Multi-Link Analysis: Brain Network Comparison via Sparse Connectivity Analysis

Alessandro Crimi[a,g,1], Luca Giancardo[a,b], Fabio Sambataro[d], Alessandro Gozzi[c], Vittorio Murino[a,e], Diego Sona[a,f]

[a] *Pattern Analysis and Computer Vision, Istituto Italiano di Tecnologia, Genova, Italy*
[b] *Center for Precision Health, School of Biomedical Informatics, University of Texas Health Science Center at Houston, USA*
[c] *Functional Neuroimaging Laboratory, Istituto Italiano di Tecnologia, Rovereto, Italy*
[d] *Department of Experimental and Clinical Medical Sciences, University of Udine, Italy*
[e] *Department of Computer Science, University of Verona, Italy*
[f] *Neuroinformatics Laboratory, Fondazione Bruno Kessler, Trento, Italy*
[g] *Institute of Neuropathology, University Hospital of Zürich, Switzerland*

Abstract

The analysis of the brain from a connectivity perspective is unveiling novel insights into brain structure and function. Discovery is, however, hindered by the prior knowledge used to make hypotheses. On the other hand, exploratory data analysis is made complex by the high dimensionality of data. Indeed, in order to assess the effect of pathological states on brain networks, neuroscientists are often required to evaluate experimental effects in case-control studies, with hundreds of thousand connections.

In this paper, we propose an approach to identify the multivariate relationships in brain connections that characterise two distinct groups, hence permitting the investigators to immediately discover sub-networks that contain information about the differences between experimental groups. In particular, we are interested in data discovery related to connectomics, where the connections that characterize differences between two groups of subjects are found, rather than maximizing accuracy in classification since this does not guarantee reliable interpretation of specific differences between groups. In practice, our method exploits recent machine learning techniques employing sparsity to deal with weighted networks describing the whole-brain macro connectivity. We evaluated our technique on functional and structural connectomes from human and mice brain data. In our experiments, we automatically identified disease-relevant connections in datasets with unsupervised and anatomy driven parcellation approaches using high-dimensional datasets.

Keywords:
Connectome, DTI, Tractography, fMRI, Diffusion MRI, Sparse Methods, mice, BTBR, Schizophrenia, Autism, Network Based Statistics

[1] Corresponding author at: Pattern Analysis and computer VISion department, Istituto Italiano di Technologia, Via Morego 30, 16163 Genova, Italy. Tel: +39 010 71781
E-mail: alessandro.crimi@usz.ch

Introduction

The analysis of brain networks, or connectomes, is a recent and exciting advancement in magnetic resonance imaging (MRI) which promises to identify new phenotypes for healthy, diseased or ageing brains (Sporns (2011)). A connectome is a mathematical description of the brain, which is conceived as a network, where brain areas (nodes) are connected by links (edges) (Sporns (2010)), and connections can be either given by white matter tracts between pairs of brain regions, or by an index of correlation of functional activity (Richiardi et al. (2011)). This allows for analysing the brain as a complex system of dynamically interacting components without explicitly relying on local activation or brain morphology.

Case-control studies and connectomics

Experiments with connectomes are typically designed by comparing a studied group with a control group in order to identify brain-network topological biomarkers relevant to the studied group (Rubinov and Sporns (2010)). Indeed, inter-group differences in some of these topological measures have been discovered for various neuropsychiatric disorders (Bullmore and Sporns (2009)), like Alzheimer's disease (Stam et al. (2009)), multiple sclerosis (He et al. (2009)), schizophrenia (Lynall et al. (2010); Cocchi et al. (2014)), stroke (Grefkes and Fink (2011); Bonilha et al. (2014)), major depression (Zeng et al. (2014)), etc. All these approaches use topological measures with statistical tests to assess their discrimination power in a univariate analysis framework. Alternatively, in a multi-variate framework, machine learning methods have been proposed to differentiate groups of subjects using topological measures (e.g., Iturria-Medina et al. (2011)). Surveys on graph-topological metrics using functional magnetic resonance imaging (fMRI) data and related clinical applications using structural features are given respectively in Varoquaux and Craddock (2013) and Griffa et al. (2013).

Local differences between connectomes

The main drawback of these approaches is the limited interpretability of graph statistics as they miss the local characterization of the groups in terms of differences in the connectivity, but rather employ global statistics which are difficult to be translated into clinical settings for local analysis. A method which allows insights on local connectivity patterns in case-control studies relies on Network-Based Statistics (NBS). In this approach, the connectivity between pairs of brain regions is tested for significance using univariate statistics for functional (Fornito et al. (2011)) and anatomical (Zalesky et al. (2011)) connectivity disturbances. Simpson et al. (2013) extended the NBS method using a permutation test based on Jaccard index at node level. While, Chen et al. (2015) enhanced NBS regulating the topological structures comprised. With the same aim, trying to identify discriminating regions between groups, Ng et al. (2016) and Gaonkar and Davatzikos (2013) proposed to analyse the weights resulting from trained support vector machines (SVM). In particular, in Ng et al. (2016) a projection of covariance estimates onto a common tangent space was carried out to reduce the statistical dependencies between elements. Then, these estimates were used in a SVM framework which uses its weights to find meaningful difference between two groups. Preliminary results showed that more advanced machine learning approaches based on SVM coupled with Riemannian/Grassmannian geometry can be used to discriminate groups of connectomes (Dodero et al. (2015)). Despite NBS and its extensions have been shown to outperforms other methods in comprehensive comparisons (Kim et al. (2014)), the identification of graph sub-networks is a pre-requisite which can limit the detected connections. Moreover, the threshold for the significance of the statistics has to be set, influencing considerably the results

(Baggio et al. (2018)). Searching for a more specific and localized information, van den Heuvel and Sporns (2011) proposed a sub-graph level analysis, with a specific emphasis on the potential functional importance of highly connected hubs ("rich-clubs"). Although the focus on rich-clubs is insightful, this method could leave out subtle differences between case-control groups which are not present in highly connected hubs.

Relation to previous methods

In this context, we are interested in data discovery related to connectomics, where the connections that characterize differences between two groups of subjects are found, and where maximizing accuracy does not guarantee reliable interpretation since similar accuracies can be obtained from distinct sets of features (Rondina et al. (2014)). To overcome the limitations of the univariate approaches which perform statistical tests on single connections mentioned in the previous subsection - and in particular to the most commonly used NBS (Zalesky et al. (2010)) - we use a multivariate bootstrap-like approach. Therefore, we propose a fully data driven method to identify relevant brain sub-networks in experiments with case-control design. Our approach aims at creating an hypothesis generation tool for connectomes investigations. The method is supposed to work equally well with functional and structural MRI data, and no prior knowledge about the type of connectivity is required, only examples of brain connectivity matrices of two groups are needed.

A similar method proposed by McMenamin and Pessoa (2015) implemented a two-layer dimensionality reduction technique based on principal component analysis (PCA), followed by quadratic discriminant analysis to identify clusters with altered connectivity at voxel level. However, when PCA was used for feature selection, the eigenvalues of the covariance matrix were used regardless the prior knowledge on the groups to be discriminated, and in doing so the resulting features may not be those which were really meaningful in terms of discrimination between groups. Conversely, our method directly performs a sparse version of linear discriminant analysis (LDA) that, by design, tries to optimize the feature selection step aiming at discriminating the groups. This allows the proposed method to be more specific in terms of identified discriminating connections.

More specifically, our model is based on an ensemble of sparse linear discriminant models allowing to find the networks' elements (a set of edges) able to consistently distinguish two groups, in the attempt to minimize the subset of selected connectivity features and simultaneously maximize the difference between the groups, jointly using the selected features. Essentially, the system acts as a *filter* removing the elements that are not useful to discriminate between the groups.

Other methods have already used sparsity to estimate relevant connections (Huang et al. (2010); Lee et al. (2011); Gramfort et al. (2013)). However, these methods did not focus on finding the discriminant connections between groups while performing the sparse selection. They use sparsity to reduce the number of connections regardless on the inter-class discrimination.

Paper overview

Owing to the sparsity principle driving the learning method combined with the statistical robustness of ensemble methods, our multivariate approach can scale up with the number of analysed connections, even when employing a limited number of whole-brain connectivity matrices. By virtue of being multivariate, this approach can identify brain sub-networks whose edges combination can characterize the differences between the connectomes but taken independently

cannot. Moreover, the method does not have to rely on covariance matrices. It just needs an index describing the strength of connectivity between the areas in terms of correlation, similarity, dissimilarity or other metrics. For example, in case of structural connectivity the matrix can be determined counting the number of connections between the areas.

