



Comments and Controversies

On spurious and real fluctuations of dynamic functional connectivity during rest



Nora Leonardi, Dimitri Van De Ville*

Institute of Bioengineering, Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland
 Department of Radiology and Medical Informatics, University of Geneva, Geneva, Switzerland

ARTICLE INFO

Article history:

Accepted 4 September 2014

Available online 16 September 2014

Keywords:

fMRI
 Dynamic functional connectivity
 Resting state
 Sliding-window correlation
 Non-stationarity

ABSTRACT

Functional brain networks reconfigure spontaneously during rest. Such network dynamics can be studied by dynamic functional connectivity (dynFC); i.e., sliding-window correlations between regional brain activity. Key parameters—such as window length and cut-off frequencies for filtering—are not yet systematically studied. In this letter we provide the fundamental theory from signal processing to address these parameter choices when estimating and interpreting dynFC. We guide the reader through several illustrative cases, both simple analytical models and experimental fMRI BOLD data. First, we show how spurious fluctuations in dynFC can arise due to the estimation method when the window length is shorter than the largest wavelength present in both signals, even for deterministic signals with a fixed relationship. Second, we study how real fluctuations of dynFC can be explained using a frequency-based view, which is particularly instructive for signals with multiple frequency components such as fMRI BOLD, demonstrating that fluctuations in sliding-window correlation emerge by interaction between frequency components similar to the phenomenon of beat frequencies. We conclude with practical guidelines for the choice and impact of the window length.

© 2014 Elsevier Inc. All rights reserved.

Introduction

Functional magnetic resonance imaging (fMRI) has become a key tool to probe the large-scale organization of the brain. Functional connectivity (FC), which is estimated by correlation of BOLD activity, identifies coherent brain activity in distributed and reproducible networks. FC has revealed reorganization of brain networks during cognitive tasks (Ekman et al., 2012; Lewis et al., 2009; Richiardi et al., 2011, 2013; Shirer et al., 2012), but also at rest (Allen et al., 2014; Chang and Glover, 2010; Hutchison et al., 2013b; Kang et al., 2011; Leonardi et al., 2013; Majeed et al., 2011; Smith et al., 2012). To study changes in FC over time sliding-window correlation analysis, where the correlation is estimated for brain activity during multiple, possibly overlapping temporal segments (typically 30–60 s), has been widely deployed (Allen et al., 2014; Chang and Glover, 2010; Hutchison et al., 2013a; Sakoglu et al., 2010). A caveat of analyzing dynamic FC (dynFC) by sliding-window correlation is that the small number of time points renders the estimates unreliable and might lead to spurious variability of dynFC (Hutchison et al., 2013a; Smith et al., 2012). However, there is no systematic account that perspicuously indicates the trade-off that is made by choosing the window length, and its implications for filtering of BOLD activity time series and dynFC itself.

We first break sliding-window correlation into several components to facilitate its study. Then, we present a simple yet instructive analytical model to study the emergence of spurious variability of dynFC in stationary signals. In particular, we investigate the influence of various parameters such as frequency, phase lag, and window length. Next, we introduce a small change to our analytical model to study how real variability of dynFC due to non-stationarity might arise. To provide the best possible insights for signals with many frequency components, we present a frequency-based view on dynFC. This provides an elegant explanation of how fluctuations of dynFC emerge through the interaction between different frequency components. Finally, we illustrate dynFC between two main regions of the default-mode network with experimental fMRI data.

Breaking down sliding-window correlations

We start by reformulating sliding-window correlation into simpler terms. In particular, we first look at sliding-window *covariance*, which for two time series x and y with sampling period TR is defined as follows at scan n :

$$\begin{aligned}
 c_{xy}[n] &= \text{cov}(x[n-\Delta, n+\Delta], y[n-\Delta, n+\Delta]) \\
 &= \frac{TR}{W} \sum_{i=n-\Delta}^{n+\Delta} (x_i - \bar{x}_n)(y_i - \bar{y}_n),
 \end{aligned} \tag{1}$$

* Corresponding author at: EPFL/IBI-STI, Station 17, 1015 Lausanne, Switzerland.
 E-mail address: dimitri.vandeville@epfl.ch (D. Van De Ville).

where $w = (2\Delta + 1)TR$ is the odd window length in seconds, i sums only over the scans inside the window, and

$$\bar{x}_n = \frac{TR}{w} \sum_{i=n-\Delta}^{n+\Delta} x_i$$

is the local average inside the window at position n . This calculation is then repeated for all values of n (“sliding” the window across time). After some elementary manipulations, we arrive at the following equality:

