

RESEARCH ARTICLE

Intersubject MVPD: Empirical comparison of fMRI denoising methods for connectivity analysis

Yichen Li^{1,2,3}, Rebecca Saxe³, Stefano Anzellotti^{4*}

1 Courant Institute of Mathematical Sciences, New York University, New York, NY, United States of America, **2** Department of Computer Science, New York University, New York, NY, United States of America, **3** Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA, United States of America, **4** Department of Psychology, Boston College, Chestnut Hill, MA, United States of America

* stefano.anzellotti@bc.edu



Abstract

Noise is a major challenge for the analysis of fMRI data in general and for connectivity analyses in particular. As researchers develop increasingly sophisticated tools to model statistical dependence between the fMRI signal in different brain regions, there is a risk that these models may increasingly capture artifactual relationships between regions, that are the result of noise. Thus, choosing optimal denoising methods is a crucial step to maximize the accuracy and reproducibility of connectivity models. Most comparisons between denoising methods require knowledge of the ground truth: of what is the ‘real signal’. For this reason, they are usually based on simulated fMRI data. However, simulated data may not match the statistical properties of real data, limiting the generalizability of the conclusions. In this article, we propose an approach to evaluate denoising methods using real (non-simulated) fMRI data. First, we introduce an intersubject version of multivariate pattern dependence (iMVPD) that computes the statistical dependence between a brain region in one participant, and another brain region in a different participant. iMVPD has the following advantages: 1) it is multivariate, 2) it trains and tests models on independent partitions of the real fMRI data, and 3) it generates predictions that are both between subjects and between regions. Since whole-brain sources of noise are more strongly correlated within subject than between subjects, we can use the difference between standard MVPD and iMVPD as a ‘discrepancy metric’ to evaluate denoising techniques (where more effective techniques should yield smaller differences). As predicted, the difference is the greatest in the absence of denoising methods. Furthermore, a combination of removal of the global signal and CompCorr optimizes denoising (among the set of denoising options tested).

OPEN ACCESS

Citation: Li Y, Saxe R, Anzellotti S (2019) Intersubject MVPD: Empirical comparison of fMRI denoising methods for connectivity analysis. PLoS ONE 14(9): e0222914. <https://doi.org/10.1371/journal.pone.0222914>

Editor: Pan Lin, Hunan Normal University, CHINA

Received: February 6, 2019

Accepted: September 10, 2019

Published: September 24, 2019

Copyright: © 2019 Li et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All data can be found on this website: <http://studyforrest.org/>.

Funding: Researchers were supported by a NIH grant R01 6933949 to Rebecca Saxe and a startup grant from Boston College to Stefano Anzellotti.

Competing interests: The authors have declared that no competing interests exist.

Introduction

Cognitive tasks elicit the activation of multiple brain regions [1, 2]. To understand how these regions function jointly to implement cognition, we need to investigate their connectivity.

Connectivity measures based on structural data (i.e. [3]) and connectivity measures based on functional data (i.e. [4]) have distinct advantages: the former enable inferences about whether regions are directly connected, the latter have the potential to investigate task-specific changes in the interactions between regions. Perhaps the most popular method to study interactions using functional data is ‘functional connectivity’, an analysis technique based on computing Pearson’s correlation between the average responses in two brain regions over time [4]. Functional connectivity has been used extensively to map brain networks [5] and search for biomarkers for patient populations [6].

Functional connectivity has served as the backbone of a vast literature, but the development of new methods for the study of neural responses and the emergence of new questions called for a new approach to study functional interactions between brain regions. Functional connectivity is univariate in nature: it is based on correlations between spatially-averaged responses. By contrast, it is now known that multivariate patterns of response encode rich information that is lost by averaging [7, 8]. Furthermore, functional connectivity does not lead to a predictive model that can be tested in independent data. Training a model with part of the data and testing its accuracy in independent data has become the norm for some types of analyses (i.e. MultiVoxel Pattern Analysis—MVPA, [9]), but it is not done in functional connectivity, and this makes it more susceptible to noise. A technique was needed to capture interactions between brain regions preserving multivariate information, and to implement training and testing in independent data. To satisfy these requirements, we have recently developed Multi-Variate Pattern Dependence (MVPD, [10]), and we have found that it is more sensitive than univariate connectivity techniques [10].

MVPD fits a model with a set of training data and tests it in independent data, therefore it is more robust to noise than methods lacking an independent ‘generalization’ test. However, some sources of noise may still affect MVPD results. For this reason, it is critical to use effective denoising techniques. Estimating the effectiveness of different denoising techniques for connectivity analysis has been typically challenging. Studies comparing different denoising approaches usually rely on simulations (i.e. [11]), because in simulations the ground truth is known. The simulation model determines what part of the overall response is the signal and what part is the noise, and the failure or success of the denoising methods can be determined unambiguously.

Unfortunately, applying conclusions obtained with simulation models to real fMRI data is non-trivial. Despite efforts to generate simulations that preserve many of the statistical properties of Blood-Oxygen Level Dependent (BOLD) signal, it is very difficult to generate simulations that perfectly match the statistical properties of real BOLD timecourses. This difficulty poses a major obstacle towards developing increasingly sophisticated connectivity techniques: more complex connectivity models might capture complex relationships between the noise in different regions, and thus require increasingly effective denoising. The differences between simulations and real data could thus have an even greater impact on the evaluation of denoising techniques.

In this article, we developed an approach to evaluate different denoising techniques using real fMRI data (as opposed to simulations). Our approach relies on computing interactions between brain regions between participants, following an ‘intersubject’ approach. Intersubject versions of univariate connectivity measures have been used in the previous literature to study individual differences in the responses to the same stimuli (intersubject correlations, [12–14]). We introduce here an intersubject version of MVPD (iMVPD), which predicts responses between brain regions in different participants (source code available at https://github.com/yl3506/iMVPD_denoise). Three key properties define iMVPD: 1) it is multivariate, 2) it trains and tests models on independent partitions of the data, and 3) it generates predictions that are both between subjects and between regions.

We took advantage of the unique properties of iMVPD to test the effectiveness of different denoising techniques on real fMRI data in the context of multivariate connectivity analyses. Considering a set of functionally-defined regions of interest, we computed connectivity matrices showing MVPD strength across each pair of regions both within subject and between subjects. Since different participants are likely to exhibit variation in their head movements and patterns of breathing or respiration, iMVPD should be less affected by these sources of noise. As a consequence, effective denoising techniques should yield a reduced difference between the within-subject and intersubject connectivity matrices. This gave us an index that we could use to compare different denoising approaches.

Materials and methods

Data

All analyses in this project use the 3 tesla (3T) audio-visual fMRI dataset from the StudyForrest project (<http://studyforrest.org>), which encompasses over 2 hours of scans for each of 15 participants (all right handed, mean age 29.4 years, range 21–39, 6 females). T2*-weighted echo-planar images (TR = 2s, TE = 30m, 90deg flip angle, parallel acquisition with sensitivity encoding (SENSE) reduction factor 2) were acquired using a whole-body Philips Achieva dStream MRI scanner with a 32 channel head coil. Each volume consists of 35 axial slices with a 10% inter-slice gap. Each slice comprises 80 × 80 voxels, with voxel resolution = 3 × 3 × 3mm, covering a Field of view (FoV) of 240mm (see [15] for additional details).

The dataset includes localizer sessions with stimuli from multiple categories, and the viewing of the entire movie Forrest Gump, subdivided into 8 sections presented to participants in 8 separate functional runs. This dataset was chosen for this study as it provides complex sensory input that follows the same timecourse between participants [16]. During the localizer session for this experiment, participants were shown 24 unique grayscale images from each of six stimulus categories: human faces, human bodies without heads, small objects, houses and outdoor scenes comprising of nature and street scenes, and phase scrambled images [17]. During the movie session for this experiment, participants were shown two-hour audio-visual stimuli (the movie Forrest Gump) [18].

