NeuroImage 135 (2016) 16–31

Contents lists available at ScienceDirect

NeuroImage

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

Functional connectivity and structural covariance between regions of interest can be measured more accurately using multivariate distance correlation

Linda Geerligs [a,b,*], Cam-CAN [b], Richard N. Henson [a,b]

[a] MRC Cognition and Brain Sciences Unit, 15 Chaucer Road, Cambridge CB2 7EF, UK
[b] Cambridge Centre for Ageing and Neuroscience (Cam-CAN), University of Cambridge and MRC Cognition and Brain Sciences Unit, Cambridge, UK, http://www.cam-can.com

ARTICLE INFO

Article history:
Received 29 January 2016
Revised 24 March 2016
Accepted 20 April 2016
Available online 23 April 2016

Keywords:
Structural covariance
Functional connectivity
Multivariate
Distance correlation
Resting state
Graph theory

ABSTRACT

Studies of brain-wide functional connectivity or structural covariance typically use measures like the Pearson correlation coefficient, applied to data that have been averaged across voxels within regions of interest (ROIs). However, averaging across voxels may result in biased connectivity estimates when there is inhomogeneity within those ROIs, e.g., sub-regions that exhibit different patterns of functional connectivity or structural covariance. Here, we propose a new measure based on "distance correlation"; a test of multivariate dependence of high dimensional vectors, which allows for both linear and non-linear dependencies. We used simulations to show how distance correlation out-performs Pearson correlation in the face of inhomogeneous ROIs. To evaluate this new measure on real data, we use resting-state fMRI scans and T1 structural scans from 2 sessions on each of 214 participants from the Cambridge Centre for Ageing & Neuroscience (Cam-CAN) project. Pearson correlation and distance correlation showed similar average connectivity patterns, for both functional connectivity and structural covariance. Nevertheless, distance correlation was shown to be 1) more reliable across sessions, 2) more similar across participants, and 3) more robust to different sets of ROIs. Moreover, we found that the similarity between functional connectivity and structural covariance estimates was higher for distance correlation compared to Pearson correlation. We also explored the relative effects of different preprocessing options and motion artefacts on functional connectivity. Because distance correlation is easy to implement and fast to compute, it is a promising alternative to Pearson correlations for investigating ROI-based brain-wide connectivity patterns, for functional as well as structural data.

1. Introduction

The brain is a network of a large number of regions, which may support different (cognitive) processes, but nonetheless interact with each other. In recent years, there has been much interest in the properties of this network, such as its modular structure and the existence of hub regions that help integrate information across brain regions (Bullmore and Bassett, 2011; Sporns and Betzel, 2016; Sporns et al., 2007). Such network analyses have become an important tool to characterize individual differences related to cognitive function, age and mental health (e.g. Alexander-Bloch et al., 2010; Brier et al., 2014; Crossley et al., 2014; Geerligs et al., 2014; Spreng and Turner, 2013; van den Heuvel et al., 2009). Three main, complementary techniques have been used to examine the network structure obtained from magnetic resonance imaging (MRI) of human participants. The first is diffusion-weighted MRI, which can be used to estimate the integrity of white-matter tracts between regions of interest (ROIs), but which is not considered here. Second is functional MRI (fMRI), in which connectivity within an individual is typically inferred by the correlation between time series of neuronal activity in each ROI. Third is structural MRI, from which the covariance between ROIs of a tissue property like grey matter volume or thickness can be examined across participants, which may reflect synchronized maturational changes in anatomically connected brain regions (Mechelli, 2005). In the remainder of this manuscript we will refer to these structural covariance analyses as estimates of structural connectivity.

MRI images typically contain of the order of 100,000 voxels, and there are several different parcellation schemes by which those voxels are grouped into ROIs. Some of these parcellations are adaptive, based on the data being analysed (Smith et al., 2013), but others typically come from a priori definitions, based on neuroanatomy (e.g. Tzourio-Mazoyer et al., 2002), task-based functional activations (e.g. Power

* Corresponding author at: MRC Cognition and Brain Sciences Unit, Chaucer Road, Cambridge CB2 7EF, UK
E-mail address: lindageerligs@gmail.com (L. Geerligs).
URL: http://www.cam-can.com (Cam-CAN).

http://dx.doi.org/10.1016/j.neuroimage.2016.04.047

et al., 2011) or prior functional connectivity results (e.g. Craddock et al., 2012; Gordon et al., 2014). Different studies use different parcellations, and ROIs selected on one criterion (e.g., neuroanatomy) may not respect divisions according to another criterion (e.g., functional activity). Once ROIs are defined, the relevant property of each ROI is normally reduced to a univariate measure by averaging the properties of voxels within that ROI (or by taking the first singular vector across voxels). Typically, the strength of connections between ROIs is then measured by the normalized covariance (Pearson correlation) across multiple measurements (time points or participants). Other methods have also been used, such as mutual information and time-lagged measures such as Granger causality, but these do not perform as well on typical fMRI data (Smith et al., 2011).

There are two distinct limitations to the ROI-based approach. First, important information might be lost by reducing each ROI to one dimension, given that techniques such as multi-voxel pattern analysis (MVPA) and representation similarity analysis (RSA) have demonstrated the importance of taking into account relative patterns across voxels (Kriegeskorte et al., 2008; Norman et al., 2006). This is especially likely for large ROIs, which are more likely to encompass distinct functional sub-regions (Gordon et al., 2014; Park et al., 2013). This problem is compounded when the same ROIs are used across participants, yet those participants have different functional organization, and/or there are errors in the coregistration of brain regions across participants. The second limitation is that covariance-based measures are not able to capture non-linear interactions between regions, yet previous studies have shown that shown that non-linear behaviour exists in regional interactions (Hlinka et al., 2011; Lahaye et al., 2003; Xie et al., 2008).

