtain is only an approximation, leaving the variability contained in the error $\underline{\mathbf{E}}$ unexplained.

In this study, we are mostly interested in the time courses w^i. Each of the community vectors v^i is identified as a known RSN such that each time course can be associated to a certain RSN. For this, we use templates obtained from our data, i.e. we look for communities that are valid for the whole group. We cluster the v^i obtained from all subjects into K clusters using K-means-clustering, using different numbers of F (but always the same F for all subjects). We assess the goodness of clustering via the silhouette value (de Amorim and Hennig, 2015), comparing to phase-randomized surrogate data (section 2.4) which do not contain a cluster structure. Additionally, we remove noise be keeping only the 2% highest dFC values and binarizing the resulting tensor (Glomb et al., 2016). We choose the values of F and K such that the difference between the sihouette values for surrogates and real data is maximal. This is the case for $F = 3$ and $K = 4$. Then, the cluster centers are used as templates.

Thus, for empirical data, the RSN identity is its cluster membership. For surrogate and simulated data, we use Cohen's kappa to compute how similar extracted features are to the templates (comparing each community vector to each template). This measure relies on computing confusion matrices, and we use a quantization on three levels. Then, the κ-value is computed as

$$\kappa = \frac{P_a - P_e}{1 - P_e}$$

P_a is the overlap, P_e is the expected overlap. Each feature is assigned to a template according to its maximum κ-value.

2.3 Measures of RS dynamics

We use three measures to quantify global dynamics of RS. 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, using the same sliding windows as for the tensors (width 120s, i.e. 60 frames, increments of one frame, giving $W = 601$ windows per measure per subject) in order to be able to relate the resulting time courses to the RSN time courses.

The FC-related measures rely on pairwise correlations, i.e. as for the tensors (section 2.2), we compute dFC-matrices for each window, containing pair-wise FC values, but this time, using correlation instead of mutual information. Again, the dFC matrices are symmetrical, and their diagonals are filled with ones. The *instantaneous average correlation* is just the average over all unique pairs, i.e. using the upper or lower triangle of the dFC(w) matrix, excluding the diagonal:

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

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. In other words, this is the average FC present in each slice of a tensor (without thresholding).

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. We call this measure *similarity to avFC*. This is done by simply correlating the upper or lower triangle of the matrices with each other after flattening them into a vector.

$$\text{sim to avFC}_s(w) = \text{corr}\left\{\mathbf{vec}\left[\mathbf{U}(\tanh(dFC(w)))\right], \mathbf{vec}\left[\mathbf{U}(\tanh(avFC))\right]\right\} \tag{3}$$

$\mathbf{vec}(\cdot)$ stands for the vectorization of a matrix, $\mathbf{U}$ is the upper triangle of a matrix, and 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 $t + 59$ (in frames) and average over ROIs, resulting in a measure for the mean variance of the

signal, which we refer to as *instantaneous BOLD variance*.

$$\text{BOLD var}_s(w) = \frac{1}{N}\sum_{i=1}^{N} \text{var}_{t:t+59}(\mathbf{x}_i) \tag{4}$$

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

2.4 Surrogates

Surrogate data are constructed by randomizing an aspect of the empirical data with the goal to produce data that preserve all characteristics except the one of interest. In this case, this characteristic is nonstationarity, corresponding to the null hypothesis that observed dynamics are a result of the long-term correlations and power spectrum, but that no additional structure is present. 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 $x_1, x_2, ..., x_T$ and $y_1, y_2, ..., y_T$, 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 $X_1, X_2, ..., X_N$ and $Y_1, Y_2, ..., Y_N$ in frequency space, a random phase vector $\varphi_{r,1}, \varphi_{r,2}, ..., \varphi_{r,N}$, 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.

2.5 Model and simulations

We use a dynamic mean-field (DMF) model (Wong and Wang, 2006; Deco et al., 2014) where each of the 66 cortical regions are modelled by a node. Each node i contains one excitatory and one inhibitory pool which are described by their population firing rates r_i, currents I_i, and synaptic

gating variables S_i according to the following coupled differential equations:

$$
\begin{aligned}
I_i^{(E)} &= I_{0,E} + w_{\mathrm{EE}} S_i^{(E)} + G \sum_j C_{ij} S_j^{(E)} - w_{\mathrm{EI},i} S_i^{(I)} \\
I_i^{(I)} &= I_{0,I} + w_{\mathrm{IE}} S_i^{(E)} - w_{\mathrm{II}} S_i^{(I)} \\
r_i^{(E)} &= H^{(E)}(I_i^{(E)}) = \frac{a_E I_i^{(E)} - b_E}{1 - \exp(-d_E(a_E I_i^{(E)} - b_E))} \\
r_i^{(I)} &= H^{(I)}(I_i^{(I)}) = \frac{a_I I_i^{(I)} - b_I}{1 - \exp(-d_I(a_I I_i^{(I)} - b_I))} \\
\frac{\mathrm{d}S_i^{(E)}(t)}{\mathrm{d}t} &= -\frac{S_i^{(E)}}{\tau_E} + (1 - S_i^{(E)})\gamma r_i^{(E)} + \sigma \nu_i(t) \\
\frac{\mathrm{d}S_i^{(I)}(t)}{\mathrm{d}t} &= -\frac{S_i^{(I)}}{\tau_I} + r_i^{(I)} + \sigma \nu_i(t)
\end{aligned}
\tag{5}
$$

The model is illustrated in figure 1A. Locally, excitatory (indexed by E) and inhibitory (indexed by I) pools are connected to each other, using weights w_{EE}, $w_{\mathrm{EI},i}$, w_{IE}, and w_{II}. Most of these weights are constant ($w_{EE} = 0.21$, $w_{II} = 1$, $w_{IE} = 0.15$), except for the feedback inhibition $w_{\mathrm{EI},i}$ which is adjusted prior to running simulations in order to keep mean firing rates in each population between 3 and 10 Hz. This is necessary because in this asynchronous regime, the model only has one attractor and is thus stationary (Deco and Jirsa, 2012). This model includes three types of synapses: 1) NMDA for local excitatory connections with $\tau_E = 100$ ms, 2) GABA for local inhibitory connections with $\tau_I = 10$ ms, and 3) AMPA for long-range connections between excitatory pools, governed by the connectivity matrix $\mathbf{C}$ with entries C_{ij} setting the weights between nodes i and j. The dynamics of these synapses are not modelled directly because they reach their steady state almost instantaneously (Deco et al., 2014). Additionally, each population receives a constant background current, $I_{0,E} = 0.382$ nA and $I_{0,I} = 0.267$ nA. C_{ij} are entries in the adjacency matrix governing the weights between the nodes. These weights are scaled by the factor G, which is a global coupling parameter, and the only free parameter of the model.

Here, $\mathbf{C}$ is the effective connectivity (EC) matrix, shown in figure 1B, derived by combining DTI-based structural connectivity (SC) with a noise diffusion (ND) model (Gilson et al., 2016). The weights in the EC are tuned such that time-shifted covariance matrices are reproduced when simulating the ND model. We previously showed (Glomb et al., 2016) that, when the nodes in the DMF model are connected according to this matrix, communities can be extracted from the resulting dFC time courses that approximate the RSNs found in empirical data. The EC matrix is asymmetric, contains homotopic connections, and exhibits a nearly uniform degree distribution which promotes stable simulations over a wide range of values of G. The results we show here are obtained by simulating with the value of G which leads to the maximum overlap with RSNs obtained from empirical data (see section 2.2).

