

NETWORK NEURO SCIENCE

an open access  journal



Citation: Renard, F., Heinrich, C., Bouthillon, M., Schenck, M., Schneider, F., Kremer, S., & Achard, S. (2021). A covariate-constraint method to map brain feature space into lower dimensional manifolds. *Network Neuroscience*, 5(1), 252–273. https://doi.org/10.1162/netn_a_00176

DOI: https://doi.org/10.1162/netn_a_00176

Supporting Information:
https://doi.org/10.1162/netn_a_00176
<https://github.com/renardfe/CCML>

Received: 8 May 2020
Accepted: 24 November 2020

Competing Interests: The authors have declared that no competing interests exist.

Corresponding Author:
Félix Renard
felixrenard@gmail.com

Handling Editor:
James Shine

Copyright: © 2020
Massachusetts Institute of Technology
Published under a Creative Commons
Attribution 4.0 International
(CC BY 4.0) license



RESEARCH

A covariate-constraint method to map brain feature space into lower dimensional manifolds

Félix Renard¹, Christian Heinrich², Marine Bouthillon², Maleka Schenck^{3,4}, Francis Schneider^{3,4,5}, Stéphane Kremer^{2,6}, and Sophie Achard¹

¹Université Grenoble Alpes, CNRS, Inria, Grenoble, France

²iCube, Université de Strasbourg, CNRS, Illkirch, France

³Service de Médecine Intensive Réanimation, CHU de Strasbourg, France

⁴Faculté de Médecine FMTS, Strasbourg, France

⁵U1121, Université de Strasbourg, France

⁶Imagerie 2, CHU de Strasbourg, Université de Strasbourg, France

Keywords: Graphs, Machine learning, Connectomes, Hub disruption index

ABSTRACT

Human brain connectome studies aim to both explore healthy brains, and extract and analyze relevant features associated with pathologies of interest. Usually this consists of modeling the brain connectome as a graph and using graph metrics as features. A fine brain description requires graph metrics computation at the node level. Given the relatively reduced number of patients in standard cohorts, such data analysis problems fall in the high-dimension, low-sample-size framework. In this context, our goal is to provide a machine learning technique that exhibits flexibility, gives the investigator an understanding of the features and covariates, allows visualization and exploration, and yields insight into the data and the biological phenomena at stake. The retained approach is dimension reduction in a manifold learning methodology; the originality is that the investigator chooses one (or several) reduced variables. The proposed method is illustrated in two studies. The first one addresses comatose patients; the second one compares young and elderly populations. The method sheds light on the differences between brain connectivity graphs using graph metrics and potential clinical interpretations of these differences.

AUTHOR SUMMARY

Human brain connectome studies aim at both exploring healthy brains, and analyzing relevant features associated to pathologies of interest. This consists in modeling the brain connectome as a graph and in using graph metrics as features. Such data analysis problems fall in the high-dimension low sample size framework. Our goal is to provide a machine learning technique that exhibits flexibility, gives the investigator grip on the features and covariates, allows visualization, and yields insight into the biological phenomena at stake. The retained approach is dimension reduction in a manifold learning methodology, the originality lying in that several reduced variables be chosen by the investigator. The method sheds light on the differences between brain connectivity graphs using graph metrics and potential clinical interpretations of these differences.

INTRODUCTION

Brain modeling and understanding is a very active field of research involving different disciplines, such as neuroscience, image and signal processing, statistics, physics, and biology. These last years, neuroimaging modalities have been developed to explore the brain for both structural and functional features. It is now recognized that these images are providing very promising noninvasive observations of the brain (Bullmore & Sporns, 2009; Mwangi, Tian, & Soares, 2014; Richiardi, Achard, Bunke, & Van De Ville, 2013). One consequence of the availability of such massive datasets is the need to develop more and more sophisticated models to unravel the possible alteration of brains due to the impact of different pathologies. In this context, representing the brain as a global system is capital. This may be achieved using a *network* (Bullmore & Sporns, 2009). A brain network is a graph where nodes correspond to specific regions, and edges describe interactions and links between those regions. Different kinds of links and interactions may be of interest. Anatomical tracts are identified using diffusion imaging (Sporns, Tononi, & Kötter, 2005) and used in anatomical connectivity studies, where the whole set of links is called an *anatomical connectome*. Functional interactions are identified in functional imaging studies, whether in resting state or in task performing (Fallani, Richiardi, Chavez, & Achard, 2014; Rosazza & Minati, 2011), and used in functional connectivity studies. The whole set of functional links is called a *functional connectome*. In the functional case, brain networks are particularly adequate in encapsulating both spatial and temporal information in a single model. Indeed, brain networks are constructed using brain parcellation, namely spatial features, and time series interactions, namely temporal features. This model has attracted lots of attention these last 20 years by providing both very intuitive and spatial maps of brain networks.

Brain networks can be quantified using graph metrics such as minimum path length, clustering (Watts & Strogatz, 1998), global and local efficiency (Latora & Marchiori, 2001), modularity (Newman, 2006), and assortativity (Newman, 2002), among others. As these metrics are associated with specific network features, it is often possible to find the appropriate metrics to use given specific neuroscience hypotheses of the study. For the study of brain disorders, these metrics have been used in order to extract biomarkers for pathologies such as for example Alzheimer's disease (Supekar, Menon, Rubin, Musen, & Greicius, 2008), schizophrenia (Lynall et al., 2010), and multiple sclerosis (Filippi et al., 2014). Extracting quantitative parameters of brain networks is compulsory to conduct any statistical analysis. In this framework, statistical and machine learning approaches on graph metrics on all nodes allow the quantification of differences between groups (Richiardi et al., 2013).

For any dataset, any graph metric can be computed either at the global level with one value for an entire network or at the nodal level with one value for each node and a vector of values for the entire network. It has already been shown that global values may not discriminate two groups of subjects (Achard et al., 2012), which shows their limits as biomarkers. Few attempts have been made to directly use distances between networks such as the edit distance (Mokhtari & Hossein-Zadeh, 2013), or network similarities (Mheich et al., 2017). However, nodal-level approaches are challenging since hundreds of brain areas can be extracted whereas the number of subjects is generally small. This corresponds to the high-dimension, low-sample-size (HDLSS) configuration and falls under the curse of dimensionality (Bellman, 1961). In particular, standard classification and regression algorithms are not robust anymore in such a context (Hastie, Tibshirani, & Friedman, 2001, Chapter 2, Section 5, and Chapter 18).

Dimension-reduction techniques tackle curse of dimensionality issues (Hastie et al., 2001). In this framework, feature selection, where a subset of the original variables is considered, and

High-dimension, low-sample-size (HDLSS) framework:
HDLSS is a mathematical concept to characterize a key challenge in data analysis, mainly the difference between the number of parameters and the number of observations.

ISOMAP methodology:
ISOMAP is a dimension-reduction methodology consisting of projecting the initial observations to a reduced space, where it is simpler to describe and interpret the results.

κ or hub disruption index:
 κ is a graph metric to quantify the differences of locations of hubs and non-hubs between a patient and a group of controls.

feature extraction, where the original variables are transformed to a smaller set, may be envisaged (Webb, 2002). We resort here to the ISOMAP methodology, which is a well-known non-linear feature extraction algorithm generalizing principal component analysis (PCA) dimension reduction (Huo, Ni, & Smith, 2007; Tenenbaum, de Silva, & Langford, 2000). ISOMAP may be seen as a manifold learning approach, where the degrees of freedom of the data are captured by the latent variables, and where the structure of points in the latent space (the reduced space) mimics the structure of data in the original space. Nevertheless, ISOMAP raises two issues: interpreting the latent variables and determining the effect a change in the latent variables incurs in the data space; that is, the corresponding changes in brain networks and the underlying neuroscience hypotheses at stake in the case of the present study.

Dimension reduction is not new in the field of brain connectivity studies. Several methods have been proposed to extract nodal features at the level of brain regions. Using the hub disruption index (the κ index) to analyze a set of brain networks may be considered a feature extraction approach: This is a user-defined transformation of the original space to a 1D latent space (Achard et al., 2012). Principal component analysis was previously applied on graph metrics in Robinson, Hammers, Ericsson, Edwards, and Rueckert (2010) with vectors representing brains at the nodal level. We proposed in Renard, Heinrich, Achard, Hirsch, and Kremer (2012) to use kernel PCA, a nonlinear version of PCA. Moreover, interpreting latent variables may be addressed by correlating the reduced space with clinical data (Gerber, Tasdizen, Thomas Fletcher, Joshi, & Whitaker, 2010). Covariates may also be mapped or regressed on the reduced space as proposed in Aljabar, Wolz, and Rueckert (2012), thus shedding light on latent variables. Dimension-reduction methods have also been applied to connectivity matrices (Ktena et al., 2018; Kumar, Toews, Chauvin, Colliot, & Desrosiers, 2018; Yamin et al., 2019) or to the voxel time series (Saggar et al., 2018), mainly for classification purposes. It is indeed difficult using the whole connectivity matrices or voxel time series to give an interpretation at the nodal or voxel level (Gallos & Siettos, 2017; Haak, Marquand, & Beckmann, 2018; Laurienti et al., 2019). Network embedding framework can be viewed as a dimension-reduction method and was also applied to brain connectivity graphs (Rosenthal et al., 2018).

The objective of this article is to integrate all features cited above in one method: working at the nodal level, applying dimension-reduction techniques, and mapping covariates to ease interpretation. In addition, a new methodology is proposed to incorporate interesting network features already identified in specific datasets directly in the manifold learning approach. Contrary to statistical tests at nodal levels where each feature is treated independently of others, our approach based on machine learning is able to analyze joint variations between local descriptors.

