

Journal Pre-proof

Multi-dimensional connectivity: a conceptual and mathematical review

Alessio Basti , Hamed Nili , Olaf Hauk , Laura Marzetti ,
Richard N. Henson

PII: S1053-8119(20)30665-0
DOI: <https://doi.org/10.1016/j.neuroimage.2020.117179>
Reference: YNIMG 117179



To appear in: *NeuroImage*

Received date: 6 February 2020
Revised date: 13 July 2020
Accepted date: 14 July 2020

Please cite this article as: Alessio Basti , Hamed Nili , Olaf Hauk , Laura Marzetti , Richard N. Henson , Multi-dimensional connectivity: a conceptual and mathematical review, *NeuroImage* (2020), doi: <https://doi.org/10.1016/j.neuroimage.2020.117179>

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2020 Published by Elsevier Inc.
This is an open access article under the CC BY-NC-ND license.
(<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

Multi-dimensional connectivity: a conceptual and mathematical review

Alessio Basti^{a,b1}, Hamed Nili^{c1}, Olaf Hauk^d, Laura Marzetti^{a,b}, Richard N. Henson^{d,e}

¹ Joint first and corresponding authors

Affiliations

^a Department of Neuroscience, Imaging and Clinical Sciences, University of Chieti-Pescara, IT

^b Institute for Advanced Biomedical Technologies, University of Chieti-Pescara, IT

^c Wellcome Centre for Integrative Neuroimaging, University of Oxford, UK

^d Medical Research Council Cognition and Brain Sciences Unit, University of Cambridge, UK

^e Department of Psychiatry, University of Cambridge, UK

Abstract

The estimation of functional connectivity between regions of the brain, for example based on statistical dependencies between the time series of activity in each region, has become increasingly important in neuroimaging. Typically, multiple time series (e.g. from each voxel in fMRI data) are first reduced to a single time series that summarises the activity in a region of interest, e.g. by averaging across voxels or by taking the first principal component; an approach we call one-dimensional connectivity. However, this summary approach ignores potential multi-dimensional connectivity between two regions, and a number of recent methods have been proposed to capture such complex dependencies. Here we review the most common multi-dimensional connectivity methods, from an intuitive perspective, from a formal (mathematical) point of view, and through a number of simulated and real (fMRI and MEG) data examples that illustrate the strengths and weaknesses of each method. The paper is accompanied with both functions and scripts, which implement each method and reproduce all the examples.

1 Introduction

Neuroimaging research has demonstrated that the characterisation of functional interactions among regions of the brain is vital for a deeper comprehension of the brain's functioning. Indeed, some have claimed that cognitive functions are more closely related to the synergic cooperation among brain regions than to responses in individual brain regions (Bressler and Menon 2010). Functional connectivity is normally inferred by analysing the statistical dependency between time series (e.g. fMRI, EEG or MEG signals) associated with regions of interest (ROIs). There are many different methods for estimating this dependency (e.g. Hipp et al. 2012, Bressler and Seth 2010, Van Den Heuvel et al. 2010, Stam et al. 2007, Nolte et al. 2004, Lachaux et al. 1999, Biswal et al. 1995). However, one aspect is ignored in most of these methods: the fact ROIs are typically composed of multiple voxels whose associated time series contain important information about the complex dependencies among the regions themselves. A concern is that this information is potentially lost when reducing the original multi-dimensional data to a representative one-dimensional time series (Basti et al. 2018). This concern has led to the development of various "multi-dimensional connectivity" methods. Note that these methods have also been called "multivariate connectivity" (e.g. Anzellotti and Coutanche 2018; Geerligs et al., 2016), but the latter term may also refer to the estimation of multiple (one-dimensional) connections between all pairs of ROIs within a network, e.g. exploiting multivariate autoregressive modelling (Baccala and Sameshima,

2001; Harrison et al., 2003); so to avoid confusion, we use “multi-dimensional connectivity” here to refer to the estimation of a single, pairwise connection between two ROIs, but where those ROIs consist of multiple measurements (e.g. voxels).

While Anzellotti and Coutanche (2018) provided an informal and historical review of various multi-dimensional connectivity methods, our approach here is more pedagogical, in providing 1) some toy (2D) examples that help in illustrating some of the main concepts, such as mapping from voxel-space to pattern-distance space, 2) a formal (i.e. mathematical) definition of the main methods, 3) a number of simulations that illustrate the strengths and weaknesses of each method, 4) empirical examples (on both fMRI and MEG data) and 5) a traditional summary with “ten rules” for conducting multi-dimensional connectivity analysis. All of the examples are accompanied by MATLAB functions, including implementation of each metric used, which are available here: <https://github.com/RikHenson/MultivarCon/>.

2 Simple 2D illustration

The purpose of this section is to provide an intuitive understanding of some of the methods and concepts, before their formal description in the next section. The MATLAB script that can be used to reproduce these results is *Example2D.m*.

Consider a simple case with two ROIs with $N_x = N_y = 2$ voxels and $N_t = 4$ time points. Figure 1A-C shows the data plotted for each ROI, where each axis (x, y) represents one voxel, and the four time points are numbered. In Figure 1A, the dominant variance (between points 1 or 2 and 3 or 4) in both ROIs is along the direction $y = x$. Since averaging across the two voxels is equivalent to projecting onto the line $y = x$, simple averaging within each ROI followed by correlation across ROIs reveals a Pearson’s correlation coefficient of $R = 1$. Thus, this is a situation where the currently-dominant approach of one-dimensional connectivity, based on the mean ROI response, works well. By contrast, consider the case in Figure 1B, where the data in ROI2 are reflected across the y -axis with respect to those in ROI1, i.e. the dominant direction of variance in ROI2 is now orthogonal to that in ROI1. When the data are averaged across voxels, the projection of the data in ROI2 onto the line $y = x$ produces a different order of time point values than when the same projection is done in ROI1, such that the resulting Pearson correlation is $R = 0$.

One common solution to this problem is to summarise the ROI’s time series by their first temporal mode (or “eigenvector”), rather than their average. The temporal mode can be calculated

using a singular-value decomposition (SVD) of the ROI's matrix of voxels-by-timepoints (SVD is also the basis of principal component analysis). A SVD is defined formally in the next section, but it basically provides the direction (in voxel space) that captures the dominant variance in the data, and the temporal mode is the projection of the data onto that direction. In the current example, the data are projected onto $y = x$ for ROI1 (i.e. the first temporal mode is identical to the average for this ROI), but for ROI2, the data are projected onto $y = -x$ instead. Now, the relative ordering of time points on these two principal axes then becomes the same, so the Pearson correlation is now $R = 1$ again (if the direction of the principal axis is towards the upper left¹).

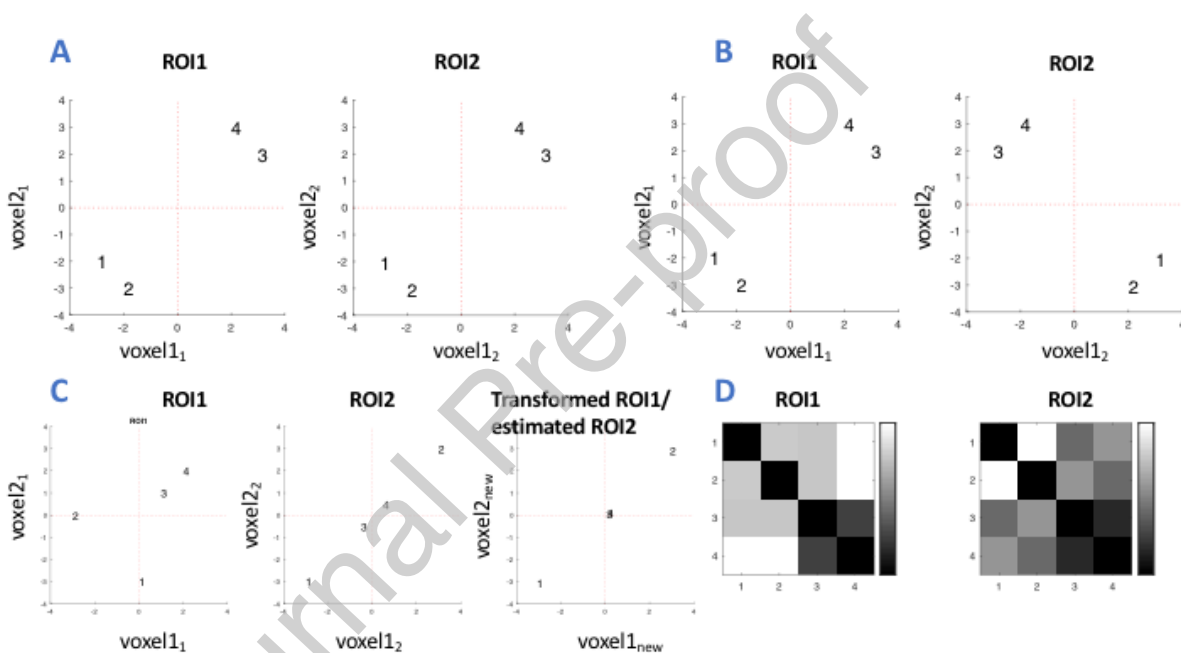


Figure 1. Toy example with 2 voxels in each of 2 ROIs, with each voxel providing data for 4 time points. Panels A-C show three different patterns of (actual and transformed) data, while Panel D shows the similarity across voxels between every pair of time points in panel C (a so-called “RDM”; see Section 2.1.4). In Panel A the dominant variance between points 1 or 2 and 3 or 4 in both ROIs is along $y = x$, while in Panel B the dominant direction in ROI2 is orthogonal to that in the first ROI. Unlike in the previous two panels, in Panel C there is a considerable variance along more than one direction, causing the one-dimensional connectivity methods to perform worse than multi-dimensional connectivity ones.

¹ The sign of the singular vectors following SVD is somewhat arbitrary, so can be flipped to produce correlations of either $R=1$ or $R=-1$. Thus, the SVD approach normally means one does not care about the sign of the connectivity; just its magnitude.

Figure 1C shows the most important case where no one-dimensional connectivity method works, and a multi-dimensional connectivity method is needed. This is because reducing the data to one dimension (even using SVD) does not capture important dependencies between the two ROIs. In particular, the first principal direction for ROI1 is again along the line $y = x$, even though there is also considerable variance along the orthogonal direction $y = -x$ (between points 1 and 2). This means that, when projecting both ROIs onto $y = x$, the order of the four time points does not match, such that the Pearson correlation following SVD is $R = 0.04$. However, if you look more closely, you can still see some dependency between the ROIs in terms of the Euclidean distances (dissimilarities) between pairs of time points. For example, points 1 and 2 are far apart in both ROIs, whereas points 3 and 4 are close together. This becomes apparent when you calculate the distance (in voxel-space) between all pairs of time points to produce the 4×4 “representational dissimilarity matrices” (RDMs) shown in Figure 1D. Although not identical, these (symmetric) matrices share properties like small (dark) values between time points 3 and 4, and relatively larger (lighter) values between time points 1 and 2. This means that when you calculate the Pearson correlation between these matrices (or just their upper right triangular elements), the coefficient is reasonably high, $R = 0.48$. This projection from voxel space to “pattern-distance space” (as expanded in the next Section) is essentially the way that many multi-dimensional connectivity methods work, such that the correlation in pattern space can be higher than in the original voxel space. In other words, even though the patterns across voxels for each time point can differ dramatically between ROIs, the similarities between those patterns can be similar across ROIs.

Finally, the way that other connectivity measures identify the multi-dimensional statistical dependency in Figure 1C is by explicitly estimating the (linear) transformation of the two axes that maximises the dependency between the two ROIs, e.g. by considering both directions $y = x$ and $y = -x$ in ROI1, in order to map the four points as closely as possible onto the same four points on the (dominant axis) of $y = x$ in ROI2. With this basic introduction in mind, we now introduce the main MD-connectivity measures more formally.

3 Formalisation of one- and multi-dimensional connectivity methods

In this section, we describe a range of one-dimensional (1D) and multi-dimensional (MD) connectivity methods, using both time-domain and frequency-domain measures. For the time-domain, we start with the two 1D-connectivity methods described in the previous section, namely a temporal correlation between the single time series resulting after either taking the average or first temporal mode across voxels. We then consider five different, time-domain MD-connectivity methods: “canonical

correlation” (Hotelling 1936), “multivariate pattern dependence” (Anzellotti et al. 2017b), “distance correlation” (Geerligs et al. 2016), “representational connectivity analysis” (Kriegeskorte et al. 2008) and “linearly predicted representational dissimilarity” (Basti et al. 2019). These represent the prototypical cases of all the major methods, of which we are aware, that have been proposed and currently used in fMRI. For the frequency-domain, we focus on phase-coupling methods. In particular, we define two 1D phase-coupling methods, i.e. “imaginary coherency” (Nolte et al. 2004) and “lagged coherence” (Pascual-Marqui 2007a), and their MD-generalisations termed “multivariate interaction measure” (Ewald et al. 2012) and “multivariate lagged coherence” (Pascual-Marqui 2007b). These are the main MD phase-coupling methods of which we are aware. Nonetheless it would also be possible to consider amplitude-based methods or information-theoretic approaches that do not disentangle between phase- and pure amplitude coupling (Barrett et al. 2010).