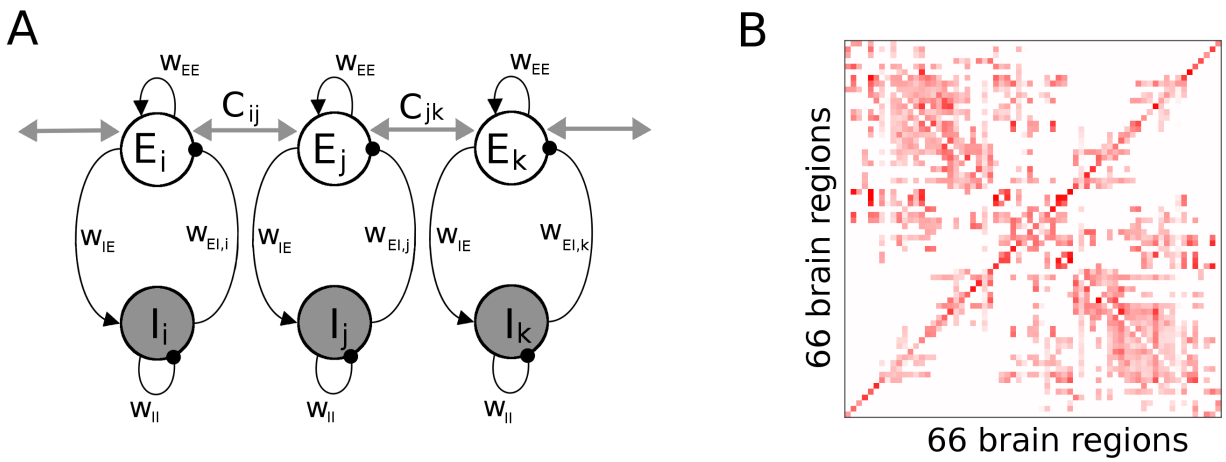

Figure 1: **A** Illustration of DMF model used to produce stationary simulated data. Each node contains an excitatory (E) and an inhibitory (I) pool which are locally connected by GABA (black lines, spheres) and NMDA (black lines with arrows) synapses and E pools have long range AMPA connections (gray lines) set by entries C_{ij} of the connectivity matrix. **B** Effective connectivity matrix. Note the asymmetry and the even distribution of the weight magnitudes. Weights are in arbitrary units.

3 Results

3.1 Resting state networks coactivate strongly

One of the most prominent approaches is to measure pairwise functional connectivity (FC) inside of overlapping sliding time windows, thus obtaining a sequence of dFC matrices (dFCMs, figure 2A, B).

We were interested in shedding light on the basis of these fluctuations, given that it remains unclear whether RSNs are separate in time (Baker et al., 2014; Betzel et al., 2016a) or not (Mitra et al., 2014; Karahanoğlu and Van De Ville, 2015; Glomb et al., 2016).

Figure 5 illustrates these two scenarios schematically.

We use 22 minute long resting state (RS) scans from $S = 24$ healthy participants (Schirner et al., 2015), parcellated according to the Desikan-Killiany atlas (Desikan et al., 2006), obtaining $N = 66$ large-scale brain regions covering the entire cortex. We use nonnegative canonical polyadic decomposition (NCP), a highly interpretable version of tensor decomposition (Cichocki, 2013; Gauvin et al., 2014; Ponce-Alvarez et al., 2015) to obtain RSN time courses. This approach approximates a sequence of dynamic FC matrices (dFCMs) as a linear superposition of a small number of static FC patterns. The patterns' weights and thus, the degree to which they contribute to dFC, is time-dependent (figure 2). In contrast to the most frequently used approach, i.e. independent component analysis (ICA), these networks are allowed to overlap in both space and time. We obtain a small number of networks that generalize across subjects and find that they resemble RSNs (figure 3): default mode network, somatomotor network, control networks, visual network.

Their associated time courses (figure 4) are interpreted as activation time courses. Our first observation is that the activations fluctuate considerably, and that these fluctuations seem to be strongly modulated on an infraslow time scale of 50 to 100 seconds. However, what is not immediately clear is to what degree the RSNs overlap in time. In order to quantify their coactivation, we consider the "contribution" of each RSN to the overall activation in each time window, which is just the ratio between its activation and the sum of all activations in this window. If this value is close to 1, this

