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, Alzheimer, ADHD, Autism, ASD, 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)), autism spectrum disorder (Crimi et al. (2017)), 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 the aforementioned 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, Mastrovito et al. (2018); 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, and these estimates were used in a SVM framework. Then, both in Gaonkar and Davatzikos (2013) and Ng et al. (2016) meaningful weights were found through t-tests. While in Mastrovito et al. (2018) recursive feature elimination (RFE) is employed to keep features which are relevant for the classification starting from the whole set of features. 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)). 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. Lastly, 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 and the t-tests are carried out in a univariate manner. Moreover, the choice of related statistics can influence considerably the results (Baggio et al. (2018)).

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 followed by a stability selection step. 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. Furthermore, sparse models follow a feature selection agenda to subselect among existing variables, whereas PCA dimensionality reduction follows a feature engineering agenda to generate a set of new variables. A feature selection by sparse model is indeed similar to the RFE used by Mastrovito et al. (2018). However, the stability of RFE approach depends heavily on the type of used model for feature ranking at each iteration, and as shown empirically, using regularized ridge regression jointly to stability selection criteria can provide more stable results in terms of *stability selection* of features if used combined with a stability criterion, as this combination yields finite sample familywise error control and dramatically improves feature selection (Meinshausen and Bühlmann (2010); Ye et al. (2012)). 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 (Clemmensen et al. (2011)). Essentially, the system acts as a *filter* removing the elements that are not useful to discriminate between the groups. First by enforcing sparsity at individual level. Then, by performing a second stage of feature filtering

across the dataset to assure stability selection. It is acknowledged that such a feature selection can be applied to arbitrary high-dimensional, multivariate datasets. Nevertheless, recent studies showed that sparsity based approach can be particularly useful in graph/connectome analysis as they can highlight significant connections when prior knowledge is missing (Deligianni et al. (2013); Xie and He (2012); Coloigner et al. (2017); Kim et al. (2014)).

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 structural connectivity matrices from a publicly available dataset of patients affected by Azheimer's disease, where connectivity is also defined by tractography. The final experiment was carried out on a large functional 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. Those results are compared to the results obtained by using NBS. NBS and MLA select discriminative features in different ways. NBS performs univariate t-tests among the features while MLA performs a sparse multivariate regression. Nevertheless, NBS is the commonly used algorithm for this type of analysis and considered the state-of-art.

Methods and Data

This section first describes the two types of data used to test the proposed method: a mice dataset with high dimensionality, and two publicly available human 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 (n=16): BTBR T+tf/J mice (n=8) 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 (n=8) 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). Figure .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 freely distributed [2]. 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 Structural Connectivity Alzheimer Data

The himan experiments have been performed on the Alzheimers Disease Neuroimaging Initiative (ADNI) dataset publicly available[3]. Only baseline scans were used to avoid confounding factors as advanced brain atrophy and treatment in the Alzheimer patients. This cohort comprised 51 Alzheimer patients (age: 76.5 ± 7.4 years), and 49 normal elderly subjects (77.0±5.1) matched by age. The used data were DTI, and T1-Weighted obtained by using a GE Signa scanner 3T (General Electric, Milwaukee, WI, USA). The T1-weighted scans were acquired at with voxel size = $1.2 \times 1.0 \times 1.0$ mm^3 TR = 6.984 ms; TE = 2.848 ms; flip angle=11°). DTI were acquired at voxel size = $1.4 \times 1.4 \times 2.7$ mm^3, scan time = 9 min, and 46 volumes (5 T2-weighted images with no diffusion sensitization b0 and 41 diffusion-weighted images b=1000 s/mm^2). For each subject, DTI and T1 have been acquired and co-registered.

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[4], 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.

[2]https://www.iit.it/code/multi-link-analysis
[3]http://adni.loni.usc.edu/
[4]http://umcd.humanconnectomeproject.org

The fMRI volumes were acquired with a Siemens Allegra 3T, with TR/TE 2000/15 ms and voxel size $3 \times 3 \times 3mm^3$. This dataset is used as a real case study where NBS is not useful in proving local connectivity differences.

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 $1 \times 1 \times 1mm^3$, 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.

