
CONTROLLING FOR SPURIOUS NONLINEAR DEPENDENCE IN CONNECTIVITY ANALYSES

A PREPRINT

Craig Poskanzer^{1,*}, Mengting Fang¹, Aidas Aglinskas¹, and Stefano Anzellotti¹

¹Department of Psychology and Neuroscience, Boston College, Boston, MA 02467

*Corresponding Author: poskanzc@bc.edu

December 4, 2020

Keywords: Nonlinear, connectivity, denoising, CompCor

ABSTRACT

Recent analysis methods can capture nonlinear interactions between brain regions. However, noise sources might induce spurious nonlinear relationships between the responses in different regions. Previous research has demonstrated that traditional denoising techniques effectively remove noise-induced *linear* relationships between brain areas, but it is unknown whether these techniques can remove spurious nonlinear relationships. Among existing denoising methods, CompCor has been hypothesized to remove noise in BOLD responses that is nonlinearly related to its source. In this paper, we investigated whether CompCor additionally removes spurious nonlinear interactions between different brain regions. To test this, we analyzed fMRI responses while participants watched the film *Forrest Gump* using both linear and nonlinear Multivariate Pattern Dependence Networks (MVPN). We found nonlinear interactions between the nondenoised responses in face-selective regions and nondenoised responses in the anterior frontal and temporal lobes. CompCor denoising removed these nonlinear interactions. We then asked whether information contributing to the removal of nonlinear interactions is localized to specific anatomical regions of no interest or to specific principal components. We denoised the data 8 separate times by regressing out 5 principal components extracted from combined white matter (WM) and cerebrospinal fluid (CSF), each of the 5 components separately, 5 components extracted from WM only, and 5 components extracted solely from CSF. In all cases, denoising was sufficient to remove the observed nonlinear interactions. Finally, we replicated our results using different types of neural networks as the bases of MVPN, demonstrating that CompCor’s ability to remove nonlinear interactions is independent of network architecture.

Impact Statement: With the growing popularity of nonlinear connectivity analyses, this research provides key insight into the ability of common denoising approaches to reduce spurious nonlinear interactions between brain regions. This work can help guide the selection of data analysis procedures for connectivity studies across multiple domains of cognition.

1 Introduction

Cognitive tasks recruit multiple brain regions, and studying how these regions interact is key to understanding the neural bases of behavior. In recent years, novel techniques have been introduced that take advantage of multivariate information in fMRI response patterns to reveal statistical dependence between the responses in different regions [[1]]. Using multivariate information can improve sensitivity and identify interactions between regions that could not have been detected with univariate methods [[2, 3, 4]].

Interactions between brain regions are likely not only multivariate, but also nonlinear in nature [[5, 6, 7]]. At an individual level, responses of single neurons are nonlinear [[8, 9]]; on a larger scale, evidence suggests these nonlinear transformations are reflected at the population level [[10, 5]], which undercuts the argument that the nonlinear responses

of individual neurons could be used to approximate population level linear transformations. Furthermore, computational models of behavior benefit from incorporating nonlinear transformations. For example, nonlinearities are key for deep neural networks that achieve good performance in object recognition tasks and yield accurate models of neural responses in both humans [[11]] and macaques [[12]], suggesting that understanding nonlinear interactions between responses in different brain regions is key to understanding the neural bases of cognition.

Novel methods leveraging differential equations, information theory and machine learning techniques have been introduced to capture nonlinear interactions between the responses in different brain regions [[13, 14, 5]]. The ability to capture nonlinear interactions, however, comes with the risk of inflated susceptibility to noise in the BOLD signal [[15, 16]]. FMRI noise is particularly problematic for studies of functional connectivity. Head movements, heart rate and respiration can lead to global BOLD fluctuations [[16, 17, 18]], which in turn can yield inflated estimates of statistical dependence (‘connectivity’) [[19]].

A variety of data driven approaches have been proposed to remove non-neural fluctuations in the BOLD signal, including regressing out fluctuations in the global signal [[20]], band-pass filtering of slow trends [[21]], extracting and regressing out motion parameters [[17]], and regressing out principal components extracted from regions of no interest [[22]]. Notably, most of these methods rely on linear regression. This raises the question of whether, and to what extent, existing denoising methods might be able to remove noise-induced nonlinear statistical dependence. Addressing this question is essential in order to interpret the results obtained with novel analysis techniques that can capture nonlinear dependence between brain regions.

The performance of denoising methods has been thoroughly evaluated in previous studies [[23, 24, 25, 26, 27, 28]]. Their effectiveness has been assessed largely on the basis of their impact on measures like ‘functional connectivity’, which are linear in nature and therefore would not be susceptible to noise-induced nonlinear dependence. Nonlinear models of connectivity are more sensitive than their linear counterparts, and, thus, able to identify noise-induced nonlinear interactions that the linear models would not be susceptible to. As a consequence, it remains unclear whether existing denoising methods based on linear regression are able to remove spurious nonlinear dependence between brain regions introduced by noise.

Among existing denoising methods, CompCor [[22]] offers unique advantages. CompCor applies principal component analysis (PCA) to data from white matter (WM) and cerebrospinal fluid (CSF), and uses the timecourses of the principal components as regressors of no interest [[22]]. In addition to capturing cardiac and respiratory artifacts, CompCor may exhibit greater sensitivity in modeling motion based fluctuations in the BOLD signal than using motion parameters to develop nuisance regressors [[29]]. Since CompCor derives regressors from measured BOLD responses in regions of no interest (instead of using motion parameters), it has the potential to remove noise which is nonlinearly related to the motion parameters themselves [[30]; [29]]. In fact, if a given set of motion parameters θ produce whole-brain noise $f(\theta)$, where f is nonlinear, extracting signals from WM and CSF can offer a direct estimate of $f(\theta)$ – which could not be obtained by applying a linear transformation to the motion parameters θ . In a 2019 study comparing the effectiveness

of different denoising techniques with intersubject linear multivariate pattern dependence (iMVPD), Li et al. found that CompCor was the most efficient among the methods tested at removing artifactual linear dependence [[28]]. Moreover, using CompCor has the additional benefit of maintaining the global signal, which can be particularly useful in the study of clinical populations [[31]]. As the increasingly popular studies of network dynamics among clinical populations [[32]] broaden their use of nonlinear models of connectivity, it is essential to make use of a denoising method that both maintains discriminative information while also removing spurious nonlinear interactions.

Taken together, these considerations single out CompCor as an especially promising candidate to remove noise-induced nonlinear statistical dependence between brain regions. However, CompCor is based on linear regression. While it is able to effectively remove noise-based *linear* interactions between a brain region and another brain region, and *nonlinear* interactions between a noise source and a brain region, it is not yet known whether it is also able to capture *nonlinear* interactions between two brain regions. Establishing this is essential to determine whether CompCor is a viable method to mitigate noise for nonlinear analyses of statistical dependence.