As shown in Figure 2, let the two ROIs be called X and Y, associated with two multiple time series composed of N_X and N_Y spatially-distinct signals, respectively. The signals within each ROI might correspond, for example, to voxels in an fMRI experiment, or source-reconstructed cortical vertices in a MEG experiment, or discrete electrodes in an extracranial or intracranial EEG experiment; though we call them “voxels” below for simplicity. The time points might be real-time samples, or they could also be estimates of event-related responses to successive trials in an fMRI experiment, or the same pre-stimulus time across successive trials in an MEG/EEG experiment (later, we distinguish between analyses based on continuous time series versus those based on multiple trials, but for simplicity, we use the term “time points” below). Thus, the multiple time series associated with X and Y can be represented as two matrices whose rows and columns denote time points and voxels, respectively. Below, we refer to the vector of values over voxels (at a single time point) as a pattern. The presence of multiple runs is important for some of the MD-connectivity metrics below, in which patterns in one run are used to predict those in another run.

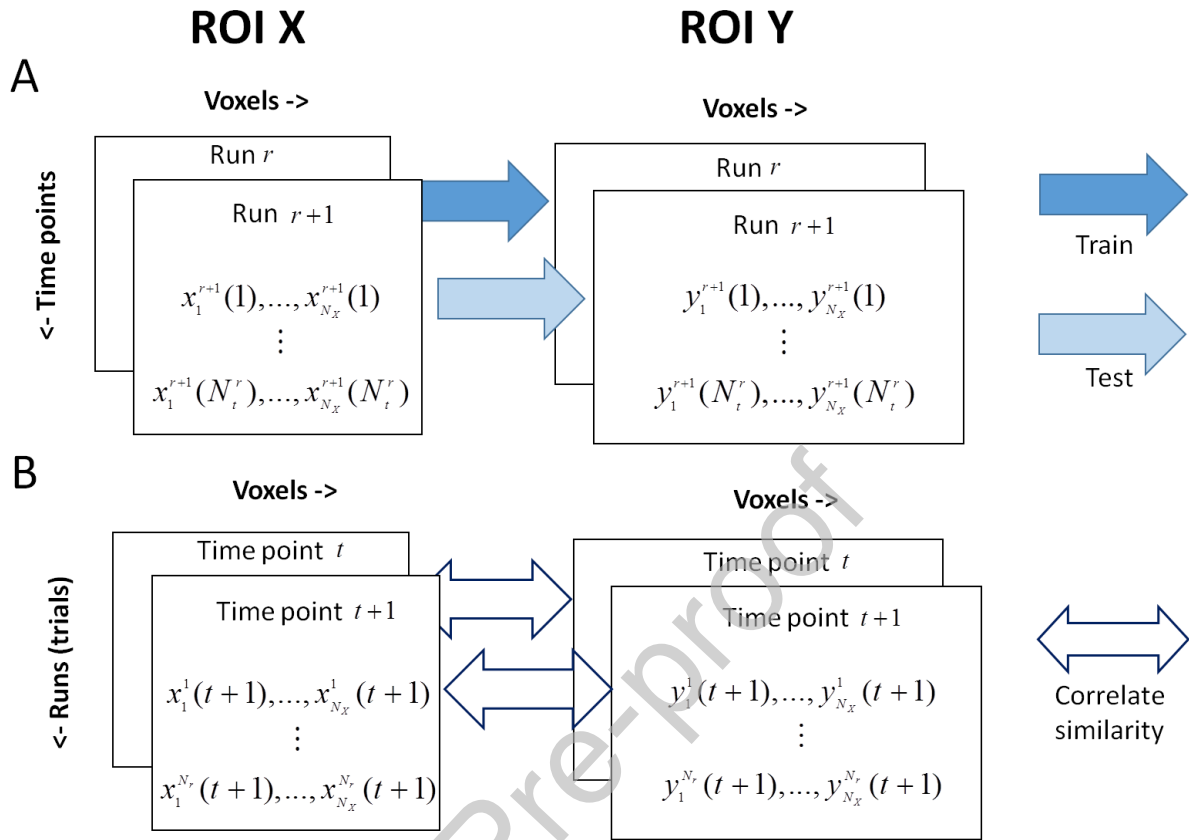


Figure 2. Estimating MD-connectivity via relationships between voxel patterns across time points (A) or across runs/trial (B). Panel A also illustrates explicit estimation of the multi-dimensional mapping between ROIs X and Y by training on one set of runs and testing on another (cross-validation), while Panel B illustrates an alternative of correlating pattern similarity kernels (e.g. representational dissimilarity matrices, RDMs; see text). The superscripts and subscripts for each time series denote runs/trials and voxels, respectively.

3.1 Time-domain methods

We start with the case of a single time series for each voxel, as in Figure 2A.

3.1.1 Pearson correlation between two single time series

To reduce the multiple time series across voxels to a single time series, the simplest summary is to take the mean across voxels. This is the most common approach in the fMRI connectivity literature, and it assumes functional homogeneity within each ROI. If this assumption is true, the averaging enhances the signal-to-noise ratio (SNR).

An alternative is to perform a singular value decomposition (SVD²) of the matrices X and Y , and use the first singular vector along the time dimension (first temporal mode) as a summary time series for each ROI (as is default for example in the SPM software package, www.fil.ion.ucl.ac.uk/spm; see Friston et al 2006). A SVD of the matrix X is defined as: $svd(X) = USV'$, where, in this case, U is a matrix whose columns are time series (temporal modes) and V is a matrix whose columns are weights across voxels (spatial modes). The matrix S is a diagonal matrix of “singular values”, in which the square of the diagonal elements relates to how much variance is explained by each pair of temporal and spatial modes, ordered from most to least. Thus, the normalised value of the first singular value (normalised by the sum of all singular values) indicates the portion of variance captured by the first temporal and spatial modes, i.e. how “successful” the dimension reduction has been. (Principal Component Analysis, PCA, is simply the SVD of the covariance matrix $X'X$.) In the example code provided, SVD is implemented in the function *dimreduction.m*.

If the ROIs are functionally homogeneous (i.e. contain identical signal time series for each voxel, plus noise that is independent across voxels), then the first spatial mode should give equal weight to each voxel, such that the first temporal mode is equivalent to the mean across voxels. However, the advantage of SVD arises if the ROIs are not functionally homogeneous (as demonstrated in Section 2, where the spatial mode can give different weights (even different signs) to voxels, and thus the first temporal mode can be very different from the mean time series. Of course, one could use SVD to define ROIs in the first place, i.e. select contiguous voxels such that they do have a common time series (i.e. high normalised first singular value). While SVD might be appropriate (for directly defining or) in case of ROIs defined through functional/adaptive approaches (see e.g. Farahibozorg et al. 2018 for adaptive approaches in MEG), it may not be appropriate if ROIs are defined a priori (e.g. based on standard anatomical parcellations). Moreover, this SVD approach is particularly valuable with signed data like in EEG/MEG, where the sign of the data depends on the sensor or source orientation, which is often not of interest, and such that straight averaging can produce values close to zero.

Whichever way the dimensionality of the ROI data is reduced to one dimension, with time series represented by the vectors \bar{x} and \bar{y} , the simplest measure of connectivity is the Pearson correlation:

$$R = \text{corr}(\bar{x}, \bar{y}) = \frac{\text{cov}(\bar{x}, \bar{y})}{\sigma_{\bar{x}}\sigma_{\bar{y}}}, \quad (1)$$

² In principle, SVD is a way of providing an alternative (orthogonalised) representation of the data, but in functional connectivity it is usually applied for the purpose of dimensionality reduction.

where cov and σ denote the covariance and the standard deviation, respectively (nonparametric measures also exist, such the Spearman rank correlation).

Another method that can be exploited in order to select the direction in the two ROIs along which to evaluate the Pearson correlation is "canonical correlation analysis" (CCA, Hotelling 1936). Even if this method applies a dimensionality reduction, it can be still considered as a MD-connectivity method. Indeed, rather than averaging across voxels or to applying a within-ROI SVD, CCA finds the two, single time series (termed "canonical variates"), \bar{x} and \bar{y} , that have the maximal 1D Pearson correlation (even if those time series do not capture the majority of variance within each ROI). That is, CCA evaluates $\bar{x} = Xu$ and $\bar{y} = Yv$, where the two directions (termed "canonical vectors") are $(u, v) = \text{argmax}_{(\tilde{u}, \tilde{v})} \{\text{corr}(X\tilde{u}, Y\tilde{v})\}$. This can be achieved by taking a SVD of the covariance matrix between X and Y (Uurtio et al. 2017). Here we use the correlation value between the first pair of canonical variables as a measure of the linear statistical dependency between X and Y , but it is worth noting that other pairs of canonical variables (which are all orthogonal) could also be selected and considered (see Wang et al. 2020 for a broader application of CCA in neuroscience).

By contrast, all the other (MD) connectivity methods below aim to reduce any dimension-reduction of the data, thus minimising the loss of potential information.

3.1.2 Multivariate pattern dependence

Multivariate pattern dependence (MVPD), as proposed by Anzellotti et al. (2017b), estimates the statistical dependency between the two matrices of time series, e.g. how well the time series in Y can be predicted using X . For linear dependence, the simplest way to do this is to use ordinary least squares (OLS) estimation, i.e. multi-dimensional regression, estimating (training) the regression weights from one set of runs and testing these weights on the remaining set of runs (i.e. cross-validation).

Since OLS estimation entails matrix inversion, the product $X'X$ needs to be invertible, one requirement for which is that X has fewer (independent) columns than rows, i.e. in this case, fewer voxels than time points. If this is not the case, one can perform a SVD and retain only the first p modes, where p is less than or equal to the number of time points (and hopefully where the p modes capture the majority of variance in X and Y).

For leave-one-run-out cross-validation, the first step is to concatenate the multiple time series across all the runs except one. For the sake of simplicity, assume we have only two runs, and use the first for training and second for testing. If a SVD is applied to $X^{r=1}$ and $Y^{r=1}$ (the superscript "r=1"

denotes the first run), then assume the dimensionality of the ROIs is reduced from N_X to K_X (and from N_Y to K_Y). The transformation matrix T between $X^{r=1}$ and $Y^{r=1}$ (or between their dimensionality reduced version) is then estimated via OLS, i.e. by solving the minimization problem $T = \operatorname{argmin}_B \{ \|X^{r=1} B - Y^{r=1}\|_F^2 \}$, where $\|\cdot\|_F$ denotes the Frobenius norm, defined as the square-root of the sum of all the squared entries of the matrix. For this minimization problem (if it is well defined) there is a unique solution that can be written as $T = (Y^{r=1})' X^{r=1} ((X^{r=1})' X^{r=1})^{-1}$, where $(\cdot)^{-1}$ denotes the inverse operation. The idea behind MVPD is to apply the transformation T in order to estimate the multiple time series $Y^{r=2} = (y_1^{r=2}, \dots, y_{K_Y}^{r=2})$ associated with the left-out run via $\hat{Y}^{r=2} = X^{r=2} T$. MVPD is then defined as the weighted-average across voxels and runs of the correlation coefficient between the actual single time series $y_i^{r=2}$ and the estimated one, $\hat{y}_i^{r=2}$, i.e.

$$\left\{ \begin{array}{l} \text{MVPD} = \frac{1}{2} \sum_{i,j} \operatorname{corr}(y_i^{r=j}, \hat{y}_i^{r=j}) w_i^{r=j} \\ w_i^{r=j} = \frac{(\sigma_{y_i^{r=j}})^2}{\sum_k (\sigma_{y_k^{r=j}})^2} \end{array} \right. \quad (2)$$

The weights $w_i^{r=j}$ are given by the portion of the overall variance explained by the i -th single time series. One advantage of MVPD is that the transformation between ROIs is explicitly represented as part of estimating connectivity. While a linear transformation has been assumed above, a different estimator (or nonlinear projection of the ROI data) could be used to estimate nonlinear transformations (Anzellotti et al. 2017a).

More generally, one does not need to perform dimensionality reduction of each ROI first in order to apply OLS, but instead use a regularised solution to the above minimization problem, such as ridge regression³ for example (see the Linearly Predicted Representational Dissimilarity method below, Basti et al. 2019). This has the potential advantage of detecting temporal components that covary strongly between X and Y , even if they only explain a small percentage of variance within X and Y separately (and therefore might be removed by the prior dimensionality reduction). Both SVD and regularised approaches are implemented in the associated function `data2mvpd_lprd_fc.m`.

³ OLS and ridge regression are just two of many linear methods to capture a statistical dependency between two matrices, e.g. voxel-by-time point matrices in the case of MD-connectivity. Other classical approaches are canonical correlation analysis (CCA), partial least squares (PLS) and least absolute shrinkage and selection operators (Kherif et al. 2002, Tibshirani 1996); however, these are beyond the present remit.

3.1.3 Distance Correlation

Distance Correlation, dCor (Székely et al 2007, Geerligs et al. 2016), is a MD-connectivity method that measures both linear and non-linear dependencies between ROIs. Conceptually, if each time point is viewed as a point in the N_X - and N_Y -dimensional voxel spaces, dCor will indicate high connectivity when a large distance between two time points within one space (ROI) is mirrored by a large distance between the same time points in the other space (ROI), and conversely when small distances in one space are mirrored by small distances in the other (as in Figure 1D).

The first step is to compute the Euclidean distances between each pair of time points s and t in voxel-space, for each ROI, i.e. $A_X(t, s) = \|X(t) - X(s)\|_2$ and $B_Y(t, s) = \|Y(t) - Y(s)\|_2$. Second, a correction, called U-centring, is applied to the time-by-time distance matrices A_X and B_Y in order to ensure that the correlation estimates are not biased by the different dimensions of the two ROIs. Finally, the distance correlation is defined as:

$$dCor = \sqrt{\max\{\text{corr}(A_X, B_Y), 0\}} \quad (3)$$

where $\text{corr}(A_X, B_Y)$ is computed between the vectorised U-centred matrices A_X and B_Y ; the term $\max\{\text{corr}(A_X, B_Y), 0\}$ included in the above equation ensures positive estimates. This is implemented in the function *data2dcor.m*.