$$\begin{aligned} c_{xy}[n] &= \frac{TR}{w} \sum_{i=n-\Delta}^{n+\Delta} (x_i - \bar{x}_n)(y_i - \bar{y}_n) \\ &= \frac{TR}{w} \sum_{i=n-\Delta}^{n+\Delta} x_i(y_i - \bar{y}_n) - TR \frac{\bar{x}}{w} \sum_{i=n-\Delta}^{n+\Delta} (y_i - \bar{y}_n) \stackrel{(a)}{=} \frac{TR}{w} \sum_{i=n-\Delta}^{n+\Delta} x_i(y_i - \bar{y}_n) \quad (2) \\ &= \frac{TR}{w} \sum_{i=n-\Delta}^{n+\Delta} x_i y_i - \bar{y}_n \frac{TR}{w} \sum_{i=n-\Delta}^{n+\Delta} x_i = \underbrace{\frac{TR}{w} \sum_{i=n-\Delta}^{n+\Delta} x_i y_i}_I - \underbrace{\bar{y}_n \bar{x}_n}_{II}, \end{aligned}$$

where (a) simplifies as the second term equals zero. Thus, $c_{xy}[n]$ can be separated into two terms, which are the local average of the cross-product xy (I) minus the product of the local averages of x and y (II).

The sliding-window correlation is then obtained by normalizing at each window by the local variances:

$$\rho_{xy}[n] = \frac{c_{xy}[n]}{\sqrt{c_{xx}[n]c_{yy}[n]}}. \quad (3)$$

Having identified the components that constitute sliding-window correlation, we can now analyze and understand dynFC more easily.

Spurious fluctuations in dynFC

Effect of the window length

We want to understand how spurious fluctuations of dynFC might arise even for deterministic signals with a fixed relationship; i.e., we consider two pure sinusoidal signals that are phase-locked. Specifically, we take

$$x_i = \sqrt{2}\cos(2\pi f i TR), \quad y_i = \sqrt{2}\cos(2\pi f i TR + \theta), \quad (4)$$

where the factor $\sqrt{2}$ normalizes both signals for variance equal to one per time unit. This normalization makes the sliding-window covariance comparable to sliding-window correlation as a first approximation; i.e., we have the asymptotic equivalence $\lim_{w \rightarrow \infty} \rho_{xy}[n] = c_{xy}[n]$.

To investigate the influence of the key parameters frequency f , phase lag θ , and window length w , we derive the analytical form of $c_{xy}[n]$ for the signals of Eq. (4). First, we approximate \bar{y}_n by integration as follows:

$$\begin{aligned} \bar{y}_n &= \frac{TR}{w} \sum_{i=n-\Delta}^{n+\Delta} \sqrt{2}\cos(2\pi f i TR + \theta) \approx \frac{\sqrt{2}}{w} \int_{(n-\Delta)TR}^{(n+\Delta)TR} \cos(2\pi f t + \theta) dt \\ &= \frac{\sqrt{2}}{w} \left[\frac{1}{2\pi f} \sin(2\pi f t + \theta) \right]_{(n-\Delta)TR}^{(n+\Delta)TR} = \frac{\sqrt{2}}{w2\pi f} (\sin(2\pi f(n+\Delta)TR + \theta) \\ &\quad - \sin(2\pi f(n-\Delta)TR + \theta)) = \frac{\sqrt{2}}{w\pi f} \cos(2\pi f n TR + \theta) \sin(2\pi f \Delta TR). \end{aligned}$$

On similar grounds, we also find $\bar{x}_n = \frac{\sqrt{2}}{w\pi f} \cos(2\pi f n TR) \sin(2\pi f \Delta TR)$. Therefore, the second term $\bar{x}_n \bar{y}_n$ of Eq. (2) reverts to

$$\bar{x}_n \bar{y}_n = \frac{2}{w^2 \pi^2 f^2} \cos(2\pi f n TR) \cos(2\pi f n TR + \theta) \sin^2(2\pi f \Delta TR).$$

To estimate the first term of Eq. (2), we use the product-to-sum trigonometric identity

$$2\cos(2\pi f i TR) \cos(2\pi f i TR + \theta) = \cos(4\pi f i TR + \theta) + \cos(\theta),$$

which, after integration, leads to

$$\cos(\theta) + \frac{1}{w\pi f} \cos(2\pi f n TR + \theta) \sin(2\pi f \Delta TR).$$

By combining both terms, we retrieve the expression

$$\begin{aligned} c_{xy}[n] &= \cos(\theta) + \frac{1}{w\pi f} \cos(2\pi f n TR + \theta) \sin(2\pi f \Delta TR) \\ &\quad - \frac{2}{w^2 \pi^2 f^2} \cos(2\pi f n TR) \cos(2\pi f n TR + \theta) \sin^2(2\pi f \Delta TR). \quad (5) \end{aligned}$$

As a sanity check, we see that in the limit of stationary covariance (i.e., infinite window length), we have

$$\lim_{w \rightarrow \infty} c_{xy}[n] = \cos(\theta).$$

We now use this expression to efficiently trace $c_{xy}[n]$ as function of frequency f , phase lag θ , window length w , and window position n . In Fig. 1a, $c_{xy}[n]$ is plotted for $f = 0.025$ Hz and zero phase lag, as a function of window length w . The dashed lines are for different window positions n , and the thick line corresponds to the mean $\bar{c}_{xy} = E[c_{xy}[n]]$. We observe considerable fluctuations of $c_{xy}[n]$ for short window lengths, and crossings with the true value (i.e., 1) exactly for multiples of the window length because the term $\sin(2\pi f \Delta TR)$ in Eq. (5) vanishes for $2\Delta TR = 1/f$. Importantly, only when the window length is larger than the first crossing, which corresponds to the wavelength $1/f = 40$ s, fluctuations of $c_{xy}[n]$ diminish and converge to the true value of $\cos(\theta)$.