Alzheimer Dataset Processing and Encoding

Tractographies for all subjects have been generated processing the DTI data with a deterministic Euler approach, stemming from 2,000,000 seed-points and stopping in case of FA smaller than 0.25. Tracts shorter than 3cm were discarded during the connectome construction. Structural connectivity matrices were constructed by counting the number of fibers connecting two regions of interest (ROIs) of the registered Harvard-Oxford atlas (Makris et al. (2006)). This atlas defines 96 ROIs, it is freely available with several brainimage analysis platforms, and it has been used in several structural studies including Alzheimer (e.g. inDagley et al. (2017)). The Harvard-Oxford atlas is a probabilistic atlas. However, we used the version where each voxel is associated to the ROI with highest probability.

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 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 according to a stability selection criterion.

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 a nested cross-validation manner to produce a jackknife-like classification. In practice a nested leave-one-out procedure was employed. For 1000 iterations, a sample was removed from the training dataset and used as test-set to find the optimal value within a range, then the performances were evaluated on another sample also removed prior the optimization from the training set and used as validation. In presence of a plateau of identical optimal values, the value generating less connections was taken. 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 second stage exploiting the ensemble of low-stability algorithms to produce a more stable feature selection. In practice, we perform further feature selection to ensure that the features are stable across subjects.

In the specific case, the SDA classifier was trained with a nested leave-one-out approach. 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. Then we refine this ensemble of models 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. In this way, we ensure stability selection of the features (Meinshausen and Bühlmann (2010)). 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 structural and functional connectivity collected to study respectively *Alzheimer's* and *ADHD*.

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, 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. The mean misclassification varying the parameter α resulting by the cross-validation is shown in Figure .2 (a).

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(b) 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.

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 $\alpha = 110$ estimated by nested cross-validation which also gave 100% accuracy).

Human Structural and Functional Data

We tested the algorithm also on a publicly available dataset based on human MRI recorded from patients with Alzheimer against normal elderly subjects, and ADHD against TD subjects. The Alzheimer dataset was in particular used to investigate the influence of the regularization parameter regarding the features which are obtained by using the proposed method and NBS. In practice, to test which features are detected by the proposed method and not by the NBS and vice-versa varying α.

The NBS algorithm with threshold equivalent to p-value =0.05 detected 14 connections which used in nested cross-validation classification task produced 65% accuracy. As shown in Figure .4, by increasing the α parameter, the number of included features by MLA and not by NBS increases and the opposite decreases with a break-even-point at $\alpha = 15$. Interestingly, the best classification score with the features detected by the MLA was optimal for the α values of 31 where MLA was including more features than the NBS algorithm and producing 80% accuracy. It has to be noted that with the second stage of feature selection the number of retained features by MLA can be slightly less than the parameter α. Moreover, it is worthwhile to mention that generally features were symmetric. Namely, the connection from a ROI a to a ROI b and vice-versa from a ROI b to a ROI a were detected. The resulting features produced by using the MLA when $\alpha = 31$ are depicted in Figure .5. Here, NBS produced 14 connections, which interestingly were all included in those found by the proposed algorithm. The identified connections were mostly ipsilateral connections within the two temporal lobes.

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 with threshold equivalent to p-value = 0.05 did not produce 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 Alzheimer and ADHD experiment. Connections set to zero by MLA are not depicted. The Alzheimer and ADHD samples analysis took respectively less than 1 second and about 30 seconds on 2.6 GHz machine with 4GB of RAM.

Table 1: Structural connections differentiating patients with Alzheimer from normal elderly individuals detected by MLA. Pairs of source and target regions and p-values of the univariate t-test computed on NBS (Zalesky et al. (2010)) using the t-test threshold corresponding to p-values < 0.05 are reported. "Not detected" (N.D.) means not significant difference between the two areas.

