NeuroImage 122 (2015) 399–407

Contents lists available at ScienceDirect

NeuroImage

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

Estimation of dynamic functional connectivity using Multiplication of Temporal Derivatives

James M. Shine [a,b,*], Oluwasanmi Koyejo [b], Peter T. Bell [a], Krzysztof J. Gorgolewski [b], Moran Gilat [a], Russell A. Poldrack [b]

[a] Parkinson's Disease Research Clinic, Brain and Mind Research Institute, The University of Sydney, NSW, Australia
[b] Department of Psychology, Stanford University, Stanford, CA, USA

ARTICLE INFO

Article history:
Received 15 February 2015
Accepted 23 July 2015
Available online 29 July 2015

ABSTRACT

Functional connectivity provides an informative and powerful framework for exploring brain organization. Despite this, few statistical methods are available for the accurate estimation of dynamic changes in functional network architecture. To date, the majority of existing statistical techniques have assumed that connectivity structure is stationary, which is in direct contrast to emerging data that suggests that the strength of connectivity between regions is variable over time. Therefore, the development of statistical methods that enable exploration of dynamic changes in functional connectivity is currently of great importance to the neuroscience community. In this paper, we introduce the 'Multiplication of Temporal Derivatives' (MTD) and then demonstrate the utility of this metric to: (i) detect dynamic changes in connectivity using data from a novel state-switching simulation; (ii) accurately estimate graph structure in a previously-described 'ground-truth' simulated dataset; and (iii) identify task-driven alterations in functional connectivity. We show that the MTD is more sensitive than existing sliding-window methods in detecting dynamic alterations in connectivity structure across a range of correlation strengths and window lengths in simulated data. In addition to the temporal precision offered by MTD, we demonstrate that the metric is also able to accurately estimate stationary network structure in both simulated and real task-based data, suggesting that the method may be used to identify dynamic changes in network structure as they evolve through time.

Introduction

A recent paradigm shift in functional brain imaging has led to dramatic advances in our understanding of the functional organization of the human brain. These advances place *connectivity* at the core of brain organization and emphasize the fundamental importance of inter-regional communication for global brain function (Bullmore and Sporns, 2009). Given that functional connectivity provides an informative methodology for exploring brain organization, the development and optimization of analytic methods to describe statistical relationships between regions of the brain is currently of great importance to the neuroscientific community (Hutchison et al., 2013).

To date, the vast majority of statistical techniques used to calculate functional connectivity have assumed that the connectivity structure is stationary through time. In direct contrast to this assumption, emerging methodologies have demonstrated that the strength of connectivity between regions is highly dynamic, enabling the emergence and dissolution of transient functional repertoires through time (see Hutchison et al., 2013 for review). Previous attempts to describe dynamic patterns of functional connectivity have utilized two main approaches: (i) estimation of pairwise variations in inter-regional covariance (Allen et al., 2014; Chang and Glover, 2010; Handwerker et al., 2012; Hutchison et al., 2012; Jones et al., 2012; Kiviniemi et al., 2011; Sakoğlu et al., 2010; Zalesky et al., 2014), and (ii) the identification of changing patterns of covariance at a multivariate level (Calhoun et al., 2014; Cribben et al., 2012; Damaraju et al., 2014; Leonardi et al., 2013, 2014; Lindquist et al., 2007; Liu and Duyn, 2013; Magnuson et al., 2010; Majeed et al., 2011; Robinson et al., 2014; Smith et al., 2012; Tagliazucchi et al., 2012). Although these techniques have provided a number of insights into the temporal dynamics of the brain, the methods have generally been limited by a lack of appropriate temporal resolution.

Here, we introduce and validate a novel metric, Multiplication of Temporal Derivatives (MTD), that can be used to interrogate dynamic functional connectivity in fMRI time series data. In the following, we demonstrate the results of a three-part experiment — demonstrating the ability of this novel metric to calculate dynamic and stationary functional connectivity structure in both real and simulated data. In Experiment 1, we utilize a novel simulation framework to assess the sensitivity and specificity of MTD to detect switches in covariance

* Corresponding author at: Department of Psychology, Stanford University, Stanford, CA, USA.
E-mail address: mac.shine@sydney.edu.au (J.M. Shine).

http://dx.doi.org/10.1016/j.neuroimage.2015.07.064
1053-8119/

structure within a simulated time series dataset. In Experiment 2, we demonstrate that the metric also provides a robust method for describing stationary patterns of functional connectivity by comparing its performance against common methodologies using a 'gold-standard' simulated dataset (Smith et al., 2011). In Experiment 3, we apply the metric to a high-quality task-based dataset from the Human Connectome Project (Barch et al., 2013) in order to demonstrate that the method can identify task-driven changes in network reconfiguration in real fMRI data. Together, the results of these highlight the potential of MTD in estimating dynamic and stationary functional connectivity structure.

Methods

Multiplication of Temporal Derivatives

To estimate connectivity at each time point of functional neuroimaging data using the MTD metric, we first calculate the temporal derivative (dt) of each time series (ts) of length t by performing a first-order differencing (i.e. subtracting the BOLD intensity at time point $t - 1$ from the intensity at time point t; see Eq. 1).

$$dt_{it} = ts_{it} - ts_{it-1} \tag{1}$$

Equation 1: the temporal derivative of node i at time t is calculated as the temporal difference between the intensity of the signal at time t and $t - 1$ at node i.

For each of n nodes with t unique time points, we calculate a $t - 1$ vector of temporal derivative values, and then normalize each data point by dividing each dt by the standard deviation (σ) of the dt, calculated over the entire time course. To calculate the MTD score at each time point, the dt for each pair of nodes ij are then multiplied, creating a $t - 1 \times n \times n$ matrix, in which the value in each cell reflects the degree of functional coupling between the ith and jth nodes of the network (see Eq. 2). Positive MTD scores thus reflect 'coupling' in the same direction of signal change across nodes (that is, signal either both increasing or both decreasing together), whereas negative scores reflect 'anti-coupling' (that is, signal in one node increasing while the other is decreasing). The mean MTD value can then be calculated over the course of an entire experiment to estimate the stationary functional coupling between two nodes, i and j (see http://github.com/macshine/coupling/ for code).

$$MTD_{ijt} = \frac{(dt_{it} \times dt_{jt})}{(\sigma_i \times \sigma_j)} \tag{2}$$

Equation 2: Multiplication of Temporal Derivatives metric at time t between nodes i and j.

Any method that estimates connectivity on a single data point will, by definition, be more susceptible to fluctuations in high frequency noise. For this very reason, there is currently debate in the literature regarding the ideal window length to use for dynamic connectivity analyses — long window lengths are relatively insensitive to rapid alterations in connectivity, whereas short windows are more susceptible to fluctuations in high frequency noise (Leonardi and Van De Ville, 2015; Zalesky and Breakspear, 2015). In addition, methodological considerations also play a role, as certain techniques (such as sliding window Pearson correlation analysis) require relatively large temporal windows to improve estimates of covariance (Hutchison et al., 2013). To assess the effect of window size in our experiment, we calculated a simple moving average (SMA) of the MTD metric, such that data are averaged surrounding a point in time within a window, w (Eq. 3), a technique commonly employed in the literature when using sliding window analyses (Hutchison et al., 2013). The window length can then be flexibly manipulated depending on the nature of the hypothesis being interrogated (i.e. short window lengths for change points analysis and long window lengths for robust estimation of positive covariance structure).

$$SMA_{ijt} = \frac{1}{2w+1}\sum_{t-w}^{t+w} MTD \tag{3}$$

Equation 3: simple moving average (SMA) of the Multiplication of Temporal Derivatives (MTD) score for window length w over time t.