In this study, we used nonlinear multivariate pattern dependence based on artificial neural networks (nonlinear MVPN) to test the ability of CompCor to remove nonlinear statistical dependence between brain regions. We found that CompCor led to a substantial reduction in nonlinear dependence between functionally defined face-selective regions and regions in frontal and anterior temporal cortex which are most affected by head motion [[33]], indicating that CompCor can contribute to the removal of spurious nonlinear dependence.

2 Methods

This project consists entirely of analyses of a publicly available dataset (<http://studyforrest.org>), therefore it did not require IRB approval.

2.1 Data

The data in this study were collected as part of the *studyforrest* project [[34]], a publicly available dataset. No new data were collected for this manuscript. Fifteen volunteers (6 female, ages 21-39, mean = 29.4) took part in the study. Informed consent procedures are described in [34]. One participant was excluded from the analysis because the data failed to pass the fMRIPrep preprocessing pipeline [[35]]. The data were acquired with a 3 Tesla Philips Achieva dStream MRI scanner using a 32-channel head coil. BOLD fMRI responses were recorded at high resolution (3 x 3 x 3mm) with T2*-weighted echo-planar (2 sec repetition time (TR)) imaging sequence (see [34] for a detailed description of the protocol).

A total of 8 functional runs (lasting approximately 15 minutes each) were acquired while participants watched the film *Forrest Gump*. In addition, participants also completed 4 runs of a functional localizer, in which they were shown 24 grayscale images from each of six categories: faces, bodies (without heads), small objects, houses, outdoor scenes of nature and streets, and phase scrambled images [[36]]. During the functional localizer, participants performed a

one-back task in which they were instructed to press a button if the current figure was a mirror image of the previous image on display. During each of the 4 experimental blocks, participants saw 2 sets of 16 unique and randomly ordered images from each stimulus category (faces, bodies, objects, houses, scenes, and scrambled images). Images were presented for 900ms each, interleaved with 100ms blank background displays [[36]]. In this study, data from the first localizer run was used to identify category-selective areas responding preferentially to faces for each individual subject (contrast: faces > bodies, objects, houses, scenes, and scrambled images).

2.2 Preprocessing

All fMRI data were preprocessed with the fMRIPrep preprocessing pipeline [[35]]. First, the T1-weighted (T1w) anatomical images were smoothed and skull-stripped using ANTs before using FAST (FSL) [[37]] to segment the brain white matter (WM), gray matter (GM), and cerebrospinal fluid (CSF). Next, functional scans were corrected for head-movement using MCFLIRT (FSL) [[38]] before being aligned with the corresponding T1w reference image using boundary-based co-registration implemented in FLIRT (FSL). Data from one subject was excluded after it failed to pass this processing pipeline (see also [28]).

2.3 Denoising

To investigate how denoising affects estimates of the linear and nonlinear interactions between brain regions, we applied 8 distinct denoising procedures to the data, and compared them to a control analysis in which no denoising was applied. All denoising procedures were based on CompCor [[22]], a method based on regressing out principal components extracted from regions of no interest from each voxel's responses.

In a first analysis, we extracted 5 principal components (PCs) from a combination of the BOLD signal generated by WM and CSF. Next, we tested the relative contributions of CSF and WM to denoising, evaluating the results after regressing out 5 PCs generated from the WM alone, and after regressing out 5 PCs generated from the CSF alone. Finally, we assessed the contribution of individual PCs by extracting 5 components from the combined WM and CSF data and performing 5 distinct denoising procedures, regressing out a single PC at a time.

Throughout the remainder of the article, we will refer to these different denoising pipelines as follows: nondenoised, full denoised (using 5 PCs extracted from a combination of WM and CSF), WM denoised (using 5 PCs extracted from the WM), CSF denoised (using 5 PCs extracted from the CSF), 1st PC denoised, 2nd PC denoised, 3rd PC denoised, 4th PC denoised, and 5th PC denoised (each denoised with a consecutive PC generated from a combination of WM and CSF data).

2.4 Regions of Interest

Using the first run of the block-designed initial localizer task, we estimated a standard GLM using FSL FEAT [[39]] for each subject in order to identify regions that maximized the contrast between responses to face and non-face images.

The model included a predictor for each object category. Predictors were convolved with a gamma hemodynamic response function (HRF) as implemented in FSL.

We identified three ROIs (the fusiform face area (FFA), occipital face area (OFA), and the face-selective superior temporal sulcus (STS)) for each participant by creating a 9mm sphere around the voxel exhibiting peak t-value for contrast between faces and non-faces. Each ROI was defined to be 80 voxels within that sphere that showed the strongest selectivity for face images. Finally, we created a union of the three ROIs for each subject to create a 240 voxel mask of face-selective regions [[40]].

In addition to the individualized masks of face-selective regions, we also created a group-level grey matter mask by averaging across the grey matter probability maps of each participant (with a threshold of $p > .1$) created during segmentation.

2.5 Multivariate Pattern Dependence

In order to evaluate the effects of denoising on linear and nonlinear models of the interactions between brain regions, we analyzed the data using linear and nonlinear multivariate pattern dependence [MVPD, [3, 28]]. MVPD is a sensitive technique that captures the multivariate statistical dependence between response patterns in different brain regions, and has been shown to be a meaningful measure of functional interactions [[3]]. Specifically, in this article we used an extension of MVPD based on artificial neural networks [MVPN, [40]], which allowed us to directly compare linear and nonlinear models with the same architecture. Artificial neural networks can implement dimensionality reduction [[41]], eliminating the need for PCA.

In MVPN [[40]], fMRI data for each participant are first divided into a training set and a testing set. A brain region is chosen as the predictor, and another region as the target of prediction. The training data is used to train an artificial neural network to predict the response patterns in the target region using as input the response patterns in the predictor region. At this point, the artificial neural network’s weights are frozen, and its accuracy is evaluated by calculating the proportion of variance explained in the testing data.

In this study, 7 of the 8 runs were used for training, and the remaining run was used for testing. This leave-one-run-out procedure was iterated 8 times, one for each possible choice of the left-out run (8-fold cross validation). We used the union of all three face-selective regions as the predictor, and the gray matter ROI as the target of prediction. This allows us to see which areas of the brain are statistically dependent on activity in face-selective regions.

Multivariate pattern dependence was estimated using artificial neural networks with three distinct architectures, each with a linear and nonlinear version. The first architecture was a fully connected network with one hidden layer consisting of 100 hidden nodes. This architecture was selected for its simplicity, and previous work has shown that it performs well at predicting global gray matter activity from BOLD responses in face-selective ROIs [[40]]. The nonlinear version of the network used the same architecture as the linear version, with the addition of rectified linear units (ReLU) as activation functions. In order to evaluate the robustness of results across different network architectures, we additionally

tested a ‘standard’ feedforward network with five layers (with connections only between successive layers), and a five layer DenseNet [[42]]. For both networks, each layer consisted of 100 hidden nodes [see also [40]].

