

Functional connectivity in resting-state fMRI: Is linear correlation sufficient?

Jaroslav Hlinka^a, Milan Paluš^a, Martin Vejmelka^a, Dante Mantini^{b,c}, Maurizio Corbetta^{c,d,e}

^a*Institute of Computer Science, Academy of Sciences of the Czech Republic, Pod vodarenskou vezi 2, 18207 Prague, Czech Republic*

^b*Laboratory for Neuro- and Psychophysiology, Katholieke Universiteit Leuven, 3000 Leuven, Belgium*

^c*Institute for Advanced Biomedical Technologies, G. D'Annunzio University Foundation, G. D'Annunzio University, 66013 Chieti, Italy*

^d*Department of Radiology, Washington University, St. Louis, Missouri, USA*

^e*Department of Neurology, Washington University, St. Louis, Missouri, USA*

Abstract

Functional connectivity (FC) analysis is a prominent approach to analyzing fMRI data, especially acquired under the resting state condition. The commonly used linear correlation FC measure bears an implicit assumption of Gaussianity of the dependence structure. If only the marginals, but not all the bivariate distributions are Gaussian, linear correlation consistently underestimates the strength of the dependence. To assess the suitability of linear correlation and the general potential of nonlinear FC measures, we present a framework for testing and estimating the deviation from Gaussianity by means of comparing mutual information in the data and its Gaussianized counterpart. We apply this method to 24 sessions of human resting state fMRI. For each session, matrix of connectivities between 90 anatomical parcel time series is computed using mutual information and compared to results from its multivariate Gaussian surrogate that conserves the correlations but cancels any nonlinearity. While the group-level tests confirmed non-Gaussianity in the FC, the quantitative assessment revealed that the portion of mutual information neglected by linear correlation is relatively minor – on average only about 5% of the mutual information already captured by the linear correlation. The marginality of the non-Gaussianity was confirmed in comparisons using clustering of the parcels – the disagreement between clustering obtained from mutual information and linear correlation was attributable to random error. We conclude that for this type of data, practical relevance of nonlinear methods trying to improve over linear correlation might be limited by the fact that the data are indeed almost Gaussian.

Keywords: fMRI, functional connectivity, Gaussianity, nonlinearity, correlation, mutual information

1. Introduction

Neuroimaging methods play an important role in the process of extending our understanding of brain structure and function. In this regard, functional Magnetic Resonance Imaging (fMRI) holds a central position, particularly due to its non-invasiveness and its relatively high spatial resolution. Despite the large number of neuroimaging studies conducted to delineate brain activations related to a wide range of cognitive functions, the fundamental issue of how the brain regions communicate to each other still remains open.

The concept of brain functional connectivity

(Friston, 1994) is central to the understanding of the organized behavior of cortical regions beyond the simple mapping of their activity. Functional connectivity analyses of fMRI data as an approach to study temporal coherence in activity of distant brain areas, either during task performance or in resting state, has been widely applied in the neuroscience research. Much recent work has focused on measuring and interpreting the spontaneous signal fluctuations typically seen during rest (Fox and Raichle, 2007). The initial reservations regarding the neural origin of these fluctuations (and the related connectivity) are now mostly overcome - some of the strongest evidence for relevance of these sig-

nal fluctuations is based on the observed direct or indirect electrophysiological correlates (Laufs, 2008; Mantini et al., 2007; Miller et al., 2009), although there is still much to be done in the field of methods for separation of neural and nonneural sources of resting signal variation. In the context of this paper, we stress that in resting state studies, functional connectivity analysis commonly plays a central role (Fox and Raichle, 2007; van Dijk et al., 2010).

The most widely spread method of measuring functional connectivity between a pair of regions is computing a linear correlation of activity time series derived from these regions by e.g. simple spatial averaging across all the voxels in the regions. Linear correlation is also widely used to obtain so-called correlation maps by correlating the seed voxel or seed region signal with signal from all the other voxels in the brain, or constrained to gray matter area. From all possible bivariate measures of association, linear correlation is clearly a method of first choice, reflecting the assumption that the relationship between the fMRI time series can be suitably approximated by a multivariate Gaussian white noise process. Additionally, linear correlation is a well-known statistical concept, sufficiently simple to allow wide use and easy communication of results between researchers of diverse backgrounds.

On the other hand, from the mid-1980s, nonlinear approaches to analysis of brain signals are getting increased interest of researchers who consider non-linearity as an intrinsic property of brain dynamics, see e.g. (Stam, 2005) for a review. Hemodynamic nonlinearities are known to affect the blood oxygenation level-dependent (BOLD) fMRI signal (de Zwart et al., 2009). More specifically, non-linearity of dependence between fMRI time series during resting state has been reported (Lahaye et al., 2003). Use of non-linear measures of functional connectivity for the analysis of resting state data has been proposed (Deshpande et al., 2006; Xie et al., 2008; Maxim et al., 2005), particularly including measures based on analysis of chaotic non-linear dynamical systems to analyze resting state data, suggesting that the assumption of linearity might be oversimplifying.

The question arises, to what extent and in what context is it justified and beneficial to use non-linear measures of functional connectivity and more generally any indices such as those motivated by the assumption of non-linear dynamical system as the underlying mechanism behind the resting state

fMRI signal. When linear correlation is used as a measure of functional connectivity, there are some implicit assumptions made. The first is that the information in the temporal order of the samples can be ignored (both within each timeseries and the mutual interaction). While the extent of justifiability of this assumption deserves exploration of its own, we keep this interim assumption for the purposes of this paper, not least in order to keep the comparison of linear correlation to nonlinear measures fair. Accepting for now the assumption of no role of temporal order of samples, we ask if the instantaneous (zero-lag) dependence between the time series, expressed in the probability distribution $p(X, Y)$, is fully captured by the linear correlation $r(X, Y)$. We answer that this is true under the assumption of bivariate Gaussianity of the distribution. Indeed, a multivariate normal (Gaussian) distribution is uniquely defined by its correlation - up to linear shifts and rescaling. To be more precise, bivariate normal distribution is fully characterised by its mean $\mu = (\mu_x, \mu_y)$ and its 2×2 covariance matrix $\text{Cov}(X, Y)$ - if we allow for linear shifting and scaling, the remaining invariant parameter characterizing fully the distribution is indeed the correlation $r(X, Y)$. For a bivariate Gaussian distribution, the correlation also uniquely defines the mutual information shared between the two variables X, Y which can be computed as $I(X, Y) = I_{Gauss}(r) \equiv -\frac{1}{2} \log(1 - r^2)$.