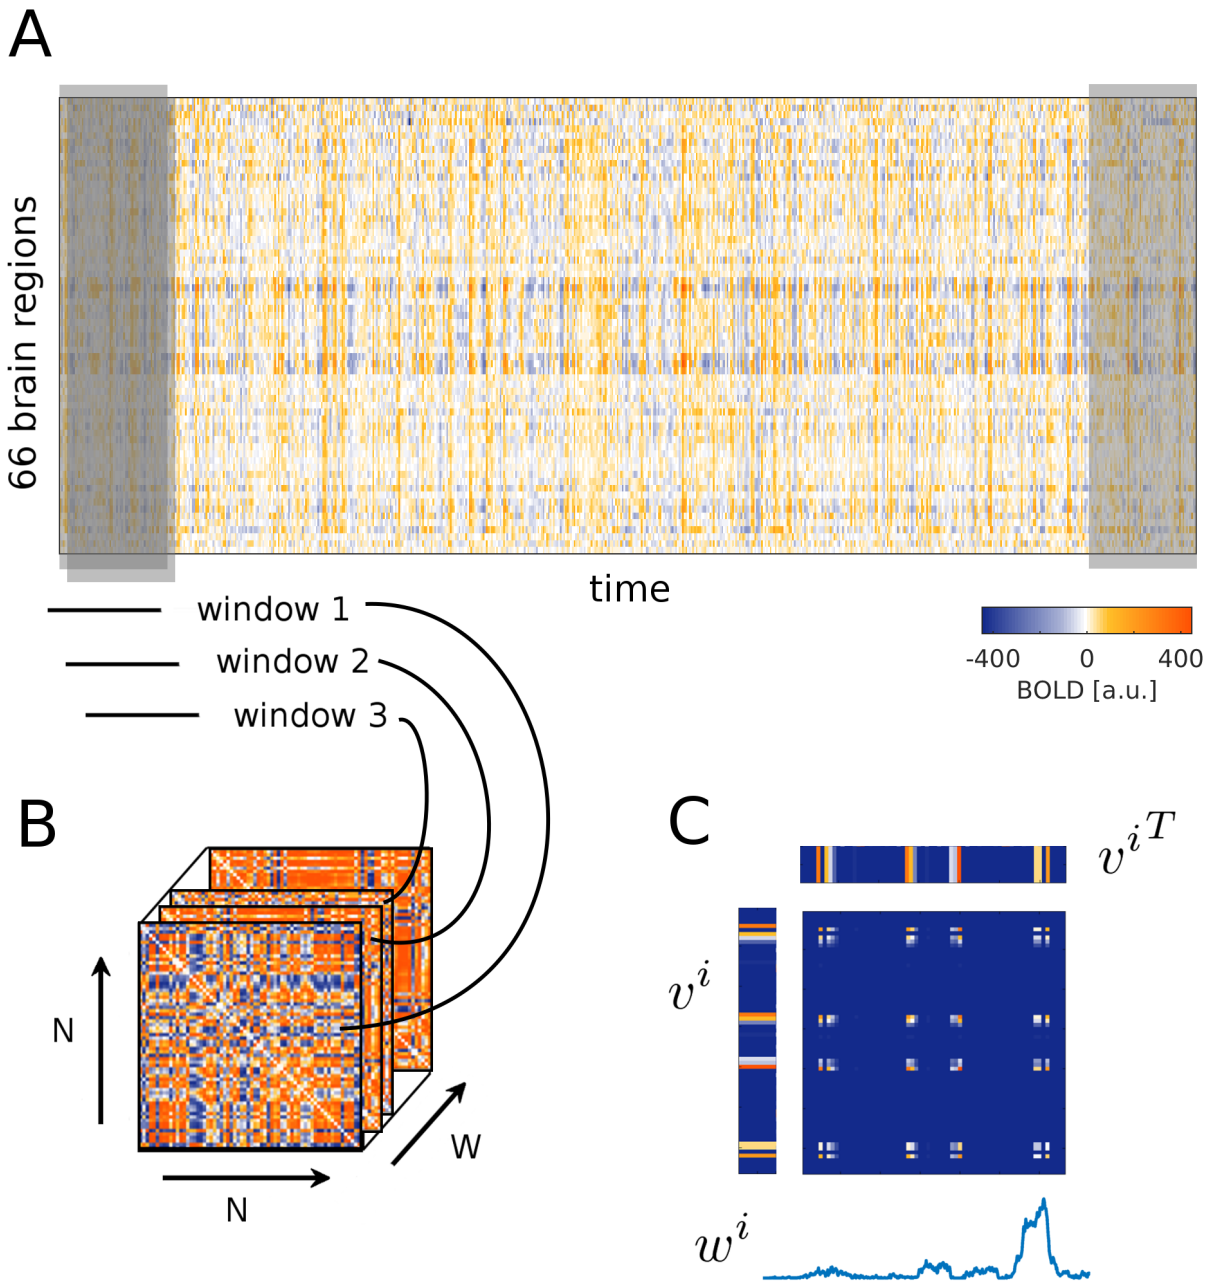

Figure 2: Illustration of sliding windows dynamic FC and tensor decomposition. **A** Example time courses for one subject. The slliding window are illustrated in grey: it is moved alond the entire time course, resulting in W windows, each of which covers a portion of the BOLD time courses. **B** The data points that fall into any given window are used to compute a dynamic FC matrix. DFC matrices from all W windows are stacked to form the tensor. **C** The tensor is decomposed into a given number F of sets of three vectors. Each set i contains two spatial features v^i, i.e., the communities; they are identical due to the symmetry of the dFC-matrices. The third vector w^i is the time course of weights associated with this community. Taking the outer product $v^i \circ v^{i^T}$ results again in a matrix. Any given dFC matrix is approximated by adding all F matrices given in this way for the tensor, weighed by the value $w^i(t)$.

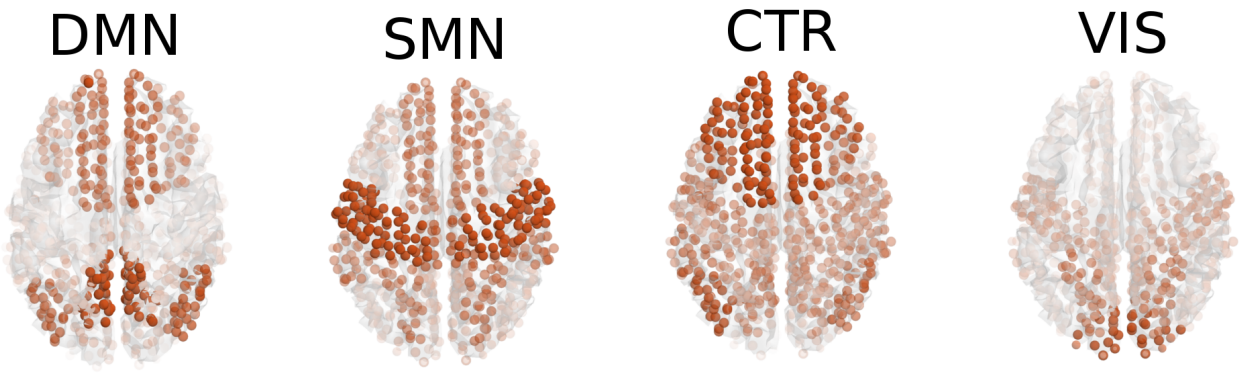

Figure 3: Community extraction from empirical data leads to recovery of RSNs. Each subject's tensor is decomposed into three features, and the three community vectors are pooled across all subjects. Pooled community vectors are clustered. Using the cluster centers as templates and projecting them back to the cortex surface, we recover four known RSNs: DMN - default mode network, SMN - sensorimotor network, CTR - control network, VIS - visual network.

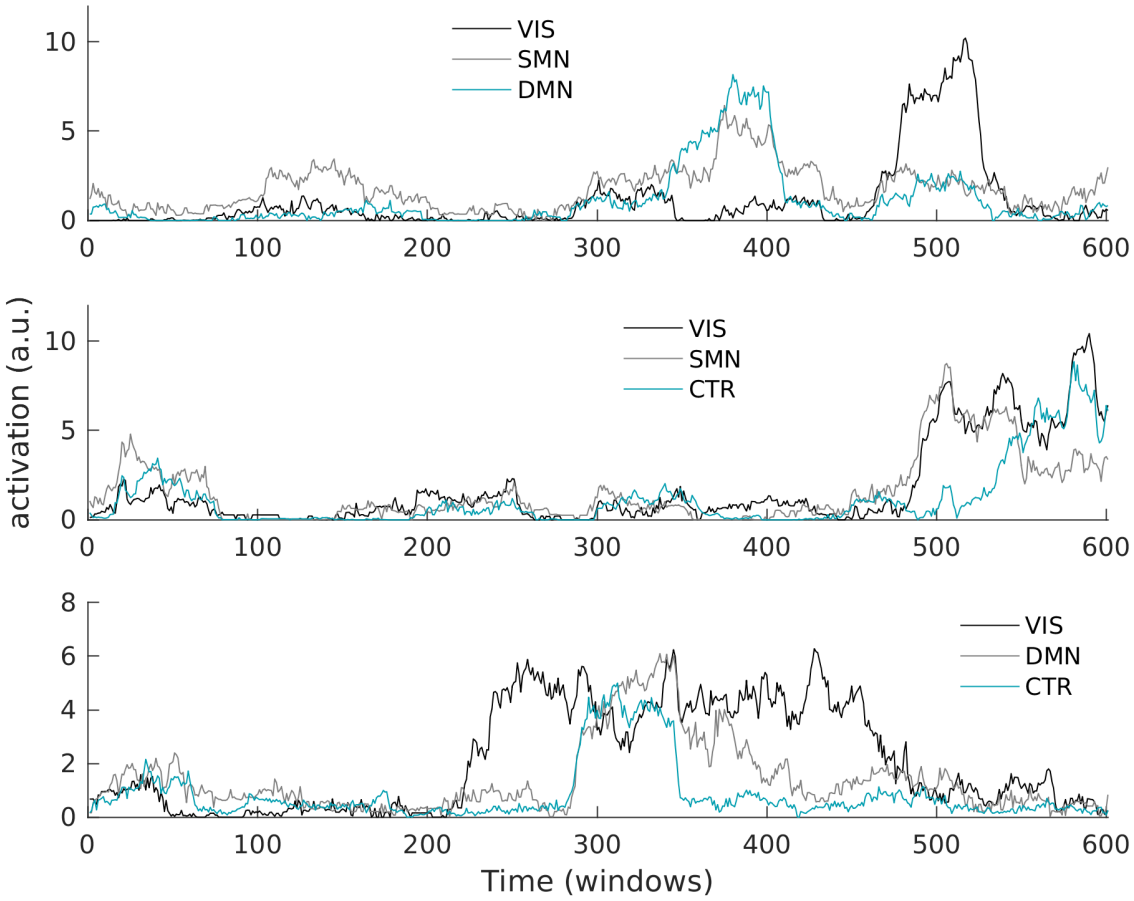

Figure 4: Example time courses of RSNs for three subjects. The time courses are obtained from the tensor decomposition and describe the degree to which each network contributes to the dFCM in each window and thus, how much it is activated.

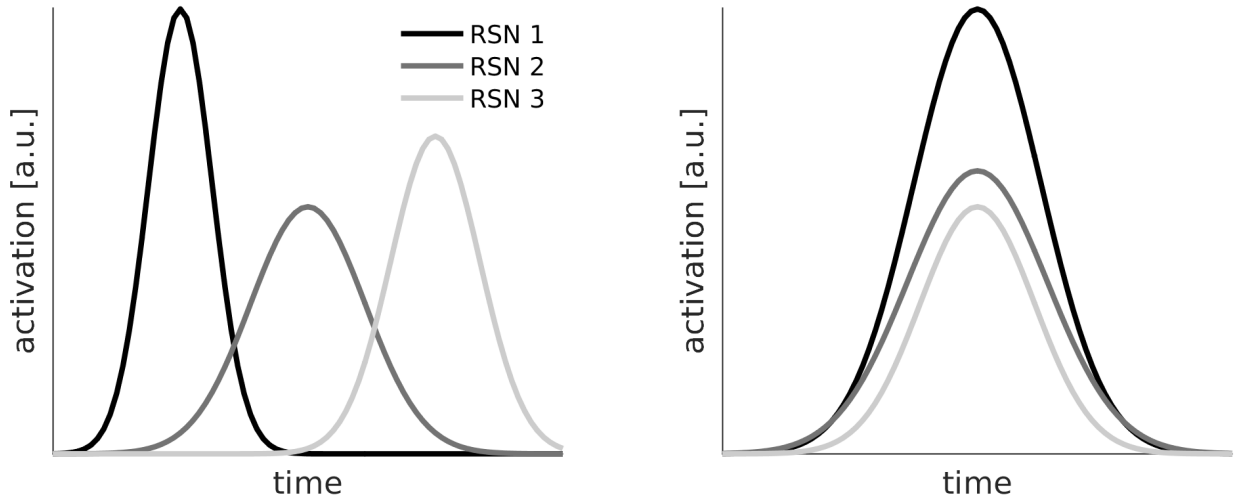

Figure 5: Illustration of two alternative scenarios for RSN activation time courses. (a) RSNs could activate in sequence. At any point in time, one RSN's weight would be much higher than the other two. (b) RSNs could coactivate. At any point in time, combinations of RSNs would be present.

network dominates in the given time window. If the value is close to $1/F$, i.e. one third in this case, there is a large degree of co-activation; and if the value is very small, the network is not active at all.