This paper is focusing on two already published datasets. The first one consists of fMRI datasets on 20 healthy controls and 17 coma patients from Achard et al. (2012). The second one is based on Achard and Bullmore (2007), where 15 young healthy subjects and 11 elderly healthy subjects were scanned using resting-state fMRI. Our first experiment compares data-driven approaches such as linear discriminant analysis (LDA) and random forests (RF) to an ad hoc description such as the hub disruption index κ . This allows us to compare classical machine learning approaches, in which the interpretability of the results is often difficult, with approaches resorting to descriptors constructed using neuroscientific hypotheses. This first experiment can be seen as preliminaries of the sequel of the paper, where a feature is extracted for each individual in order to optimize classification of the two groups using either classical machine learning approaches or ad hoc descriptors. The second experiment consists of constructing a data-driven manifold, ISOMAP, using the graph metrics as features. ISOMAP

is providing a compact representation of brain connectomes in a reduced space, where it is straightforward to map the available covariates. In addition, we may interpret changes in connectomes by regressing covariables like κ on the reduced space using latent variables.

This representation allows a visualization of each subject relative to the whole population, which is crucial in clinical studies, for example in order to better understand brain changes for each specific subject. In addition, κ has been shown to be both a meaningful descriptor and a good classifying feature for brain connectomes of coma patients. Therefore, we propose a new method based on a covariate-constrained manifold learning (CCML) using κ as an input of ISOMAP. This allows us to propose a new generative model based on our new data representation, to better predict the variation in each patient given the changes of covariables. Based on the results of the first experiment, the choice of the covariate, κ in this work, can be adjusted to the studied datasets.

MATERIALS AND METHODS

Resting-State fMRI Data

Comatose study. The data were acquired in a previous study aimed at characterizing resting-state connectivity brain networks for patients with consciousness disorders. The description of the data and results is reported in Achard et al. (2012). The patients were scanned a few days after major acute brain injury, when sedative drug withdrawal allowed for spontaneous ventilation. Therefore, all patients were spontaneously ventilating and could be safely scanned at the time of fMRI. The causes of coma are patient-dependent: 12 had cardiac and respiratory arrest due to various causes; 2 had a gaseous cerebrovascular embolism; 2 had hypoglycemia; and 1 had extracranial artery dissection. A total of 25 patients were scanned (age range, 21–82 years; 9 men). Data on eight patients were subsequently excluded because of unacceptable degrees of head movement. The coma severity for each patient was clinically assessed using the 62 items of the WHIM scale: scores range from 0, meaning deep coma, to 62, meaning full recovery. Six months after the onset of coma, 3 patients had totally recovered, 9 patients had died, and 5 patients remained in a persistent vegetative state. The normal control group is composed of 20 healthy volunteers matched to the group of patients for sex (11 men) and approximately for age (range, 25–51 years). This study was approved by the local Research Ethics Committee of the Faculty of Health Sciences of Strasbourg on October 24, 2008, (CPP 08/53) and by the relevant healthcare authorities. Written informed consent was obtained directly from the healthy volunteers and from the next of kin for each of the patients. Resting-state data were acquired for each subject using gradient echo planar imaging technique with a 1.5-T MR scanner (Avanto; Siemens, Erlangen, Germany) with the following parameters: relaxation time = 3 s, echo time = 50 ms, isotropic voxel size = $4 \times 4 \times 4 \text{ mm}^3$, 405 images, and 32 axial slices covering the entire cortex. The preprocessing of the data is detailed in our previous study (Achard et al., 2012).

Young and elderly study. The data used in this study have already been analyzed in two papers (Achard & Bullmore, 2007, and Meunier, Achard, Morcom, & Bullmore, 2009). The goal of these papers was to identify the changes in brain connectomes for elderly subjects in terms of topological organization of brain graphs. The data consist of 15 young subjects aged 18–33 years, mean age = 24 and 11 elderly subjects aged 62–76 years. Each subject was scanned using resting-state fMRI as described in Achard and Bullmore (2007, Wolfson Brain Imaging Centre, Cambridge, UK). For each dataset, a total of 512 volumes were available with number of slices, 21 (interleaved); slice thickness, 4 mm; interslice gap, 1 mm; matrix size, 64×64 ;

flip angle, 90°; repetition time (TR), 1,100 ms; echo time, 27.5 ms; in-plane resolution, 3.125 mm.

Preprocessing and Wavelet Graph Estimation

Brain network graphs were determined following Achard et al. (2012) for the comatose study and Achard and Bullmore (2007) for the young and elderly study. For each subject, data were corrected for head motion and then coregistered with each subject’s T1-weighted structural MRI. Each subject’s structural MRI was nonlinearly registered with the Colin27 template image. The obtained deformation field image was used to map the fMRI datasets to the automated anatomical labeling (AAL) or to a customized parcellation image with 417 anatomically homogeneous size regions based on the AAL template image (Tzourio-Mazoyer et al., 2002). Regional mean time series were estimated by averaging the fMRI time series over all voxels in each parcel, weighted by the proportion of gray matter in each voxel of the segmented structural MRIs. We estimated the correlations between wavelet coefficients of all possible pairs of the $N = 90$ or 417 cortical and subcortical fMRI time series extracted from each individual dataset. For the coma, only scale 3, 0.02–0.04 Hz, wavelet correlation matrices were considered. For the young and elderly, the wavelet scale considered corresponds to 0.06–0.11 Hz. The choice of these wavelet scales or frequency bands is explained precisely in the corresponding papers (Achard & Bullmore, 2007; Achard et al., 2012). To generate binary undirected graphs, a minimum spanning tree algorithm was applied to connect all parcels. The absolute wavelet correlation matrices were thresholded to retain 2.5% of all possible connections. Each subject was then represented by a graph with nodes corresponding to the same brain regions, and with the same number of edges.

Graph Metrics

The objective is to extract differences between the two groups with respect to the topological organization of the graphs. Each graph is summarized by graph metrics computed at the nodal level. Three metrics are considered here: degree, global efficiency, and clustering (Bullmore & Sporns, 2009).

The degree is quantifying the number of edges belonging to one node. Let G denote a graph with $G_{ij} = 0$ when there is no edge between nodes i and j , and $G_{ij} = 1$ when there is an edge between nodes i and j . The degree D_i of node i is computed as

$$D_i = \sum_{j \in G, j \neq i} G_{ij}. \tag{1}$$

The global efficiency measures how the information is propagating in the whole network. A random graph will have a global efficiency close to 1 for each node, and a regular graph will have a global efficiency close to 0 for each node. The global efficiency E_{glob} is defined as the inverse of the harmonic mean of the set of the minimum path lengths L_{ij} between node i and all other nodes j in the graph:

$$E_{glob_i} = \frac{1}{N-1} \sum_{j \in G} \frac{1}{L_{ij}}. \tag{2}$$

Clustering is a local efficiency measure corresponding to information transfer in the immediate neighborhood of each node, defined as

$$Clust_i = \frac{1}{N_{G_i}(N_{G_i} - 1)} \sum_{j, k \in G_i, j \neq k} \frac{1}{L_{jk}}, \tag{3}$$

where G_i is the subgraph of G defined by the set of nodes that are the nearest neighbors of node i . A high value of clustering corresponds to highly connected neighbors of each node, whereas a low value means that the neighbors of each node are rather disconnected.

Each graph metric emphasizes a specific property at the nodal level. With a view to statistical comparison, several methods have already been developed, representing data in specific spaces. Each method aims to separate classes. Usually these methods are very general and can be applied without careful inspection of the data. We used here four different methods (Richiardi, Achard, Bullmore, & Van De Ville, 2011): the κ index resulting from a careful inspection of the data, mean over graph metrics (denoted here MEAN), LDA, and feature selection (FS) by selecting the best feature based on a univariate statistical student t test. Like the κ index, each of these methods provides, for each patient, a scalar feature corresponding to a particular property of the data. Figure 1 gives an illustration of the different methods.

κ Index Definition

In our previous study (Achard et al., 2012), κ was devised to compare graph metrics obtained on each node of a subject or of a group with a reference group (see Figure 2). In classical comparisons between a group of patients and a group of healthy volunteers, the reference is the group of healthy volunteers. In the present study, for a given graph metric and two groups, we first compute the average of this metric for each node over the group of healthy volunteers, denoted as the reference. Each subject is then summarized as a vector of values of dimension of the number of nodes. Then, for each patient, κ corresponds to the slope of the regression of a nodal graph metric between the given patient minus the reference and the




Figure 1. General framework from graphs of cerebral connectomes to the different scalar features. Brain connectivity graphs are extracted from fMRI data. Graph metrics are computed at the nodal level for each subject. The matrices of graph metrics can then be analyzed using different methods: the hub disruption index based on regression analyses (κ); linear discriminant analysis (LDA); average of metrics (MEAN); and feature selection (FS). Each of these methods allows us to summarize the graph metric in one scalar for each subject in order to better differentiate the studied populations.