On the other side, when the Gaussianity assumption does not hold, the distribution cannot be fully described by the mean and covariance. More, possibly infinitely many higher order moments need to be specified to determine the distribution. As the correlation is not sufficient to describe the dependence structure, the equation for $I(X, Y)$ above cannot hold in general. Interestingly, as we detail later, the notable properties of normal distributions enable a derivation of a useful lower bound on mutual information valid for a broad class of probability distributions. In particular, for a bivariate distribution $p(X, Y)$ with standard normal marginals $p(X), p(Y)$, it holds that $I(X, Y) \geq I_{Gauss}(r) = -\frac{1}{2} \log(1 - r^2)$, where the equality holds exactly for bivariate Gaussian distributions. This allows us to quantify the deviation from Gaussianity as the difference between the total mutual information of the two variables $I(X, Y)$ and the mutual information $I_{Gauss}(r) = -\frac{1}{2} \log(1 - r^2)$ that correspond to bivariate Gaussian distribution with the observed correlation r .

While there are many potential nonlinear FC measure candidates, mutual information holds a specific position among these for its generality. In theory, it is general enough to capture an arbitrary form of dependence relation between the variables without any a priori model restrictions on its form. The properties of mutual information allow us not only to test the suitability of linear correlation through probing the Gaussianity of the fMRI time series, but also to construct a quantitative estimate of connectivity information neglected by the use of linear correlation. This gives the amount of additional information available and bounds the potential contribution of non-linear alternatives over the Pearson correlation coefficient.

We implement the outlined ideas by comparing the total mutual information between the signals with the mutual information between the signals in surrogate datasets. These surrogates are generated in a way that preserves the linear correlation, but cancels any nonlinear information by enforcing bivariate Gaussian distribution on the surrogate signal-pair. This approach allows us to both test and quantify the deviation from Gaussianity, providing a principled guide in judging the suitability of linear correlation as a measure of FC. The focus on bivariate Gaussianity as the crucial condition of suitability of use of linear correlation as FC index, along with the illustrative quantitative estimation of the deviation from Gaussianity by means of the mutual information neglected by linear correlation, are the two main contributions of this study to the discussion of fMRI functional connectivity methods. We apply the presented method to parcel-average time series obtained from resting state fMRI BOLD signal of healthy subjects, testing and quantifying the deviation from bivariate Gaussianity. We complement this by assessing the relevance of the detected non-Gaussianity through tests of agreement of clustering results obtained from original and Gaussianized data.

2. Materials and methods

2.1. Data

Twelve right-handed healthy young volunteers (5 males and 7 females, age range 20 – 31 years) participated in the study. Participants were informed about the experimental procedures and provided written informed consent. The study design was approved by the local Ethics Committee of Chieti

University. Subjects lay in a supine position and viewed a black screen with a centered red fixation point of 0.3 visual degrees, through a mirror tilted by 45 degrees. Each volunteer underwent two scanning runs of 10 minutes in a resting-state condition. Specifically, they were instructed to be relaxed, but to maintain fixation during scanning. The eye position was monitored at 120 Hz during scanning using an ISCAN eye tracker system.

Scanning was performed with a 3T MR scanner (Achieva; Philips Medical Systems) located at the Institute for Advanced Biomedical Technologies in Chieti, Italy. Functional images were obtained using T2-weighted echo-planar imaging (EPI) with blood oxygenation level-dependent (BOLD) contrast using SENSE imaging. EPIs (TR/TE=2000/35 ms) comprised 32 axial slices acquired continuously in ascending order covering the entire cerebrum (voxel size=3x3x3.5mm³). For each scanning run, initial 5 dummy volumes allowing the MRI signal to reach steady state were discarded. The next 300 functional volumes were used for the analysis. A three-dimensional high-resolution T1-weighted image (TR/TE=9.6/4.6 ms, voxel size=0.98x0.98x1.2mm³) covering the entire brain was acquired at the end of the scanning session and used for anatomical reference.

Initial data preprocessing was performed using the SPM5 software package (Wellcome Department of Cognitive Neurology, London, UK) running under MATLAB (The Mathworks). The preprocessing steps involved the following: (1) correction for slice-timing differences (2) correction of head-motion across functional images, (3) coregistration of the anatomical image and the mean functional image, and (4) spatial normalization of all images to a standard stereotaxic space (Montreal Neurological Institute, MNI) with a voxel size of 3x3x3 mm³.

Ninety parcels from the Automated Anatomical Labeling (AAL) atlas were used to extract mean BOLD time series after masking out non-gray matter voxels. The anatomical positions of the parcels are described in (Tzourio-Mazoyer et al., 2002). Every parcel time series was orthogonalized with respect to motion parameters and global mean signal and high-pass filtered at 1/120Hz.