Figure 6A shows the contribution values across windows for each RSN (DMN, somatomotor, control, visual), pooled across subjects. The distributions show a prominent peak around $1/3$, a value which corresponds to perfect coactivation in the case of having three RSNs per subject. There are also many zeroes, 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 coactivated or not active at all and away from single RSNs being dominant. This prevalence of coactivation clearly contradicts the notion of independent networks.

3.2 Modulations in RSN activation driven by underlying BOLD

In order to explain the observation that RSNs coactivate strongly, we take a step back and look at the underlying BOLD signal. Considering the example shown in figure 2A, one immediately notices "bands" or "events" as many, if not all, ROIs increase or decrease their BOLD activity simultaneously. It is fair to assume that these strong modulations will have an effect on the overall FC strength (Chawla et al., 1999; Daffertshofer and van Wijk, 2011) and eventually, on RSN dynamics themselves. We quantify this with two simple global measures. First, we take the variance of each ROI's BOLD signal, averaged over all ROIs, inside each window. In time windows when all ROIs increase or decrease their activity together, this value will be high. We term this the *instantaneous BOLD variance* (figure 7, orange). Second, we take the average over all pairwise correlations in each window (figure 7, blue) and term this *instantaneous average correlation*.

It is immediately visible that these two measures are closely correlated. Computing the correlation coefficient for all subjects, we obtain an average of 0.81, suggesting that there is an immediate need for methods that will disentangle the effects of variations in signal strength on FC. Taking this a step further, we also correlate the instantaneous BOLD variance with the previously extracted RSN

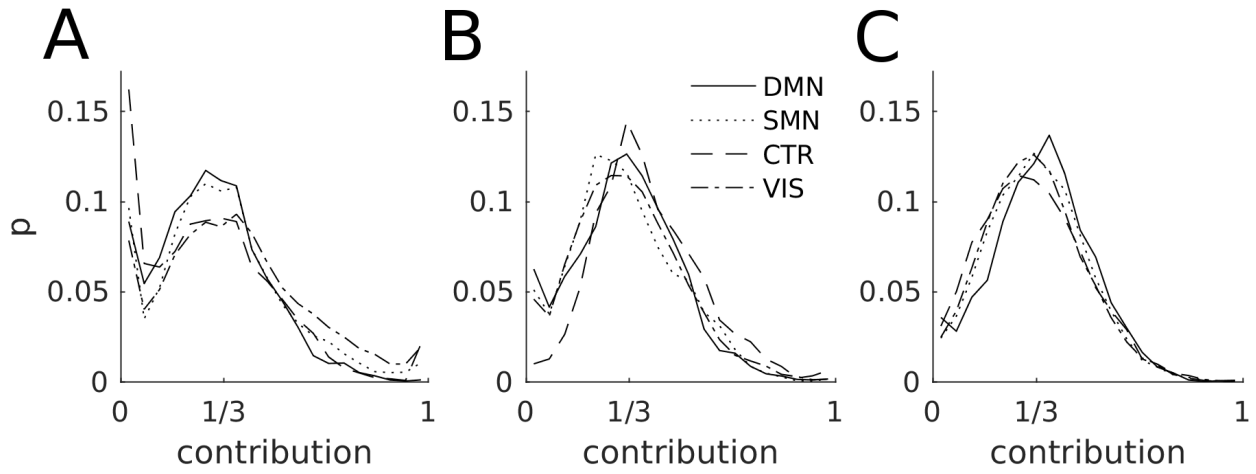

Figure 6: Contribution histogram for each RSN, pooled across subjects. **A** Empirical data, **B** surrogate data, **C** simulated data. Since there are three features per subject, a value of $1/3$ means perfect coactivation between the three RSNs. Higher values correspond to dominance of one RSN over the others. Values close to 0 indicate deactivation. The large peak around $1/3$ clearly favors the coactivation scenario in figure 5. DMN: Default mode network, SMN: somatomotor network, CTR: right and left control networks, VIS: visual network

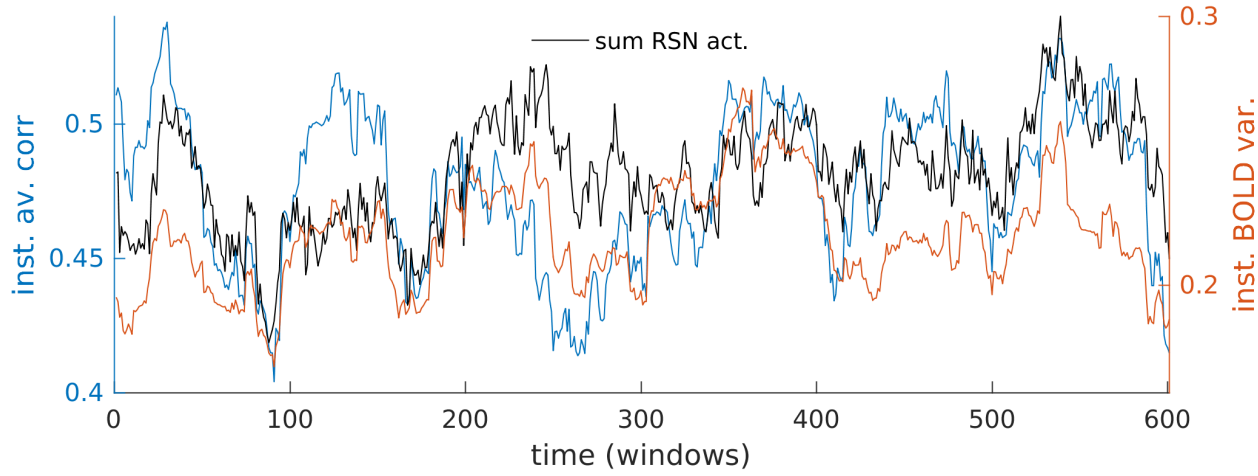

Figure 7: Global modulations in underlying BOLD signal (blue and orange curves) is tightly correlated with RSN dynamics (grey curve). Orange: *instantaneous BOLD variance*, blue: *instantaneous average correlation*, grey: sum over all RSN activations.

time courses. Again, we find a correlation of 0.81 across all subjects, indicating that variations in RSN activation (and coactivation) may to a large part be driven by changes in the variance of the underlying BOLD signal.

3.3 During peaks, the brain is closer to average FC structure

We asked the question what influence the modulations in instantaneous BOLD variance have on FC. We saw that these modulations are closely related to modulations in instantaneous average correlation. What is more, RSN activations are modulated in the same way. This suggests that windows of high BOLD variance are to be interpreted as windows in which the signal to noise-ratio (SNR) is high; alternatively, one could think that it might be low, leading to an unstructured all-to-all dFCM. To illustrate further that the former is the case rather than the latter, we employ a third global measure of temporal dynamics: We measure, for each time window, the similarity of the dFCM to the avFCM by taking the correlation between the two. We expect that the dFCM should be more similar to the avFCM when the other two measures (instantaneous BOLD variance and -average correlation) and also the RSN activations are high, and vice versa. Indeed, this is the case for almost all subjects, however, the correlation values are not as high as might be expected and are not significant for all subjects (averages correlations to sim. to avFC across subjects: inst. BOLD var.: 0.53, significant for 21/24 subjects; inst. av. corr: 0.57, 22/24; sum of RSN activities: 0.49, 22/24).