Experiment 1 — dynamic connectivity simulation

The primary goal of Experiment 1 was to compare the performance of MTD against the primary method used for calculating dynamic functional connectivity — sliding window Pearson's correlation coefficients (SWPC) — using a novel simulated dataset. To do so, we compared the sensitivity and specificity of MTD against the SWPC method across a range of correlation thresholds and window-sizes. Importantly, the use of a simulated dataset allowed us to systematically evaluate the sensitivity and specificity of each technique, as we knew the 'ground truth' state of the functional connectivity in the simulated data.

Experiment 1a — sustained alterations in connectivity structure

To assess the capacity of MTD to estimate dynamic connectivity, we performed a series of analyses using a novel simulation that used a novel State Switching model (SSM) to generate realistic BOLD time course data in which the connectivity 'state' between two regions could be dynamically manipulated between two experimental states: one involving no positive correlation between nodes (State 1), and another with a positive correlations between nodes (State 2). At each time point, the covariance structure between a pair of nodes was determined by a latent state variable, which evolved following a discrete dynamical model, with dynamics determined by a pre-defined state transition switch (Baum and Petrie, 1966). At each time point, an observation vector was generated from a multivariate Gaussian distribution with covariance associated with the hidden state (Eq. 4).

$$X_t \sim P(X|Z_t, C) = N(X; 0, C_{Z_t}) \tag{4}$$

Equation 4: State Switching Model for simulation of time series data.

Where Z indexes states, X represents observations, and C_Z is the covariance associated with state Z.

As the State identity of each matrix was known from time point to time point, we were able to evaluate the sensitivity and specificity of the MTD in capturing dynamic switches in covariance structure. To do so, we simulated the data using the state-switching model (SSM) 1000 times each across a range of correlation coefficients (Pearson's r = 0.1, 0.2, 0.3, 0.4 and 0.5) for two dynamic 'switch' scenarios. In the first scenario, we created a single sustained switch from State 1 to State 2. We then calculated an SMA of the MTD (Eq. 3) for the node-pair over a range of window lengths (with the smallest window including one time point either side of target point [i.e. a window length of 3 TRs; $w = 1$] and the largest encompassing 10 time points either side of target point [i.e. a window length of 21 TRs; $w = 10$]). As a direct comparison, we calculated SWPC, a widely utilized method for exploring dynamic functional connectivity in the literature (Hutchison et al., 2013), with the same range of variable window lengths (i.e. $w = 1$–10), shifted sequentially by 1 TR to cover the entire time course. Importantly, the two competing methods were aligned such that they had access to identical data points for each simulation.

To directly compare the performance of the MTD and SWPC over the range of window lengths, we used a 'sliding window' approach to assess the data for significant differences in mean connectivity between the two simulated nodes at each point in time. To achieve this, we performed an independent-samples t-test comparing the 10 time points prior to with the 10 time points following a given data point and then 'slid' this window across the time series. Therefore, at each point in

the data, we obtained a t-statistic that reflected the likelihood that there was a large deviation in the mean connectivity score (calculated using either the MTD or SWPC) between the two windows surrounding each point. By performing a negative $\log_{10}$ transform of the p-value for each of the 1000 simulations, we were able to estimate the 'confidence' of a switch occurring at a given data point. That is, high values in the confidence measure, corresponding to small p-values, suggested increased confidence that a switch occurred at that point in the data. After standardizing this confidence measure within each window to allow comparison across window lengths, we were able to create a measure that reflected the confidence that a switch occurred at a given data point for each method that was directly comparable across the range of windows and correlation strengths.

To further characterize the performance of the two metrics, we calculated the sensitivity and specificity of each measure to accurately detect the switch point: sensitivity was defined as the proportion of simulations in which there was a significant confidence value (confidence score >3.0, equivalent to $p < 0.001$) within 5 time points of a switching event; and specificity was defined as the inverse of the proportion of simulations in which there was a high confidence value present in a window that did not involve any switch in correlation structure (i.e. comparing two windows with matching states). This latter value reflects the extent to which the MTD or SWPC at a given window length mistakenly identified data in State 1 as associated with positive covariance structure (that is, State 2). Importantly, our objective was to create a metric with high sensitivity to switch point detection, such that the techniques were likely to be associated with a trade-off in specificity, the latter of which has recently been shown to be maximized using multivariate volatility methods (Lindquist et al., 2014).

Experiment 1b — transient alterations in connectivity

In addition to detecting changes in connectivity structure (Experiment 1a), it is also important for a dynamic connectivity metric to identify transient changes in connectivity (Experiment 1b), such as those involved in an event-related fMRI design, in which behavioral variables are presented rapidly and classically modeled as events with zero latency. To compare the effectiveness of the MTD and SWPC to estimate these dynamic fluctuations, we created a transient switch from State 1 to State 2 lasting for a duration of six time points, the number of which was estimated to most realistically model the number of time points in which a BOLD response from a single neural event would last in a typical fMRI experiment with a TR of 1 s (i.e. 6 time points). We simulated data 1000 times across the same range of correlation coefficients as in the previous experiment ($r = 0.1–0.5$) and subsequently calculated the SMA of the MTD and the SWPC scores for the same range of window sizes ($w = 1–10$). The confidence, sensitivity and specificity of each metric were calculated as per Section Experiment 1 – dynamic connectivity simulation, however, as the covariance structure shifted after six time points, we only used six data points (as opposed to ten in the previous analysis) per window for the sliding t-tests to increase temporal precision.

Experiment 1c — effect of amplitude change on estimates of connectivity

One difficulty of interpreting functional connectivity results in the context of external task demands is disentangling spurious increases in connectivity that occur due to co-activation of neural regions by task demands from underlying increases in task-related functional connectivity. In order to mimic this scenario in a simulation, we added 1.0% signal intensity (a conservative estimate of the BOLD percent signal change predicted from a standard block-design experiment) to two time series that were otherwise not correlated with one another (i.e. they were in State 1). We then calculated the MTD and SWPC values for this data as per methods described in Section Experiment 1 – dynamic connectivity simulation. As there was no underlying positive covariance between the two nodes, we reasoned that a method sensitive to actual changes in covariance but insensitive to alterations in evoked responses should not detect a positive switch in this scenario.

Experiment 1d — effect of noise addition

There is an extensive literature detailing the negative effect of spurious noise on the estimation of stationary functional connectivity (e.g. see Power et al., 2012). However, there is less clarity regarding the effect of noise on estimates of dynamic connectivity. To estimate the possible adverse effects of spurious noise to both the MTD and SWPC, we created two separate simulations. In the first simulation, we simulated a range of low and high frequency noise signals by adding a randomly phase-shifted sinusoid signal to the original data time series ($r = 0.3$) at a range of frequencies (0.001, 0.02, 0.04, 0.06, 0.08, 0.10, 0.20 and 0.40 Hz) and then calculated the MTD and SWPC estimates of covariance structure for $w = 3$.

Experiment 1e — effect of head motion

To estimate the global effects of motion, we added a series of global signal spikes (1.0% signal intensity) to each time series randomly to a proportion of time points (at the same point in each time series), varying the proportion from 1/n (i.e. one time point) up to 25% of the time series. We then calculated the MTD and SWPC estimates of covariance structure for $w = 3$. In both instances, we then compared the resultant sensitivity and specificity of each trial to the value estimated from data without noise added.

