Stable Biomarker Identification For Predicting Schizophrenia in the Human Connectome

Leonardo Gutiérrez-Gómez[1,✉], Jakub Vohryzek[2,3], Benjamin Chiêm[1], Philipp S. Baumann[4], Philippe Conus[4], Kim Do Cuenod[4], Patric Hagmann[3], and Jean-Charles Delvenne[1,5]

[1]Institute for Information and Communication Technologies, Electronics and Applied Mathematics (ICTEAM), Université catholique de Louvain, Louvain-la-Neuve, Belgium. Corresponding author: leonardo.gutierrez@uclouvain.be
[2]Department of Psychiatry, University of Oxford, Oxford, United Kingdom
[3]Department of Radiology, University Hospital Centre and University of Lausanne, Lausanne, Switzerland
[4]Service of General Psychiatry and Center for Psychiatric Neuroscience, Department of Psychiatry, Lausanne University Hospital (CHUV), Lausanne, Switzerland
[5]Center for Operations Research and Econometrics (CORE), Université catholique de Louvain, Louvain-la-Neuve, Belgium

Schizophrenia, as a mental disorder, has been well documented with both structural and functional magnetic resonance imaging. The developing field of connectomics has attracted much attention as it allows researchers to take advantage of powerful tools of network analysis in order to study structural and functional connectivity abnormalities in schizophrenia. Many methods have been proposed to identify biomarkers in schizophrenia, focusing mainly on improving the classification performance or performing statistical comparisons between groups. However, the stability of biomarkers selection has been for long overlooked in the connectomics field. In this study, we follow a machine learning approach where the identification of biomarkers is addressed as a feature selection problem for a classification task. We perform a recursive feature elimination and support vector machines (RFE-SVM) approach to identify the most meaningful biomarkers from the structural, functional, and multi-modal connectomes of healthy controls and patients. Furthermore, the stability of the retrieved biomarkers is assessed across different subsamplings of the dataset, allowing us to identify the affected core of the pathology. Further analysis of the connectivity strength of the affected core quantified with the generalized fractional anisotropy (gFA) and apparent diffusion coefficient (ADC) reveal a significant alteration in patients. Considering our technique altogether demonstrates a principled way to achieve both accurate and stable biomarkers while highlighting the importance of multi-modal approaches to brain pathology as they tend to reveal complementary information.

schizophrenia diagnosis | stable biomarkers | RFE-SVM | human connectome

1 Introduction

Schizophrenia (SZ) is a severe psychiatric disorder characterized by hallucinations and delusions, as well as impairments in memory, attention, executive and other high-order cognitive dysfunctions (1). The development of magnetic resonance imaging (MRI) has offered an effective way to examine the anatomy of the brain and has motivated numerous scientists to explore the underlying neuropathology of SZ. Over the past few years, advances in high-field structural and functional neuroimaging have made it possible to map the macroscopic neural wiring system of the human brain (2, 3). This framework allows us to combine a graph theoretical approach with functional and/or diffusion MRI to investigate the brain network alterations occurring in SZ (4–6). Studies from the graph theory perspective have shown alterations of both structural and functional brain topology in SZ characterized by a less efficient network organization and the limited capacity of functional integration. Further research using diffusion spectrum imaging reported the brain areas responsible for the loss of network integration and segregation properties (5, 7, 8). However, such findings were identified using conventional univariate strategies performing a separate statistical test at each edge of the connectome under scrutiny, thereby requiring excessively stringent corrections for multiple comparisons. On the other hand, multivariate methods are promising, although they require specialized approaches when the number of parameters dominates the observations (9). In this study, we adopt a machine learning approach that aims at discovering the most relevant set of biomarkers for discriminating subjects groups and thus quantitatively describing the group differences, both in terms of classification accuracy and stability of selected features.

A Machine learning and automatic biomarker selection

The identification of regions or connections of interest associated with a neural disorder is referred to as biomarker discovery. The identification of such biomarkers in schizophrenia could lead to clinically useful tools for establishing both diagnosis and prognosis. From a machine learning perspective, the choice of biomarkers can be addressed as a feature selection problem, aiming to find a subset of relevant features allowing us to differentiate patients from control subjects accurately.

In this work, we perform an automatic feature selection procedure in order to identify biomarkers that are relevant for the diagnosis of schizophrenia from brain connectivity data. In this context, biomarkers, therefore, correspond to structural or functional links between neural Regions of Interest (ROIs).

A key challenge in feature selection lies in the fact that diverse feature selection methods might result in different sets of retrieved features. Even when using the same technique, it may produce different results when applied to different splittings of the data. When the dimensionality of the input data is

large and exceeds the number of training examples, the complexity grows by several orders of magnitude (10, 11). As a consequence, when two subsets of features are compared the decision of which should be preferred involves uncertainty (12).

These issues underline the need to integrate the stability in the feature selection process so that the method can retrieve consistent features across random subsamplings of the dataset. This is especially true in a biomedical context where many authors have focused on improving the classification performance in several mental disorders such as schizophrenia (13), Alzheimer (14), depression (15). Even though the stability of biomarker selection has been studied mainly in genomics and proteomics (12**?**), the stability of feature selection has been overlooked in the connectomics community. Therefore, we propose a general framework for stability analysis of selected features, thereby enabling the robust identification of impaired connections in the connectome of schizophrenic patients. The proposed approach is extendable to other brain disorders as well.

In the present work, we use Support Vector Machines (SVM) as a classifier (10). This is a supervised machine learning method that aims to classify data points by maximizing the margin between classes in a high-dimensional space. This classifier offers state of the art classification performances on a wide range of applications and is particularly appropriate for high-dimensional problems with few examples. The SVM classifier has been integrated into an embedded feature selection approach (**?**). The so-called Recursive Feature Elimination with Support Vector Machine (RFE-SVM) technique was first introduced to perform gene selection for cancer diagnosis on microarray data (16). More recently, it has also been used on human brain networks to identify differences in structural connectivity related to gender (**?**). This method trains an SVM classifier removing the less important features and iteratively re-estimating the classifier with the remaining features until reaching the desired number of them. Accordingly, we adopt the RFE-SVM approach to automatically select brain connections that lead to the best discrimination between patients and controls, and consequently to highlight brain regions that are responsible for the disease.

The aim of the present work is threefold: First, we investigate the effect of structural, functional, and multi-modal (structural+functional) connectome with different resolutions in the classification performance of schizophrenia. Second, we perform a careful feature selection procedure across modalities in order to assess the robustness of the selected features providing the best trade-off between high accuracy and stability. Finally, the analysis of retrieved biomarkers allows us to identify a distributed set of brain regions engaged in the discrimination of patients and control subjects.