2.2. Analysis

As already mentioned in the Introduction, for a bivariate distribution $p(X, Y)$ with standard nor-

mal marginals $p(X), p(Y)$, it holds that

$$I(X, Y) \geq I_{Gauss} = -\frac{1}{2} \log(1 - r^2), \quad (1)$$

where the equality holds exactly for bivariate Gaussian distributions. The inequality (1) stems from the fact, that normal distribution is the *maximum* entropy distribution for a given covariance matrix (or for a given correlation, as we assume without loss of generality that $\sigma(X) = \sigma(Y) = 1$). From the relation between mutual information and entropy ($I(X, Y) = H(X) + H(Y) - H(X, Y)$) it follows that mutual information of Gaussian distribution $I_{Gauss}(r)$ is then *minimal* from all distributions of given correlation r , under the assumption of fixed marginal entropies, which is true when the marginals have standard normal distribution. Note that the assumption of normality of the marginals is far less restrictive than it might seem. First, approximate data normality is commonly assumed in areas not restricted to fMRI FC analysis. More importantly, even if we find particular data deviating strongly from normality, any sample distribution can be monotonously transformed to match normal distribution.

To assure precise non-Gaussianity estimates, we have indeed carried out this ‘normalization’ step. First, to each of the values x_i of the original variable x its sample percentile p_i is assigned. Subsequently, the values x_i of the original variable are replaced by values n_i corresponding to the respective percentiles p_i in the standard normal distribution $N(0, 1)$.

For two discrete random variables X_1, X_2 with sets of values Ξ_1 and Ξ_2 , the mutual information is defined as

$$I(X_1, X_2) = \sum_{x_1 \in \Xi_1} \sum_{x_2 \in \Xi_2} p(x_1, x_2) \log \frac{p(x_1, x_2)}{p(x_1)p(x_2)},$$

where the probability distribution function is defined by $p(x_i) = Pr\{X_i = x_i\}, x_i \in \Xi_i$ and the joint probability distribution function is $p(x_1, x_2)$ is defined analogously. When the discrete variables X_1, X_2 are obtained from continuous variables on a continuous probability space, then the mutual information depends on a partition ξ chosen to discretize the space. Here a simple box-counting algorithm based on marginal equiquantization method (Palus et al., 1993) was used, i.e., a partition was generated adaptively in one dimension (for each variable) so that the marginal bins

become equiprobable. This means that there is approximately the same number of data points in each marginal bin. In this paper we used a simple pragmatic choice of $Q = 8$ bins for each marginal variable Palus and Vejmelka (2007).

For each session, we have computed the mutual information (MI) for each pair of parcels, yielding a symmetric 90-by-90 matrix of MI values. Notably, the raw estimates of mutual information suffer from some inevitable bias due to finite sample size and discretization of the variables. While this bias does not affect the statistical testing framework, it is necessary to correct for it to allow the use of standard scales for reporting the resulting mutual information and for more accurate computation of the neglected information.

This correction is carried out with the help of sample mutual information estimates for finite-size ($N = 300$) random samples from bivariate Gaussian distributions with known correlations (and therefore known mutual informations $I_{Gauss} = -\frac{1}{2} \log(1 - r^2)$). These are computed for 50000 bivariate random samples (each with size $N = 300$) for each correlation in the range from 0 to 1 (in 200 steps of 0.005). The average of the 50000 mutual information estimates gives us an approximate expected sample mutual information estimate corresponding to each tabulated correlation/mutual information. Thus we obtain a monotonous function transforming the known true correlation (or the respective mutual information) to its expected numerical MI estimate for sample size $N = 300$. Once generated, the inverse of this monotonous function can be used to transform each estimated mutual information to obtain a more accurate bias-corrected estimate of the true mutual information. (The relationship between the true population mutual information and the expected finite-sample estimate is shown in Supplementary Figure 1.)

To compare the (total) mutual information to the portion of information conveyed in the linear correlation, for each dataset, 99 random realizations of multivariate time series preserving the linear structure but canceling the nonlinear structure were constructed, and MI was computed for these surrogates. If the original time series dependence structure was Gaussian (and therefore fully captured by the linear correlation), the MI in the surrogates should not differ from the original MI, up to some random error. The alternative case should manifest itself as a decrease in the MI in the surrogates with respect to the original data.

The surrogates were constructed as multivariate Fourier transform (FT) surrogates (Prichard and Theiler, 1994; Palus, 1997): realizations of multivariate linear stochastic process which mimic individual spectra of the original time series as well as their cross-spectrum. The multivariate FT surrogates are obtained by computing the Fourier transform of the series, keeping unchanged the magnitudes of the Fourier coefficients (the amplitude spectrum), but adding the same random number to the phases of coefficients of the same frequency bin; the inverse FT into the time domain is then performed. The multivariate FT surrogates preserve the part of dependence which can be explained by a multivariate linear stochastic process.

The idea of comparing the MI of data to MI of ‘linear’ surrogates rather than directly to linear correlation of data has two aspects. First, it allows a direct quantitative comparison of the nonlinear and linear connectivity, while correlation and mutual information estimators have generally different properties. Second, generation of the surrogates allows direct statistical testing of the difference. However, this procedure generates 99 estimates of the linear MI for each parcel pair; one for each surrogate. While these are useful for hypothesis testing, for general presentation of the difference we use the mean value of these 99 values. In the following we refer to this as ‘Gaussian’ MI, and it actually closely estimates the MI of a bivariate Gaussian distribution $I_{Gauss}(r) = -\frac{1}{2} \log(1 - r^2)$, where r stands for the correlation of the two variables (see the close match in Figure 3, red and purple line). The ‘neglected’ MI is estimated by the difference between data MI and the Gaussian MI: $I_{neglected}(X, Y) = I(X, Y) - I_{Gauss}(r)$.

2.3. Statistical tests

For each session and each parcel pair, non-Gaussianity was tested at $p = 0.05$ by comparing the data MI against the MI distribution of the multivariate FT surrogates. Notably, the parcel-level results suffer from heavy multiple comparison problem. As we are primarily interested in the bivariate non-Gaussianity in general rather than in its specific allocation to particular parcel-pairs, we further use the results of the individual tests in a higher-level analysis.