The same observations can be made from Figs. 1b and c, where we plot \bar{c}_{xy} for various frequencies, and the difference between maximal and minimal $c_{xy}[n]$ in Fig. 1d. Spurious fluctuations of $c_{xy}[n]$ occur when the window length is too short with respect to the underlying frequency component. We propose the following rule of thumb for minimal window length when observing underlying frequencies of f_{\min} or higher:

$$w \geq \frac{1}{f_{\min}}.$$

Therefore, high-pass filtering that removes frequency components below $1/w$ can be recommended; see also Smith et al. (2012) and Hutchison et al. (2013a) for similar recommendations. The cut-off frequency f_{\min} is indicated in Fig. 1. It should be noted that these plots only depend on the window length in seconds, not in TRs.

Sliding-window correlation $\rho_{xy}[n]$ (and its fluctuations) can be obtained by normalizing $c_{xy}[n]$ according to Eq. (3). In the ideal case with zero phase lag, sliding-window correlation clamps to 1; however, even a small phase lag is sufficient to introduce the same spurious fluctuations as we observed for sliding-window covariance. In Fig. 2, we plot sliding-window correlation and its extrema for phase lags of $\theta = \pi/16$ and $\theta = \pi/4$, respectively. The variability of sliding-window correlation is decreased compared to sliding-window covariance, but still the true correlation of $\cos(\theta)$ is recovered only for window lengths above w_{\min} , in accordance with the previous rule of thumb.




Fig. 1. Sliding-window covariance $c_{xy}[n]$ for two pure sinusoids without phase lag; i.e., the true covariance is always 1. (a) For frequency $f = 0.025$ Hz as a function of window length. Thick line indicates the mean \bar{c}_{xy} over all shifts n of the window position. Dashed lines indicate $c_{xy}[n]$ for different n . (b) For different frequencies $f = 0.01, 0.025, 0.05, 0.10$ Hz as a function of window length. Thick line indicates \bar{c}_{xy} , dashed lines indicate $\min_n c_{xy}[n]$ and $\max_n c_{xy}[n]$, and vertical lines indicate the minimal window lengths $w = 1/f_{\min}$. (c) Landscape of \bar{c}_{xy} as a function of frequency and window length. The dashed line indicates $w = 1/f_{\min}$. (d) Landscape of maximum — minimum of $c_{xy}[n]$ as a function of frequency and window length.

Effect of sampling and noise

Up to now, we did not take into account the effect of sampling or noise contributions, which can also drive fluctuations even if appropriate high-pass filtering has been applied. For the estimation of sliding-window correlation, we can use the 5 % confidence interval for the standard test of significant non-zero correlation:

$$\rho^* = \frac{t}{\sqrt{w/\text{TR} - 2 + t^2}},$$

where t follows Student's t -distribution with $w/\text{TR} - 2$ degrees of freedom. The confidence intervals for several TRs are shown in Fig. 2b and we list ρ^* for common choices of w and TR in the Appendix. Clearly, the number of scans included inside the window is an important factor for the estimator and thus depends on the TR for fixed window length (in seconds).

In sum, we have shown that spurious fluctuations of dynFC can arise by analyzing components with wavelengths larger than the window length, and that the limited number of data points inflates the influence of noise on the correlation estimates. The latter effect has also been referred to as “poor sampling of correlation” (Smith et al., 2012).

Real fluctuations in dynFC due to non-stationarity

Effect of modulatory component

We now first make a slight modification to the analytical model to introduce non-stationary relationship between both signals and thus true variability in dynFC. Specifically, we modify y by multiplying it with a low-frequency component ($f_0 \ll f$),

$$y_i = \sqrt{2} \cos(2\pi f i \text{TR}) \cos(2\pi f_0 i \text{TR}).$$

Then, dynFC between x and y will vary between $+1$ and -1 depending on the phase with respect to the low-frequency component. We note that the signal y is equivalent to

$$y_i = \frac{\sqrt{2}}{2} (\cos(2\pi(f + f_0)i\text{TR}) + \cos(2\pi(f - f_0)i\text{TR})),$$

which shows that low-frequency modulation is equivalent to introducing more frequency components. Repeating the analytical derivation of the previous section is possible, but the formulas rapidly become involved due to the interaction of many sinusoids. Instead, we revisit the sliding-window covariance of Eq. (2) and provide a Fourier interpretation that shows more easily how interactions between frequency components

can contribute to dynFC. Such a frequency-based view of dynFC is particularly helpful for real fMRI data with broad spectra.