Preprocessing

The format and folder structure of the dataset was modified to match the BIDS standard [19], and the data were preprocessed using fMRIPrep (<https://fmriprep.readthedocs.io/en/latest/index.html>): a preprocessing tool that takes advantage of nipy (<https://nipy.readthedocs.io/en/latest/>) to combine efficient algorithms for fMRI preprocessing from different software packages, minimizing experimenter degrees of freedom and offering a controlled environment that favors reproducibility [20]. Specifically, FSL MCFLIRT (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/MCFLIRT>) was used to estimate head motion, and each individual's functional data were coregistered to her/his anatomical scan. Segmentation and normalization were performed with ANTs (<http://stnava.github.io/ANTs/>).

Regions of interest (ROIs) definition

In order to make the results easily replicable, we analyzed two networks of category-selective brain regions that are highly reliable between participants and widely studied in the literature (Fig 1): the face-selective network (including the occipital face area—rOFA, fusiform face area—rFFA, anterior temporal lobe—rATL, and superior temporal sulcus—rSTS) and the scene-

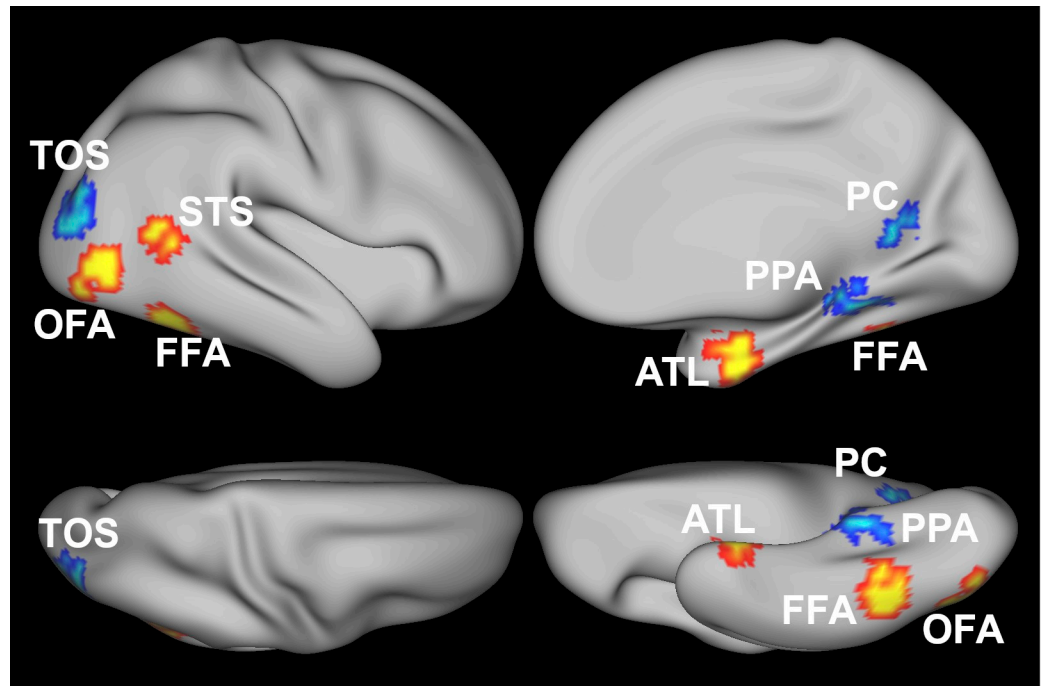


Fig 1. Regions of interest. Regions of interest for one example participant shown on an inflated cortical surface. Face-selective regions are shown in red-yellow (occipital face area—OFA, fusiform face area—FFA, superior temporal sulcus—STS, anterior temporal lobe—ATL), scene-selective regions are shown in blue (temporo-occipital sulcus—TOS, parahippocampal place area—PPA, posterior cingulate/retrosplenial cortex—PC).

<https://doi.org/10.1371/journal.pone.0222914.g001>

selective network (including transverse occipital sulcus—rTOS, parahippocampal place area—rPPA, and posterior cingulate—rPC).

Given the similarity in the responses in these regions across hemispheres, we restricted our analysis to the right hemisphere. First-level general linear models (GLM) were estimated in each participant with FSL FEAT (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FEAT>) using the independent localizer data. Face-selective regions were identified in individual participants using the contrast faces > scenes in an independent functional localizer. Anatomical parcels generated from a large number of independent subjects were used as a search space to identify contrast peaks for each of the regions. We then defined a 9mm sphere centered in the peak and selected the 80 voxels showing strongest face-selectivity within the sphere, as assessed with the contrast t-map. An analogous procedure was adopted for the scene regions, with the exception that the search space mask was defined by visually determining the group-level peaks for the contrast scenes > faces. Regions of interest were visualized on the cortical surface with Connectome Workbench (<https://www.humanconnectome.org/software/connectome-workbench>, Fig 1).

Denosing methods

We compared the effectiveness of four different denoising approaches for fMRI data: regression of slow trends [21], commonly adopted to remove ‘scanner drift’; low frequency fluctuations attributed to physiological noise and subject motion [22]; regression of the six motion parameters generated during motion correction [23], which attempts to remove noise that is linearly related to translations and rotations of the head; removal of the global signal [24], which discards the variability in a voxel’s responses that is shared with the fluctuation of the

average signal in the entire brain; and CompCorr [25], which extracts principal components from the signal in the white matter and cerebrospinal fluid and regresses them out from each voxel. All denoising methods were implemented by first generating one (i.e. in the case of global signal) or more (i.e. in the case of motion parameters) predictors. The predictors were then regressing out of each voxel's responses, and the residuals of the regression were used as the 'denoised' signal. Translation and rotation predictors and predictors for the removal of slow trends were obtained from the fMRIPrep outputs. Predictors for CompCorr were computed by generating an eroded mask of the white matter and cerebrospinal fluid using the segmented anatomical data from individual participants, and extracting 5 principal components. The predictor for the global signal was computed averaging the responses in all voxels in a subject-specific gray matter mask generated by fMRIPrep during segmentation. In addition to testing the effectiveness of denoising methods taken individually, we assessed combinations of multiple denoising methods to identify optimal noise removal approaches. In total, we compared 11 different denoising pipelines.

Data quality

The StudyForrest project (<http://studyforrest.org>) applied various approaches to ensure data quality [15]. They estimated participant movement by aligning all time series images to a common participant-specific brain template and obtained good results. Besides, the temporal signal-to-noise ratio was also measured, and was found to be reliably above the level of 40 within gray matter [17], which has been identified with simulations as the minimal level to reliably detect effects in fMRI [26]. In addition, two strategies were employed to demonstrate that the data captured rich and reliable signal: BOLD responses reflected the emotions presented in the video, and BOLD responses to the audio-visual movie correlated with previously collected responses to an audio-only version of the same movie. We analyzed data only from participants who completed all the sessions (4 runs of the object category localizer, and 8 runs of the movie, as well as anatomical data). One more subject failed to pass the FMRIprep preprocessing pipeline and therefore it was also removed. Based on within-subject analyses (in which we used a default denoising approach consisting of removal of scanner drift followed by CompCorr), we identified and removed three outlier subjects who had low within-subject prediction accuracy across all region pairs (see S1 Fig). After data quality control, the number of subjects left is 11.

iMVPD: Modeling representational spaces

For both within subject MVPD, and between subject iMVPD, the data were divided into a 8 partitions; each analysis used 7 partitions for training and one for testing, iteratively. Using only the training data, dimensionality reduction was applied to the multivariate responses in each region using PCA. Second, a linear function was estimated to predict the responses along the PCs in the target region from the responses in the predictor region (still using only the training data). Finally, the testing data were projected on the PC dimensions estimated with the training data, and the function f estimated with the training data was applied to the testing data in the predictor region to generate a prediction for the testing data in the target region. The proportion of variance explained by this prediction in the observed testing data in the target region was computed and used as a measure of multivariate statistical dependence between the two regions. The key difference between MVPD and between subjects iMVPD was that in iMVPD the target regions' data were drawn from a different participant.