Downloaded from http://direct.mit.edu/neh/article-pdf/5/1/252/1911574/neh_a_00176.pdf by guest on 07 September 2023

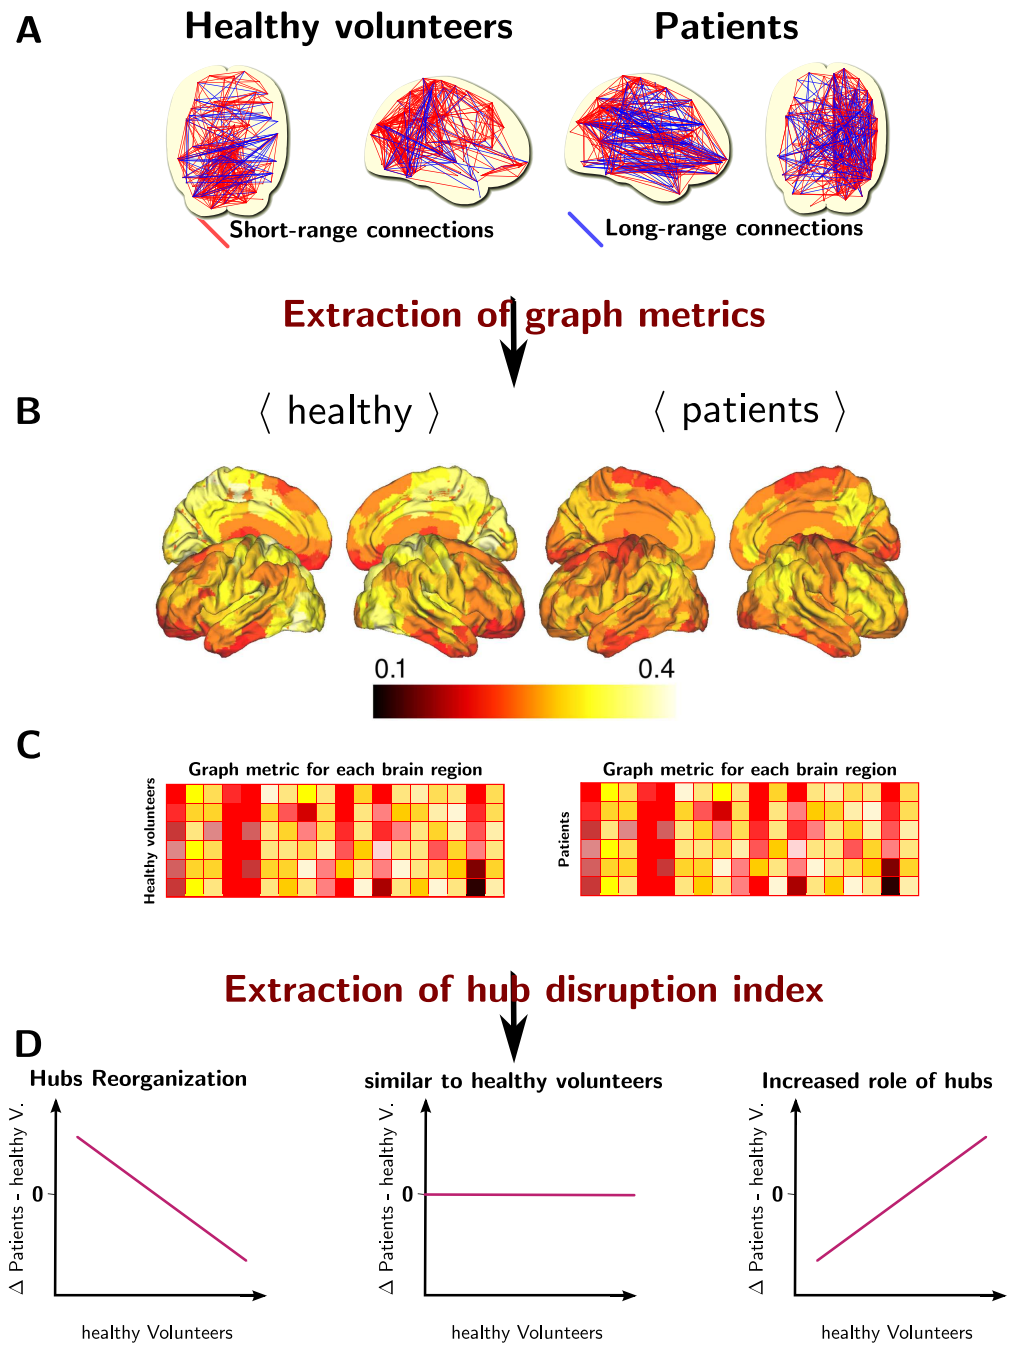


Figure 2. Extraction of hub disruption index κ . (A) Brain connectomes inferred for each subjects. (B) For each brain connectome, extraction of graph metrics for each region of the brain. (C) Matrix representation of the graph metrics, where a row corresponds to a subject and a column corresponds to a brain region. (D) Computation of the hub disruption index by regressing the average of brain metrics of the difference of patients and average of healthy volunteers against the average of healthy volunteers. The hub disruption index corresponds to the slope coefficient. We give several illustrations following the sign of this coefficient.

reference. Let N denote the number of nodes in the graph, n_p the number of patients, and n_c the number of controls. Let $(m_1, \dots, m_{n_p}) \in \mathbb{R}^{N \times n_p}$ denote a matrix of graph metric extracted given the graphs of patients, for $j, 1 \leq i \leq n_p, m_j \in \mathbb{R}^N$. For each j , m_j is equal to one graph metric such as D , $Eglob$, or $Clust$. Let us also define a similar matrix for the controls, $(h_1, \dots, h_{n_c}) \in \mathbb{R}^{N \times n_c}$. Let us define the average metric for controls, for each $i, 1 \leq i \leq N$,

$$\bar{h}_i = \frac{1}{n_c} \sum_{j=1}^{n_c} h_{ij}. \quad (4)$$

κ is defined by the following regression:

$$m_i - h_i = \kappa \bar{h}_i + \epsilon_i, \quad (5)$$

where ϵ_i is the classical error term in linear regression. In order to give a simple interpretation of κ , we assume that the global graph metric computed as an average over the nodes is the same in both groups. A value of 0 for κ is showing that the graph metric obtained at the node level is the same for the patient and the reference. A positive value of κ is indicating that the hubs and non-hubs of the patient in comparison to the reference are located on the same nodes. However, the values of the graph metrics are increased for the hubs and decreased for the non-hubs. Finally, when the value of κ is negative, the hubs of the reference are no longer hubs of the patient, and the non-hubs of the reference are hubs for the patient. In Achard et al. (2012), we showed that the κ index is able to discriminate both groups (coma patients and healthy volunteers) while the global metric is unable to identify any significant difference. Instead of averaging the graph metrics, the κ index is capturing a joint variation of the metrics computed for each node.

Mean Over the Nodes (MEAN)

For each graph metric, the mean over the nodes of the graph captures a global property of the network. These global metrics have been previously used to discriminate two populations of networks, for example for Alzheimer's disease (Supekar et al., 2008) and for schizophrenia (Lynall et al., 2010). Such a coefficient can discriminate two networks well when their topologies are really different. However, such metrics do not take into account the specificity of the nodes. Indeed, when permuting the nodes of the graph, the global metric is not changed, but the hubs of the graph are not associated with the same nodes anymore. Therefore, a graph reorganization cannot be detected using such global metrics.

Linear Discriminant Analysis (LDA)

LDA (Fisher, 1936) is a classification method that aims to identify the linear projection optimally separating two groups. It can be considered a gold standard for linear group discrimination. It is not specific to the analysis of networks.

LDA has been previously used for network discrimination in Robinson et al. (2010). This algorithm amounts to computing a scalar for each graph. However, there is no simple clinical interpretation of the discriminating parameter.

Feature Selection (FS)

As for LDA, FS determines the features yielding the best separation of the two groups. Several features may be used simultaneously. In order to establish a fair comparison with the other methods, we choose to extract the single feature yielding the best separation. Several methods

exist for FS. We choose univariate FS implemented in Pedregosa et al. (2011). An advantage of FS is that it is capturing discriminative features at the node level. As the selected features are extracted directly from the data, it is usually possible to derive a clinical interpretation. However, joint variations are not modeled, and on the comatose study FS is not able to yield results of the same quality as those obtained using κ .

Modeling Populations of Networks with Manifold Learning

ISOMAP (Tenenbaum et al., 2000) is used as a manifold learning approach to describe population networks. We propose an original approach based on ISOMAP, where we constrain one variable of the reduced space (the latent space) to correspond to a covariate.

Manifold learning using ISOMAP. ISOMAP devises a reduced-dimension version of the original set of points. Interpoint distances in the reduced space reproduce as much as possible interpoint distances in the original space. Euclidean and geodesic distances are respectively used. Principal component analysis may be seen as a particular case of ISOMAP, where Euclidean distances are used in the original space, instead of geodesic distances. The reader is referred to Tenenbaum et al. (2000) for details about the algorithm.

In our case, the original data correspond to a vector of graph metrics for each subject, the dimension of the vector being the number of nodes times the number of metrics. For each analysis, only one metric is considered here. However, this method could be applied by jointly using several metrics. Covariates may be regressed on the reduced space. In the present work, this was achieved using a classic radial basis function interpolation.

The choice of the ISOMAP is twofold: First, the estimated reduced space is a smooth manifold, and preserves the global structure of the dataset. Notably, the reduced space exhibits a continuum function of subjects. Second, the cost function of the ISOMAP allows the integration of additional constrained scores. ISOMAP was performed by computing a nearest neighbor graph connecting the four nearest neighbors according to the Euclidean distance. This distance correctly reflects the local topology of the graph metrics space. The choice of four neighbors is driven by the relatively small number of subjects in the study.