All neural networks were trained using stochastic gradient descent on a mean squared error (MSE) loss function with a learning rate of 0.001 and a momentum of 0.9. The networks were trained over the course of 5000 epochs using a batch size of 32, with batch normalization applied to the input of each layer. The MVPN code can be found at <https://github.com/pandamt/MVPN>.

Multivariate statistical dependence was calculated as the proportion of variance explained by the artificial neural network in the test data. For each left-out run i and for each voxel (j), variance explained was calculated as:

$$\text{varExpl}_i(j) = \max \left\{ 0, 1 - \frac{\text{var}(Y_i(j) - f_j(X_i))}{\text{var}(Y_i(j))} \right\}, \quad (1)$$

where $Y_i(j)$ is the response in voxel j of the target region during the test run i , X_i is the response pattern in the predictor region during the test run i , and $f_j(X_i)$ is the voxel-wise prediction of response $Y_i(j)$ generated by MVPN.

2.6 Statistical Analysis

For each voxel, we calculated the difference between the prediction accuracy with the nonlinear model minus the accuracy with the linear model. Statistical non-parametric mapping (SnPM) was used to assess the significance of the difference [[43]], as implemented in the SnPM toolbox in MATLAB [[43]; [44]]. SnPM calculates significance using a permutation test. In all voxels, the distribution of the null hypothesis is estimated randomly swapping the accuracies between conditions (linear and nonlinear) for each subject. This procedure is repeated for 5000 iterations. At each permutation, a t-statistic is computed for the comparison between the nonlinear and linear accuracies. The p-value is obtained as the proportion of values of the t-statistic in the permutation iterations that are greater or equal to the t-statistic computed without swapping the accuracies between conditions. In order to control for multiple comparisons, we used cluster correction with a threshold of $p = .002$.

3 Results

3.1 Nonlinear MVPD on nondenoised data reveals nonlinear dependence

First, we aimed to identify nonlinear interactions between BOLD responses in nondenoised data. In order to detect nonlinear dependence between face-selective regions and the rest of the brain, we computed the difference between the proportion of variance explained by the 1-layer, nonlinear MVPN and the proportion of variance explained by 1-layer, linear MVPN for all voxels in gray matter (see Materials and Methods for details). This analysis revealed that nonlinear MVPN explained greater variance in anterior regions of cortex, including prefrontal cortex (both medial and lateral) and the anterior temporal lobes bilaterally. By contrast, linear MVPN performed on par with nonlinear MVPN in posterior regions, including early visual cortex and posterior temporal cortex [Figure 1, A: Nondenosed].

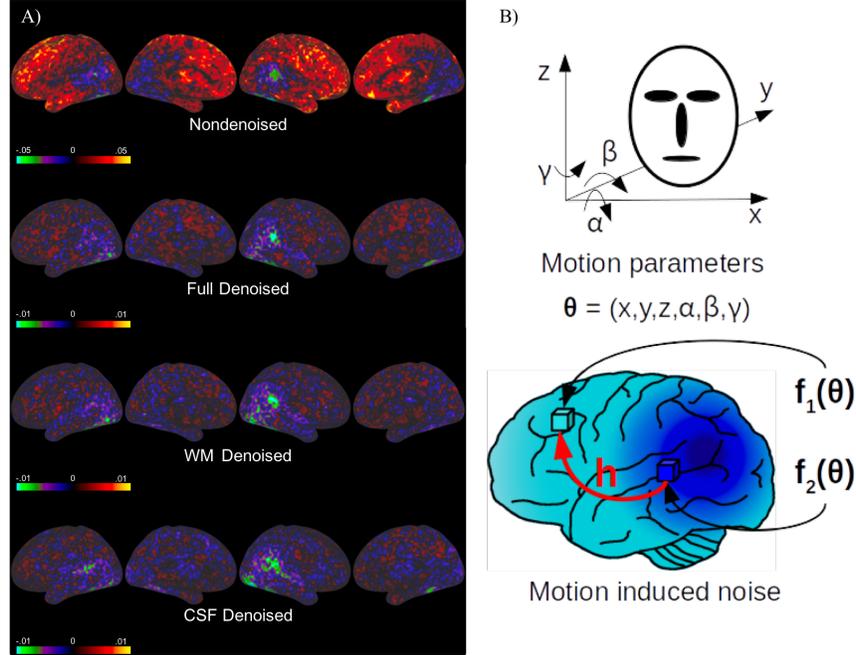


Figure 1: **A:** Difference between the proportion of variance explained by nonlinear and linear MVPN. Regions where nonlinear MVPN explains more variance are shown in red. Difference maps for nondenoised data, data denoised with CompCor using both white matter (WM) and cerebrospinal fluid (CSF), and data denoised with CompCor using WM alone or CSF alone are shown on an inflated cortical surface. **B:** Spurious nonlinear dependence between brain (h) can arise from a single noise source (i.e. head motion θ), if the noise-induced BOLD fluctuations in different brain regions are related to the source via different nonlinear functions (f_1, f_2). In addition, multiple noise sources affect the BOLD signal, including heart rate, respiration, etc.

Statistical analysis with SnPM revealed 10 significant clusters (between 7 and 13 voxels each) across the frontal lobe and anterior temporal lobe for which the variance explained was significantly larger for the nonlinear network than that of the linear network (all $p < .05$, FWE corrected, as assessed with SnPM). There were no significant clusters in which the linear network outperformed the nonlinear network.

3.2 CompCor denoising removes nonlinear dependence

After determining the presence of nonlinear dependence between face-selective regions and anterior frontal and temporal regions, we investigated whether CompCor attenuates or removes the advantage of nonlinear dependence over linear dependence. A common noise source (i.e. head motion) could differentially induce artifactual responses in different regions. More specifically, the noise source could be related to the noise-induced signal in different regions through different nonlinear functions, leading to spurious nonlinear dependence between those regions' responses [Figure 1, B]. In addition, the BOLD signal is affected by multiple sources of physiological noise (besides head motion), including heart rate and respiration, and multiple sources of measurement noise, including scanner drift. The interplay between the effect of these different noise sources on fMRI responses can further contribute to the emergence of spurious nonlinear dependence between brain regions.

Evaluating the performance of linear and nonlinear networks on the fully denoised data, the difference in performance between the two networks was greatly reduced. Across all voxels, the linear and nonlinear networks explained similar amounts of variance in the responses [Figure 1, Full Denoised]. Importantly, the pattern of nonlinear interactions across the anterior regions of the brain was eliminated, with performance of nonlinear MVPN on par with that of linear MVPN. Further analysis using SnPM showed that there were no significant clusters in which the nonlinear network performed better across subjects. The analysis identified one significant cluster of 7 voxels at the right STS in which the linear net explained greater variance than the nonlinear MVPN ($p_{FWE-corr} = .0018$). This cluster corresponds to one of the seed regions, suggesting that in seed regions nonlinear MVPN might lead to overfitting.