On a session-level, the number of significant parcel-pairs in a given session was tested against the null hypothesis that the number of individual significant entries has a binomial distribution

$B(n = 4005, p = 0.05)$, where $n = 4005 = \frac{90(90-1)}{2}$ is the number of all parcel pairs and $p = 0.05$ is the single-entry false positive rate under the condition of pure Gaussianity of the bivariate distributions.

As it may be argued that the assumption of pair independence is too lenient, but the exact level of dependence is difficult to establish, we also carried out robust group level tests. The group level tests involved one-sample t-tests. The first was a test of the percentages of significant pairs in the 24 sessions against the null hypothesis of mean of 5% of significant parcel-pair results; the second was a test of the neglected information session-wise averages (across all parcel-pairs) against the null hypothesis of zero mean.

To explicitly control for any potential bias in the numerical generation of the surrogate distributions, the group-level tests were complemented by paired t-tests of the same two quantities against their values obtained from a control set of linear, ‘shadow’ datasets. For each session, a shadow dataset was created as a multivariate FT surrogate of the marginally normalized original dataset. Thus the shadow dataset preserved only the linear (correlation) structure of the original dataset of the respective session. Subsequently, each shadow dataset has undergone the same procedure as original data, including the initial normalization, generation of multivariate surrogates, computation of MI and statistical testing of pair-wise MI against surrogates. In this way, we have mimicked the full procedure of processing the original data using a ‘purely linear FC’ shadow dataset, accounting for any potential slight bias in the detection rate introduced by numerical properties of the algorithm. Apart from the percentages, we have also tested the mean neglected information from data versus shadow datasets by mean of a paired t-test. All the group-level tests used a $p=0.05$ significance threshold, but we also report the attained significance level.

2.4. Relevance for clustering

To explore the relevance of the studied non-Gaussianity for further data analysis, we have assessed the difference in clustering of anatomical parcels based on two functional connectivity methods: mutual information and linear correlation. To ensure a fair comparison, two preparatory steps were carried out here. Firstly, to prevent artificial mismatch due to merely different scaling of the two functional connectivity measures, we have

monotonously rescaled the correlation values by $r \mapsto -\frac{1}{2} \log(1 - r^2)$, effectively yielding an estimate of the ‘linear’ mutual information I_{Gauss} . Although transformed to share the scale of mutual information, this connectivity measure reflects only the linear correlation between the parcel time series, unlike the true mutual information computed by the equiquantization method. Secondly, due to random error in estimation of functional connectivity measures, we could expect some random level of disagreement between connectivity matrices and therefore between clustering based on different functional connectivity measures. To control for this effect, we have tested the specific effect of non-linearity included in the original data by comparing the agreement between MI and correlation matrices (represented by I_{Gauss}) computed from data with the agreement of MI and correlation computed from a multivariate linear surrogate dataset. This way we effectively test whether the clustering disagreement includes a contribution due to nonlinearity, or is a mere result of the imperfect estimation of the connectivity measures from finite-size samples.

2.4.1. Identification of clusters

Functional networks in the parcel time series were identified using a spectral clustering technique called Multiclass Spectral Clustering (Yu and Shi, 2003). In this technique clusters (corresponding to the sought functional networks) are found by optimizing the Multiway Normalized Cut objective, which is based on the Normalized Cut objective (Shi and Malik, 2000). Formally the applied method operates on weighted graphs where the vertices of the graph correspond to parcels and the edges of the graph are weighted by the functional connectivity between the respective parcels. Intuitively a good clustering is represented by a partition of all vertices to disjoint sets (clusters) which satisfies that the sum of weights of edges between vertices in different sets is small and the sum of weights of edges between vertices in the same set is large. Both of these properties are optimized at the same time by the applied algorithm - we refer the reader to (Yu and Shi, 2003) for details. Thus the optimal partition computed by the algorithm has clusters which are more internally connected and less connected with each other. This corresponds to a grouping of the parcels so that the time series inside one cluster are similar and at the same time series of parcels in different clusters are not similar.

2.4.2. Agreement of clusterings

After obtaining the two clusterings of the parcels from the two connectivity methods, we have to quantify their agreement. A natural procedure is to view the cluster assignment of a randomly chosen parcel as a discrete random variable, with integer values ranging from one to the total number of clusters. Then we can express the agreement between clusterings obtained for the two connectivity methods as the mutual information between the two discrete random variables corresponding to the two generated clusterings. As the mutual information between two variables is bounded by entropy of each of them, we further normalize the mutual information, dividing it by the minimum of the entropies of the two variables. Thus, if the clusterings are identical, then their normalized mutual information will be 1, while normalized mutual information of 0 corresponds to completely independent cluster assignments.

To statistically test the nonlinearity effect, we compared the agreements obtained for the original data and linearized ‘shadow’ data by means of a paired t-test. Visual inspection of the clustering of the group-average connectivity matrices revealed that the choice of $N=10$ clusters provides satisfactory results — see Figure 1 for the example clustering results. Therefore, $N=10$ clusters were used for processing each session. For robustness, the analysis was repeated for all other cluster counts in the range from $N=2$ to $N=20$ with $FDR < 0.05$ correction for multiple comparisons.

3. Results

3.1. Descriptive assessment

In descriptive terms, the data MI has proved very similar to the Gaussian MI (see Figure 2). In particular, averaging across all parcel pairs, the data MI ranged between 0.04 and 0.10 bits for different sessions, while the neglected MI was more than an order of magnitude smaller (0.0005-0.0068 bits). Nevertheless, the neglected MI was consistently positive, which was not the case for shadow datasets (ranging from -0.0007 to 0.0016 bits).