The classification score of the ISOMAP was computed using a nonlinear support vector machine (SVM) approach with radial basis function kernel in the reduced space (Hearst, Dumais, Osuna, Platt, & Scholkopf, 1998).

Covariate constrained manifold learning. One drawback of manifold learning algorithms is the difficulty of interpreting the reduced coordinates because they are usually meaningless. The original method proposed in this work consists of constraining one coordinate of the reduced space to correspond to a specific covariate. The other coordinates are left unconstrained, as in classical ISOMAP. Such a procedure requires special care regarding the optimization aspect. We apply a strategy proposed in Brucher, Heinrich, Heitz, and Armspach (2008), where points are introduced one by one.

Moreover, a scale factor α is considered for the axis corresponding to the covariate. This parameter, obtained by optimization, balances the scales of the different axes.

The reduced point $\tilde{\mathbf{x}}_i$ is defined by $\tilde{\mathbf{x}}_i = [\alpha c_i; \mathbf{x}_i]^T$, where c_i is the chosen covariate and \mathbf{x}_i are the other coordinates. The cost function E is defined as

$$E = \sum_{i,i < j} (\|\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j\|^2 - \|\mathbf{y}_i - \mathbf{y}_j\|^2)^2, \tag{6}$$

where $\{\mathbf{y}_i\}_{i=1,\dots,N}$ is the graph metric vectors over N graph nodes. For an incoming data point i , the cost function E is optimized three times with regard to (a) \mathbf{x}_i as $\min_{\mathbf{x}_i} E$, (b) α as $\min_{\alpha} E$, and (c) \mathbf{x}_j for each point that has already been included as $\min_{\{\mathbf{x}_j\}_{j=1,\dots,i-1}} E$. We consider $i < j$ in the sum of the cost function to avoid counting the errors between two samples twice.

The distance in the cost function is the Euclidean one. Since the samples are added sample by sample, this distance reflects only the local neighborhood of the new added one.

To facilitate optimization and to avoid possible local minima, instead of inserting the samples at random, we choose the sample to be incorporated next as the one with the largest geodesic distance to the samples already incorporated. Indeed, interpolation problems are always easier than extrapolation problems where greater uncertainty may occur. We initialize the procedure by taking the two samples with the largest geodesic distance. The first two samples are used as landmarks of the border of the reduced space, and the insertion of new samples will generate only small displacements of the already inserted samples.

The algorithm is described in Algorithm 1 and available at <https://github.com/renardfe/CCML> (Renard, 2020).

Algorithm 1 Covariate-constrained manifold learning (CCML)

Input – dataset: N vectors (samples) $\{\mathbf{y}_i\}_{i=1,\dots,N}$ of a graph metric over n graph nodes
Result: Reduced-space representation $\{\tilde{\mathbf{x}}_i\}_{i=1,\dots,N}$ of the dataset, where the first coordinate of each $\tilde{\mathbf{x}}_i$ corresponds to the covariate.

Initialization: Select the two most distant samples.

Determine their reduced coordinates by minimizing E with $\alpha = 1$.

Update the scale α by minimizing E wrt α , \mathbf{x}_i fixed, as $\min_{\alpha} E$.

While all points are not included **do** the following:

(1) **Select the most distant sample \mathbf{y}_k** to the already selected samples.

(2) **Compute \mathbf{x}_k** by minimizing E wrt \mathbf{x}_k (α and other \mathbf{x}_i 's fixed) as $\min_{\mathbf{x}_k} E$.

(3) **Update the scale α** by minimizing E wrt α (\mathbf{x}_i 's fixed) as $\min_{\alpha} E$.

(4) **Update \mathbf{x}_j 's of samples already included** by minimizing E as $\min_{\{\mathbf{x}_j\}_{j=1,\dots,k-1}} E$.

End while

Application: A Generative Model for the Prediction of the Variation in a Subject with Regard to the Changes of a Covariate

From the obtained embedding, a generative model

$$\hat{\mathbf{y}} = f(\tilde{\mathbf{x}}) \tag{7}$$

can be devised, where $\hat{\mathbf{y}}$ is a vector in the original space (the connectome space), $\tilde{\mathbf{x}}$ is a vector from the manifold embedding, and f is a regression function. Multivariate adaptive regression splines (MARS) (Friedman, 1991) is chosen for the regression function f for its nice properties (one regression for each coordinate of f , i.e., n regressions): locally linear and globally

nonlinear. The parameters of f can be determined using the dataset $\{\mathbf{y}_i\}_{i=1,\dots,N}$ and the corresponding reduced vectors $\{\tilde{\mathbf{x}}_i\}_{i=1,\dots,N}$ using the following equation:

$$\mathbf{y}_i = f(\tilde{\mathbf{x}}_i) + \epsilon_i = \hat{\mathbf{y}}_i + \epsilon_i, \tag{8}$$

where ϵ_i is the residual between a sample and its prediction $\hat{\mathbf{y}}_i$. The residuals allow us to evaluate the accuracy of the regression function.

This kind of model is not original, PCA being the most well-known case where the model is defined as $\mathbf{y} = \mathbf{A} \tilde{\mathbf{x}} + \epsilon$ (see, e.g., Lawrence, 2004; Sfikas & Nikou, 2016, for references). Such a generative model used in the CCML framework allows us to determine changes in the original space (the connectome space) generated by a displacement in the reduced space, for example along the covariate axis.

RESULTS

The different algorithms have been implemented in the Python language using the scikit-learn toolbox (Pedregosa et al., 2011). When left unspecified, coma data are used. The use of the young and elderly data is explicitly stated.

Local Analysis Using Dimension Reduction

Permutation tests are performed on the κ index and on the three other measures (LDA, FS, MEAN) to assess the ability of those four metrics to discriminate two populations. More precisely, for each coefficient separately, the difference of the means of the two populations is determined for the observed populations. The labels of the samples are then shuffled and the difference of the means of the shuffled two populations is determined. This latter step is performed 10^4 times. This method allows us to avoid making any assumptions on the distribution of the statistic. Simultaneously, the correlations between the observed κ index and the other coefficients are estimated.

The results corresponding to the different methods that aim to discriminate the two groups (control and coma) are given in Table 1. As expected, the machine learning algorithms (i.e., LDA and FS) show good performances in separating the two groups for different graph metrics. This is consistent with the fact that these methods have been tailored to classify the two groups. The results with the κ index show similar performances in separating the two groups. The large

Table 1. P value of permutation tests comparing the mean of the two groups (10^4 permutations, which bounds the p values). The correlation scores are estimated between the κ index and the three other measures (LDA, FS, and MEAN).

Mean diff. or correlation	<i>Eglob</i> (p value)	corr. (p value)	<i>Clust</i> (p value)	corr. (p value)	<i>D</i> (p value)	corr. (p value)
κ	0.79 ($< 10^{-4}$)		0.75 ($< 10^{-4}$)		0.81 ($< 10^{-4}$)	
LDA	-2.89 ($< 10^{-4}$)	0.88 ($< 10^{-4}$)	-1.78 ($< 10^{-4}$)	0.87 ($< 10^{-4}$)	-2.39 ($< 10^{-4}$)	0.88 ($< 10^{-4}$)
FS	0.12 ($< 10^{-4}$)	0.60 (10^{-4})	-0.50 ($< 10^{-4}$)	-0.66 (10^{-4})	21.93 ($< 10^{-4}$)	0.60 (10^{-4})
MEAN	0.14 (0.58)	0.25 (0.78)	-0.01 (0.43)	-0.19 (0.85)		

correlations between machine learning algorithms on the one hand and the graph metric κ on the other hand show retrospectively that similar performances were to be expected.

Moreover, a strong relationship can be observed between κ and LDA (correlation scores greater than 0.87 for each metric). The FS correlation scores are lower than the LDA correlation ones. The difference between the two methods is that LDA considers a linear combination of features, with a global perspective, whereas FS selects one feature and acts locally. Since κ reflects a global reorganization of the brain, it is expected that the correlation score of LDA will be greater than the FS one. Finally, MEAN scores reveal that this measure is not appropriate in this study.

Standard ISOMAP Manifold Learning and the κ Index

In this section, the goal is to link the reduced space obtained by manifold learning to different covariates such as the κ index. We want to assess whether a given covariate varies smoothly across the reduced space, and is therefore predictable using this space.

Figure 3 represents the reduced space obtained using standard ISOMAP, as opposed to using CCML. The values of the different covariates are color coded. The reduced-space representation allows us to separate both populations. By visual inspection of the color-coded maps, it appears that those regression maps are capturing features corresponding to κ and MEAN.