This paper is organized as follows: Section 2 introduces the properties of the dataset, the procedure for connectomes estimation as well as the general protocol we used in biomarkers identification. In section 3, we present the results on stability, classification performances, and identification of brain areas indicative of the pathology. Finally, in section 4, we lead a discussion on our findings and conclusions.

2 Materials and methods

A Subjects

For this study, two age-balanced groups were considered. The cohort consisted of a schizophrenic group of 37 subjects with a mean age of 40.9 ± 9.4 years and a control group of 37 healthy subjects with a mean age of 32.3 ± 7.6 years. The patients in the schizophrenic group were recruited from the psychiatric clinic at the Lausanne University Hospital. They met DSM-IV criteria for schizophrenic and schizoaffective disorders (American psychiatry association, 2000 ref). We obtained written consent from all the subjects following the institutional guidelines approved by the Ethics Committee of Clinical Research of the Faculty of Biology and Medicine, University of Lausanne, Switzerland.

B Brain network estimation

Magnetic resonance imaging

All subjects were scanned on the 3 Tesla Siemens Trio scanner with a 32-channel head coil. Three acquisition protocols were part of the MRI session: 1) magnetization-prepared rapid acquisition gradient echo (MPRAGE) sequence with in-plane resolution of 1 mm, slice thickness of 1.2 mm of total voxel number of 240x257x160 and TR, TE and TI were 2300, 2.98 and 900 ms respectively, 2) diffusion spectrum imaging (DSI) sequence with 128 diffusion-weighted images of b0 as a reference image and a maximum b-value of 8000 s/mm2. The time of acquisition was 13 min and 27s. The number of voxels was 96x96x34 with a resolution of 2.2x2.2x3.3 mm, and TR and TE were 6100 and 144 ms respectively. The issue of motion- artifacts linked to signal drop-outs was dealt with by visually inspecting the signal, and no subject had to be excluded as a result of this (17).

Structural networks

Structural and diffusion MRI data were used to estimate the weighted and undirected structural connectivity matrices in the Connectome Mapping Toolkit (7, 18, 19). Firstly, white matter, grey matter, and cerebrospinal fluid segmentation was performed on the structural data and further linearly registered to the b0 volumes of the DSI dataset. Secondly, the first three scales of the Lausanne multi-scale atlas were used to parcellate the grey matter. In detail, the first scale consisted of 68 cortical brain regions and 14 subcortical regions with scale two and three subdividing the first scale into 114 and 219 cortical regions (19). Further, deterministic streamline tractography, estimating 32 diffusion directions per voxels, was used to reconstruct the structural connectivities from the DSI data(20). The normalized connection density quantified the structural connectivity between brain regions and is defined as follows,

$$w_{ij} = \frac{2}{S_i S_j} \sum_{f \in E_f} \frac{1}{\ell(f)} \qquad \textbf{(1)}$$

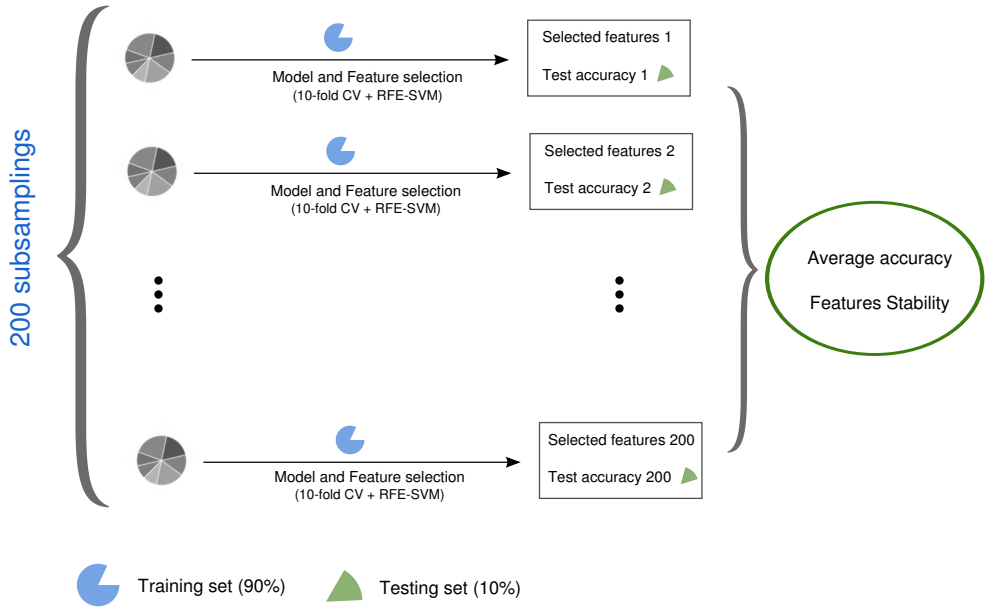

Fig. 1. Overview of the proposed method. Feature selection is performed systematically across different partitions of the original dataset. Robustness of selected biomarkers is assessed from the output of the RFE-SVM algorithm, and the final accuracy is averaged over subsamplings estimations.

where w represents an edge between brain regions i and j, S_i and S_j are the surface areas of regions i and j, $f \in E_f$ represents a streamline f in the set of streamlines E and $\ell(f)$ is a length of a given streamline f (2, 7). The normalisation by brain region surfaces accounts for their slightly varying size and the streamline length normalisation accounts for a bias towards longer connections imposed by the tractography algorithm.

Functional Networks

Functional connectivity matrices were computed from fMRI BOLD time-series. Firstly, the first four-time points were excluded, yielding the number of time points to be $T = 276$ (21). Rigid-body registration was applied to individual timeslices for motion-correction. The signal was then linearly detrended and corrected for physiological confounds and further motion artifacts by regressing white-matter, cerebrospinal fluid, and six motion (translations and rotations) signals. Lastly, the signal was spatially smoothed and bandpass-filtered between $0.01 - 0.1$ Hz with Hamming windowed sinc FIR filter. Linear registration was performed between the average fMRI and MPRAGE images to obtain the ROIs timeseries(22). An average timecourse for each brain region was computed for the three atlas scales. In order to obtain the functional matrices, the absolute value of the Pearson correlation was computed between individual brain regions' timecourses. All of the above was performed in subject native space with Connectome Mapper Toolkit and personalized Python and Matlab scripts(18)(23).

