NeuroImage 197 (2019) 37–48

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

Investigating the impact of autocorrelation on time-varying connectivity

Hamed Honari [a], Ann S. Choe [b,c], James J. Pekar [b,c], Martin A. Lindquist [d,*]

[a] Department of Electrical and Computer Engineering, Johns Hopkins University, USA
[b] Russell H. Morgan Department of Radiology and Radiological Science, Johns Hopkins University School of Medicine, USA
[c] F.M. Kirby Research Center for Functional Brain Imaging, Kennedy Krieger Institute, USA
[d] Department of Biostatistics, Johns Hopkins University, USA

ARTICLE INFO

Keywords:
Dynamic functional connectivity
Time-varying functional connectivity
Resting-state fMRI
Autocorrelation
Sliding-window
Prewhitening

ABSTRACT

In recent years, a number of studies have reported on the existence of time-varying functional connectivity (TVC) in resting-state functional magnetic resonance imaging (rs-fMRI) data. The sliding-window technique is currently one of the most commonly used methods to estimate TVC. Although previous studies have shown that autocorrelation can negatively impact estimates of static functional connectivity, its impact on TVC estimates is not well known at this time. In this paper, we show both theoretically and empirically that the existence of autocorrelation within a time series can inflate the sampling variability of TVC estimated using the sliding-window technique. This can in turn increase the risk of misinterpreting noise as true TVC and negatively impact subsequent estimation of whole-brain time-varying FC profiles, or "brain states". The latter holds as more variable input measures lead to more variable output measures in the state estimation procedure. Finally, we demonstrate that prewhitening the data prior to analysis can lower the variance of the estimated TVC and improve brain state estimation. These results suggest that careful consideration is required when making inferences on TVC.

1. Introduction

It was previously assumed that functional connectivity (FC) among brain regions did not vary substantially during the course of a single resting-state functional magnetic resonance imaging (rs-fMRI) run. However, recent studies have begun uncovering the possible existence of dynamic changes in FC taking place on a considerably shorter time frame (i.e., seconds instead of entire runs lasting several minutes) (Chang and Glover, 2010; Hutchison et al., 2013a; Preti et al., 2016; Tagliazucchi et al., 2012; Thompson et al., 2013; Allen et al., 2014). These studies of time-varying FC (TVC) have allowed for the classification of whole-brain time-varying FC profiles into distinct "brain states", which are defined as recurring whole-brain connectivity profiles that are reliably observed across subjects during a resting-state run (Calhoun et al., 2014; Lurie et al., 2018). Results have demonstrated that while the patterns of connectivity describing each brain state are reliable across groups and individuals (Yang et al., 2014; Choe et al., 2017), other characteristics such as the amount of time spent in specific brain states and the number of transitions between brain states can vary as a function of individual differences such as age (Hutchison and Morton, 2015) or disease status (Damaraju et al., 2014; Rashid et al., 2014).

To date, a number of different analytic techniques have been proposed for quantifying and investigating the presence of TVC. These include the widely-used sliding-window approach (Tagliazucchi et al., 2012; Chang and Glover, 2010; Hutchison et al., 2013a, b), change point analysis (Cribben et al., 2012, 2013; Xu and Lindquist, 2015), point process analysis (Tagliazucchi et al., 2011), Co-Activation Patterns (CAPs) (Liu and Duyn, 2013), transient-based CAPs (Karahanoğlu and Van De Ville, 2015), time series models (Lindquist et al., 2014), time-frequency analysis (Chang and Glover, 2010), Hidden Markov Models (HMM) (Eavani et al., 2013; Vidaurre et al., 2017), Hidden Semi-Markov Models (HSMM) (Shappell et al., 2019), and Sparse Coupled Hidden Markov Model (SCHMM) (Bolton et al., 2018). Despite development of these promising approaches, estimation of TVC still remains a difficult endeavor due to the low signal-to-noise ratio (SNR) of the blood oxygen level dependent (BOLD) signal and the presence of image artifacts and nuisance confounds (Hutchison et al., 2013a; Lindquist et al., 2014; Laumann et al., 2016).

One potentially significant confound that has not yet received much attention in the literature is the presence of autocorrelation *within* the specific time series used to estimate TVC. It is well known that even in the absence of explicit neuronal activation BOLD fMRI noise exhibits positive autocorrelation, which is generally attributed to unmodeled nuisance signal (Lund et al., 2006; Lenoski et al., 2008; Zarahn et al., 1997; Purdon

* Corresponding author.
E-mail address: mlindqui@jhsph.edu (M.A. Lindquist).

https://doi.org/10.1016/j.neuroimage.2019.04.042
Received 27 February 2019; Received in revised form 10 April 2019; Accepted 15 April 2019
Available online 22 April 2019

and Weisskoff, 1998). In the context of task fMRI, this existence of positive autocorrelation is generally accounted for within the general linear model (GLM) framework, using either an AR or ARMA model after removing the effects of the task from the signal (Lindquist et al., 2008). In contrast, analytic concerns related to the existence of such autocorrelation in rs-fMRI data has received much less attention, although recent studies have reported that autocorrelation in rs-fMRI data can negatively impact the estimates of static functional connectivity (Christova et al., 2011; Georgopoulos and Mahan, 2013). Arbabshirani et al. (2014) did attempt to investigate this issue in depth and showed that while autocorrelation influences point estimation of the cross-correlation coefficient, it has minimal effect on hypothesis tests of group-level differences in correlation. The authors concluded that since hypothesis testing is the primary inferential technique used in most FC studies, the impact of autocorrelation on the estimate of the correlation coefficient was not necessarily detrimental to the findings of most studies.

However, when studying TVC, this argument is less suitable as it is typically the point estimate (for example, the cross-correlation coefficient at each time point) that is of primary importance. For example, these estimates are typically used as an input in a secondary analysis to estimate brain states. Therefore, it is critical to re-visit the issue of autocorrelation in this context, and to investigate the influence of autocorrelation on TVC estimation. In this study, we illustrate these issues using the sliding-window technique, due to its popularity. However, it should be noted that the issue of autocorrelation is not unique to the sliding-window technique and is shared by other TVC estimation techniques.

The problem is particularly complicated in the rs-fMRI setting as it is not clear that the autocorrelation can be assumed to be purely artifactual from a neuroscientific point of view. Here, much like in the task fMRI setting, it is instead reasonable to assume that the autocorrelation present in the signal is partially due to physiological noise uncoupled from neural activity and partially due to brain hemodynamics (i.e., the temporal convolution of underlying neuronal activity with the hemodynamic response function (HRF)). In general, it is of interest for researchers to remove the impact of autocorrelation arising from the former, while retaining the influence of the latter. While in the context of task fMRI this is generally performed within a GLM framework, similar strategies are not directly applicable to resting-state data due to difficulties in postulating *a priori* the expected BOLD responses due to hemodynamics. We explore this topic further in simulation studies.

We begin with the problem set-up and introduction of the sliding-window technique. We continue by theoretically illustrating how the presence of autocorrelation *within* a rs-fMRI time series can lead to potential false positive findings, and discuss possible solutions to the problem. Finally, we empirically illustrate the autocorrelation problem through the use of a series of simulations and a test-retest rs-fMRI data set.

2. Methods

In this section we introduce the problem. We illustrate by studying bivariate relationships because it makes the problem simple and easy to understand, with closed-form solutions to many of the necessary expressions. In the case of the sliding window correlations the jump to the multivariate case is trivial. Here the multivariate case is equivalent to computing bivariate correlations for each pair of time courses. This holds because the Pearson correlation between the i^{th} and j^{th} region does not depend on any of the other time courses included in the analysis. We do note that if one is interested in computing partial correlations this no longer holds, and a multivariate treatment would be required. However, this is not the focus of the current work.

2.1. Problem set-up

Suppose we are interested in studying the relationship between two different time series, denoted $y_{1,t}$ and $y_{2,t}$, which are measured over two separate regions of interest (ROIs) at equally spaced time points $t = 1, \ldots n$. Let $\boldsymbol{y}_t = (y_{1,t},\ y_{2,t})^T$ be a vector containing the values of both time series at a specific time t. Further, assume that we can separate $\boldsymbol{y}_t$ into a signal ($\boldsymbol{\mu}_t$) and a noise term ($\boldsymbol{\varepsilon}_t$), that model the signal related to brain hemodynamics and nuisance terms uncoupled from neural activity (e.g., physiological noise), respectively. Hence, we can write:

$$\boldsymbol{y}_t = \boldsymbol{\mu}_t + \boldsymbol{\varepsilon}_t, \tag{1}$$

where $\mathbb{E}_{t-1}[\boldsymbol{y}_t] = \boldsymbol{\mu}_t = (\mu_{1,t},\ \mu_{2,t})^T$ represents the conditional mean of $\boldsymbol{y}_t$ using all information in the time series observed up to time $t-1$, and $\boldsymbol{\varepsilon}_t = (\varepsilon_{1,t},\ \varepsilon_{2,t})^T$ represents the noise at time t. Assume the latter has mean zero and a conditional covariance matrix at time t that can be expressed as

$$\boldsymbol{\Sigma}_t = \begin{pmatrix} \sigma^2_{1,t} & \sigma_{12,t} \\ \sigma_{12,t} & \sigma^2_{2,t} \end{pmatrix}. \tag{2}$$

Here the diagonal terms of the matrix represent the conditional variance of $y_{i,t}$ obtained using all information in the time course observed up to time $t-1$, or $\sigma^2_{i,t} = \mathbb{E}_{t-1}[(y_{i,t} - \mu_{i,t})^2]$ for $i = 1, 2$. The off-diagonal term represents the conditional covariance, or $\sigma_{12,t} = \sigma_{1,t}\sigma_{2,t}\rho_t$. Here the term

$$\rho_t = \frac{\mathbb{E}_{t-1}\left[\left(y_{1,t} - \mu_{1,t}\right)\left(y_{2,t} - \mu_{2,t}\right)\right]}{\sqrt{\mathbb{E}_{t-1}\left[\left(y_{1,t} - \mu_{1,t}\right)^2\right]\mathbb{E}_{t-1}\left[\left(y_{2,t} - \mu_{2,t}\right)^2\right]}} \tag{3}$$

is the conditional correlation coefficient. Under this definition the conditional correlation at time t relies on information that is observed up to time $t-1$. In this paper, we will primarily be concerned with estimating ρ_t.