3.3 Information that removes nonlinear dependence is widespread anatomically and across PCs

Surprisingly, although CompCor uses linear regression to remove noise, its application eliminated the advantage of the nonlinear network over the linear network. One possible explanation for the ability of CompCor to differentially reduce nonlinear dependence is that the signal from the individual regions of no interest (e.g. CSF and WM) could separately correlate with noise in different regions of the brain. For example, BOLD signal from CSF might correlate with noise-induced BOLD fluctuations in the predictor regions, while responses from WM could correlate with noise-induced BOLD signal in anterior frontal and temporal target regions. In this scenario, the removal of either WM or CSF could lead to a reduction of the nonlinear relationship *between* those two brain areas. In one case, the reduction of the nonlinear relationship would be due to the removal of variance in the predictor regions; in the other, it would be due to the removal variance in the target regions (see the Discussion section for more details). For this reason, both denoising procedures could remove nonlinear interactions in the anterior brain areas through different mechanisms.

To test the relative contributions of WM and CSF to denoising, we trained both linear and nonlinear MVPN using the WM denoised as well as the CSF denoised data [Figure 1, WM denoised and CSF denoised respectively]. In both sets of denoised data, the nonlinear network performed on par with the linear network across the brain. In parallel with our findings in the full denoised data, the nonlinear network did not significantly outperform the linear network in either the CSF denoised or the WM denoised data. In the CSF denoised data, however, we did find one significant cluster of 5 voxels near the STS in which the linear network explained a greater variance the nonlinear network ($p_{FWE-corr} = .004$). As mentioned above, this cluster corresponds to one of the seed regions.

After testing how separately removing WM and CSF affects the reduction of nonlinear interactions, we aimed to determine whether a subset of the 5 PCs derived from the combined WM and CSF data plays a predominant role for the removal of nonlinear dependence. We hypothesized that a subset of the PCs might explain a higher proportion of noise in anterior regions of the brain, playing a predominant role in the reduction of nonlinear dependence between the predictor regions and the frontal cortex. To examine this possibility, we calculated the proportion of variance explained by each individual component for all voxels and visualized the average values across subjects [Figure 2]. Each PC exhibited a distinct pattern of explained variance; PC1 and PC2 showed greater variance explained in the frontal lobe

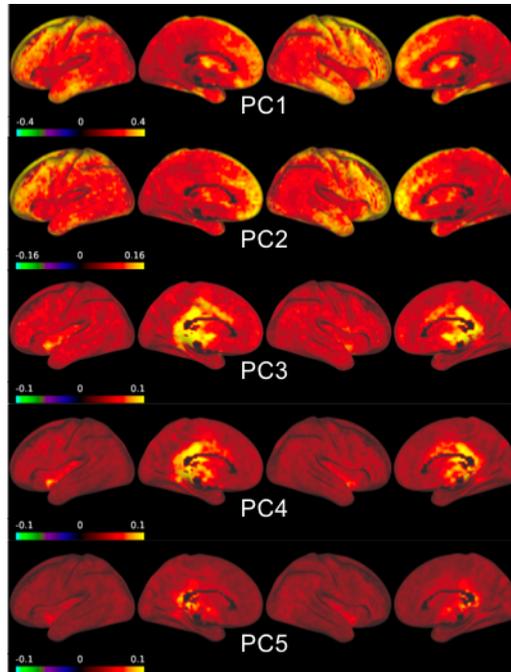


Figure 2: Proportion of variance explained in voxels' timecourses by principal components extracted from combined white matter and cerebrospinal fluid.

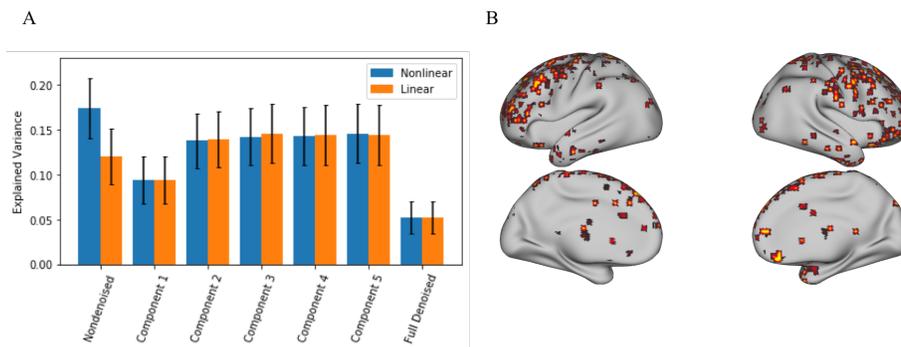


Figure 3: **A:** Proportion of variance explained by linear and nonlinear MVPN for different denoising procedures in the 1000 voxels showing the largest nonlinear > linear difference in the nondenoised data. Error bars represent the standard error of the mean across participants. **B:** Location of the 1000 voxels displaying the highest difference in nonlinear vs. linear MVPN performance on the nondenoised data, visualized on an inflated cortical surface.

and the anterior temporal lobe, matching the initial pattern of nonlinear interactions in the nondenoised data. PCs 3, 4, and 5, in contrast, showed peaks of explained variance in posterior cingulate and in subcortical areas [Figure 2].

To study in more detail the contribution of individual PCs to denoising, we examined the impact of removing one component at a time on the reduction of nonlinear MVPN between the predictor regions and anterior brain regions. In a first analysis, we focused on a region of interest including the 1000 voxels which showed the largest difference in explained variance between nonlinear and linear MVPN (nonlinear > linear) from the nondenoised data [Figure 3, B]. We examined the performance of the nonlinear and linear networks at predicting responses in these 1000 voxels from the responses in the seed regions, using the data denoised with one PC only, for each of the five PCs [Figure 3,

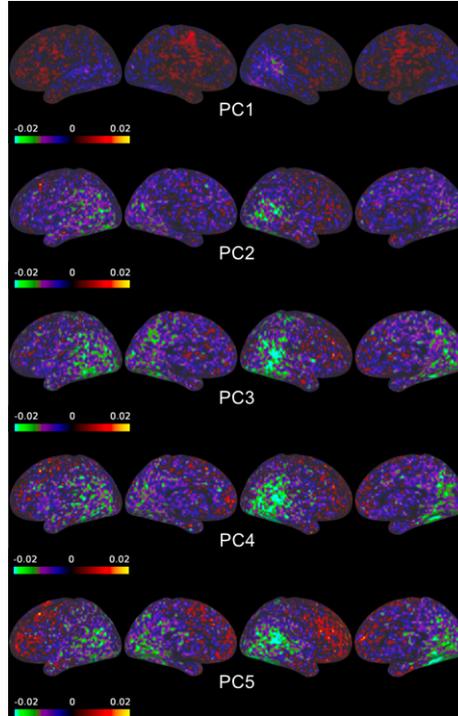


Figure 4: Difference between the proportion of variance explained by nonlinear and linear MVPN in data denoised with PCs 1-5. Regions better predicted by nonlinear MVPN are shown in red, regions better predicted by linear MVPN are shown in green.