Experiment 1f — effect of filtering

In conventional sliding-window FC, it is common practice to apply a high-pass filter to timeseries that has a cut-off about the reciprocal of the window length (e.g., 100 s requires an 0.01 Hz HPF cut-off) to avoid spurious fluctuations due to aliasing (for instance, see Leonardi and van de Ville, 2015). To determine the extent to which the choice of high-pass filter affected estimates of covariance, we ran a series of simulations in which we altered the extent of high-pass filtering on the original time series, using a range of filters (0.005, 0.006, 0.007, 0.008, 0.01 and 0.013 Hz) and then calculated the MTD and SWPC on the resultant time series ($r = 0.3$ and $w = 3$). In addition, we also ran a further simulation in which the data were band-pass filtered in a similar to the window utilized for many resting state analyses (between 0.001 and 0.01 Hz). In a similar fashion to experiment 1d, the sensitivity and specificity of each trial was then normalized to the value estimated from data without filtering.

Experiment 2 — stationary network detection

In Experiment 2, our aim was to demonstrate that the MTD is able to accurately estimate stationary connectivity structure in a well-characterized fMRI dataset, demonstrating the capacity of the MTD to capture meaningful patterns of dynamic changes in connectivity from fMRI time series data. We achieved this aim by directly comparing the performance of MTD against a variety of existing methods using a previously published gold-standard stimulated data set (Smith et al., 2011), which was obtained from FMRIB (http://www.fmrib.ox.ac.uk/analysis/netsim). Briefly, this dataset consists of 28 simulations of BOLD data in 50 realizations (TR = 1.5–3 s; 200–1000 individual time points; 5–50 nodes; and differing levels of noise and hemodynamic response function variability). Each simulated dataset was created using an fMRI forward model based on dynamic causal modeling (DCM; (Friston et al., 2011)), combined with a nonlinear balloon model (Buxton et al., 1998) to simulate vascular dynamics (see (Smith et al., 2011) for details of the simulation). The authors simulated realistic BOLD data with known network structure (i.e. 'ground truth'), enabling

the evaluation of a range of popular connectivity models to estimate the 'ground truth' network structure.

To test for the strength of functional connectivity, the authors calculated 'c-sensitivity', which was defined as the fraction of true positive connectivity estimates (that is, the estimated connection strength between nodes that were connected in the true network) with a connectivity strength greater than two standard deviations above the mean of the estimated scores for edges that are not present in the true network (that is, true negative connections). Smith and colleagues found that strong methods performed at greater than 80% c-sensitivity, however there was a great deal of variability amongst the different simulations. In this study, we applied the same evaluation criteria as (Smith et al., 2011). We calculated the mean c-sensitivity for the entire group using the MTD metric for each of the 28 simulations from the Smith et al. (2011) dataset. For clarity, results are presented without the application of SMA to MTD data, however c-sensitivity scores for the MTD metric did improve substantially with the application of a SMA with moderate window size ($w = 3$–4; see Fig. S1).

Experiment 3 — task-based functional connectivity

In the final experiment, we provide further validation of the potential utility of MTD by demonstrating that the metric is able to provide insights into the functional organization of the brain using real fMRI time series data. To assess task-based functional connectivity, preprocessed data from 40 unrelated subjects collected while performing a visually-based working memory task were downloaded from the Human Connectome Project server (Glasser et al., 2013). Briefly, the task consisted of interleaved blocks of a high and low load working memory task (2-back and 0-back, respectively), paired with an object recognition task (four separate conditions: places, faces, body parts and tools). In this experiment, we examined two contrasts: firstly, working memory load was interrogated by comparing patterns of activity and connectivity during all four '2-back' blocks and all four '0-back' blocks; and secondly, the differential pattern of 'face' and 'place' object recognition was interrogated by comparing patterns during both high and low load 'face' blocks with patterns during both 'place' blocks.

Preprocessing of timeseries data involved distortion correction and linear head movement correction, following which all data was projected onto two 32,492 vertex cortical sheets (one per hemisphere). The mean time series was then extracted from each of 333 surface parcels (166 in the left hemisphere and 167 in the right hemisphere), which were defined according to patterns of homogenous vertex-wise stationary resting state connectivity (Gordon et al., 2014). To estimate patterns of connectivity in a targeted fashion, the MTD metric was calculated for a reduced set of 18 parcels (Fig. S3), which were identified via reverse inference of the terms: 'working memory' (8 parcels), 'faces' (5 parcels) and 'places' (5 parcels) using the NeuroSynth database (Yarkoni et al., 2011). A time series of the MTD score ($w = 3$, the window length with the highest sensitivity, specificity and confidence across a range of correlation strengths from our previous simulations; see Sections 3.1.1 and 3.1.2) was calculated for each of the 153 unique connections and then entered into a mixed-effects general linear model analysis (Mumford and Poldrack, 2007) with task-onset regressors from the working memory task using customized code implemented in Python. Results from the mixed-effect analyses were visualized on the cortical surface by performing a negative $\log_{10}$ transformation of the corresponding p-value for each parcel-wise connection using the Connectome Workbench toolbox (see Fig. S3).

Results

Experiment 1a — sustained alterations in connectivity

The MTD outperformed the conventional SWPC approach when detecting dynamic switches in connectivity States in simulated timeseries data (Fig. 1). When detecting a sustained switch in connectivity, the MTD was associated with a higher degree of confidence in the switch point across a range of correlation strengths (Fig. 1a), particularly when using finer temporal window ($w = 2$–4; Table 1). Across all simulations, the SWPC showed poorer discrimination of switch points,

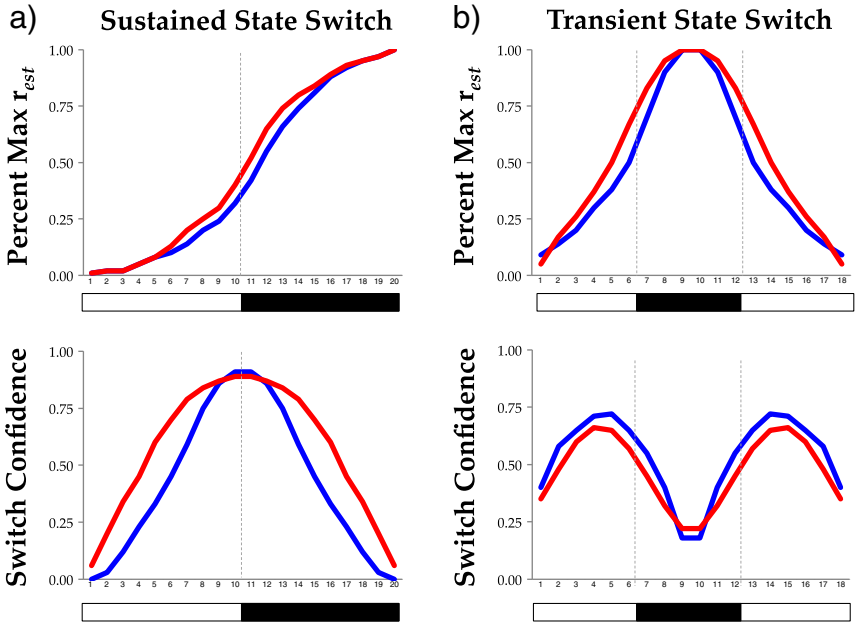

Fig. 1. Sustained alterations in covariance structure. The Multiplication of Temporal Derivatives (MTD; shown in blue) was able to detect the 'switch' between State 1, in which there was no covariance structure (shown in white below the plot), and State 2, in which sparse covariance structure was present in the data (shown in black) with more precision than the sliding window Pearson's correlation technique (SWPC; shown in red) across a range of connectivity strengths (top panel). Values reflect the mean connectivity strength (standardized to the maximum estimated value of the correlation strength for each simulation) across all window lengths and all correlation strengths. The bottom panel shows the average 'confidence' in the occurrence of a switch at a given point in the time series, averaged across window length and correlation strength. For both sustained switches (a) and transient switches (b), the MTD was able to estimate alterations in connectivity with more precision, as evidenced by steeper gradients of change in correlation strength estimation (top panel) and change point detection (bottom panel).