In order to further explore this, we plot in the left panel of figure 8 the joint distribution of instantaneous BOLD variance and similarity to avFC, averaged over all subjects. It is evident that this relationship is not linear. Furthermore, there seem to be very many time windows in which the instantaneous BOLD variance and the similarity to avFC are both very low (prominent cluster at the bottom left). The same, although a bit less strongly, goes for the opposite end of both ranges (cluster at the top right). In the next step, we will use surrogate data to test whether this observations really hints at nonstationarity, as could be suspected.

3.4 Stationary dynamics can explain a lot, but not everything

It has been suggested that the modulations apparent in dFC are well within what is predicted by a multivariate Gaussian process, adjusted for the autocorrelations inherent in the BOLD signal (Handwerker et al., 2012; Hindriks et al., 2015). This is equivalent to the null hypothesis that observed modulations can be fully explained by the presence of long-term correlations and thus, by noisy fluctuations around the average FC structure. We test this hypothesis by constructing appropriate surrogate data. We use a method that randomizes the phases while preserving the pairwise correlations and the power spectrum (Handwerker et al., 2012; Hindriks et al., 2015), resulting in any nonstationary properties being removed (see figure 9A and C for an example).

It is evident that also the surrogates possess the bands/events noted before, and thus, these synchronized movements are not per se an indication of any active process. Rather, they are the consequence of the average correlations and the 1/f power spectrum, which is shown in figure 9E to be reproduced by the surrogates. As a result, when using the same three measures as before, modulations are apparent also in the time courses derived from surrogate data, and the correlations between them are preserved as well.

In many studies on dFC, the focus is on large deviations from the average correlation, referred to in some instances as "hypersynchronization" (Hutchison et al., 2012). As we have shown here, the peaks of the modulation do not in fact coincide with large deviations from the average correlation

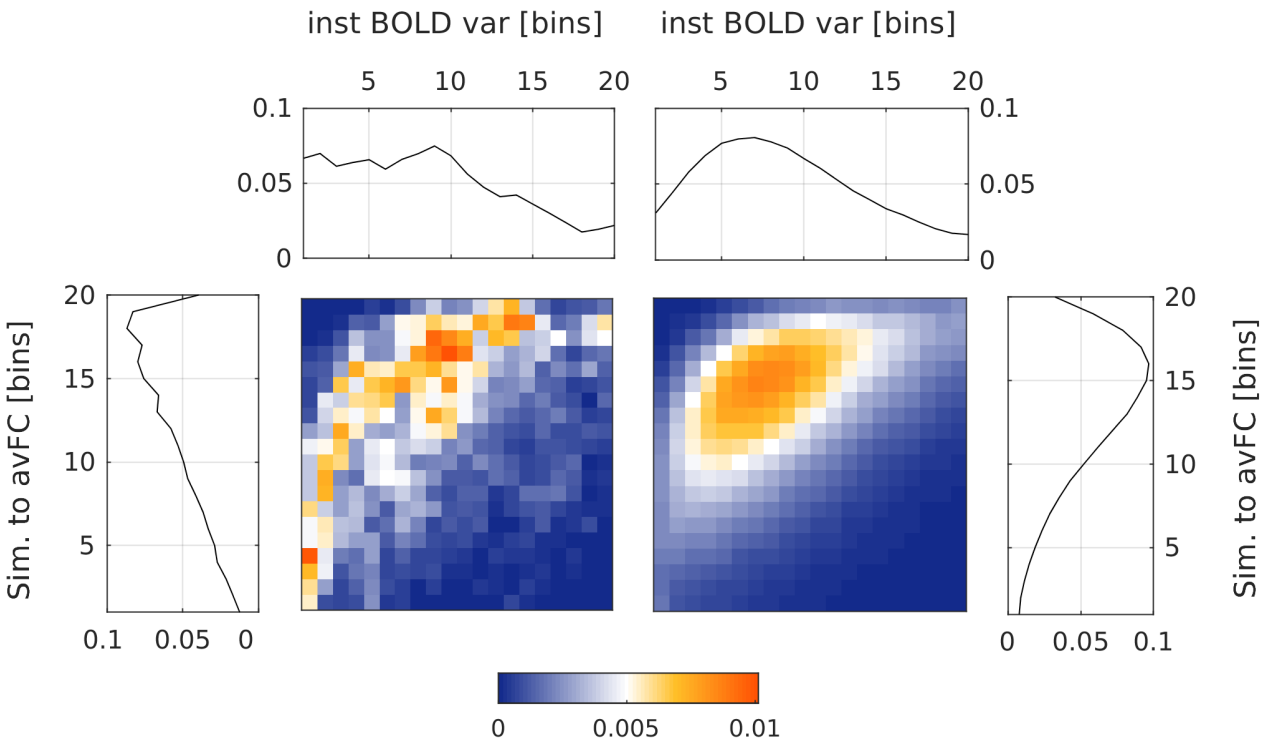

Figure 8: Joint distribution of instantaneous BOLD variance (x-axes) and similarity to avFC (y-axes) for empirical data (left) and surrogates (right). The marginal distributions are shown at the top and sides. There is a clear and significant difference between the distributions, see section... (surrogates), hinting at longer dwell times in the extremes of the distributions.