A]. In each set of separately denoised data, the performance of linear and nonlinear networks were indistinguishable, suggesting that despite the fact that components differed in terms of the whole-brain patterns of explained variance [Figure 2], noise information that contributed to the removal of nonlinear dependence was not isolated to a specific subset of components, but rather was distributed across all components.

In a second analysis, we investigated whether denoising information is localized to individual PCs looking at regions beyond the 1000 voxels ROI. For each of the 5 denoising procedures using a single PC, we computed the difference of nonlinear and linear networks across the whole brain to observe the fine-grained differences in performance pattern across denoising methods [Figure 4]. Examining network differences in each set of denoised data, we found that, although removing each component eliminated significant differences in explained variance across networks, the components varied in the extent and location in which they removed most of the nonlinear vs. linear difference. Specifically, data denoised with PCs 2, 3, and 4 showed greater reduction in the nonlinear interactions as compared to the linear interactions in the frontal lobe and anterior temporal lobe.

3.4 Noise-dependent nonlinear interactions and their removal by CompCor are not specific to one particular network architecture

To further test whether the ability of CompCor to remove nonlinear interactions in nondenoised data was dependent upon the initial choice of neural network architecture used for MVPN, we replicated some of our analyses using a 5

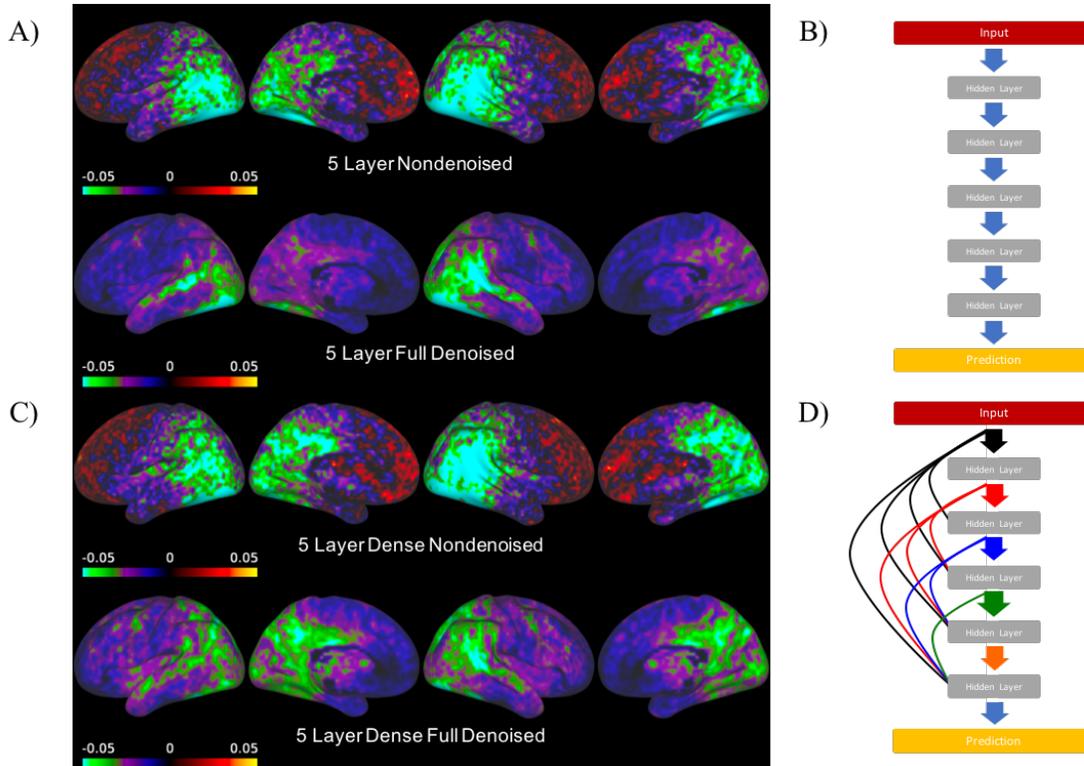


Figure 5: **A:** Difference in the proportion of variance explained by nonlinear and linear MVPN using a 5 layer feedforward neural network with 100 nodes per layer. Regions better predicted by nonlinear MVPN are shown in red, regions better predicted by linear MVPN are shown in green. **B:** Schematic of the 5 layer feedforward architecture. **C:** Difference in the proportion of variance explained by nonlinear and linear MVPN using a 5 layer densely connected neural network (DenseNet) with 100 nodes per layer. **D:** Schematic of the 5 layer DenseNet architecture.

layer feedforward architecture (with connections only between adjacent layers), and with a 5 layer dense architecture [[42]]. Estimating MVPN with these more complex networks on the nondenoised data, we were able to replicate the pattern of nonlinear interactions in anterior regions [Figure 5, 5 Layer Nondenoised and 5 Layer dense Nondenoised]. After denoising the data using 5 PCs extracted from the combined WM and CSF, nonlinear interactions dramatically decreased for both architectures [Figure 5, 5 Layer Full Denoised and 5 Layer dense Full Denoised]. In contrast to the 1 layer networks, it is worth noting that in both the 5 layer and the 5 layer dense architectures, the linear version outperformed the nonlinear version in a larger array of voxels in both the nondenoised *and* fully denoised data.

4 Discussion

Using functionally defined face-selective regions as a seed (OFA, FFA, pSTS), we found that, across network architectures, nonlinear MVPN outperformed linear MVPN in anterior temporal and frontal regions in nondenoised data. However, this effect was removed by CompCor denoising. Furthermore, regressing out any of the first five CompCor dimensions was sufficient to eliminate the advantage of nonlinear MVPN over linear MVPN in these anterior regions. These results demonstrate that the noise information that CompCor is able to capture is sufficient to mitigate false

positives in the study of nonlinear dependence between brain regions. We conclude that CompCor can be an important addition to the toolkit of researchers who aim to investigate nonlinear interactions across the brain.

4.1 What mechanisms could give rise to spurious nonlinear dependence?

Spurious nonlinear dependence between different brain regions might arise from a variety of mechanisms. One possible source of spurious nonlinear dependence is head motion. A common denoising approach consists of computing translation and rotation parameters from a reference frame, and regressing the change in these parameters out from the BOLD responses from one volume to the next (DVARs, see [45, 23]). However, at each time point, a given set of translation and rotation parameters produces different amounts of displacement for different voxels, depending on the type of motion (i.e. yaw vs. roll vs. pitch rotations) and the location of a voxel in the brain [[46, 47, 33]].

While the voxel-specific displacement resulting from head rotation is a linear function of the voxel’s distance from the center of rotation (the length of an arc is given by the radius times the angle in radians), previous studies suggest a nonlinear relationship between framewise displacement and its impact on measured BOLD responses [i.e. [23], Fig. 2]. As a consequence, if some regions are subjected to larger displacement than others, even if the amounts of displacement in the regions are linearly related the measured BOLD signal in these regions could be nonlinearly related.