Table 1
Sensitivity and Specificity values for sustained shifts in correlation using both the Multiplication of Temporal Derivatives and Sliding Window Pearson's Correlation across a range of correlation strengths (r = 0.1 − 0.5) and window lengths (w = 1 − 10).

Window (w)	r = 0.1		r = 0.2		r = 0.3		r = 0.4		r = 0.5	
	Sensitivity	Specificity	Sensitivity	Specificity	Sensitivity	Specificity	Sensitivity	Specificity	Sensitivity	Specificity
Multiplication of Temporal Derivatives										
1	0.40	0.94	0.48	0.94	0.52	0.95	0.58	0.94	0.62	0.93
2	0.59	0.85	0.64	0.84	0.74	0.84	0.77	0.85	0.81	0.83
3	0.69	0.78	0.75	0.77	0.81	0.80	0.86	0.77	0.88	0.74
4	0.75	0.72	0.80	0.71	0.87	0.70	0.91	0.71	0.91	0.69
5	0.76	0.70	0.81	0.69	0.87	0.69	0.91	0.68	0.93	0.67
6	0.75	0.70	0.80	0.68	0.87	0.68	0.90	0.68	0.93	0.67
7	0.73	0.70	0.79	0.69	0.85	0.69	0.90	0.67	0.93	0.67
8	0.71	0.71	0.78	0.70	0.84	0.70	0.90	0.67	0.92	0.69
9	0.70	0.71	0.77	0.69	0.85	0.69	0.89	0.68	0.92	0.69
10	0.69	0.70	0.78	0.69	0.84	0.69	0.90	0.68	0.92	0.69
Sliding window Pearson's correlation										
1	0.28	0.97	0.35	0.97	0.38	0.96	0.44	0.97	0.46	0.94
2	0.50	0.89	0.6	0.89	0.67	0.89	0.73	0.89	0.77	0.85
3	0.64	0.83	0.71	0.82	0.79	0.81	0.84	0.80	0.89	0.77
4	0.71	0.76	0.79	0.76	0.85	0.76	0.89	0.75	0.93	0.71
5	0.74	0.73	0.81	0.74	0.87	0.73	0.91	0.71	0.95	0.68
6	0.74	0.72	0.81	0.73	0.86	0.71	0.92	0.70	0.96	0.68
7	0.74	0.72	0.81	0.72	0.86	0.72	0.92	0.71	0.96	0.67
8	0.70	0.72	0.8	0.72	0.86	0.73	0.91	0.70	0.96	0.67
9	0.70	0.72	0.81	0.73	0.86	0.72	0.91	0.69	0.96	0.68
10	0.69	0.72	0.8	0.72	0.87	0.73	0.91	0.69	0.96	0.68

consistently predicting switches earlier than they occurred and for longer periods of time (Fig. 1a). This relative lack of precision was also evidenced by a decreased sensitivity to the switch point in the SWPC method compared to the MTD (Table 1), however the MTD was associated with a slight decrease in specificity over the range of simulations (Table 1). Across the range of window lengths, a SMA using $w = 3$ struck the most effective balance between high sensitivity and specificity (Table 1).

Experiment 1b — transient alterations in connectivity

Across the range of correlation coefficients, the MTD showed a consistently higher sensitivity to transient alterations in connectivity than the SWPC analysis (Fig. 1b), albeit with a similar small tradeoff in specificity (Table 2). In a similar fashion to the sustained alterations in connectivity, a window length of 3 (i.e. 7 total time points) consistently provided the most effective balance of high sensitivity and specificity (Table 2).

Experiment 1c — effect of amplitude change on estimates of connectivity

The results of our final dynamic simulation show an important circumstance in which the estimates of connectivity from the MTD and SWPC differ in their estimates of dynamic connectivity. Although there was no underlying positive connectivity structure between the two nodes, the targeted addition of an evoked response (modeled by an increase in signal amplitude) led to a large, sustained increased in the estimation of connectivity using the SWPC approach across all

Table 2
Sensitivity and Specificity values for transient shifts in correlation using both the Multiplication of Temporal Derivatives and Sliding Window Pearson's Correlation across a range of correlation strengths (r = 0.1 − 0.5) and window lengths (w = 1 − 10).

Window (w)	r = 0.1		r = 0.2		r = 0.3		r = 0.4		r = 0.5	
	Sensitivity	Specificity	Sensitivity	Specificity	Sensitivity	Specificity	Sensitivity	Specificity	Sensitivity	Specificity
Multiplication of Temporal Derivatives										
1	0.40	0.95	0.42	0.94	0.41	0.95	0.49	0.94	0.55	0.94
2	0.67	0.82	0.69	0.79	0.71	0.82	0.78	0.81	0.80	0.82
3	0.72	0.76	0.74	0.75	0.75	0.76	0.81	0.74	0.84	0.76
4	0.68	0.77	0.73	0.76	0.74	0.76	0.80	0.75	0.82	0.79
5	0.68	0.79	0.72	0.78	0.72	0.77	0.80	0.77	0.81	0.79
6	0.69	0.78	0.73	0.79	0.70	0.77	0.80	0.78	0.81	0.79
7	0.69	0.78	0.73	0.79	0.69	0.77	0.79	0.78	0.80	0.78
8	0.70	0.79	0.72	0.79	0.66	0.78	0.78	0.79	0.78	0.78
9	0.69	0.79	0.71	0.79	0.63	0.76	0.75	0.78	0.77	0.77
10	0.69	0.78	0.69	0.78	0.63	0.76	0.75	0.78	0.73	0.79
Sliding window Pearson's correlation										
1	0.29	0.96	0.30	0.96	0.27	0.96	0.37	0.97	0.41	0.96
2	0.56	0.87	0.56	0.87	0.56	0.88	0.67	0.87	0.73	0.89
3	0.66	0.83	0.67	0.81	0.66	0.82	0.77	0.82	0.81	0.84
4	0.65	0.83	0.70	0.82	0.68	0.82	0.77	0.82	0.82	0.83
5	0.65	0.81	0.68	0.82	0.69	0.81	0.77	0.81	0.78	0.84
6	0.66	0.82	0.69	0.82	0.66	0.80	0.78	0.82	0.79	0.82
7	0.66	0.82	0.70	0.81	0.66	0.80	0.75	0.82	0.78	0.82
8	0.66	0.82	0.68	0.81	0.63	0.81	0.72	0.82	0.76	0.82
9	0.64	0.82	0.69	0.82	0.63	0.81	0.70	0.82	0.74	0.81
10	0.65	0.81	0.65	0.81	0.60	0.80	0.69	0.81	0.72	0.81

Table 3
Effect of noise addition and removal.