Independently of the strength of coupling, the data MI was moreover typically within the range of surrogate MI, as illustrated in Figure 3. Here, each blue dot corresponds to MI of one parcel pair; the surrogate distribution is represented by red (light blue, green) lines for mean (1st percentile, 99th

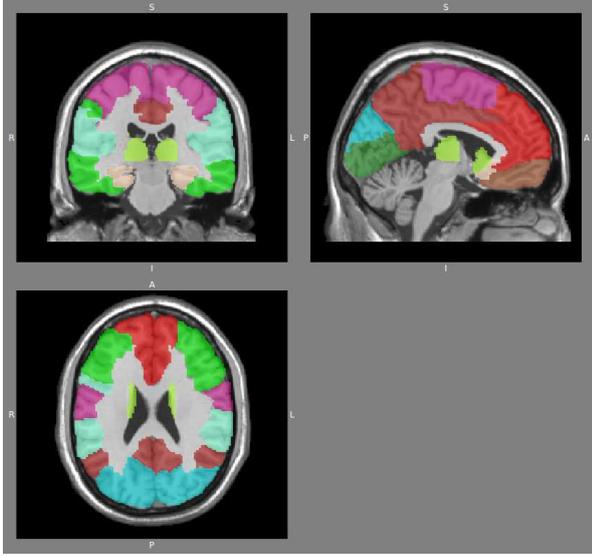


Figure 1: Clustering of the parcels for the choice of $N=10$ clusters based on average linear correlation matrix. The same colour is used for all parcel within a given cluster; an MNI template anatomical map is underlaid. The obtained clusters generally correspond to known functional networks.

percentile) of the surrogate distribution. Although the session with the most non-Gaussianity is depicted here, the distribution of computed MI for data and the corresponding shadow dataset (Figure 4) are almost indiscernible. Also, apart from the random error due to MI estimation from short time series, which is shared by data and shadow data, both scatters follow well the theoretical prediction of dependence of MI on linear correlation ($I_{Gauss} = -\frac{1}{2} \log(1-r^2)$, valid exactly under Gaussianity, purple line), which is closely approximated by the surrogate mean.

3.2. Statistical tests

The percentage of parcel pairs with significant non-Gaussianity was slightly elevated in all sessions above the 5% expected under the null hypothesis (ranging from 5.3 to 10.0% of significant pairs in different sessions). If all the parcel pairs were considered independent this would constitute significant percentage for all but 5 of the sessions considered (comparing to binomial distribution $B(n = 4005, p = 0.05)$, where $n = 4005 = \frac{90(90-1)}{2}$ is the number of all parcel pairs).

While the assumption of independence of the pairs might be too lenient, yielding false positive results, we also carried out group level tests that

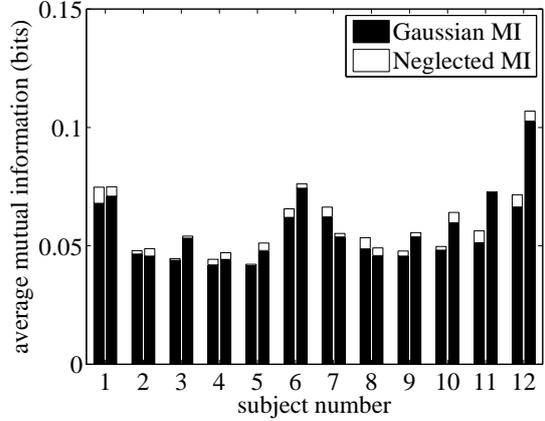


Figure 2: Comparison of the average Gaussian and neglected information. Each stackbar represents values for one session, averaged across all parcel pairs.

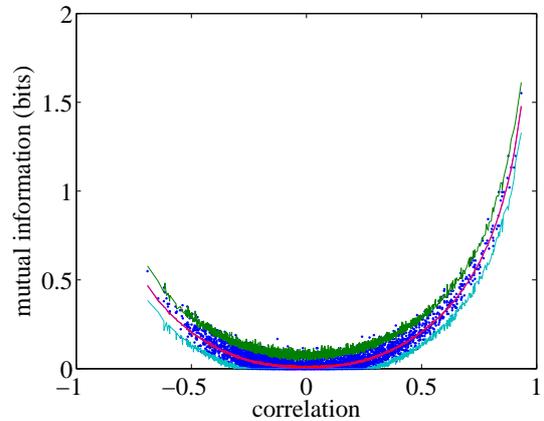


Figure 3: Mutual information as function of correlation in an example dataset. The session with the most non-Gaussianity is depicted. Each blue dot corresponds to MI of one parcel pair; red (light blue, green) lines correspond to mean (1st percentile, 99th percentile) of the surrogate distribution. The purple line shows the theoretical mutual information of an exactly Gaussian distribution with the given correlation $I_{Gauss}(r)$ - it is not well visible as it closely matches the mean of the surrogate distribution.

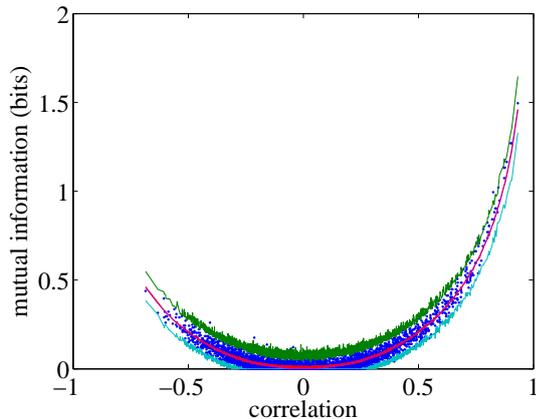


Figure 4: Mutual information as function of correlation in an example surrogate. Each blue dot corresponds to MI of one parcel pair; red (light blue, green) lines correspond to mean (1st percentile, 99th percentile) of the surrogate distribution. The session with the most non-Gaussianity is depicted.

confirmed the statistical deviation from Gaussianity. In particular, the counts of pairs with significant nonlinearity were significantly higher than the 5% expected by chance, as tested by one-sided t-test ($t = 6.95, df = 23, p < 10^{-6}$). Also, the neglected information averaged over parcel pairs was significantly biased above the mean zero value corresponding to full Gaussianity ($t = 8.52, df = 23, p < 10^{-7}$).