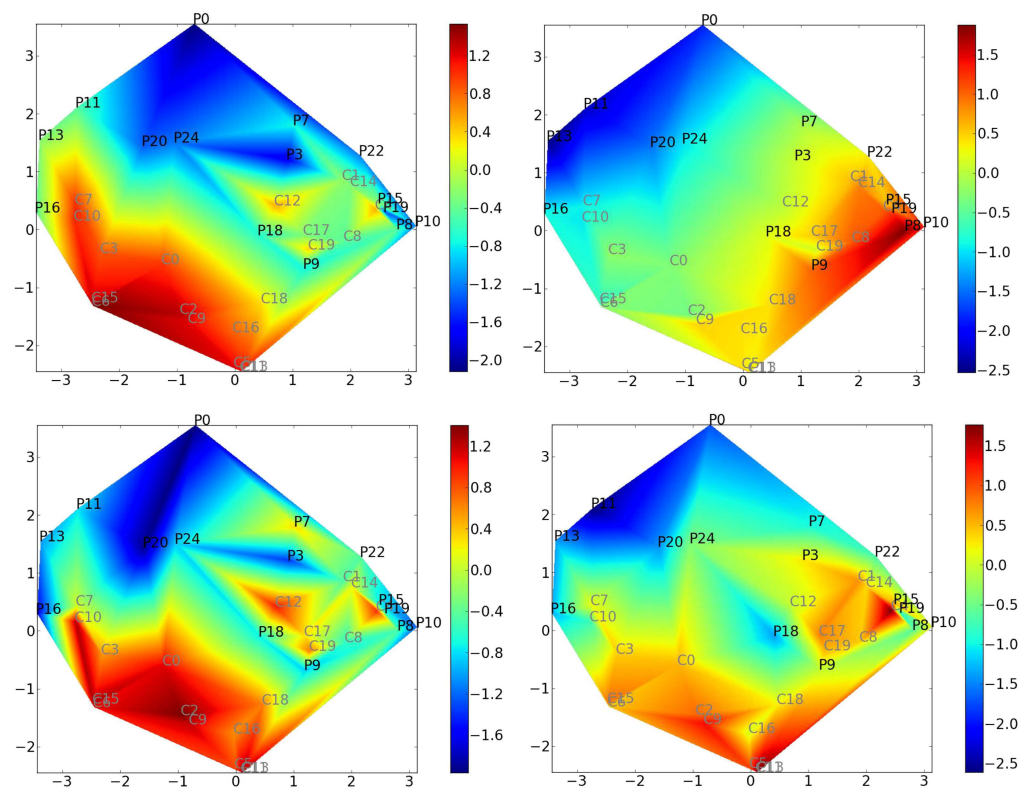


Figure 3. Standard ISOMAP reduced-space representation of the original dataset. P_i : (comatose) patient # i ; C_j : control # j . Covariates are mapped onto the reduced space (covariate value is color coded). Top left: κ index mapping; top right: MEAN mapping; bottom left: LDA mapping; bottom right: FS mapping.

Table 2. Assessment of the regression of covariates on the reduced space. Three different reduced spaces are at stake, one for each graph metric. Root mean square error (RMSE) and maximal error M are displayed. The MEAN strategy is not relevant for the degree D since the degrees D for all graphs are equal (the number of edges is set to be the same for all graphs).

Covariate	<i>Eglob</i>	<i>Eglob</i>	<i>Clust</i>	<i>Clust</i>	<i>D</i>	<i>D</i>
	RMSE	M	RMSE	M	RMSE	M
κ	0.51	2.44	0.68	2.59	0.26	2.14
LDA	0.97	7.97	0.90	5.44	0.66	4.13
FS	0.64	2.20	1.10	9.12	0.83	7.49
MEAN	0.15	0.73	1.02	5.44		

In order to quantify these visual observations, covariates are regressed on the reduced space. The root mean square error (RMSE) and the maximum error M are determined in a leave-one-out procedure. The results are given in Table 2. It can be noted that the MEAN strategy gives the same values for the graph metric degree D for all graphs since the number of edges is set to be the same for all graphs.

In this table, the lower the values, the better the adequacy with the reduced space. It appears that κ is the best choice across all metrics, except for *Eglob* where it is outperformed by MEAN. In the case of *Eglob*, this suggests that both κ and MEAN scores correspond to degrees of freedom of the intrinsic manifold of the functional connectivity graphs.

Figure 4 displays the probability of belonging to a specific class computed using logistic regression on the reduced space stemming from standard ISOMAP. The probability of belonging to the comatose class is color coded in Figure 4.




Figure 4. Logistic regression using reduced space stemming from standard ISOMAP. The color codes the probability of belonging to the comatose class.

Table 3. Correlation scores between the probabilistic mapping and the different coefficients (κ index and MEAN measure). A high correlation score of the κ index indicates a good fitting between the reduced-space representation and the classification of the two groups. The p values are all smaller than 10^{-12} .

Coefficients	<i>Eglob</i>	<i>Clus</i>	<i>D</i>
κ	0.87	0.87	0.81
MEAN	0.55	0.42	

This probability estimation using logistic regression is compared with covariates such as κ or MEAN, in Table 3. A high correlation score is observed between κ and logistic regression probability. The correlation score between the MEAN coefficient and the probabilistic mapping is lower than the one with κ , as expected.

Taken together, these observations demonstrate the importance of κ in the classification of these populations. Obviously, for the special case of the global efficiency metric, the MEAN score correctly describes the reduced space, but it does not correspond to the classification pattern.

Covariate-Constrained Manifold Learning

Comatose population. First we evaluate the convergence of the optimization problem (Algorithm 1). To assess the difficulty of the optimization problem, we ran it with 500 random initializations. Only 63% of the runs converged to the same solution, whereas 37% of the runs converged to a local (worse) optimum. It thus appears that our initialization procedure addresses the local optimum issue. Nevertheless, this optimization problem would probably deserve further investigation, which is outside the scope of this paper.

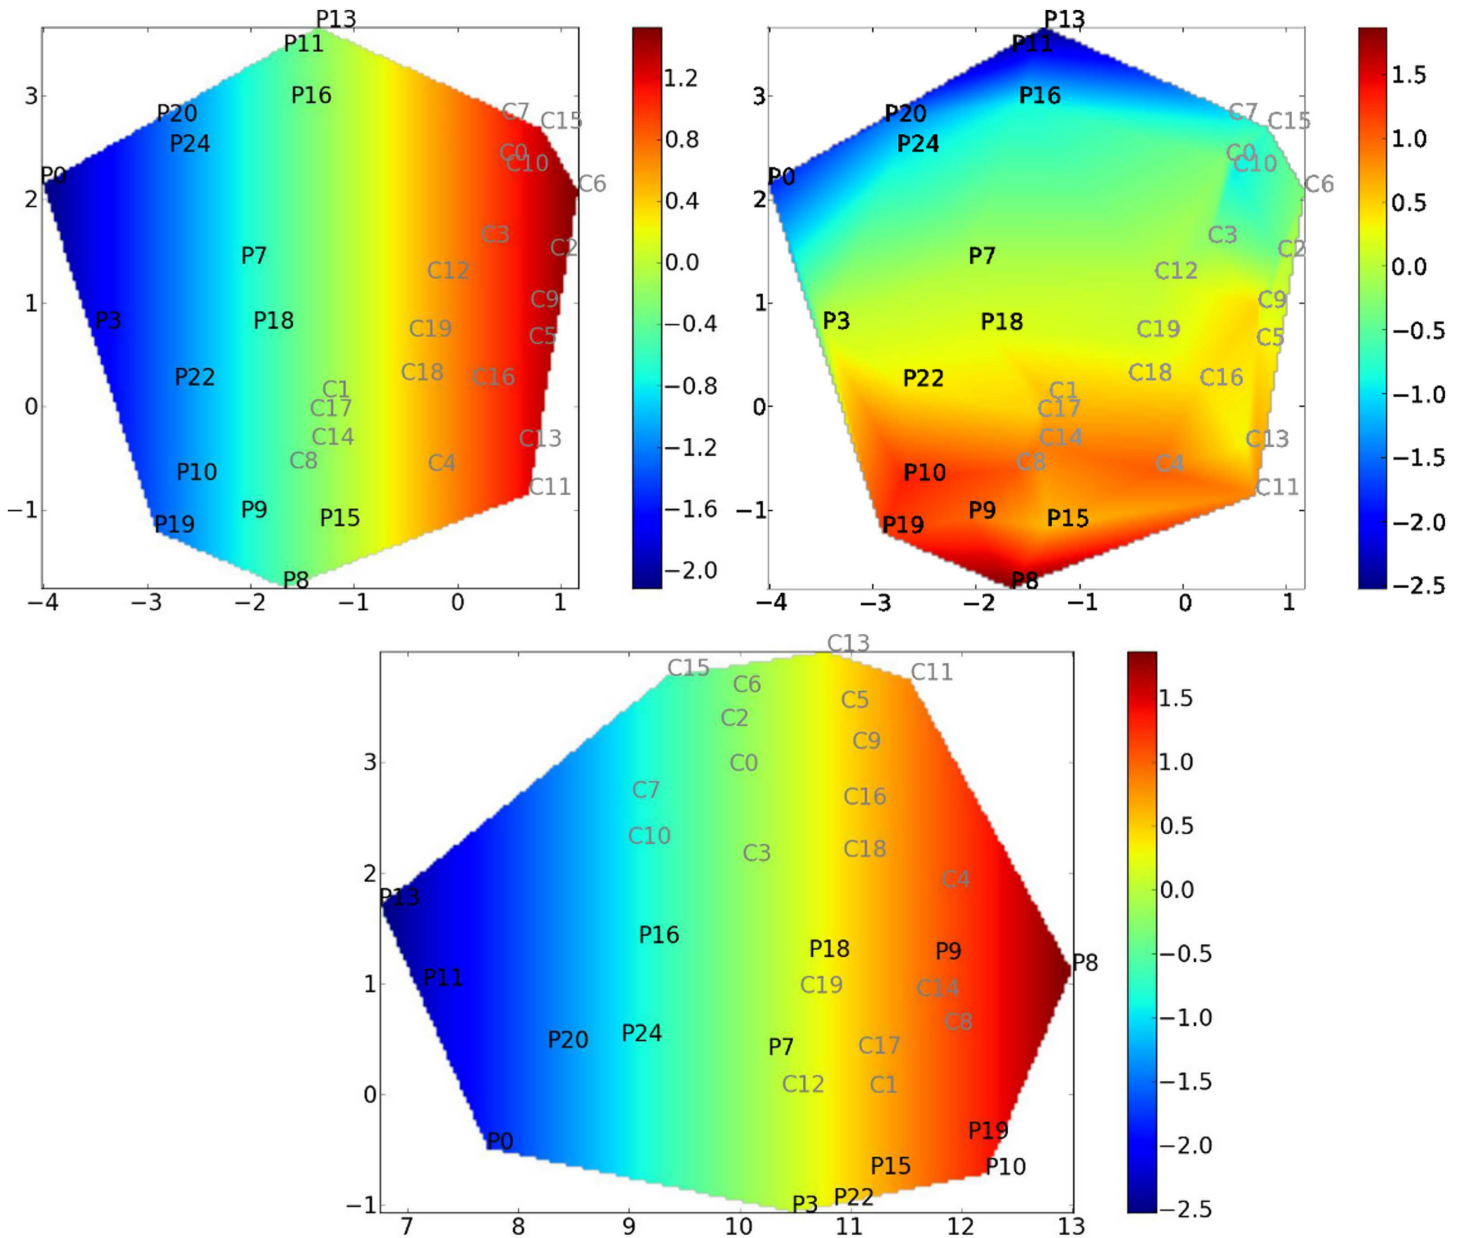
In Figure 5, we display the reduced space corresponding to our new manifold learning algorithm. We can observe that the two populations are well discriminated in the case of κ , but not for the MEAN coefficient. This is quantified by applying a classical SVM procedure in the reduced space. We obtained the following results: for CCML, 1; for ISOMAP, 0.86; for LDA, 0.91; for κ , 0.89; and for MEAN and FS, 0.57.