3.1.4 Representational connectivity analysis (RCA)

U-centring aside, dCor can be seen as a special case of representational connectivity analysis (RCA, Kriegeskorte et al. 2008). In RCA, a (dis)similarity metric is estimated between every pair of patterns (i.e. between every pair of time points/trials), and then dissimilarity values for all pairwise comparisons are compared across ROIs. This resembles the “kernel method” (Theodoridis and Koutroumbas, 1998) in machine-learning, in which the original (voxel) data are transformed into a feature (pattern) space via a dissimilarity function (kernel). Indeed, dCor is a special case of RCA when the dissimilarity measure is Euclidean distance. However, other measures of dissimilarity are also possible, such as 1 minus the Pearson correlation (which is what we use here), or more generally, the Mahalanobis distance (Mahalanobis 1936).

The time point by time point matrices whose entries describe the dissimilarity (of voxel patterns) between each pair of time points are called Representational Dissimilarity Matrices (RDMs), as

in Figure 1D. The RCA between the RDMs for ROI X and Y, given by M_X and M_Y , is simply the correlation between them:

$$\text{RCA} = \text{corr}(M_X, M_Y) \quad (4)$$

where the correlation is computed between the vectorised upper (or lower) matrices M_X and M_Y . This is implemented in the function *data2rc.m*.⁴

A related idea is to train a classifier to discriminate between the patterns associated with, say, two classes of stimuli, in each ROI separately (e.g. on one subset of runs), then compute each ROI's classifier performance at each time point (in the remaining runs). The correlation between the resulting classification performances of the two ROIs across time points can then be used to indicate whether the information in the two ROIs is coupled ("informational connectivity", Coutanche and Thompson-Schill 2013). Though this provides a narrow perspective on shared information between the two regions (in that only one particular distinction is being tested), it is potentially more powerful when a single hypothesis is under consideration.

One limitation with the RCA approaches (and hence with dCor as well) is that they do not explicitly estimate the transformation between the pattern spaces represented within the ROIs. This leads to our last time-domain measure, the linearly predicted representational dissimilarity (Basti et al. 2019).

3.1.5 Linearly predicted representational dissimilarity

Finally, a time-domain measure that combines the estimation of the transformation between the multiple time series and the concept of representational connectivity is the linearly predicted representational dissimilarity (LPRD), defined in Basti et al. (2019). Similar to MVPD, an estimate of the matrix transformation T between the (Z-scored, i.e. de-means and scaled by standard deviation) multiple time series X and Y can be obtained for each run by using regularised least-squares estimation, i.e. $T = \text{argmin}_B \{ \|XB - Y\|_2^2 + \lambda \|B\|_F^2 \}$, where λ denotes a regularisation parameter. When the norm F corresponds to the (Frobenius) 2-norm of the matrix (ridge regression), there is a unique solution to the estimation of the transformation $T = (Y)'X(X'X + \lambda Id_X)^{-1}$. The regularization parameter λ can be

⁴ Note that RCA could also be performed across runs, i.e. M_X and M_Y could correspond to RDMs from two separate scanning runs. This is also implemented in the script, however, our simulations generate data independently in different runs and therefore are not suited for exploring between-run RCA.

obtained in several ways, including (nested) leave-one-out cross-validation. In this case, the optimal value is $\lambda = \operatorname{argmin}_{\alpha} \left\{ \sum_t \|X(t)B_{\alpha,t} - Y(t)\|_2^2 \right\}$, where $B_{\alpha,t}$ is the transformation obtained from ridge regression of the multiple time series X and Y , but in which the t -th time point has been removed, and with a regularisation parameter equal to α (Basti et al. 2019, Golub et al 1979).

Finally, LPRD is defined as the average (over runs) of the correlation between the RDM M_Y associated with the actual Y and the RDM $M_{\hat{Y}}$ of the estimated \hat{Y} (whose patterns $\hat{Y}(t)$, for every t , are defined as $X(t)T_t$ with the subscript denoting the removal of the time point t in the estimation of the transformation), i.e.

$$\text{LPRD} = \operatorname{corr}(M_Y, M_{\hat{Y}}) \quad (5)$$

(an across run approach similar to MVPD can be also performed). The latter is implemented in the function `data2mvpd_lprd_fc.m`.

3.2 Frequency-domain methods

Continuing with the single time series per voxel in Figure 2A, we can also calculate connectivity in the frequency-domain instead.

3.2.1 Imaginary part of coherency

The frequency-domain version of temporal correlation is coherence. In particular, coherence can be thought as the magnitude of a normalized version of the Fourier transform of the cross-covariance, i.e. the covariance between one time series and a second one shifted by l , over all the time-lags l (Brillinger, 1981). The complex version of the coherence, i.e. the one obtained without computing the magnitude, is simply called coherency (Nolte et al. 2004).

In the context of MEG/EEG data, the zero-lag component between two sensors or estimated sources can be artifactual, caused by field spread from the same true source (Marzetti et al. 2019). In the frequency-domain, this corresponds to phase differences of 0 or π , which in terms of complex number representation corresponds to situations in which the imaginary component of coherency is 0. To exclude such possibly artifactual connections, a common proposal is to use only the imaginary part of coherency (ImCoh, Nolte et al. 2004), instead of relying on its magnitude, thus reducing the zero-lag contributions associated with its real part. Using the two single time series \bar{x} and \bar{y} above to refer to a

dimensionality-reduced version of the multiple time series X and Y , the ImCoh at a frequency of interest f is defined as:

$$\text{ImCoh} = \Im \left(\frac{c_{\bar{x},\bar{y}}(f)}{\sqrt{c_{\bar{x},\bar{x}}(f) c_{\bar{y},\bar{y}}(f)}} \right), \quad (6)$$

where $\Im(\cdot)$ is the imaginary part of a complex number and e.g. $c_{\bar{x},\bar{y}}(f)$ denotes the Fourier transform (at frequency f) of the cross-covariance between \bar{x} and \bar{y} (the so called cross-spectrum). Note ImCoh is still a 1D-connectivity measure (and can be combined with a SVD to reduce each ROI time series to one dimension).

3.2.2 Multivariate interaction measure

The generalisation of ImCoh to MD-connectivity has been called a multivariate interaction measure (MIM, Ewald et al. 2012). MIM can be applied directly to X and Y without the need of dimensionality reduction. Let $C_{XY}(f)$ be the cross-spectral matrix between X and Y , i.e. the matrix whose entry in the i -th row and the j -th column denotes the cross-spectrum, $c_{x_i y_j}(f)$, between the single time series x_i and y_j . MIM is thus defined as:

$$\text{MIM} = \text{Tr} \left(\left(C_{XX}^{\Re}(f) \right)^{-1} C_{XY}^{\Im}(f) \left(C_{YY}^{\Re}(f) \right)^{-1} \left(C_{XY}^{\Im}(f) \right)' \right) \quad (7)$$

where $(\cdot)^{\Re}$ and $(\cdot)^{\Im}$ are the real and the imaginary parts of the cross-spectral matrices, and $\text{Tr}(\cdot)$ is the trace of a matrix, i.e. the sum of the diagonal entries. Similar to ImCoh, MIM also avoids inflation of the estimated values caused by artificial zero-lag connectivity. Moreover, its value is invariant under rotation of the data X and Y , making MIM independent from the choice of the physical reference frame (e.g. for source-localised MEG/EEG data).⁵

⁵ Interestingly, there is a strong mathematical relation between MIM and CCA (Kherif et al. 2002, Hotelling 1936). Indeed, most of the linear algebra steps used in Ewald et al. (2012) to introduce MIM matrix (i.e. the matrix prior the application of the trace operator) from the real part of cross-spectrum can be directly applied to the covariance for defining the canonical correlation matrix between the so-called canonical vectors. Thus, the application of the trace operator to this canonical correlation matrix defines a further time-domain measure of the total multi-dimensional (linear) statistical dependency between two sets of time series (i.e. a time-domain version of MIM).

3.2.3 Lagged coherence

Similar to imaginary coherency, Pascual-Marqui (2007a) introduced the lagged coherence (LagCoh) as a method that, by completely removing the zero-lag contributions from the signals, is only sensitive to lagged coupling. LagCoh is defined as:

$$\text{LagCoh} = \frac{[\Im(c_{\bar{x},\bar{y}}(f))]^2}{c_{\bar{x},\bar{x}}(f) c_{\bar{y},\bar{y}}(f) - [\Re(c_{\bar{x},\bar{y}}(f))]^2}, \quad (8)$$

where $\Re(\cdot)$ is the real part of a complex number. Indeed, this formulation is equivalent to the corrected imaginary coherence (Ewald et al. 2012). Note that there are yet other phase-coupling methods, such as imaginary part of the phase locking value (Palva and Palva 2012; Lachaux et al 1999) or weighted and unweighted phase lag index (Vinck et al. 2011; Stam et al 2007), but they do not yet have multi-dimensional generalisations.

3.2.2 Multivariate lagged coherence

We call the multi-dimensional generalisation of LagCoh “multivariate lagged coherence” (MVLagCoh, Pascual-Marqui 2007b), defined as:

$$\text{MVLagCoh} = \ln \frac{\det\left(\Re \begin{bmatrix} C_{YY}(f) & C_{YX}(f) \\ C_{XY}(f) & C_{XX}(f) \end{bmatrix}\right) / \det\left(\Re \begin{bmatrix} C_{YY}(f) & 0 \\ 0^T & C_{XX}(f) \end{bmatrix}\right)}{\det\left(\begin{bmatrix} C_{YY}(f) & C_{YX}(f) \\ C_{XY}(f) & C_{XX}(f) \end{bmatrix}\right) / \det\left(\begin{bmatrix} C_{YY}(f) & 0 \\ 0^T & C_{XX}(f) \end{bmatrix}\right)} \quad (9)$$

where 0 is a matrix of zeros and $\det(\cdot)$ indicates the determinant of a matrix. Similar to its 1D-version, it is sensitive to (linear) lagged interactions but insensitive to instantaneous coupling; in fact, the denominator of the above equation characterises that instantaneous coupling. Like MIM, MVLagCoh is invariant under rotation of X and Y. Since MVLagCoh is not bounded (unlike the other connectivity measures above), we use a bounded version here, i.e. $1 - 1/\exp(\text{MVLagCoh})$.

3.3 Similarity of patterns between runs/trials for each time point

Thus far, we have estimated a single measure of connectivity by using multiple time points (replications) – what might be called static connectivity (Figure 2A). Of course, one could use moving time windows to estimate connectivity changes over time. However, another common approach is to leverage the fact that, particularly in trial-based EEG/MEG experiments and single-cell recordings, there can be multiple trials, each containing the same time points with respect to the onset of a trial (e.g. epochs of -100 ms to +500 ms locked to stimuli presented every 1000 ms). In this case, one can take a single time point, and construct a RDM of similarities between all pairs of trials/stimuli for that time point (Figure 2B). In other words, one can calculate the RCA between the RDMs for a single time point and repeat across time points in order to construct a time series of dynamic connectivity between ROIs. Indeed, once one has a single time series of MD-connectivity, one can even use techniques like Granger causality on the two RCA time series, to infer which ROI is “driving” the other ROI (Goddard et al 2016; Kietzmann et al., 2019). Alternatively, one can correlate each ROI’s RDM at a given time point with an independent (model-based) RDM – based on, for example, a theoretical division of stimuli into two classes, akin to the “informational connectivity” (Coutanche and Thompson-Schill, 2013) mentioned earlier – thereby producing a single time series for each ROI that captures the similarity of that ROI’s pattern with the theoretical RDM, e.g. tracks out the dynamics of when a theoretically-relevant distinction emerges.

More generally, the above examples illustrate the flexibility of pattern-based approaches. Thus, while measures like dCor and RCA considered in the previous sections quantify the extent to which two regions have similar temporal dynamics, in terms of the geometries of the time points in each ROI’s voxel space, the same logic of correlating pattern similarity kernels (RDMs) can be applied to compare the geometry of points in each ROI’s voxel space when each point now reflects a run, trial or stimulus, rather than time point.

In the next section, we will focus on some simulation cases, showing some possible pitfalls and caveats associated with all the above connectivity methods.

4 Simulation Examples

The purpose of this section is to illustrate some of the advantages and disadvantages associated with the MD-connectivity metrics defined above. The MATLAB script for all examples is *demo.m* in <https://github.com/RikHenson/MultivarCon>.

For the first seven examples, the time series in the ROI2 for the r -th run, Y^r , is a function of those in the ROI1, X^r , i.e. for each time point t : $Y^r(t) = f(X^r(t), T^r)$, where T^r is the functional mapping that may, or may not, change across runs. For linear mappings, f is a multiplication and T^r is a $N_X \times N_Y$ matrix, leading to $Y^r(t) = X^r(t)T^r$. Finally, to simulate measurement noise, independent Gaussian noise (E^r) with mean of 0 and standard deviation equal to σ is added to both $X^r(t)$ and $Y^r(t)$.

An important property of the voxels in ROI1 (that in our examples are being mapped to ROI2) is the covariance matrix of their time series, C^r . If this matrix is such that the time series are highly positively correlated between all pairs of voxels, i.e. the ROI is functionally homogeneous (or the data are spatially smooth), then the mean time series over voxels can be a sufficient summary of activity in that ROI. Indeed, if there is additional independent noise on each time series in a ROI, then averaging is an effective way of attenuating that noise. If the voxel time series in the ROI2 are also positively correlated (which here depends on the properties of the functional mapping, T^r , e.g. whether a uniform mapping), then connectivity can be captured by a 1D metric, as shown in Example 1. However, if ROI1 is not functionally homogeneous, or if the functional mapping is not uniform, then the remaining examples illustrate the merits of using multi-dimensional connectivity metrics instead.