As described in subsection 2.3, we have also carried out another group-level tests that explicitly control for any potential biases in the surrogate method. The comparison of results from the original data with those from control (linear) ‘shadow’ datasets confirmed the detection of non-Gaussianity in the data. The counts of pairs with significant nonlinearity were significantly higher in data than similar counts obtained from shadow datasets, when compared on group level by means of a paired t-test ($t = 6.26, df = 23, p < 10^{-5}$). Also, the neglected information in data averaged over parcel pairs was positive for all sessions and on average had value 0.0029 bits while the neglected information in the shadow datasets fluctuated around zero with mean of 0.0006 bits. This difference was also clearly statistically significant ($t = 6.51, df = 23, p < 10^{-5}$).

3.3. Relevance for clustering

We tested whether the neglected information had a significant effect on the overall structure of

the connectivity, as captured in clustering of the parcels. $N=10$ clusters were computed for each session using the MI and rescaled linear correlation FC matrices and the agreement of the two clusterings was computed. This was carried out both for normalized data and their linearized ‘control’ versions. The agreements were generally quite high (0.80 ± 0.07 on the scale from 0 to 1) and did not differ significantly from those obtained from the ‘control’ linearized data ($t = 0.39, df = 23, p = 0.70$). This suggest that the mismatch between the clusterings based on the two measures was (similarly as in the case of the ‘control’ dataset) attributable generally to random error in estimation of the functional connectivity indices from finite size samples. Repeating the analysis for all cluster counts in the range $N = 2$ to $N = 20$ confirmed the observation of no significant difference ($p < 0.05$, FDR corrected).

4. Discussion

The presented study reveals that the bivariate dependence structure of the fMRI BOLD regional time series is captured very well by linear correlation. Indeed, average mutual information was only several percent higher than the mutual information in surrogate data that contained only the linear part of the dependence. This gives explicit and quantifiable argument for the intuitive choice of linear correlation as a measure of functional connectivity for fMRI time series.

Nevertheless, we have shown that there is a statistically significant contribution of non-Gaussian dependencies in the data, although the effect is so subtle that testing across many pairs or even across many sessions was needed to acquire sufficient power for such tests. The detection of non-Gaussian coupling is not surprising in the light of the fact that the dynamics of brain activity as well as the hemodynamic response of the vasculature contain many nonlinearities.

We have also carried out tests of the relevance of the neglected information for applications: the marginality of the difference between full mutual information and correlation as FC measures was confirmed in comparisons using clustering of the parcels – the disagreement between clustering obtained from mutual information in data and its correlation-based approximation did not exceed the disagreement due to random error in estimation of the FC.

For testing Gaussianity of the bivariate dependences, we have used a framework based on a comparison of mutual information evaluated on data and on surrogate timeseries. Below we discuss the motivation of the method choice from the variety of methods available in the wider context of multivariate normality testing. Importantly, the problem of testing multivariate normality is more complex than that of testing univariate normality and for principal reasons there is not a unique agreement in the statistical community on a single best method for testing multivariate normality. For a recent review of the topic we refer the reader to Mecklin and Mundfrom (2004). Many alternative tests for multivariate normality could therefore be considered, which may to some extent differ in the power against particular alternatives. However, the framework used in this paper has several important features that make it well suited for the current study. The first reason is practical; mutual information is used both for testing and direct quantification of the deviation from Gaussianity. Crucially, as mutual information is a widely used and theoretically grounded assumption-free measure of dependence, this quantification also has a general direct and intuitive interpretation. Secondly, many available multivariate normality tests assume independence of the samples. This is clearly not valid in the case of fMRI time series. The use of multivariate linear surrogates in our approach controls any bias due to temporal autocorrelation and sidesteps the need to estimate and correct for the reduced degrees of freedom of the test. Also, the test based on mutual information has the advantage of being consistent in the sense of being sensitive to virtually any departure from Gaussianity of the dependence. This advantageous property is not shared by many commonly used tests for multivariate normality such as those based on testing the deviation of skewness or kurtosis of the distribution from the values valid for normal distribution. Finally, the used framework allows to test the dependence structure itself, independently of the normality of the univariate marginal distributions. This is also not in general true for other multivariate normality tests, although adaptations may be possible.

Another topic deserving a discussion is the adopted multi-level testing strategy and its potential alternatives. As we mentioned, the binomial test used to assess the Gaussianity null hypothesis on session-level might be too lenient due to dependence among the 4005 lower-level tests. We have

avoided speculative estimation of the test interdependence level within each session and rather robustly tested the Gaussianity on the group level. Importantly, this was sufficient to statistically reject the Gaussianity null hypothesis. If further evidence on the level of individual parcel-pairs or sessions was required, the use of criteria for hypothesis testing such as the family-wise error rate and the false discovery rate to obtain more detailed and potentially more powerful tests could be considered. Nevertheless, their implementation may prove prohibitively cumbersome in the current context. In particular, these methods generally involve comparing the smallest (or each) p-value against the corrected significance level $p = 0.05/4005 \sim 10^{-6}$. Such a low p-value is clearly below the resolution of the sample histograms capturing the null hypothesis with 99 surrogates; the minimal required number of surrogates would be in the order of millions for each test, which is technically impractical to infeasible. Fitting the null distribution from smaller number of surrogates by e.g. imposing an assumption of approximate normality of the MI distribution may seem promising, but may be too crude and even slight departure from the true null distribution might crucially affect the true significance levels, particularly for the critical values of the order of 10^{-6} .