More specifically, let's consider two participants A , B and two brain regions (a 'predictor' region in participant A , and a 'target' region in participant B). For participant A , we extracted

multivariate timecourses of response X_1^A, \dots, X_n^A in the predictor region, and for participant B we extracted multivariate timecourses Y_1^B, \dots, Y_n^B in the target region (where $n = 8$ is the number of runs).

The multivariate timecourse X_i^A in the predictor region for participant A in run i was a matrix of size $T_i \times N_X^A$, where T_i was the number of time points of the i^{th} run, and N_X^A was the number of voxels in the predictor region for participant A . Analogously, the multivariate timecourse Y_i^B in the target region for participant B in run i was a matrix of size $T_i \times N_Y^B$, where T_i was the number of time points of the i^{th} run, and N_Y^B was the number of voxels in the target ROI for participant B .

Just as in standard MVPD [10], for each choice of a testing run i , data in the remaining runs were concatenated, obtaining the training predictor:

$$X_{train}^A = (X_1^A, \dots, X_{i-1}^A, \dots, X_{i+1}^A, \dots, X_n^A) \tag{1}$$

The training target will be obtained by the same procedure, applied to the data from the predictor region from a different subject B :

$$Y_{train}^B = (Y_1^B, \dots, Y_{i-1}^B, \dots, Y_{i+1}^B, \dots, Y_n^B) \tag{2}$$

After the generation of the training set, principal component analysis (PCA) was applied to X_{train}^A and Y_{train}^B :

$$X_{train}^A = U_{X,train} S_{X,train} V_{X,train}^T \tag{3}$$

$$Y_{train}^B = U_{Y,train} S_{Y,train} V_{Y,train}^T \tag{4}$$

Dimensionality reduction was implemented projecting X_{train}^A and Y_{train}^B on lower dimensional subspaces spanned by the first k_X and k_Y principal components respectively:

$$\tilde{X}_{train}^A = X_{train}^A V_{X,train}^{[1, \dots, k_X]} \tag{5}$$

$$\tilde{Y}_{train}^B = Y_{train}^B V_{Y,train}^{[1, \dots, k_Y]} \tag{6}$$

where $V_{X,train}^{[1, \dots, k_X]}$ was the matrix formed by the first k_X columns of $V_{X,train}$, and $V_{Y,train}^{[1, \dots, k_Y]}$ was the matrix formed by the first k_Y columns of $V_{Y,train}$. For each region, each dimension obtained with PCA was a linear combination of the voxels in the region, whose weights defined a multivariate pattern of response over voxels. Considering as an example the predictor region, the scores of a dimension k were encoded in the k^{th} column of \tilde{X}_{train}^A , and represented the intensity with which the multivariate pattern corresponding to dimension k was activated over time. In previous work, we found that using 3 components outperformed using 1 or 2 components [10]. To test this observation in this new dataset, we compared the MVPD results obtained from choosing 1, 2, and 3 components using a default denoising method—removal of slow trends + CompCorr. Consistently with prior findings [10], the model with 3 components explained the most independent variance (see the section ‘Variance Explained’): averaging between all region pairs we obtained the following values of ‘multivariate independent correlation index’ (square root of the independent variance explained): 0.2169, 0.2560, 0.2774 respectively for 1, 2 and 3 principal components. Therefore, we used 3 components for the rest of our analyses.

iMVPD: Modeling statistical dependence and predicting multivariate timecourses

The mapping f from the dimensionality-reduced timecourses in the predictor region \tilde{X}_{train}^A to the dimensionality-reduced timecourses in the target region \tilde{Y}_{train}^B was modeled with multiple linear regression:

$$\tilde{Y}_{train}^B = W_{train} \tilde{X}_{train}^A + E_{train} \tag{7}$$

where the model parameters were estimated using ordinary least squares (OLS).

The extension of multivariate analysis methods to intersubject formulations can pose unique challenges. For example, in MVPA, classification between participants is usually poor unless functional alignment techniques are not first used to identify a common response space between participants [27, 28]. In the case of iMVPD, despite the brain region used as predictor and the brain region that is the target of prediction are in different participants, the learned prediction function is always applied to (independent subsets of) data from the same subject during both training and testing. This ensures that the function learned during training can be applied at the testing stage without the need of additional hyperalignment.

After the estimation of parameters W_{train} , predictions for the multivariate responses in the left out run i were computed by 1) projecting the predictor region’s data in the test run i on the principal components of the predictor region estimated with the other runs (the training runs), and 2) multiplying the dimensionality-reduced testing data by the parameters estimated using data from the other runs (the training runs). More formally, for each run i , we generated dimensionality reduced responses in the predictor region:

$$\tilde{X}_{test}^A = X_{test}^A V_{X,train}^{[1,\dots,k_X]} \tag{8}$$

Then, we calculated the predicted responses in the seed region in run i :

$$\hat{Y}_{test}^B = W_{train} \tilde{X}_{test}^A \tag{9}$$

using the parameters W_{train} independently estimated with the training runs.

Variance explained

To compute the proportion of variance explained in the test data we adopted the following procedure: 1) we applied principal component analysis to the test data, 2) we projected the observed and predicted responses on the principal components obtained from the test data, 3) we calculated the proportion of variance explained by the prediction in each component, and 4) we weighted these proportions by the proportion of variance each component explained in the overall observed response.

More specifically, we calculated

$$Y_{test}^B = U_{Y,test} S_{Y,test} V_{Y,test}^T \tag{10}$$

(step 1 above) and we projected the predicted and observed responses on this common orthogonal basis set (step 2):

$$\hat{Y}_{proj} = \hat{Y}_{test}^B V_{Y,train}^{[1,\dots,k_Y]T} V_{Y,test} \tag{11}$$

(where $\hat{Y}_{test}^B V_{Y,train}^{[1,\dots,k_Y]T}$ projects the predicted responses from the PCA space computed with the training data to voxel space, and the product by $V_{Y,test}$ projects them on the PCA space

computed with the testing data) and

$$Y_{proj} = Y_{test}^B V_{Y_{test}} \tag{12}$$

The proportion of variance explained in each component i was computed as

$$\text{varExpl}_i = 1 - \frac{\text{var}(Y_{proj}^i - \hat{Y}_{proj}^i)}{\text{var}(Y_{proj}^i)} \tag{13}$$

where Y_{proj}^i denotes the i -th column of matrix Y_{proj} , which contains the timecourse along component i of the PCA space computed with the testing data (step 3). Finally, the total variance explained was calculated as a weighted sum of the variance explained along each dimension (step 4), where the weights are given by the proportion of variance explained by that dimension in the overall response (which can be calculated using the PCA eigenvalues):

$$\text{varExpl} = \sum_{i=1}^{N_Y^B} \frac{\lambda_i}{\sum_{j=1}^{N_Y^B} \lambda_j} \text{varExpl}_i \tag{14}$$

For all analyses involving comparisons between within-subject MVPD and iMVPD, we did not compute predictions from a region to itself (i.e. values on the diagonal) because they are not meaningful for the within-subject analyses, and thus cannot be compared across methods.

Comparing denoising methods

To compare denoising models, we noted that an individual participant’s observed multivariate responses in a region $Y(t)$ can be decomposed as

$$Y(t) = Y_c(t) + Y_i(t) + \epsilon_c(t) + \epsilon_i(t) + \eta(t) \tag{15}$$

where the term $Y_c(t)$ is the part of the true response that is common across individuals, $Y_i(t)$ is the part of the true response that is specific to the particular individual, ϵ_c is the whole-brain noise that is shared across individuals and $\epsilon_i(t)$ is the whole-brain noise that is specific to that individual, and $\eta(t)$ is region-specific noise.