We validated the approach on three real datasets. In a first experiment, we used the structural connectivity, based on the tractography extracted from diffusion tensor imaging (DTI). Specifically, we compared a group of acallosal BTBR mice (a well-characterized model of autism) with a group of control normocallosal and normosocial C57BL/6J mice (Sforazzini et al. (2014a); Squillace et al. (2014)). Performing this experiment with a simple and well known connectivity dysfunction, without the use of any prior anatomical parcellation to avoid any prior bias, we empirically validated the approach, which was able to retrieve the expected dissimilarity between the two groups

A further experiment was conducted on functional connectivity matrices from a publicly available dataset of patients affected by schizophrenia (Zalesky et al. (2010)). The final experiment was carried out on a larger dataset of attention deficit hyperactivity disorder (ADHD) children compared to typically developing (TD) children. In all cases, our method successfully detected inter group differences relevant to the medical condition investigated.

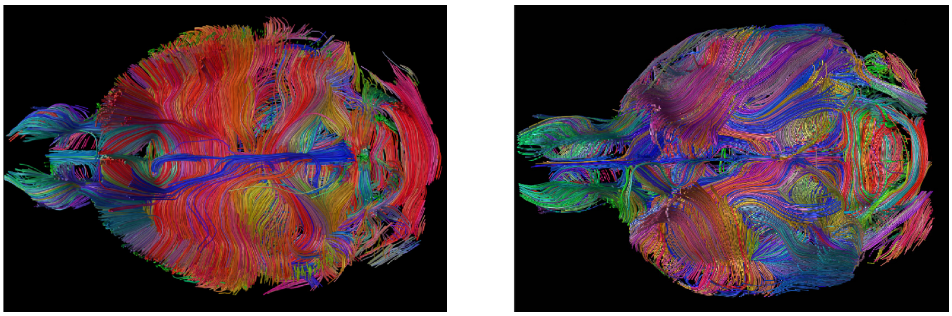

(a) (b)

Figure 1: Example of axial section tractography of (a) a normo-callosal C57BL/6J control and acallosal BTBR (b) mouse respectively, where the different anatomical structures are apparent but difficult to understand. In particular, the lack of corpus callous in (b) is visible.

Methods and Data

This section first describes the two types of data used to test the proposed method: a mice DTI dataset with high dimensionality, and two publicly available human fMRI datasets. Afterwards the pre-processing and the proposed computational model for discriminating patterns in whole-brain analysis are described.

Data

Mouse Structural Connectivity Data

The mice cohort was composed of two groups of 22-26 weeks old male subjects: BTBR T+tf/J mice which share analogies to all diagnostic symptoms of autism and characteristic functional

and structural features of the brain (Ren et al. (2007); Dodero et al. (2013); Fenlon et al. (2015)), and C57BL/6J mice which are characterised by normal sociability and represent the control group. The animal preparation protocol is described in Sforazzini et al. (2014b) and Dodero et al. (2013). Fig. 1 depicts an example of the expected difference between the BTBR and C57BL/6J mice groups. In particular, BTBR mice lack the corpus callosum differently from the C57BL/6J mice.

Briefly, brains were imaged inside intact skulls to avoid post-extraction deformations. Ex-vivo high-resolution DTI and T2-weighted images were acquired on paraformaldehyde-fixed specimens with a 7 Tesla Bruker Pharmascan MRI scanner (Billerica, MA, USA). T2-weighted MR anatomical images were acquired using a RARE sequence with the following imaging parameters: TR/TE = 550/33 ms, RARE factor = 8, echo spacing 11 ms, and a voxel size of 90 μm isotropic. DTI volumes were acquired using 4 scans at b0 and 81 scans with different gradient directions (b=1262 s/mm^2), with resolution 130×130 μm^2, using a 4-shot EPI sequence with TR/TE = 5500/26 ms. Anatomical and DTI sequences were acquired sequentially at the same centre with the same scanner.

This dataset is used to show that the algorithm is able to identify difference between the groups which are expected to be found as a proof of concept.

Human Functional Connectivity Schizophrenia Data

Functional connectivity was firstly investigated on a previously described resting-state fMRI dataset used in a study aiming at identifying differences in brain networks (Zalesky et al. (2010)). This dataset has been chosen to allow direct comparison to the state-of-art method NBS (Zalesky et al. (2010)). 12 patients with schizophrenia (all males, mean age 32.8 ± 9.2 years) and 15 controls (all males, mean age 33.3 ± 9.2 years) participated in this study. The two groups were matched for age, gender and years of education. Patients were diagnosed according to standard operational criteria in the Diagnostic and Statistical Manual of Mental Disorders (American Psychiatric Association et al. (1980)). All subjects underwent blood-oxygenation-level dependent (BOLD) fMRI scanning at rest. All patients were treated at the time of the scan but did not receive any medication on the day of scanning. The functional data of these individuals were acquired by using a 1.5 Tesla GE Signa scanner (General Electric, Milwaukee, WI) with the following parameters for BOLD acquisition: TR/TE 2/40 s, flip angle 70 degrees, voxel size $3.05 \times 3.05 \times 7$ mm, and 512 frame/volume. These data are publicly available (Zalesky et al. (2010)).

Human Functional Connectivity ADHD Data

Functional connectivity was also investigated on a larger resting-state fMRI dataset comprising ADHD and TD subjects (Castellanos et al. (2008)). In particular, we used the publicly available New York University Child Center dataset[2], which is the main cohort of this study. The dataset comprised 95 ADHD subjects (67 male and 28 female, mean age 11.4 ± 2.7) which were either inattentive, or hyperactive or both, and 92 healthy TD (45 male and 47 female, mean age 12.4 ± 3.1) which represents the control group.

The fMRI volumes were acquired with a Siemens Allegra 3T, with TR/TE 2000/15 ms and voxel size $3 \times 3 \times 3$ mm.

[2]http://umcd.humanconnectomeproject.org

Methods

Mouse Dataset Processing and Encoding

Deterministic tractography was performed on the DTI volumes after eddy current corrections, by using the Fiber Assignment by Continuous Tracking (FACT) algorithm (Mori et al. (1999)). Fibres were reconstructed in the original volumes following the 2nd-order Runge-Kutta integration scheme (Lazar et al. (2003)) starting from the centre of each voxel and following the main direction of the tensor. The tracking was stopped when the fibre made a sharp turn ($> 35°$) or entered a voxel with fractional anisotropy (FA) < 0.15.

To allow inter-subject comparisons, registration matrices to a common space were computed for each subject by using affine transformation (12 degrees of freedom). The obtained registration matrices were then applied to the endpoints of each fibre. This allowed the tractography algorithm to work on the original volume space without warping the tensors.

To enable a purely data-driven inter-group comparisons without the use of anatomical priors, the brain volumes were split into 3D cubes of size $1 \times 1 \times 1$ mm, without considering any atlas. Each cube was a node in the graph and the connectivity matrix was built counting the fibres starting and ending into two distinct cube elements of the grid, avoiding the inclusion of u-fibers. This resulted in defining 42,704 edges.

The advantage of this approach was that the result of the proposed analysis method was nearly independent from the size of parcellation. Indeed, not considering the anatomy nor the physiology of the brain might result in bundles of fibers split into "sub-bundles" connecting adjacent cubes. However, if there is a difference between the two groups it is retrieved for all sub-bundles, hence the overall bundles are then reconstructed. Yet the choice of using a fine grid or an atlas is arbitrary.

Schizophrenia Dataset Processing and Encoding

Brain volumes were already pre-processed according to the steps indicated in Zalesky et al. (2010), and made available on-line as precomputed functional connectivity matrices. Briefly, data were pre-processed by using the Cambridge Brain Activation (CamBA) software[3]. All volumes were motion corrected, filtered with a pass-band leaving only frequencies in the range $0.03 < f < 0.06$ Hz, and skull stripped. Brain volumes were also registered by using affine transformation (12 degrees of freedom). 74 regions of interest (nodes) were selected following Zalesky et al. (2010) within the automated anatomical labelling (AAL) atlas (Tzourio-Mazoyer et al. (2002)). Regions for which node-averaged time series could not be accurately estimated were excluded. In particular, olfactory cortex and gyrus rectus were removed because of the signal drop-outs.

The average signal within each node was calculated. Fluctuations of nuisance signals of no interest were reduced via linear regression against reference time series extracted from the seed regions defined by the white matter and cerebrospinal fluid. The functional connectivity was defined for each subject by temporal correlation among the time series resulting in a 74×74 connectivity matrix.

ADHD Dataset Preprocessing and Encoding

This dataset has been pre-processed as described in Colby et al. (2012), and the final connectivity matrices are publicly available. In brief, resting-state fMRI data were preprocessed following

[3] http://www.bmu.psychiatry.cam.ac.uk/software/

these steps: Removal of first 4 EPI volumes, slice timing correction, motion correction, and then applying the regressors for WM, CSF, motion time courses and a low order polynomial detrending. A band-pass filter of $0.009 < f < 0.08$ Hz was also applied. Lastly, the data were blurred using a 6-mm Full Width at Half Maximum Gaussian filter. The functional region of interests were obtained using the method described in Craddock et al. (2012) for 200 areas.