2.2. Sliding-window technique

The sliding-window technique is the most widely-used technique for estimating TVC. Here, a temporal window of fixed length ℓ is chosen, and data points within the window are used to compute correlation coefficients between pairwise time series of ROIs. New correlation coefficients are then calculated by sliding the window across time points.

Here, a general form of the correlation coefficient is given by:

$$\widehat{\rho}_t = \frac{\sum_{s=t-\ell-1}^{t-1}\left(y_{1,s} - \widehat{\mu}_{1,s}\right)\left(y_{2,s} - \widehat{\mu}_{2,s}\right)}{\sqrt{\left(\sum_{s=t-\ell-1}^{t-1}\left(y_{1,s} - \widehat{\mu}_{1,s}\right)^2\right)\left(\sum_{s=t-\ell-1}^{t-1}\left(y_{2,s} - \widehat{\mu}_{2,s}\right)^2\right)}} \tag{4}$$

where $\widehat{\mu}_{1,s}$ and $\widehat{\mu}_{2,s}$ represent the estimated time-varying mean of the two time series. This equation describes a boxcar window that gives equal weight to all observations that lie within the window and zero weight to all other observations. Note it is common to alternatively estimate the correlation at time t using a window over the interval $[t - \ell/2,\ t + \ell/2]$, that is using the $\ell/2$ points before and after t. We note that as long as we keep track of the reference point in the window used in the sliding window analysis, and use this reference consistently with regards to any other simultaneously recorded variables, then it is not important if this reference point is chosen as the middle of the window or at the beginning of the window.

2.3. The autocorrelation problem - theoretical illustration

In this section we investigate the effects of autocorrelation related to the nuisance signal (i.e., the portion assumed uncoupled from neural activity) on TVC estimation using the sliding-windows technique. Let us assume that $\mu_{i,t} = 0$ and $\varepsilon_{i,t}$ follows a univariate AR(1) process for $i = 1$,

2. This implies that we are dealing with two time series consisting solely of the autocorrelated error term. In matrix form we can write this as follows:

$$\boldsymbol{y}_t = \boldsymbol{\varepsilon}_t \tag{5}$$

$$\boldsymbol{\varepsilon}_t = \boldsymbol{\Phi}\boldsymbol{\varepsilon}_{t-1} + \boldsymbol{w}_t \tag{6}$$

where $\boldsymbol{w}_t = (w_{1,t}, w_{2,t})^T$ are white noise processes with mean zero and variance $\boldsymbol{\Sigma}_w$, and

$$\boldsymbol{\Phi} = \begin{pmatrix} \phi_{11} & 0 \\ 0 & \phi_{22} \end{pmatrix}. \tag{7}$$

Consider for the sake of simplicity and without loss of generality, that the AR(1) parameters are: $\phi_{11} = \phi_{22} = \phi$. If we let $w_{1,t}$ and $w_{2,t}$ be independent, the covariance matrix $\boldsymbol{\Sigma}_w$ can be written:

$$\boldsymbol{\Sigma}_w = \begin{pmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{pmatrix} \tag{8}$$

Thus, $\boldsymbol{y}_t$ at time t can be written:

$$\boldsymbol{y}_t = \boldsymbol{\Phi}\boldsymbol{y}_{t-1} + \boldsymbol{w}_t. \tag{9}$$

Hence, the terms of $\boldsymbol{y}_t$ can be considered two univariate AR(1) processes. This set-up corresponds to a null case where the true correlation between the two time series is 0 at all time points.

Now, the theoretical mean of each time series is $\mathbb{E}[y_{1,t}] = \mathbb{E}[y_{2,t}] = 0$ and the variance can be written as:

$$Var(y_{i,t}) = \frac{\sigma_i^2}{1-\phi_{ii}^2} = \frac{\sigma_i^2}{1-\phi^2}; \quad i = 1, 2 \tag{10}$$

It can be shown (Bartlett, 1935) that the variance of the sample correlation between these two time series can be approximated by:

$$Var(\rho) = \frac{1}{n}\frac{1+\phi^2}{1-\phi^2} \tag{11}$$

where n is the sample size; which in the context of sliding-window technique, is the number of observations used in the window.

The above equation suggests that the variance of the sample correlation between two AR(1) processes depends heavily on the value of ϕ, even when the true correlation between the two processes is zero. In particular, as ϕ approaches 1, the variance of the sample correlation inflates rapidly even though the two time series are truly uncorrelated. This emphasizes the importance of taking into account the autocorrelation before making any statements about the significance of correlations between two time series.

The inflation of the variance of the sample correlation due to autocorrelation is particularly important in the context of sliding-window technique as it implies that the variance of the TVC estimate at time t when using the window of length ℓ is given by:

$$Var(\rho_t) = \frac{1}{\ell}\frac{1+\phi^2}{1-\phi^2}. \tag{12}$$

In other words, as the window length gets smaller, the variance increases. Previous studies have shown that such increased variance has the risk of being misinterpreted as important information about TVC driven by neuronal fluctuations, when it is instead entirely consistent with the correlation being static (Lindquist et al., 2014; Hindriks et al., 2016). Importantly, although often neglected, is the fact that as the autocorrelation *within* each time course increases, the variance of the correlation *between* the two time series will similarly increase, and thus further increase the possibility of observing large correlations in the null setting.

2.4. Prewhitening

In order for inference to be valid, the autocorrelation that is present in most fMRI time series needs to be properly taken into account. This can either be done by incorporating the additional uncertainty associated with the autocorrelation into the inferential procedure (i.e., compensating for the autocorrelation when estimating the standard error and degrees of freedom), or by attempting to remove its influence prior to analysis (i.e., prewhitening the data). Examples of the former approach include bootstrapping from the residuals of the fitted model (Zalesky et al., 2014; Chang and Glover, 2010), and performing phase-randomization (Hindriks et al., 2016). Here we focus on an approach for removing the effects of autocorrelation. Prewhitening is a common strategy for removing temporal autocorrelation from a signal, particularly in the context of the GLM (Monti, 2011). Therefore, in this study, we investigate whether prewhitening, applied to the data prior to using the sliding-window technique, also improves the estimation of TVC.

To illustrate, let us consider a single univariate times series $\boldsymbol{y} = (y_1 \ldots y_n)^T$, where n represents the total number of time points. Assume further that $\boldsymbol{y}$ has covariance $\boldsymbol{V}\sigma^2$, which is formulated according to the chosen autocorrelation model. If we assume that the value of $\boldsymbol{V}$ is known, we can pre-multiply $\boldsymbol{y}$ by $\boldsymbol{W} = \boldsymbol{V}^{-1/2}$. Now $\boldsymbol{Wy}$ has covariance matrix $\sigma^2\boldsymbol{WVW} = \sigma^2\boldsymbol{I}$. Hence, the effects of autocorrelation have been removed, and the transformed data set $\tilde{\boldsymbol{y}} = \boldsymbol{Wy}$ is independent and identically distributed. By applying the sliding-window technique after prewhitening has been performed on both time series, one can thereby avoid the increased variance in the estimated time-varying connectivity that arises due to autocorrelation; see Eq. (12).

In practice, we generally do not know the value of $\boldsymbol{V}$. Hence, we need to estimate it from the data using either the Yule-Walker equations for AR(p) models or maximum-likelihood methods for ARMA models (Shumway and Stoffer, 2011). In our simulations and application, we will assume an AR(1) model throughout and use the Yule-Walker equation approach towards estimating the model parameters (see Appendix A). In addition, we will explore both the use of a fixed value of p, and an optimal order determined using the Bayesian Information Criteria (*BIC*).

To further characterize the AR(1) filter used in the simulations and experiment, we can study its spectral density. It can be expressed in closed form as follows:

$$S_i(f) = \frac{\sigma_i^2}{(1 - 2\phi cos(2\pi f) + \phi^2)} \tag{13}$$

where the ϕ, σ_i, f, and $S_i(f)$ represent the AR parameter, variance of w_i, frequency, and power spectral density of y_i, respectively. This illustrates that in the case of fMRI time series, since $\phi > 0$, the power spectrum density is governed mostly by low frequency components. This result is derived in Appendix B.

2.5. The autocorrelation problem - empirical illustration using simulations

Here, we performed a series of simulation studies in order to evaluate the impact of autocorrelation on the estimates of TVC. Specifically, we began by performing three simulation studies, where TVC was estimated between two time series with: (1) no correlation; (2) periodic, slowly varying correlation; and (3) correlation that quickly switches between a no-correlation state to an enhanced positive correlation state. Details of how each time series pair was modeled and the type of analysis performed to analyze the time series pair are presented below. The simulations here are done for a bivariate case for the sake of simplicity and transparency.

The number of time points n was set to 600 for all three simulation studies. For each simulation, we generated random data $\boldsymbol{y}_t = (y_{1,t}, y_{2,t})^T$ using a mean-zero multivariate normal distribution. In each case, the

covariance matrix at a given time point was set to

$$\mathbf{\Sigma}_t = \begin{pmatrix} 2 & p(t) \\ p(t) & 3 \end{pmatrix}$$

where the covariance term, $p(t)$, was allowed to vary across time for $t = 1,\ldots,n$. This allowed us to control the dynamic relationship between the two time series $y_{1,t}$ and $y_{2,t}$. In addition, each time series was simulated to follow an AR(1) process with parameter ϕ; see Appendix C for additional details. This enabled us to control the degree of autocorrelation within each time series. For each simulation, the values of $p(t)$ and ϕ were set as follows:

Simulation 1. $p(t) = 0$ for all values of $t = 1,\ldots n$. Here the time series are uncorrelated across the entire time course, corresponding to null data. Values of ϕ varied from 0 to 0.9 in increments of 0.1.
Simulation 2. $p(t) = sin(t/\Delta)$ for $t = 1,\ldots n$, $\Delta = 1024/(2^k)$ and $k = 1, \ldots, 4$. Here the correlation changes in a slowly varying periodic manner. Values of ϕ were set to 0.3, 0.4, and 0.5.
Simulation 3. $p(t)$ is equal to a Gaussian kernel with mean 250 and standard deviation $15k$ for $t = 1, \ldots, n$ and $k = 1, \ldots, 4$. Thus, $p(t)$ is non-zero in an interval approximately within ± 3 standard deviations from 250. Here the data moves from a state of no correlation to an enhanced state of positive correlation, and back again. Values of ϕ were set to 0.3, 0.4, and 0.5.