A model of connectivity within individual can aspire to explain an amount of variance given by

$$\text{varExpl}_{within} = \frac{\text{var}(Y_c(t) + Y_i(t) + \epsilon_c(t) + \epsilon_i(t))}{\text{var}(Y_c(t) + Y_i(t) + \epsilon_c(t) + \epsilon_i(t) + \eta(t))} \tag{16}$$

by contrast, an intersubject model can aspire to explain an amount of variance given by

$$\text{varExpl}_{between} = \frac{\text{var}(Y_c(t) + \epsilon_c(t))}{\text{var}(Y_c(t) + Y_i(t) + \epsilon_c(t) + \epsilon_i(t) + \eta(t))} \tag{17}$$

If a denoising method is removing in equal proportions signal and noise, (i.e. by a multiplicative factor $0 < \lambda < 1$), we obtain new proportions of variance explained:

$$\text{varExpl}_{within}^* = \frac{\text{var}(\lambda Y_c(t) + \lambda Y_i(t) + \lambda \epsilon_c(t) + \lambda \epsilon_i(t))}{\text{var}(\lambda Y_c(t) + \lambda Y_i(t) + \lambda \epsilon_c(t) + \lambda \epsilon_i(t) + \lambda \eta(t))} \tag{18}$$

$$= \frac{\lambda^2 \text{var}(Y_c(t) + Y_i(t) + \epsilon_c(t) + \epsilon_i(t))}{\lambda^2 \text{var}(Y_c(t) + Y_i(t) + \epsilon_c(t) + \epsilon_i(t) + \eta(t))} = \text{varExpl}_{within} \tag{19}$$

and similarly

$$\text{varExpl}_{between}^* = \frac{\text{var}(\lambda Y_c(t) + \lambda \epsilon_c(t))}{\text{var}(\lambda Y_c(t) + \lambda Y_i(t) + \lambda \epsilon_c(t) + \lambda \epsilon_i(t) + \lambda \eta(t))} \tag{20}$$

$$= \frac{\lambda^2 \text{var}(Y_c(t) + \epsilon_c(t))}{\lambda^2 \text{var}(Y_c(t) + Y_i(t) + \epsilon_c(t) + \epsilon_i(t) + \eta(t))} = \text{varExpl}_{between}. \tag{21}$$

By contrast, if a denoising method disproportionately reduces within-subject whole brain noise, we will have that the numerator of varExpl_{within} gets proportionally smaller than the numerator of $\text{varExpl}_{between}$. This is because the term that is disproportionately reduced by denoising is $\epsilon_i(t)$, which only appears in the numerator of varExpl_{within} . This effectively reduces the difference

$$\Delta \text{varExpl} = \text{varExpl}_{within}^* - \text{varExpl}_{between}^*. \tag{22}$$

Relying on this logic, we compared the alternative denoising methods using the index $\Delta \text{varExpl}$, where values closer to zero indicate better denoising performance. We computed the index $\Delta \text{varExpl}$ for each pair of brain regions, obtaining ‘connectivity-matrix-like’ figures in which each cell depicts the $\Delta \text{varExpl}$ index for the corresponding connection.

The average $\Delta \text{varExpl}$ index across all regions provides a measure of the efficacy of a denoising method (lower is better). Furthermore, we compared the pattern of denoising between regions for different denoising methods calculating the correlation between $\Delta \text{varExpl}$ matrices, obtaining a measure of similarity between denoising methods in terms of the set of connections on which they are more or less effective.

Results

Validating intersubject MVPD

In a first analysis, we tested whether intersubject MVPD is able to identify expected patterns of statistical dependence between different brain regions. Specifically, we tested whether intersubject MVPD, like within-subject MVPD, shows stronger interactions between regions that belong to the same network (i.e. face-selective regions vs scene-selective regions). We found that interactions among regions in the same network were indeed stronger also when using intersubject MVPD, a comparison of connectivity matrices is shown in Fig 2 (see sections ‘Denoising methods’ and ‘iMVPD: modeling representational spaces’). To measure the similarity between the patterns of interactions between regions across different region pairs for within-subject and intersubject connectivity we computed the Pearson’s correlation between the connectivity matrices generated with the two methods, obtaining an r value of 0.8596. In both within-subject MVPD and iMVPD, strong statistical dependence was found across pairs of region within a same network (defined based on response selectivity in an independent localizer): on average, pairs of face-selective regions showed stronger statistical dependence with other face-selective regions than with scene-selective regions, and viceversa (Fig 2). Relatively accurate predictions were also observed when face-selective regions were used as predictors for some scene-selective regions (in particular the rTOS), but not viceversa. This effect was observed with within-subject and also with intersubject MVPD.

To evaluate the effectiveness of denoising techniques, we did a matched t-test on the mean of the discrepancy metric before applying any denoising techniques and after applying denoising techniques. Since our work includes an intersubject analysis, the different pairs of subjects are not all independent, because two different subject pairs might have one subject in

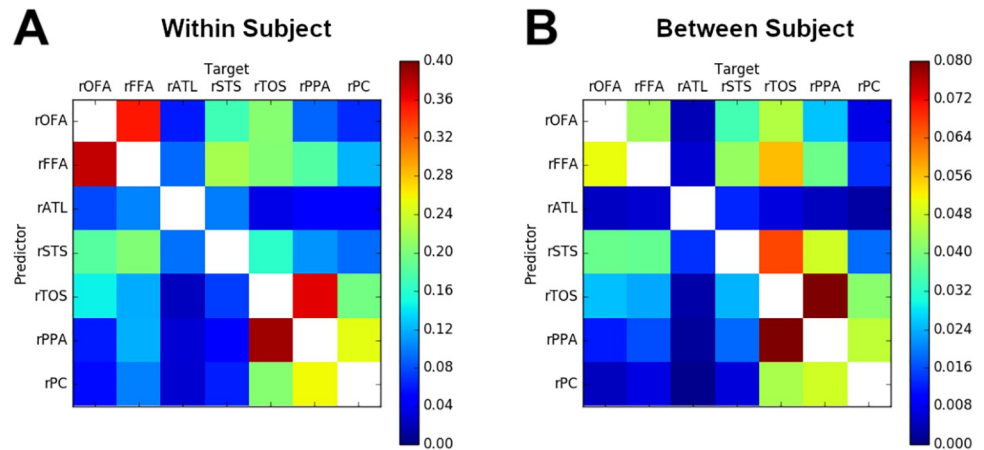


Fig 2. MVPD within- and between-subject. Amount of raw variance explained (as defined in Section ‘Variance explained’) for within-subject MVPD (A) and between-subject MVPD (B). Each element in a matrix represents the average between participants and cross-validation iterations of the proportion of the variance explained by predicting the multivariate responses in the target region from the the multivariate responses in the predictor region. The similarity between the two matrices was assessed with Pearson’s correlation, yielding $r = 0.8596$.

<https://doi.org/10.1371/journal.pone.0222914.g002>

common. Therefore, to estimate the variance and the mean while ensuring independence across pairs, we bootstrapped subsets of the subject pairs that are independent of each other. This bootstrapping approach leaves us with a very low number of degrees of freedom, therefore we only tested the difference between the most promising approach (removal of global signal and CompCorr) vs lack of denoising. Despite the low degrees of freedom, we still found significant differences in the mean of the discrepancy metric before and after denoising (one-tailed $t(4) = 3.07$, $p = 0.0186$), providing evidence removal of global signal and CompCorr significantly reduces the noise.

In addition, to evaluate the impact of the amount of data on the results, we computed the differences between the discrepancy metric obtained with no denoising, and the discrepancy metric obtained with removal of global signal and CompCorr using half of the data (4 runs) and one quarter of the data (2 runs). As expected, using less data yielded noisier estimates of the differences between the denoising methods (see S2 Fig). Furthermore, when using only two runs (one for training and one for testing MVPD), the differences in the discrepancy metric were lower (S2 Fig), suggesting that MVPD generated poorer predictions overall when using only one run for training.