	MTD		SWPC	
	Sensitivity	Specificity	Sensitivity	Specificity
Raw time series				
w = 3	0.81	0.80	0.77	0.80
Sinusoid waves (Hz)				
0.001	0.81	0.80	0.78	0.80
0.02	0.81	0.80	0.68	0.80
0.04	0.81	0.80	0.65	0.80
0.06	0.79	0.80	0.62	0.80
0.08	0.77	0.80	0.57	0.80
0.1	0.78	0.80	0.55	0.80
0.2	0.66	0.80	0.53	0.80
0.4	0.71	0.76	0.36	0.80
Head motion spikes				
1%	0.82	0.80	0.77	0.80
5%	0.81	0.80	0.76	0.80
10%	0.80	0.80	0.73	0.80
15%	0.77	0.79	0.72	0.79
20%	0.76	0.78	0.68	0.79
25%	0.74	0.76	0.63	0.79
High-pass filtering				
0.005	0.81	0.80	0.77	0.80
0.006	0.81	0.80	0.77	0.80
0.007	0.81	0.80	0.77	0.80
0.008	0.81	0.80	0.77	0.80
0.010	0.81	0.80	0.77	0.80
0.013	0.81	0.80	0.77	0.80
Band-pass filtering				
0.001–0.1	0.90	0.80	0.77	0.80

The effect of noise addition to the sensitivity and specificity of the Multiplication of Temporal Derivatives (MTD) and Sliding Window Pearson's Correlation (SWPC) values. Values reflect the sensitivity and specificity of each metric at r = 0.3 and w = 3, normalized against the values obtained for data without additional noise. We simulated four scenarios associated with noise: a) randomly phase-shifted sinusoid waves at a range of frequencies were added to the data to simulate low and high frequency noise; b) random fixed-amplitude 'spikes' were added to the data at different proportions to simulate sporadic head motion; c) data were high pass filtered at different frequencies (ranging from 0.005–0.13 Hz); and d) data were band pass filtered between 0.001 and 0.1 Hz, in order to simulate a preprocessing strategy common to resting state data analysis. The MTD was more resistant to the addition of sinusoidal noise than the SWPC across a range of frequencies. Unsurprisingly, both methods were susceptible to decreased accuracy with extensive head motion, however the MTD was more robust to this phenomenon at the mid-range (5–15%). High-pass filtering did not have a measurable effect on either measure, although band pass filtering improved the sensitivity of the MTD to change-point detection, possibly due to the removal of high-frequency noise from the data.

window lengths (Fig. 3). In contrast, the MTD was relatively insensitive to this change in the data, only estimating a change in connectivity in small window lengths ($w = 1$), although the effects were relatively transient. These results suggest that the MTD is more resilient to spurious correlations driven by task-induced co-activation of otherwise independent nodes.

Table 4
Task-based connectivity results.

Contrast	Seed	Target	P value
2-back			
	L DLPFC (108)	L Lat vOcc (136)	0.001
	R DLPFC (236)	R_Med_vOcc (298)	0.001
	R DLPFC (236)	L vOcc (229)	0.003
	L DLPFC (108)	L vOcc (229)	0.004
	L PPC (51)	L Lat vOcc (136)	0.017
	R DLPFC (236)	L Lat Occ (98)	0.025
Faces			
	R PPC (211)	L Lat vOcc (136)	0.005
Places			
	R DLPFC (236)	L vOcc (229)	0.002
	R DLPFC (273)	R Med vOcc (299)	0.007

Numbers in parentheses represent the parcel identification number associated with each region and p-values reflect the results from a mixed-effects general linear model analysis. Key: DLPFC — dorsolateral prefrontal cortex; PPC — posterior parietal cortex; Lat — lateral; Med — medial; vOcc — ventral occipital cortex; Occ — occipital cortex.

Experiment 1d — effect of noise addition

The addition of sinusoidal noise across a range of frequencies led to marked detriment in the performance of the SWPC, particularly with high frequency noise (Table 3). In contrast, the MTD was relatively impervious to the additional of sinusoidal noise, only displaying a deficit in change-point detection with relatively high frequency noise (0.2 and 0.4 Hz).

Experiment 1e — effect of head motion

Both methods showed decreased accuracy with extensive head motion (20–25% of trial with sporadic global 'spikes'), however the MTD was more robust to this phenomenon at the mid-range (5–15%; Table 3). Importantly, both the MTD and SWPC were able to detect a relative difference between State 1 and State 2 in the highest noise states.

Experiment 1f — effect of filtering

The choice of high-pass filtering threshold did not have a measurable effect on either the MTD or SWPC, which may be due to the fact that the temporal differencing of the time series acts like a HPF (see Fig. S4). In contrast, band pass filtering improved the sensitivity of the MTD to change-point detection, possibly due to the removal of high-frequency noise from the data (Table 3).

Experiment 2 — stationary network detection

The MTD metric was associated with a mean c-sensitivity of 0.76 +/− 0.2 (see Fig. S2), which was greater than the mean c-sensitivity achieved by the methods reported in Smith et al. (0.63 +/− 0.2; $t = 6.7$, $p < 0.001$), but not as sensitive as the best methods (0.86 +/− 0.2). Importantly, the functional coupling metric defined at the individual level performed excellently in the first four simulations (0.92 +/− 0.1), which were designed to reflect 'normal' parameters used in 3 T fMRI BOLD imaging experiments. The MTD metric also performed well with respect to common issues associated with

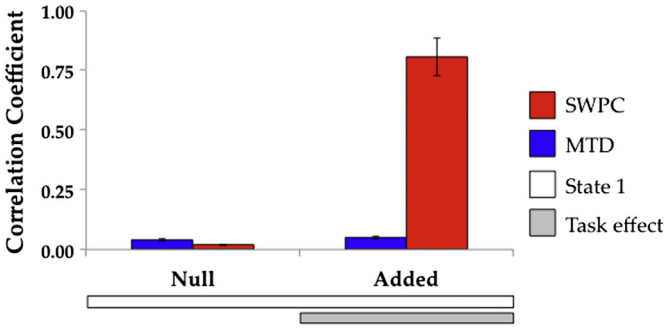

Fig. 2. Effect of amplitude change on estimates of dynamic connectivity. The addition of a 1% amplitude change (shown in gray below the graph) to the signals that did not otherwise share any covariance structure (i.e. State 1; shown in white) was associated with a large increase in estimated functional connectivity by the average sliding window Pearson's correlation (SWPC). In contrast, the average Multiplication of Temporal Derivatives (MTD) metric did not estimate significant alterations in connectivity during this perturbation. This example highlights the insensitivity of the MTD to correlations that are driven by task transients, which have long been noted as a problem for functional connectivity (e.g. Friston, 1994). Values reflect the mean connectivity strength (standardized to the maximum estimated value of the correlation strength for the r = 0.5 simulations) across all window lengths and all correlation strengths.

fMRI data time series, which were shown to be problematic for several other measures of functional connectivity (Fig. S2). Overall, the MTD metric displayed similar c-sensitivity scores to the full Pearson's correlation (mean c-sensitivity: 0.70), perhaps unsurprising due to the formal similarity between two metrics. The MTD metric did not perform as well as partial correlation (mean c-sensitivity: 87.9 +/− 16%), or sparse inverse covariance (mean c-sensitivity: 87.6 +/− 17%), however the method did out-perform the majority of methods tests across the range of simulations (mean c-sensitivity: 63.1 +/− 17%), suggesting that the MTD can be used as a tool for estimating static connectivity structure in fMRI time series data. Overall, these results suggest that this method is a useful alternative for estimating undirected static graph structure in BOLD time series data.

Experiment 3 — task-based connectivity

During blocks of high working memory load, we observed a number of significant patterns of task-based functional connectivity when using the MTD (Table 4 and Fig. 3). During blocks of high versus low cognitive load, we observed an increase in functional connectivity between the bilateral dorsolateral prefrontal cortex and the ventral occipital cortex (Fig. 3a), consistent with the visual nature of working memory task. In addition, there was also increased coupling between the posterior parietal cortex and ventral occipital cortex, as well as functional anti-coupling between the bilateral ventral occipital cortices. When contrasting patterns of face versus place recognition, we observed preferential connectivity between frontoparietal cortical regions and the ventral occipital regions responsible for the processing of either face-related (lateral) or place-related (medial) information (Fig. 3b).