In our study, we found that the nonlinear dependence removed by CompCor was localized predominantly to anterior brain regions [Figure 1]. Previous studies found that these regions show the greatest amount of displacement due to head movements [[33], Figure 1]. In particular, anterior regions are the most affected when participants ‘nod their head’ (pitch), and this movement was found to be the most frequent type of head movement during typical fMRI sessions [[33]].

Importantly, the seed regions we used in this study are more posterior, and thus they were likely subjected to a smaller amount of motion-induced displacement. Therefore, the nonlinear relationship we observed between the non-denoised data in the seed regions and in prefrontal and anterior temporal regions might be at least in part the result of head movements. However, we cannot rule out that other noise sources in addition to head movements (e.g., heart rate, respiration, scanner drift) could contribute to the spurious nonlinear dependence. Furthermore, preprocessing steps could play a role in generating spurious interactions, and indeed, previous research has demonstrated that increased activity in response to a task can induce systematic errors in functional connectivity estimates [[48]]. Importantly, regardless of the *source* of the noise, these results illustrate that CompCor is able to effectively mitigate confounds and reduce spurious findings of nonlinear dependence. Future studies could attempt to disentangle the relative contributions of head movements, heart rate, and respiration to spurious nonlinear dependence between brain regions.

4.2 How can CompCor, a linear method, remove nonlinear dependence?

The results of this study demonstrate that CompCor can remove spurious nonlinear dependence between brain regions. This might seem surprising given that CompCor is based on linear regression. One possible explanation for this

surprising result is that different PCs used in CompCor could be selectively related to noise in anterior vs. posterior regions. In this case, some PCs might have contributed differentially to removing noise in posterior regions, while other PCs might have contributed differentially to removing noise in anterior regions. As a consequence, even though the modeled relationship between the PCs and the measured BOLD signal is linear, CompCor could have contributed to removing nonlinear dependence between the posterior seed regions and anterior portions of the brain.

Our results show that, indeed, different components were characterized by different spatial distributions across cortex: the first two components explained comparatively more variance in anterior brain regions, while the remaining three components explained comparatively more variance in posterior regions [Figure 2]. However, in apparent contrast with the proposed mechanism, nonlinear dependence between the seed regions and anterior brain regions was largely reduced even when we regressed out each of the PCs taken individually [Figure 4].

To understand this phenomenon, it is important to note that in order to remove statistical dependence between two regions, it is not necessary to remove signal from both regions: removing signal from either the predictor or the target of prediction is sufficient to eliminate the nonlinear relationship between the respective regions. Consider spurious nonlinear dependence in which a target region’s noise y is a nonlinear function $f(x)$ of a predictor region’s noise x . If noise y is removed, function f will no longer predict responses in the target region even if noise x is not removed. Likewise, removing noise x will be sufficient to remove the nonlinear dependence between the predictor and the target regions, even if noise y is still present.

Previous research has suggested that, due to its reliance on the BOLD signal in regions of no interest (rather than on the noise sources directly), CompCor denoising could be able to remove noise that is nonlinearly related to the original sources of noise [[30]; [29]]. The current results show that CompCor can even effectively remove noise that shows nonlinear relationships across different brain regions. This finding is especially significant given that measures of statistical dependence, such as MVPD and functional connectivity, are often applied to the study of development [[49]], aging [[50]], and patient populations [[51]], in combination with the fact that head motion is affected by age [[52, 53]], impulsivity [[54]], autism [[55]], and stroke [[56]]. This is particularly problematic considering that even small differences in head motion can produce differences in estimates of statistical dependence [[57]]. Denoising procedures such as CompCor are essential to mitigate possible false positives in studies of group differences in the interactions between regions, for both linear and nonlinear measures of statistical dependence.

4.3 Does using a different architecture protect from spurious nonlinear dependence?

Additional analyses using a 5 layer network and with a 5 layer DenseNet [[42]] revealed that the observed nonlinear dependence with anterior regions in nondenoised data is not unique to one-hidden-layer neural network architectures, and instead appear robustly across the choice of network architecture. This result is in line with previous findings suggesting that the type of neural network architecture used in MVPN has a quantitative, but not a qualitative, impact on the results [[40]]. Importantly, CompCor also effectively removed nonlinear dependence with anterior regions for

these more complex network architectures, indicating that its contribution to removing spurious nonlinear dependence extends beyond one-layer MVPN architectures.

5 Conclusion

As methods for the study of connectivity and statistical dependence become increasingly sophisticated, they also become increasingly susceptible to noise. Novel methods have the ability to capture nonlinear interactions between brain regions [[14, 1]] - inspiring the question: are existing denoising approaches adequate to eliminate spurious nonlinear dependence between brain regions? In this study, we found (perhaps surprisingly) that a widely used linear denoising approach - CompCor [[22]] - can contribute to removing spurious nonlinear dependence between brain regions, robustly across different choices of the nonlinear neural networks used to model such dependence.

By repeatedly testing the effects of various noise regressors on the nonlinear relationship between the neural activity in different brain areas, we were able to determine that the noise information giving rise to spurious nonlinear dependence is distributed across WM and CSF, and across multiple PCs extracted from these regions of no interest. These findings indicate that CompCor can play an important role in mitigating false positives even in studies aiming to model nonlinear interactions between brain regions, and highlight the importance of incorporating this denoising method into the standard preprocessing pipeline.

Availability of Data and Material

All analyses in this study were performed using data from the publicly available *studyforrest* project, which can be accessed at <http://studyforrest.org/>.

Code Availability

The MVPN code can be found at <https://github.com/pandamt/MVPN>.

Acknowledgments

We would like to thank Emily Schwartz and Tony Chen for their invaluable contributions to this manuscript. We would also like to thank the researchers involved in the *studyforrest* project for sharing their data.

Conflicts of Interest/Competing Interests

Declarations of interest: none.

Funding

This work was funded by the Department of Psychology and Neuroscience at Boston College. Aidas Aglinskas and Stefano Anzellotti were supported by a grant from the Simons Foundation/SFARI (Grant Number: 614379, to S.A. and J.H.).

Author Contribution Statement

Craig Poskanzer: Methodology, Formal Analysis, and Writing- Original Draft; **Mengting Fang:** Methodology and Writing- Review and Editing; **Aidas Aglinskas:** Writing- Review and Editing; **Stefano Anzellotti:** Conceptualization, Supervision, Methodology, Writing- Original Draft.