Multi-link Analysis (MLA)

The interpretation of differences in brain networks is not always straightforward given individual variability and the high dimensionality of data (Sporns (2012)). Moreover, the internal structure of the brain connectivity with cross-relationships and dependencies in the feature space (the edges) may prevent a full retrieval of groups' differences using univariate analysis. Machine learning and dimensionality reduction techniques are designed to solve these issues, and hence these methods are a natural choice for addressing this discrimination task. We propose a two-stage feature selection process. At the first stage a classifier reinforcing sparsity is employed to select discriminant features, then only features which are consistent across dataset are kept.

An approach simultaneously implementing both techniques in a common sparse framework is *sparse logistic regression*, which has been already used to select relevant voxels for decoding fMRI activity patterns (Yamashita et al. (2008); Ryali et al. (2010)). Alternatively, in case of Gaussian-distributed data, the well known *linear discriminant analysis* has been extended to the sparse case with the *sparse discriminant analysis* (SDA) model (Clemmensen et al. (2011); Witten and Tibshirani (2011)). In particular, the method by Clemmensen et al. (2011) implements the elastic net regression with the ℓ_1-norm on the feature weights that indirectly sets the number of selected features. In all our experiments we resorted to this SDA model with the assumption that data in each group have a Gaussian distribution.

For all the experiments, the connectivity matrices are vectorized and ordered as rows in a $n \times p$ data-matrix $\mathbf{X}$, with n being the number of observations and p their dimensionality. The corresponding classification of objects is encoded into the $n \times K$ indicator matrix $\mathbf{Y}$, where each cell $\mathbf{Y}_{ik}$ indicates whether observation i belongs to class k. The SDA proposed by Clemmensen et al. (2011) then finds the discriminant vectors β_k for each class k and the vector of scores θ_k by the convex optimization given by the following regularized linear discriminant formulation

$$\begin{cases} \min_{\beta_k,\theta_k} \|\mathbf{Y}\theta_k - \mathbf{X}\beta_k\|^2 + \eta \|\beta_k\|_1 + \gamma\, \beta_k^T \mathbf{\Omega}\, \beta_k, \\ \\ \text{subject to } \frac{1}{n}\theta_k^T \mathbf{Y}^T \mathbf{Y}\theta_k = 1, \\ \qquad\qquad \theta_k^T \mathbf{Y}^T \mathbf{Y}\theta_l = 0 \;\forall l < k. \end{cases} \tag{1}$$

where $\mathbf{\Omega}$ is an arbitrary positive definite matrix, which allows to calculate a smooth discriminant vectors β_k even if the number of samples is smaller than the number of features ($n \ll p$). In our experiments we used $\mathbf{\Omega} = \mathbf{I}$ which makes the formulation an elastic net problem. The non-negative parameters η and γ control respectively the ℓ_1 and ℓ_2 regularization.

The advantage of the proposed sparse method is its capability of managing high-dimensional data thanks to the ℓ_2 regularization. Moreover, the ℓ_1 regularization term allows the model to select a small subset of features for the linear discrimination. This might result in a loss of predictive power while however reducing the over-fitting problem. In contrast, the ℓ_2 penalty term enjoys the grouping effect property, i.e., it works keeping small and comparable the weights of correlated predictors (Zou and Hastie (2005)). Moreover, ℓ_2 penalty term is much better at minimizing the prediction error than ℓ_1 regularization. As a result, their combination allows

to determine a good trade-off between an optimal classifier and a minimal selection of relevant predictors.

Although it is acknowledged that this type of model can be affected by the choice of η which represents the advantage and limitation of the method, indeed we noticed that the algorithm was satisfactorily discriminating the two classes on a wide range of η values. In this work we were mostly interested on discriminant features rather than finding an optimal classification. However, the results are shown using the values which allow better accuracy estimated in cross-validation manner and using the minimum number of features according to compactness criterion. The parameter γ was not optimized for the same reason and it was fixed as 10^{-6} to guarantee a minimal regularization. Lastly, the parameter η can also be reformulated as the desired number of variables selected by the model. In the following we will refer to this number as α instead of η. We address the reader to Zou and Hastie (2005) for a description of the relation between η and α, and further details on the algorithm are given in the Appendix.

Although this model is very powerful in determining small and good subset of features allowing to linearly discriminate the classes, it suffers from a stability problem (Meinshausen and Buhlmann (2010)), i.e., small changes in the data can drastically change the result of a single run. To cope with this stability issue, in order to improve the robustness of SDA, we have introduced a *leveraging* principle exploiting the ensemble of low-stability algorithms to produce a more stable feature selection (Giancardo et al. (2012)). In particular this leveraging principle consisted in a bagging-like approach without randomization during the bootstrap. Indeed this combines the output of multiple weak learners to create a stronger one (Meir and Rätsch (2003)). Boosting can be seen as a specific instance of leveraging algorithm (Mason et al. (2000)).

In our specific case, the SDA classifier was trained with a balanced leave-one-out approach, i.e.: in turn we removed one sample, we trained the model on the remaining sample, and tested only the sample that were left out. The process was repeated for all samples generating a cross-validation statistics. This ended up in an ensemble of models each one with a subset of "relevant" features (connections), selected so to maximize the discrimination between the two groups.

Statistics over the ensemble of all models were then refined by occurrence validation, where only features which were frequently selected during cross-validation were retained, i.e., features occurring in less than a pre-defined percentage of runs were discarded. In all our experiments reported below this threshold was determined as half the number of subjects in the corresponding dataset.

The weights of the selected features β_k were then used to evaluate their importance, namely, the higher the weight, the more important the feature is.

Results

In this section we report the results for the experiments with high-dimensional structural connectivity with murine data, and two datasets of human functional connectivity collected to study respectively *schizophrenia* and *autism*.

Mice Structural Connectivity Data

In order to prove the discriminative power of our approach, we tested its ability to correctly distinguish the structural connectomes of two groups of mice (C57BL/6J and BTBR) characterized by previously described white matter alterations, i.e., the presence/absence of the two major neocortical intra-hemispheric tracts: the corpus callosum and the dorsal hippocampal commissure (Sforazzini et al. (2014b)). Being the structural alteration in the BTBR mice well known,

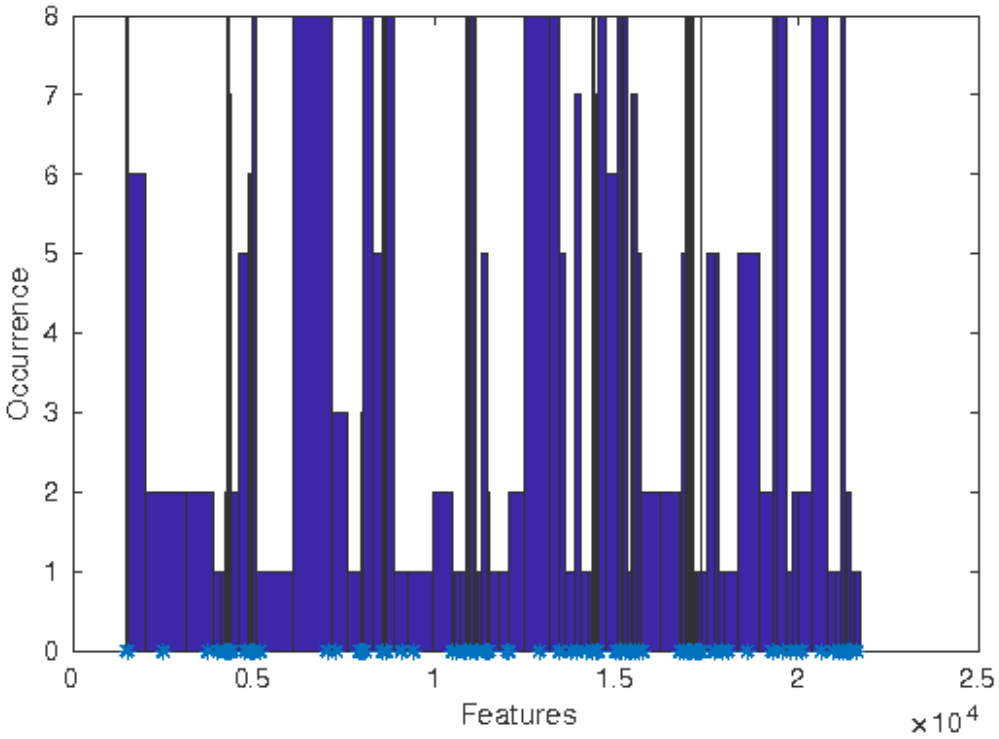

Figure 2: Histogram describing the occurrences of features selected for the mouse experiment as discriminative over all runs of the model within the ensemble framework. Higher values indicate those edges that characterize the differences between BTBR and control mice in many classifiers. This histogram is used to select a sub-set of "relevant" features. Namely, the most frequent features highlighted by the histogram are kept.

this dataset is used to validate the proposed method. Indeed, the BTBR mice model represents a ground truth of expected differences between the two groups. Over and above, more than the discrimination between the groups, we are interested in empirically assessing the ability of our approach to correctly identify white matter tracts differences in the two groups.

Indeed, by using the proposed algorithm, the model correctly classified all samples in a cross-validation schema, and structural differences - as the lack of corpus callus - were found as expected from literature.