The repeated calculation of local averages by shifting the window can be written as the convolution with a rectangular window h

$$c_{xy} = (xy) * h - (x * h)(y * h), \quad (6)$$

where $*$ denotes convolution and h is defined as

$$h[n] = \frac{TR}{w} \text{rect}\left(\frac{nTR}{w}\right) = \begin{cases} \frac{TR}{w}, & \text{for } |n| \leq \Delta, \\ 0, & \text{otherwise.} \end{cases}$$

For a tapered window, the local sums become weighted sums.

Using the convolution theorem,¹ the discrete Fourier transform (DFT) of Eq. (6) can be obtained as

$$C_{xy} = (X * Y)H - (XH) * (YH), \quad (7)$$

where capital letters denote the discrete Fourier transform (DFT) of the signals, in particular:

$$X[k] = F\{x[n]\} = \frac{1}{N} \sum_{n=0}^{N-1} x[n] \exp\left(-j2\pi \frac{kn}{N}\right), \quad k = 0, 1, \dots, N-1,$$

where N is the full length of the signal. For a signal sampled with period TR , the corresponding frequencies in Hz are given by $f_k = k/(NTR)$.

The DFT of the rectangular window h is well-known to be the Dirichlet kernel

$$H[k] = \frac{TR \sin(\pi kw/(NTR))}{w \sin(\pi k/(NTR))},$$

which can be seen as the discrete version of the common sinc-function. Fig. 3a shows the window function and Fig. 3b its spectrum, where the width of the main lobe is $1/w$ [Hz]. The convolution of x , y and xy can thus be seen as low-pass filtering operations.

The DFT of our signals x and y can be obtained as²

$$X[k] = \frac{\sqrt{2}}{2} (\delta_{f_{NTR}}[k] + \delta_{-f_{NTR}}[k]), \\ Y[k] = \frac{\sqrt{2}}{4} (\delta_{(f+f_0)NTR}[k] + \delta_{-(f+f_0)NTR}[k] + \delta_{(f-f_0)NTR}[k] + \delta_{-(f-f_0)NTR}[k]),$$

where δ is the Kronecker-delta function (i.e., $\delta_0[k] = 1$ for $k = 0$, and 0 otherwise). Exemplary signals x and y and their frequency spectra are shown in Figs. 3c and d. Since we deal with real-valued signals, all amplitude spectra are Hermitian symmetric and we only depict positive frequencies in the plots. The convolution $X * Y$ then redistributes the delta functions as

$$(X * Y)[k] = \frac{1}{2} (\delta_{f_0 NTR}[k] + \delta_{-f_0 NTR}[k]) \\ + \frac{1}{4} (\delta_{(2f-f_0)NTR}[k] + \delta_{(2f+f_0)NTR}[k] + \delta_{(-2f-f_0)NTR}[k] + \delta_{(-2f+f_0)NTR}[k]),$$

which is illustrated in Fig. 3f. Assuming that the window length has been chosen according to the rule of thumb (i.e., we have $w > 1/f_{\min}$ and thus also $f > f_{\min}$), frequency components at $\pm 2f \pm f_0$ are well suppressed by the filtering operation $(X * Y)H$, as well as those at $\pm f$

and $\pm f \pm f_0$ in XH and YH , respectively. Consequently, the second term of Eq. (7) vanishes and the first term simplifies, leading to

$$C_{xy}[k] = \frac{TR \sin(\pi f_0 w)}{w \sin(\pi f_0)} \left(\frac{\delta_{f_0 NTR}[k] + \delta_{-f_0 NTR}[k]}{2} \right),$$

where only $k = \pm f_0 NTR$ survives. The remaining frequency component at f_0 is also illustrated in Fig. 3d. C_{xy} can be identified as the DFT of

$$c_{xy}[n] = \frac{TR \sin(\pi f_0 w)}{w \sin(\pi f_0)} \cos(2\pi f_0 nTR),$$

shown in Fig. 3e. Because fluctuations in xy are low-pass filtered by the convolution with h with cut-off frequency $1/w = f_{\min}$, the slow modulation term—which in this case is a true fluctuation of dynFC—is recovered as long as $f_0 < f_{\min}$ and $f - f_0 \approx f > f_{\min}$. The influence of the window length on its low-pass filtering effect has previously been noted by Handwerker et al. (2012) and less variable dynFC with longer windows is a well documented empirical observation (e.g., Chang and Glover, 2010; Hutchison et al., 2013b; Leonardi et al., 2013). The spectral selectivity of the windowing operation can be improved by using tapering; e.g., Hamming filter (Handwerker et al., 2012), Gaussian filter (Allen et al., 2014), or other windows with smooth roll-off at the edges (Smith et al., 2012). In such case, the window length should be replaced by the “equivalent window length” that corresponds to the cut-off wavelength of the tapered window. It is essential to note that the frequency component f_0 emerges by “interaction” between both spectra and is not present as such in the original spectra. Mathematically, the beat frequency is recovered by multiplication in the time-domain, or, equivalently, convolution in the Fourier domain.