For each choice of $p(t)$ and ϕ, simulations 1 through 3 were repeated 1000 times. A window length of 30 TRs (equivalent to about one minute of scan time) was chosen for the sliding-window technique. Previous studies have suggested that a window length of one minute may be optimal for the analysis of rs-fMRI (Leonardi and Van De Ville, 2015).

For all simulations we computed the maximum estimated correlation across time, as well as the mean squared error (MSE) between the estimated correlation and the true value (i.e., $\rho(t) = p(t)/\sqrt{6}$). Note that in Simulation 1, the MSE corresponds to the variance of the correlation across time, which is a common summary of TVC. We then repeated the three simulation studies using the prewhitened pairs of time series. Specifically, we used the Yule-Walker equations to estimate the parameters of the univariate AR(1) model, prewhitened the data, and applied the sliding-window technique to the prewhitened data.

Finally, we performed a fourth simulation to generate a data set with known number and transition patterns of brain states. We then used the data set to assess our ability to recover the underlying brain states.

Simulation 4. We used the SimTB toolbox (Erhardt et al., 2012), with the same simulation settings described in Allen et al. (2014). Data was generated for 10 ROIs using 4 discrete states to describe their underlying connectivity structure. In total there were five transisitions between states, with State 2 repeating. The simulation parameters (TR $= 2\,s$, 148 vol) were roughly matched to experimental data.

Importantly, this simulation creates an autocorrelation in the data that is partially based on brain hemodynamics, as SimTB mimics neuronal activation using a linear convolution with a canonical hemodynamic response function.

We were also interested in determining whether the use of prewhitening may remove the effects of hemodynamics, and thereby negatively affect the estimation of brain states. We thus estimated the TVC using two different variants of the simulated data: (*i*) non-prewhitened; and (*ii*) prewhitened using an AR(1) model. The sliding-window technique with window length of $\ell = 30$ was applied to each variant, yielding a $10 \times 10 \times 118$ array ($n - \ell = 148 - 30 = 118$) of time-varying correlation matrices. Next, we used *k*-means clustering with 4 centroids to estimate the true underlying brain states. As input we used the 45×118 matrix consisting of the lower triangular portion of the correlation matrices ($10(10-1)/2$ elements) at each of the 118 time points. We used the cluster centroids as estimates of true underling brain states.

2.6. The autocorrelation problem - empirical illustration using experimental data

2.6.1. Image acquisition

We used the Multi-Modal MRI Reproducibility Resource from the F.M. Kirby Research Center; colloquially referred to as Kirby21. A detailed description of the acquisition protocol can be found in Landman et al. (2011), and is briefly summarized here. The resource includes data from 21 healthy adults scanned on a 3T Philips Achieva scanner designed to achieve 80 mT/m maximum gradient strength with body coil excitation and an eight channel phased array SENSitivity Encoding (SENSE) (Pruessmann et al., 1999) head-coil for reception. All participants completed two scanning sessions on the same day, between which they briefly exited the scan room. A T1-weighted (T1w) Magnetization-Prepared Rapid Acquisition Gradient Echo (MPRAGE) structural run was acquired during both sessions (acquisition time = 6 *min*, TR/TE/TI = 6.7/3.1/842 *ms*, resolution = $1 \times 1 \times 1.2\ mm^3$, SENSE factor = 2, flip angle = 8°). A multi-slice SENSE-EPI pulse sequence (Stehling et al., 1991; Pruessmann et al., 1999) was used to acquire two rs-fMRI runs during each session. Each run consisted of 210 vol sampled every 2 *s* at 3 *mm* isotropic spatial resolution (acquisition time: 7 *min*, TE = 30 *ms*, SENSE acceleration factor = 2, flip angle = 75°, 37 axial slices collected sequentially with a 1 *mm* gap). Participants were instructed to rest comfortably while remaining still, and no other instruction was provided. One participant was excluded from further analyses due to excessive head motion.

2.6.2. Image processing

The data was preprocessed using SPM8 (Wellcome Trust Centre for Neuroimaging, London, United Kingdom) (Friston et al., 1994) and custom MATLAB (The Mathworks, Inc., Natick, MA) scripts. Four volumes were discarded at acquisition to allow for the stabilization of magnetization, and an additional volume was discarded prior to preprocessing. Slice-time correction was performed using the slice acquired at the middle of the TR as a reference, and rigid body realignment transformation was performed to adjust for head motion. Structural runs were registered to the first functional frame and spatially normalized to Montreal Neurological Institute (MNI) space using SPM8's unified segmentation-normalization algorithm (Ashburner and Friston, 2005). The estimated rigid body and nonlinear spatial transformations were applied to the rs-fMRI data, which were high-pass filtered using a cutoff frequency of 0.01 *Hz*. The rs-fMRI data was spatially smoothed using a 6 *mm* full-width-at-half-maximum (FWHM) Gaussian kernel, which is twice the nominal size of the rs-fMRI acquisition voxel.

The Group ICA of fMRI toolbox (GIFT) (http://mialab.mrn.org/software/gift; Medical Image Analysis Lab, Albuquerque, New Mexico) was used to estimate the number of independent components (ICs) present in the data, perform data reduction via principal component analysis (PCA) prior to independent component analysis (ICA), and perform group independent component analysis (GICA) (Calhoun et al., 2001) on the PCA-reduced data. Estimation of the number of ICs was performed using the minimum description length (MDL) criterion (Li et al., 2007). Across subjects and sessions, 56 was the maximum estimated number of ICs and 39 the median. Prior to GICA, the image mean was removed from each time point for each session, and three steps of PCA were performed. Individual session data were first reduced to 112 principal components (PCs), and the reduced session data were then concatenated within subjects in the temporal direction and further reduced to 56 PCs. Finally, the data were concatenated across subjects and reduced to 39 PCs. The dimensions of the individual session PCA (112) was chosen by doubling the estimated maximum IC number (56), to ensure robust back-reconstruction (Allen et al., 2011, 2012) of subject- and

session-specific spatial maps and time courses from the group-level independent components. ICA was repeated on these 39 group-level principal components 10 times, utilizing the Infomax algorithm with random initial conditions (Bell and Sejnowski, 1995). The resulting 390 ICs were clustered across iterations using a group average-link hierarchical strategy, and 39 aggregate spatial maps were defined as the modes of the clusters. Subject- and session-specific spatial maps and time courses were generated from these aggregate ICs using the GICA3 algorithm.

The spatial distribution of each of the group-level, aggregate ICs was compared to a publicly available set of 100 unthresholded t-maps of ICs estimated using rs-fMRI data collected from 405 healthy participants (Allen et al., 2014). These t-maps have already been classified as resting-state networks (RSNs) or noise by a group of experts, and the 50 components classified as RSNs have been organized into seven large functional groups: visual (Vis), auditory (Aud), somatomotor (SM), default mode (DMN), cognitive-control (CC), sub-cortical (SC) and cerebellar (Cb) networks. For each of the group-level spatial maps, we calculated the percent variance explained by the seven sets of RSNs. The functional assignment of each Kirby component was determined by the set of components that explained the most variance, and if the top two sets of RSNs explained less than 50% of the variance in a Kirby component, the component was labeled as noise. In total 21 of the 39 components were determined to belong to a RSN. Subject- and run-specific time series from these components then served as input for the dynamic FC analyses described below.

2.6.3. Prewhitening

TVC was estimated using three variants of the experimental data. Specifically, (*i*) the original (non-prewhitened) data, (*ii*) prewhitened data assuming an AR(1) model, and (*iii*) pre-whitened data using the optimal AR order chosen based on the Bayesian Information Criteria (*BIC*).

2.6.4. Time varying functional connectivity analysis

The sliding-window technique with $\ell = 30$ was applied to the data for each subject and run. This gave rise to a $21 \times 21 \times 210$ array of correlation matrices for each time point. Following the method of Allen and colleagues (Allen et al., 2014), we used *k*-means clustering to estimate recurring brain states across subjects for each run. First, we reorganized the lower triangular portion of each subject's dynamic correlation data into a matrix with dimensions 210×210; here the row dimension corresponds to the number of elements in the lower triangular portion of the matrix (i.e., $21(21-1)/2$), and the column dimension corresponds to the number of time points. Then we concatenated the data from all subjects into a matrix with row dimensions 210 and column dimensions $(210 \times 20 = 4200)$. Finally, we applied *k*-means clustering to the concatenated data set, where each of the resulting cluster centroids were assumed to represent a recurring brain state. Based on prior analysis of the data set (Choe et al., 2017), we chose the number of centroids to be equal to two, representing two distinct brain states. The same analysis was performed for both of the imaging sessions.

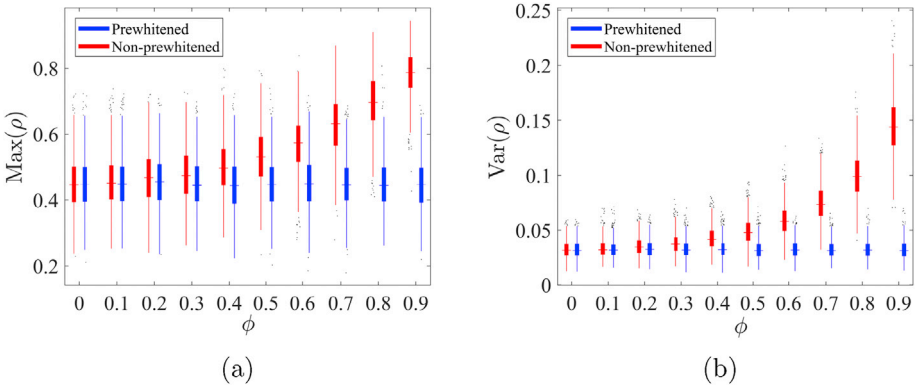

Fig. 1. Results of Simulation 1. Box plots of **(a)** the maximum sample correlation and **(b)** the variance of sample correlation for Simulation 1, where the time series are uncorrelated across the entire time course, corresponding to null data. Values of AR parameter, ϕ, varied from 0 to 0.9 in increments of 0.1.