To this aim the proposed approach returns a statistics of the relevance of features, by counting the amount of occurrences of the features selected by the ensemble of models. Figure 2 shows the occurrence of the detected features for the experiment with mice structural connectomes, some of which are present in all the runs, indicating a strong relevance for the problem at hand. Interestingly, the edges identified by the algorithm showed the expected characteristic features of the BTBR strain, including the agenesis of the corpus callosum and the presence of rostral-caudal rearrangement of white matter. Figure 3 shows how our algorithm (MLA) and NBS identify the parts of the corpus callosum which are known to be missing. This experiment confirms that our new approach and NBS are able to identify the acallosal connections in the BTBR models.

NBS and MLA select discriminative features in different ways. NBS performs univariate t-

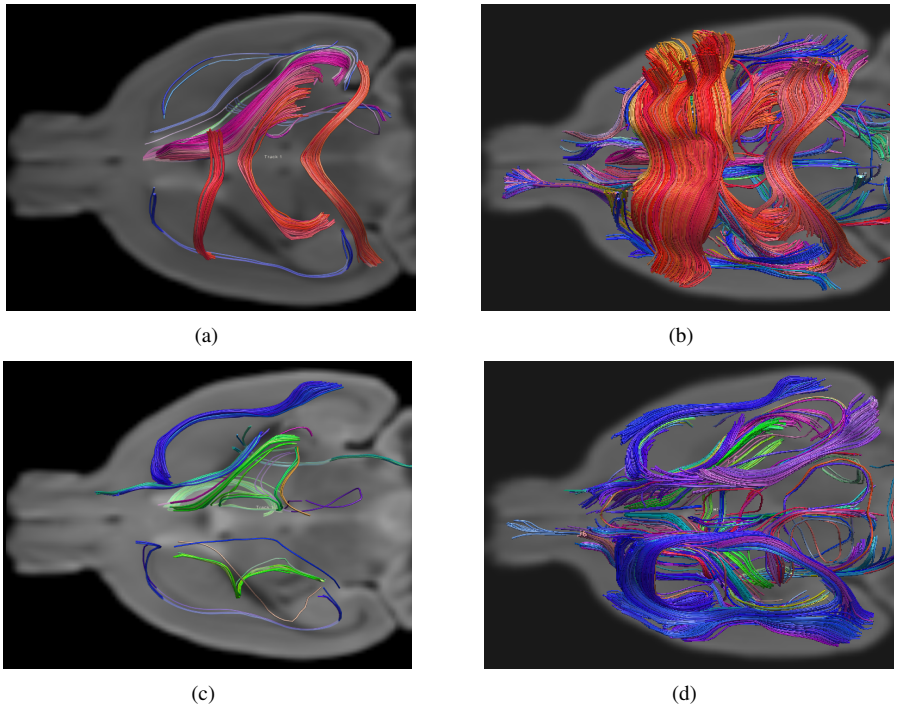

Figure 3: Graphical representation of the most significant features charactering the structural connectome of the two populations: the axial views of a randomly selected subject from the C57BL/6J control population (a) using our algorithm ($\alpha = 50$) and (b) the NBS algorithm using as a threshold p-value 0.01. While (c) and (d) are the axial views and of a randomly selected subject from the BTBR population respectively for our algorithm and NBS. As expected BTBR mice show a lack of corpus callosum and hippocampal commissure and an increased intra-hemispheric ipsilateral connectivity.

tests among the features while MLA performs a sparse multivariate regression.

The whole analysis from raw DTI data to tracts selection of the 10 subjects, by using Matlab Mathworks 2014, took less than 40 minutes on a 2.6 GHz machine with 4GB of RAM. However, the five rounds of MLA analysis required only less than 1 sec (with ℓ_1 parameter $\alpha = 50$ estimated by cross-validation which also gave 100% accuracy).

Human Functional Data

We tested the algorithm also on a publicly available dataset based on human functional MRI recorded from patients with schizophrenia and healthy controls (Zalesky et al. (2010)) and ADHD against TD subjects. By using the parameter $\alpha = 30$ estimated by cross-validation in the schizophrenia experiment, the algorithm highlighted significant differences between the case and control groups, which are reported both in Table 1 and 2 with NBS p-values, and in the histogram of features occurrence in Figure 7 respectively for the Schizophrenia experiment. Both our algorithm and NBS evidenced similar disconnected functional sub-networks in the schizophrenia group, namely in the fronto-temporal, parietal and occipital regions. Furthermore, some connections such as those between the precentral gyrus and the fusiform gyrus, the frontal mid-orbital gyrus and the talamus, and between the putamen and pallidum, would have been discarded by

Table 1: Functional connections differentiating patients with schizophrenia from controls identified by MLA (our algorithm). Pairs of source and target regions and p-values of the univariate t-test computed on NBS (Zalesky et al. (2010)) are reported. "Not detected" (N.D.) means not significant difference between the two areas.

#	Region 1	Region 2	p-value NBS
1	Frontal-Mid-Orb-R	Frontal-Inf-Tri-L	0.04
2	Precentral-R	Fusiform-R	N.D.
3	Rolandic-Oper-L	Postcentral-L	0.0001
4	Putamen-L	Pallidum-R	N.D.
5	Frontal-Mid-Orb-R	Thalamus-R	N.D.
6	Occipital-Mid-R	Thalamus-R	0.04
7	Frontal-Inf-Orb-L	Heschl-L	0.0004
8	Heschl-L	Temporal-Sup-L	0.0004

NBS as having non-significant p-values at univariate level as shown in Table 1. To rule out the hypothesis that this difference was due to artefacts of the proposed method, we further investigated whether the non significant connections at univariate level have any correlations with the significant connections. These connections were related to the connection Frontal-Inf-Orb-L/Hescl-L with an R^2 score of 0.7 of regression (which has a value of 1 in case of a perfect correlation). Suggesting that those connections discarded by NBS are indeed meaningful. As an example of this correlation, Figure 4 shows the Pearson correlation between Frontal-Inf-Orb-L/Hescl-L and one of those connections with non significant univariate statistics, specifically, Precentral-R/Fusiform-R for all subjects. Although it cannot be ruled out that some of these correlated features capture noise in some associated variables, the figure highlights the presence of a correlation between these two functional connections. In particular, each point in the graph identify the two weights for one subject and the correlation has been computed for each group separately.

In the experiment with the ADHD data, the cross-validation found the optimal value for the the MLA algorithm as $\alpha = 9$, which highlighted 7 discriminant connections across the groups with an accuracy of 66%, but the NBS did not produced any significant value. This might be due to the fact that the first key step of NBS is to identify candidate subnetworks which are then tested for their relevance using permutation testing. These candidate subnetworks are only selected if the nodes are explicitly connected between each others. Connectomes derived from data with high dimensional parcellations (as the used dataset with 200 areas) are more likely to have areas that are not connected, either because of low signal or for over parcellation. This is a problem for methods expecting a connected graph, like NBS, but it is not for our approach that does not have any prior on the types of connectivity expected. The connections among the areas detected by the proposed algorithm are also reported graphically in Figure 5 and 6 respectively for the schizophrenia and ADHD experiment. The schizophrenia and ADHD samples analysis took respectively less than 1 second and about 30 seconds on 2.6 GHz machine with 4GB of RAM. The classification task on the schizophrenia dataset had an accuracy of 84% in a cross-validation setting.

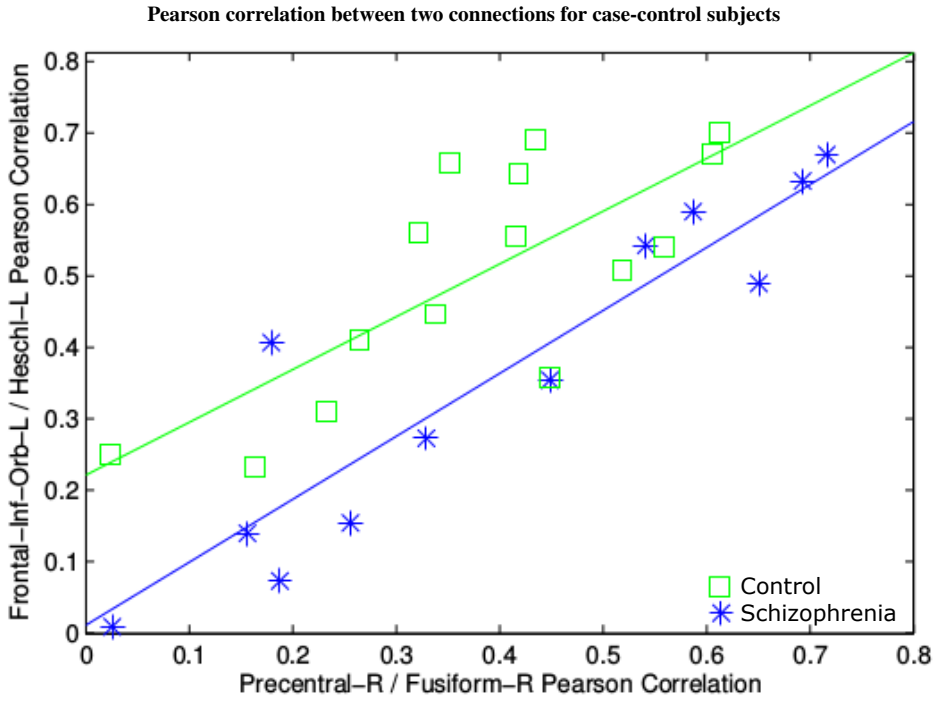

Figure 4: Scatter plot of the relationship between Frontal-Inf-Orb-L/Hescl-L and Precentral-R/Fusiform-R connections in patients and controls. Each element (either star or square) represents the correlation features for the same subject. Pearson's correlation between these two connections was computed separately in each individual. Blue stars represent the patients with schizophrenia sample, while green squares are control sample. Best-fit line are indicated for each sample. While t-test rejected the significance of the difference between the diagnostic group, MLA identified each individual as belonging to a specific group as highlighted by correlation line.