We observe strong interaction between κ and the MEAN in the top right of Figure 5, where the reduced space on one axis is κ and the mapping corresponds to the MEAN. To quantify this, in the case of κ , we estimate the correlation between the second coordinate and the MEAN score. The obtained score equals 0.92, which confirms the intrinsic relationship between κ and the MEAN coefficient.

Elderly and young population. The elderly and young groups are investigated in this section.

Figure 6 displays the manifold obtained by standard ISOMAP. It is interesting to highlight that the κ index is not a pertinent feature to discriminate the old from the young, whereas the MEAN is a better discriminating feature. In both cases, the interpretation of the mapping is complex since it is not smooth.

Figure 7 displays results from CCML for the MEAN coefficient. We can observe that the MEAN mapping discriminates the two groups.



Downloaded from http://direct.mit.edu/nn/article-pdf/5/1/252/1911574/nn_1_00176.pdf by guest on 07 September 2023

Figure 5. Covariate mapping onto the reduced space given by our method, CCML, using global efficiency as a graph metric. The reduced space is computed using E_{glob} as a graph metric (the y_i 's). Covariate value is color coded. For each subfigure, the coordinates correspond to $[ac_i; x_i]^T$, where c_i is the constrained variable and x_i is the free parameter. Top left: κ mapping with a κ -constrained reduced space. Top right: MEAN mapping with a κ -constrained reduced space. Bottom: MEAN mapping with a MEAN-constrained reduced space. As expected, we can observe that the mappings correlate well with the first coordinate by construction (top left and bottom). It is also clear that using a κ -constrained reduced space is facilitating the discrimination between the two populations. Indeed, the controls and patients are not covering the same part of the reduced space. On the contrary, as expected using the MEAN-constrained reduced space, the method is not providing a very clear discrimination between patients and controls. Especially, patients 9 and 18 are very close to controls. Finally, the top right figure is showing a correlation between the second coordinate of CCML and the MEAN.




Figure 6. Covariate mapping onto the standard ISOMAP reduced space using global efficiency as a graph metric. Left: κ mapping with the standard ISOMAP reduced space. Right: MEAN mapping with the same reduced space using global efficiency as a graph metric. The old controls (resp. young controls) are labeled O (resp. Y). For these groups, the κ index cannot discriminate the two groups. However, the MEAN index behaves better for the discrimination between the two groups. In each case, the interpretation is complex since the mapping of the covariate is not linear.




Figure 7. CCML approach using global efficiency as graph metric. Left: κ mapping with the κ -constrained reduced space. Right: MEAN mapping with the MEAN-constrained reduced space. The old controls (resp. young controls) are labeled O (resp. Y). Two manifolds (one for each constraint) have been determined. However, only the MEAN-constrained manifold discriminates both groups.

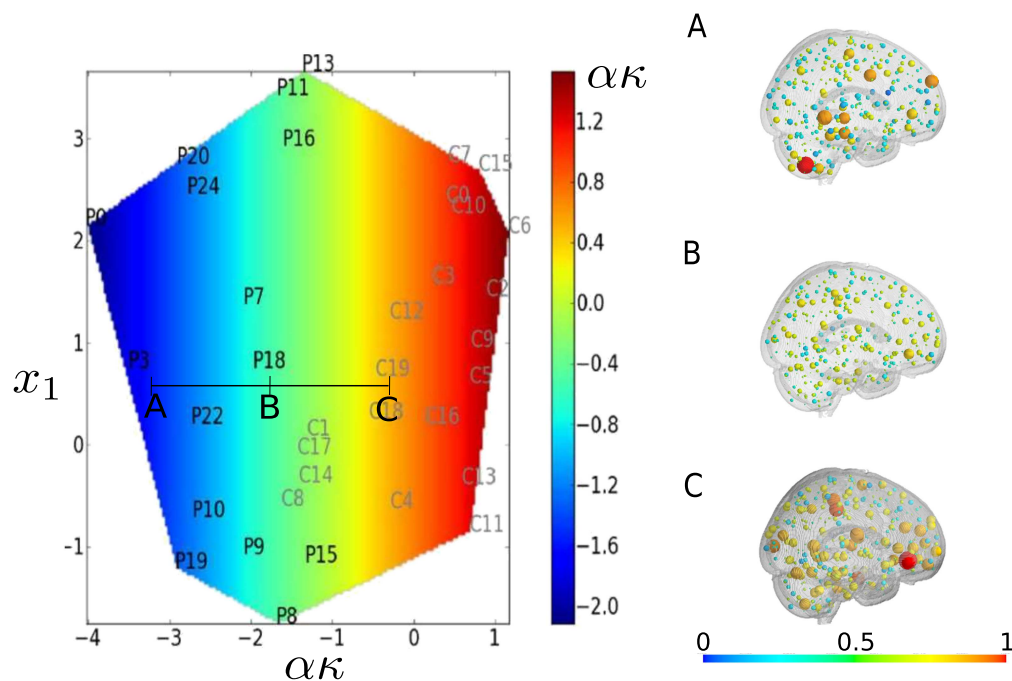


Figure 8. Variation in one patient along the covariate axis using global efficiency as a graph metric. The size and the color of the nodes on the right part of the figure are proportional to the graph metric, global efficiency, at each node. We consider the patient P18 in state B in the reduced space (left part of the figure). We predict its changes when the κ index is decreased (point A) or increased (point C). Intuitively, point A and point C correspond respectively to a degradation and to an improvement of the health of the patient. Interestingly, these trends can be directly observed in the graph space for clinical insights (right part of the figure). It can be noticed that the variations are not linear.

Application: A Generative Model for the Prediction of the Variation in a Subject with Regard to the Changes of a Covariate

Using the algorithm detailed in the last section, a map of the population is determined, with one of the reduced coordinates corresponding to a chosen covariate. To highlight the potential of the proposed method, we compute the transformation of a patient with regard to the changes of a covariate by creating a map from the reduced space back to the original space, where we can make brain-related interpretations. This gives us insight into the effect of the covariate on the patient. To perform this analysis, a regression is used to map the reduced space to the initial space (we used MARS regression [Friedman, 1991], coordinate-wise). This application is illustrated in Figure 8.

DISCUSSION

Assessment of Graph Metric Descriptors

The handcrafted design of graph metric descriptors is interesting since such descriptors carry straightforward physical meaning, like the κ index for hub reorganization. However, in a classification framework, such new scalar descriptors may not be optimal. To assess the pertinence of a new ad hoc graph metric descriptor for a classification task, it will be interesting to confront it with specific scalar coefficients used in standard classification algorithms (like LDA or FS), and examine whether there is some correlation between the scalars at stake. It is

also worth mentioning that other approaches are dedicated to classifications directly using the connectivity matrices (Dadi et al., 2019; Richiardi et al., 2013). Usually, the objective of these approaches is to obtain classification scores, that is assign a subject to a group, patient, or control. To our knowledge, these approaches do not enable exploration of the underlying brain mechanisms at the nodal level using the results of classification (Ktena et al., 2018; Kumar et al., 2018; Yamin et al., 2019). Graph metrics are known to capture important topological indicators in networks that are impossible to capture using classical data-mining approaches (Zanin et al., 2016) or network embedding (Rosenthal et al., 2018).

No Free Lunch for Graph Metric Descriptors

We hypothesize that there is no best descriptor adapted to all datasets. The optimality depends on (a) the kind of the data and (b) the kind of question/task addressed. This idea is known as the “no free lunch theorem” (Wolpert, 1996): If an algorithm performs well on a certain class of problems, then it has necessarily poorer performances on the set of all remaining problems.

In the present study, we showed that the κ index yields good classification performances in separating a comatose population from a healthy population. However, the MEAN index better describes the groups of elderly and young people (see Figure 3): For this dataset, the κ index cannot separate the two groups, but the MEAN score can. It is interesting to notice that several descriptors can correctly map a population, while providing different information.