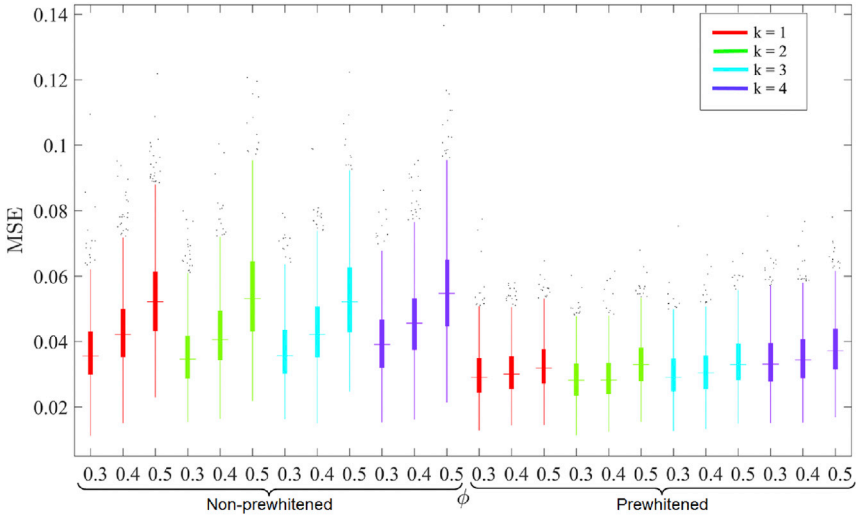

Fig. 2. Results of Simulation 2. Box plots of the mean squared error (MSE) between the estimated correlation and the true value of sample correlation for Simulation 2, where $p(t) = sin(t/\Delta)$ with $\Delta = 1024/(2^k)$. Values of AR parameter, ϕ, were set to 0.3, 0.4, and 0.5.

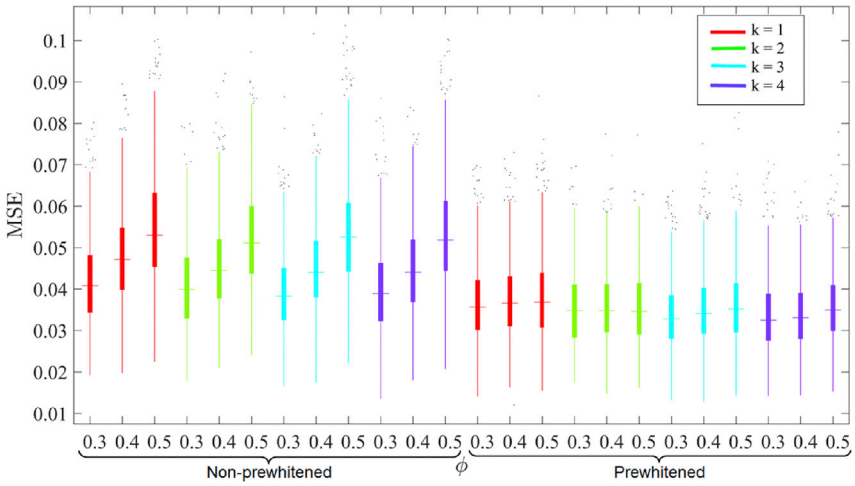

Fig. 3. Results of Simulation 3. Box plots of the mean squared error (MSE) between the estimated correlation and the true value of sample correlation for Simulation 3, where $p(t)$ is a Gaussian kernel with mean 250 and standard deviation of 15k. Values of AR parameter, ϕ, were set to 0.3, 0.4, and 0.5.

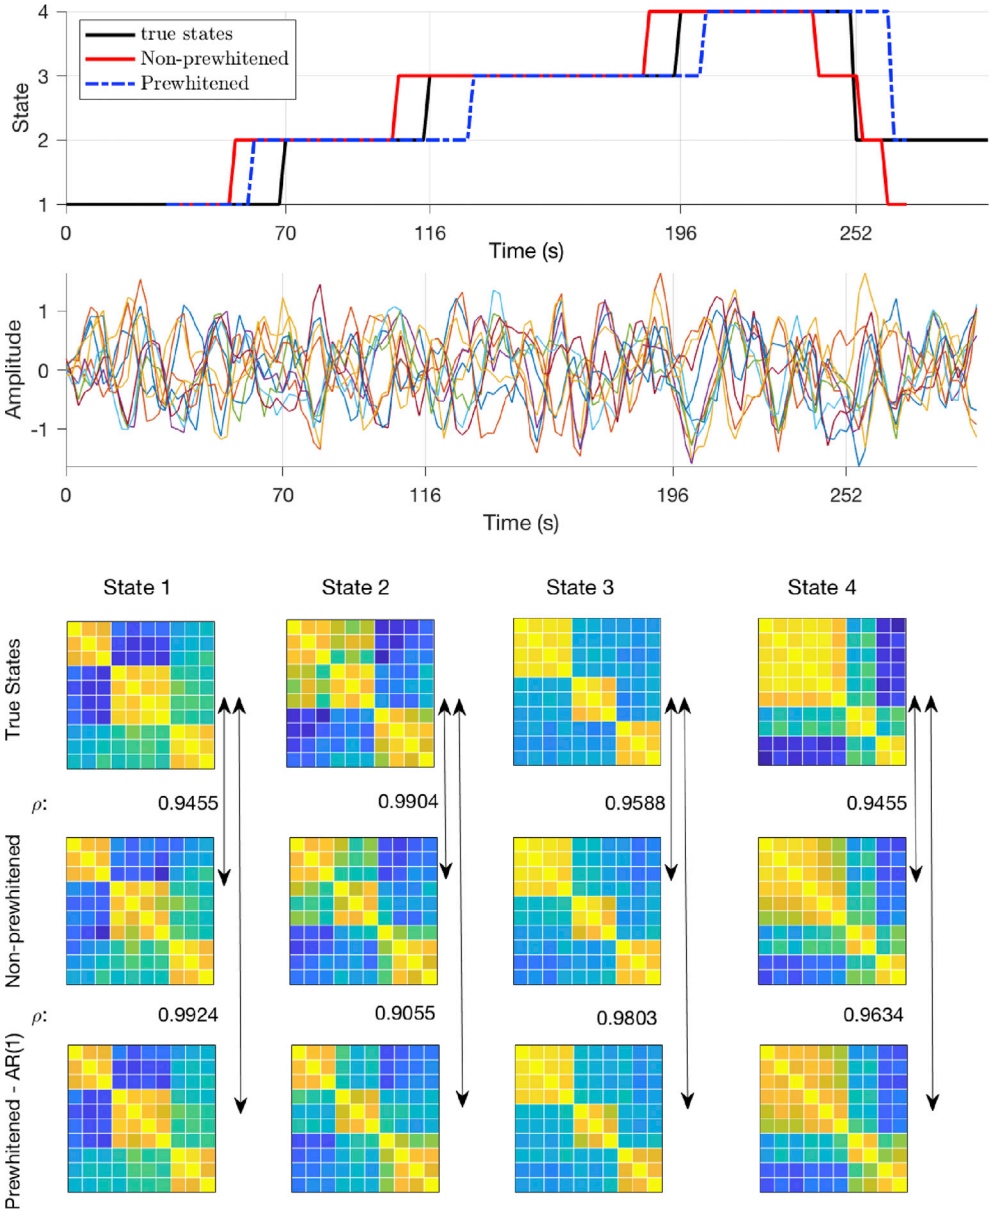

Fig. 4. Results of Simulation 4. Time varying connectivity was generated for 10 ROIs whose connectivity structure altered between 4 states, labeled here as "True States" (middle panel). The timing of the transitions are shown in black in the top panel and the generated time courses are shown in the panel below. Using the sliding-window technique with $w = 30$ and k-means clustering, the states were estimated both using the raw (non-prewhitened) and prewhitened data. The correlation between each estimated state and the corresponding true state are indicated by arrows. Consistently these correlations are higher for the prewhitened data, except for State 2. The estimated states using the prewhitened data (top panel, dashed blue) showed less transitions compared to the non-prewhitened data (top panel, red) and estimated state-transitions closer to the true states (top panel, black). In addition, the state-transitions estimated from the raw data consistently starts to occur earlier then the true state-transition timings.

3. Results

3.1. Simulations

3.1.1. Simulation 1

For the first simulation, $p(t)$, which controls the true dynamic correlation between the pair of time series, is set to 0. Hence, the time series are uncorrelated across time. Fig. 1a shows box plots of the maximum sample correlation across simulations for both the non-prewhitened data and the prewhitened data. As expected, the maximum correlation increases as a function of ϕ for the non-prewhitened data. In particular, for values of ϕ above 0.3 the correlations begins to steadily increase, which is consistent with the theoretical results. In contrast, prewhitening the data prevents similar increase of correlation as a function of the value of ϕ. Similarly, Fig. 1b shows that the variance of the sample correlation (equivalent to the MSE in this case) also increases as a function of ϕ for the non-prewhitened data, while prewhitening lowers the estimated variance. Together these results suggest that prewhitening lessens the risk of observing extreme correlations in the null setting.

3.1.2. Simulation 2

For the second simulation study, we set $p(t) = sin(t/\Delta)$, where $\Delta = 1024/(2^k)$ and $k = 1, \ldots, 4$ to simulate a correlation that changes in a slowly varying periodic manner. As seen in Fig. 2, the MSE between the estimated correlation and the true value of $p(t)$ increases as a function of ϕ. In contrast, prewhitening the data circumvents these issues as expected. This simulation indicates that when a true correlation exists between two time series, prewhitening the data reduces the MSE and hence allows us to obtain a better estimate of the true underlying correlation.