Single denoising methods

After determining that iMVPD yielded similar patterns of statistical dependence across region pairs as within-subject MVPD, we proceeded to assess the efficacy of different denoising approaches (when considered individually) using the difference between the proportion of variance explained in the within-subject analysis minus the proportion of variance explained in the intersubject analysis ($\Delta \text{varExpl}$, see Section ‘Comparing denoising methods’).

As a sanity check, we calculated $\Delta \text{varExpl}$ without applying any denoising method, and indeed we observed a higher value of $\Delta \text{varExpl}$ than the ones observed after applying any of the other denoising methods (Fig 3A and 3F). Additionally, we observed differences in the $\Delta \text{varExpl}$ index across the denoising methods, with CompCorr achieving the best performance (as assessed by $\Delta \text{varExpl}$).

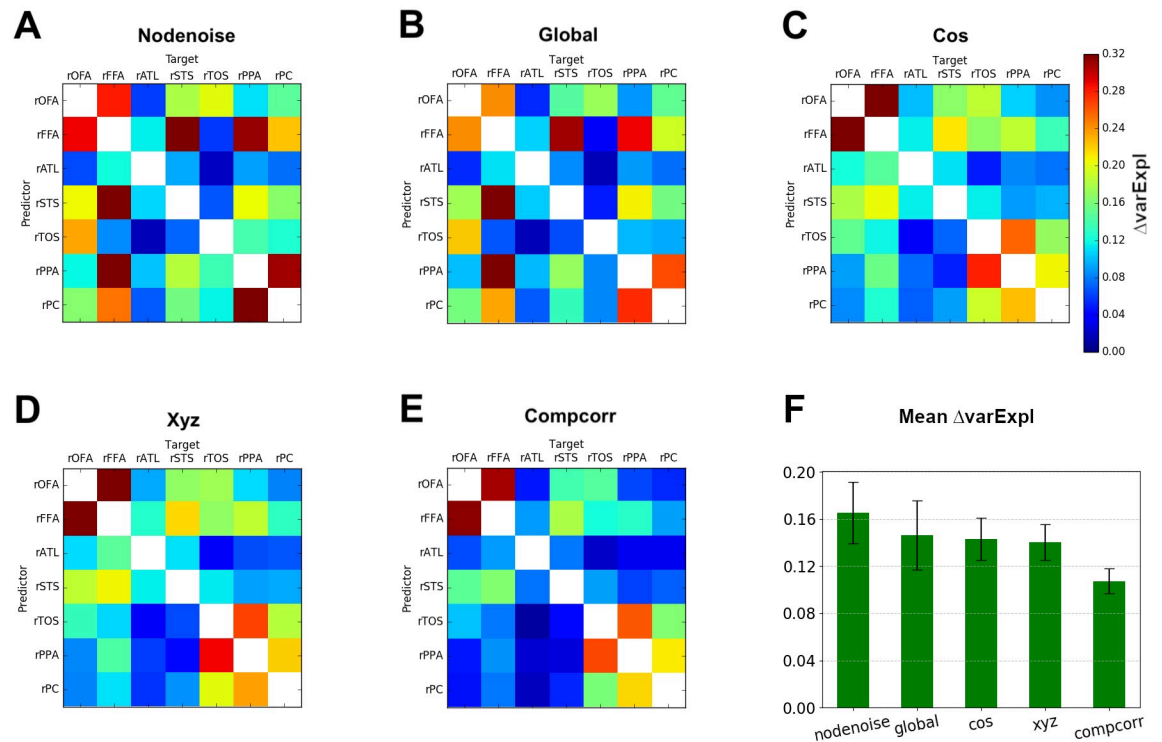


Fig 3. Difference in variance explained within and between participants for different denoising methods. Matrices showing the difference between the variance explained for within-subject MVPD minus the variance explained for iMVPD ($\Delta \text{varExpl}$) of A) data without any denoising applied, B) removal of global signal, C) removal of slow trends, D) regression of head translation and rotation parameters (motion parameters), E) CompCorr. Figure F) shows the mean $\Delta \text{varExpl}$ and the bootstrapped standard deviation of the previous five denoising methods in descending order.

<https://doi.org/10.1371/journal.pone.0222914.g003>

Different denoising methods varied not only in terms of the average value of $\Delta \text{varExpl}$, but also in the pattern of $\Delta \text{varExpl}$ across different region pairs. In order to perform a quantitative evaluation of the similarity between denoising methods in terms of their pattern of $\Delta \text{varExpl}$ across region pairs (that is, in terms of whether they removed noise from similar or distinct sets of connections), we calculated the Pearson’s correlation between the $\Delta \text{varExpl}$ matrices obtained for the different methods. This analysis revealed largely similar patterns for the removal of slow trends and the regression of translation and rotation parameters. The pattern of noise removal for CompCorr was also relatively similar, but the pattern for removal of global signal was markedly distinct, affecting different pairs of brain regions (Fig 4).

Comparing combinations of denoising methods

Multiple denoising methods can be used jointly to improve the efficacy of noise-removal. However, using too many denoising predictors may lead to the risk of removing meaningful variability in the signal. Therefore, we aimed to test different combinations of denoising methods to identify a minimal and effective noise-removal procedure which reduces $\Delta \text{varExpl}$ without including unnecessary predictors.

To this end, we compared several combinations of denoising methods. Since the removal of slow trends is widely used as a denoising method, we investigated the performance obtained by combining it with the other methods we assessed in the previous analyses (Fig 5A–5C). First, we observed that adding motion regressors (translation and rotation) did not improve denoising appreciably (compare Figs 5A to 3C). Note that this is not a trivial consequence of