It is important to keep in mind that the observed deviations from Gaussianity might not reflect only a stationary non-Gaussianity in neuronal connectivity. In the presented framework, deviation from the null hypothesis could be caused also by nonstationarity of the signal. Vice versa, non-Gaussianity might lead to false detection of nonstationarity if the test assumed a linear Gaussian generating process as in (Chang and Glover, 2010). Technically, a precise isolation of these two effects is extremely challenging. A practical consideration allows to reasonably reconcile the two alternative interpretations: heavy nonstationarity could indeed increase the estimated non-Gaussianity, but it would also on its own invalidate the use of simple linear correlation, leading to the same conclusion about its suitability. On another note, we also should not forget that we are working on the level of fMRI BOLD signal rather than with neuronal activations - both the Gaussian and non-Gaussian contributions to FC are likely to be affected to some extent by non-neural sources of signal variation. Nevertheless, in the end the estimation of whether this would play in favor of or against the use of linear correlation (and how

much) seems to be entirely speculative at this point.

It can be argued that the amount of non-linear mutual information detected is likely to depend on the preprocessing of the data and the acquisition parameters. In this study, we have used standard acquisition and preprocessing, maybe with the exception of not using any explicit temporal low-pass filtering. The rationale for omission of this step was to avoid introduction of temporal averaging that would further diminish an already small nonlinearity in the data. The choice of working with several tens of gray matter parcels per hemisphere was motivated by approach of (Lahaye et al., 2003) as well as computational feasibility and allowing reproducibility and comparability across subjects and datasets. On the other side, incoherence of voxel-signals within the parcels might lead to losing some of the nonlinearity on the level of parcel-averaging, and therefore working on the level of temporal signal of spatial Independent Component Analysis (ICA) components or single voxel signal might in theory be more sensitive to nonlinearity. Exploring the dependence of the non-Gaussian contribution to FC on the time series extraction method is a subject of future work. Also, the analysis presented in this paper did not consider any time-lags. This is consistent with comparing to linear correlation as a measure of FC. The time lags would have to be included were we interested in assessment of causal or ‘effective connectivity’ measures (linear versus nonlinear) – this defines a natural follow-up of the recent study.

Apart from some minor technical aspects such as normalisation of the timeseries, our study differs from the previous probes into the potential of non-linear fMRI FC such as (Xie et al., 2008; Maxim et al., 2005; Lahaye et al., 2003; Deshpande et al., 2006) mainly in that we explicitly focus on the bivariate Gaussianity rather than the linearity assumption as the condition of suitability of use of linear correlation as functional connectivity index. The deviation from this condition allowed us to quantify the potential available for arbitrary ‘non-linear’ connectivity measures. This general interpretation is allowed by the use of a very general dependence measure – mutual information. In theory, this is able to capture virtually any form of statistical dependence. Of course, some practical limitations stem from the inevitably finite sample size, forcing us to summarize the results across parcel pairs and subjects. Last but not least, we provide an illustrative quantitative estimation of the deviation

from Gaussianity by means of the mutual information neglected by linear correlation, that should give a theoretical upper bound on any improvement to be made by an arbitrary nonlinear connectivity measure. We note that an earlier attempt towards quantification of the FC ‘nonlinearity’ was made in (Lahaye et al., 2003), who reported also the explained variances by the higher order (nonlinear up to order 5) terms in the predictor set within a linear regression framework – although these variances were not corrected for the effect of extra number of regressors in the model, and also the considered model could not capture a more general form of dependence. A similar approach as in Lahaye et al. (2003) was previously also applied to assessment of the non(linear) character of connectivity in EEG data by Pijn et al. (1990). The h^2 coefficient applied in that study expresses the proportion of variance in one variable explained by another variable using a polynomial fit; the aim was to provide more robust detection of interactions than simple linear correlation in case when the linearity assumption is violated.

Despite the differences in the aims and methodology, our observations agree qualitatively with the previous studies such as (Lahaye et al., 2003; Xie et al., 2008) in that we also reject the specific hypothesis of stationary multivariate linear Gaussian process as the structure of resting state fMRI signal.

For clarity, we stress again that the above mentioned Gaussianity condition, rather than mere linearity of the process, is the key assumption for suitability of use of linear correlation. To fully acknowledge this consider that as stated e.g. in (Bickel and Bühlmann, 1996), a linear stationary process $(X_t)_{t \in \mathbb{Z}}$ is usually defined by

$$X_t = \sum_{j=0}^{\infty} \phi_j \epsilon_{t-j}; (t \in \mathbb{Z}), \quad (2)$$

where ϵ_t is i.i.d with $E[\epsilon_t] = 0, E|\epsilon_t|^2 < \infty$ and $\sum_{j=0}^{\infty} \phi_j^2 < \infty$.

While the definitions of linear process may differ in details, in most cases they are general enough to include processes with *non-Gaussian* distribution, which are then not fully described by their correlation structure. This may lead to some confusion, pointed out in (Nagarajan, 2006), who showed that some widely used surrogate-based linearity tests such as those used in (Xie et al., 2008) are actually sensitive to non-Gaussianity. Careful examination

of the assumptions of these tests reveals that the null hypothesis there is that the data come from a process that is equivalent to a linearly filtered white *Gaussian* noise. Of course, if the general definition of linearity is complemented by the condition that the generating noise is Gaussian, such a linear process is indeed also Gaussian, with probability distribution fully characterised by the mean and covariance structure. Discussion of the potential causes of rejection of the null hypothesis of ‘linearity’ in the FT-surrogate framework by processes that are still linear in the general sense of having the form (2) is also provided in (Palus, 2007).

For completeness, we note that linearity is also often discussed as an alternative to nonlinear, potentially chaotic deterministic dynamical systems. In this context caution is warranted with the interpretation of many ‘chaotic’ characteristics such as fractional correlation dimension or Lyapunov exponents when the underlying system might be of stochastic (non)linear nature rather than deterministic (non)linear dynamical system, and particularly when short time series such as those acquired from fMRI are being analyzed.