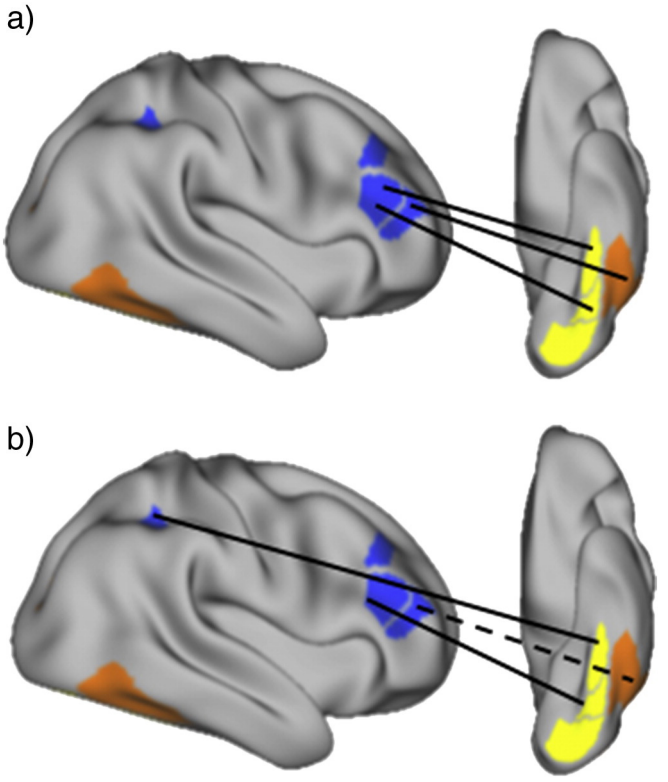

Fig. 3. Task-based functional coupling. Task-based functional coupling: a) During the performance of a visually based working memory task, we observed significant functional coupling between key frontoparietal regions (shown in blue) and ventral occipital cortical regions specialized for either the recognition of faces (orange) or places (yellow). Specifically, periods of high cognitive load, irrespective of visual object, were associated with a preferential increased coupling between the bilateral dorsolateral prefrontal cortex and the ventral occipital cortex, along with parieto-occipital coupling and inter-hemispheric anti-coupling between the ventral occipital cortex; b) The recognition of faces was predominantly associated with increased coupling between the dorsolateral prefrontal cortex and the lateral ventral occipital cortex, as well as within the lateral ventral occipital cortex (shown as a dotted line). In contrast, the recognition of places was associated with increased coupling between the lateral frontal cortex and the medial ventral occipital cortex (solid lines). $p < 0.001$ (survived correction for multiple comparisons using FDR $p < 0.05$).

Discussion

In this manuscript, we have: (i) introduced the Multiplication of Temporal Derivatives (MTD) metric as an alternative method for calculating dynamically evolving statistical relationships between nodes using BOLD time series data; (ii) demonstrated the utility of this novel functional coupling metric in detecting dynamic changes in connectivity over time using a novel simulated dataset; (iii) demonstrated that this metric can accurately estimate network structure in a previously described 'ground-truth' simulated dataset (Smith et al., 2011); and (iv) demonstrated the utility of the metric in detecting significant patterns of task-based functional connectivity in a high-quality dataset. Together, these results suggest that the functional coupling metric may be applied to functional neuroimaging data to interrogate patterns of connectivity across space and time.

A major advantage of the MTD technique lies in the superior temporal sensitivity to dynamic changes in connectivity structure compared to existing methods (see Figs. 1 & 2; Tables 1 and 2). Using data simulated using a State-Switching Model, we have demonstrated that the simple moving average of the MTD metric detects 'switches' in connectivity states with a higher sensitivity than SWPC analysis (Fig. 1a & Table 1). In addition, the MTD metric was more resilient to the introduction of signals intended to mimic the presence of task-related transient increases in BOLD signal (modeled by increasing the amplitude of signal fluctuations; Fig. 2), suggesting that the method may be more resistant to spurious correlations between independent regions that are coincidentally co-activated during task performance, which have long been noted as a problem for functional connectivity (e.g. Friston et al., 1994). Finally, the MTD was also able to detect transient changes in weak covariance structure with higher sensitivity than SWPC (Fig. 1b & Table 2), suggesting that the MTD can be applied to task-based data with high sensitivity to changes in correlation strength over time. Although there may be a theoretical risk that temporal derivatives (such as those used to estimate dynamic connectivity in the MTD; Eq. 2) might amplify noise in the data, our simulations provide strong evidence that with an appropriate window size (w), this is not the case. Specifically, the MTD was more robust to the addition of high and low frequency noise (Experiment 1d) and was also more robust to the spurious effects of head motion (Experiment 1e). Together, these analyses suggest that the MTD has the potential to interrogate dynamic alterations in functional connectivity. Together, these results highlight the utility of the MTD in estimating dynamic connectivity, but also suggest that existing dynamic connectivity work in which correlations are estimated using relatively long time windows (on the order of 30–45 time points), may be insensitive to important fluctuations in connectivity structure, although it bears mention that long windows demonstrate an improvement in specificity, particularly using the SWPC technique (Tables 1 and 2). Importantly, recent work has questioned whether dynamic alterations in connectivity can indeed be appropriately estimated using current models (Lindquist et al., 2014) and as such, future experiments should seek to clarify the nature of dynamic alterations in connectivity in a wide range of real fMRI data. An alternative approach could utilize wavelet transformation to separate data into different frequency sub-bands, with different estimates of connectivity calculated on each resultant sub-band (Achard and al, 2006). Indeed, it will be interesting to determine whether connectivity estimates created with the MTD method will be qualitatively similar to estimates of connectivity created using wavelet transformed data.

Although the MTD and SWPC are based on similar mathematical estimates of covariance over time, there are important differences

between the two methods. For instance, the MTD shows a higher sensitivity to subtle alterations in connectivity structure (Fig. 1), possibly due to an increased sensitivity to small changes in signal over time afforded by the calculation of the first temporal derivative of the time series, which acts like a HPF (see Fig. S4). In addition, as SWPC are calculated on demeaned data, the SWPC technique will necessarily over-estimate patterns of connectivity in circumstances where two independent regions show similar increases in signal amplitude, such as in task-based experiments where two independent regions may be coincidentally co-activated by task demands (Fig. 2). These results have important implications for estimates of functional connectivity between brain regions in task-based experiments, where external manipulations may independently drive increased activity in neural systems that themselves are not effectively communicating with one another (e.g. Friston et al., 1994). Together, these results suggest that by using conventional methods, such as SWPC, researchers may potentially be detecting spurious changes in covariance structure within networks as identified by task-based fMRI. However, further experiments are required to determine whether such circumstances occur in a meaningful way during the resting state (e.g. during 'switches' in internal states (Allen et al., 2014; Gorgolewski et al., 2014)).

In addition to providing increased sensitivity to dynamic alterations in connectivity, the MTD was also able to robustly estimate undirected network structure across a wide range of parameters in a 'gold-standard' simulated dataset (Smith et al., 2011). In particular, the MTD performed extremely well when analyzing imaging parameters commonly used amongst modern functional MRI experiments (i.e. 5–50 nodes, 10 min scan time, mean signal-to-noise ratio), suggesting that the MTD provides an accurate measure for assessing functional connectivity over the course of a resting state session. Together, these results suggest that, in addition to providing temporal sensitivity, the MTD is also able to effectively estimate the presence or absence of significant network edges over longer periods of time.