Discussion

The proposed method performs a global multivariate analysis characterizing local differences between networks. As this method is based on sparsity principles, it is particularly suited for those experiments with high-dimensional data and small sample size. Moreover, the analysis based on multivariate statistics allows to retrieve sub-networks based on feature dependencies. The limitation of NBS in detecting univariate differences is visible in the experiment with human functional data. In fact, the proposed algorithm detects some connections which are very often selected by the ensemble of learners, as seen in the histogram in Figure 2 and 7, but if considered with the univariate analysis, some edges are discarded as producing non-significant p-values (e.g., see rows 2, 4 and 5 in Table 1 for the schizophrenia experiment). Nevertheless, both MLA and NBS gave similar results confirming the initial hypothesis given by the anatomical differences.

The stability of the selected features is an important characteristic of the algorithm. As assessed empirically, by increasing the value of α, only new features are added without removing the previous ones. The parameter α does not need to be set automatically, but actually it represents the strength and limitation of the method. In fact, despite in the reported experiments

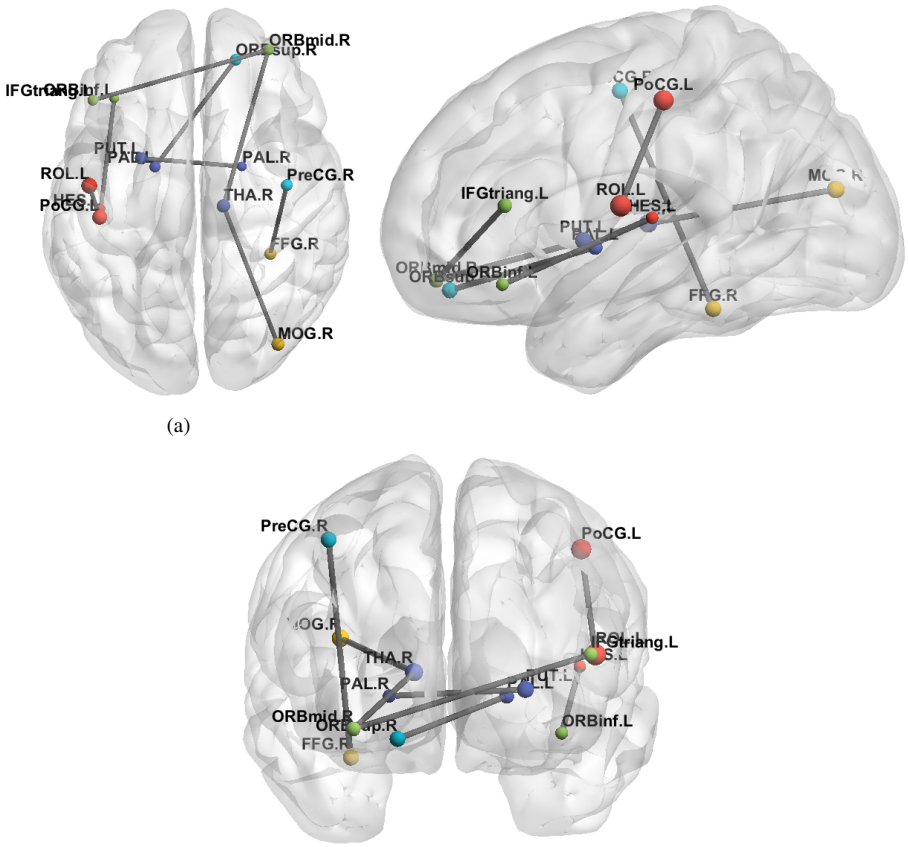

(a)

Figure 5: Functional connections differentiating patients with schizophrenia from controls with MLA (our algorithm) using $\alpha = 30$. From left to right, axial (a), sagittal (b), coronal (c) views of the brain indicate significant connections for at least one method are depicted. Each line represents a specific functional connection. For details on the statistics and acronyms, see Table 1.

the used α was estimated by cross-validation, the sensitivity of the algorithm can be manually adjusted trough this single parameter which allows the neuroscientists to decide how strong the class characterisation should be and it is directly related to the number of connections shown. A similar user-guided approach with methods based on sparsity has been previously described (Huang et al. (2010); Lee et al. (2011)). When MLA was applied to the acallosal BTBR mice, a mouse model of autism (Sforazzini et al. (2014a); Squillace et al. (2014)), as shown in Figure 3, the tracts detected as discriminant were those with altered white matter connectivity in BTBR mice with respect to control mice. These results are in line with previous results in literature, including the lack of corpus callosum and hippocampal commissure (Wahlsten et al. (2003); Ren et al. (2007); Fenlon et al. (2015)), and increased intra-hemispheric ipsilateral connectivity (Dodero et al. (2013); Meyer and Röricht (1998)), also observed in human patients with autism spectrum disorder (ASD) (Frazier and Hardan (2009); Casanova et al. (2011)). This demonstrates that the algorithm is able to identify the known differences between groups.

Table 2: Functional connections differentiating patients with ADHD from TD individuals. Pairs of source and target regions and p-values of the univariate t-test computed on NBS (Zalesky et al. (2010)) are reported. "Not detected" (N.D.) means not significant difference between the two areas. ADHD= Attention-Deficit/Hyperactivity Disorder, TD= Typically developed.

#	Region 1	Region 2	p-value NBS
1	Temp-Pole-L	Inf-Temp-Gyrus-posterior-division-L	N.D.
2	Temp-Fusiform-Cortex-anterior-division-L	Temp-Pole-L	N.D.
3	Frontal-Orb-Cortex-L	Supramarginal-Gyrus-posterior-division-L	N.D.
4	Temporal-Pole-L	Supramarginal-Gyrus-posterior-division-L	N.D.
5	Supramarginal-Gyrus-posterior-division-L	Parahippocampal-Gyrus-anterior-division-L	N.D.
6	Cerebellum-Vermis VI	Inf-Occipital-Cortex-R	N.D.
7	Middle-Temp-Gyrus-anterior-division-R	Lateral-Occipital-Cortex-inferior-division-L	N.D.

Similar results were obtained with the functional dataset on schizophrenia. Altered brain connectivity both at the microstructural and macrocircuitry levels has been described in this disorder in particular in the cortical areas (Pettersson-Yeo et al. (2011); Karbasforoushan and Woodward (2012)). As depicted in Figure 5, we found widespread connectivity cortical and sub-cortical differences in patients with schizophrenia compared to controls, some of which could have been missed by using the NBS as shown by the p-values in Table 1. Specifically, the inter-hemispheric connections of the frontal lobe as well as frontal sub-cortical connections differed. Reduced inter-hemispheric connectivity has been identified early in the course of the disorder, it correlates with severity of the disorder, and it is independent of drug treatment (Mwansisya et al. (2013)). Altered cortical-sub-cortical connectivity has been frequently reported in schizophrenia (Salvador et al. (2010)) and this is consistent with the limbic hyperactivity associated with the positive symptoms of schizophrenia. Also, local connectivity disturbances of the Rolandic operculum and postcentral gyrus have also been reported (Pu et al. (2014)). Furthermore, disturbed fronto-temporal connectivity has been associated with the genetic risk for schizophrenia (Winterer et al. (2003)), and it has been associated with positive and negative symptoms by converging imaging techniques (John et al. (2009)). In particular, fronto-temporal connectivity between Heschl's gyrus and fronto-parietal region could be associated with auditory hallucinations and is increased in patients with chronic auditory hallucinations (Shinn et al. (2013)). Consistently, DTI studies showed also disruptions of tracts connecting frontal, temporal and occipital regions (Ellison-Wright and Bullmore (2009)), thus suggesting that these structural alterations may contribute to functional connectivity changes in schizophrenia.

The interesting aspect of the proposed method was the capacity to retain several connections which would have been discarded by NBS, due to non-significant p-values as shown in Table 1. As it was shown, those connections discarded by NBS but not by the proposed method are correlated to significant connections, and used in a multivariate setting can help better to discriminate groups. Despite our experiments were conducted with the same small dataset previously used by Zalesky et al. (2010), to allow comparison with their results and the proposed method showed to be more robust to small sample size issues.