3.1.3. Simulation 3

For the third simulation study, we let $p(t)$ correspond to a Gaussian kernel with mean 250 and standard deviation of $15k$ with $k = 1, 2, 3, 4$. Fig. 3 shows how the MSE between the estimated correlation and the true value of $p(t)$ increases as a function of ϕ. Again, prewhitening the data reduces the MSE and hence we obtain a better estimate of the true correlation.

3.1.4. Simulation 4

For the fourth simulation study, we investigated two variants of the data: (*i*) non-prewhitened; and (*ii*) prewhitened using an $AR(1)$ model. Fig. 4 shows the true underlying states (middle panel). The timing of the transitions are shown in black in the top panel and the generated time courses (prior to prewhitening) are shown in the panel below. As shown in Fig. 4, the estimates of the brain states and the state-transitions displayed a high degree of reliability for both variants. Interestingly, the prewhitened data led to better estimates of both the states and state-transitions. The correlation between each estimated state and the corresponding true state are shown along the arrows in the bottom panels. Except for State 2, the correlations are consistently higher for the prewhitened data. In addition, the estimated states using the raw data showed more transitions (top panel, red) than the prewhitened data (top panel, dashed blue), which estimated the transitions closer to the truth (top panel, black). In general, the state-transitions estimated from the raw data consistently occurs earlier in comparison to the actual state-transition timings.

3.2. Experimental data

We analyzed three different variants of the preprocessed Kirby21 data: (*i*) non-prewhitened; (*ii*) prewhitened assuming an AR(1) model; and (*iii*) prewhitened using an optimal AR model determined using the *BIC*.

Fig. 5 shows the estimated AR(1) parameters for each RSN and subject. The average AR(1) parameter for this data set was 0.5994. Hence, the inflation of variance of the sample correlation and maximum estimated correlation stands to be considerable according to the results shown both in Eq. (11) and Fig. 1. It is interesting to note that there are substantial differences in autocorrelation between subjects, regions, and sessions. Fig. 6 demonstrates the optimal AR order that were selected using the *BIC*. The mean of the orders used for Session 1 was 2.74 and Session 2 was 2.481. This indicates an average AR order of approximately 3. However, note again there are significant differences across both subjects and regions.

Throughout the study, we concentrated on identifying two distinct brain states by setting the number of clusters in the *k*-means algorithm to 2. Fig. 7 shows the resulting estimated brain states for Sessions 1 and 2, for the 3 variants of the Kirby 21 data. Using the between-session spatial correlation as a measure of reliability, we found that the two brain states were highly reliable across sessions regardless of the type of prewhitening that was performed. State 2 was characterized by stronger correlations (both positive and negative) relative to State 1. Moderate to strong negative correlations between sensory systems (auditory, somatomotor, and visual) components were present in State 2 but were reduced in State 1.

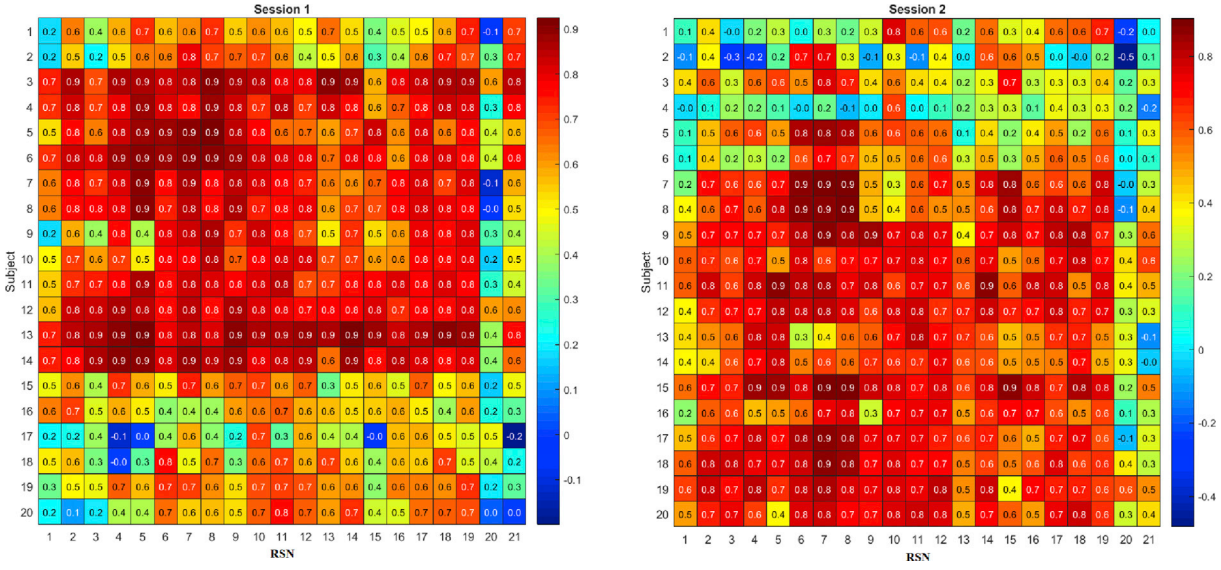

Fig. 5. Estimates of the AR(1) parameter ϕ for the Kirby 21 data. Each row consists of a subject and each column a RSN. Results are shown separately for each session.

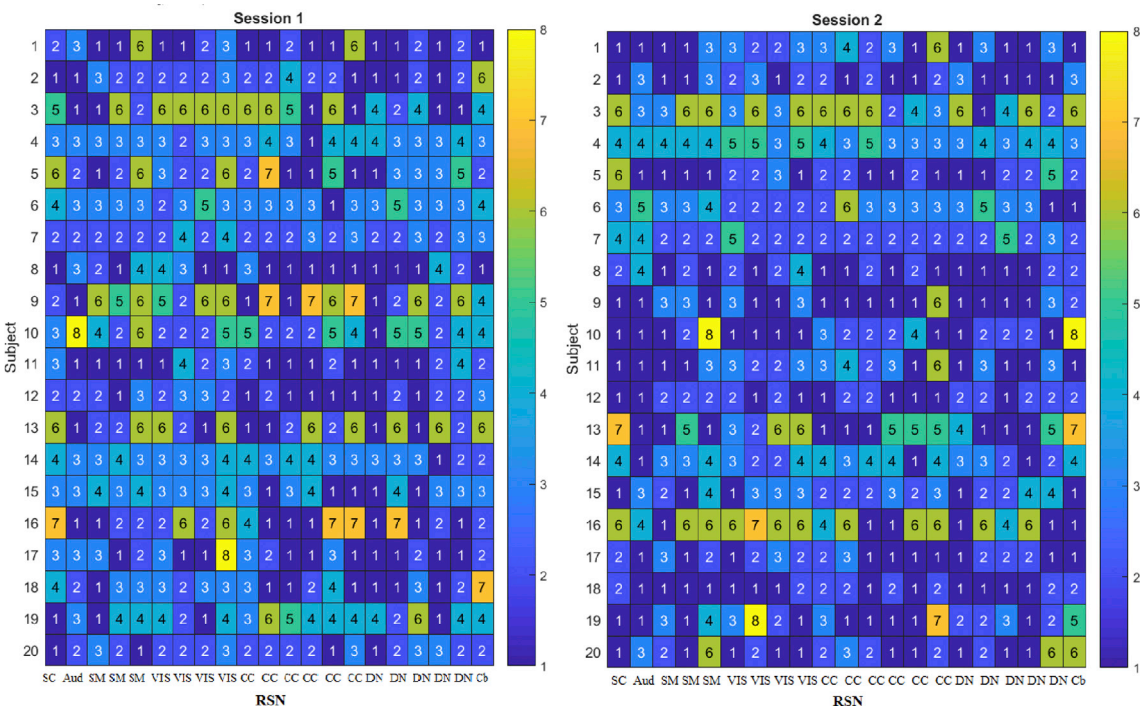

Fig. 6. The optimal AR orders identified using BIC for each session, subject, and RSN of the Kirby21 dataset.