For the first six examples below, we assume 50 voxels for ROI1 and 60 for the ROI2, each with 400 time points, generated for two independent runs (we need more than one run in order to estimate MVPD). In the last example, we assume 12 voxels for ROI1 and 10 for the ROI2, each with 2 runs of 15,360 time points. We simulated data from 20 participants with the measurement noise in both ROIs having a standard deviation of 1, i.e. $\sigma = 1$.

4.1 Positively correlated activities in ROI1 and one-to-one voxel mapping

Figure 3A shows a covariance matrix for ROI1 that produces positively correlated time series, while Figure 3B shows a mapping matrix that produces a linear one-to-one mapping between the N_X voxels in ROI1 and the first N_X of the N_Y voxels in ROI2 (with the remaining voxels in ROI2 therefore being just random noise). Figure 3C and 3D show the first 50 time points of two voxels in each ROI and their positive correlation. Figure 3E shows the mean values for the connectivity methods described in the previous section, together with error bars for their standard deviation, across 20 simulated subjects (grey bars); the transparent blue bars show the mean scaled by the standard deviation, analogous to a Z-

statistic. Figure 3F shows their normalised values, which arise after subtracting the mean value of each metric when there is no true connectivity and then normalising by the standard deviation of the null estimates. Removing this potential bias is important because some of the metrics do not have an expected value of 0 when there is no connectivity (e.g. dCor is always positive) and also different measures do not necessarily have the same scale, thereby allowing effect sizes to be compared across the different metrics. Similar to panel 3E, the transparent blue bars give the Z-statistic for the normalised values which can be used to infer reliability of connectivity across subjects (versus zero). For each simulated subject, this null-connectivity was estimated by computing the measures from null data obtained by permuting the time points randomly for every voxel (20 times).⁶

In terms of the basic (unnormalised) values (Figure 3E), Pearson-CCA and both the two 1D metrics are best, achieving values close to 1. Pearson-CCA always selects the directions in the two ROIs that maximise the correlation and, in terms of raw performance, it does better than the other two metrics by definition. Instead, the advantage of the two 1D-connectivity methods with respect to the other MD-connectivity ones arises from the intrinsic low dimensionality of the signals generated by the high covariance among the voxels in ROI1, thus allowing an enhancement of the SNR when using a dimensionality reduction approach. The SVD approach actually does better than simply taking the average: though the voxel weights of the first spatial mode in ROI1 are fairly uniform (analogous to taking the average), they can vary sufficiently to capture small differences in the randomly-drawn time points that make some time series more or less similar to others. This difference is less evident when considering the normalised values (Figure 3F). Notably, Pearson-CCA produces lower performance than 1D-connectivity methods, owing to an overestimation of the correlation in case of no connectivity. This behaviour will reoccur in some of the next examples. Though the remaining MD-connectivity methods are not as sensitive as the two 1D-connectivity methods in this low-dimensional case, they do nonetheless produce significant (non-zero) Z-statistics, i.e. also able in principle to detect the presence of connectivity.

⁶ Different time points were independent from each other in our simulations, however, in real data there is often temporal autocorrelation. Random permuting of the time points does not preserve the temporal autocorrelation and therefore permuting for real data should preserve this autocorrelation, e.g. by using Fourier phase scrambling instead.

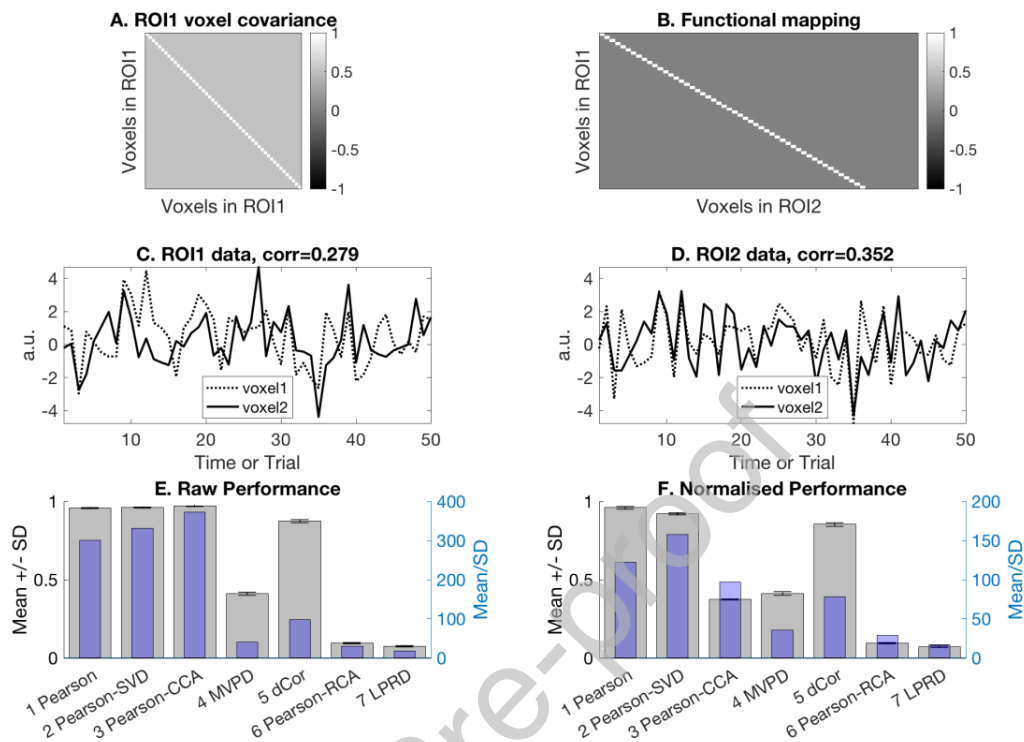


Figure 3. We simulated a one-to-one voxel mapping in which every voxel in ROI1 has a clone (except for an additive noise component) in ROI2. Panel A shows the covariance matrix of ROI1. Panel B shows the functional mapping between the two regions. Panel C and D show the first 50 time points of two voxels in each ROI and their correlation. Panel E and F respectively show the unnormalised and normalised mean (and standard deviation) values for the connectivity methods across the 20 simulated subjects. The high covariance among the voxels of ROI1 leads to intrinsic low dimensional time series, and thus to higher performance of the one-dimensional connectivity methods than the multi-dimensional connectivity ones.

4.2 Anticorrelated activities within ROI1 and one-to-one voxel mapping

Though the functional mapping is the same as the one used in the previous example (Figure 4B), this second example illustrates the presence of two functional subdivisions within ROI1, which are negatively correlated one another, as indicated in Figure 4A. This pattern has been seen in real fMRI data for example (Geerligts et al. 2016), arising perhaps when the chosen ROIs do not respect the true functional anatomy of the brain. This “structure” can be seen in Figure 4C-D, which show the time series (for the first 50 time points) for all voxels, now in an “image” format, rather than the line plots in Figure 3C-D. In this case, due to anticorrelations, averaging over voxels in ROI1 (and ROI2) destroys most of the signal, leaving just noise, and so Pearson correlation between the mean values across voxels is close to

zero (Figure 4E-F). Taking the first temporal mode from a SVD, however, recovers sensitivity. This is because some voxel weights of the dominant spatial mode are positive whereas the other voxel weights are negative, such that the combined signal is not cancelled out. Nonetheless, the signal is still one dimensional (indicated by the dominance of the first singular value of the SVD), meaning that the 1D-correlation is still more sensitive than the MD-connectivity methods.

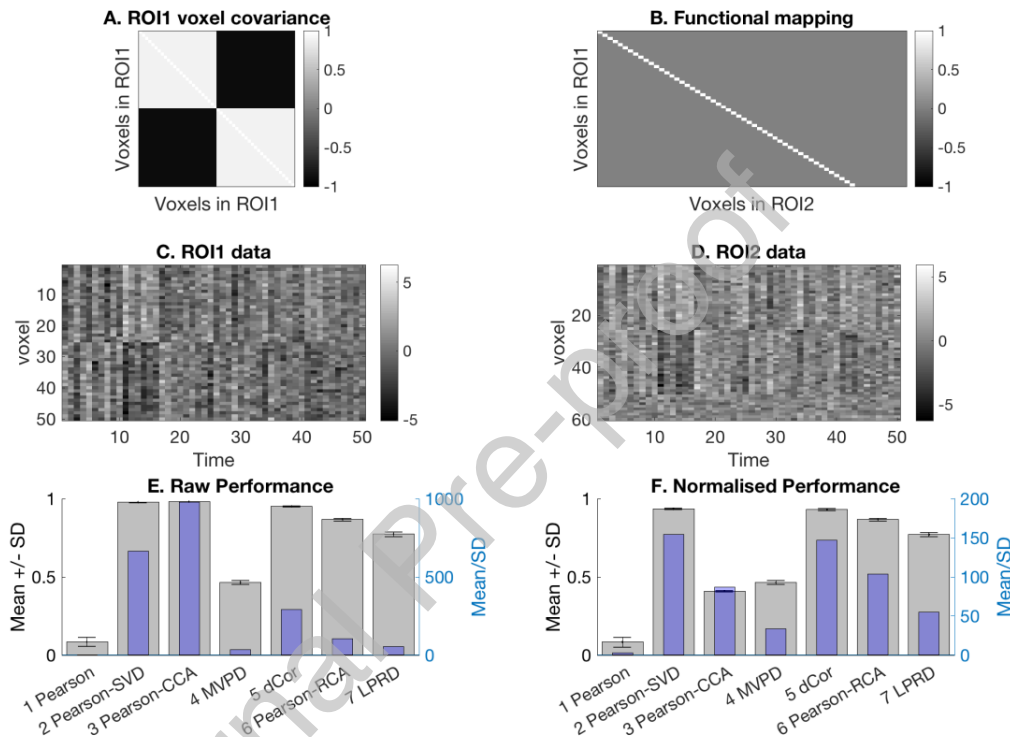


Figure 4. We simulated a one-to-one voxel mapping in which every voxel in ROI1 has a clone (except for an additive noise component) in ROI2. Panel A shows the covariance matrix of ROI1. Panel B shows the functional mapping between the two regions. Panel C and D show the first 50 time points of all the time series in each ROI. Panel E and F respectively show the unnormalised and normalised mean (and standard deviation) values for the connectivity methods across the 20 simulated subjects. In this case, ROI1 is composed of two anti-correlated subpopulations time series, such that simple averaging over voxels destroys the signal. However, by taking the first temporal mode from a SVD instead, connectivity can still be detected by a 1D Pearson's correlation.

4.3 Uncorrelated activities and multi-dimensional mapping

Figure 5A shows a covariance matrix that produces uncorrelated time series in ROI1, while Figure 5B shows a functional mapping whose entries are drawn randomly from a Gaussian distribution. The independent time series generated in ROI1 (Figure 5C) and the time series in ROI2 (Figure 5D) obtained through the application of the complex MD-mapping can no longer be optimally reduced to a

single time series by taking the average or first temporal modes (better performance is obtained by Pearson-CCA, since the two single time series are not separately selected for each ROI, but rather chosen together by taking into account the covariance between ROIs). Thus, both 1D-connectivity methods performed poorly compared to all MD-connectivity methods (Figure 5E and 5F).




Figure 5. In this example, the covariance matrix is the identity matrix that generates uncorrelated time series in ROI1. The functional mapping from ROI1 to ROI2 is generated such that its entries followed standard Gaussian distribution. The combination of these two features leads to low performance of both the 1D-connectivity methods, while the MD-connectivity methods obtained high performance.

We can use this example to compare the various metrics in their sensitivity to noise (the number of time points is 400 and the size of the two ROIs X and Y is 50 and 60, respectively). Figure 6A shows how each metric performs as the signal-to-noise ratio (SNR) increases (i.e. noise decreases from left to right). The y-axis plots the normalized performance, such that the lines and shaded areas show the mean and standard error of the difference between true and null (permuted) values. As expected, the 1D-connectivity methods perform poorly, though for high SNRs, they perform above chance, most likely because some dependency between the ROIs can still be captured by projecting onto a single dimension. In this case, selecting only one single direction per ROI based on CCA is not sufficient to

reach the performance obtained by other MD-connectivity methods, which can use multiple dimensions. Interestingly, dCor seems least sensitive to noise, and RCA is most sensitive to noise, while LPRD and MVPD show maximal performance at high SNRs. However, as apparent in subsequent examples below, there are situations other than the current case of uncorrelated activities and MD-mapping (i.e a fixed, linear mapping T plus random noise) where some MD-connectivity metrics are better than others, regardless of the SNR.