C Biomarker evaluation protocol

Our evaluation methodology is based on Abeel et al 2010 (24) used for biomarker identification in cancer diagnosis on microarray data. In order to assess the robustness of the biomarker selection process, we generate slight variations of the dataset and compare the outcome of selected features across these variations. Therefore, for a stable marker selection algorithm, small variations in the training set should not produce important changes in the retrieved set of features. Concretely, we generate 200 datasets from the original one, repeating 20 times 10-fold cross-validation (CV). Subsequently, for each CV partition, 90% of folds are used as a training set for selecting the model and features, and the remaining 10% as the testing set to provide an unbiased evaluation of a final model and assess the performance of the classifier. Therefore, the overall accuracy is given by the average testing accuracy across subsamplings. See Figure 1 for a schematic view of the methodology.

Assessing the stability of feature selection

We consider the vectorized connectivity matrices of the connectomes as input features for the biomarker selection process. Therefore, for a given connectome, one structural feature refers to the normalized connection density between two linked brain regions, whereas a functional feature refers to the Pearson correlation between two individual brain time courses.

We consider a dataset with $M = 74$ subjects and N features. If we denote the considered connectome resolution as $d \in \{83, 129, 234\}$, the number of features is $N = \frac{d(d+1)}{2}$ because of the symmetry of the connectivity matrices. Drawing $k = 200$ subsamplings and after applying a feature selection procedure (RFE-SVM) in the 90% of each subsampling, we obtain a respective feature signature, i.e. sequence of indices of selected features. Considering two signatures f_i and f_j obtained from different subsamplings i and j, the stability index (25) between f_i and f_j is defined as:

$$KI(f_i, f_j) = \frac{rN - s^2}{s(N-s)} = \frac{r - (s^2/N)}{s - (s^2/N)} \quad \textbf{(2)}$$

where $r = |f_i \cap f_j|$ and $s = |f_i| = |f_j|$, the size of the signature, is the number of selected features. This quantity measures the consistency between pairs of features. For s and N fixed KI increases when increases r, reaching its maximum at 1 when the two subsets are identical. The minimum value is bounded from below by -1 when the subsets are perfectly disjoint and the signature size of $N/2$. The overall stability index for a sequence of signatures is defined as the average of all pairwise stability indices on k subsamplings:

$$I_{tot} = \frac{2}{k(k-1)} \sum_{i=1}^{k-1} \sum_{j=i+1}^{k} KI(f_i, f_j) \quad (3)$$

Given that I_{tot} is bounded between -1 and 1, the greater this value the better the agreement between the selected subsets of features. In particular, a negative value for I_{tot} indicates that the potential agreement between the selected biomarkers is mostly due to chance. In the sequel we will refer to the overall stability I_{tot} as the Kuncheva index (25).

D Measuring the classification performance

Because our dataset is class balanced and the task is a binary classification problem, we adopt the accuracy as the metric to quantify the performance of classification. This metric is defined as the ratio of correct classifications to the number of classifications done as follows:

$$ACC = \frac{TP + TN}{TP + TN + FP + FN} \quad (4)$$

where TP is true positive (number of control subjects classified correctly), TN is the true negative (patients classified correctly), FP is false positives (number of patients classified as control subjects) and FN is false negatives (number of control subjects misclassified).

E Embedded feature selection (RFE-SVM)

In this work we use a linear SVM classifier (10). SVM has proven state of the art performance in computational biology (26) in particular with problems of very high dimension, scaling very well as a function of the number of examples. Given a set of data examples, SVM aims to find the hyperplane that has the largest distance to the nearest training data points of any class. The solution of the optimization problem provides the coefficients of such a hyperplane as an affine combination of the support vectors, i.e., points lying on the max-margin hyperplane of separation between classes, and the training examples. These coefficients can be interpreted as a strength or contribution of each feature to the decision of the hyperplane. As a consequence, the square value of each coefficient (or weight) can be used as a score to rank features from the most to the least important for the selection process. Recursive feature elimination SVM (RFE-SVM) (16) is an iterative algorithm integrating a ranking criterion for eliminating features in a backward fashion. Starting with the whole set of features, a linear SVM is estimated using the training set, and their features are ranked according to the weights assigned by the algorithm. Consequently, the least important features are removed and the remaining ones are used to train a new model, repeating the process until reaching a desired minimal subset of features.

The RFE-SVM algorithm has a set of internal parameters influencing the computational complexity and the accuracy of the method. The fraction E of features to remove at each step of RFE (also called step size) is critical for the running time. Dropping one feature at a time allows a finer selection but with a prohibitive computational cost. Following the work of Abeel et al. (24), we drop 20% of the least relevant features at each iteration by default. Yet, a stopping criterion is needed to finish the iterative process. Thus, in our experiments, we dropped features until reaching a minimum of $s \in \{0.5, 1, 2, 5, 10, 25, 50\}$ percentage of selected features (stopping criterion).

Another critical parameter is the regularization constant C of the SVM. The C parameter controls the misclassification rate of the classifier. A larger value makes the optimization choose a smaller margin hyperplane, losing generalization capabilities. The smaller the C, the larger the margin of separation, yielding more misclassified points. Therefore this parameter influences the classification accuracy of the model. We cross-validate the optimal C using only the training set, e.g., the 90% of each subsampling.

3 RESULTS

A Connectomes classification and features stability

First, we investigated the effect of different brain connectivity modalities and different scales in the discrimination of patients and normal controls. For each case, we control the step size and the percentage of selected features of the RFE-SVM algorithm, assessing their impact on the classification accuracy and the stability of the selected features.

Figures 2, 3 and 4 show the average classification accuracy after performing RFE-SVM as well as the stability of selected biomarkers across modalities and scales. It can be seen that across scales, the functional connectivity matrices (Fig. 3) achieve better accuracies than the structural matrices (Fig. 2), but conversely, structural matrices are more stable than functional. However, when combining the two modalities, i.e., by concatenating features of both modalities and letting the algorithm choose a blend of structural and functional features we achieve the best performances (Fig. 4) in terms of both accuracy and stability. Note that the multi-modal curves overlap more than the structural and functional ones in the same resolutions. This shows that the multi-modal matrices provide similar performance in both accuracy and stability with respect to the number of dropped features in the RFE-SVM algorithm. It is to be noted that in the multi-modal case, the percentage of finally selected features is divided by two since we combine twice as many features as in the case of structural or functional connectivity alone. In doing so, the stopping criterion will be the same for all modalities. We observe in all cases that the stability increases with the percentage of selected features, which is expected since the overlapping be-