#	Region 1	Region 2	p-value NBS
1	Insula-L (INS.L)	Frontal Orbital Cortex-L (FOC.L)	0.0002
2	Insula-L (INS.L)	Inferior Frontal Gyrus, pars opercularis-L (F3t.L)	0.0002
3	Superior Frontal Gyrus-L (FS.L)	Inferior Frontal Gyrus, pars opercularis-L (F3t.L)	N.D.
4	Superior Frontal Gyrus-L (FS.L)	Parahippocampal Gyrus, ant. div.-L (PHa.L)	0.01
5	TemporalPole-L (TP.L)	Parahippocampal Gyrus, ant. div.-L (PHa.L)	0.0003
6	TemporalPole-R (TP.R)	Frontal Operculum-R (FO.R)	N.D.
7	TemporalPole-R (TP.R)	Planum Polare-R (PP.R)	0.0003
8	Superior Temporal Gyrus, post. div.-L (T1p.L)	Angular Gyrus-L (AG.L)	N.D.
9	Superior Temporal Gyrus, post. div.-L (T1p.L)	Planum Temporale-L (PT.L)	0.001
10	Superior Temporal Gyrus, post. div.-R (T1p.R)	Parietal Operculum-R (PO.R)	0.0003
11	Inferior Temporal Gyrus, temporooccipital-R (TO3.R)	Temporal Occipital Fusiform-R (TOF.R)	N.D.
12	Angular Gyrus-R (AG.R)	Parietal Operculum-R (PO.R)	N.D.
13	Cuneal Cortex-L (CN.L)	Frontal Operculum-R (OF.R)	N.D.
14	Insula-R (INS.R)	Frontal Orbital Cortex-R (FOC.R)	0.006
15	Parahippocampal Gyrus, ant. div.-L (PHa.L)	Temporal Fusiform Cortex, ant.r div. (TFa.L)	0.001
16	Temporal Fusiform Cortex, ant. div. (TFa.L)	Parahippocampal Gyrus, post. div.-L (PHp.L)	0.01

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, but if considered with the univariate analysis, some edges are discarded as producing non-significant p-values (e.g., the features detected by using the ADHD data). Nevertheless, both MLA and NBS gave similar results for the mice data experiment 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

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. For this experiment, no statistically significant features were obtained by the NBS algorithm using the t-test threshold corresponding to p-value = 0.05.

#	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.

the previous ones, and satisfying accuracy is achieved for a large range of values as shown in Figure .2 (a). Nevertheless, 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 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.

Similar results were obtained with the structural dataset on Alzheimer. Altered brain connectivity both at the microstructural and macrocircuitry levels has been described in this disorder due to the amyloid plaques (Sheline et al. (2010)). As depicted in Figure .5, we found widespread connectivity temporal and para-hippocampal differences in patients with Alzheimer compared to healthy elderly subjects. This is inline with several studies about functional connectivity which highlighted decreased functional connectivity between the temporal gyrus and neighbouring regions (e.g. Huang et al. (2010); Supekar et al. (2008)). While pathways between the hippocampus, the parahippocampal gyrus and neocortical regions are considered to be the first affected in Alzheimer's patients (de LaCoste and White (1993); Allen et al. (2007)). Neurodegeneration and loss of connectivity between frontal areas, the insula and within the areas of the frontal gyrus has been also documented (Seeley et al. (2009)). The found connection between the cuneal cortex left and the frontal opercolum right were possible through fibers bifurcating from the corpus callous. Studies have been shown that in particular the functional activities at the cuneal cortex is reduced with the progress of Alzheimer Zhang et al. (2010), this can explain visual field defects

seen in some patients Berisha et al. (2007).

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

The connectivity differences for the human dataset, shown in Figure .5 and .6, and reported in Table1 and 2 were those detected by MLA, where some of them were gone miss by NBS. This shows that different approaches - one based on sparsity and one based on family-wise error rate- can produce similar results but the proposed method can find more features than the NBS. NBS and MLA select discriminative features in different ways. NBS performs univariate t-tests among the features while MLA performs a sparse multivariate regression. It is acknowledged that the methods achieve different objectives. However, it has been shown that sparse model jointly to stability selection criteria introduced in the second stage of the proposed method can lead to more robust results than false discovery rate and family-wise error rate based model Hofner et al. (2015). Moreover, without a ground truth, MLA can be preferred for exploratory purposes as it can provide more insights (more detected features) to neuroscientists looking for case-control differences.

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 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 structural and functional connections in human Alzheimer's and ADHD 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

Allen, G., Barnard, H., McColl, R., Hester, A.L., Fields, J.A., Weiner, M.F., Ringe, W.K., Lipton, A.M., Brooker, M., McDonald, E., et al., 2007. Reduced hippocampal functional connectivity in Alzheimer disease. Archives of neurology 64, 1482–1487.