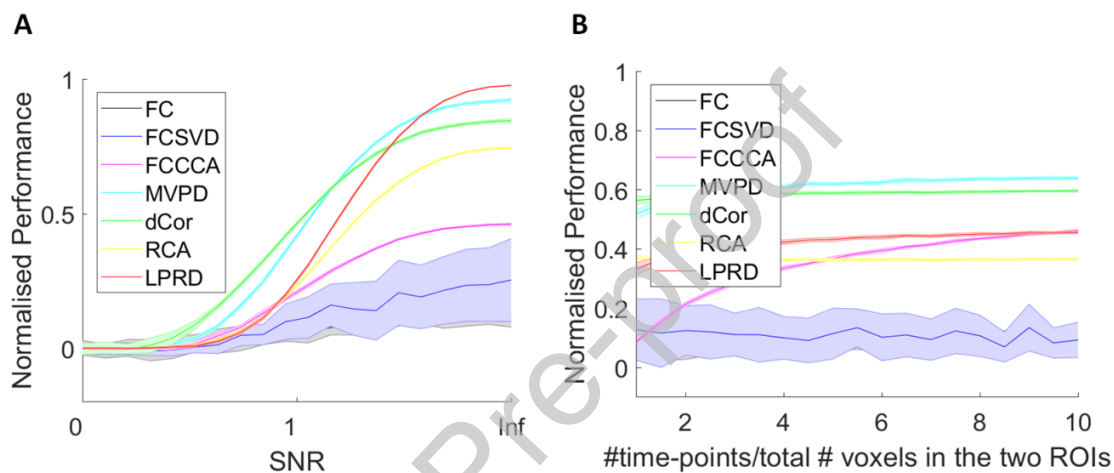


Figure 6. Panel A shows normalised performance (true minus null) of each metric for Example 3 (uncorrelated activities and multi-dimensional mapping) as signal-to-noise ratio (SNR) increases left to right. SNR is defined as the variance of signal over the variance of noise. 0 corresponds to no signal and pure noise and very large SNR means that there is very little noise variance compared to signal variance. Panel B shows normalised performance as the ratio between the number of time points and the total number of voxels in the two ROIs increases left to right.

Figure 6B shows the same normalized performance plotted now against the number of time points (relative to the total number of voxels across the two ROIs). Thus, for the same level of noise used in example 3 (i.e. a SNR approximately equal to 1), this plot shows how the metrics perform as the amount of data (duration of recording) increases. Again, MVPD does best when there are lots of data, but dCor can do better when there is little data (at least relative to the size of the ROIs). Interestingly, RCA and LPRD seem relatively robust to the amount of data, while CCA shows the greatest dependence on data duration, only approaching performance of the other metrics when the ratio of time points to voxels is high.

4.4 Multi-dimensional mapping that changes across runs

This example is identical to the previous one – in that the voxel time series are uncorrelated (Figure 7A) and the functional mapping is random (Figure 7B) – except that the functional mapping now changes across runs (Figure 7B is an example from only one run). This might happen if the voxel-wise sampling of the underlying neurons changes across runs, e.g. due to uncorrectable head motion, or if there are effects of learning across runs that change the functional connectivity. Alternatively, it might happen if different runs contain different stimuli (where each time point represents one trial with one stimulus, and there are complex interactions between neurons in the two ROIs that depend on the specific stimuli). For example, in an experiment with both auditory and visual stimuli, it might well be the case that regions have different connectivity for visual and for auditory stimuli. In any case, changes in T across runs detrimentally affect MVPD, because the MD-mapping is trained on one run and tested on others. However, the within-run measures of Pearson-CCA, dCor, RCA and LPRD remain sensitive (Figure 7F).

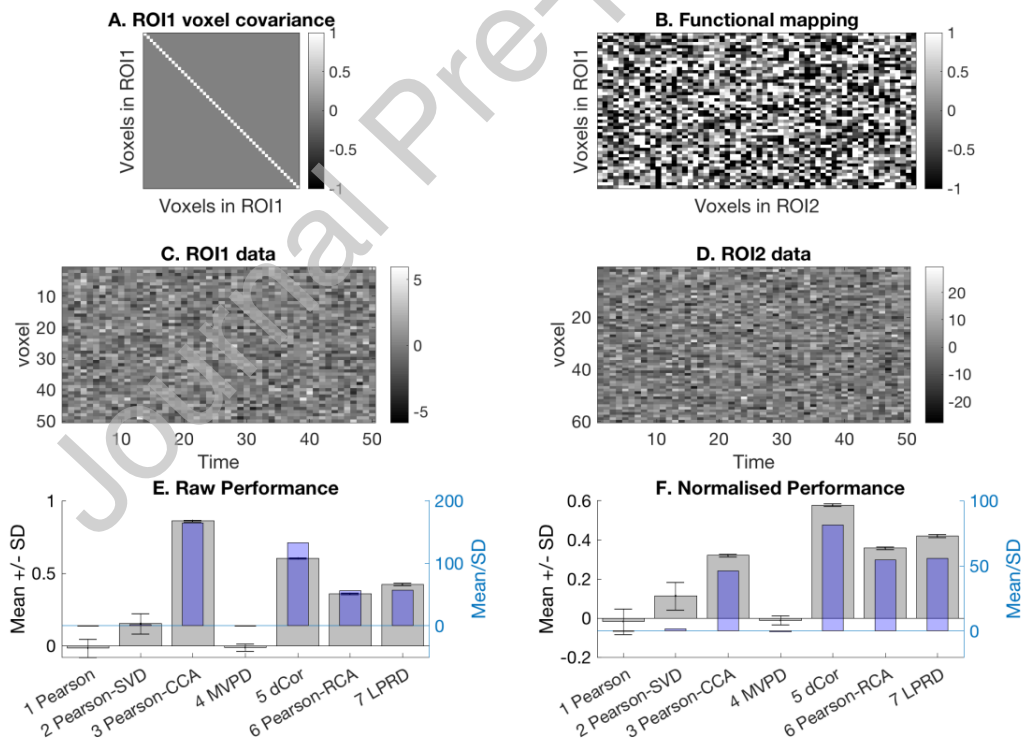


Figure 7. Inconsistent linear MD-mappings between the two regions cannot be detected by MVPD. In this example, like the previous one, the activity of each voxel in ROI2 is a weighted combination of activities of all voxels in ROI1 with some additive noise (destroying 1D-connectivity). Importantly the weights change for independent measurements, i.e. runs. Panel B shows the weights for one of the runs only (it would be different for any other

run). This severely affects MVPD. However, distance correlation, RCA and LPRD can detect these types of interactions.

4.5 Nonlinear mapping

In Figure 8, we return to the positive voxel covariance in ROI1 (Figure 8A) and the one-to-one functional mapping matrix T^r (Figure 8B) that were used in Example 1. However, the time series in ROI2 are now a nonlinear function of those in ROI1 and T^r , here illustrated by taking the absolute value of $X^r(t)T^r$. Thus, whereas the time series in ROI1 (Figure 8C) are centered around zero, the time series in ROI2 (Figure 7D) are generally above zero (except for the additive Gaussian noise). This now abolishes connectivity according to all measures but dCor, which can handle such nonlinearity (Figure 8F). Note however that this is because dCor uses a Euclidean metric of similarity of voxel-patterns between time points (trials) – if we change the similarity measure in RCA from (Pearson) correlation to Euclidean, then RCA can also produce significant connectivity just like dCor.

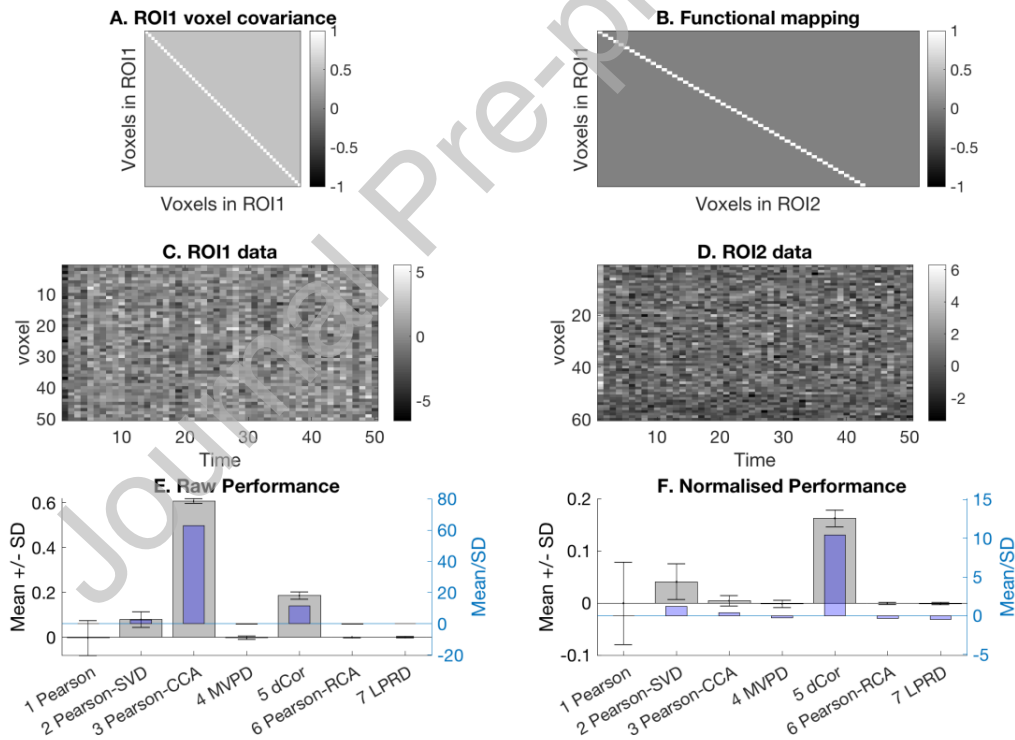


Figure 8. In this example, the covariance matrix of ROI1 and the functional mapping is exactly the same as in Example 1. The only difference to Example 1 is that the mapping has an extra nonlinearity that is applied after the linear transformation (taking the absolute value). Only dCor can detect MD-relationships in this case.

4.6 Structured noise in ROI2

In Figure 9, the functional mapping matrix (Figure 9B) is identical to that in Example 2. However, additional structured noise has been added to ROI2, which is identical across voxels (producing the coherent, vertical-bar pattern in Figure 9D). This reduces performance of dCor (Figure 9F), which uses a Euclidean measure of similarity between time points (trials), but e.g. not (the current version of) RCA, which uses a correlational measure that is invariant to any constant offsets in the voxel-patterns. Structured noise (that affects all voxels) can be present in real data due to various reasons, such as subject movement or change in the subject's alertness/vigilance across time points. Another way to eliminate this would be to normalize (z-score) the activity patterns in each time point (across voxels) before computing connectivity measures.

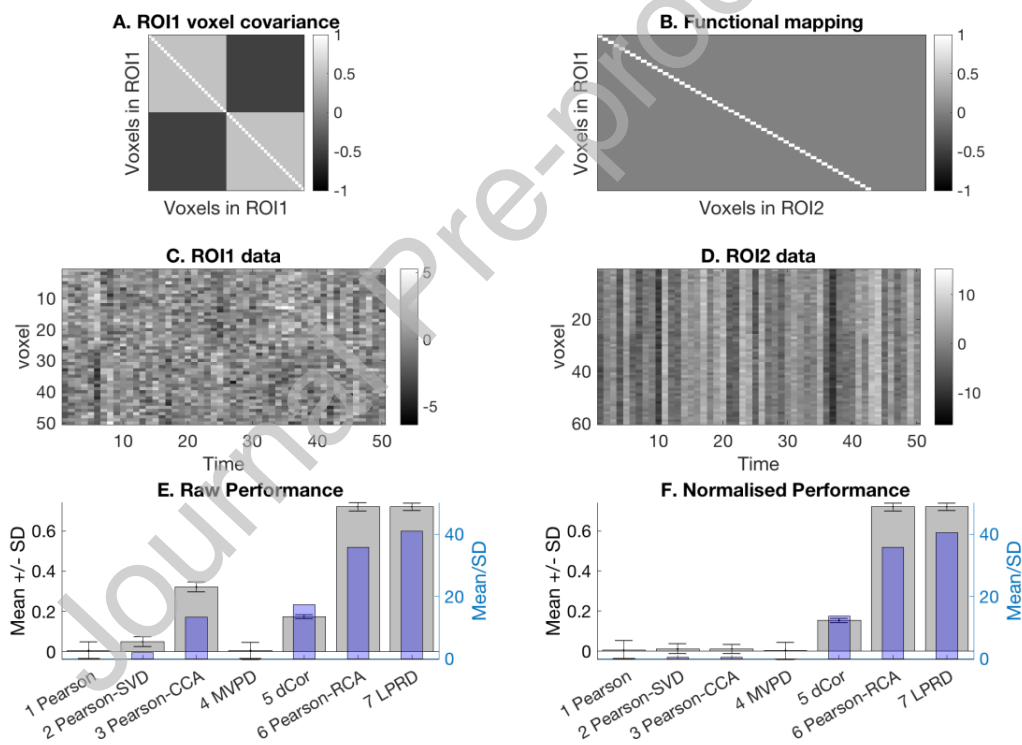


Figure 9. The presence of structured noise in ROI2 (identical across voxels but changing across time points) causes 1D-connectivity measures and dCor and MVPD to perform poorly, though Pearson-RCA and LPRD can still successfully detect the connectivity of the two ROIs.

4.7 Multi-dimensional lagged interaction

Whereas the previous six examples showed some advantages and disadvantages of connectivity methods in capturing the presence of instantaneous interaction between multidimensional signals, here we focus on lagged interactions (e.g. when time point N in ROI2 depends on time point $N-1$ in ROI1). As explained in Section 2.2, coherency can capture such lagged interactions, of which imaginary coherency (ImCoh) and lagged coherence (LagCoh) are two frequency-dependent measures that do not lead to artifactual phase-coupling due to zero-lag interactions (or more precisely, signals in phase or anti-phase) between sensors or sources in MEG/EEG data. However, these two methods were originally defined for pairs of single time series; the multivariate interaction measure (MIM) and multivariate lagged coherence (MVLagCoh) are their extensions that can capture multi-dimensional dependencies.