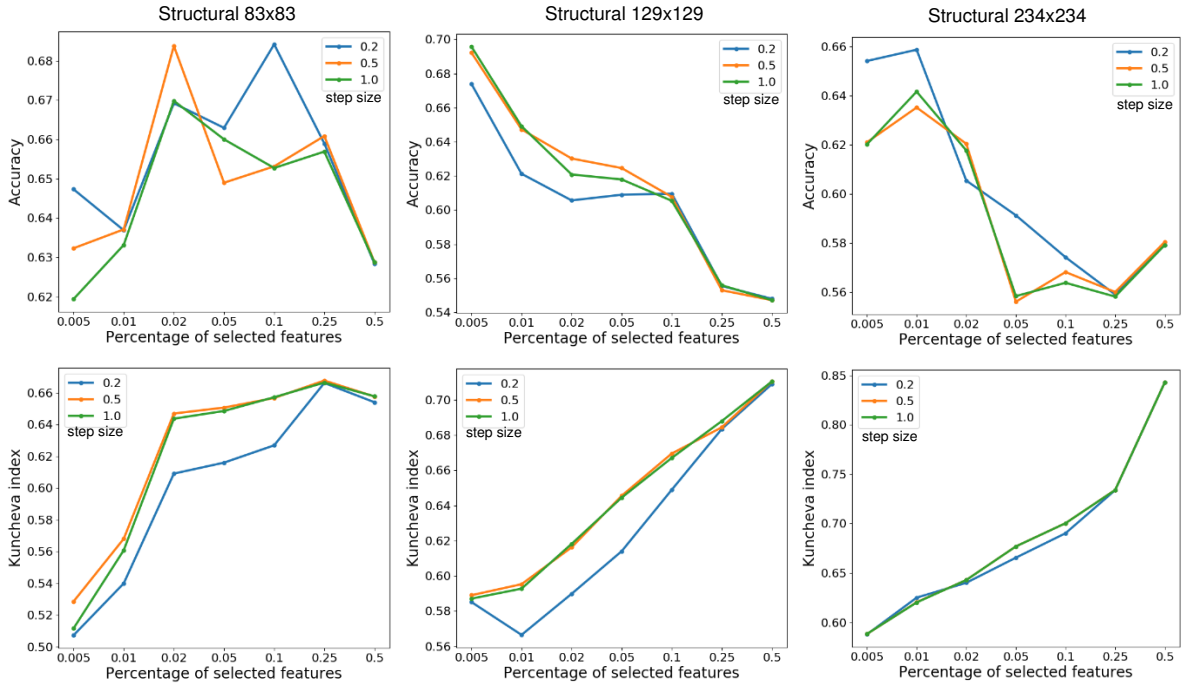

Fig. 2. Each column represents the average classification accuracy (top) and stability (bottom) for a given scale of the structural connectome. The colors correspond to the percentage of features dropped (step size) at each step of the RFE-SVM algorithm.

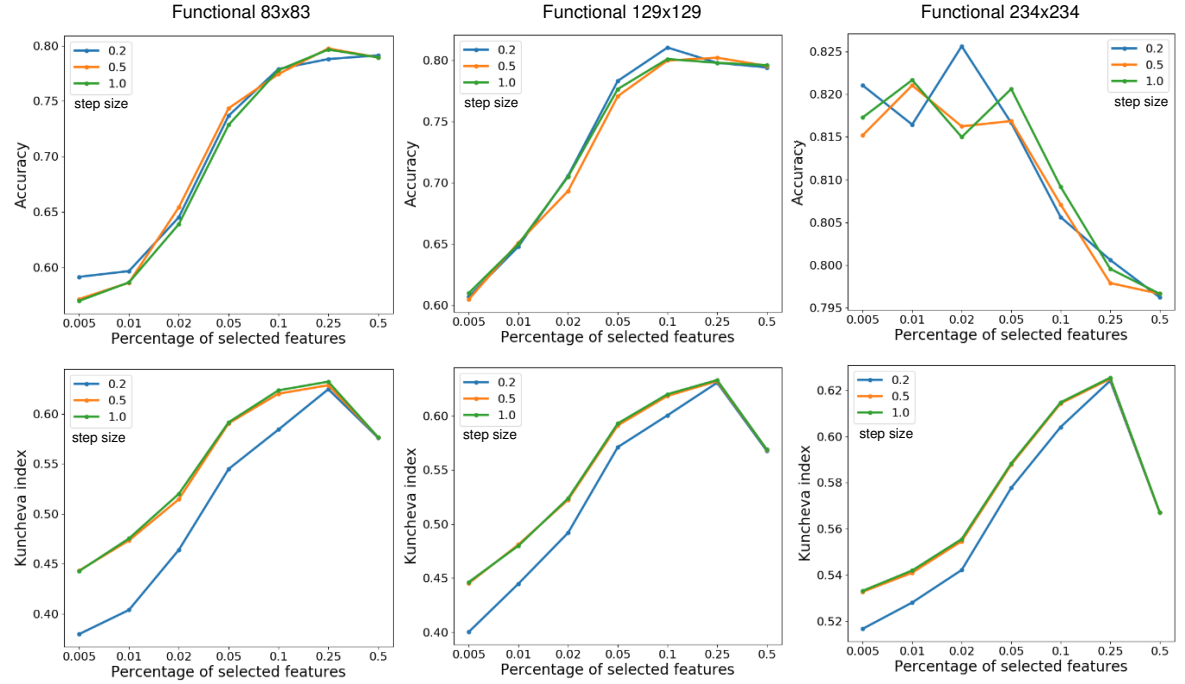

Fig. 3. Each column represents the average classification accuracy (top) and stability (bottom) for a given scale of the functional connectome. The colors correspond to the percentage of features dropped (step size) at each step of the RFE-SVM algorithm.

tween signatures in Eq. 3 is more likely when more features are considered.

To further investigate the effect of modalities and resolutions on both classification accuracy and stability, we extract the best scores from Figs. 2, 3 and 4, and plot them in Fig. 5.

As demonstrated, the best trade-off between both metrics is achieved by 129×129 and 83×83 multi-modal resolutions. The 234×234 multi-modal resolution has the lowest stability, which can be explained by the fact that finer resolutions are subsamplings of smaller ROIs introducing redundant and correlated edges in the connectome.