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. @articlemastrovito2018differences, title=Differences in atypical resting-state effective connectivity distinguish autism from schizophrenia, author=Mastrovito, Dana and Hanson, Catherine and Hanson, Stephen Jose, journal=NeuroImage: Clinical, volume=18, pages=367–376, year=2018, publisher=Elsevier . Human Brain Mapping .

Berisha, F., Feke, G.T., Trempe, C.L., McMeel, J.W., Schepens, C.L., 2007. Retinal abnormalities in early Alzheimers disease. Investigative ophthalmology & visual science 48, 2285–2289.

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.

Coloigner, J., Phlypo, R., Coates, T.D., Lepore, N., Wood, J.C., 2017. Graph lasso-based test for evaluating functional brain connectivity in sickle cell disease. Brain connectivity 7, 443–453.

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.

Crimi, A., Dodero, L., Murino, V., Sona, D., 2017. Case-control discrimination through effective brain connectivity, in: Biomedical Imaging (ISBI 2017), 2017 IEEE 14th International Symposium on, Ieee. pp. 970–973.

Dagley, A., LaPoint, M., Huijbers, W., Hedden, T., McLaren, D.G., Chatwal, J.P., Papp, K.V., Amariglio, R.E., Blacker, D., Rentz, D.M., et al., 2017. Harvard aging brain study: dataset and accessibility. NeuroImage 144, 255–258.

Deligianni, F., Varoquaux, G., Thirion, B., Sharp, D.J., Ledig, C., Leech, R., Rueckert, D., 2013. A framework for inter-subject prediction of functional connectivity from structural networks. IEEE transactions on medical imaging 32, 2200–2214.

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.

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.

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.
Hofner, B., Boccuto, L., Göker, M., 2015. Controlling false discoveries in high-dimensional situations: Boosting with stability selection. BMC bioinformatics 16, 144.
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.
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.
de LaCoste, M.C., White, C.L., 1993. The role of cortical connectivity in alzheimer's disease pathogenesis: a review and model system. Neurobiology of Aging 14, 1–16.
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.
Makris, N., Goldstein, J.M., Kennedy, D., Hodge, S.M., Caviness, V.S., Faraone, S.V., Tsuang, M.T., Seidman, L.J., 2006. Decreased volume of left and total anterior insular lobule in schizophrenia. Schizophrenia research 83, 155–171.
Mastrovito, D., Hanson, C., Hanson, S.J., 2018. Differences in atypical resting-state effective connectivity distinguish autism from schizophrenia. NeuroImage: Clinical 18, 367–376.
McMenamin, B.W., Pessoa, L., 2015. Discovering networks altered by potential threat (anxiety) using quadratic discriminant analysis. NeuroImage 116, 1–9.
Meinshausen, N., Bühlmann, P., 2010. Stability selection. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 72, 417–473.
Meinshausen, N., Buhlmann, P., 2010. Stability selection. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 72, 417–473.
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.
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.
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.
Seeley, W.W., Crawford, R.K., Zhou, J., Miller, B.L., Greicius, M.D., 2009. Neurodegenerative diseases target large-scale human brain networks. Neuron 62, 42–52.
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.
Sheline, Y.I., Raichle, M.E., Snyder, A.Z., Morris, J.C., Head, D., Wang, S., Mintun, M.A., 2010. Amyloid plaques disrupt resting state default mode network connectivity in cognitively normal elderly. Biological psychiatry 67, 584–

587.
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.
Supekar, K., Menon, V., Rubin, D., Musen, M., Greicius, M.D., 2008. Network analysis of intrinsic functional brain connectivity in alzheimer's disease. PLoS computational biology 4, e1000100.
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.
Witten, D., Tibshirani, R., 2011. Penalized classification using Fisher's linear discriminant. Journal of the Royal Statistical Society: Series B (Statistical Methodology) .
Xie, T., He, Y., 2012. Mapping the alzheimers brain with connectomics. Frontiers in psychiatry 2, 77.
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.
Ye, J., Farnum, M., Yang, E., Verbeeck, R., Lobanov, V., Raghavan, N., Novak, G., DiBernardo, A., Narayan, V.A., 2012. Sparse learning and stability selection for predicting mci to ad conversion using baseline adni data. BMC neurology 12, 46.
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.
Zhang, H.Y., Wang, S.J., Liu, B., Ma, Z.L., Yang, M., Zhang, Z.J., Teng, G.J., 2010. Resting brain connectivity: changes during the progress of alzheimer disease. Radiology 256, 598–606.
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.