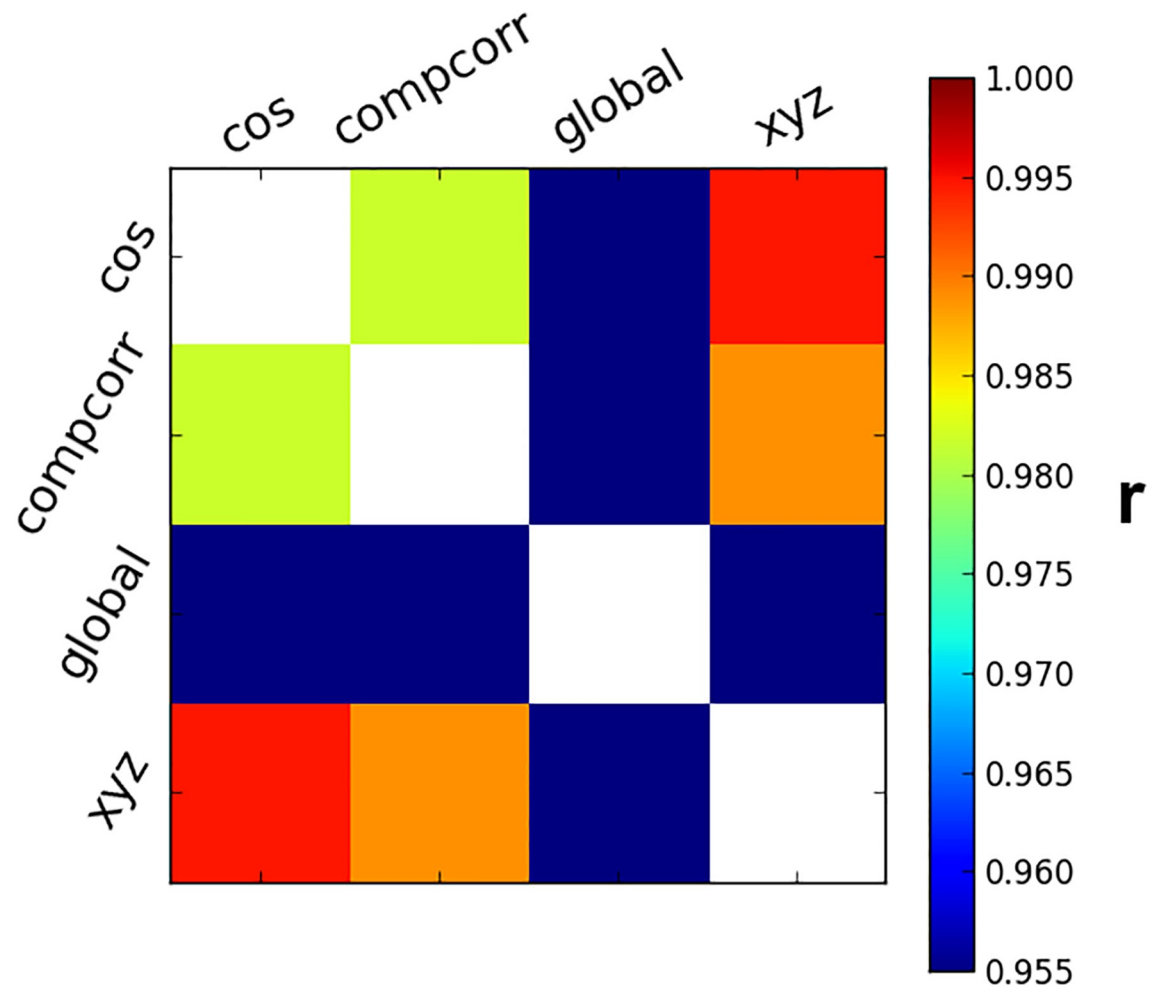


Fig 4. Correlation between denoising patterns. Pearson's correlation (r) of the Δ varExpl across all region pairs of the single denoising methods showed in Fig 3.

<https://doi.org/10.1371/journal.pone.0222914.g004>

the high correlation between the Δ varExpl matrices for these two methods, because distinct denoising methods might remove non-overlapping variance in similar amounts from the same set of brain regions, thus yielding high correlations between Δ varExpl matrices without redundancy in their denoising contribution. For example, two different denoising methods might remove two noise sources that are independent of each other in the timecourse of response within a region, but these independent noise sources might nonetheless have a similar distribution in terms of the proportion of noise they generate across region pairs.

Adding CompCorr to the removal of slow trends (Fig 5B) performed better than using the removal of slow trends in isolation (Fig 3C), however, it did not improve appreciably over using CompCorr alone (Fig 3E), indicating that the noise eliminated by the removal of slow trends is a subset of the noise captured by CompCorr.

A combination of removal of slow trends and removal of the global signal (Fig 5A) improved on both methods used in isolation, suggesting that they account for independent sources of noise.

Considering this pattern of results, we expected that using jointly CompCorr and the removal of the global signal would yield optimal results, even as compared to models

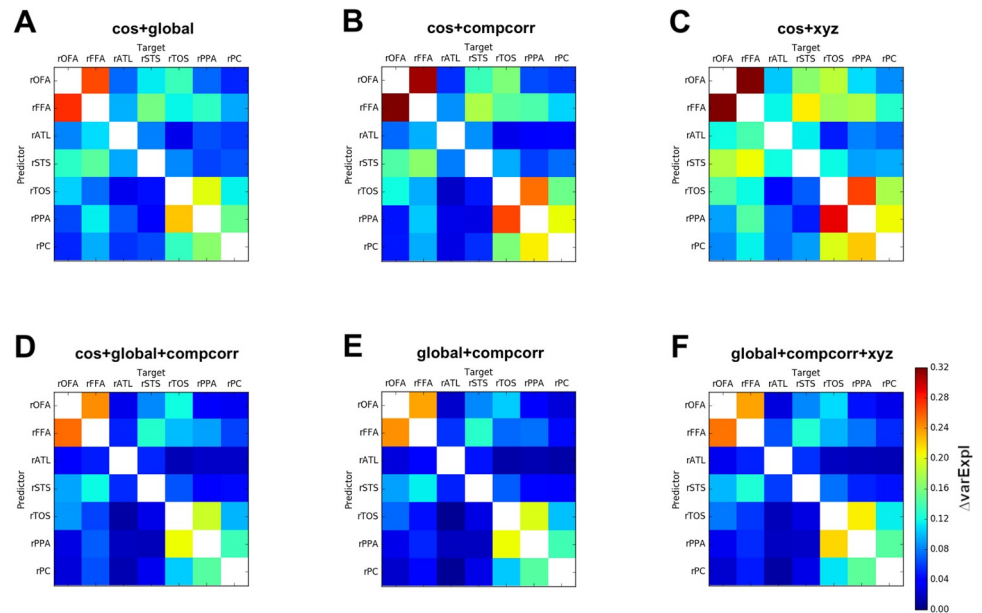


Fig 5. Difference in variance explained within minus between participants for combinations of denoising methods. Matrices showing the difference between the variance explained for within-subject MVPD ($\Delta \text{varExpl}$) of A) removal of slow trends followed by removal of global signal, B) removal of slow trends followed by CompCorr, C) removal of slow trends followed by regression of head translation and rotation parameters (motion parameters), D) removal of slow trends followed by removal of global signal followed by CompCorr, E) removal of global signal followed by CompCorr, F) removal of global signal followed by CompCorr followed by regression of head translation and rotation parameters (motion parameters).

<https://doi.org/10.1371/journal.pone.0222914.g005>

additionally including removal of slow trends or removal of motion regressors. To test this we computed $\Delta \text{varExpl}$ matrices for the combination of CompCorr and removal of global signal and found that the combination of these two methods performs better than either CompCorr alone (Fig 3E) or the removal of global signal alone (Fig 3B). Furthermore, by comparing Fig 5E with Fig 5D or Fig 5F, we found that adding the removal of slow trends or the removal of motion regressors to CompCorr plus the removal of global signal did not further improve $\Delta \text{varExpl}$. The combination of CompCorr and removal of global signal achieved parsimonious and effective noise removal.

Discussion

In this article we have introduced iMVPD, and have shown that it produces similar patterns of statistical dependence between brain regions as standard MVPD [10]. We have then introduced the difference between the proportion of variance explained within and between participants (hereafter referred to as the ‘discrepancy metric’) and motivated its use as a measure of the effectiveness of denoising methods.

Other methods for studying multivariate interactions exist (i.e. informational connectivity, [29, 30], see [31] for a review), and might have the potential to be extended to perform intersubject analyses. However, the unique combination of 1) attempting to capture most of the variance in each region’s responses, and 2) testing models of connectivity in independent data made MVPD an ideal choice for this study.

Multivariate variants of intersubject connectivity based on canonical correlation analysis have been introduced in previous work [32]. iMVPD differs substantially from these approaches in that 1) it tests prediction accuracy in independent data, and 2) enables modeling

of statistical dependence not only between-subjects within-region, but also between-subjects between-regions.

As predicted, we found that the discrepancy metric was greatest in the absence of denoising (see Fig 3A and 3F). Different denoising methods varied both in terms of their overall effectiveness (as measured by the average discrepancy metric across all pairs of connections), as well as in terms of the pattern of noise removal across different region pairs. This suggests that different pairs of regions may be differentially affected by distinct sources of noise, and that methods that are very effective at reducing the noise in the interactions between a pair of brain regions may be less effective at reducing noise for other pairs of regions.

In the present study, we used three principal components relying on previous observations indicating that components after the third have a minimal impact on the proportion of variance explained. However, developing a systematic approach for data-driven selection of the optimal number of components is an important direction for future work. Techniques such as Bayesian PCA [33] could be used for this purpose, offering principled approaches to adaptively select the number of components based on the available measurements.

Correlation between the discrepancy metric matrices (Fig 4) revealed that removal of the global signal disproportionately affected a different set of connections as compared to removal of slow trends, removal of motion regressor, and CompCorr. This finding suggests that removal of the global signal might target different sources of noise in fMRI data. Specifically, as compared to CompCorr, removal of the global signal disproportionately reduced the discrepancy metric for pairs of regions within a same network (i.e. across pairs of face-selective regions or across pairs of scene-selective regions).