B Identification of brain regions in schizophrenia diagnosis

We proceed with the identification of brain areas involved in the classification of patients and controls. For simplicity in the identification of brain regions and comparison with other authors, we analyze the results for the multi-modal 83×83

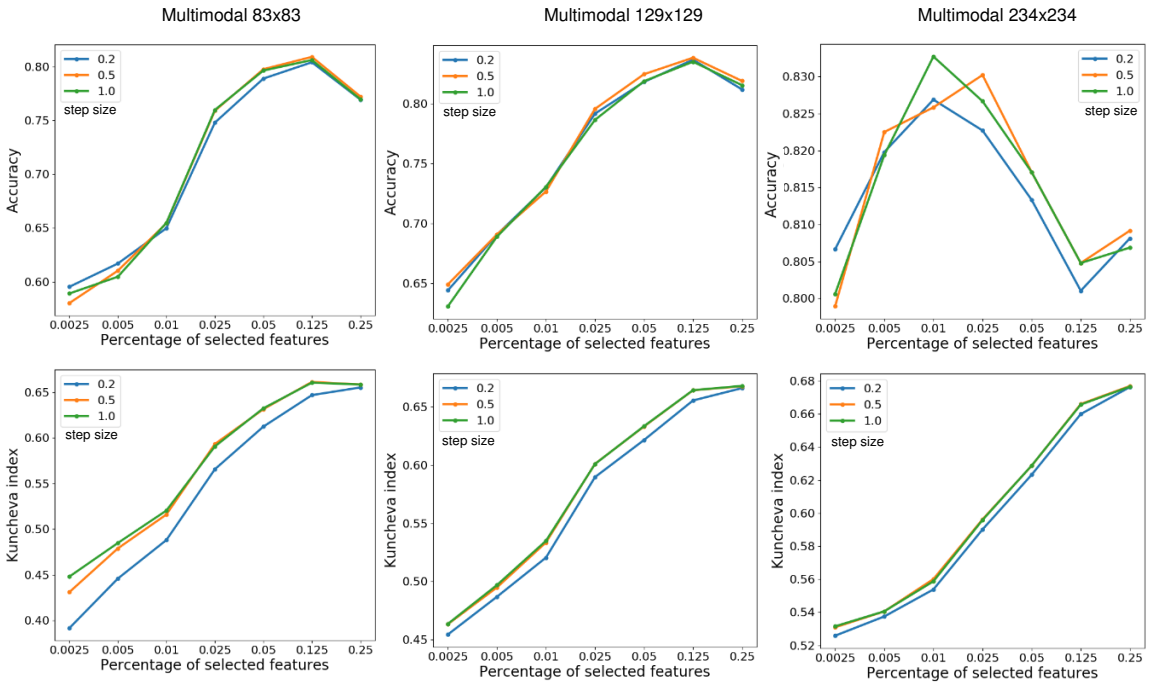

Fig. 4. Each column represents the average classification accuracy (top) and stability (bottom) for a given scale of the multimodal connectome. The colors correspond to the percentage of features dropped (step size) at each step of the RFE-SVM algorithm.

connectome.

Selected features in the graph space correspond to links representing either connection densities in the structural matrices or Pearson correlations in the functional connectome. Furthermore, inspecting the frequency of each selected feature across subsamplings informs us about the overall relevance of the edge in the classification. In other words, the frequency of an edge is indicative of the importance of the associated ROIs in the classification task.

Given a brain connectivity matrix at a resolution r we define W as the $r \times r$ matrix where the element $w_{i,j}$ encodes the frequency at which the edge (i, j) is selected as relevant across subsamplings. Thus, the degree of relevance of an ROI i reads:

$$d_i = \sum_{j \in N_i} w_{i,j} \qquad \textbf{(5)}$$

Figure 7 represents the degree of relevance of brain regions for the multi-modal 83×83 connectome sorted in decreasing order. We defined the affected core (a-core) to be composed of brain areas with a degree of relevance higher than the overall average. These findings overlap with the brain regions in the definition of the affected core of Griffa et al. 2015 (7). We plot the brain surface in Figure 9, normalizing both FC and SC by the sum of all their connections and plotted the regions above the median distribution.

C Characterizing intergroup differences

We also investigate the influence of the selected biomarkers in the brain connectivity networks for patients and healthy subjects. To do so, the connectivity strength inside and outside the identified a-core was estimated and compared between both groups. The connectivity strength between

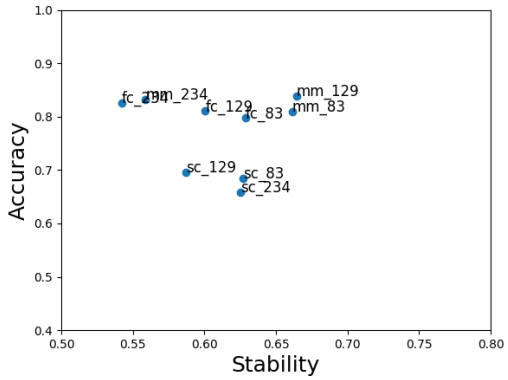

Fig. 5. Accuracy versus stability for all considered modalities and resolutions.

ROIs was quantified with two metrics: the inverse apparent diffusion coefficient (iADC) and the generalized fractional anisotropy (gFA). These measures were weighted by the tract size expressed as the number of fibers. Subsequently, they were averaged over all connectivity strengths between a-core nodes and all connectivity strengths between regions outside the a-core. We reported the $p-$values from the Mann–Whitney–Wilcoxon (MWW) test for between-group differences of the averaged connectivity strengths, with a significance level $\alpha = 0.05$.

As can be seen in Figure 6, the weighted gFA and weighted iADC were altered in patients compared with controls. When averaging within the a-core the gFA was decreased with $p-$value $= 0.00017$ as well the iADC with $p-$value $= 0.01845$. None difference between groups was found in the outside a-core network of averaged gFA and iADC.