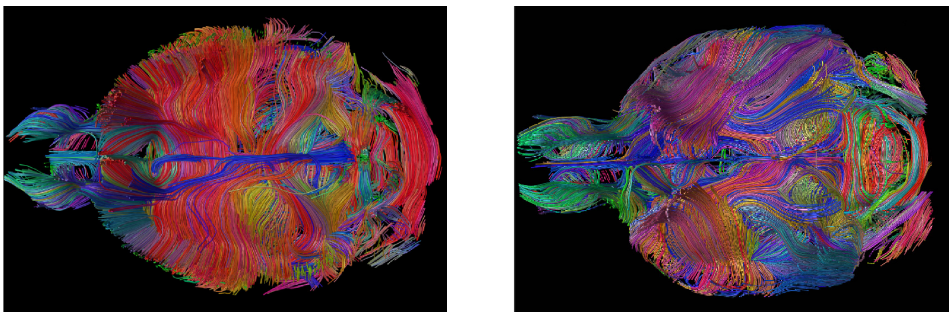

(a) (b)

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.

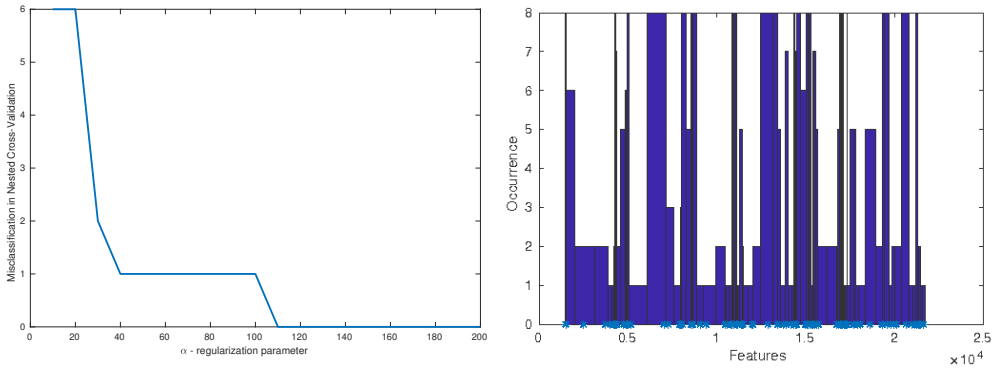

(a) (b)

Figure .2: Mice data experiment: (a) Misclassification error as function of regularization parameter computed with the nested cross-validation. The misclassification error reaches a plateau after $\alpha = 110$. (b) Histogram describing the occurrences of features (i.e. brain connections) selected in the mouse experiment. Higher values indicate connections that characterize the differences between BTBR and control mice in the classifiers within our ensemble framework. This information is used to automatically select a sub-set of "relevant" features. Namely, the most frequent features highlighted by the histogram are kept.