Further studies could investigate the extent to which the different predictors generated by different denoising methods correlate with a variety of physiological measures (i.e. head movement, heart rate, respiration, eye movements) and susceptibility measures to determine whether the noise removed by the global signal has an independent origin.

One possible concern is that removal of the global signal might be subtracting meaningful variation in the fMRI responses. While it is difficult to rule out this possibility entirely, this in itself could not account for the reduction of the discrepancy metric within network: in fact, removal of meaningful signal would reduce both the absolute variance explained within subject as well as the total amount of variance within subject, leaving the proportion of variance explained unchanged (see Eq 21).

Testing different combinations of denoising methods revealed that the noise removed by the regression of slow trends and by the regression of motion parameters was mostly a subset of the noise removed by CompCorr. As a consequence, the use of regression of slow trends and regression of motion parameters in addition to CompCorr was redundant. In addition, CompCorr outperformed both of those methods. The effect of motion as measured by rotation and translation parameters on the observed BOLD signal is likely dependent on the anatomical location of voxels, and it is affected by complex phenomena like spin-history effects (in which a voxels move to some extent between slices and spins are no longer excited at regular intervals), and movement across non-uniform regions of the static magnetic field and of the radio-frequency (RF) fields [34]. These effects induce nonlinear relationships between motion parameters and the resulting noise, and may be better captured by CompCorr. Across all combinations of denoising methods tested, the discrepancy metric indicated that jointly removing the global signal and applying CompCorr was the optimal denoising pipeline.

Prior studies using simulated data have found that motion artifact causes variable duration of disruptions in signals [11]. [11] also found that proximal correlations are increased by artifact motions more than distal correlations, and that motion regressors have a limited effect in removing motion-related variance, even with voxel-specific regressors or when including a

large set of motion regressors. This finding is in line with our result in Fig 3D that motion parameters have limited efficacy in removing noise. [11] also found that mean white matter or mean ventricle signals are of modest utility as regressors, while fractionation of these signals via ANATICOR [35] or aCompCorr [36] or other methods may provide additional benefit. With iMVPD and the discrepancy metric we find that aCompCorr using multiple dimensions extracted from white matter and the cerebrospinal fluid (Fig 3E) provides considerable effectiveness in removing noise as compared with motion regressors only (Fig 3D). [11] pointed out that regressing the global signal is an effective processing step. In our findings (Fig 3B), removal of global signal is disproportionately effective in removing noise within-network (among face-selective region pairs or among scene-selective region pairs), and less effective in removing noise from between-network region pairs.

While these findings advance our understanding of denoising methods for fMRI, they are nonetheless affected by some limitations. First, optimal denoising methods are expected to minimize the discrepancy metric, however, it is worth noting that even perfect denoising would not reduce the discrepancy metric to zero. This is because individual differences between participants are expected to reduce iMVPD as compared to standard within-subject MVPD: the discrepancy metric is a result of both noise and individual differences. The discrepancy metric is a useful tool to compare the performance of different denoising methods in real fMRI data. However, it is important to note that while it provides a relative measure that can be used to compare different denoising methods, it is not an absolute measure of denoising. Therefore, the discrepancy metric cannot be used to determine exactly how much noise is left in the data. Second, our study was necessarily restricted to a subset of the vast set of denoising approaches that have been developed for fMRI. For example, we did not consider interpolation [37], ANATICOR [35], group-level covariates and partial correlations. Future studies could use the discrepancy metric to test additional denoising approaches that have not been included in the present investigation. Another effective strategy we did not test in this study consists in censoring motion-contaminated data. In the context of intersubject analysis, censoring would require removing all timepoints that show excessive motion in either of two subjects, further reducing the available data. Despite we did not test this approach in the present study, it can be applied in conjunction with the removal of global signal plus CompCorr approach we recommend. Finally, it will be important to extend the present results to other datasets employing different tasks and studying different brain regions to test the generality of the present findings.

In this study, we have used iMVPD as an instrument to define a discrepancy metric for denoising methods in order to assess their performance on ‘real’ (non-simulated) fMRI data. Among the methods tested, a combination of CompCorr and removal of the global signal was the most effective. Beyond the scope of this study, iMVPD could be used for a variety of applications. In ongoing work, we are using iMVPD to study individual differences across participants. Another potential area of application for iMVPD is the analysis of imaging modalities in which it is challenging to acquire data simultaneously from multiple regions, such as electrocorticography (ECoG) or primate electrophysiology. More generally, iMVPD can be applied as a conservative approach to study multivariate interactions between brain regions.

Supporting information

S1 Fig. Data quality control. The figure shows the within-subject prediction accuracy of all subjects who completed all the experimental sessions and passed the FM RIPREP preprocessing pipeline. Each dot shows the mean variance explained across all region pairs within a subject using a default denoising approach consisting of removal of scanner drift followed by

CompCorr. Three outlier subjects were identified due to their low within-subject prediction accuracy.

(TIF)

S2 Fig. Data length and experimental results. The figure shows the differences between the discrepancy metric obtained with no denoising, and the discrepancy metric obtained with removal of global signal and CompCorr using different amount of data. Each dot represents a subject pair. From left to right: yellow dots show the differences using one quarter of the data (2 runs), green dots show the differences using half of the data (4 runs), and blue dots show the differences using full data (8 runs).

(TIF)

Acknowledgments

We would like to thank the researchers who contributed to the studyforrest project [15, 17] for sharing their data, Jörn Diedrichsen for his comments on an earlier version of the manuscript, and the developers of fmripiprep [20] for their assistance with the fmripiprep preprocessing pipeline. Researchers were supported by a NIH grant R01 6933949 to Rebecca Saxe and a startup grant from Boston College to Stefano Anzellotti.

Author Contributions

Conceptualization: Yichen Li, Stefano Anzellotti.

Data curation: Yichen Li.

Formal analysis: Yichen Li.

Funding acquisition: Rebecca Saxe.

Investigation: Yichen Li, Stefano Anzellotti.

Methodology: Yichen Li, Stefano Anzellotti.

Resources: Rebecca Saxe.

Software: Yichen Li, Stefano Anzellotti.

Supervision: Rebecca Saxe, Stefano Anzellotti.

Validation: Stefano Anzellotti.

Visualization: Yichen Li, Stefano Anzellotti.

Writing – original draft: Yichen Li, Rebecca Saxe, Stefano Anzellotti.

Writing – review & editing: Yichen Li, Rebecca Saxe, Stefano Anzellotti.

References

1. Ishai A. Let's face it: it's a cortical network. *Neuroimage*. 2008; 40(2):415–419. <https://doi.org/10.1016/j.neuroimage.2007.10.040> PMID: 18063389
2. Anzellotti S, Caramazza A. From parts to identity: invariance and sensitivity of face representations to different face halves. *Cerebral Cortex*. 2015; 26(5):1900–1909. <https://doi.org/10.1093/cercor/bhu337> PMID: 25628344
3. Assaf Y, Pasternak O. Diffusion tensor imaging (DTI)-based white matter mapping in brain research: a review. *Journal of molecular neuroscience*. 2008; 34(1):51–61. <https://doi.org/10.1007/s12031-007-0029-0> PMID: 18157658