Example of experimental fMRI data

The frequency-based view is particularly instructive for experimental BOLD data because they have broad spectra that are not easily understood in terms of single frequency components. We illustrate dynFC between two key regions of the default-mode network: the posterior cingulate cortex (PCC) and left angular gyrus (AG). Changes of FC over time between these regions have been previously demonstrated (e.g. Chang and Glover, 2010).

Two regionally-averaged time series were extracted from a 10-minute long resting-state fMRI scan (data acquisition and preprocessing as described in Shirer et al. (2012)). We choose the window length to be $w = 50$ s, or 25 scans. Consequently, we high-pass filtered the time series with a cut-off frequency of $1/w = 0.02$ Hz; see Figs. 4a and b. Next, the point-wise multiplication of both time series and c_{xy} (~ low-pass filtered xy) are computed, shown in Fig. 4c. Clearly, these time series are similar during a majority of the scan, but their similarity is strongly diminished during two periods in time (around 90 and 270 s) and another small dip is visible at 420 s. In the Fourier domain, shown in Fig. 4d, this variability is apparent from new low-frequency components, notably a peak at 0.006 Hz, which appears to be approximately the wavelength of the interaction cycles. Because of the broad spectra of X and Y multiple frequency components contribute to these new low-frequency components.

In Figs. 4c and d, we also show sliding-window correlation ρ_{xy} in time and frequency. While the normalization reduces some fluctuations, its main characteristics are the same as for c_{xy} and are determined by $(X * Y)H$.

Conclusion

Taking advantage of an analytical model, we have derived some important properties of dynFC that explain the emergence of spurious fluctuations due to a mismatch between the choice of the window length and high-pass filtering of the original timecourses, as well as how potentially real fluctuations can arise due to modulatory components. We

¹ Convolution in the time domain corresponds to multiplication in the frequency domain and vice versa.

² Without any loss of generality, we have assumed here that $fNTR$ corresponds to an integer number.

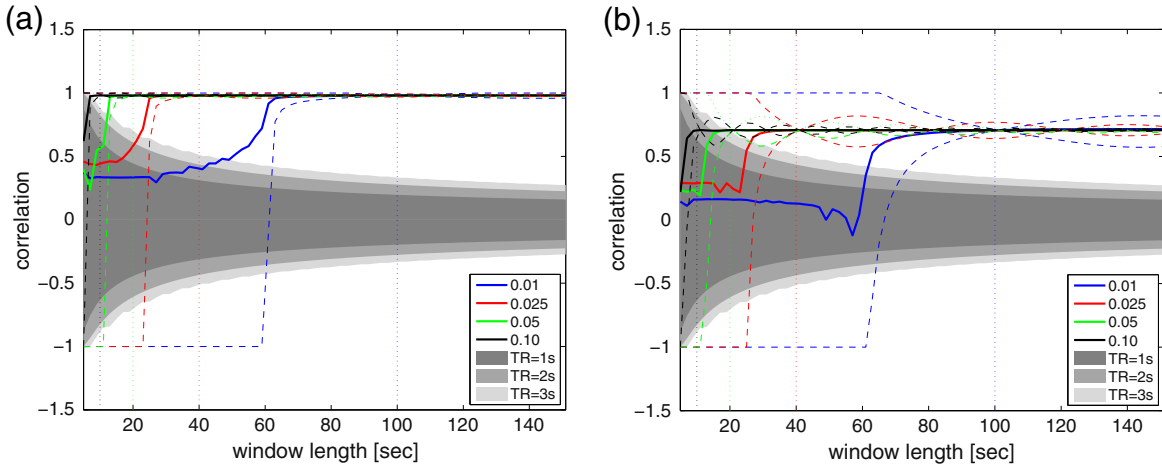


Fig. 2. (a) Sliding-window correlation $\rho_{xy}[n]$ for two pure sinusoids with phase lag $\pi/16$. (b) Sliding-window correlation $\rho_{xy}[n]$ for phase lag $\pi/4$. The gray shaded area indicates the 5% confidence interval of significant correlation for different TRs = 1, 2, 3 s.

conclude that the window length w , specified in seconds, is the key parameter that needs to be chosen carefully as it sets following trade-offs:

1. Spurious fluctuations of dynFC due to the nature of the estimation method are limited by high-pass filtering of the original time series with cut-off frequency $1/w$.
2. Remaining fluctuations of dynFC are low-pass filtered with cut-off frequency $1/w$.

3. Variability due to the influence of sampling and noise on the estimator increases with smaller window lengths or longer TR.

