NeuroImage xxx (2015) xxx–xxx

Contents lists available at ScienceDirect

NeuroImage

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

Comments and Controversies

Towards a statistical test for functional connectivity dynamics

Andrew Zalesky [a,b,*], Michael Breakspear [c,d]

[a] Melbourne Neuropsychiatry Centre and Melbourne Health, The University of Melbourne, Victoria 3010, Australia
[b] Melbourne School of Engineering, The University of Melbourne, Victoria 3010, Australia
[c] QIMR Berghofer Medical Research Institute, Brisbane, Queensland 4029, Australia
[d] The Royal Brisbane and Womens Hospital, Brisbane, Queensland 4029, Australia

ARTICLE INFO

Article history:
Received 26 November 2014
Accepted 18 March 2015
Available online xxxx

Keywords:
Functional connectivity
Dynamic connectivity
Sliding window
Non-stationarity
Time-resolved networks

ABSTRACT

Sliding-window correlation is an emerging method for mapping time-resolved, resting-state functional connectivity. To avoid mapping spurious connectivity fluctuations (false positives), Leonardi and Van De Ville recently recommended choosing a window length exceeding the longest wavelength composing the BOLD signal, usually assumed to be ~100 s. Here, we provide further statistical support for this rule of thumb. However, we demonstrate that non-stationary fluctuations in functional connectivity can in theory be detected with much shorter window lengths (e.g. 40 s), while maintaining nominal control of false positives. We find that statistical power is near-maximal for window lengths chosen according to Leonardi and Van De Ville's rule of thumb. Furthermore, we lay some foundations for a parametric test to identify non-stationary fluctuations in functional connectivity, also noting limitations of the sinusoidal model upon which our work, and the work of Leonardi and Van De Ville, is based. Most notably, our analytical results pertain to covariances, as does our statistical test, whereas functional connectivity is more commonly measured using correlations.

Resting-state functional brain connectivity is typically represented as a static network. This is an oversimplification, since any dynamic patterns of neural synchronization are reduced to time averages. In view of this, interest is growing in mapping time-resolved, or dynamic, functional brain connectivity using resting-state functional magnetic resonance imaging (fMRI) data (Fornito et al., 2013). Mapping brain networks as a function of time is becoming a widespread approach in the field (Calhoun et al., 2014; Handwerker et al., 2012; Kopell et al., 2014).

Sliding-window correlation is the most common analysis strategy for mapping dynamic functional connectivity (Hutchison et al., 2013), although temporal independent component analysis (Smith et al., 2012), model-based approaches (Lindquist et al., 2014), time-frequency coherence analysis (Chang and Glover, 2010) and change-point detection methods to identify stationary time segments (Cribben et al., 2012) have been used. The ubiquity of sliding-window correlation is probably owing to its simplicity: Pearson correlation is computed between regionally-averaged fMRI data falling within a fixed-length time window. Sliding the window in time and repeating yields a continuous series of snapshots tracking functional connectivity dynamics. Qualitatively, if the window length is too short, noise and intrinsic changes in the blood-oxygen-level dependent (BOLD) signal manifest as fast but spurious fluctuations in functional connectivity. These fluctuations represent false positives. However, if the window length is too long, genuine dynamic patterns of neural synchronization are smoothed within the length of a single window, and therefore might remain undetected. These undetected dynamics represent false negatives. The choice of window length can therefore be cast as a problem of maximizing statistical power.

In a recent commentary, Leonardi and Van De Ville (in press-a) analytically quantified this tradeoff in the choice of window length by modeling the BOLD signal with a sinusoid. They derived a closed-form expression approximating the sliding-window covariance between two identical sinusoids as a function of window position and length. The true covariance between identical sinusoids with unity root mean square power is unity. Their approximation accurately predicts larger deviations in the covariance from this true value as the window length is shortened.[1]

Based on their sinusoidal model, Leonardi and Van De Ville recommend using a window length exceeding $1/f_{min}$, where f_{min} is the slowest frequency component in the BOLD signal. Given that resting-state functional connectivity is thought to reflect slow (~0.01 Hz), spontaneous fluctuations in the BOLD signal (Biswal et al., 1995), their recommendation implies a window length of $1/0.01 = 100$ s. It follows from their

* Corresponding author. Level 3, Alan Gilbert Building, The University of Melbourne, Victoria 3010, Australia.
E-mail addresses: azalesky@unimelb.edu.au (A. Zalesky), michael.breakspear@qimrberghofer.edu.au (M. Breakspear).