For this example, the time series in ROI2 is a lagged function of ROI1, i.e for each time point t , we have that $Y^r(t) = X^r(t - \tau)T^r$. Here, τ denotes the time delay in the interaction, which we set to $\tau = 10$ (which e.g. corresponds to ≈ 40 ms when the sampling frequency is set to 256Hz). To reduce computation time, we assume 12 voxels for ROI1 and 10 voxels for ROI2, each with 2 runs of 15,360 time points (typical for MEG/EEG). The elements of the transformation matrix T were again drawn randomly from a Gaussian distribution (like in examples 3-4 above). Finally, the zero-mean Gaussian additive measurement noise $E^r(t)$ was assumed independent across voxels with a standard deviation equal to 0.1 (note that results similar to those described below can also be obtained by using spatially-correlated noise, such as the one induced by field spread/volume conduction in MEG/EEG).

We compared MIM and MVLagCoh measures with ImCoh and LagCoh, with the latter applied to the first temporal modes of each ROI (labelled ImCoh-SVD and LagCOh-SVD in Figure 10). Unlike the MD-connectivity methods, both 1D-connectivity methods perform close to chance. This is because projecting the time series onto a single dimension results in smaller, and thus less detectable, phase differences. Conversely, MIM and MVLagCoh do not require the application of the dimensionality reduction and consider all the dimensions, thus resulting in higher performance. In this specific example, MVLagCoh performance overcomes the one obtained by MIM. However, it is possible that this difference is due to the specific simulation settings (e.g. type of noise and delay in the coupling).

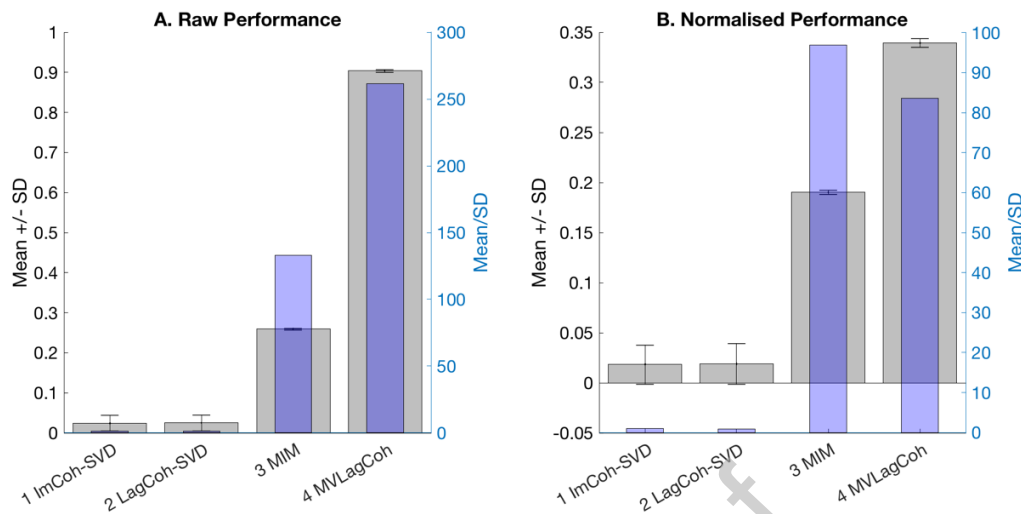


Figure 10. The simulated multi-dimensional lagged interaction between ROIs means that the 1D-connectivity methods perform poorly. On the contrary, both MD-connectivity methods show good performance.

5 Empirical examples

We compared 1D and MD measures of connectivity on real fMRI and MEG resting-state data. These data came from the 20 youngest participants (aged 18-21) in the CamCAN cohort (www.camcan.org) and from 100 randomly-chosen subjects in the Human Connectome Project (HCP) dataset (<http://www.humanconnectome.org>). The CamCAN participants only have 1 run (of fMRI and of MEG data), so are only used below for metrics that do not require cross-validation; the HCP participants have 4 runs (of fMRI data), so are used to test the cross-validated metrics.

The raw and preprocessed CamCan data are available on request here: (<https://camcan-archive.mrc-cbu.cam.ac.uk/dataaccess/>), while the final ROI-based time series are available here <ftp://ftp.mrc-cbu.cam.ac.uk/personal/rik.henson/MDcon/CamCAN>. The fMRI data and MEG data were recorded (in separate sessions) while subjects rested with their eyes closed for approximately 8.5mins. We used the fMRI data to test the non-lagged (time-domain) connectivity metrics, and the MEG data to test the lagged (frequency-domain) metrics. Since we do not know the ground truth when dealing with real data, we chose a measure of quality as the difference between connectivity estimates for homologous versus non-homologous connections between cortical ROIs: one generally expects higher connectivity between homologous ROIs. More precisely, for each ROI in the left hemisphere, we compared its connectivity to its homologous ROI in the right hemisphere, with the average of

connectivity estimates to all other right hemisphere ROIs (note that, because all our metrics are symmetrical, this is equivalent to comparing each right hemisphere ROI with all corresponding left hemisphere ROIs). We refer to this as the “homology score”, with higher scores assumed to be better. More precisely, for the CamCAN data, we used this approach for the 48 cortical, anatomical ROIs in the Harvard-Oxford Atlas (HOA atlas, <https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/Atlases>). For the HCP data, given their much larger size, we focused (for computational reasons) on the connectivity between two, sub-cortical ROIs: left and right hippocampus, given that functional heterogeneity has been proposed within the hippocampus (particularly along its anterior-posterior extent, Ranganath & Ritchey, 2012).

5.1 fMRI (zero-lag, time-domain) analyses

5.1.1 Non-cross-validated (single run) CamCAN results

The CamCAN fMRI data consisted of 261 scans (time points) every 1.97s. The original images were corrected for motion and slice-timing, warped to a standard MNI space and wavelet de-spiked (see Taylor et al., 2017, for fuller description of acquisition details and preprocessing). Time series were then extracted for each 3x3x3mm voxel within each of the 96 cortical ROIs in the HOA, which typically varied from 25 to 1643 voxels (median = 230). The script for estimating the fMRI connectivity metrics is called “test_fmri.m” in the GitHub directory.

In order to compare across metrics, each metric was normalised by 20 phase-scrambled versions of the data, producing a Z-score (similar to what was done for simulations, though unlike the random permuting of time points done there, here the scrambling of phases preserved the power spectrum, using the `phase_rand.m` function available in the GitHub directory). The difference between (noise-normalised) homologous versus non-homologous connectivity estimates was calculated for each participant. These homology scores were averaged over all 48 connections, and divided by the standard deviation across subjects, as shown in Figure 11A. As can be seen, all the metrics show a positive homology effect, as predicted, though the 1D Pearson correlation did best, with the MD dCor metric doing next best. Pearson-CCA did worst, which is likely to reflect the relatively low number of time points, or more specifically, the low mean ratio across ROIs (0.57) of time points (261) to voxels (2x230); see Figure 6B.

Figure 11A only shows average performance across ROIs. To further explore these findings, instead of averaging over connections, we averaged over subjects, and calculated the connectivity

metric that produced the best (normalised) homology score for each of the 48 connections. The results are shown in Figure 11B, which reveal a somewhat different pattern. The left axis shows the count of the number of ROIs for which that metric showed the highest normalised homology effect. As can be seen, Pearson-RCA had the highest score for the largest number of ROIs, suggesting that, even if they do not do much better on average over connections, MD methods are better for an appreciable proportion of ROIs in this fMRI dataset. The right axis shows the mean number of dimensions required by a SVD to explain 95% of the variance in the (left-hemisphere) ROIs associated with the set of connections for which that metric does best. As expected, the ROIs for which the MD Pearson-RCA (and dCor) metric did best tended to have higher dimensionality than those for which 1D Pearson did best.

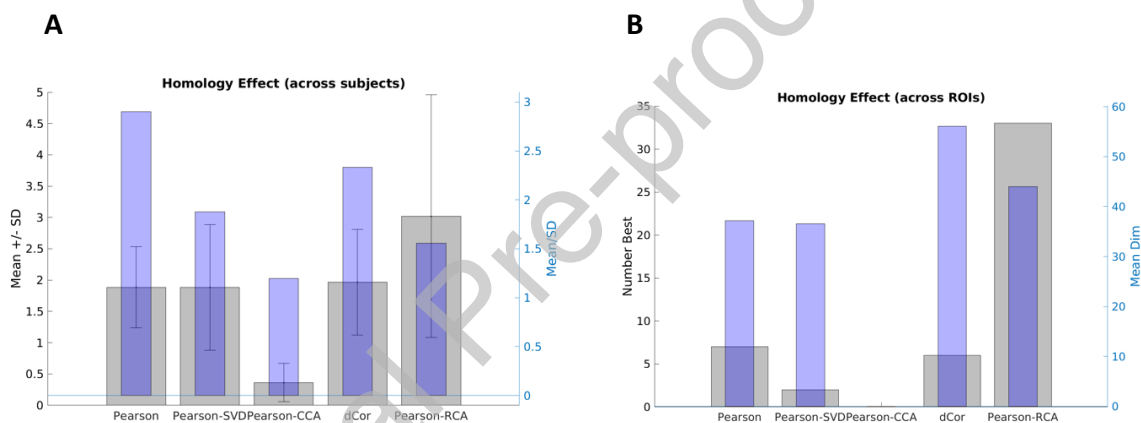


Figure 11. Performance of the non-cross-validated, zero-lag (time-domain) metrics on the single run of fMRI data from 20 subjects in the CamCAN dataset. Panel A shows mean and standard deviation across subjects of the noise-normalized connectivity values (left axis) and corresponding Z-score across subjects (right axis). The left axis of Panel B shows the number of connections (out of 48) for which each metric produced the highest noise-normalized value (on average across subjects), while the right axis shows the average number of dimensions needed to capture 95% of the variance in the left hemisphere ROI associated with each connection.

5.1.2 Cross-validated (multiple runs) HCP results

The HCP fMRI data consisted of 1200 scans acquired every 0.72s (Glasser et al., 2016, for more details of acquisition and preprocessing). The data were extracted from left and right hippocampus parcels (Glasser et al., 2016). Voxel sizes were 2mm isotropic and there were on average 230 (left) and 213 (right) voxels within the hippocampal ROIs. These data were used to test the MVPD metric that requires between-run cross-validation (which was done four-fold across the four runs, i.e. leave-one-out), though non-cross-validated metrics were also included, and their within-run results just averaged

over runs. Note that the LPRD metric was not included because our implementation additionally requires leave-one (stimulus/time point)-out cross-validation within-runs, which is difficult when time points within a run are highly auto-correlated (in this case it is preferable to use other within-run cross-validation approaches that take into account the temporal auto-correlation). While we do not know the relative SNR of the CamCAN and HCP datasets, it is worth noting (in relation to Figure 6B) that the ratio of time points to voxels in the HCP dataset ($1200/443 = 2.71$) is higher than in the CamCAN dataset (median value of 0.57 across ROIs).

The HCP results are shown in Figure 12, normalized by 20 phase-scrambled versions and averaged across subjects, as described for Figure 11A in the CamCAN data above. The results are similar to those Figure 11A, though in this case, the MD method dCor actually did better than 1D Pearson. Pearson-RCA (MD) and Pearson-SVD (1D) did reasonably well, but MVPD and Pearson-CCA did not do better than chance. The chance-performance of MVPD is similar to what was shown in simulation Example 3, where the mapping T varies between runs. It is possible that the HCP data also entail a mapping that changes across runs (e.g. owing to residual motion-related differences). Pearson-CCA does not perform well, perhaps because the ratio of time points to total number of voxels is still relatively small, even in the HCP dataset (Figure 6B).




Figure 12. Performance of the time-domain metrics across 4 runs of fMRI data from 100 subjects in the HCP dataset. See Figure 11A legend for more details.

5.2 MEG (lagged, frequency-domain) analyses

The MEG data were acquired from 102 magnetometers and 204 planar gradiometers, sampled at 1kHz. Details of the preprocessing can again be found in Taylor et al. (2017), but in brief, environmental noise was removed using signal-space separation, bad (outlying) 2s epochs were excluded, and the data bandpass-filtered from 1-48Hz. A single-shell head model was constructed from each subject's structural MRI and a scalar beamformer used to project the sensor data to approximately 8000 points on a 6mm grid within the brain. Time series (scaled by beamformer weights to adjust for depth) were then extracted for each grid point within each of the 96 cortical HOA ROIs, which typically varied from 2 to 339 "voxels" (median = 41; note these 6x6x6mm voxels are considerably larger than the fMRI ones above). To make the data size more manageable (e.g, keep the file size below 2MB), data were then downsampled to 100Hz, and only the middle half (260s) of the data extracted. The resulting time series therefore contained 26,000 time points every 10ms (100Hz). The script for estimating the fMRI connectivity metrics is called "test_meg.m" in the GitHub directory.

The frequency-domain methods were applied to a broad frequency range (4-48Hz), which encompasses the dominant brain oscillations and because cross-hemisphere coupling has been observed across many such frequencies (Vidaurre et al 2018). When normalising by 20 phase-scrambled versions, none of the MEG lagged measures (1D or MD) showed a homology effect that was significantly above 0. This may reflect the lower SNR of MEG data than fMRI data. While this might be resolved by running many more phase-scrambled versions, this becomes computationally prohibitive. Therefore, we examined the non-normalised results. Although these do not allow direct comparison across metrics, because we are subtracting non-homologous connections from homologous ones, the expected chance value of the homology score for each metric is still zero. These non-normalised results, averaged over the 48 connections, are shown in Figure 13. Importantly, the two MD metrics (MIM and MVLagCoh) both did better than their 1D equivalents (ImCoh and LagCoh respectively). In addition, the Lagged Coherence metrics did better than the Imaginary Coherence metrics (possibly due to the fact that LagCoh and MVLagCoh are better at partialling out the zero-lag contributions, which might increase the number of false negatives, than are ImCoh and MIM).