Regarding the discriminant connections detected by the MLA algorithm for the ADHD dataset,

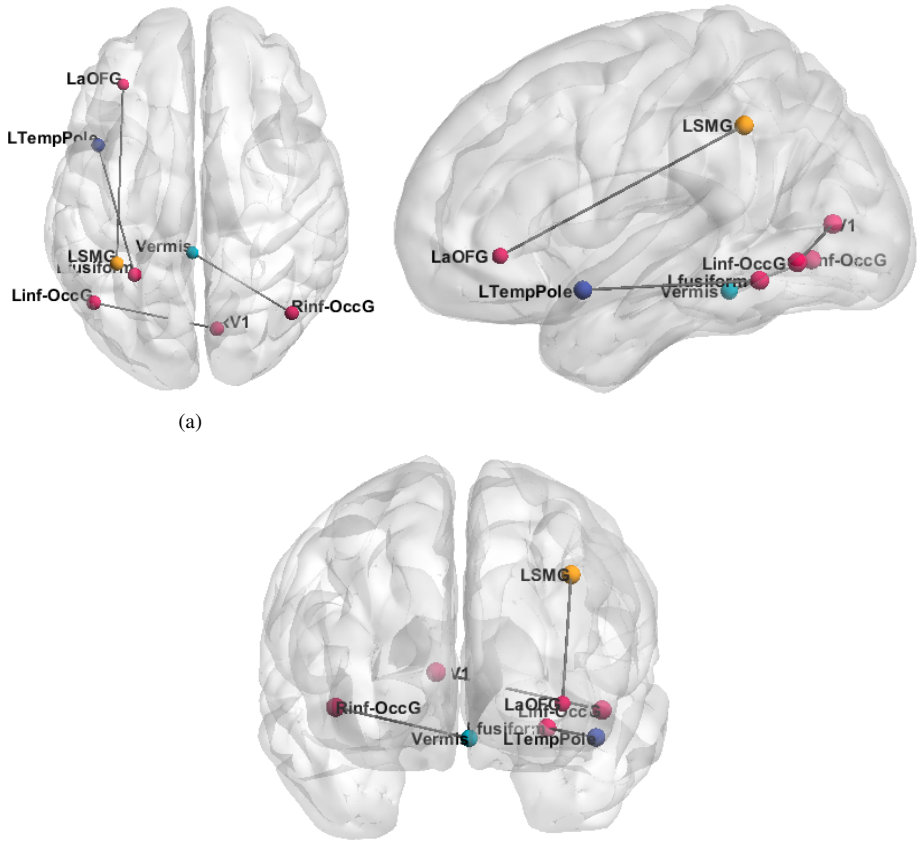

Figure 6: Functional connections differentiating ADHD from TD subjects obtained by using the proposed method (MLA) using $\alpha = 9$. From left to right, axial (a), sagittal (b), coronal (c) views of the brain indicate significant connections for at least one method are depicted. Each line represents a specific functional connection. For details on the statistics and acronyms, see Table 2.

among the detected areas using $\alpha = 9$, there were the connections between the Frontal Pole and the Cingulate Gyrus, and the Frontal Pole and Angular Gyrus, which are the main Functional hubs of the default mode network (DMN). The DMN is known to be altered in ADHD subjects (Colby et al. (2012); Uytun et al. (2016)). As it has been hypothesized that ADHD subjects may have diminished ability to inhibit the default processing of the DMN (Fassbender et al. (2009)). The other detected connections could be explained as dorsal medial and medial temporal systems still related to the DMN (Andrews-Hanna et al. (2014).

Conclusions

In this manuscript, a fully automated method to characterise brain connectivity in case-control studies was reported. The method based on a sparse learning classification, has been tested on structural and functional connectivity data. The approach is able to identify brain areas of

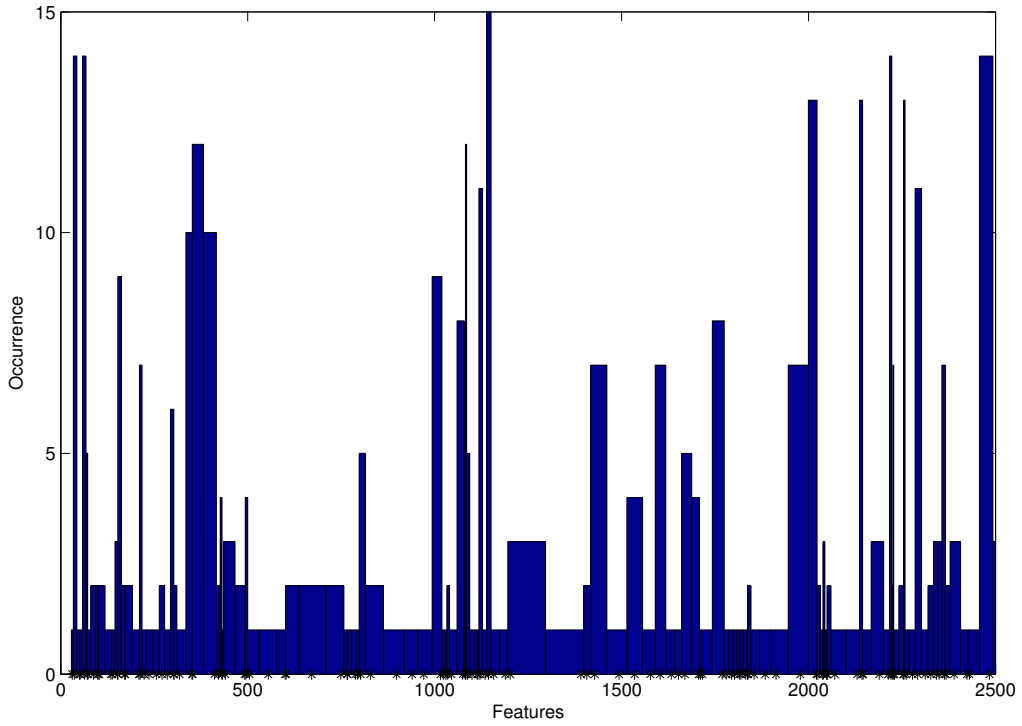

Figure 7: Histogram showing the occurrence of each feature for the human experiment after all the runs of the model for the schizophrenia experiment. Higher values indicate the edges that identify better the differences between patients with schizophrenia and controls. This histogram can be used to select a sub-set of "relevant" features, where the most frequent features are kept.

interests that can be further analysed with standard seed based approaches or through histological white matter validation.

The algorithm successfully highlighted some known structural white matter differences in acallosal mice, and identified previously reported functional connections in human schizophrenia patients with respect to control subjects. The developed software is freely distributed as a Matlab toolbox at the url *https://www.iit.it/code/multi-link-analysis*. Our approach can help highlighting differences in connectional features, and in generating hypotheses that can complement univariate techniques.

Acknowledgements

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

American Psychiatric Association, A., Association, A.P., et al., 1980. Diagnostic and statistical manual of mental disorders .

Andrews-Hanna, J.R., Smallwood, J., Spreng, R.N., 2014. The default network and self-generated thought: component processes, dynamic control, and clinical relevance. Annals of the New York Academy of Sciences 1316, 29–52.

Baggio, H.C., Abos, A., Segura, B., Campabadal, A., Garcia-Diaz, A., Uribe, C., Compta, Y., Marti, M.J., Valldeoriola, F., Junque, C., 2018. Statistical inference in brain graphs using threshold-free network-based statistics. Human Brain Mapping .

Bonilha, L., Rorden, C., Fridriksson, J., 2014. Assessing the clinical effect of residual cortical disconnection after ischemic strokes. Stroke 45, 988–993.

Bullmore, E., Sporns, O., 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. Nature Reviews Neuroscience 10, 186–198.

Casanova, M.F., El-Baz, A., Elnakib, A., Switala, A.E., Williams, E.L., Williams, D.L., Minshew, N.J., Conturo, T.E., 2011. Quantitative analysis of the shape of the corpus callosum in patients with autism and comparison individuals. Autism 15, 223–238.

Castellanos, F.X., Margulies, D.S., Kelly, C., Uddin, L.Q., Ghaffari, M., Kirsch, A., Shaw, D., Shehzad, Z., Di Martino, A., Biswal, B., et al., 2008. Cingulate-precuneus interactions: a new locus of dysfunction in adult attention-deficit/hyperactivity disorder. Biological psychiatry 63, 332–337.

Chen, S., Kang, J., Xing, Y., Wang, G., 2015. A parsimonious statistical method to detect groupwise differentially expressed functional connectivity networks. Human brain mapping 36, 5196–5206.

Clemmensen, L., Hastie, T., Witten, D., Ersbøll, B., 2011. Sparse discriminant analysis. Technometrics 53, 406–413.