These different criteria should be balanced well. For example, windows as short as 20 s would require the removal of low-frequency components up to 0.05 Hz, which are typically of interest in resting-state studies, and the confidence interval for significant correlations would be very high ($\rho > 0.63$ for a TR of 2 s). Given these considerations, typical

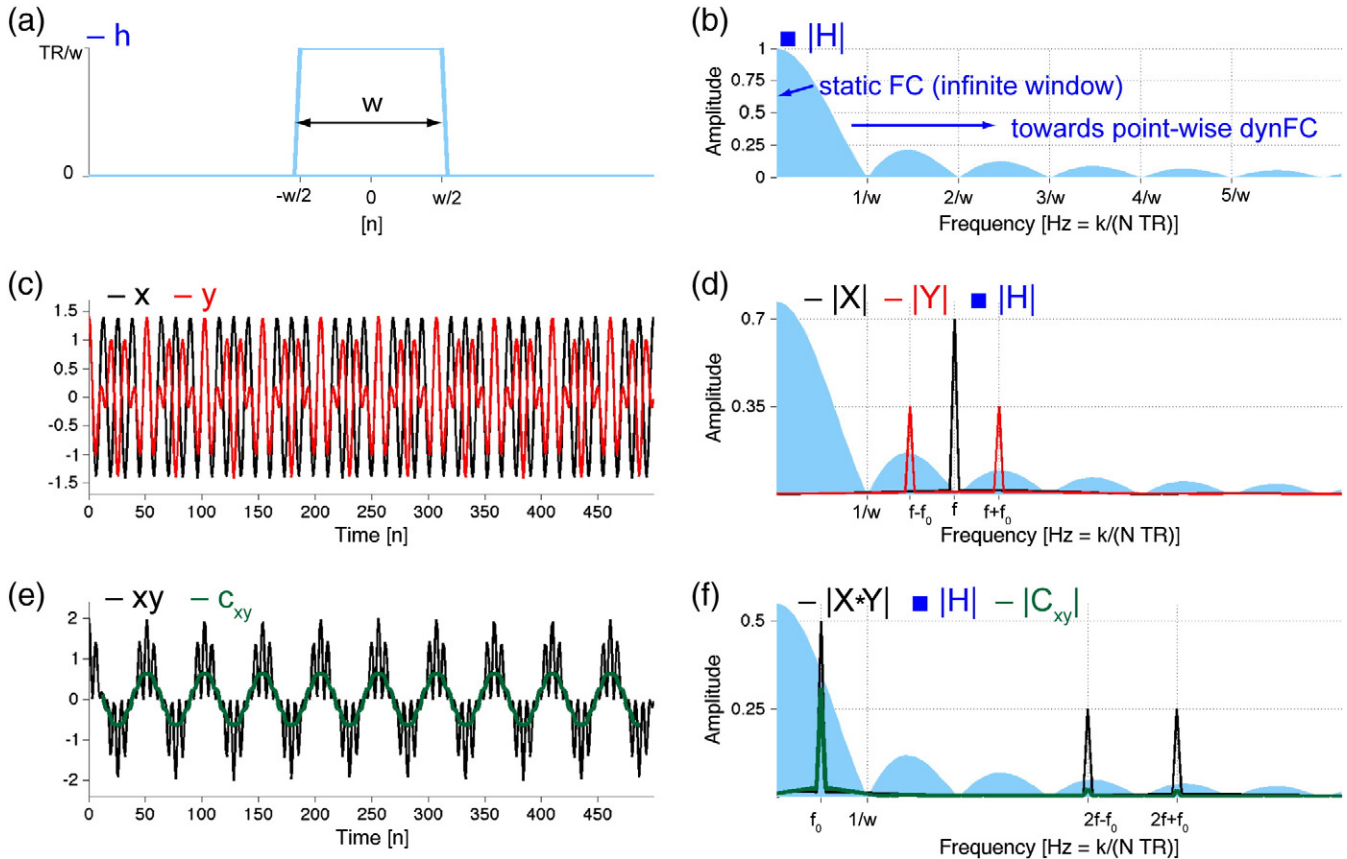


Fig. 3. Window h in time (a) and its frequency spectrum (b). The frequency selection of the filter is governed by $1/w$. Two cosines x and y , where y is modulated by a low-frequency cosine at f_0 , in time (c) and in frequency domain (d). The modulation of y corresponds to two slightly shifted peaks in Y at $f - f_0$ and $f + f_0$. The peaks at $f - f_0$, f and $f + f_0$ will be filtered out by x if they are beyond the main side lobe of the filter H . The point-wise product xy and sliding-window covariance c_{xy} in time (e) and frequency (f). Three new peaks appear in the convolution $X * Y$ from the interaction of the frequency components: f_0 , $2f - f_0$, and $2f + f_0$. Only the peak at f_0 is retained by the sliding-window covariance.