To summarize, we have assessed the suitability of linear correlation as functional connectivity measure for fMRI time series by testing and quantifying the deviation from bivariate Gaussianity using mutual information. The results were as follows: firstly, the quantitative assessment revealed that the portion of mutual information neglected by using linear correlation instead of considering an arbitrary non-linear form of instantaneous dependence is minor. Secondly, formal group-level test revealed that the percentage of parcel-pairs with significant non-Gaussian dependence contribution is indeed above random. Overall we conclude that linear correlation of normalized data well captures the full functional connectivity: although existence of non-Gaussian contribution to functional connectivity can be established, practical relevance of non-linear methods trying to improve over linear correlation might be limited by the fact that the data are indeed almost Gaussian. This was documented by a quantitative estimate of mutual information neglected due to the use of linear correlation for typical resting state fMRI data, as well as by complementary comparison of results of clustering analysis.

Acknowledgements

This work has been supported by the EC FP7 project BrainSync (EC: HEALTH-F2-2008-200728, CR: MSM/7E08027); and in part by the Institutional Research Plan AV0Z10300504. DM was partly supported by the Research Foundation Flanders (grant FWO A 4/5 SDS 15387).

References

- Bickel, P., Bühlmann, P., 1996. What is a linear process? Proceedings of the National Academy of Sciences of the United States of America 93, 12128–12131.
- Chang, C., Glover, G.H., 2010. Time-frequency dynamics of resting-state brain connectivity measured with fMRI. *NeuroImage* 50, 81–98.
- Deshpande, G., LaConte, S., Peltier, S., Hu, X.P., 2006. Connectivity analysis of human functional MRI data: From linear to nonlinear and static to dynamic. *Medical Imaging and Augmented Reality* 4091, 17–24.
- van Dijk, K.R.A., Hedden, T., Venkataraman, A., Evans, K.C., Lazar, S.W., Buckner, R.L., 2010. Intrinsic functional connectivity as a tool for human connectomics: Theory, properties, and optimization. *Journal of Neurophysiology* 103, 297–321.
- Fox, M.D., Raichle, M.E., 2007. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat Rev Neurosci* 8, 700–711.
- Friston, K., 1994. Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping* 2, 56–78.
- Lahaye, P.J., Poline, J.B., Flandin, G., Dodel, S., Garnero, L., 2003. Functional connectivity: studying nonlinear, delayed interactions between BOLD signals. *NeuroImage* 20, 962–974.
- Laufs, H., 2008. Endogenous brain oscillations and related networks detected by surface EEG-combined fMRI. *Human Brain Mapping* 29, 762–769.
- Mantini, D., Perrucci, M.G., Del Gratta, C., Romani, G.L., Corbetta, M., 2007. Electrophysiological signatures of resting state networks in the human brain. Proceedings of the National Academy of Sciences of the United States of America 104, 13170–13175.
- Maxim, V., Sendur, L., Fadili, J., Suckling, J., Gould, R., Howard, R., Bullmore, E., 2005. Fractional Gaussian noise, functional MRI and Alzheimer’s disease. *NeuroImage* 25, 141–158.
- Mecklin, C.J., Mundfrom, D.J., 2004. An appraisal and bibliography of tests for multivariate normality. *International Statistical Review* 72, 123–138.
- Miller, K.J., Weaver, K.E., Ojemann, J.G., 2009. Direct electrophysiological measurement of human default network areas. *Proc. Natl. Acad. Sci. U. S. A.* 106, 12174–12177.
- Nagarajan, R., 2006. Surrogate testing of linear feedback processes with non-gaussian innovations. *Physica A* 366, 530–538. *Physica A - statistical mechanics and its applications*.
- Palus, M., 1997. Detecting phase synchronization in noisy systems. *Physics Letters A* 235, 341–351.
- Palus, M., 2007. From nonlinearity to causality: statistical testing and inference of physical mechanisms underlying complex dynamics. *Contemporary Physics* 48, 307–348.

- Palus, M., Albrecht, V., Dvorak, I., 1993. Information theoretic test for nonlinearity in time series. *Physics Letters A* 175, 203–209.
- Palus, M., Vejmelka, M., 2007. Directionality of coupling from bivariate time series: How to avoid false causalities and missed connections. *Physical Review E* 75, 056211.
- Pijn, J.P.M., Vijna, P.C.M., da Silva, F.H.L., Boas, W.V.E., Blanes, W., 1990. Localization of epileptogenic foci using a new signal analytical approach. *Neurophysiologie Clinique/Clinical Neurophysiology* 20(1), 1–11.
- Prichard, D., Theiler, J., 1994. Generating surrogate data for time series with several simultaneously measured variables. *Physical Review Letters* 73, 951.
- Shi, J., Malik, J., 2000. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22, 888–905.
- Stam, C., 2005. Nonlinear dynamical analysis of EEG and MEG: Review of an emerging field. *Clinical Neurophysiology* 116, 2266–2301.
- Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., Joliot, M., 2002. Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage* 15, 273–289.
- Xie, X.P., Cao, Z.T., Weng, X.C., 2008. Spatiotemporal nonlinearity in resting-state fMRI of the human brain. *NeuroImage* 40, 1672–1685.
- Yu, S.X., Shi, J., 2003. Multiclass spectral clustering, in: *ICCV '03: Proceedings of the Ninth IEEE International Conference on Computer Vision*, IEEE Computer Society, Washington, DC, USA. p. 313.
- de Zwart, J.A., van Gelderen, P., Jansma, J.M., Fukunaga, M., Bianciardi, M., Duyn, J.H., 2009. Hemodynamic nonlinearities affect BOLD fMRI response timing and amplitude. *NeuroImage* 47, 1649 – 1658.