Cocchi, L., Harding, I.H., Lord, A., Pantelis, C., Yucel, M., Zalesky, A., 2014. Disruption of structure–function coupling in the schizophrenia connectome. NeuroImage: Clinical 4, 779–787.

Colby, J.B., Rudie, J.D., Brown, J.A., Douglas, P.K., Cohen, M.S., Shehzad, Z., 2012. Insights into multimodal imaging classification of ADHD. Frontiers in systems neuroscience 6, 59.

Craddock, R.C., James, G.A., Holtzheimer, P.E., Hu, X.P., Mayberg, H.S., 2012. A whole brain fMRI atlas generated via spatially constrained spectral clustering. Human brain mapping 33, 1914–1928.

Dodero, L., Damiano, M., Galbusera, A., Bifone, A., Tsaftsaris, S.A., Scattoni, M.L., Gozzi, A., 2013. Neuroimaging evidence of major morpho-anatomical and functional abnormalities in the BTBR T+ TF/J mouse model of autism. PLoS One 8, e76655.

Dodero, L., Sambataro, F., Murino, V., Sona, D., 2015. Kernel-based analysis of functional brain connectivity on grassmann manifold, in: Medical Image Computing and Computer-Assisted Intervention (MICCAI), 18th International Conference on, Springer. pp. 604–611.

Ellison-Wright, I., Bullmore, E., 2009. Meta-analysis of diffusion tensor imaging studies in schizophrenia. Schizophrenia research 108, 3–10.

Fassbender, C., Zhang, H., Buzy, W.M., Cortes, C.R., Mizuiri, D., Beckett, L., Schweitzer, J.B., 2009. A lack of default network suppression is linked to increased distractibility in ADHD. Brain research 1273, 114–128.

Fenlon, L.R., Liu, S., Gobius, I., Kurniawan, N.D., Murphy, S., Moldrich, R.X., Richards, L.J., 2015. Formation of functional areas in the cerebral cortex is disrupted in a mouse model of autism spectrum disorder. Neural development 10, 1.

Fornito, A., Yoon, J., Zalesky, A., Bullmore, E.T., Carter, C.S., 2011. General and specific functional connectivity disturbances in first-episode schizophrenia during cognitive control performance. Biological psychiatry 70, 64–72.

Frazier, T.W., Hardan, A.Y., 2009. A meta-analysis of the corpus callosum in autism. Biological psychiatry 66, 935–941.

Gaonkar, B., Davatzikos, C., 2013. Analytic estimation of statistical significance maps for support vector machine based multi-variate image analysis and classification. Neuroimage 78, 270–283.

Giancardo, L., Sona, D., Gozzi, A., Bifone, A., Murino, V., Migliarini, S., Pacini, G., Pelosi, B., Pasqualetti, M., 2012. Automatic tractography analysis through sparse networks in case-control studies, in: Pattern Recognition in NeuroImaging (PRNI), 2012 International Workshop on, IEEE. pp. 77–80.

Gramfort, A., Thirion, B., Varoquaux, G., 2013. Identifying predictive regions from fMRI with TV-L1 prior, in: Pattern Recognition in Neuroimaging (PRNI), 2013 International Workshop on, IEEE. pp. 17–20.

Grefkes, C., Fink, G.R., 2011. Reorganization of cerebral networks after stroke: new insights from neuroimaging with connectivity approaches. Brain , awr033.

Griffa, A., Baumann, P.S., Thiran, J.P., Hagmann, P., 2013. Structural connectomics in brain diseases. Neuroimage 80, 515–526.

Hand, D.J., et al., 2006. Classifier technology and the illusion of progress. Statistical science 21, 1–14.

He, Y., Dagher, A., Chen, Z., Charil, A., Zijdenbos, A., Worsley, K., Evans, A., 2009. Impaired small-world efficiency in structural cortical networks in multiple sclerosis associated with white matter lesion load. Brain 132, 3366–3379.

Huang, S., Li, J., Sun, L., Ye, J., Fleisher, A., Wu, T., Chen, K., Reiman, E., Initiative, A.D.N., et al., 2010. Learning brain connectivity of Alzheimer's disease by sparse inverse covariance estimation. NeuroImage 50, 935–949.

Iturria-Medina, Y., Perez Fernandez, A., Valdes Hernandez, P., Garcia Penton, L., Canales-Rodriguez, E.J., Melie-Garcia, L., Lage Castellanos, A., Ontivero Ortega, M., 2011. Automated discrimination of brain pathological state attending to complex structural brain network properties: the shiverer mutant mouse case. PLoS One 6, e19071.

John, J.P., et al., 2009. Fronto-temporal dysfunction in schizophrenia: A selective review. Indian journal of psychiatry 51, 180.

Karbasforoushan, H., Woodward, N., 2012. Resting-state networks in schizophrenia. Current topics in medicinal chemistry 12, 2404–2414.

Kim, J., Wozniak, J.R., Mueller, B.A., Shen, X., Pan, W., 2014. Comparison of statistical tests for group differences in brain functional networks. NeuroImage 101, 681–694.

Lazar, M., Weinstein, D.M., Tsuruda, J.S., Hasan, K.M., Arfanakis, K., Meyerand, M.E., Badie, B., Rowley, H.A., Haughton, V., Field, A., et al., 2003. White matter tractography using diffusion tensor deflection. Human brain mapping 18, 306–321.

Lee, H., Lee, D.S., Kang, H., Kim, B.N., Chung, M.K., 2011. Sparse brain network recovery under compressed sensing. IEEE Transactions on Medical Imaging 30, 1154–1165.

Lynall, M.E., Bassett, D.S., Kerwin, R., McKenna, P.J., Kitzbichler, M., Muller, U., Bullmore, E., 2010. Functional connectivity and brain networks in schizophrenia. The Journal of Neuroscience 30, 9477–9487.

Mason, L., Baxter, J., Bartlett, P.L., Frean, M.R., 2000. Boosting algorithms as gradient descent, in: Advances in neural information processing systems, pp. 512–518.

McMenamin, B.W., Pessoa, L., 2015. Discovering networks altered by potential threat (anxiety) using quadratic discriminant analysis. NeuroImage 116, 1–9.

Meinshausen, N., Buhlmann, P., 2010. Stability selection. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 72, 417–473.

Meir, R., Rätsch, G., 2003. An introduction to boosting and leveraging, in: Advanced lectures on machine learning. Springer, pp. 118–183.

Meyer, B.U., Röricht, S., 1998. In vivo visualisation of the longitudinal callosal fascicle (probsts bundle) and other abnormalities in an acallosal brain. Journal of Neurology, Neurosurgery & Psychiatry 64, 138–139.

Mori, S., Crain, B.J., Chacko, V., Van Zijl, P., 1999. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. Annals of neurology 45, 265–269.

Mwansisya, T.E., Wang, Z., Tao, H., Zhang, H., Hu, A., Guo, S., Liu, Z., 2013. The diminished interhemispheric connectivity correlates with negative symptoms and cognitive impairment in first-episode schizophrenia. Schizophrenia research 150, 144–150.

Ng, B., Varoquaux, G., Poline, J.B., Greicius, M., Thirion, B., 2016. Transport on Riemannian manifold for connectivity-based brain decoding. IEEE transactions on medical imaging 35, 208–216.

Pettersson-Yeo, W., Allen, P., Benetti, S., McGuire, P., Mechelli, A., 2011. Dysconnectivity in schizophrenia: where are we now? Neuroscience & Biobehavioral Reviews 35, 1110–1124.

Pu, W., Rolls, E.T., Guo, S., Liu, H., Yu, Y., Xue, Z., Feng, J., Liu, Z., 2014. Altered functional connectivity links in neuroleptic-naive and neuroleptic-treated patients with schizophrenia, and their relation to symptoms including volition. NeuroImage: Clinical 6, 463–474.

Ren, T., Zhang, J., Plachez, C., Mori, S., Richards, L.J., 2007. Diffusion tensor magnetic resonance imaging and tract-tracing analysis of probst bundle structure in netrin1-and dcc-deficient mice. The Journal of Neuroscience 27, 10345–10349.

Richiardi, J., Eryilmaz, H., Schwartz, S., Vuilleumier, P., Van De Ville, D., 2011. Decoding brain states from fMRI connectivity graphs. Neuroimage 56, 616–626.

Rondina, J.M., Hahn, T., de Oliveira, L., Marquand, A.F., Dresler, T., Leitner, T., Fallgatter, A.J., Shawe-Taylor, J., Mourao-Miranda, J., 2014. Scorsa method based on stability for feature selection and mapping in neuroimaging. IEEE transactions on medical imaging 33, 85–98.

Rubinov, M., Sporns, O., 2010. Complex network measures of brain connectivity: uses and interpretations. Neuroimage 52, 1059–1069.

Ryali, S., Supekar, K., Abrams, D.A., Menon, V., 2010. Sparse logistic regression for whole-brain classification of fMRI data. NeuroImage 51, 752–764.