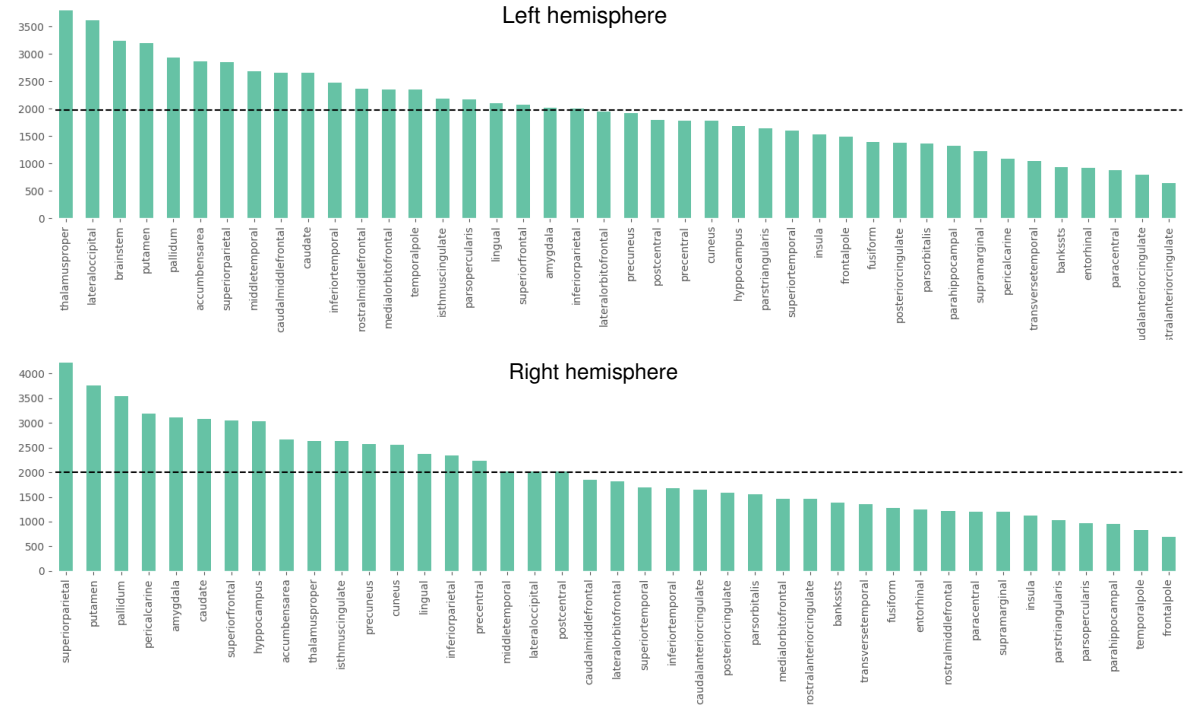

Fig. 7. Degree of relevance for ROIs in the structural mode of the multimodal 83×83 connectome. The horizontal line is the average degree of relevance.

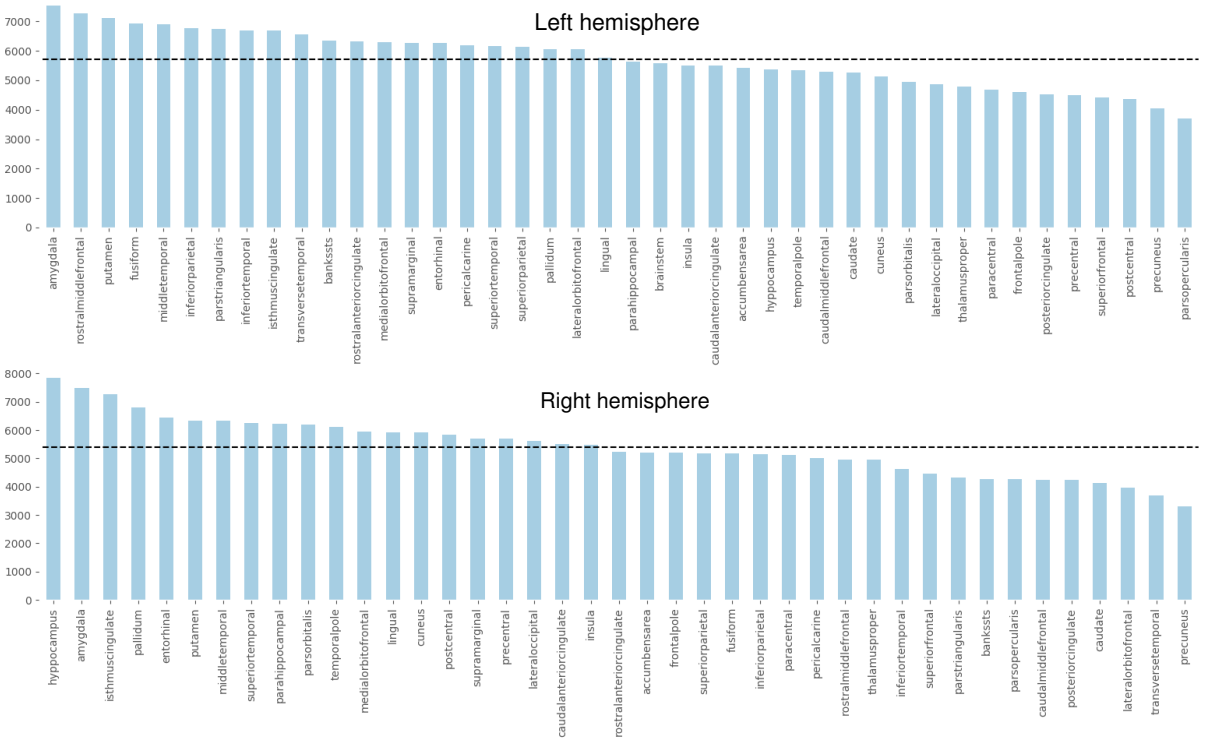

Fig. 8. Degree of relevance for ROIs in the functional mode of the multimodal 83×83 connectome. The horizontal line is the average degree of relevance.

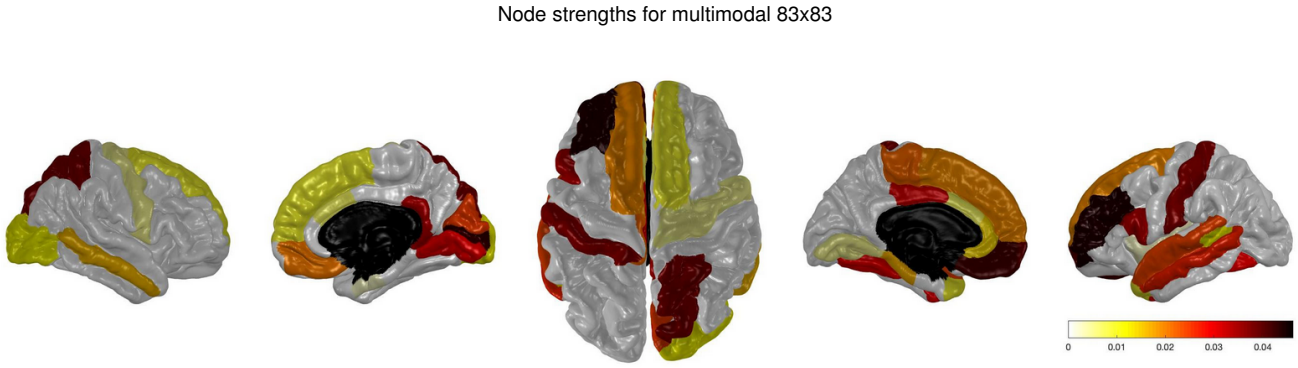

Fig. 9. Brain surface representation of brain areas with higher relevance degree than the average for the 83×83 resolution of the multi-modal structural and functional mode.