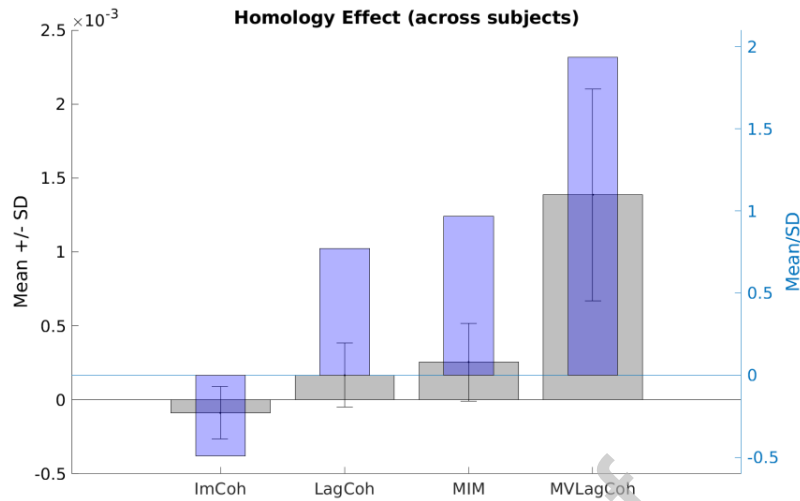


Figure 13. Performance of the non-cross-validated, lagged frequency-domain metrics on the single run of source-reconstructed MEG data from 20 subjects in the CamCAN dataset. See Figure 11A legend for more details.

6 Discussion

Neuroscience research benefits from constant advances in brain measurement and analysis techniques. For example, recent developments in machine learning and pattern analysis in particular have helped researchers gain insight into the information that is carried by populations of neurons, or that is distributed across many response channels (e.g. EEG electrodes, fMRI voxels, etc). While pattern analysis is now commonly used for information mapping for individual ROIs in the brain, research on functional connectivity between ROIs has not benefitted to the same extent. The main premise of pattern information analysis is that a region's response cannot be characterized by a one-dimensional (1D) response (e.g. average response) but can be thought of as points in high-dimensional response spaces. Recently, a number of methods have been proposed that bridge from classical 1D-functional connectivity analysis to multi-dimensional (MD) pattern analysis, which capture the shared information in two regions. In this paper, we reviewed a number of MD-connectivity methods, in a manner that we hope is helpful to both those who want an intuitive introduction and those who want a formal characterisation. We also simulated special cases that were constructed in order to illustrate the main strengths and weaknesses of each approach. The take home messages from our simulations are summarised in Table 1.

Table 1. Summary of characteristics of different functional connectivity measures tested here.

Method	1D/MD	Provide a model of transformation	Require more than one run	Handle run-specific MD-mapping (see example 4)	Able to detect nonlinear coupling (see example 5)	Sensitive to structured noise (see example 6)	Dim. reduction	Time or frequency domain (sensitive to lagged coupling)
Pearson	1D	No	No	No	No	Yes	Yes	Time (No)
Pearson-SVD	1D	No	No	No	No	Yes	Yes	Time (No)
Pearson-CCA	MD	No	No	Yes	No	Yes	Yes	Time (No)
MVPD	MD	Yes	Yes	No	No	Yes	Yes/No	Time (No)
dCor	MD	No	No	Yes	Yes	Yes	No	Time (No)
Pearson-RCA	MD	No	No	Yes	No	No	No	Time (No)
LPRD	MD	Yes	Yes/No	Yes/No	No	No	Yes/No	Time (No)
ImCoh-SVD	1D	No	No	No	No	Yes	Yes	Frequency (Yes)
LagCoh-SVD	1D	No	No	No	No	Yes	Yes	Frequency (Yes)
MIM	MD	No	No	Yes	No	Yes	No	Frequency (Yes)
MVLagCoh	MD	No	No	Yes	No	Yes	No	Frequency (Yes)

The main messages in Table 1 include the fact that multivariate pattern dependence (MVPD, Anzellotti et al. 2017b) and linearly predicted representational dissimilarity (LPRD, Basti et al. 2019) have the advantage that they provide a generative model of regional interactions (also at the level of subspaces corresponding to the dominant principal components), whereas distance correlation (Geerligts et al. 2016) and representational connectivity analysis (Kriegeskorte et al. 2008) quantify shared information without providing a model (i.e. a MD-mapping from one region to another). However, estimating this mapping requires that MVPD is given multiple datasets (runs) with which to cross-validate the mapping. This is problematic for MVPD if the mapping changes across runs (e.g. due to head motion in fMRI). LPRD (as it is implemented here) overcomes this issue by estimating the transformation in a within-run cross-validated framework, at the expense of potential biased estimates of the mapping if the noise is correlated. Also, these mappings are normally assumed to be linear, and while nonlinear mappings could be trained (e.g. using neural networks, Anzellotti et al. 2017a), this would generally require more data.

By contrast, distance-based approaches like dCor and RCA can more easily handle nonlinear mappings, e.g. through using Euclidean or other metrics of dissimilarity between patterns (indeed, dCor can be seen as a special case of RCA that uses a Euclidean metric). However, a downside of metrics like Euclidean distance (and hence dCor) is that they can be more sensitive to structured noise (e.g. between runs) than are correlational metrics (like Pearson-CCA) that normalise by the variance in each time series. An alternative solution is to Z-transform the time series prior to estimating connectivity, though this can in turn miss what might be true differences in the magnitude of dissimilarity between patterns (which is why we did not Z-transform in our examples, although our code provides this option).

Some MD-connectivity methods have the option of dimension reduction, such as using SVD to reduce to $N > 1$ components that capture the majority of variance within each ROI. However, a danger here is that there may be important covariance *between* ROIs (the primary focus of connectivity analysis) that does not represent a large proportion of the variance *within* one of the ROIs, and therefore is lost by dimension reduction. This is exactly the problem that CCA solves, by selecting one time series per ROI that simultaneously maximises the correlation between ROIs, rather than selecting one time series that summarises each ROI independently. Another solution (that could be applied to LPRD and MVPD) is to regularise estimation of the multi-dimensional mapping between ROIs, which encourages dimension reduction at the same time as capturing the important covariance between ROIs.

In Section 4.3, we also explored the sensitivity of each metric to noise. While this plot only applies when T is a fixed, linear mapping and the noise is independent across voxels, it suggests that some measures like dCor are more robust to noise than others like RCA. This robustness will depend on the details of the algorithm (e.g. measure of dissimilarity used in RCA), which is beyond the current remit, but the simulation results suggest that one might want to consider a range of MD-connectivity metrics on a (held-out) dataset, to see which is most robust to the type and level of noise in that dataset.

Finally, of the methods considered here, only the multivariate interaction measure (MIM, Ewald et al. 2012) and multivariate lagged coherence (MVLagCoh, Pascual-Marqui 2007b) can currently capture time-lagged, multi-dimensional phase-coupling between ROIs. We suspect this will be particularly important for source-reconstructed MEG/EEG data, where ROIs contain multiple time series from dipoles with different orientations (given that the sign of the data depends on the dipole orientation), which means that simple averaging will never be appropriate.

When applied to real data, using relatively large, anatomically-defined ROIs, we found that, while 1D metrics, specifically Pearson correlation, often performed as well as MD metrics like dCor and Pearson-RCA, when averaging across connections (ROIs), there was a significant proportion of connections (in the CamCAN fMRI data) for which MD metrics did better. The advantage of MD metrics was even clearer for the (lagged) metrics applied to source-reconstructed MEG data, where MD metrics also outperformed their 1D equivalents on average across all connections. This may be because the signal in each “voxel” in source-localised MEG depends on the dominant orientation of the underlying electrical currents in that voxel, which may result in a higher dimensionality of dominant signal variance within the ROIs.

Note that, while we used fMRI to demonstrate non-lagged measures and MEG to demonstrate lagged measures, it is possible to apply lagged measures to fMRI data (e.g. after adjusting for the potential confound of different haemodynamic response functions in different ROIs) and non-lagged measures to MEG data (e.g. after adjusting for zero-lag effects of field-spread). Note also that, even with non-lagged measures, regional variability in the haemodynamic response function can still confound fMRI measures of functional connectivity (Rangaprakash et al. 2018).

It is important to remember that one could define ROIs as ones in which their components (e.g. voxels) are functionally homogeneous, and therefore ones where averaging and 1D-connectivity is the correct approach by definition. For example, ROIs could be defined functionally by contiguous collections of voxels that have similar time series, or even similar 1D-connectivity (Craddock et al. 2015). This would negate the need for any MD-connectivity. However, in many situations, the homogeneity of ROIs is not known, for example when they are defined anatomically, or when their functional definition in one population (e.g. young people) may not generalize to another (e.g. older people; Geerligs et al. 2017). Relatedly, it is worth noting that investigating the information content of activity patterns can uncover processes that are not accessible by 1D-connectivity methods. For example, while average single activations for Gabor patches of various orientations might be equivalent in early visual areas, activity patterns in those areas nonetheless demonstrate explicit information about the orientation of the stimuli. Therefore, it may be the case that averaging activations across response channels masks the fact that two regions have different response properties, even though they are called “connected” according to 1D analysis. This is the converse to what we show in our examples, i.e. that two regions might have shared information in their MD-responses but no 1D-connectivity.

6.1 Ten rules for multi-dimensional functional connectivity

Below we offer ten rules for someone considering application of MD-connectivity measures:

1. If you are confident that all your ROIs are homogeneous, then 1D-connectivity methods can be more sensitive than MD-connectivity methods. In this case, the mean across voxels is a sufficient summary (though SVD should produce very similar answers). The degree of 1D-connectivity might be estimated by linear methods like Pearson correlation coefficient, or Spearman rank correlation (e.g. if data distribution contains outliers) or nonlinear methods like mutual information or (e.g. for EEG/MEG) time-lagged measures like the imaginary part of coherency, imaginary part of phase locking value, lagged coherence or phase lag indices (Marzetti et al. 2019). However, homogeneous ROIs are rarely guaranteed, and examining the singular values of a SVD can help give an idea of homogeneity.⁷
2. If you are confident that a ROI can be considered as one-dimensional, but variation along that dimension might be captured better by some voxels than others (e.g. owing to differing noise levels) or even expressed differently across voxels (e.g. some may be activated and some deactivated), then consider applying 1D-connectivity measures based on SVD rather than those based on the mean across voxels. Alternatively, if you think that the connectivity between ROIs is one-dimensional, but that each ROI contains a lot of independent noise (which would affect a SVD performed within each ROI), then consider using Pearson-CCA (assuming that there is a sufficient number of time points/scans).
3. If a ROI contains multiple dimensions of variation (e.g. SVD does not suggest a single dominant mode), then consider a MD-connectivity measure. If you do not care about the explicit form of the mapping between ROIs in fMRI, but wish to allow for nonlinear as well as linear mappings, then consider dCor. dCor does not require cross-validation across runs and appears quite robust to noise. Note that dCor is a special case of RCA; a case that uses a Euclidean measure of similarity plus some extra steps to centre the time series. However, be careful about structured forms of noise (noise that is correlated across voxels), which can reduce the sensitivity of dCor. Some of the effects of this structured noise can be ameliorated by normalising your time series or switching to a normalised measure of similarity like a correlation coefficient (the most common choice in RCA).
4. As opposed to MD-connectivity methods that only test for the presence of shared information, MVPD and LPRD give insight into the nature of that shared information, by providing an estimate of the

⁷ Of course, homogeneity at the level of measurement (e.g. voxels in fMRI) does not imply functional homogeneity at finer spatial scales: e.g. voxels could average over many functionally different neuronal populations, effectively reducing a multi-dimensional neural space to one dimension of data.

mapping between ROIs. Thus, if you want to estimate the functional transformation between two regions (besides simply testing for the presence of MD-connectivity), you should consider using one of these two metrics. It is worth noting that the mapping estimated by MPVD is conceptually different from the mapping estimated by LPRD: the former method relies on a transformation between two subsets of temporal modes, while the latter directly relies on a MD-mapping between voxels (in particular on a voxel-by-voxel linear transformation). Finally, both MVPD and LPRD will only detect linear mappings (though nonlinear mappings could be captured by projecting the data to basis functions or using more advanced nonlinear estimation techniques such as neural networks, Anzellotti et al. 2017a).

5. If you want to investigate the presence of MD phase-coupling (e.g. in MEG or EEG) between ROIs, then consider MIM or MVLagCoh (or multivariate phase slope index for directed multi-dimensional interactions, Basti et al. 2018). Instead, if you are interested in amplitude-coupling (the other main coupling mode, Siems and Siegel 2020, Engel et al. 2013), an extension/modification of these measures is needed. For instance, one strategy could be to average all the possible pairwise (orthogonalised) amplitude envelope correlations between either the time series or the temporal modes associated with the two ROIs. Instead, if you want to investigate MD directed lagged interaction without disentangling phase- from pure amplitude-coupling, then consider applying properly (e.g. handling of the artificial zero-lag contribution) a Granger-causality method to the sets of multiple time series (Barrett et al. 2010).

6. If you are interested in the representational geometry of ROIs (for example, if you have multiple stimuli across different trials for each time point), then consider RCA, specifically the explicit construction of a RDM for each ROI, which may reveal structure in your stimulus patterns. If you have an explicit hypothesis about that structure, you could estimate the degree to which each ROI's RDM matches a model RDM that encapsulates that hypothesis (or use some other form of multivariate classifier) and correlate the degree of that match across time points (informational connectivity, Coutanche and Thompson-Schill 2013).

7. If you either have only one run (or trial) or it is possible that the mapping between ROIs changes across runs, you cannot use either MVPD or the between-run version of LPRD (Basti et al. 2019), since they require cross-validation of the estimates across independent datasets. If the samples (e.g. time points) are independent, you could consider within-run cross-validation of LPRD. If you do have multiple runs, and the mapping is stable across those runs, but the presence of correlated noise cannot be

excluded (Henriksson et al. 2015), only MVPD (and the between-run version of LPRD) can control false positives, via cross-validation across the runs.