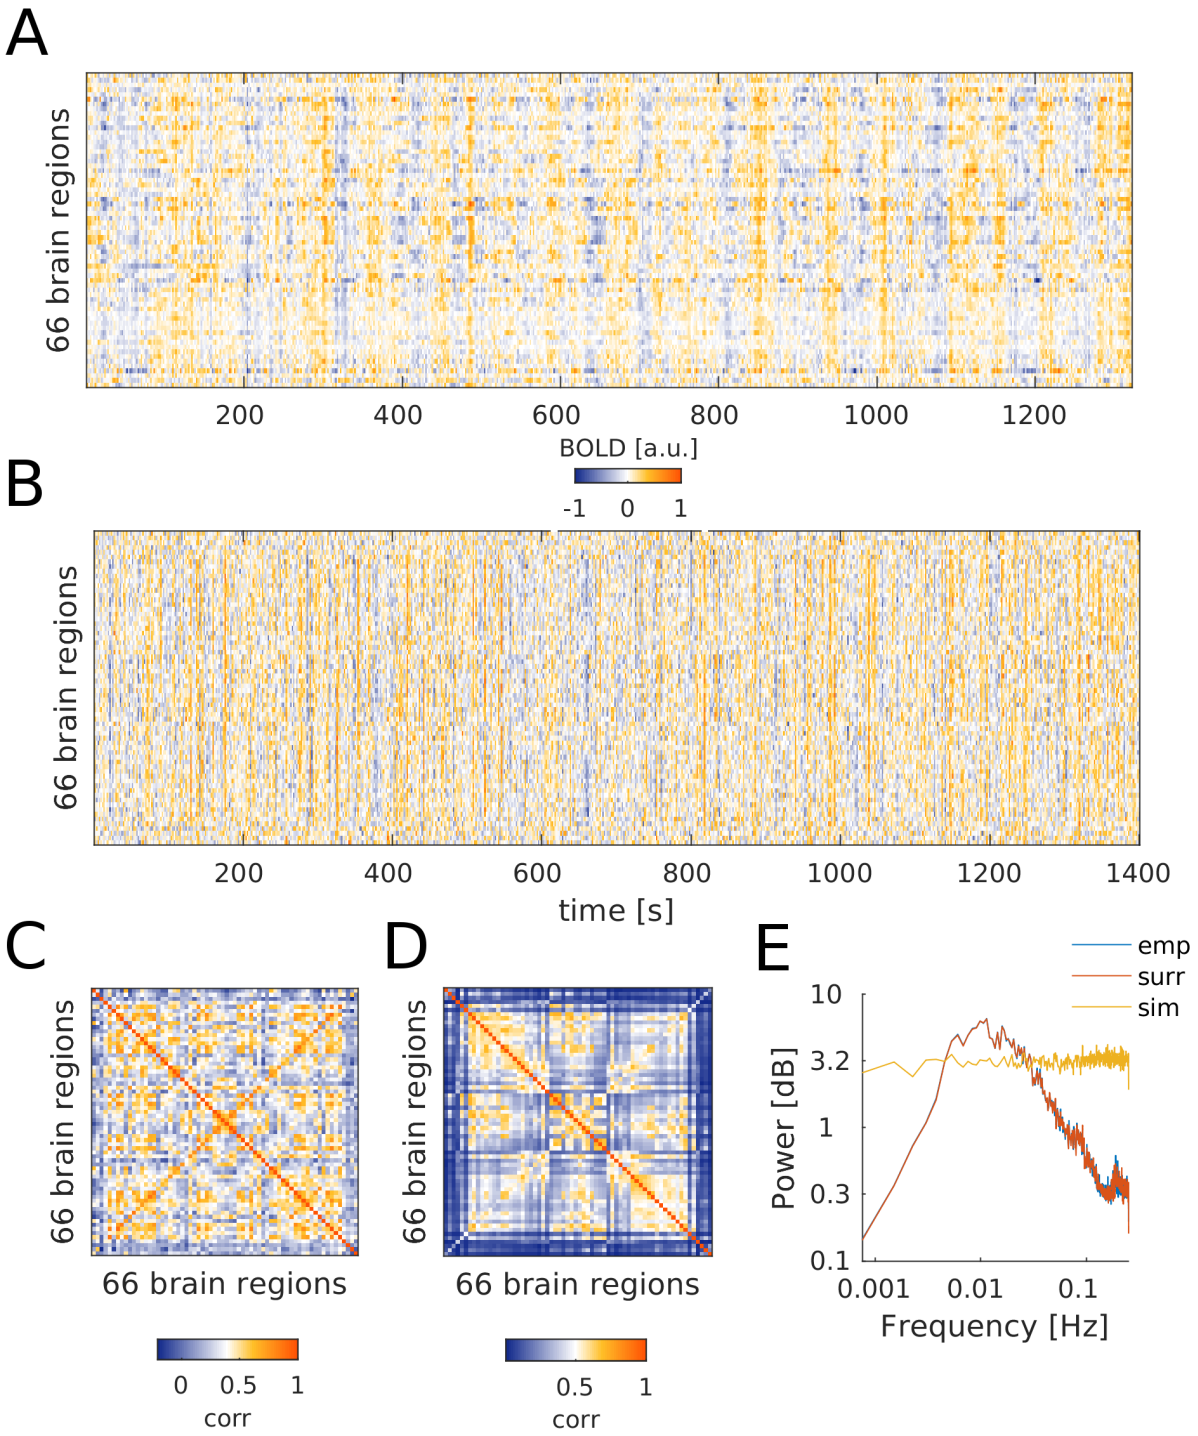

Figure 9: Stationary null data used in this study. **A** Phase randomized BOLD time courses from one subject. **B** Simulated data which were passed through the Balloon-Windkessel model. **C** Average FC matrix obtained from the surrogate time courses shown in **A**. It is nearly identical to that of the empirical data from which the surrogates where derived. **D** Average FC matrix obtained from the simulated time courses shown in **B**. The correlation between the two FC matrices is about 0.6. **E** Average power spectra of empirical, surrogate, and simulated data. Empirical and surrogate spectra lie almost exactly on top of each other.

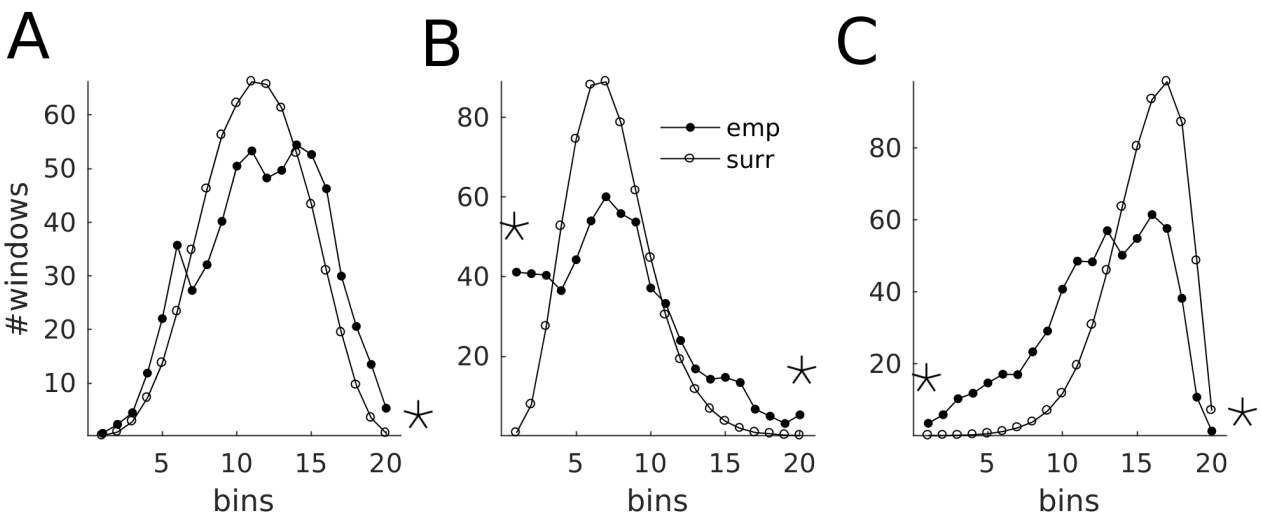

Figure 10: Comparison of empirical (filled circles) and surrogate (empty circles) histograms. The asterisks indicate instances where the real data have significantly more or fewer windows in the extreme bins than the surrogate data. **A** Instantaneous average correlation, **B** instantaneous BOLD variance, **C** similarity to avFC.

structure - on the contrary, dFCMs tend to be more similar to the avFCM at the peaks. Furthermore, the fact that brief instances in time are sufficient to extract RSNs can from this point of view be explained with the prominent overlap of RSNs, which leads to a situation where they are present together in those windows with "up"-modulation. Still, it is possible that the deviations are large in empirical data as compared to what can be explained by stationary dynamics.

In order to test this, we record, for each subject, both the minimum and the maximum of the modulations in all three measures (instantaneous BOLD variance, instantaneous average correlation, similarity to avFC). We compute 100 sets of surrogate data sets and count how many times they reach a higher (lower) value than the original data. We find that many surrogates reach higher maxima (inst. av. corr: 21.8%, inst. BOLD var.: 16.8%, sim. to avFC: 81.0%) or lower minima (inst. av. corr: 52.4%, inst. BOLD var.: 6.9%, sim. to avFC: 11.6%) than the real data, meaning that surrogates can fully explain the size of the modulations.

On the other hand, it might still be possible that the system visits these states more (or less) often than could be explained by stationary dynamics, as is indeed suggested by the clusters of windows seen in figure 8 (left). In order to test this, we count the number of windows falling into the extreme bins of the distributions (both minimum and maximum). Figure 10 illustrates this with the average distributions. We take the average number of windows across subjects and, again, compute 100 sets of surrogates. Here we find that the real data have more windows of high instantaneous average correlation, both low and high instantaneous BOLD variance, and of low similarity to avFC than the surrogates, i.e. none of the surrogates reached or surpassed the number of windows present in the real data. On the other hand, real data have *fewer* windows with high similarity to avFC than the surrogates, i.e. all surrogates reached or surpassed the number of windows present in the real data.

Returning to figure 8 and now comparing the two panels (empirical vs. surrogates), it becomes clear that the modulations themselves are explained by the power spectrum and the presence of average correlations, but that the prevalence of extreme modulations (both up and down), or in other words, the *dwell times* in those states, is *not* explained by these stationary features of the data.