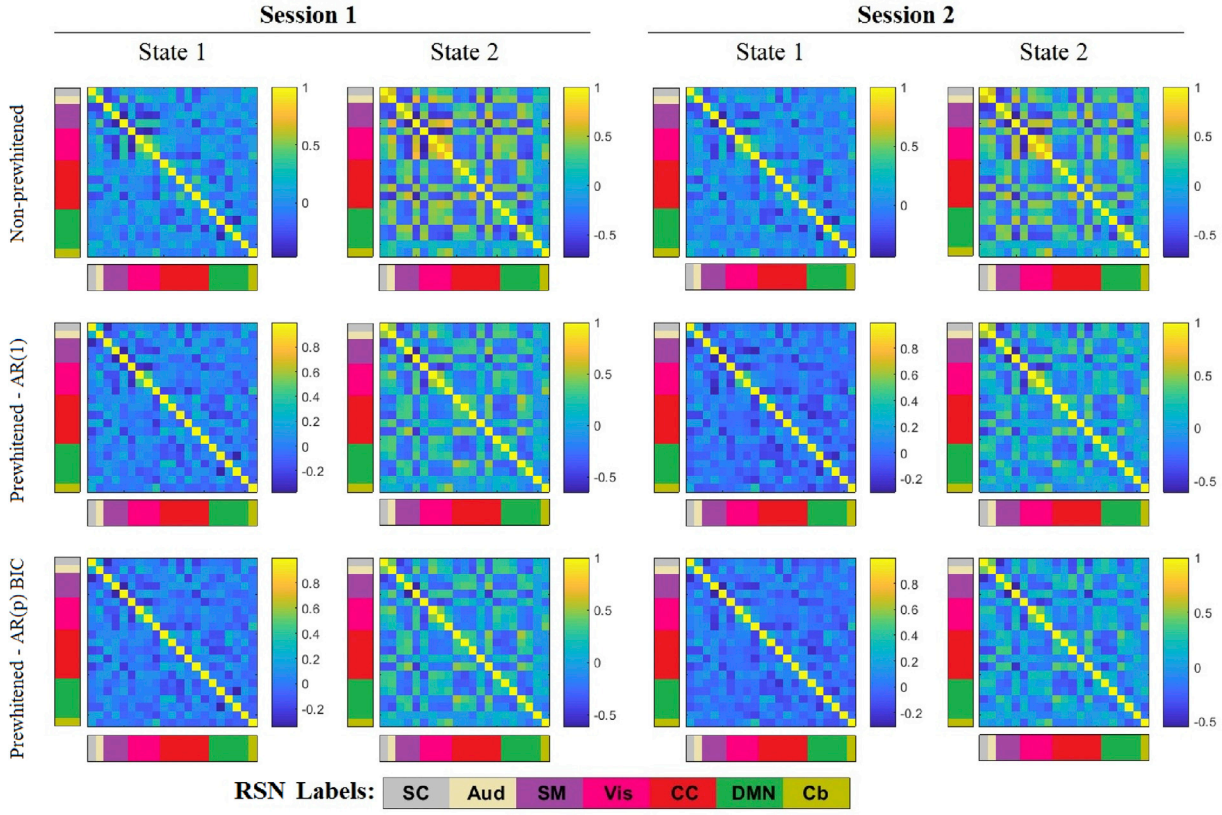

Fig. 7. Brain states estimated using sliding-windows + *k*-means of the Kirby21 data set. For each session two brain states are estimated. The results are represented for the non-prewhitened data (top), the data prewhitened using an AR(1) model (middle), and the data prewhitened using an optimal AR(p) model (bottom). The brain states are organized into seven large functional groups: visual (Vis), auditory (Aud), somatomotor (SM), default mode (DMN), cognitive-control (CC), sub-cortical (SC) and cerebellar (Cb) networks.

Though the ground truth is not known, we can see in Fig. 8 that clear differences between the two prewhitened cases and the non-prewhitened cases exist. The differences appear largest in the sensory systems components, with differences in certain RSNs taking values up to 0.3. However, the impact of prewhitening seems to be much greater than that of the choice of *AR*-model. For each variant, the estimates of the brain states are similar across sessions, indicating a high degree of reliability in brain state estimates.

4. Discussion

There is increasing evidence that TVC provides important information about brain function (Chang and Glover, 2010; Hutchison et al., 2013a; Preti et al., 2016; Tagliazucchi et al., 2012; Thompson et al., 2013; Allen et al., 2014). To date, a number of different methods have been proposed for estimating TVC, with the sliding-window technique currently being the most commonly used approach. Although multiple previous studies have performed in-depth investigations of the effects of nuisance confounders (e.g., motion) on TVC estimates, to date there has not been a similarly thorough investigation of the effects of autocorrelation on TVC estimates. In this paper, we show both theoretically and empirically that the presence of autocorrelation leads to increased estimates of time-varying correlation, even in the null setting, which in turn increases the risk of false positive findings and can potentially have detrimental effects on the estimation of brain states.

When assessing the quality of an estimator it is important to consider its sampling variability. This refers to how much the estimated statistic (i.e., TVC) varies from sample to sample. The lower the sampling variability, the more precise the subsequent inference (e.g., confidence intervals and hypothesis tests). As in any statistical inference problem, the estimated TVC can be sensitive to the procedure used to obtain the estimate. For the sliding-window technique it is well known that the sampling variability depends on the window length, with smaller windows giving rise to larger sampling variability. In this study, we show that the sampling variability is also affected by the existence of autocorrelation within time series (Bartlett, 1935), with increased autocorrelation associated with increased sampling variability. Because the sliding-window technique estimates a series of correlations across time, we can think of them as 'repeated samples' from the sampling distribution for a standard correlation. Therefore, as the length of the time series increases, the estimates will eventually sample values across the entire range of the sampling distribution. Thus, even in the absence of a true correlation, we will obtain estimates across the range of feasible values under the null distribution. As this range increases so does the risk for obtaining artificially large estimates.

Our simulations show that when the time series are truly uncorrelated, the variance of the sample correlation increases as a function of the autocorrelation (Fig. 1). However, when prewhitening is used to remove the effects of autocorrelation, the variance of the sample correlation is reduced and no longer positively correlated to the value of ϕ. Further, our simulations illustrate that in the presence of a true TVC, prewhitening still allows for a more accurate estimation of TVC (Figs. 2 and 3).

The impact of autocorrelation has been extensively studied in the context of task-fMRI (e.g., Bullmore et al. (1996); Purdon and Weisskoff (1998); Woolrich et al. (2001); Lund et al. (2006)). Here the data is prewhitened in order to make the residuals from a general linear model (GLM) analysis, after removal of task-related neuronal activity, uncorrelated. The use of prewhitening in the analysis of rs-fMRI data is less explored (Christova et al., 2011; Arbabshirani et al., 2014), and also more complex. The reason is that in contrast to task-fMRI, it is not possible to remove important neuronal activity prior to performing prewhitening. Thus, in the context of rs-fMRI it is reasonable to assume that

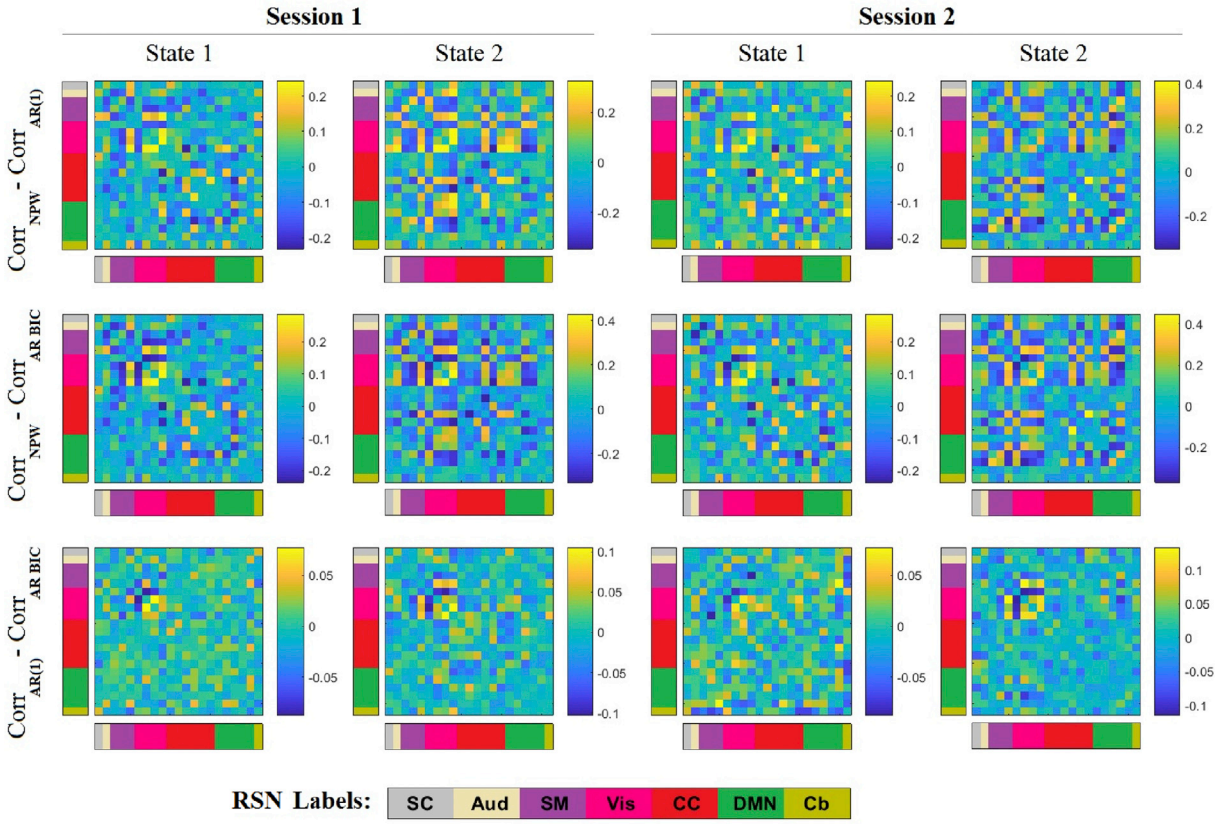

Fig. 8. The difference between the estimated brain states shown in Fig. 7. The top row shows the difference between the non-prewhitened (NPW) data and the data prewhitened using an AR(1) model for each brain state and session. The middle shows similar results comparing the non-prewhitened data and the data prewhitened using an optimal AR(p) model. Finally, the bottom row compares the two prewhitened data sets.

the autocorrelation present in the signal is partially due to physiological noise and partially due to brain hemodynamics. Therefore, it is not clear that one actually wants to make the signal entirely uncorrelated, as this may remove the effects of important neuronal information related to the hemodynamic response function (HRF). Simulation 4 was constructed to illustrate such a situation, and it appears that performing prewhitening did not have detrimental effects on the results (Fig. 4). This may perhaps also help explain why using AR models of different orders did not provide substantially different results (Figs. 7 and 8), as partial removal of autocorrelation may be sufficient. This is a topic of future work, and interestingly a paper recently appeared that showed that HRF variability can potentially confound resting-state fMRI functional connectivity (Rangaprakash et al., 2018).

As the ground truth is not known, the investigation of the TVC estimation in experimental data is more challenging. However, we observe clear differences between estimates of brain states (see Fig. 8) that we hypothesize are due to the manner in which the autocorrelation was handled. To illustrate, Fig. 5 shows that the estimated AR(1) parameter for each RSN and subject, takes an average value equal to 0.5994. These results indicate that the level of autocorrelation in the data reach levels where the variability in estimated time-varying correlation is large, which can potentially increase the risks of false positive findings if the autocorrelation is not properly accounted for. This can be seen by comparing the empirically estimated values of ϕ with the equivalent values in Fig. 1. Clearly, we are in a range where our simulations have shown increased variability that can potentially be addressed using prewhitening. We believe this is leading to the observed differences in brain states, and that the prewhitened states are closer to the ground truth. Interestingly, there are substantial differences in autocorrelation between subjects, regions, and sessions. Hence the sampling variability will differ depending on which RSNs are being compared.