8. If the data might be contaminated by structured noise, i.e. a global signal across all voxels that changes with time (e.g. artefactual differences across trials in overall activation, when using single-trial fMRI estimates), then among the MD-connectivity methods considered in this review, only Pearson-RCA and LPRD are insensitive to such noise. In order to make other metrics like MVPD able to cope with this issue, one may consider first normalising (e.g. Z-scoring) the multiple time series within each voxel.

9. For all the MD-connectivity measures (e.g. dCor and RCA) for which the expected value in the case of no true connectivity is not zero, it is fundamental to test for significance of MD-connectivity values with some form of randomisation testing (e.g. phase-shuffling of time series data).

10. In general, there is no single, best MD-connectivity measure, and we showed that each has pros and cons. Researchers could try all of them (or a subset according to the properties we explained here) and then use a nested cross-validation scheme for choosing the “right measure” for their dataset and choice of ROIs (or correct for multiple hypothesis testing when using multiple measures).

Exploring MD-connectivity is still only in its infancy. It is likely that there is a unified mathematical framework in which the current examples are special cases. Moreover, here we have only considered methods for estimating single (pairwise) connections, but there are methods (usually called “multivariate connectivity methods”) that simultaneously estimate all connections in a network, which is needed to properly account for the fact that a “direct” connection between regions X and Y needs to partial out the contributions of “indirect” connections via a third region Z (Sanchez-Romero and Cole 2019, Stramaglia et al. 2014, Barrett et al. 2010;). Future work might enable a combined “multivariate, multi-dimensional” approach also for other types of methods (e.g. pattern-distance metrics), where multiple direct dependencies are estimated among multiple ROIs, where each ROI is itself represented by multiple response channels.

Finally, it is important to note that we have focused on the methods, using artificial examples, rather than the possible underlying neuronal truth, i.e. how neuronal populations might communicate via multi-dimensional interactions (or to use the terminology of Reid et al., 2019, we have also focused more on the “map” than the “territory”). One way to address the latter is to use sophisticated (generative), neurophysiological models of neuronal interactions, perhaps using meso-scale or mean-field approximations, together with a “forward” or “observation” model” that captures the

measurement process (e.g. fMRI or MEG), and then investigate the extent to which the current MD methods capture those interactions, compared to 1D connectivity methods. If so, this would cement the importance of considering how best to measure MD-connectivity in the brain.

Author contributions

AB, HN and RNH designed the review, wrote the manuscript and the MATLAB code (<https://github.com/RikHenson/MultivarCon/>) used for the simulations and for the real data analyses. All the authors critically reviewed and edited the manuscript.

Acknowledgements

AB has been partially funded by the project Departments of Excellence of the Italian Ministry of Education, University and Research (MIUR). AB and LM are partially supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 810377). OH and RNH are supported by UK Medical Research Council programme grant SUAG/058 G101400 and SUAG/046 G101400, respectively. We thank Stephen Smith and Saad Jbabdi for useful discussions about the time-domain connectivity methods. We would also like to thank all the reviewers for their useful comments and suggestions.

References

- Anzellotti, S., Fedorenko, E., Kell, A. J., Caramazza, A., and Saxe, R. (2017a). Measuring and modeling nonlinear interactions between brain regions with fMRI. *bioRxiv*, 074856.
- Anzellotti, S., Caramazza, A., and Saxe, R. (2017b). Multivariate pattern dependence. *PLoS computational biology*, 13(11), e1005799.
- Anzellotti, S., and Coutanche, M. N. (2018). Beyond functional connectivity: investigating networks of multivariate representations. *Trends in cognitive sciences*, 22(3), 258-269.
- Baccalá, L. A., and Sameshima, K. (2001). Partial directed coherence: a new concept in neural structure determination. *Biological cybernetics*, 84(6), 463-474.
- Barrett, A. B., Barnett, L., and Seth, A. K. (2010). Multivariate Granger causality and generalized variance. *Physical Review E*, 81(4), 041907.

- Basti, A., Pizzella, V., Chella, F., Romani, G. L., Nolte, G., and Marzetti, L. (2018). Disclosing large-scale directed functional connections in MEG with the multivariate phase slope index. *NeuroImage*, *175*, 161-175.
- Basti, A., Mur, M., Kriegeskorte, N., Pizzella, V., Marzetti, L., and Hauk, O. (2019). Analysing linear multivariate pattern transformations in neuroimaging data. *PloS one*, *14*(10).
- Biswal, B., Zerrin Yetkin, F., Haughton, V. M., and Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic resonance in medicine*, *34*(4), 537-541
- Bressler, S. L., and Menon, V. (2010). Large-scale brain networks in cognition: emerging methods and principles. *Trends in cognitive sciences*, *14*(6), 277-290.
- Bressler, S. L., and Seth, A. K. (2011). Wiener–Granger causality: a well established methodology. *Neuroimage*, *58*(2), 323-329.
- Brillinger, D. R. (1981). *Time series: data analysis and theory* (Vol. 36). Siam.
- Coutanche, M. N., and Thompson-Schill, S. L. (2013). Informational connectivity: identifying synchronized discriminability of multi-voxel patterns across the brain. *Frontiers in human neuroscience*, *7*, 15.
- Craddock, R. C., Tungaraza, R. L., and Milham, M. P. (2015). Connectomics and new approaches for analyzing human brain functional connectivity. *Gigascience*, *4*(1), s13742-015.
- Engel, A. K., Gerloff, C., Hlilgetag, C. C., and Nolte, G. (2013). Intrinsic coupling modes: multiscale interactions in ongoing brain activity. *Neuron*, *80*(4), 867-886.
- Ewald, A., Marzetti, L., Zappasodi, F., Meinecke, F. C., and Nolte, G. (2012). Estimating true brain connectivity from EEG/MEG data invariant to linear and static transformations in sensor space. *Neuroimage*, *60*(1), 476-488.
- Farahibozorg, S. R., Henson, R. N., & Hauk, O. (2018). Adaptive cortical parcellations for source reconstructed EEG/MEG connectomes. *NeuroImage*, *169*, 23-45.
- Friston, K. J., Rotshtein, P., Geng, J. J., Sterzer, P., and Henson, R. N. (2006). A critique of functional localisers. *Neuroimage*, *30*(4), 1077-1087.
- Geerligs, L., and Henson, R. N. (2016). Functional connectivity and structural covariance between regions of interest can be measured more accurately using multivariate distance correlation. *NeuroImage*, *135*, 16-31.
- Geerligs, L., Tsvetanov, K. A., and Henson, R. N. (2017). Challenges in measuring individual differences in functional connectivity using fMRI: the case of healthy aging. *Human brain mapping*, *38*(8), 4125-4156.
- Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., ... & Smith, S. M. (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, *536*(7615), 171-178.

- Goddard, E., Carlson, T. A., Dermody, N., and Woolgar, A. (2016). Representational dynamics of object recognition: Feedforward and feedback information flows. *Neuroimage*, *128*, 385-397.
- Golub, G. H., Heath, M., and Wahba, G. (1979). Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, *21*(2), 215-223.
- Harrison, L., Penny, W. D., and Friston, K. (2003). Multivariate autoregressive modeling of fMRI time series. *Neuroimage*, *19*(4), 1477-1491.
- Henriksson, L., Khaligh-Razavi, S. M., Kay, K., and Kriegeskorte, N. (2015). Visual representations are dominated by intrinsic fluctuations correlated between areas. *NeuroImage*, *114*, 275-286.
- Hipp, J. F., Hawellek, D. J., Corbetta, M., Siegel, M., and Engel, A. K. (2012). Large-scale cortical correlation structure of spontaneous oscillatory activity. *Nature neuroscience*, *15*(6), 884.
- Hotelling, H. (1936). Relations Between Two Sets of Variates. *Biometrika*, *28*(3/4), 321-377. doi:10.2307/2333955
- Kherif, F., Poline, J. B., Flandin, G., Benali, H., Simon, O., Dehaene, S., and Worsley, K. J. (2002). Multivariate model specification for fMRI data. *Neuroimage*, *16*(4), 1068-1083.
- Kietzmann, T. C., Spoerer, C. J., Sörensen, L. K., Cichy, R. M., Hauk, O., and Kriegeskorte, N. (2019). Recurrence is required to capture the representational dynamics of the human visual system. *Proceedings of the National Academy of Sciences*, *116*(43), 21854-21863.
- Kriegeskorte, N., Mur, M., and Bandettini, P. A. (2008). Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, *2*, 4.
- Lachaux, J. P., Rodriguez, E., Martinerie, J., and Varela, F. J. (1999). Measuring phase synchrony in brain signals. *Human brain mapping*, *8*(4), 194-208.
- Mahalanobis, P. C. (1936). On the generalized distance in statistics. National Institute of Science of India.
- Marzetti, L., Basti, A., Chella, F., D'Andrea, A., Syrjala, J., and Pizzella, V. (2019). Brain functional connectivity through phase-coupling of neuronal oscillations: a perspective from magnetoencephalography. *Frontiers in neuroscience*, *13*, 964.
- Nolte, G., Bai, O., Wheaton, L., Mari, Z., Vorbach, S., and Hallett, M. (2004). Identifying true brain interaction from EEG data using the imaginary part of coherency. *Clinical neurophysiology*, *115*(10), 2292-2307.
- Palva, S., and Palva, J. M. (2012). Discovering oscillatory interaction networks with M/EEG: challenges and breakthroughs. *Trends in cognitive sciences*, *16*(4), 219-230.
- Pascual-Marqui, R. D. (2007a). Coherence and phase synchronization: generalization to pairs of multivariate time series, and removal of zero-lag contributions. *arXiv preprint arXiv:0706.1776*.

- Pascual-Marqui, R. D. (2007b). Instantaneous and lagged measurements of linear and nonlinear dependence between groups of multivariate time series: frequency decomposition. *arXiv preprint arXiv:0711.1455*.
- Ranganath, C., & Ritchey, M. (2012). Two cortical systems for memory-guided behaviour. *Nature Reviews Neuroscience*, *13*(10), 713-726.
- Rangaprakash, D., Wu, G. R., Marinazzo, D., Hu, X., & Deshpande, G. (2018). Hemodynamic response function (HRF) variability confounds resting-state fMRI functional connectivity. *Magnetic resonance in medicine*, *80*(4), 1697-1713.
- Reid, A. T., Headley, D. B., Mill, R. D., Sanchez-Romero, R., Uddin, L. Q., Marinazzo, D., ... and Calhoun, V. (2019). Advancing functional connectivity research from association to causation. *Nature neuroscience*, *1*(10).
- Sanchez-Romero, R., and Cole, M. W. (2019). Combining multiple functional connectivity methods to improve causal inferences. bioRxiv.
- Siems, M., and Siegel, M. (2020). Dissociated neuronal phase-and amplitude-coupling patterns in the human brain. *NeuroImage*, 116538.
- Stam, C. J., Nolte, G., and Daffertshofer, A. (2007). Phase lag index: assessment of functional connectivity from multi channel EEG and MEG with diminished bias from common sources. *Human brain mapping*, *28*(11), 1178-1193.
- Stramaglia, S., Cortes, J. M., and Marinazzo, D. (2014). Synergy and redundancy in the Granger causal analysis of dynamical networks. *New Journal of Physics*, *16*(10), 105003.
- Székely, G. J., Rizzo, M. L., and Bakirov, N. K. (2007). Measuring and testing dependence by correlation of distances. *The annals of statistics*, *35*(6), 2769-2794.
- Taylor, J.R., Williams, N., Cusack, R., Auer, T., Shafto, M.A., Dixon, M., Tyler, L.K., Cam-CAN, Henson, R.N. (2017). The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) data repository: Structural and functional MRI, MEG, and cognitive data from a cross-sectional adult lifespan sample. *NeuroImage*. *144*, 262-269. doi: 10.1016/j.neuroimage.2015.09.018
- Theodoridis, S., and Koutroumbas, K. (1999, July). Pattern recognition and neural networks. In *Advanced Course on Artificial Intelligence* (pp. 169-195). Springer, Berlin, Heidelberg.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, *58*(1), 267-288.
- Uurtio, V., Monteiro, J. M., Kandola, J., Shawe-Taylor, J., Fernandez-Reyes, D., & Rousu, J. (2017). A tutorial on canonical correlation methods. *ACM Computing Surveys (CSUR)*, *50*(6), 1-33.
- Van Den Heuvel, M. P., and Pol, H. E. H. (2010). Exploring the brain network: a review on resting-state fMRI functional connectivity. *European neuropsychopharmacology*, *20*(8), 519-534.

Vidaurre, D., Hunt, L. T., Quinn, A. J., Hunt, B. A., Brookes, M. J., Nobre, A. C., & Woolrich, M. W. (2018). Spontaneous cortical activity transiently organises into frequency specific phase-coupling networks. *Nature communications*, *9*(1), 1-13.

Vinck, M., Oostenveld, R., Van Wingerden, M., Battaglia, F., and Pennartz, C. M. (2011). An improved index of phase-synchronization for electrophysiological data in the presence of volume-conduction, noise and sample-size bias. *Neuroimage*, *55*(4), 1548-1565.

Wang, H. T., Smallwood, J., Mourao-Miranda, J., Xia, C. H., Satterthwaite, T. D., Bassett, D. S., & Bzdok, D. (2020). Finding the needle in a high-dimensional haystack: Canonical correlation analysis for neuroscientists. *NeuroImage*, 116745.

Journal Pre-proof