[1] We noticed that their approximation yielded negative covariances for some window positions, but the covariance between two identical signals is necessarily nonnegative. This algebraic error was corrected in a subsequent erratum (Leonardi and Van De Ville, in press-b).

http://dx.doi.org/10.1016/j.neuroimage.2015.03.047
1053-8119/

recommendation that window lengths of 30 s, for example, necessitate 0.03 Hz high-pass filtering of the BOLD signal to suppress intrinsic low frequency fluctuations from manifesting as spurious functional connectivity dynamics. However, we suggest that this kind of high-pass filtering may suppress the very frequencies that give rise to genuine dynamic fluctuations in functional connectivity.

Here, we argue from a statistical viewpoint that the $1/f_{min}$ recommendation is a good rule of thumb for selecting window lengths. To this end, we lay some foundations for a statistical test to identify dynamic fluctuations in resting-state functional connectivity. The null hypothesis considered here is that fluctuations in connectivity from window to window merely reflect finite sample length effects due to the window length. Our statistical analysis of Leonardi and Van De Ville's sinusoidal model suggests that dynamic connectivity owing to spontaneous BOLD fluctuations at frequency f can be detected with window lengths substantially shorter than $1/f$, particularly when the signal-to-noise ratio (SNR) of the fMRI data is moderate and/or the functional connectivity dynamics are slow relative to spontaneous fluctuations in the BOLD signal. We suggest that window lengths as short as 40 s (Shirer et al., 2012) may be feasible, without suppression of BOLD frequencies beyond 0.01 Hz, as long as appropriate statistical testing is performed. Although we argue that window lengths substantially shorter than $1/f_{min}$ can be used to detect connectivity dynamics, we also suggest that statistical power is approximately maximized with $1/f_{min}$. We thus conclude that Leonardi and Van De Ville's recommendation is supported statistically, although there is some discretion for researchers to choose shorter windows.

Towards a statistical test

Leonardi and Van De Ville's $1/f_{min}$ recommendation corresponds to the smallest window length for which the sliding window covariance between two identical sinusoids is constant across all shifts in the window position; or in other words, the smallest window length for which there are absolutely no spurious fluctuations in covariance over time (see Fig. 1A). However, from a statistical viewpoint, some level of spurious fluctuations can be tolerated, particularly in the presence of system noise.

To demonstrate this, we revisit Leonardi and Van De Ville's sinusoidal model, but this time we add Gaussian white noise to each sinusoid. In particular, we consider the sliding-window covariance $c_{xy}^{\sigma}[n]$ for shifts in the window position n, where $x_i = a\cos(2\pi f i \mathrm{TR}) + \sigma\varepsilon_i, y_i = a\cos(2\pi f i\mathrm{TR} + \theta) + \sigma\varepsilon_i$ and ε is a standard normal random variable. Here TR denotes the sampling period, σ^2 is the noise variance and f is the "carrier" frequency; that is, the frequency component of the BOLD signal responsible for generating covariance fluctuations. We set $a = \sqrt{2}$ to give unity root mean square power. Given x_i and y_i are identical up to a phase shift θ, any fluctuations in $c_{xy}^{\sigma}[n]$ are solely owing to the carrier frequency, or noise, and can thus be considered spurious (i.e. false positives).

It can be shown that $c_{xy}^{\sigma}[n]$ is approximately normal with mean $\mu_n = c_{xy}^{\sigma=0}[n]$, as given by Eq. (5) in Leonardi and Van De Ville (in press-a), and with variance that can be bounded from above as follows,

$$\mathrm{var}\left(c_{xy}^{\sigma}[n]\right) \le \frac{\sigma^2}{w}\left(1 + \frac{1}{w^2}\right)\left(\sigma^2 + 2a^2\right), \tag{1}$$

where w is the number of time points comprising the window (see Fig. 1B). This result can be used to place confidence intervals on the extent to which $c_{xy}^{\sigma}[n]$ deviates, due to noise, from its mean value μ_n. We have $\mathbb{P}(|c_{xy}^{\sigma}[n] - \mu_n| \le C_{\alpha/2}) \le 1 - \alpha$, where $C_{\alpha/2}$ is the $(1 - \alpha/2)$-quantile for a zero-mean Gaussian with variance given by Eq. (1).