3.5 RSN temporal dynamics in surrogates and stationary model

RSNs can be extracted in the exact same way from the surrogate as from the empirical data. When computing the histograms of the contributions in order to evaluate the propensity for coactivation (figure 6, middle panel), we find again strong temporal overlap as in the empirical data. However, as a consequence of the longer dwell times, the peaks to the left and right are missing, i.e. what is reduced in the surrogate RSN time courses are both periods in which networks are inactive and in which one network dominates.

Our next step is to explicitly test the hypothesis that the observed modulations can be explained by correlated fluctuations around a single single fixed point - that is, by the average functional connectivity. We use a dynamic mean field model which possesses exactly this structure and which has been shown before (Glomb et al., 2016) to reproduce RSNs. Nodes in the model, representing brain regions, are connected according to effective connectivity obtained by combining diffusion tensor imaging results with a simple model (Gilson et al., 2016).

Figure 9B and D show an example for a simulated signal and the average FC matrix computed from it. Figure 9E also shows the average power spectrum of all simulations when using the value of G which produces the best match between empirical and simulated avFC. Clearly, while the correlation structure is captured well, the power spectrum is consistent with white noise, while empirical (and thus surrogate) data exhibit a $1/\mathrm{f}$ spectrum (Biswal et al., 1995; Niazy et al., 2011).

As for the empirical data and the surrogate data, we compute time courses of our three global measures of temporal dynamics, and find that they exhibit very similar dynamics. Likewise, computing the coactivation histograms for the model data (figure 6, right panel) shows again strong coactivation, but fewer windows in the extremes of the distributions, even more prominently than in the surrogate data.

In conclusion, we have shown that

1. the modulations observed in BOLD temporal dynamics and RSNs are *stereotypical* in the sense that peaks correspond to strong RSN coactivation, high BOLD variance, high average correlation and high similarity of dFCM to avFCM (as opposed to the notion that each peak represents a different state),

2. the presence and the size of these modulations are explained by the power spectrum and the average correlations present in the data, as shown by using surrogates and a stationary dynamic mean field model,

3. nonetheless the dwell times in the maxima and, most prominently, the minima of the modulations are longer than can be explained by stationary dynamics, suggesting the presence of an active, nonstationary process.

4 Discussion

In this study, we used three simple measures to characterize modulations in the BOLD signal and their relationship to dFC: (1) instantaneous BOLD variance, (2) instantaneous average correlation, (3) similarity to avFC. We related these measures to RSN activation time courses obtained from empirical data via tensor decomposition. We found that a good part of the modulations we observe can be described as stereotypical events which evolve from a state where all three measures and RSN activations are low, to a state where these measures are high. The dynamics of the transitions

between "down" and "up" exhibit nonstationary features: There are many more "down" windows than expected, and also slightly more "up" windows.

The main finding of this study is that modulations in the activation strength of RSNs can be largely explained by changes in the variance of the underlying BOLD signal. The highest BOLD variance ("up") occurs in windows which have both synchronized increase in BOLD activity (orange band in figure 2A) as well as a decrease (blue band). In the following, we summarize this as a "cycle" for illustration purposes, without claimin that there is an actual oscillation going on. At the beginning of a "cycle", there is a large amount of activity that propagates through the entire brain network, shaped by the underlying effective connectivity and the node dynamics. At this stage, the avFC is expressed clearly because the signal-to-noise-ratio is favorable. When the BOLD activity drops, pairwise dFC is still high because of synchronized decrease of activity. It is due to these high pairwise dFC values that we find RSNs in these windows, and due to the synchronized increase and decrease, that the RSNs overlap in time. On the other hand, RSNs are derived from the avFC, which seems at odds with the lower correlations of similarity to avFC with RSN activation, as well as with other measures. This could be due to the fact that during the decrease, the avFC as a whole is not as well expressed because the SNR is lower, but thresholding in the tensor before decomposition depends only on the magnitude of the dFC values.

Additionally to these descriptive results, we used surrogate data from which nonstationarities were removed, as well as simulated data obtained from a stationary model, to investigate the mechanisms behind our observations. We showed that many statistical properties of spontaneous fluctuations in the resting state BOLD signal are explained by the presence of long-term correlations and the characteristic power spectrum of these data. Both the range of modulations and the correlations between them are reproduced by the surrogates and the simulated data. However, when considering the full distributions of the three measures we established, it is revealed that dynamics of the empirical data exhibit nonstationary properties in the amount of time the system spends in the extrema of the modulations.

4.1 Relationship between changes in activity and in FC

We have shown very clearly that there is a strong positive correlation between changes in BOLD variance and changes in FC, in parallel to findings on a finer spatial and temporal scale (Chawla et al., 1999; De La Rocha et al., 2007) and for phase relationships (Daffertshofer and van Wijk, 2011). While it may be clear that this relationship is expected in the sense that it is explained by the long-term correlations, it is not necessarily intuitive or without alternative. Naively, one could hypothesize that fluctuations in BOLD act as noise which would lead to globally high correlations and thus, 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 straightforward to define FC in terms of statistical dependency of two full time courses, it seems that to what exactly it refers in terms of dynamical changes is not as clear. Effective connectivity might be a useful complementary approach.

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.

4.2 Time scales and possible origins of fluctuations

One limitation of our methods here is that we cannot make clear predictions about the origin of the observed time scales: In the surrogate data, the Fourier spectra of the time courses are preserved as per construction, and the simulated data, while preserving many properties of the dynamics, do not exhibit the same kind of power spectrum at all (figure 9E). Recent work has begun to illuminate the relationship between network structure and power spectrum (He, 2014; Chaudhuri et al., 2016; Ocker et al., 2017). Elsewhere, 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.

Fluctuations of our three global measures occur on an even slower time scale ($\approx$0.01 Hz). 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 or metastability.

4.3 RSNs overlap in time

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 being separate in time. Here, we find that they coactivate a lot of the time, which was also reported by de Pasquale et al. (2012); Smith et al. (2012); Karahanoğlu and Van De Ville (2015). We find that RSNs are most strongly activated and coactivated during periods of strong synchronization (i.e., high average correlation). 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.

These findings go hand in hand with the result that dFC structure is most of the time close to avFC and that most modulations can be explained by fluctuations around this average structure. While this may be surprising, it might indeed be a sign of optimal organization of the brain network, as suggested by recent work combining dFC with control theory (Gu et al., 2015; Betzel et al., 2016b). From the point of view proposed there, it is desirable to remain in a regime from which frequently visited states are easily reached, both in terms of energy cost and time. In this light, it would actully be surprising to see large excursions from the average during RS.

4.4 Conclusions and Outlook

We have shown that no additional assumptions about dynamics are necessary to model a community structure that evolves in time (figure 3, Glomb et al. (2016)), but rather, that this property can be obtained from combining a meaningful underlying connectivity (EC in our case) with simple node dynamics. RSNs can be extracted from brief intervals which correspond to the peaks in the global measures used here. This is explained by the fact that, in terms of distributions of dFC values, the BOLD signal on the longer time scale of the modulations studied here (about 50-100s) contains much of the same information as that on the shorter time scale used for the EC, and both time scales are predicted by the avFC.

However, we do find interesting properties in the dynamics not explained by a fluctuating avFC. In terms of RSNs, the transition from "down" to "up" means a transition from none of them being activated to all of them being coactivated. Here, the nonlinear relationship is evident from the histograms of the contributions (figure 6, left), where both strong coactivation ("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.

The model used here is able to reproduce some of the hallmarks of dFC due to the temporal structure of BOLD signals, but lacks the non-stationarities. 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, they 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. Another property that the DMF model used here does not reproduce is the power spectrum. We find that the time scale of modulations is not explained in spite of the fact that the hemodynamic signal is modelled explicitly (Friston et al., 2000), i.e. the $1/f$ power spectrum of real BOLD signals is not shared by the simulated data. Whether the Hopf model could reproduce the power spectrum should be included in the investigations.

References

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.

Babo-Rebelo M, Richter CG, Tallon-Baudry C (2016) Neural responses to heartbeats in the default network encode the self in spontaneous thoughts. *Journal of Neuroscience* 36:7829–7840.