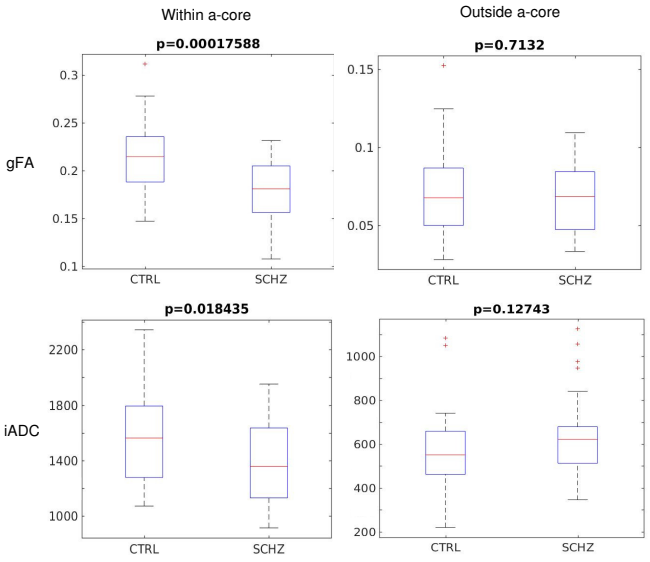

Fig. 6. Boxplots of the distribution of connectivity strength averaged within (left) and outside the a-core. Both considered metrics, the generalized fractional anisotropy (gFA) and inverse apparent diffusion coefficient (ADC) were weighed by the size of the tracts.

4 CONCLUSIONS

This paper has investigated the effect of different connectivity modes of the human connectome in the identification of robust biomarkers for the diagnosis of schizophrenia. We perform an automatic feature selection process on the edge space aiming to retrieve a compact subset of meaningful biomarkers performing accurately on the diagnosis of schizophrenia. Besides, we analyze the robustness of the retrieved features concerning the sample variation, based on the fact that stable biomarkers will not change dramatically in different subsamplings of the underlying dataset (24). It turns out that combining structural and functional connectivity matrices as a multi-modal representation of connectomes provides the best trade-off between high accurate and stable biomarkers.

Importantly, the fact that the best accuracy and stability was achieved by combining SC and FC features reflects the relevance of both approaches to studying and classifying brain disorders. It contributes to the literature of the relationship between structural and functional large-scale networks (23, 27) and highlights the importance of understanding both approaches to explain complex neural disorders at least insofar the prediction of them is concerned.

Based on the frequency from which the RFE-SVM algorithm selects relevant edges, we can map edges to the node space by looking at the strengths of connection densities or correlation between brain regions. The degree of relevance for nodes allows us to identify the affected core as the brain regions that are highly active during the training phase and therefore, are selected more often for relevant edges. Our findings overlap results from Griffa et al. 2015 (7) and provide further evidence of brain regions involved in the pathology. Furthermore, they offer a unique perspective in defining the affected core both from the influence of SC and FC and hence defining an SC-FC a-core which might give a richer description of the regions that are affected in SCHZ (6, 7).

The alterations of gFA and iADC values within and outside the a-core depicted in Fig. 6, suggest a micro-structural impairment of the a-core regions given by an interruption of white matter across the affected core circuits.

ACKNOWLEDGEMENTS

This work was supported by Concerted Research Action (ARC) supported by the Federation Wallonia-Brussels Contract ARC 14/19-060, Flagship European Research Area Network (FLAG-ERA) Joint Transnational Call “FuturICT 2.0”; Swiss National Science Foundation grant numbers: 156874 and 158776 to which are gratefully acknowledged.

Bibliography

1. Jim van Os, Gunter Kenis, and Bart P. F. Rutten. The environment and schizophrenia. *Nature*, 468:203 EP –, Nov 2010. Perspective.
2. Patric Hagmann, Leila Cammoun, Xavier Gigandet, Reto Meuli, Christopher J. Honey, J. Van Wedeen, and Olaf Sporns. Mapping the structural core of human cerebral cortex. *PLoS Biology*, 6(7):1479–1493, 2008. ISSN 15449173. doi: 10.1371/journal.pbio.0060159.
3. Olaf Sporns, Giulio Tononi, and Rolf Kötter. The human connectome: A structural description of the human brain. *PLoS computational biology*, 1(4):e42, 9 2005. ISSN 1553-7358. doi: 10.1371/journal.pcbi.0010042.
4. Alex Fornito, Andrew Zalesky, Christos Pantelis, and Edward T. Bullmore. Schizophrenia, neuroimaging and connectomics. *NeuroImage*, 62(4):2296 – 2314, 2012. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2011.12.090. Connectivity.
5. Alessandra Griffa, Philipp S. Baumann, Jean Philippe Thiran, and Patric Hagmann. Structural connectomics in brain diseases. *NeuroImage*, 80:515–526, 2013. ISSN 10538119. doi: 10.1016/j.neuroimage.2013.04.056.
6. Mary-Ellen Lynall, Danielle S. Bassett, Robert Kerwin, Peter J. McKenna, Manfred Kitzbichler, Ulrich Muller, and Ed Bullmore. Functional connectivity and brain networks in schizophrenia. *Journal of Neuroscience*, 30(28):9477–9487, 2010. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.0333-10.2010.
7. Alessandra Griffa, Philipp Sebastian Baumann, Carina Ferrari, Kim Quang Do, Philippe Conus, Jean Philippe Thiran, and Patric Hagmann. Characterizing the connectome in schizophrenia with diffusion spectrum imaging. *Human Brain Mapping*, 36(1):354–366, 2015. ISSN 10970193. doi: 10.1002/hbm.22633.
8. Alessandra Griffa, Philipp S. Baumann, Paul Klauser, Emeline Mullier, Martine Cleusix, Raoul Jenni, Martijn P. van den Heuvel, Kim Q. Do, Philippe Conus, and Patric Hagmann.

Brain connectivity alterations in early psychosis: from clinical to neuroimaging staging. *Translational Psychiatry*, 9(1):62, 2019. ISSN 2158-3188. doi: 10.1038/s41398-019-0392-y.