This immediately implies a parametric test for stationarity: We reject the null hypothesis of stationarity if we observe any fluctuations exceeding the interval,

$$I_\alpha := \left[\min_n \mu_n - C_{\alpha/2}, \max_n \mu_n + C_{\alpha/2}\right].$$

Taking the minimum/maximum of μ_n ensures control of familywise errors across the set of all shifts in window position. When calculating μ_n, the phase shift θ should be set such that $\cos(\theta)$, the covariance between x_i and y_i as $w \to \infty$, matches the observed covariance. The noise variance σ^2 should be set in accordance with the SNR of the fMRI data and depends on the region-of-interest size (Zalesky et al., 2010). And the sinusoidal amplitude a should be set to the amplitude of the BOLD signal.

We have thus far focussed on ensuring our test controls false positives at a nominal level. To demonstrate the sensitivity of our test, we apply it to the same sinusoidal model as above, but this time with one of the sinusoids modulated by a slow frequency component $f_0 < f$. The slow frequency component models a non-stationarity. Following Leonardi and Van De Ville (in press-a), we consider $\bar{y}_i = a\cos(2\pi f i\mathrm{TR})\cos(2\pi f_0 i\mathrm{TR})$, in which case the sliding window covariance $c_{x\bar{y}}[n]$ is non-stationary at the timescale of $1/f_0$. We can therefore determine the minimum window length necessary to reject the null hypothesis of stationarity for a given significance level and observed noise variance. Fig. 2 shows $c_{x\bar{y}}[n]$ as a function of window length for $f_0 = f, f/2, f/3$ and $f/4$, where $f = 0.01$ Hz. Black lines represent the covariance at different window positions. Thick red, green and blue lines are $\alpha = 0.05$ cutoffs for SNRs of 2, 4 and 6, respectively. We set $\theta = \pi/2$ when calculating μ_n since the covariance between x_i and $\bar{y}_i$ is zero as $w \to \infty$. The null hypothesis is rejected at a significance of α when there exists a shift in window position n such that $c_{x\bar{y}}[n] \notin I_\alpha$; that is, when at least one black line exceeds the appropriate colored line.

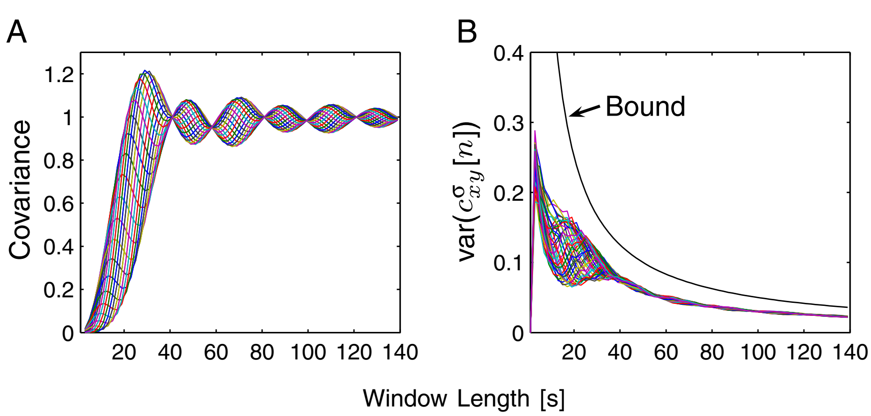

Fig. 1. A) Covariance between two identical sinusoids of frequency 0.025 Hz as a function of window length. Colored lines represent the covariance at different window positions. B) Upper bound on the variance of $c_{xy}^{\sigma}[n]$. Colored lines represent empirical variances at different window positions for $\sigma = 1, \theta = 0$ and $f = 0.025$ Hz. Black line represents our upper bound.

Please cite this article as: Zalesky, A., Breakspear, M., Towards a statistical test for functional connectivity dynamics, NeuroImage (2015), http://dx.doi.org/10.1016/j.neuroimage.2015.03.047