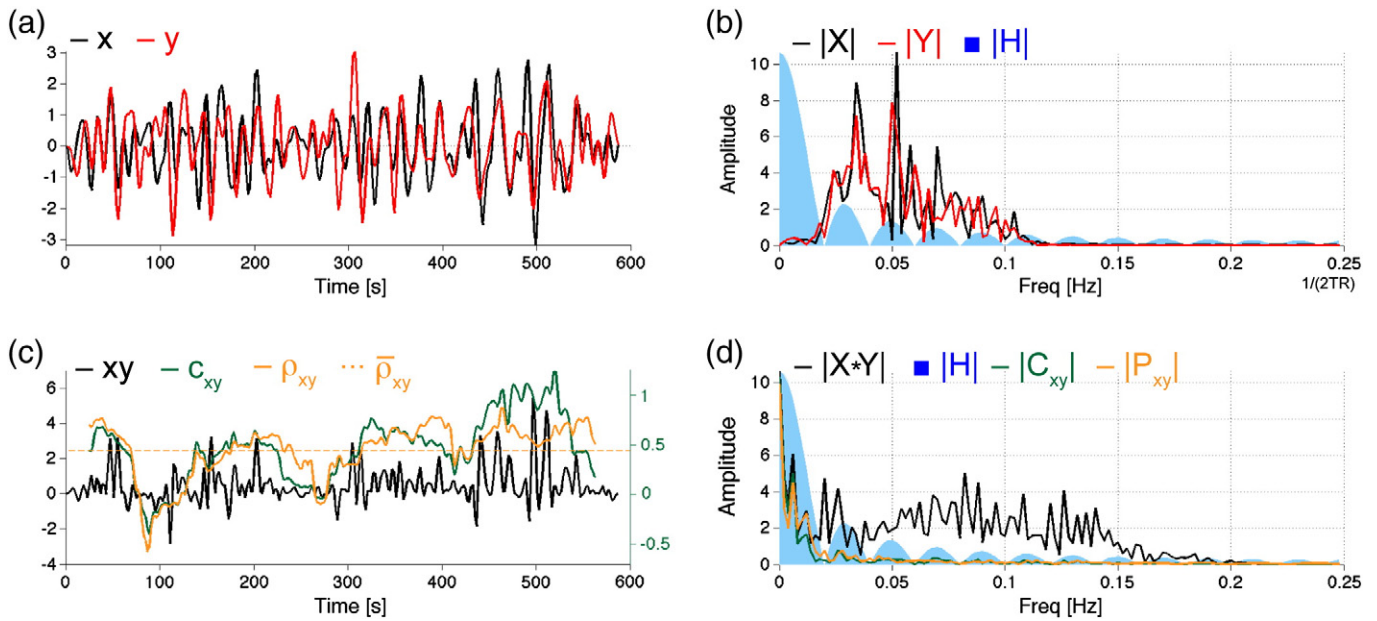


Fig. 4. (a–b) PCC and left AG time series and their frequency spectra. (c–d) Illustration of sliding-window covariance and correlation estimation in time and frequency. xy is the point-wise multiplication of the two time series, c_{xy} the sliding-window covariance and ρ_{xy} the sliding-window correlation. Note that the Fourier spectra of c_{xy} and ρ_{xy} are similar. $\bar{\rho}_{xy}$ is the average of dynFC across time, corresponding to a static FC estimate.

choices of window lengths (30–60 s) appear reasonable and these lengths have also been backed up by empirical studies that discriminated between cognitive states (Gonzalez-Castillo et al., 2013; Shirer et al., 2012). Therefore, when interpreting dynFC spectra, we suggest to focus on the frequency interval $[0 - 1/w]$ Hz because of the low-pass filtering effect of the window on dynFC, which means that the modulatory effects that can be observed are relatively slow; e.g., up to 0.16 Hz for a 60 s window length.

One promising future avenue to overcome the choice of a fixed window length is the use of the wavelet transform, which would allow to conveniently focus on particular frequencies (scales). In particular, the wavelet transform coherence (WTC) has been suggested as one alternative to estimate instantaneous “correlation coefficients” at different frequency bands, with a window length adjusted to the frequency content of the signal (Chang and Glover, 2010; Hutchison et al., 2013a; Torrence and Webster, 1999). In particular, the signal at each wavelet scale has been band-pass filtered with a high-pass cut-off according to the rule of the thumb, and, consequently, temporal variations in the scale-dependent correlations are limited to the same frequency. While this gives access to a rich amount of information, one remaining issue is how to combine scales into a concise and meaningful summarizing measure. Recent approaches in EEG analysis have employed a combination of windowed Fourier analysis and principal component analysis (PCA) for that purpose (Mehrkanoon et al., 2014).

Fluctuations of dynFC can be driven by noise, and, therefore, should be tested for significance using parametric testing as suggested before, or, alternatively, surrogate data using autoregressive models (Chang and Glover, 2010) or phase randomization (Handwerker et al., 2012; Prichard and Theiler, 1994) that preserve temporal correlation properties. Finally, it is important to note that while fluctuations of dynFC might be driven by true non-stationarities and interactions of the time series, the origin of these signals could be both neurological and non-neurological. Ongoing and future research should further validate to what extent these origins can be disentangled (Chang et al., 2013); e.g., using concurrent measurements such as electroencephalography (EEG) and non-neurophysiological signals.

Acknowledgments

The Matlab code to generate the plots of this letter is available from the authors’ website at <http://miplab.epfl.ch/software/>.