The “no free lunch theorem” also applies to manifold learning algorithms. The underlying question is the one of choosing an interpoint distance in the data space. A given interpoint distance will yield a specific structure of samples in the reduced space. Therefore, the retained interpoint distance chosen will depend on the final goal, such as mimicking the structure of the initial data points, enhancing class separation with a view to achieving better classification performances, or focusing on a specific property of the data. The proposed algorithm CCML aims to mimic the structure of the initial data points by explicitly using a particular characteristic, chosen and imposed by the investigator.

Manifold Learning for Brain Graph Analysis

Manifold learning is well suited for brain graph analysis for several reasons. First, global descriptors of graph metrics represent an entire graph by a scalar value, which is generally ultimately insufficient to correctly model the complexity of a graph population. Manifold learning is better suited to capture the complexity and variability of a given graph population, since more degrees of freedom are structuring the reduced space.

Second, connectomes have been studied for their capacity to represent the brain as a system and not merely as a juxtaposition of independent anatomical/functional regions. Classical statistical tests are not adapted to analyze joint variations between local descriptors of graph metrics since those tests assume independence between features. Brain graph manifold learning for comparing groups of graphs is promising because joint variations are accounted for.

Third, manifold learning may be turned to a generative model, when resorting to a mapping from the reduced space to the data space. Brain graph manifold learning can be seen as a trade-off between global and local brain graph metrics analysis. In other words, manifold learning is considered to be a model at the level of the group while preserving the information of the individuals. However, this technique is hard to interpret on its own. The addition of explicative covariables as proposed with the CCML method can provide an understandable

and generative model of population with the possibility of focusing at the individual level (Costa et al., 2015). Using either global or local metrics is usually a hard task to appropriately link these features to clinical information. Statistical approaches suffer from a lack of interpretability where null and alternative hypotheses are tested. This is usually the case for coma studies where simple univariate statistical tests are computed on graph metrics (Demertzi et al., 2014; Malagurski et al., 2019). Using manifold learning, as illustrated in Figure 8, it is possible to provide a smooth description of the changes of the brain graphs of the patient in the reduced space. This approach can adequately relate the changes in brain connectivity along with the changes in clinical features. More generally, manifold learning can be an interesting solution for personalized and predictive medicine purposes. In our paper, we illustrate the result of our new proposed approach, CCML, on one graph metric, namely global efficiency. However, including several graph metrics is also a possibility and would lead to a more accurate description of the data, maybe at the cost of interpretability.

CONCLUSION

The originality and contribution of this paper is the devising of a nonlinear manifold model of brain graph metrics. The essence of the approach is the capture of a metric through all nodes of a graph, across all graphs of an entire population. A population of graphs is represented by a population of vectors, each vector holding the variation of a metric through the nodes at stake. The population is then represented in a reduced space, as opposed to the standard representation of a given brain graph by a mere scalar.

The proposed approach has several advantages. First and foremost, the data are represented with several degrees of freedom, corresponding to the dimensions of the reduced space. The structure of the original dataset is captured by a compact representation. This allows us to account for the complex variability of populations of brain graphs. Second, such an approach naturally offers analysis of joint variations of those brain graph metrics. In addition, the investigator has the possibility to analyze the data at the population scale and simultaneously at the individual scale.

The investigation tool corresponding to the proposed approach allowed us to retrospectively assess the hub disruption index (HDI), denoted κ , and proposed in one of our former works. Earlier work showed that κ is a very good candidate for discriminating patients and controls in the case of coma. Retrospectively, its performance is assessed here in comparison with machine learning methods dedicated to linear group classification, such as LDA. Besides yielding nice classification performances, the present study showed that an advantage of κ , put in the perspective of a manifold model, is to give clinical clues related to the pathology mechanism.

We observed strong relationships between scalar coefficients such as κ and MEAN, and the coordinates of the manifold. It is important to notice that MEAN, which can separate groups in several pathologies (Lynall et al., 2010; Supekar et al., 2008), is not able to discriminate the comatose patients from the normal population. However, it brings additional information in terms of description of the population. The manifold at stake shows that a scalar coefficient cannot capture all of the information encapsulated in the graphs. One interesting aspect of manifold learning, and more specifically our new proposed method, is its ability to reach a new level of interpretation of the brain graph metrics and the interaction between them.

ACKNOWLEDGMENTS

We want to thank the anonymous reviewers for their suggestions, which significantly improve the quality of our paper.

SUPPORTING INFORMATION

Supporting information for this article is available at https://doi.org/10.1162/netn_a_00176.
The algorithm is available at <https://github.com/renardfe/CCML>.

AUTHOR CONTRIBUTIONS

Félix Renard: Conceptualization; Methodology; Software; Validation; Visualization; Writing – original draft; Writing – review & editing. Christian Heinrich: Conceptualization; Formal analysis; Methodology; Software; Supervision; Validation; Visualization; Writing – original draft; Writing – review & editing. Marine Bouthillon: Data curation; Formal analysis; Methodology. Maleka Schenck: Conceptualization; Data curation; Funding acquisition; Investigation; Resources. Francis Schneider: Conceptualization; Data curation; Funding acquisition; Investigation; Resources. Stéphane Kremer: Conceptualization; Data curation; Funding acquisition; Investigation; Project administration; Resources. Sophie Achard: Conceptualization; Formal analysis; Investigation; Methodology; Project administration; Software; Supervision; Validation; Visualization; Writing – original draft; Writing – review & editing.

REFERENCES

- Achard, S., & Bullmore, E. (2007). Efficiency and cost of economical human brain functional networks. *PLoS Computational Biology*, 3, e17. DOI: <https://doi.org/10.1371/journal.pcbi.0030017>, PMID: 17274684, PMCID: PMC1794324
- Achard, S., Delon-Martin, C., Vértes, P. E., Renard, F., Schenck, M., Schneider, F., ... Bullmore, E. T. (2012). Hubs of brain functional networks are radically reorganized in comatose patients. *Proceedings of the National Academy of Sciences*, 109(50), 20608–20613. DOI: <https://doi.org/10.1073/pnas.1208933109>, PMID: 23185007, PMCID: PMC3528500
- Aljabar, P., Wolz, R., & Rueckert, D. (2012). Manifold learning for medical image registration, segmentation and classification. In K. Suzuki (Ed.), *Machine learning in computer-aided diagnosis: Medical imaging intelligence and analysis* (pp. 351–372). IGI Global. DOI: <https://doi.org/10.4018/978-1-4666-0059-1.ch017>
- Bellman, R. E. (1961). *Adaptive control processes - A guided tour*. Princeton University Press. DOI: <https://doi.org/10.1515/9781400874668>
- Brucher, M., Heinrich, C., Heitz, F., & Armspach, J. (2008). A metric multidimensional scaling-based nonlinear manifold learning approach for unsupervised data reduction. *EURASIP Journal on Advances in Signal Processing*, 862015. DOI: <https://doi.org/10.1155/2008/862015>
- Bullmore, E., & Sporns, O. (2009). Complex brain networks: Graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3), 186–198. DOI: <https://doi.org/10.1038/nrn2575>, PMID: 19190637
- Costa, L., Smith, J., Nichols, T., Cussens, J., Duff, E. P., & Makin, T. R. (2015). Searching multiregression dynamic models of resting-state fMRI networks using integer programming. *Bayesian Analysis*, 10(2), 441–478. DOI: <https://doi.org/10.1214/14-BA913>
- Dadi, K., Rahim, M., Abraham, A., Chyzyk, D., Milham, M., Thirion, B., & Varoquaux, G. (2019). Benchmarking functional connectome-based predictive models for resting-state fMRI. *NeuroImage*, 192, 115–134. DOI: <https://doi.org/10.1016/j.neuroimage.2019.02.062>, PMID: 30836146
- Demertzi, A., Gomez, F., Crone, J. S., Vanhauudenhuysse, A., Tshibanda, L., Noirhomme, Q., ... Soddu, A. (2014). Multiple fMRI system-level baseline connectivity is disrupted in patients with consciousness alterations. *Cortex*, 52, 35–46. DOI: <https://doi.org/10.1016/j.cortex.2013.11.005>, PMID: 24480455
- Fallani, F. D. V., Richiardi, J., Chavez, M., & Achard, S. (2014). Graph analysis of functional brain networks: Practical issues in translational neuroscience. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1653), 20130521. DOI: <https://doi.org/10.1098/rstb.2013.0521>, PMID: 25180301, PMCID: PMC4150298
- Filippi, M., Valsasina, P., Sala, S., Martinelli, V., Ghezzi, A., Veggioni, P., ... Rocca, M. (2014). Abnormalities of the brain functional connectome in pediatric patients with multiple sclerosis. *Neurology*, 82(10 Suppl.).
- Fisher, R. A. (1936). The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7(2), 179–188. DOI: <https://doi.org/10.1111/j.1469-1809.1936.tb02137.x>
- Friedman, J. (1991). Multivariate adaptive regression splines. *Annals of Statistics*, 19(1), 1–67. DOI: <https://doi.org/10.1214/aos/1176347963>
- Gallos, I., & Siettos, C. (2017). Classification of fMRI resting-state maps using machine learning techniques: A comparative study. *AIP Conference Proceedings*, 1906, 200001. DOI: <https://doi.org/10.1063/1.5012477>
- Gerber, S., Tasdizen, T., Thomas Fletcher, P., Joshi, S., & Whitaker, R. (2010). Manifold modeling for brain population analysis. *Medical Image Analysis*, 14(5), 643–653. DOI: <https://doi.org/10.1016/j.media.2010.05.008>, PMID: 20579930, PMCID: PMC3020141
- Haak, K. V., Marquand, A. F., & Beckmann, C. F. (2018). Connectopic mapping with resting-state fMRI. *NeuroImage*, 170, 83–94. DOI: <https://doi.org/10.1016/j.neuroimage.2017.06.075>, PMID: 28666880