An important application of the MTD is its ability to estimate task-based functional connectivity. By combining the functional coupling metric with a high-quality working memory task dataset from the Human Connectome Project (Barch et al., 2013), we were able to demonstrate robust patterns of task-based connectivity during 2-back versus 0-back blocks in a visually based working memory task from the Human Connectome Project (Fig. 3a). Interestingly, the MTD also discovered separable patterns of task-based connectivity between frontoparietal regions and both medial and lateral ventral occipitotemporal cortex (Fig. 3b), which are thought to be selective for the processing of faces and places, respectively (Bell et al., 2009; Kriegeskorte et al., 2007, 2008). These results are strongly aligned with previous hypothetical (Gazzaley and Nobre, 2012; Ranganath and D'Esposito, 2005) and neuroimaging work (Axmacher et al., 2008; Zanto et al., 2011), providing evidence that effective visual working memory performance involves the flexible reconfiguration of select nodes with frontoparietal and visual networks. Despite these promising results, further work is required to determine the most efficient task-design for estimation of dynamic patterns of connectivity, as designs classically used to estimate dynamic patterns of neural activity (e.g. event related potentials in electrophysiology) bear little resemblance to the rapid event-related designs currently favored by many fMRI experiments. In addition, it is not clear whether the low pass filter of the BOLD response would render such experiments ineffective, irrespective of task design. Despite these potential issues, the simulation experiment provides evidence that the MTD contains is useful as a tool for dynamically exploring patterns of covariance across time series.

The combination of high temporal sensitivity, effective network edge detection and ease of computation suggest a number of other potential uses for the MTD in the neuroscience literature. For instance, the MTD could be combined with resting state data and a model-free clustering approach in an attempt to detect unique 'states' within patterns of brain connectivity (e.g. Allen et al., 2014; Calhoun et al., 2014; Damaraju et al., 2014; Yang et al., 2014). The MTD could also be used in combination with graph-theoretical metrics in an effort to describe the functional connectome as it evolves over time (Bullmore and Sporns, 2009). Alternatively, the MTD could be combined with task data (as shown above) to estimate task-related connectivity, similar to a Psychophysiological Interaction analysis (Friston et al., 1997). Indeed, the ability to represent the functional connectivity between two regions over time into a single vector provides an avenue for interesting questions, including the manner in which patterns of connectivity interact with one another over time, potentially allowing for a higher level of description of dynamic neural interaction that is currently under-explored. Finally, the MTD could be used to test hypotheses relating to functional connectivity impairments in patient cohorts, both in resting-state and task data (e.g. see Shine et al., 2013).

There is currently debate in the literature regarding the principled choice of window length for dynamic estimates of connectivity using BOLD fMRI. Using simulated data, Leonardi and Van De Ville (2015) have proposed that the minimum window for estimating connectivity should be larger than the reciprocal of the slowest frequency component in the signal of interest (for instance, using a 0.01 Hz filter for data fluctuating once every 100 s). In response, Zalesky and Breakspear (2015) have since demonstrated that although temporal windows of greater than $1/f_{min}$ maximize statistical power, temporal windows of this duration may be over conservative in moderate signal to noise conditions. This is an important finding, as most published works in the field have utilized a sliding window approach with a window length of approximately 30–40 s (e.g. Allen et al., 2014; Yang et al., 2014; Zalesky et al., 2014). In this manuscript, we have shown that both the MTD and SWPC can effectively estimate alterations in connectivity using much smaller windows than currently employed in the literature, with the MTD showing the highest sensitivity to abrupt changes in covariance in small windows (Fig. 1). In addition, we have shown that the MTD is relatively robust to both low and high frequency noise (Experiment 1d) and is not adversely affected by the choice of high-pass filter (Experiment 1f). Together, these results suggest that dynamic patterns of functional connectivity can be interrogated in much smaller windows than currently applied in the neuroimaging literature.

Conclusion

Together, the results our experiments suggest that the MTD provides a robust and powerful method for estimating patterns of dynamic functional connectivity and network structure in multivariate time series data. There are a number of obvious benefits and utilities for this methodology, which will help to clarify the precise dynamic relationship between neural regions in both rest and task-based experiments and assist in the definition of the brains' functional connectome. Future experiments could potentially combine the MTD with multivariate volatility methods (Lindquist et al., 2014) or metrics that calculate connectivity based on patterns of coincident supra-threshold activity (Chen et al., 2015) to help improve the sensitivity of estimations of functional connectivity in both simulated and real fMRI data, which in turn will help to solidify our understanding of the dynamic interactions between neural regions that define the functioning of the human brain.

Acknowledgments

We would like to thank Vanessa Sochat for her critical input to the manuscript. The study was not funded by an external agency.

Appendix A. Supplementary data

Supplementary data to this article can be found online at http://dx.doi.org/10.1016/j.neuroimage.2015.07.064.

References

Achard, S., Salvador, R., Whitcher, B., Suckling, J., Bullmore, E., 2006. A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. J. Neurosci. 26 (1), 63–72.

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. http://dx.doi.org/10.1093/cercor/bhs352.

Axmacher, N., Schmitz, D.P., Wagner, T., Elger, C.E., Fell, J., 2008. Interactions between medial temporal lobe, prefrontal cortex, and inferior temporal regions during visual working memory: a combined intracranial EEG and functional magnetic resonance imaging study. J. Neurosci. 28, 7304–7312. http://dx.doi.org/10.1523/JNEUROSCI.1778-08.2008.

Barch, D.M., Burgess, G.C., Harms, M.P., Petersen, S.E., Schlaggar, B.L., Corbetta, M., Glasser, M.F., Curtiss, S., Dixit, S., Feldt, C., Nolan, D., Bryant, E., Hartley, T., Footer, O., Bjork, J.M., Poldrack, R., Smith, S., Johansen-Berg, H., Snyder, A.Z., Van Essen, D.C., WU-Minn HCP Consortium, 2013. Function in the human connectome: task-fMRI and individual differences in behavior. NeuroImage 80, 169–189. http://dx.doi.org/10.1016/j.neuroimage.2013.05.033.

Baum, L.E., Petrie, T., 1966. Statistical inference for probabilistic functions of finite state Markov chains. Ann. Math. Stat. 37, 1554–1563.

Bell, A.H., Hadj-Bouziane, F., Frihauf, J.B., Tootell, R.B.H., Ungerleider, L.G., 2009. Object representations in the temporal cortex of monkeys and humans as revealed by functional magnetic resonance imaging. J. Neurophysiol. 101, 688–700. http://dx.doi.org/10.1152/jn.90657.2008.

Bullmore, E., Sporns, O., 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. Nat. Rev. Neurosci. 10, 186–198. http://dx.doi.org/10.1038/nrn2575.

Buxton, Wong, Frank, 1998. Dynamics of blood flow and oxygenation change during brain activation: The balloon model. Magn. Reson. Med. 39, 855–864.

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

Chang, C., Glover, G.H., 2010. Time-frequency dynamics of resting-state brain connectivity measured with fMRI. NeuroImage 50, 81–98. http://dx.doi.org/10.1016/j.neuroimage.2009.12.011.

Chen, J.E., Chang, C., Greicius, M.D., Glover, G.H., 2015. Introducing co-activation pattern metrics to quantify spontaneous brain network dynamics. NeuroImage (e-pub ahead of print).

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 (4), 907–920.

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: Clin. 5, 298–308. http://dx.doi.org/10.1016/j.nicl.2014.07.003.

Friston, K.J., 1994. Functional and effective connectivity in neuroimaging: A synthesis. Hum. Brain Mapp. 2, 56–78.

Friston, K.J., Ungerleider, L.G., Jezzard, P., Turner, R., 1994. Characterizing modulatory interactions between areas V1 and V2 in human cortex: a new treatment of functional MRI data. Hum. Brain Mapp. 2, 211–224. http://dx.doi.org/10.1002/hbm.460020403.