In fMRI research, it is commonly assumed that the order of the AR process is fixed. In this study, we also computed the optimal AR order using the BIC, and the results show that the average AR order is approximately 3 (Fig. 6). However, significant differences across subjects and regions were observed suggesting that the optimal AR model may in fact not be fixed. Interestingly, while we can see clear differences between estimated brain states when the data has been prewhitened versus when it has not, the effect of which AR model selection procedure was used is relatively small (Figs. 7 and 8). This indicates that perhaps the type of prewhitening performed is less important than that the procedure is actually performed at all.

A general take-home point of the simulations and application to experimental data is that in order for inference based on fMRI data to be valid, the autocorrelation needs to be properly taken into account. One way to do this is to incorporate the additional uncertainty related to the autocorrelation into the subsequent inferential procedure (i.e., taking into consideration the fact that the time-varying connectivity is estimated with a higher degree of variability, and with fewer degrees of freedom). For example, one can compensate for the autocorrelation when estimating the standard error of the point estimate. However, this aproach becomes less optimal when performing a secondary analysis on the estimated TVC (e.g., k-means clustering to determine brain states) as the input data will be noisier and therefore risks giving rise to noisier output as well. Alternatively, one can attempt to remove its influence prior to analysis by prewhitening the data. This is the approach we have taken in this work.

Acknowledgments

The work presented in this paper was supported in part by NIH grants R01 EB016061, R01 EB026549, and P41 EB015909 from the National Institute of Biomedical Imaging and Bioengineering, R21 NS104644 from the National Institute of Neurological Disorders and Stroke, and Craig H. Neilsen Foundation (Project Number 338419).

Appendix A

Here we describe the Yule-Walker approach towards estimating the parameters in a univariate AR process. Let us assume we are working with an AR(p) process, given by the following equation:

$$y_t = \sum_{i=1}^{p} \phi_i y_{t-i} + \varepsilon_t \tag{14}$$

Note that here, without loss of generality and for the sake of simplification, we assume a zero-mean model; nevertheless, the estimation of mean is straightforward.

By multiplying both sides of Eq. (14) by y_{t-h}, for $h = 0, \ldots p$, and taking the expected value, we obtain $p + 1$ difference equations written as follows:

$$\gamma_0 = \phi_1 \gamma_1 + \ldots + \phi_p \gamma_p + \sigma^2 \tag{15}$$

$$\gamma_h = \phi_1 \gamma_{h-1} + \ldots + \phi_p \gamma_{h-p}, \qquad h = 1, 2, \ldots, p \tag{16}$$

where γ_h represents the lag h autocorrelation.

Thus, the Yule-Walker equations can be expressed as follows:

$$\gamma_h = \sum_{i=1}^{p} \phi_i \gamma_{h-i} + \sigma^2 \delta_{h,0}, \qquad h = 1, 2, \ldots, p \tag{17}$$

where $\delta_{0,0} = 1$ and $\delta_{h,0} = 0$ for $h > 0$. This equation can be written in matrix form as follows:

$$\begin{bmatrix} \gamma_1 \\ \gamma_2 \\ \gamma_3 \\ \vdots \\ \gamma_p \end{bmatrix} = \begin{bmatrix} \gamma_0 & \gamma_{-1} & \gamma_{-2} & \cdots \\ \gamma_1 & \gamma_0 & \gamma_{-1} & \cdots \\ \gamma_2 & \gamma_1 & \gamma_0 & \cdots \\ \vdots & \vdots & \vdots & \ddots \\ \gamma_{p-1} & \gamma_{p-2} & \gamma_{p-3} & \cdots \end{bmatrix} \begin{bmatrix} \phi_1 \\ \phi_2 \\ \phi_3 \\ \vdots \\ \phi_p \end{bmatrix} \tag{18}$$

or alternatively,

$$\boldsymbol{\gamma}_p = \boldsymbol{\Gamma}_p \boldsymbol{\Phi}. \tag{19}$$

The solution can be obtained by computing:

$$\widehat{\boldsymbol{\Phi}} = \widehat{\boldsymbol{\Gamma}}_p^{-1} \widehat{\boldsymbol{\gamma}}_p \tag{20}$$

Once $\boldsymbol{\Phi}$ is estimated, we can now pre-whiten the data by computing:

$$\widetilde{y}_t = y_t - \sum_{i=1}^{p} \widehat{\phi}_i y_{t-i}. \tag{21}$$

In the case of AR(1) this can be written as:

$$\widetilde{y}_t = y_t - \widehat{\phi} y_{t-1} \tag{22}$$

Hence, in the prewhitening process, the identification of the AR(1) model is first required, followed by the computation of Eq. (22).

Appendix B

In this section, we provide a the proof of expression shown in Eq. (13). Recall that the spectral density of a time series can be written as follows:

$$S(f) = \sum_{h=-\infty}^{+\infty} \gamma(h) e^{-2\pi j f h} \tag{23}$$

For an AR(1) process,

$$\gamma(h) = \sigma^2 \frac{\phi^{|h|}}{(1-\phi^2)} \tag{24}$$

Thus, the spectral density can be written as:

$$\begin{aligned}
S(f) &= \frac{\sigma^2}{(1-\phi^2)} \sum_{h=-\infty}^{+\infty} \phi^{|h|} e^{-2\pi j f h} \\
&= \frac{\sigma^2}{(1-\phi^2)} \left(1 + \sum_{h=1}^{+\infty} \phi^h \left(e^{-2\pi j f h} + e^{2\pi j f h}\right)\right) \\
&= \frac{\sigma^2}{(1-\phi^2)} \left(1 + \frac{\phi e^{-2\pi j f h}}{1-\phi e^{-2\pi j f h}} + \frac{\phi e^{2\pi j f h}}{1-\phi e^{2\pi j f h}}\right) \\
&= \frac{\sigma^2}{(1-\phi^2)} \frac{1-\phi e^{-2\pi j f h}\phi e^{2\pi j f h}}{(1-\phi e^{-2\pi j f h})(1-\phi e^{2\pi j f h})} \\
&= \frac{\sigma^2}{(1-2\phi cos(2\pi f)+\phi^2)}.
\end{aligned}$$

The corresponding pre-whitening filter for this process will then be a high-pass filter ($= \sqrt{(1/S(f))}$). In Supplemental Material Fig. 1, we show examples of the pre-whitening filter for various choices of ϕ. Clearly, as value of ϕ increases the amount of low-frequency components removed by the filter similarly increases.

Appendix C

Here we describe how data is generated in our simulation studies. Assume $y_{1,t}$ and $y_{2,t}$ both follow an AR(1) process:

$$\begin{aligned}
y_{1,t} &= \phi y_{1,t-1} + w_{1,t}; \qquad w_{1,t} \sim \mathcal{N}\left(0, \sigma_1^2\right) \\
y_{2,t} &= \phi y_{2,t-1} + w_{2,t}; \qquad w_{2,t} \sim \mathcal{N}\left(0, \sigma_2^2\right)
\end{aligned} \tag{25}$$

In matrix form this can be written:

$$\begin{pmatrix} y_{1,t} \\ y_{2,t} \end{pmatrix} = \begin{pmatrix} \phi & 0 \\ 0 & \phi \end{pmatrix} \begin{pmatrix} y_{1,t-1} \\ y_{2,t-1} \end{pmatrix} + \begin{pmatrix} w_{1,t} \\ w_{2,t} \end{pmatrix} \tag{26}$$

or

$$\boldsymbol{y}_t = \boldsymbol{\Phi} \boldsymbol{y}_{t-1} + \boldsymbol{w}_t, \tag{27}$$

where $\boldsymbol{\Phi}$ is (2×2) coefficient matrix and $\boldsymbol{w}_t$ is an (2×1) zero mean white noise vector process (serially uncorrelated or independent) with time invariant covariance matrix $\boldsymbol{\Sigma}_w$.

In the generation of the synthetic data, we first fix the values of the covariance matrix $\mathbf{\Sigma}_w$ of the innovations driving the AR-process and the coefficient matrix $\mathbf{\Phi}$, therafter we generate $\mathbf{w}_t$ for all t, and finally we sequentially generate y_t according to Eq. (27).

Appendix D. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.neuroimage.2019.04.042.

References

Allen, E.A., Damaraju, E., Plis, S.M., Erhardt, E.B., Eichele, T., Calhoun, V.D., 2014. Tracking whole-brain connectivity dynamics in the resting state. Cerebr. Cortex 24, 663–676.

Allen, E.A., Erhardt, E.B., Damaraju, E., Gruner, W., Segall, J.M., Silva, R.F., Havlicek, M., Rachakonda, S., Fries, J., Kalyanam, R., et al., 2011. A baseline for the multivariate comparison of resting-state networks. Front. Syst. Neurosci. 5.

Allen, E.A., Erhardt, E.B., Wei, Y., Eichele, T., Calhoun, V.D., 2012. Capturing inter-subject variability with group independent component analysis of fmri data: a simulation study. Neuroimage 59, 4141–4159.

Arbabshirani, M.R., Damaraju, E., Phlypo, R., Plis, S., Allen, E., Ma, S., Mathalon, D., Preda, A., Vaidya, J.G., Adali, T., et al., 2014. Impact of autocorrelation on functional connectivity. Neuroimage 102, 294–308.

Ashburner, J., Friston, K.J., 2005. Unified segmentation. Neuroimage 26, 839–851.

Bartlett, M., 1935. Some aspects of the time-correlation problem in regard to tests of significance. J. R. Stat. Soc. 98, 536–543.

Bell, A.J., Sejnowski, T.J., 1995. An information-maximization approach to blind separation and blind deconvolution. Neural Comput. 7, 1129–1159.

Bolton, T.A., Tarun, A., Sterpenich, V., Schwartz, S., Van De Ville, D., 2018. Interactions between large-scale functional brain networks are captured by sparse coupled hmms. IEEE Trans. Med. Imaging 37, 230–240.