The authors thank Michael Greicius and Will Shirer for making the data available and Steve Smith for insightful discussions during the Third International Workshop on Pattern Recognition in Neuroimaging (PRNI2013). This work was supported in part by the Swiss National Science Foundation (grant PP00P2-146318), and in part by the Center for Biomedical Imaging (CIBM).

Conflict of interest

The authors declare no conflict of interest.

Appendix

Table 1
 ρ^* for significant non-zero correlation (5 % confidence level). Different window lengths and TRs (w/ TR was rounded to the nearest integer).

Window length	TR		
	1 s	2 s	3 s
20 s	0.44	0.63	0.75
30 s	0.36	0.51	0.63
40 s	0.31	0.44	0.55
50 s	0.28	0.40	0.48
60 s	0.25	0.36	0.44
120 s	0.18	0.25	0.31

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. *Cereb. Cortex* 24, 663–676.
 Chang, C., Glover, G., 2010. Time–frequency dynamics of resting-state brain connectivity measured with fMRI. *NeuroImage* 50, 81–98.
 Chang, C., Liu, Z., Chen, M.C., Liu, X., Duyn, J.H., 2013. EEG correlates of time-varying BOLD functional connectivity. *Neuroimage* 72, 227–236.

- Ekman, M., Derrfuss, J., Tittgemeyer, M., Fiebach, C.J., 2012. Predicting errors from reconfiguration patterns in human brain networks. *Proc. Natl. Acad. Sci. U. S. A.* 109, 16714–16719.
- Gonzalez-Castillo, J., Hoy, C., Handwerker, D.A., Robinson, M., Band, 2013. Detection of consistent cognitive processing at the single subject level using whole-brain functional connectivity. Society for Neuroscience, San Diego.
- Handwerker, D.A., Roopchansingh, V., Gonzalez-Castillo, J., Bandettini, P.A., 2012. Periodic changes in fMRI connectivity. *NeuroImage* 63, 1712–1719.
- 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., Handwerker, D.A., Keilholz, S., Kiviniemi, V., Leopold, D.A., de Pasquale, F., Sporns, O., Walter, M., Chang, C., 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.
- Kang, J., Wang, L., Yan, C., Wang, J., Liang, X., He, Y., 2011. Characterizing dynamic functional connectivity in the resting brain using variable parameter regression and Kalman filtering approaches. *NeuroImage* 56, 1222–1234.
- Leonardi, N., Richiardi, J., Gschwind, M., Simioni, S., Annoni, J.M., Schluep, M., Vuilleumier, P., Van De Ville, D., 2013. Principal components of functional connectivity: a new approach to study dynamic brain connectivity during rest. *Neuroimage* 83, 937–950.
- Lewis, C.M., Baldassarre, A., Committeri, G., Romani, G.L., Corbetta, M., 2009. Learning sculpts the spontaneous activity of the resting human brain. *Proc. Natl. Acad. Sci. U. S. A.* 106, 17558–17563.
- Majeed, W., Magnuson, M., Hasenkamp, W., Schwarb, H., Schumacher, E.H., Barsalou, L., Keilholz, S.D., 2011. Spatiotemporal dynamics of low frequency BOLD fluctuations in rats and humans. *NeuroImage* 54, 1140–1150.
- Mehrkanoun, S., Breakspear, M., Boonstra, T.W., 2014. Low-dimensional dynamics of resting-state cortical activity. *Brain Topogr.* 27, 338–352.
- Prichard, D., Theiler, J., 1994. Generating surrogate data for time series with several simultaneously measured variables. *Phys. Rev. Lett.* 73, 951–954.
- Richiardi, J., Eryilmaz, H., Schwartz, S., Vuilleumier, P., Van De Ville, D., 2011. Decoding brain states from fMRI connectivity graphs. *NeuroImage* 56, 616–626.
- Richiardi, J., Achard, S., Bunke, H., Van De Ville, D., 2013. Machine learning with brain graphs. *IEEE Signal Process. Mag.* 30, 58–70.
- Sakoglu, U., Pearlson, G.D., Kiehl, K.A., Wang, Y.M., Michael, A.M., Calhoun, V.D., 2010. A method for evaluating dynamic functional network connectivity and task-modulation: application to schizophrenia. *Magn. Reson. Mater. Phys.* 351–366.
- Shirer, W.R., Ryali, S., Rykhlevskaia, E., Menon, V., Greicius, M.D., 2012. Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cereb. Cortex* 22, 158–165.
- Smith, S.M., Miller, K.L., Moeller, S., Xu, J., Auerbach, E.J., Woolrich, M.W., Beckmann, C.F., Jenkinson, M., Andersson, J., Glasser, M.F., Van Essen, D.C., Feinberg, D.A., Yacoub, E.S., Ugurbil, K., 2012. Temporally-independent functional modes of spontaneous brain activity. *Proc. Natl. Acad. Sci. U. S. A.* 109, 3131–3136.
- Torrence, C., Webster, P.J., 1999. Interdecadal changes in the ENSO–monsoon system. *J. Clim.* 12, 2679–2690.