Babo-Rebelo M, Wolpert N, Adam C, Hasboun D, Tallon-Baudry C (2016) Is the cardiac monitoring function related to the self in both the default network and right anterior insula? *Phil. Trans. R. Soc. B* 371:20160004.

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.

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.

Bassett DS, Porter MA, Wymbs NF, Grafton ST, Carlson JM, Mucha PJ (2013) Robust detection of dynamic community structure in networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 23:013142.

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.

Beckmann CF, Smith SM (2005) Tensorial extensions of independent component analysis for multisubject fMRI analysis. *NeuroImage* 25:294–311.

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.

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.

Betzel RF, Fukushima M, He Y, Zuo XN, Sporns O (2016a) Dynamic fluctuations coincide with periods of high and low modularity in resting-state functional brain networks. *NeuroImage* 127:287–297.

Betzel RF, Gu S, Medaglia JD, Pasqualetti F, Bassett DS (2016b) Optimally controlling the human connectome: the role of network topology. *Scientific Reports* 6.

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.

Britz J, Van De Ville D, Michel CM (2010) BOLD correlates of EEG topography reveal rapid resting-state network dynamics. *NeuroImage* 52:1162–1170.

Cabral J, Hugues E, Sporns O, Deco G (2011) Role of local network oscillations in resting-state functional connectivity. *NeuroImage* 57:130–139.

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.

Calhoun VD, Miller R, Pearlson G, Adalı T (2014) The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery. *Neuron* 84:262–274.

Chang C, Glover GH (2010) Time–frequency dynamics of resting-state brain connectivity measured with fMRI. *NeuroImage* 50:81–98.

Chang C, Metzger CD, Glover GH, Duyn JH, Heinze HJ, Walter M (2013) Association between heart rate variability and fluctuations in resting-state functional connectivity. *Neuroimage* 68:93–104.

Chaudhuri R, He B, Wang XJ (2016) Random recurrent networks near criticality capture the broadband power distribution of human ecog dynamics. *bioRxiv* p. 036228.

Chawla D, Lumer ED, Friston KJ (1999) The relationship between synchronization among neuronal populations and their mean activity levels. *Neural Computation* 11:1389–1411.

Cichocki A (2013) Tensor decompositions: A new concept in brain data analysis? *arXiv* pp. 507–517.

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.

Daffertshofer A, van Wijk B (2011) On the influence of amplitude on the connectivity between phases. *Frontiers in neuroinformatics* 5:6.

Damaraju E, Allen E, Belger A, Ford J, McEwen S, Mathalon D, Mueller B, Pearlson G, Potkin S, Preda A et al. (2014) Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia. *NeuroImage: Clinical* 5:298–308.

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.

Davison EN, Schlesinger KJ, Bassett DS, Lynall ME, Miller MB, Grafton ST, Carlson JM (2015) Brain network adaptability across task states. *PLoS Comput Biol* 11:e1004029.

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.

De La Rocha J, Doiron B, Shea-Brown E, Josić K, Reyes A (2007) Correlation between neural spike trains increases with firing rate. *Nature* 448:802–806.

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.

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.

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.

Deco G, Jirsa VK (2012) Ongoing cortical activity at rest: criticality, multistability, and ghost attractors. *The Journal of Neuroscience* 32:3366–3375.

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.

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 Neuroscience* 33:11239–11252.

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.

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.

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.

Fries P (2005) A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in cognitive sciences* 9:474–480.

Friston KJ, Mechelli A, Turner R, Price C (2000) Nonlinear responses in fmri: the balloon model, volterra kernels, and other hemodynamics. *NeuroImage* 12:466–477.

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.

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* .

Glomb K, Ponce-Alvarez A, Gilson M, Ritter P, Deco G (2016) Robust extraction of spatio-temporal patterns from resting state fmri. *bioRxiv* p. 089516.

Gu S, Pasqualetti F, Cieslak M, Telesford QK, Alfred BY, Kahn AE, Medaglia JD, Vettel JM, Miller MB, Grafton ST et al. (2015) Controllability of structural brain networks. *Nature communications* 6.

Handwerker DA, Roopchansingh V, Gonzalez-Castillo J, Bandettini PA (2012) Periodic changes in fMRI connectivity. *NeuroImage* 63:1712–1719.

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.

He BJ (2014) Scale-free brain activity: past, present, and future. *Trends in cognitive sciences* 18:480–487.

He BJ, Raichle ME (2009) The fMRI signal, slow cortical potential and consciousness. *Trends in cognitive sciences* 13:302–309.

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.

Horovitz 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.

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.

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.

Karahanoğlu FI, Van De Ville D (2015) Transient brain activity disentangles fmri resting-state dynamics in terms of spatially and temporally overlapping networks. *Nature communications* 6.

Kiviniemi V, Kantola JH, Jauhiainen J, Hyvärinen A, Tervonen O (2003) Independent component analysis of nondeterministic fMRI signal sources. *NeuroImage* 19:253–260.

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.

Kopell NJ, Gritton HJ, Whittington MA, Kramer MA (2014) Beyond the connectome: the dynome. *Neuron* 83:1319–1328.

Kraskov A, Stögbauer H, Grassberger P (2004) Estimating mutual information. *Physical Review E* 69:066138.

Kucyi A, Davis KD (2014) Dynamic functional connectivity of the default mode network tracks daydreaming. *Neuroimage* 100:471–480.

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.

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* .

Leonardi N, Van de Ville D (2013a) 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.

Leonardi N, Van De Ville D (2013b) Tight wavelet frames on multislice graphs. *IEEE Transactions on Signal Processing* 61:3357–3367.

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.

Liu X, Chang C, Duyn JH (2013) Decomposition of spontaneous brain activity into distinct fMRI co-activation patterns. *Frontiers in Systems Neuroscience* 7:101.

Lowe M, Mock B, Sorenson J (1998) Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. *NeuroImage* 7:119–132.

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.

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.

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.

Mitra A, Snyder AZ, Hacker CD, Raichle ME (2014) Lag structure in resting-state fMRI. *Journal of Neurophysiology* 111:2374–91.

Niazy R, Xie J, Miller K, Beckmann C, Smith S (2011) Spectral characteristics of resting state networks In Van Someren E, Van Der Werf Y, Roelfsema P, Mansvelder H, Lopes da Silva F, editors, *Progress in brain research*, Vol. 193, chapter 17, pp. 259–276.

Ocker GK, Hu Y, Buice MA, Doiron B, Josić K, Rosenbaum R, Shea-Brown E (2017) From the statistics of connectivity to the statistics of spike times in neuronal networks. *arXiv preprint arXiv:1703.03132* .

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.

Preti MG, Bolton TA, Van De Ville D (2016) The dynamic functional connectome: State-of-the-art and perspectives. *NeuroImage* .

Richter CG, Babo-Rebelo M, Schwartz D, Tallon-Baudry C (2016) Phase-amplitude coupling at the organism level: the amplitude of spontaneous alpha rhythm fluctuations varies with the phase of the infra-slow gastric basal rhythm. *NeuroImage* .

Ritter P, Schirner M, McIntosh AR, Jirsa VK (2013) The virtual brain integrates computational modeling and multimodal neuroimaging. *Brain connectivity* 3:121–145.

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.

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.

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.

Tournier JD, Calamante F, Connelly A (2007) Robust determination of the fibre orientation distribution in diffusion mri: non-negativity constrained super-resolved spherical deconvolution. *NeuroImage* 35:1459–1472.

Tournier JD, Calamante F, Gadian DG, Connelly A (2004) Direct estimation of the fiber orientation density function from diffusion-weighted mri data using spherical deconvolution. *NeuroImage* 23:1176–1185.

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.

Wong KF, Wang XJ (2006) A recurrent network mechanism of time integration in perceptual decisions. *The Journal of Neuroscience* 26:1314–1328.

Yang Z, Craddock RC, Margulies DS, Yan CG, Milham MP (2014) Common intrinsic connectivity states among posteromedial cortex subdivisions: Insights from analysis of temporal dynamics. *Neuroimage* 93:124–137.

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.

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.