- Hastie, T., Tibshirani, R., & Friedman, J. (2001). *The elements of statistical learning*. Springer. DOI: <https://doi.org/10.1007/978-0-387-21606-5>
- Hearst, M. A., Dumais, S. T., Osuna, E., Platt, J., & Scholkopf, B. (1998). Support vector machines. *IEEE Intelligent Systems and Their Applications*, 13(4), 18–28. DOI: <https://doi.org/10.1109/5254.708428>
- Huo, X., Ni, X. S., & Smith, A. K. (2007). A survey of manifold-based learning methods. In *Recent advances in data mining of enterprise data*. (pp. 691–745). World Scientific. DOI: https://doi.org/10.1142/9789812779861_0015
- Ktena, S. I., Parisot, S., Ferrante, E., Rajchl, M., Lee, M., Glocker, B., & Rueckert, D. (2018). Metric learning with spectral graph convolutions on brain connectivity networks. *NeuroImage*, 169, 431–442. DOI: <https://doi.org/10.1016/j.neuroimage.2017.12.052>, PMID: 29278772
- Kumar, K., Toews, M., Chauvin, L., Colliot, O., & Desrosiers, C. (2018). Multi-modal brain fingerprinting: A manifold approximation based framework. *NeuroImage*, 183, 212–226. DOI: <https://doi.org/10.1016/j.neuroimage.2018.08.006>, PMID: 30099077
- Latora, V., & Marchiori, M. (2001). Efficient behavior of small-world networks. *Physical Review Letters*, 87, 198701. DOI: <https://doi.org/10.1103/PhysRevLett.87.198701>, PMID: 11690461
- Laurienti, P. J., Bahrani, M., Lyday, R. G., Casanova, R., Burdette, J. H., & Simpson, S. L. (2019). Using low-dimensional manifolds to map relationships between dynamic brain networks. *Frontiers in Human Neuroscience*, 13, 430. DOI: <https://doi.org/10.3389/fnhum.2019.00430>, PMID: 31920590, PMCID: PMC6914694
- Lawrence, N. D. (2004). Gaussian process latent variable models for visualisation of high dimensional data. In *Advances in neural information processing systems* (Vol. 16, pp. 329–336). MIT Press.
- 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. *Journal of Neuroscience*, 30(28), 9477–9487. DOI: <https://doi.org/10.1523/JNEUROSCI.0333-10.2010>, PMID: 20631176, PMCID: PMC2914251
- Malagurski, B., Péran, P., Sarton, B., Vinour, H., Naboulsi, E., Riu, B., ... Silva, S. (2019). Topological disintegration of resting state functional connectomes in coma. *NeuroImage*, 195, 354–361. DOI: <https://doi.org/10.1016/j.neuroimage.2019.03.012>, PMID: 30862533
- Meunier, D., Achard, S., Morcom, A., & Bullmore, E. (2009). Age-related changes in modular organization of human brain functional networks. *NeuroImage*, 44(3), 715–723. DOI: <https://doi.org/10.1016/j.neuroimage.2008.09.062>, PMID: 19027073
- Mheich, A., Hassan, M., Khalil, M., Gripon, V., Dufor, O., & Wendling, F. (2017). Siminet: A novel method for quantifying brain network similarity. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(9), 2238–2249. DOI: <https://doi.org/10.1109/TPAMI.2017.2750160>, PMID: 28910755
- Mokhtari, F., & Hossein-Zadeh, G.-A. (2013). Decoding brain states using backward edge elimination and graph kernels in fMRI connectivity networks. *Journal of Neuroscience Methods*, 212(2), 259–268. DOI: <https://doi.org/10.1016/j.jneumeth.2012.10.012>, PMID: 23142223
- Mwangi, B., Tian, T. S., & Soares, J. C. (2014). A review of feature reduction techniques in neuroimaging. *Neuroinformatics*, 12(2), 229–244. DOI: <https://doi.org/10.1007/s12021-013-9204-3>, PMID: 24013948, PMCID: PMC4040248
- Newman, M. E. J. (2002). Assortative mixing in networks. *Physical Review Letters*, 89, 208701. DOI: <https://doi.org/10.1103/PhysRevLett.89.208701>, PMID: 12443515
- Newman, M. E. J. (2006). Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23), 8577–8582. DOI: <https://doi.org/10.1073/pnas.0601602103>, PMID: 16723398, PMCID: PMC1482622
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.
- Renard, F. (2020). Code for the estimation of the covariate constrained manifold learning. GitHub, <https://github.com/renardfe/CCML>
- Renard, F., Heinrich, C., Achard, S., Hirsch, E., & Kremer, S. (2012). Statistical kernel-based modeling of connectomes. In *International Workshop on Pattern Recognition in NeuroImaging* (pp. 69–72). DOI: <https://doi.org/10.1109/PRNI.2012.22>
- Richiardi, J., Achard, S., Bullmore, E., & Van De Ville, D. (2011). Classifying connectivity graphs using graph and vertex attributes. In *International Workshop on Pattern Recognition in NeuroImaging* (pp. 45–48). DOI: <https://doi.org/10.1109/PRNI.2011.18>
- Richiardi, J., Achard, S., Bunke, H., & Van De Ville, D. (2013). Machine learning with brain graphs: Predictive modeling approaches for functional imaging in systems neuroscience. *IEEE Signal Processing Magazine*, 30(3), 58–70. DOI: <https://doi.org/10.1109/MSP.2012.2233865>
- Robinson, E. C., Hammers, A., Ericsson, A., Edwards, A. D., & Rueckert, D. (2010). Identifying population differences in whole-brain structural networks: A machine learning approach. *NeuroImage*, 50(3), 910–919. DOI: <https://doi.org/10.1016/j.neuroimage.2010.01.019>, PMID: 20079440
- Rosazza, C., & Minati, L. (2011). Resting-state brain networks: Literature review and clinical applications. *Neurological Sciences*, 32(5), 773–785. DOI: <https://doi.org/10.1007/s10072-011-0636-y>, PMID: 21667095
- Rosenthal, G., Váša, F., Griffa, A., Hagmann, P., Amico, E., Goñi, J., ... Sporns, O. (2018). Mapping higher-order relations between brain structure and function with embedded vector representations of connectomes. *Nature Communications*, 9(1), 1–12. DOI: <https://doi.org/10.1038/s41467-018-04614-w>, PMID: 29872218, PMCID: PMC5988787
- Saggar, M., Sporns, O., Gonzalez-Castillo, J., Bandettini, P. A., Carlsson, G., Glover, G., & Reiss, A. L. (2018). Towards a new approach to reveal dynamical organization of the brain using topological data analysis. *Nature Communications*, 9(1), 1–14. DOI: <https://doi.org/10.1038/s41467-018-03664-4>, PMID: 29643350, PMCID: PMC5895632
- Sikas, G., & Nikou, C. (2016). Bayesian multiview manifold learning applied to hippocampus shape and clinical score data. In *Medical computer vision and Bayesian and graphical models for biomedical imaging – MICCAI 2016* (pp. 160–171). Springer. DOI: https://doi.org/10.1007/978-3-319-61188-4_15

- Sporns, O., Tononi, G., & Kötter, R. (2005). The human connectome: A structural description of the human brain. *PLoS Computational Biology*, 1(4), e42. DOI: <https://doi.org/10.1371/journal.pcbi.0010042>, PMID: 16201007, PMCID: PMC1239902
- 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*, 1–11. DOI: <https://doi.org/10.1371/journal.pcbi.1000100>, PMID: 18584043, PMCID: PMC2435273
- Tenenbaum, J. B., de Silva, V., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500), 2319–2323. DOI: <https://doi.org/10.1126/science.290.5500.2319>, PMID: 11125149
- Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., . . . Joliot, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage*, 15(1), 273–289. DOI: <https://doi.org/10.1006/nimg.2001.0978>, PMID: 11771995
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of “small-world” networks. *Nature*, 393(6684), 440. DOI: <https://doi.org/10.1038/30918>, PMID: 9623998
- Webb, A. (2002). *Statistical pattern recognition* (2nd ed.). Wiley. DOI: <https://doi.org/10.1002/0470854774>
- Wolpert, D. H. (1996). The lack of a priori distinctions between learning algorithms. *Neural Computation*, 8(7), 1341–1390. DOI: <https://doi.org/10.1162/neco.1996.8.7.1341>
- Yamin, A., Dayan, M., Squarcina, L., Brambilla, P., Murino, V., Diwadkar, V., . . . Sona, D. (2019). Comparison of brain connectomes using geodesic distance on manifold: A twins study. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)* (pp. 1797–1800). DOI: <https://doi.org/10.1109/ISBI.2019.8759407>
- Zanin, M., Papo, D., Sousa, P. A., Menasalvas, E., Nicchi, A., Kubik, E., & Boccaletti, S. (2016). Combining complex networks and data mining: Why and how. *Physics Reports*, 635, 1–44. DOI: <https://doi.org/10.1016/j.physrep.2016.04.005>