Salvador, R., Sarro, S., Gomar, J.J., Ortiz-Gil, J., Vila, F., Capdevila, A., Bullmore, E., McKenna, P.J., Pomarol-Clotet, E., 2010. Overall brain connectivity maps show cortico-subcortical abnormalities in schizophrenia. Human brain mapping 31, 2003–2014.

Sforazzini, F., Bertero, A., Dodero, L., David, G., Galbusera, A., Scattoni, M.L., Pasqualetti, M., Gozzi, A., 2014a. Altered functional connectivity networks in acallosal and socially impaired BTBR mice. Brain Structure and Function , 1–14.

Sforazzini, F., Schwarz, A.J., Galbusera, A., Bifone, A., Gozzi, A., 2014b. Distributed BOLD and CBV-weighted resting-state networks in the mouse brain. Neuroimage 87, 403–415.

Shinn, A.K., Baker, J.T., Cohen, B.M., Öngür, D., 2013. Functional connectivity of left Heschl's gyrus in vulnerability to auditory hallucinations in schizophrenia. Schizophrenia research 143, 260–268.

Simpson, S.L., Lyday, R.G., Hayasaka, S., Marsh, A.P., Laurienti, P.J., 2013. A permutation testing framework to compare groups of brain networks. Frontiers in computational neuroscience 7.

Sporns, O., 2010. Networks of the Brain. MIT press.
Sporns, O., 2011. The human connectome: a complex network. Annals of the New York Academy of Sciences 1224, 109–125.
Sporns, O., 2012. Discovering the human connectome. MIT press.
Squillace, M., Dodero, L., Federici, M., Migliarini, S., Errico, F., Napolitano, F., Krashia, P., Di Maio, A., Galbusera, A., Bifone, A., et al., 2014. Dysfunctional dopaminergic neurotransmission in asocial BTBR mice. Translational psychiatry 4, e427.
Stam, C., De Haan, W., Daffertshofer, A., Jones, B., Manshanden, I., Van Walsum, A.V.C., Montez, T., Verbunt, J., De Munck, J., Van Dijk, B., et al., 2009. Graph theoretical analysis of magnetoencephalographic functional connectivity in Alzheimer's disease. Brain 132, 213–224.
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.
Uytun, M.C., Karakaya, E., Oztop, D.B., Gengec, S., Gumus, K., Ozmen, S., Doğanay, S., Icer, S., Demirci, E., Ozsoy, S.D., 2016. Default mode network activity and neuropsychological profile in male children and adolescents with attention deficit hyperactivity disorder and conduct disorder. Brain Imaging and Behavior , 1–10.
van den Heuvel, M.P., Sporns, O., 2011. Rich-club organization of the human connectome. J Neurosci 31, 15775–15786.
Varoquaux, G., Craddock, R.C., 2013. Learning and comparing functional connectomes across subjects. NeuroImage 80, 405–415.
Wahlsten, D., Metten, P., Crabbe, J.C., 2003. Survey of 21 inbred mouse strains in two laboratories reveals that BTBR T/+ tf/tf has severely reduced hippocampal commissure and absent corpus callosum. Brain research 971, 47–54.
Winterer, G., Coppola, R., Egan, M.F., Goldberg, T.E., Weinberger, D.R., 2003. Functional and effective frontotemporal connectivity and genetic risk for schizophrenia. Biological psychiatry 54, 1181–1192.
Witten, D., Tibshirani, R., 2011. Penalized classification using Fisher's linear discriminant. Journal of the Royal Statistical Society: Series B (Statistical Methodology) .
Yamashita, O., Sato, M.a., Yoshioka, T., Tong, F., Kamitani, Y., 2008. Sparse estimation automatically selects voxels relevant for the decoding of fMRI activity patterns. NeuroImage 42, 1414–1429.
Zalesky, A., Fornito, A., Bullmore, E.T., 2010. Network-based statistic: identifying differences in brain networks. Neuroimage 53, 1197–1207.
Zalesky, A., Fornito, A., Seal, M.L., Cocchi, L., Westin, C.F., Bullmore, E.T., Egan, G.F., Pantelis, C., 2011. Disrupted axonal fiber connectivity in schizophrenia. Biological psychiatry 69, 80–89.
Zeng, L.L., Shen, H., Liu, L., Hu, D., 2014. Unsupervised classification of major depression using functional connectivity MRI. Human brain mapping 35, 1630–1641.
Zou, H., Hastie, T., 2005. Regularization and variable selection via the elastic net. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 67, 301–320.

Sparse Discriminant Analysis

The general formulation of ℓ_1 regularization or *lasso* is used in regression frameworks to minimize the problem $min_\beta\{||\mathbf{y} - \mathbf{X}\beta||^2 + \eta||\beta||_1\}$, where $\mathbf{X}$ is a data matrix, $\mathbf{y}$ is the output vector, and β is the regressor vector. Similarly, the elastic net is given as $min_\beta\{||\mathbf{y} - \mathbf{X}\beta||^2 + \eta||\beta||_1 + \gamma||\beta||_2\}$. In these equations η and γ are tuning parameters which are used to yield sparse coefficient vector estimation (Zou and Hastie (2005)). The parameter η can also be reformulated as the number of desired variables which are left in the model, when used in this context we refer to it as α (Zou and Hastie (2005)).

There are several extension to the linear discriminant analysis (Hand et al. (2006)) which comprises lasso and elastic net, as explained by Witten and Tibshirani (2011). Our experiments are based on the formulation proposed by Clemmensen et al. (2011). More specifically, given the matrix data $\mathbf{X}$ with n observations p-dimensional for K=2 classes each of them defined as $\mathbf{x_i}$, it is possible to define the mean for a specific class as $\mu_\mathbf{k}$. Then, it is also possible to define the within-class covariance matrix common to all classes as $\mathbf{\Sigma_w} = \frac{1}{n}\sum_{k=1}^{K}\sum_i(\mathbf{x_i} - \mu_\mathbf{k})(\mathbf{x_i} - \mu_\mathbf{k})^T$, and the between-class covariance matrix $\mathbf{\Sigma_b} = \sum_{k=1}^{K}\pi_k\mu_\mathbf{k}\mu_\mathbf{k}{}^T$, where π_k is the prior probability for each class to belong to the class k. The prior probability is generally given by the size of respective classes.

A Fischer discriminant analysis can classify to which class a sample belongs by using discriminant vectors whose directions $\beta_{\mathbf{k}}$ which maximize

$$\max_{\beta_{\mathbf{k}}} \beta_{\mathbf{k}}^{\mathbf{T}} \mathbf{\Sigma_b} \beta_{\mathbf{k}}, \tag{.1}$$

subject to $\beta_{\mathbf{k}}^{\mathbf{T}} \mathbf{\Sigma_b} \beta_{\mathbf{k}} = 1$ and $\beta_{\mathbf{k}}^{\mathbf{T}} \mathbf{\Sigma_b} \beta_{\mathbf{l}} = 0 \; \forall l < k$.

Very often, as in our case, the previous maximization process is ill-posed, as the matrix $\mathbf{\Sigma_w}$ might not be full rank as the number of features is far larger than the number of available samples. A possible solution, proposed by Witten and Tibshirani (2011), is given by using the lasso or elastic net regularization as

$$\max_{\beta_{\mathbf{k}}} \beta_{\mathbf{k}}^{\mathbf{T}} \mathbf{\Sigma_b} \beta_{\mathbf{k}} - \eta ||\beta_k||_1 - \gamma ||\beta_k||_2 \tag{.2}$$

subject to $\beta_{\mathbf{k}}^{\mathbf{T}} \mathbf{\Sigma_b} \beta_{\mathbf{k}} = 1$ and $\beta_{\mathbf{k}}^{\mathbf{T}} \mathbf{\Sigma_b} \beta_{\mathbf{l}} = 0 \forall l < k$. Alternatively we used the minimization formulation of Clemmensen et al. (2011), where the pair given by $\beta_{\mathbf{k}}$ and the vector of scores $\theta_{\mathbf{k}}$ solves the problem

$$\begin{cases} min_{(\beta_k, \theta_k)} ||\mathbf{Y}\theta_{\mathbf{k}} - \mathbf{X}\beta_{\mathbf{k}}||^2 + \eta \, ||\beta_k||_1 + \gamma \, \beta_k^T \mathbf{\Omega} \beta_k, \\ \text{subject to } \frac{1}{n} \theta_{\mathbf{k}}^{\mathbf{T}} \mathbf{Y}^{\mathbf{T}} \mathbf{Y} \theta_{\mathbf{k}} = 1, \theta_{\mathbf{k}}^{\mathbf{T}} \mathbf{Y}^{\mathbf{T}} \mathbf{Y} \theta_{\mathbf{l}} = 0 \; \forall l < k, \end{cases} \tag{.3}$$

where $\mathbf{\Omega}$ is an arbitrary positive matrix, η and γ are non negative tuning parameters, and $\mathbf{Y}$ is a $n \times K$ matrix of dummy variables for the K classes. This formulation of LDA as a regression problem introduces sparsity, and allows its use when the number of features very large compared to the number of available samples.