References

- [1] Stefano Anzellotti and Marc N Coutanche. Beyond functional connectivity: investigating networks of multivariate representations. *Trends in cognitive sciences*, 22(3):258–269, 2018.
- [2] Marc N Coutanche and Sharon L Thompson-Schill. Informational connectivity: identifying synchronized discriminability of multi-voxel patterns across the brain. *Frontiers in human neuroscience*, 7:15, 2013.
- [3] Stefano Anzellotti, Alfonso Caramazza, and Rebecca Saxe. Multivariate pattern dependence. *PLoS Computational Biology*, 13(11):e1005799, November 2017.
- [4] Linda Geerligs, Richard N Henson, et al. Functional connectivity and structural covariance between regions of interest can be measured more accurately using multivariate distance correlation. *NeuroImage*, 135:16–31, 2016.
- [5] Stefano Anzellotti, Evelina Fedorenko, Alexander J. E. Kell, Alfonso Caramazza, and Rebecca Saxe. Measuring and modeling nonlinear interactions between brain regions with fmri. *bioRxiv*, page 074856, 2017.
- [6] Takuya Ito, Luke Hearne, Ravi Mill, Carrisa Cocuzza, and Michael W Cole. Discovering the computational relevance of brain network organization. *Trends in cognitive sciences*, 24(1):25–38, 2020.
- [7] Maximilian Riesenhuber and Tomaso Poggio. Hierarchical models of object recognition in cortex. *Nature Neuroscience*, 2:1019–1025, 1999.
- [8] Srdjan Ostojic and Nicolas Brunel. From spiking neuron models to linear-nonlinear models. *PLoS Comput Biol*, 7(1):e1001056, 2011.
- [9] Duane G Albrecht, Wilson S Geisler, and Alison M Crane. Nonlinear properties of visual cortex neurons: Temporal dynamics, stimulus selectivity, neural performance. *The visual neurosciences*, 1:747–764, 2003.
- [10] Kendrick N Kay, Jonathan Winawer, Aviv Mezer, and Brian A Wandell. Compressive spatial summation in human visual cortex. *Journal of neurophysiology*, 110(2):481–494, 2013.

- [11] Seyed-Mahdi Khaligh-Razavi and Nikolaus Kriegeskorte. Deep supervised, but not unsupervised, models may explain it cortical representation. *PLoS computational biology*, 10(11):e1003915, 2014.
- [12] Daniel LK Yamins, Ha Hong, Charles F Cadieu, Ethan A Solomon, Darren Seibert, and James J DiCarlo. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences*, 111(23):8619–8624, 2014.
- [13] Klaas Enno Stephan, Lars Kasper, Lee M Harrison, Jean Daunizeau, Hanneke EM den Ouden, Michael Breakspear, and Karl J Friston. Nonlinear dynamic causal models for fmri. *Neuroimage*, 42(2):649–662, 2008.
- [14] Joseph T Lizier, Jakob Heinzle, Annette Horstmann, John-Dylan Haynes, and Mikhail Prokopenko. Multivariate information-theoretic measures reveal directed information structure and task relevant changes in fmri connectivity. *Journal of computational neuroscience*, 30(1):85–107, 2011.
- [15] Chia-Shang J Liu, Atsushi Miki, Justin Hulvershorn, Luke Bloy, Eugene E Gualtieri, Grant T Liu, John S Leigh, John C Haselgrove, and Mark A Elliott. Spatial and temporal characteristics of physiological noise in fmri at 3t. *Academic Radiology*, 13(3):313–323, 2006.
- [16] G Krüger and G H Glover. Physiological noise in oxygenation-sensitive magnetic resonance imaging. *Magnetic Resonance in Medicine*, 46(4):631–637, 2001.
- [17] Karl J. Friston, Steven Williams, Robert Howard, Richard S. J. Frackowiak, and Robert Turner. Movement-related effects in fmri time-series. *Magnetic Resonance in Medicine*, 35(3):346–355, 1996.
- [18] Thomas T Liu. Noise contributions to the fmri signal: An overview. *NeuroImage*, 143:141–151, 2016.
- [19] J Michael Tyszka, Daniel P Kennedy, Lynn K Paul, and Ralph Adolphs. Largely typical patterns of resting-state functional connectivity in high-functioning adults with autism. *Cerebral cortex*, 24(7):1894–1905, 2014.
- [20] Paul M. Macey, Katherine E. Macey, Rajesh Kumar, and Ronald M. Harper. A method for removal of global effects from fmri time series. *Neuroimage*, 22(1):360–366, 2004.
- [21] K. J. Friston, O. Josephs, E. Zarahn, A. P. Holmes, S. Rouquette, and J. B. Poline. To smooth or not to smooth?: Bias and efficiency in fmri time-series analysis. *Neuroimage*, 12(2):196–208, 2000.
- [22] Yashar Behzadi, Khaled Restom, Joy Liau, and Thomas T. Liu. A component based noise correction method (compcor) for bold and perfusion based fmri. *NeuroImage*, 37(1):90–101, 2007.
- [23] Jonathan D. Power, Kelly A. Barnes, Abraham Z. Snyder, Bradley L. Schlaggar, and Steven E. Petersen. Spurious but systematic correlations in functional connectivity mri networks arise from subject motion. *NeuroImage*, 59:2142–2154, 2012.
- [24] Jonathan D Power, Bradley L Schlaggar, and Steven E Petersen. Recent progress and outstanding issues in motion correction in resting state fmri. *Neuroimage*, 105:536–551, 2015.
- [25] Jonathan D Power, Mark Plitt, Timothy O Laumann, and Alex Martin. Sources and implications of whole-brain fmri signals in humans. *Neuroimage*, 146:609–625, 2017.