Here, we propose to use a different metric of connectivity that overcomes some of these limitations (though ultimately there is no substitute for good ROI definition). This metric is "distance correlation" (Székely et al., 2007), which estimates the multivariate dependence between high dimensional vectors, allowing for both linear and non-linear dependencies. Distance correlation therefore does not require reducing ROIs to one dimension, e.g., by averaging. We start with simulations showing how distance correlation out-performs Pearson correlation in the face of inhomogeneous ROI, and how it behaves according to ROI size, noise levels and temporal autocorrelation. We then apply distance correlation to real data, calculating ROI-by-ROI connectivity matrices for both functional and structural connectivity, and compared them with matrices obtained using the more standard Pearson correlation. More specifically, for functional connectivity, we compared the i) reliability across two scanning visits per participant, ii) similarity across a large number of individuals, iii) robustness to different sets of ROIs, and iv) robustness to different types of preprocessing and to effects of motion. For structural connectivity, we also compared reliability across two scanning visits, and furthermore, we compared the similarity of structural connectivity matrices with functional connectivity matrices.

2. Materials and methods

2.1. Participants

A sample of 236 participants (18–88 years old, M = 53.8, SD = 17.8, 119 males and 117 females) were taken from Stage 3 of the population-based sample of the Cambridge Centre for Ageing and Neuroscience (CamCAN). Participants were included if no brain abnormalities were detected, and if they completed both (f)MRI testing sessions. Participants had no contraindications to MRI, were native English-speakers, had normal or corrected-to-normal vision and hearing, scored 25 or higher on the mini mental state exam (MMSE; Folstein et al., 1975) and had no neurological disorders (see Shafto et al., 2014, for further details). Ethical approval for the study was obtained from the Cambridgeshire 2 (now East of England - Cambridge Central) Research Ethics Committee. Participants gave written informed consent.

2.2. fMRI data and image acquisition

Eyes-closed resting state functional magnetic resonance imaging (fMRI) data were collected in two separate scanning sessions, which were between three months and three years apart. MR data were collected as part of more extensive scanning sessions in a 3 T Siemens TIM Trio, with a 32 channel head-coil. The first scan lasted 8 min and 40 s (261 volumes) and the second scan lasted 5 min (152 volumes). Each volume contained 32 axial slices (acquired in descending order), with slice thickness of 3.7 mm and interslice gap of 20% (for whole brain coverage including cerebellum; TR = 1970 ms; TE = 30 ms; flip angle = 78 degrees; FOV = 192 mm × 192 mm; voxel-size = 3 mm × 3 mm × 4.44 mm). A high-resolution (1 mm × 1 mm × 1 mm) T1-weighted Magnetization Prepared RApid Gradient Echo (MPRAGE) image was acquired in both sessions. In the first session, we additionally acquired a T2-weighted structural image (1 mm × 1 mm × 1 mm) using a Sampling Perfection with Application optimized Contrasts using different flip angle Evolution (SPACE) sequence.

2.3. Data pre-processing

Pre-processing was performed using the SPM12 software (http://www.fil.ion.ucl.ac.uk/spm), as called by the automatic analysis (AA) batching system (http://imaging.mrc-cbu.cam.ac.uk/imaging/AA). For full details, see Taylor et al. (in press). In brief, fieldmaps were used to undistort the functional EPI images, which were then motion-corrected and slice-time corrected. For the first session, the T1 and T2 images were combined in order to segment various tissue classes using unified segmentation, including grey matter (GM), white matter (WM) and cerebrospinal fluid (CSF). For the second session, only the T1 images were used for segmentation. The GM and WM segments for each participant were used to create a sample-specific anatomical template, using the DARTEL procedure to optimize inter-participant alignment, separately for each session. The template for each session was subsequently transformed into MNI space, using a 12-parameter affine mapping. The EPI images were then coregistered to the T1 image, and the DARTEL flowfields and MNI transformation applied to the EPI images. The segmented images were also used to create WM and cerebrospinal fluid (CSF) masks for each participant by selecting only voxels with less than 1% of grey matter and more than 80% of WM/CSF. For the EPI images and the WM and CSF segments, we applied the DARTEL deformations and MNI transformation to the original images; for the structural connectivity analysis, we applied an addition modulation step (modulating by the Jacobean of the deformations) in order to preserve the amount of signal in the images (similar to voxel-based morphometry analyses; Ashburner and Friston, 2000).

2.4. Extended pre-processing and ROI extraction

To reduce the effects of motion on the functional connectivity results, we used a combination of approaches. The first of these was to apply the Wavelet Despike method for removing motion artefacts from fMRI data without the need for data scrubbing (Patel et al., 2014). The method detects irregular events at different frequencies by identifying chains of outlying wavelet coefficients, and projects these out of the voxel time series. The algorithm can remove both prolonged motion artefacts, such as spin-history effects, as well as higher frequency events such as spikes. The total amount of despiking performed on a dataset is quantified by the percentage of voxels containing a spike in that volume of data. Participants with an average spike percentage, in any of the mental states, of two standard