4. Biswal B, Zerrin Yetkin F, Haughton VM, Hyde JS. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic resonance in medicine*. 1995; 34(4):537–541. <https://doi.org/10.1002/mrm.1910340409> PMID: 8524021
5. Thomas Yeo B, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of neurophysiology*. 2011; 106(3):1125–1165. <https://doi.org/10.1152/jn.00338.2011>
6. Drysdale AT, Grosenick L, Downar J, Dunlop K, Mansouri F, Meng Y, et al. Resting-state connectivity biomarkers define neurophysiological subtypes of depression. *Nature medicine*. 2017; 23(1):28. <https://doi.org/10.1038/nm.4246> PMID: 27918562
7. Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*. 2001; 293(5539):2425–2430. <https://doi.org/10.1126/science.1063736> PMID: 11577229
8. Haynes JD, Rees G. Neuroimaging: decoding mental states from brain activity in humans. *Nature Reviews Neuroscience*. 2006; 7(7):523. <https://doi.org/10.1038/nrn1931> PMID: 16791142
9. Norman KA, Polyn SM, Detre GJ, Haxby JV. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends in cognitive sciences*. 2006; 10(9):424–430. <https://doi.org/10.1016/j.tics.2006.07.005> PMID: 16899397
10. Anzellotti S, Caramazza A, Saxe R. Multivariate pattern dependence. *PLoS computational biology*. 2017; 13(11):e1005799. <https://doi.org/10.1371/journal.pcbi.1005799> PMID: 29155809
11. Power JD, Schlaggar BL, Petersen SE. Recent progress and outstanding issues in motion correction in resting state fMRI. *Neuroimage*. 2015; 105:536–551. <https://doi.org/10.1016/j.neuroimage.2014.10.044> PMID: 25462692
12. Hasson U, Nir Y, Levy I, Fuhrmann G, Malach R. Intersubject synchronization of cortical activity during natural vision. *science*. 2004; 303(5664):1634–1640. <https://doi.org/10.1126/science.1089506> PMID: 15016991
13. Hasson U, Furman O, Clark D, Dudai Y, Davachi L. Enhanced intersubject correlations during movie viewing correlate with successful episodic encoding. *Neuron*. 2008; 57(3):452–462. <https://doi.org/10.1016/j.neuron.2007.12.009> PMID: 18255037
14. Wilson SM, Molnar-Szakacs I, Iacoboni M. Beyond superior temporal cortex: intersubject correlations in narrative speech comprehension. *Cerebral cortex*. 2007; 18(1):230–242. <https://doi.org/10.1093/cercor/bhm049> PMID: 17504783
15. Hanke M, Adelhöfer N, Kottke D, Iacovella V, Sengupta A, Kaule FR, et al. A studyforrest extension, simultaneous fMRI and eye gaze recordings during prolonged natural stimulation. *Scientific data*. 2016; 3:160092. <https://doi.org/10.1038/sdata.2016.92> PMID: 27779621
16. Labs A, Reich T, Schulenburg H, Boennen M, Mareike G, Golz M, et al. Portrayed emotions in the movie “Forrest Gump”. *F1000Research*. 2015; 4. <https://doi.org/10.12688/f1000research.6230.1> PMID: 25977755
17. Sengupta A, Kaule FR, Guntupalli JS, Hoffmann MB, Häusler C, Stadler J, et al. A studyforrest extension, retinotopic mapping and localization of higher visual areas. *Scientific data*. 2016; 3:160093. <https://doi.org/10.1038/sdata.2016.93> PMID: 27779618
18. Joshi AA, Chong M, Li J, Choi S, Leahy RM. Are you thinking what I’m thinking? Synchronization of resting fMRI time-series across subjects. *NeuroImage*. 2018; 172:740–752. <https://doi.org/10.1016/j.neuroimage.2018.01.058> PMID: 29428580
19. Gorgolewski KJ, Auer T, Calhoun VD, Craddock RC, Das S, Duff EP, et al. The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. *Scientific Data*. 2016; 3:160044. <https://doi.org/10.1038/sdata.2016.44> PMID: 27326542
20. Esteban O, Markiewicz C, Blair RW, Moodie C, Isik AI, Aliaga AE, et al. FMRIPrep: a robust preprocessing pipeline for functional MRI. *bioRxiv*. 2018; p. 306951.
21. Friston K, Josephs O, Zarahn E, Holmes A, Rouquette S, Poline JB. To smooth or not to smooth?: Bias and efficiency in fmri time-series analysis. *NeuroImage*. 2000; 12(2):196–208. <https://doi.org/10.1006/nimg.2000.0609> PMID: 10913325
22. Smith AM, Lewis BK, Ruttimann UE, Frank QY, Sinnwell TM, Yang Y, et al. Investigation of low frequency drift in fMRI signal. *NeuroImage*. 1999; 9(5):526–533. <https://doi.org/10.1006/nimg.1999.0435> PMID: 10329292
23. Friston KJ, Williams S, Howard R, Frackowiak RS, Turner R. Movement-related effects in fMRI time-series. *Magnetic resonance in medicine*. 1996; 35(3):346–355. <https://doi.org/10.1002/mrm.1910350312> PMID: 8699946
24. Macey PM, Macey KE, Kumar R, Harper RM. A method for removal of global effects from fMRI time series. *NeuroImage*. 2004; 22(1):360–366. <https://doi.org/10.1016/j.neuroimage.2003.12.042> PMID: 15110027

25. Behzadi Y, Restom K, Liu J, Liu TT. A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *Neuroimage*. 2007; 37(1):90–101. <https://doi.org/10.1016/j.neuroimage.2007.04.042> PMID: 17560126
26. Murphy K, Bodurka J, Bandettini PA. How long to scan? The relationship between fMRI temporal signal to noise ratio and necessary scan duration. *Neuroimage*. 2007; 34(2):565–574. <https://doi.org/10.1016/j.neuroimage.2006.09.032> PMID: 17126038
27. Guntupalli JS, Haxby JV. Inter-subject hyperalignment of neural representational space. *Chance*. 2010; 1:7.
28. Haxby JV, Guntupalli JS, Connolly AC, Halchenko YO, Conroy BR, Gobbini MI, et al. A common, high-dimensional model of the representational space in human ventral temporal cortex. *Neuron*. 2011; 72(2):404–416. <https://doi.org/10.1016/j.neuron.2011.08.026> PMID: 22017997
29. Coutanche MN, Thompson-Schill SL. Informational connectivity: identifying synchronized discriminability of multi-voxel patterns across the brain. *Frontiers in human neuroscience*. 2013; 7:15. <https://doi.org/10.3389/fnhum.2013.00015> PMID: 23403700
30. Crowe DA, Goodwin SJ, Blackman RK, Sakellaridi S, Sponheim SR, MacDonald AW III, et al. Prefrontal neurons transmit signals to parietal neurons that reflect executive control of cognition. *Nature Neuroscience*. 2013; 16(10):1484. <https://doi.org/10.1038/nn.3509> PMID: 23995071
31. Anzellotti S, Coutanche MN. Beyond Functional Connectivity: Investigating Networks of Multivariate Representations. *Trends in cognitive sciences*. 2018. <https://doi.org/10.1016/j.tics.2017.12.002> PMID: 29305206
32. Kriegeskorte N. Intersubject information mapping: revealing canonical representations of complex natural stimuli. *bioRxiv*. 2015; p. 016436.
33. Bishop CM. Bayesian pca. In: *Advances in neural information processing systems*; 1999. p. 382–388.
34. Liu TT. Noise contributions to the fMRI signal: An overview. *NeuroImage*. 2016; 143:141–151. <https://doi.org/10.1016/j.neuroimage.2016.09.008> PMID: 27612646
35. Jo HJ, Saad ZS, Simmons WK, Milbury LA, Cox RW. Mapping sources of correlation in resting state FMRI, with artifact detection and removal. *Neuroimage*. 2010; 52(2):571–582. <https://doi.org/10.1016/j.neuroimage.2010.04.246> PMID: 20420926
36. Muschelli J, Nebel MB, Caffo BS, Barber AD, Pekar JJ, Mostofsky SH. Reduction of motion-related artifacts in resting state fMRI using aCompCor. *Neuroimage*. 2014; 96:22–35. <https://doi.org/10.1016/j.neuroimage.2014.03.028> PMID: 24657780
37. Power JD, Mitra A, Laumann TO, Snyder AZ, Schlaggar BL, Petersen SE. Methods to detect, characterize, and remove motion artifact in resting state fMRI. *Neuroimage*. 2014; 84:320–341. <https://doi.org/10.1016/j.neuroimage.2013.08.048> PMID: 23994314