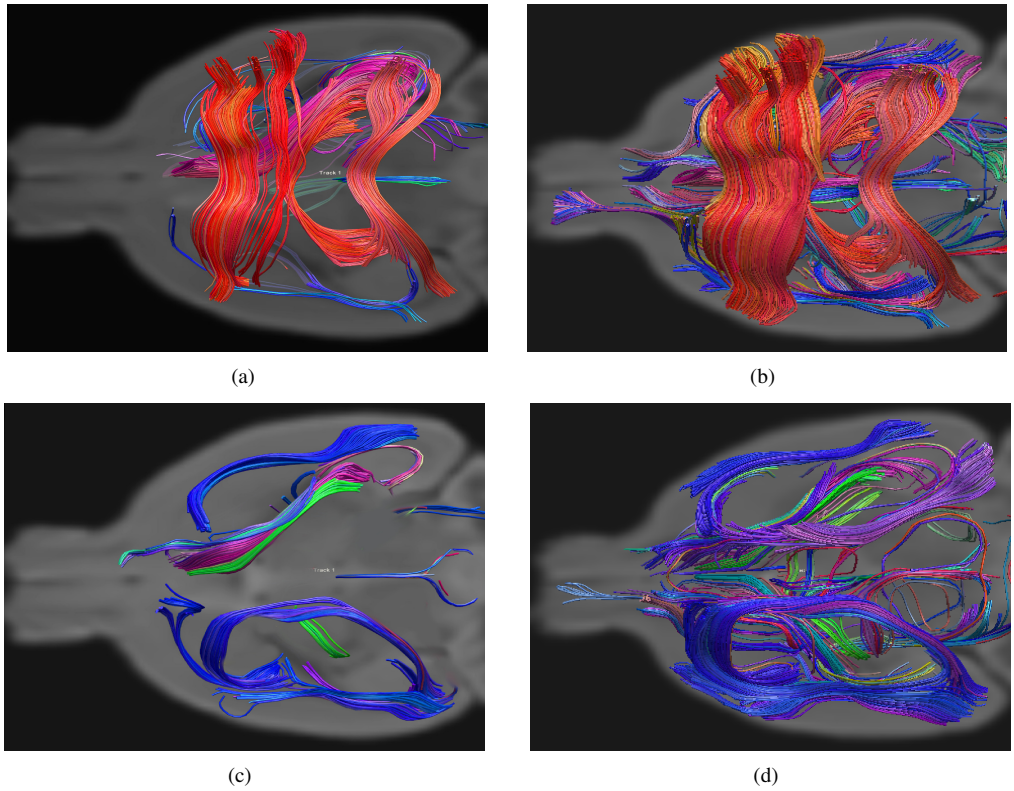

Figure .3: Graphical representation of the most significant features charactering the structural connectome of the two populations: the axial views of a randomly selected subject from the C57BL/6J control population (a) using our algorithm ($\alpha = 110$) and (b) the NBS algorithm using as a threshold p-value 0.05. 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. The depicted discriminant features detected by the proposed algorithm can be increased varying the α parameter according to the user taste.

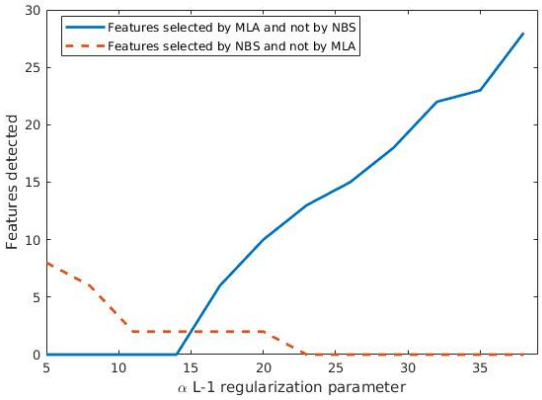

Figure .4: Number of features detected by one algorithm and not by the the other varying the α. This graph shows that by increasing the parameter, the number of features detected by the MLA is increasing.