9. Peter Bühlmann and Sara van de Geer. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer Publishing Company, Incorporated, 1st edition, 2011. ISBN 3642201911, 9783642201912.
10. Bernhard E. Boser, Isabelle M. Guyon, and Vladimir N. Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, COLT '92, pages 144–152, New York, NY, USA, 1992. ACM. ISBN 0-89791-497-X. doi: 10.1145/130385.130401.
11. Isabelle Guyon and André Elisseeff. An introduction to variable and feature selection. *J. Mach. Learn. Res.*, 3:1157–1182, March 2003. ISSN 1532-4435.
12. Zengyou He and Weichuan Yu. Stable feature selection for biomarker discovery. *Computational Biology and Chemistry*, 34(4):215 – 225, 2010. ISSN 1476-9271. doi: https://doi.org/10.1016/j.compbiolchem.2010.07.002.
13. Xiaobing Lu, Yongzhe Yang, Fengchun Wu, Minjian Gao, Yong Xu, Yue Zhang, Yongcheng Yao, Xin Du, Chengwei Li, Lei Wu, Xiaomei Zhong, Yanling Zhou, Ni Fan, Yingjun Zheng, Dongsheng Xiong, Hongjun Peng, Javier Escudero, Biao Huang, Xiaobo Li, Yuping Ning, and Kai Wu. Discriminative analysis of schizophrenia using support vector machine and recursive feature elimination on structural mri images. *Medicine (Baltimore)*, 95(30):e3973–e3973, Jul 2016. ISSN 1536-5964. doi: 10.1097/MD.0000000000003973. 27472673[pmid].
14. Zhengjia Dai, Chaogan Yan, Zhiqun Wang, Jinhui Wang, Mingrui Xia, Kuncheng Li, and Yong He. Discriminative analysis of early alzheimer's disease using multi-modal imaging and multi-level characterization with multi-classifier (m3). *NeuroImage*, 59(3):2187 – 2195, 2012. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2011.10.003.
15. Minyue Chi, Shengwen Guo, Yuping Ning, Jie Li, Haochen Qi, Minjian Gao, Xiuyong Wu, Junwei Xue, Xin Du, Jiexin Wang, Xiaowei Hu, Yangbo Guo, Yuling Yang, Hongjun Peng, and Kai Wu. Discriminative analysis of major depressive disorder and anxious depression using support vector machine. *Journal of Computational and Theoretical Nanoscience*, 12 (7):1395–1401, 2015. ISSN 1546-1955. doi: doi:10.1166/jctn.2015.3903.
16. Isabelle Guyon, Jason Weston, Stephen Barnhill, and Vladimir Vapnik. Gene selection for cancer classification using support vector machines. *Machine Learning*, 46(1):389–422, Jan 2002. ISSN 1573-0565. doi: 10.1023/A:1012487302797.
17. Anastasia Yendiki, Kami Koldewyn, Sita Kakunoori, Nancy Kanwisher, and Bruce Fischl. Spurious group differences due to head motion in a diffusion mri study. *Neuroimage*, 88:79–90, Mar 2014. ISSN 1095-9572. doi: 10.1016/j.neuroimage.2013.11.027. 24269273[pmid].
18. Alessandro Daducci, Stephan Gerhard, Alessandra Griffa, Alia Lemkaddem, Leila Cammoun, Xavier Gigandet, Reto Meuli, Patric Hagmann, and Jean Philippe Thiran. The Connectome Mapper: An Open-Source Processing Pipeline to Map Connectomes with MRI. *PLoS ONE*, 7(12):e48121, 12 2012. ISSN 19326203. doi: 10.1371/journal.pone.0048121.
19. Leila Cammoun, Xavier Gigandet, Djalel Meskaldji, Jean Philippe Thiran, Olaf Sporns, Kim Q. Do, Philippe Maeder, Reto Meuli, and Patric Hagmann. Mapping the human connectome at multiple scales with diffusion spectrum MRI. *Journal of Neuroscience Methods*, 203(2):386–397, 2012. ISSN 01650270. doi: 10.1016/j.jneumeth.2011.09.031.
20. Van J. Wedeen, Patric Hagmann, Wen Yih Isaac Tseng, Timothy G. Reese, and Robert M. Weisskoff. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magnetic Resonance in Medicine*, 54(6):1377–1386, 2005. ISSN 07403194. doi: 10.1002/mrm.20642.
21. Mark Jenkinson, Peter Bannister, Michael Brady, and Stephen Smith. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2):825–841, 2002. ISSN 10538119. doi: 10.1016/S1053-8119(02)91132-8.
22. Mark Jenkinson, Christian F. Beckmann, Timothy E.J. Behrens, Mark W. Woolrich, and Stephen M. Smith. Fsl. *NeuroImage*, 62(2):782–790, 2012. ISSN 10538119. doi: 10.1016/j.neuroimage.2011.09.015.
23. Alessandra Griffa, Benjamin Ricaud, Kirell Benzi, Xavier Bresson, Alessandro Daducci, Pierre Vandergheynst, Jean-Philippe Thiran, and Patric Hagmann. Transient Networks of Spatio-temporal Connectivity Map Communication Pathways in Brain Functional Systems. *NeuroImage*, (April):1–13, 2017. ISSN 10538119. doi: 10.1016/j.neuroimage.2017.04.015.
24. Thomas Abeel, Thibault Helleputte, Yves Van de Peer, Pierre Dupont, and Yvan Saeys. Robust biomarker identification for cancer diagnosis with ensemble feature selection methods. *Bioinformatics*, 26(3):392–398, 2010. doi: 10.1093/bioinformatics/btp630.
25. Ludmila I. Kuncheva. A stability index for feature selection. In *Proceedings of the 25th Conference on Proceedings of the 25th IASTED International Multi-Conference: Artificial Intelligence and Applications*, AIAP'07, pages 390–395, Anaheim, CA, USA, 2007. ACTA Press.
26. Asa Ben-Hur, Cheng Soon Ong, Sören Sonnenburg, Bernhard Schölkopf, and Gunnar Rätsch. Support vector machines and kernels for computational biology. *PLOS Computational Biology*, 4(10):1–10, 10 2008. doi: 10.1371/journal.pcbi.1000173.
27. P Hagmann, O Sporns, N Madan, L Cammoun, R Pienaar, V J Wedeen, R Meuli, J.-P. Thiran, and P E Grant. White matter maturation reshapes structural connectivity in the late developing human brain. *Proceedings of the National Academy of Sciences*, 107(44): 19067–19072, 11 2010. ISSN 0027-8424. doi: 10.1073/pnas.1009073107.