Friston, K.J., Buechel, C., Fink, G.R., Morris, J., Rolls, E., Dolan, R.J., 1997. Psychophysiological and modulatory interactions in neuroimaging. NeuroImage 6, 218–229. http://dx.doi.org/10.1006/nimg.1997.0291.

Friston, K.J., Li, B., Daunizeau, J., Stephan, K.E., 2011. Network discovery with DCM. NeuroImage 56 (3), 1202–1221.

Gazzaley, A., Nobre, A.C., 2012. Top-down modulation: bridging selective attention and working memory. Trends Cogn. Sci. 16, 129–135. http://dx.doi.org/10.1016/j.tics.2011.11.014.

Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M., 2013. The minimal preprocessing pipelines for the Human Connectome Project. NeuroImage 80, 105–124. http://dx.doi.org/10.1016/j.neuroimage.2013.04.127.

Gordon, E.M., Laumann, T.O., Adeyemo, B., Huckins, J.F., Kelley, W.M., Petersen, S.E., 2014. Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations. Cereb. Cortex (e-pub ahead of print).

Gorgolewski, K.J., Lurie, D., Urchs, S., Kipping, J.A., Craddock, R.C., Milham, M.P., Margulies, D.S., Smallwood, J., 2014. A correspondence between individual differences in the brain's intrinsic functional architecture and the content and form of self-generated thoughts. PLoS ONE 9, e97176. http://dx.doi.org/10.1371/journal.pone.0097176.

Handwerker, D.A., Roopchansingh, V., Gonzalez-Castillo, J., Bandettini, P.A., 2012. Periodic changes in fMRI connectivity. NeuroImage 63, 1712–1719. http://dx.doi.org/10.1016/j.neuroimage.2012.06.078.

Hutchison, R.M., Womelsdorf, T., Gati, J.S., Everling, S., Menon, R.S., 2012. Resting-state networks show dynamic functional connectivity in awake humans and anesthetized macaques. Hum. Brain Mapp. 34, 2154–2177. http://dx.doi.org/10.1002/hbm.22058.

Hutchison, R.M., Womelsdorf, T., Allen, E.A., Bandettini, P.A., Calhoun, V.D., Corbetta, M., Penna, Della, 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., 2013. Dynamic functional connectivity: promise, issues, and interpretations. NeuroImage 80, 360–378. http://dx.doi.org/10.1016/j.neuroimage.2013.05.079.

Jones, D.T., Vemuri, P., Murphy, M.C., Gunter, J.L., Senjem, M.L., Machulda, M.M., Przybelski, S.A., Gregg, B.E., Kantarci, K., Knopman, D.S., Boeve, B.F., Petersen, R.C., Jack, C.R., 2012. Non-stationarity in the "resting brain's" modular architecture. PLoS ONE 7, e39731. http://dx.doi.org/10.1371/journal.pone.0039731.

Kiviniemi, V., Vire, T., Remes, J., Elseoud, A.A., Starck, T., Tervonen, O., Nikkinen, J., 2011. A sliding time-window ICA reveals spatial variability of the default mode network in time. Brain Connect. 1, 339–347. http://dx.doi.org/10.1089/brain.2011.0036.

Kriegeskorte, N., Formisano, E., Sorger, B., Goebel, R., 2007. Individual faces elicit distinct response patterns in human anterior temporal cortex. Proc. Natl. Acad. Sci. U. S. A. 104, 20600–20605. http://dx.doi.org/10.1073/pnas.0705654104.

Kriegeskorte, N., Mur, M., Ruff, D.A., Kiani, R., Bodurka, J., Esteky, H., Tanaka, K., Bandettini, P.A., 2008. Matching categorical object representations in inferior temporal cortex of man and monkey. Neuron 60, 1126. http://dx.doi.org/10.1016/j.neuron.2008.10.043.

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

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.

Leonardi, N., Shirer, W., Greicius, M., Van De Ville, D., 2014. Disentangling dynamic networks: separated and joint expressions of functional connectivity patterns in time. Hum. Brain Mapp. 35, 5984–5995.

Lindquist, M.A., Waugh, C., Wager, T.D., 2007. Modeling state-related fMRI activity using change-point theory. NeuroImage 35 (3), 1125–1141.

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.

Liu, X., Duyn, J.H., 2013. Time-varying functional network information extracted from brief instances of spontaneous brain activity. Proc. Natl. Acad. Sci. U. S. A. 110, 4392–4397. http://dx.doi.org/10.1073/pnas.1216856110.

Magnuson, M., Majeed, W., Keilholz, S.D., 2010. Functional connectivity in blood oxygenation level-dependent and cerebral blood volume-weighted resting state functional magnetic resonance imaging in the rat brain. J. Magn. Reson. Imaging 32, 584–592. http://dx.doi.org/10.1002/jmri.22295.

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. http://dx.doi.org/10.1016/j.neuroimage.2010.08.030.

Mumford, J.A., Poldrack, R.A., 2007. Modeling group fMRI data. Soc. Cogn. Affect. Neurosci. 2, 251–257. http://dx.doi.org/10.1093/scan/nsm019.

Power, J.D., Barnes, K.A., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., 2012. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. Neuroimage. 59, 2142–2154. http://dx.doi.org/10.1016/j.neuroimage.2011.10.018.

Ranganath, C., D'Esposito, M., 2005. Directing the mind's eye: prefrontal, inferior and medial temporal mechanisms for visual working memory. Curr. Opin. Neurobiol. 15, 175–182. http://dx.doi.org/10.1016/j.conb.2005.03.017.

Robinson, L.F., Atlas, L.Y., Wager, T.D., 2014. Dynamic functional connectivity using state-based dynamic community structure: method and application to opioid analgesia. NeuroImage http://dx.doi.org/10.1016/j.neuroimage.2014.12.034.

Sakoğlu, 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. MAGMA 23, 351–366. http://dx.doi.org/10.1007/s10334-010-0197-8.

Shine, J.M., Matar, E., Ward, P.B., Frank, M.J., Moustafa, A.A., Pearson, M., Naismith, S.L., Lewis, S.J.G., 2013. Freezing of gait in Parkinson's disease is associated with functional decoupling between the cognitive control network and the basal ganglia. Brain 136, 3671–3681. http://dx.doi.org/10.1093/brain/awt272.

Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., 2011. Network modelling methods for FMRI. NeuroImage 54 (2), 875–891.

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. 109, 3131–3136. http://dx.doi.org/10.1073/pnas.1121329109.

Tagliazucchi, E., Balenzuela, P., Fraiman, D., Chialvo, D.R., 2012. Criticality in large-scale brain FMRI dynamics unveiled by a novel point process analysis. Front. Physiol. 3, 15. http://dx.doi.org/10.3389/fphys.2012.00015.

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 (Pt 1), 124–137. http://dx.doi.org/10.1016/j.neuroimage.2014.02.014.

Yarkoni, T., Poldrack, R.A., Nichols, T.E., Van Essen, D.C., Wager, T.D., 2011. Large-scale automated synthesis of human functional neuroimaging data. Nat. Methods 8, 665–670. http://dx.doi.org/10.1038/nmeth.1635.

Zalesky, A., Breakspear, M., 2015. Toward a statistical test for functional connectivity dynamics. NeuroImage 114, 466–470.

Zalesky, A., Fornito, A., Cocchi, L., Gollo, L.L., Breakspear, M., 2014. Time-resolved resting-state brain networks. Proc. Natl. Acad. Sci. 111, 10341–10346. http://dx.doi.org/10.1073/pnas.1400181111.

Zanto, T.P., Rubens, M.T., Thangavel, A., Gazzaley, A., 2011. Causal role of the prefrontal cortex in top-down modulation of visual processing and working memory. Nat. Neurosci. 14, 656–661. http://dx.doi.org/10.1038/nn.2773.