- [26] Rastko Ciric, Daniel H Wolf, Jonathan D Power, David R Roalf, Graham L Baum, Kosha Ruparel, Russell T Shinohara, Mark A Elliott, Simon B Eickhoff, Christos Davatzikos, et al. Benchmarking of participant-level confound regression strategies for the control of motion artifact in studies of functional connectivity. *Neuroimage*, 154:174–187, 2017.
- [27] Linden Parkes, Ben Fulcher, Murat Yücel, and Alex Fornito. An evaluation of the efficacy, reliability, and sensitivity of motion correction strategies for resting-state functional mri. *Neuroimage*, 171:415–436, 2018.
- [28] Yichen Li, Rebecca Saxe, and Stefano Anzellotti. Intersubject mvpd: Empirical comparison of fmri denoising methods for connectivity analysis. *PLoS ONE*, 14:e0222914, September 2019.
- [29] César Caballero-Gaudes and Richard C. Reynolds. Methods for cleaning the bold fmri signal. *NeuroImage*, 154:128–149, 2017.
- [30] John Muschelli, Mary Beth Nebel, Brian S. Caffo, Anita D. Barber, James J. Pekar, and Stewart H. Mostofsky. Reduction of motion-related artifacts in resting state fmri using acompcor. *NeuroImage*, 96:22–35, 2014.
- [31] Avital Hahamy, Vince Calhoun, Godfrey Pearlson, Michal Harel, Nachum Stern, Fanny Attar, and Rafael Malachvand Roy Salomon. Save the global: Global signal connectivity as a tool for studying clinical populations with functional magnetic resonance imaging. *Brain Connectivity*, 4:395–403, 2014.
- [32] Mark Hallett, Willem de Haan, Gustavo Deco, Reinhard Dengler, Riccardo Di Iorio, Cecile Gallea, Christian Gerloff, Christian Grefkes, Rick C. Helmich, Morten L. Kringelbach, Francesca Miraglia, Ivan Rektor, Ondřej Strýček, Fabrizio Vecchio, Lukas J. Volz, Tao Wu, and Paolo M. Rossini. Human brain connectivity: Clinical applications for clinical neurophysiology. *Clinical Neurophysiology*, 131:1388–2457, 2020.
- [33] Chao-Gan Yan, Brian Cheung, Clare Kelly, Stan Colcombe, R Cameron Craddock, Adriana Di Martino, Qingyang Li, Xi-Nian Zuo, F Xavier Castellanos, and Michael P Milham. A comprehensive assessment of regional variation in the impact of head micromovements on functional connectomics. *Neuroimage*, 76:183–201, 2013.
- [34] Michael Hanke, Nico Adelhöfer, Daniel Kottke, Vittorio Iacovella, Ayan Sengupta, Falko R. Kaule, Roland Nigbur, Alexander Q. Waite, Florian Baumgartner, and Jörg Stadler. A studyforrest extension, simultaneous fmri and eye gaze recordings during prolonged natural stimulation. *Scientific Data*, 3:160092, October 2016.
- [35] Oscar Esteban, Christopher J. Markiewicz, Ross W. Blair, Craig A. Moodie, A. Ilkay Isik, Asier Erramuzpe, James D. Kent, Mathias Goncalves, Elizabeth DuPre, Madeline Snyder, Hiroyuki Oya, Satrajit S. Ghosh, Jesse Wright, Joke Durnez, Russell A. Poldrack, and Krzysztof J. Gorgolewski. fmriprep: a robust preprocessing pipeline for functional mri. *Nature Methods*, 16:111–116, December 2019.
- [36] Ayan Sengupta, Falko R. Kaule, J. Swaroop Guntupalli, Michael B. Hoffmann, Christian Häusler, Jörg Stadler, and Michael Hanke. A studyforrest extension, retinotopic mapping and localization of higher visual areas. *Scientific Data*, 3:160093, October 2016.
- [37] Y Zhang, M Brady, and S Smith. Segmentation of brain mr images through a hidden markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Images*, 20(1):45–57, 2001.

- [38] M Jenkinson, P Bannister, J. M Brady, and S. M Smith. Improved optimisation for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2):825–841, 2002.
- [39] M. W. Woolrich, B. D. Ripley, M. Brady, and S. M. Smith. Temporal autocorrelation in univariate linear modeling of fmri data. *Neuroimage*, 14(6):1370–1386, December 2001.
- [40] Mengting Fang, Aidan Aglinskas, Yichen Li, and Stefano Anzellotti. Identifying hubs that integrate responses across multiple category-selective regions. *PsyArXiv*, November 2019.
- [41] Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural networks. *science*, 313(5786):504–507, 2006.
- [42] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [43] SnPM. Statistical non parametric mapping toolbox (snpm) 2013, 2013.
- [44] MATLAB. version 9.5.0.944444 (r2018b), 2018.
- [45] Christopher D Smyser, Terrie E Inder, Joshua S Shimony, Jason E Hill, Andrew J Degnan, Abraham Z Snyder, and Jeffrey J Neil. Longitudinal analysis of neural network development in preterm infants. *Cerebral cortex*, 20(12):2852–2862, 2010.
- [46] Marko Wilke. An alternative approach towards assessing and accounting for individual motion in fmri timeseries. *Neuroimage*, 59(3):2062–2072, 2012.
- [47] Theodore D Satterthwaite, Daniel H Wolf, James Loughhead, Kosha Ruparel, Mark A Elliott, Hakon Hakonarson, Ruben C Gur, and Raquel E Gur. Impact of in-scanner head motion on multiple measures of functional connectivity: relevance for studies of neurodevelopment in youth. *Neuroimage*, 60(1):623–632, 2012.
- [48] Michael W Cole, Takuya Ito, Douglas Schultz, Ravi Mill, Richard Chen, and Carrisa Cocuzza. Task activations produce spurious but systematic inflation of task functional connectivity estimates. *Neuroimage*, 189:1–18, 2019.
- [49] Hilary Richardson. Development of brain networks for social functions: Confirmatory analyses in a large open source dataset. *Developmental cognitive neuroscience*, 37:100598, 2019.
- [50] Linda Geerligs, Remco J Renken, Emi Saliassi, Natasha M Maurits, and Monique M Lorst. A brain-wide study of age-related changes in functional connectivity. *Cerebral cortex*, 25(7):1987–1999, 2015.
- [51] Jocelyn V Hull, Lisa B Dokovna, Zachary J Jacokes, Carinna M Torgerson, Andrei Irimia, and John Darrell Van Horn. Resting-state functional connectivity in autism spectrum disorders: a review. *Frontiers in psychiatry*, 7:205, 2017.
- [52] Jie Song, Alok S Desphande, Timothy B Meier, Dana L Tudorascu, Svyatoslav Vergun, Veena A Nair, Bharat B Biswal, Mary E Meyerand, Rasmus M Birn, Pierre Bellec, et al. Age-related differences in test-retest reliability in resting-state brain functional connectivity. *PLoS One*, 7(12), 2012.

- [53] Benjamin O Turner, Brian Lopez, Tyler Santander, and Michael B Miller. One dataset, many conclusions: Bold variability’s complicated relationships with age and motion artifacts. *Brain imaging and behavior*, 9(1):115–127, 2015.
- [54] Xiang-zhen Kong, Zonglei Zhen, Xueting Li, Huan-hua Lu, Ruosi Wang, Ling Liu, Yong He, Yufeng Zang, and Jia Liu. Individual differences in impulsivity predict head motion during magnetic resonance imaging. *PloS one*, 9(8), 2014.
- [55] Carla Caballero, Sejal Mistry, Joe Vero, and Elizabeth B Torres. Characterization of noise signatures of involuntary head motion in the autism brain imaging data exchange repository. *Frontiers in Integrative Neuroscience*, 12:7, 2018.
- [56] Emily Seto, G Sela, WE McIlroy, SE Black, WR Staines, MJ Bronskill, AR McIntosh, and SJ Graham. Quantifying head motion associated with motor tasks used in fmri. *Neuroimage*, 14(2):284–297, 2001.
- [57] Koene R.A. Van Dijk, Mert R. Sabuncu, and Randy L. Buckner. The influence of head motion on intrinsic functional connectivity mri. *NeuroImage*, 59:431–438, 2012.

6 Supplementary Materials

Average MNI Coordinates for Face Selective ROIs		
ROI	Right Hemisphere	Left Hemisphere
FFA	(42.9, -47.9, -18.5)	(-43.2, -49.8, -18.3)
Face Selective STS	(54.4, -46.7, 12.5)	(-50.6, -49.5, 11.1)
OFA	(38.4, -79.9, -12.9)	(-40.2, -78.0, -17.0)

This preprint has not undergone peer review (when applicable) or any post-submission improvements or corrections. The Version of Record of this article is published in *Neuroinformatics*, and is available online at <https://doi.org/10.1007/s12021-021-09540-9>.