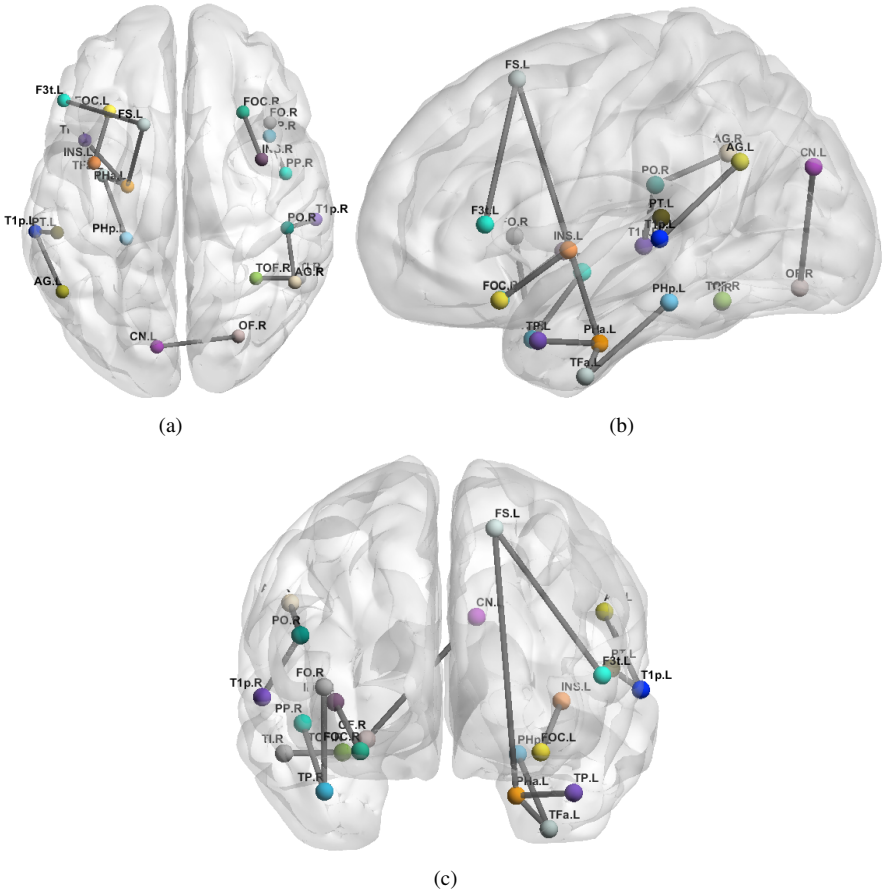

Figure .5: Structural connections differentiating patients with Alzheimer from elderly controls with MLA (our algorithm) using $\alpha = 31$. From left to right, axial (a), sagittal (b), coronal (c) views of the brain indicate significant connections not set to zero by the algorithm. Each line represents a specific structural connection. The acronymous are the same as reported in Table 1.

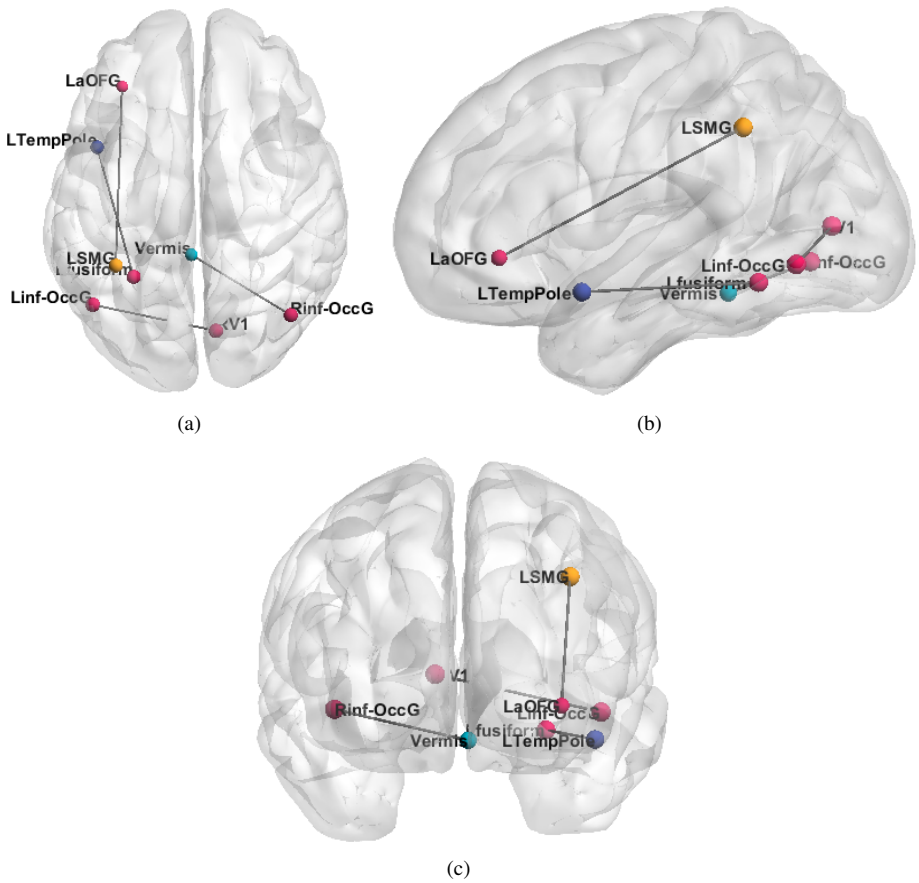

Figure .6: Functional connections differentiating ADHD from TD subjects obtained by using the proposed method (MLA) using $\alpha = 9$. From left to right, axial (a), sagittal (b), coronal (c) views of the brain indicate significant connections not set to zero by the algorithm. Each line represents a specific functional connection. For details on the statistics and name abbreviations, see Table 2.