Bullmore, E., Brammer, M., Williams, S.C., Rabe-Hesketh, S., Janot, N., David, A., Mellers, J., Howard, R., Sham, P., 1996. Statistical methods of estimation and inference for functional mr image analysis. Magn. Reson. Med. 35, 261–277.

Calhoun, V., Adali, T., Pearlson, G., Pekar, J., 2001. A method for making group inferences from functional mri data using independent component analysis. Hum. Brain Mapp. 14, 140–151.

Calhoun, V.D., Miller, R., Pearlson, G., Adal, T., 2014. The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery. Neuron 84, 262–274.

Chang, C., Glover, G.H., 2010. Time–frequency dynamics of resting-state brain connectivity measured with fmri. Neuroimage 50, 81–98.

Choe, A.S., Nebel, M.B., Barber, A.D., Cohen, J.R., Xu, Y., Pekar, J.J., Caffo, B., Lindquist, M.A., 2017. Comparing test-retest reliability of dynamic functional connectivity methods. Neuroimage 158, 155–175.

Christova, P., Lewis, S., Jerde, T., Lynch, J., Georgopoulos, A., 2011. True associations between resting fmri time series based on innovations. J. Neural Eng. 8, 046025.

Cribben, I., Haraldsdottir, R., Atlas, L.Y., Wager, T.D., Lindquist, M.A., 2012. Dynamic connectivity regression: determining state-related changes in brain connectivity. Neuroimage 61, 907–920.

Cribben, I., Wager, T.D., Lindquist, M.A., 2013. Detecting functional connectivity change points for single-subject fmri data. Front. Comput. Neurosci. 7.

Damaraju, E., Allen, E.A., Belger, A., Ford, J.M., McEwen, S., Mathalon, D.H., Mueller, B.A., Pearlson, G.D., Potkin, S.G., Preda, A., Turner, J.A., Vaidya, J.G., van Erp, T.G., Calhoun, V.D., 2014. Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia. Neuroimage: Clinical 5, 298–308.

Eavani, H., Satterthwaite, T.D., Gur, R.E., Gur, R.C., Davatzikos, C., 2013. Unsupervised learning of functional network dynamics in resting state fmri. In: International Conference on Information Processing in Medical Imaging. Springer, pp. 426–437.

Erhardt, E.B., Allen, E.A., Wei, Y., Eichele, T., Calhoun, V.D., 2012. Simtb, a simulation toolbox for fmri data under a model of spatiotemporal separability. Neuroimage 59, 4160–4167.

Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J., Frith, C.D., Frackowiak, R.S., et al., 1994. Statistical parametric maps in functional imaging: a general linear approach. Hum. Brain Mapp. 2, 189–210.

Georgopoulos, A.P., Mahan, M.Y., 2013. Fmri Data Analysis: State of Affairs and Challenges.

Hindriks, R., Adhikari, M.H., Murayama, Y., Ganzetti, M., Mantini, D., Logothetis, N.K., Deco, G., 2016. Can sliding-window correlations reveal dynamic functional connectivity in resting-state fmri? Neuroimage 127, 242–256.

Hutchison, R.M., Morton, J.B., 2015. Tracking the brain's functional coupling dynamics over development. J. Neurosci. 35, 6849–6859.

Hutchison, R.M., Womelsdorf, T., Allen, E.A., Bandettini, P.A., Calhoun, V.D., Corbetta, M., Della Penna, S., Duyn, J.H., Glover, G.H., Gonzalez-Castillo, J., et al., 2013a. Dynamic functional connectivity: promise, issues, and interpretations. Neuroimage 80, 360–378.

Hutchison, R.M., Womelsdorf, T., Gati, J.S., Everling, S., Menon, R.S., 2013b. Resting-state networks show dynamic functional connectivity in awake humans and anesthetized macaques. Hum. Brain Mapp. 34, 2154–2177.

Karahanoğlu, F.I., Van De Ville, D., 2015. Transient brain activity disentangles fmri resting-state dynamics in terms of spatially and temporally overlapping networks. Nat. Commun. 6, 7751.

Landman, B.A., Huang, A.J., Gifford, A., Vikram, D.S., Lim, I.A.L., Farrell, J.A., Bogovic, J.A., Hua, J., Chen, M., Jarso, S., et al., 2011. Multi-parametric neuroimaging reproducibility: a 3-t resource study. Neuroimage 54, 2854–2866.

Laumann, T.O., Snyder, A.Z., Mitra, A., Gordon, E.M., Gratton, C., Adeyemo, B., Gilmore, A.W., Nelson, S.M., Berg, J.J., Greene, D.J., et al., 2016. On the stability of bold fmri correlations. Cerebr. Cortex 27, 4719–4732.

Lenoski, B., Baxter, L.C., Karam, L.J., Maisog, J., Debbins, J., 2008. On the performance of autocorrelation estimation algorithms for fmri analysis. IEEE Journal of Selected Topics in Signal Processing 2, 828–838.

Leonardi, N., Van De Ville, D., 2015. On spurious and real fluctuations of dynamic functional connectivity during rest. Neuroimage 104, 430–436.

Li, Y.O., Adalı, T., Calhoun, V.D., 2007. Estimating the number of independent components for functional magnetic resonance imaging data. Hum. Brain Mapp. 28, 1251–1266.

Lindquist, M.A., Xu, Y., Nebel, M.B., Caffo, B.S., 2014. Evaluating dynamic bivariate correlations in resting-state fmri: a comparison study and a new approach. Neuroimage 101, 531–546.

Lindquist, M.A., et al., 2008. The statistical analysis of fmri data. Stat. Sci. 23, 439–464.

Liu, X., Duyn, J.H., 2013. Time-varying functional network information extracted from brief instances of spontaneous brain activity. Proc. Natl. Acad. Sci. Unit. States Am. 4392–4397, 201216856.

Lund, T.E., Madsen, K.H., Sidaros, K., Luo, W.L., Nichols, T.E., 2006. Non-white noise in fmri: does modelling have an impact? Neuroimage 29, 54–66.

Lurie, D., Kessler, D., Bassett, D., Betzel, R.F., Breakspear, M., Keilholz, S., Kucyi, A., Liégeois, R., Lindquist, M.A., McIntosh, A.R., et al., 2018. On the Nature of Resting Fmri and Time-Varying Functional Connectivity.

Monti, M.M., 2011. Statistical analysis of fmri time-series: a critical review of the glm approach. Front. Hum. Neurosci. 5, 28.

Preti, M.G., Bolton, T.A., Van De Ville, D., 2016. The dynamic functional connectome: state-of-the-art and perspectives. Neuroimage 41–54.

Pruessmann, K.P., Weiger, M., Scheidegger, M.B., Boesiger, P., et al., 1999. Sense: sensitivity encoding for fast mri. Magn. Reson. Med. 42, 952–962.

Purdon, P.L., Weisskoff, R.M., 1998. Effect of temporal autocorrelation due to physiological noise and stimulus paradigm on voxel-level false-positive rates in fmri. Hum. Brain Mapp. 6, 239–249.

Rangaprakash, D., Wu, G.R., Marinazzo, D., Hu, X., Deshpande, G., 2018. Hemodynamic response function (hrf) variability confounds resting-state fmri functional connectivity. Magn. Reson. Med. 1697–1713.

Rashid, B., Damaraju, E., Pearlson, G.D., Calhoun, V.D., 2014. Dynamic connectivity states estimated from resting fmri identify differences among schizophrenia, bipolar disorder, and healthy control subjects. Front. Hum. Neurosci. 8.

Shappell, H.M., Caffo, B.S., Pekar, J.J., Lindquist, M., 2019. Improved State Change Estimation in Dynamic Functional Connectivity Using Hidden Semi-markov Models bioRxiv , 519868.

Shumway, R.H., Stoffer, D.S., 2011. Time series regression and exploratory data analysis. In: Time Series Analysis and its Applications. Springer, pp. 47–82.

Stehling, M.K., Turner, R., Mansfield, P., 1991. Echo-planar imaging: magnetic resonance imaging in a fraction of a second. Science 254, 43–50.

Tagliazucchi, E., Balenzuela, P., Fraiman, D., Montoya, P., Chialvo, D.R., 2011. Spontaneous bold event triggered averages for estimating functional connectivity at resting state. Neurosci. Lett. 488, 158–163.

Tagliazucchi, E., Von Wegner, F., Morzelewski, A., Brodbeck, V., Laufs, H., 2012. Dynamic bold functional connectivity in humans and its electrophysiological correlates. Front. Hum. Neurosci. 6.

Thompson, G.J., Magnuson, M.E., Merritt, M.D., Schwarb, H., Pan, W.J., McKinley, A., Tripp, L.D., Schumacher, E.H., Keilholz, S.D., 2013. Short-time windows of correlation between large-scale functional brain networks predict vigilance intraindividually and interindividually. Hum. Brain Mapp. 34, 3280–3298.

Vidaurre, D., Smith, S.M., Woolrich, M.W., 2017. Brain network dynamics are hierarchically organized in time. Proc. Natl. Acad. Sci. Unit. States Am. 114, 12827–12832.

Woolrich, M.W., Ripley, B.D., Brady, M., Smith, S.M., 2001. Temporal autocorrelation in univariate linear modeling of fmri data. Neuroimage 14, 1370–1386.

Xu, Y., Lindquist, M.A., 2015. Dynamic connectivity detection: an algorithm for determining functional connectivity change points in fmri data. Front. Neurosci. 9.

Yang, Z., Craddock, R.C., Margulies, D.S., Yan, C.G., Milham, M.P., 2014. Common intrinsic connectivity states among posteromedial cortex subdivisions: insights from analysis of temporal dynamics. Neuroimage 93, 124–137.

Zalesky, A., Fornito, A., Cocchi, L., Gollo, L.L., Breakspear, M., 2014. Time-resolved resting-state brain networks. Proc. Natl. Acad. Sci. Unit. States Am. 10341–10346, 201400181.

Zarahn, E., Aguirre, G.K., D'Esposito, M., 1997. Empirical analyses of bold fmri statistics. Neuroimage 5, 179–197.