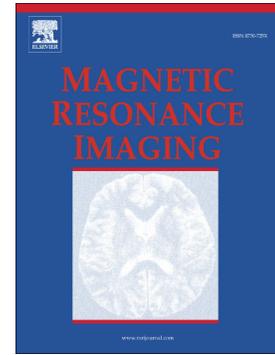


Accepted Manuscript

Machine learning in resting-state fMRI analysis

Meenakshi Khosla, Keith Jamison, Gia H. Ngo, Amy Kuceyeski,
Mert R. Sabuncu



PII: S0730-725X(18)30685-4
DOI: <https://doi.org/10.1016/j.mri.2019.05.031>
Reference: MRI 9225
To appear in: *Magnetic Resonance Imaging*
Received date: 29 December 2018
Revised date: 20 May 2019
Accepted date: 21 May 2019

Please cite this article as: M. Khosla, K. Jamison, G.H. Ngo, et al., Machine learning in resting-state fMRI analysis, *Magnetic Resonance Imaging*, <https://doi.org/10.1016/j.mri.2019.05.031>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Machine learning in resting-state fMRI analysis

Meenakshi Khosla^a, Keith Jamison^b, Gia H. Ngo^a, Amy Kuceyeski^{b,c}, Mert R. Sabuncu^{1a,d}

^a*School of Electrical and Computer Engineering, Cornell University*

^b*Radiology, Weill Cornell Medical College*

^c*Brain and Mind Research Institute, Weill Cornell Medical College*

^d*Nancy E. & Peter C. Meinig School of Biomedical Engineering, Cornell University*

Abstract

Machine learning techniques have gained prominence for the analysis of resting-state functional Magnetic Resonance Imaging (rs-fMRI) data. Here, we present an overview of various unsupervised and supervised machine learning applications to rs-fMRI. We offer a methodical taxonomy of machine learning methods in resting-state fMRI. We identify three major divisions of unsupervised learning methods with regard to their applications to rs-fMRI, based on whether they discover principal modes of variation across space, time or population. Next, we survey the algorithms and rs-fMRI feature representations that have driven the success of supervised subject-level predictions. The goal is to provide a high-level overview of the burgeoning field of rs-fMRI from the perspective of machine learning applications.

Keywords: Machine learning, resting-state, functional MRI, intrinsic networks, brain connectivity

1. Introduction

Resting-state fMRI (rs-fMRI) is a widely used neuroimaging tool that measures spontaneous fluctuations in neural blood oxygen-level dependent (BOLD) signal across the whole brain, in the absence of any controlled experimental paradigm. In their seminal work, Biswal et al. [1] demonstrated

¹Please address correspondence to: Dr. Mert R. Sabuncu (msabuncu@cornell.edu), School of Electrical and Computer Engineering, 300 Frank H. T. Rhodes Hall, Cornell University, Ithaca, NY 14853.

temporal coherence of low-frequency spontaneous fluctuations between long-range functionally related regions of the primary sensory motor cortices even in the absence of an explicit task, suggesting a neurological significance of resting-state activity. Several subsequent studies similarly reported other collections of regions co-activated by a task (such as language, motor, attention, audio or visual processing etc.) that show correlated fluctuations at rest [2, 3, 4, 5, 6, 7, 8, 9, 10, 11]. These spontaneously co-fluctuating regions came to be known as the resting state networks (RSNs) or intrinsic brain networks. The term RSN henceforth denotes brain networks subserving shared functionality as discovered using rs-fMRI.

Rs-fMRI has enormous potential to advance our understanding of the brain's functional organization and how it is altered by damage or disease. A major emphasis in the field is on the analysis of resting-state functional connectivity (RSFC) that measures statistical dependence in BOLD fluctuations among spatially distributed brain regions. Disruptions in RSFC have been identified in several neurological and psychiatric disorders, such as Alzheimer's [12, 13, 14], autism [15, 16, 17], depression [18, 19, 20], schizophrenia [21, 22], etc. Dynamics of RSFC have also garnered considerable attention in the last few years, and a crucial challenge in rs-fMRI is the development of appropriate tools to capture the full extent of this RS activity. rs-fMRI captures a rich repertoire of intrinsic mental states or spontaneous thoughts and, given the necessary tools, has the potential to generate novel neuroscientific insights about the nature of brain disorders [23, 24, 25, 26, 27, 28].

The study of rs-fMRI data is highly interdisciplinary, majorly influenced by fields such as machine learning, signal processing and graph theory. Machine learning methods provide a rich characterization of rs-fMRI, often in a data-driven manner. Unsupervised learning methods in rs-fMRI are focused primarily on understanding the functional organization of the healthy brain and its dynamics. For instance, methods such as matrix decomposition or clustering can simultaneously expose multiple functional networks within the brain and also reveal the latent structure of dynamic functional connectivity.

Supervised learning techniques, on the other hand, can harness RSFC to make individual-level predictions. Substantial effort has been devoted to using rs-fMRI for classification of patients versus controls, or to predict disease prognosis and guide treatments. Another class of studies explores the extent to which individual differences in cognitive traits may be predicted by differences in RSFC, yielding promising results. Predictive approaches can

also be used to address research questions of interest in neuroscience. For example, is RSFC heritable? Such questions can be formulated within a prediction framework to test novel hypotheses.

From mapping functional networks to making individual-level predictions, the applications of machine learning in rs-fMRI are far-reaching. The goal of this review is to present in a concise manner the role machine learning has played in generating pioneering insights from rs-fMRI data, and describe the evolution of machine learning applications in rs-fMRI. We will present a review of the key ideas and application areas for machine learning in rs-fMRI rather than delving into the precise technical nuances of the machine learning algorithms themselves. In light of the recent developments and burgeoning potential of the field, we discuss current challenges and promising directions for future work.

1.1. Resting-state fMRI: A Historical Perspective

Until the 2000s, task-fMRI was the predominant neuroimaging tool to explore the functions of different brain regions and how they coordinate to create diverse mental representations of cognitive functions. The discovery of correlated spontaneous fluctuations within known cortical networks by Biswal et al. [1] and a plethora of follow-up studies have established rs-fMRI as a useful tool to explore the brain's functional architecture. Studies adopting the resting-state paradigm have grown at an unprecedented scale over the last decade. These are much simpler protocols than alternate task-based experiments, capable of providing critical insights into functional connectivity of the healthy brain as well as its disruptions in disease. Resting-state is also attractive as it allows multi-site collaborations, unlike task-fMRI that is prone to confounds induced by local experimental settings. This has enabled network analysis at an unparalleled scale.

Traditionally, rs-fMRI studies have focused on identifying spatially-distinct yet functionally associated brain regions through seed-based analysis (SBA). In this approach, seed voxels or regions of interest are selected *a priori* and the time series from each seed is correlated with the time series from all brain voxels to generate a series of correlation maps. SBA, while simple and easily interpretable, is limited since it is heavily dictated by manual seed selection and, in its simplest form, can only reveal one specific functional system at a time.

Decomposition methods like Independent Component Analysis (ICA) emerged as a highly promising alternative to seed-based correlation analysis in the

early 2000s [29, 2, 30]. This was followed by other unsupervised learning techniques such as clustering. In contrast to seed-based methods that explore networks associated with a seed voxel (such as motor or visual functional connectivity maps), these new class of model-free methods based on decomposition or clustering explored RSNs simultaneously across the whole brain for individual or group-level analysis. Regardless of the analysis tool, all studies largely converged in reporting multiple robust resting-state networks across the brain, such as the primary sensorimotor network, the primary visual network, fronto-parietal attention networks and the well-studied default mode network. Regions in the default mode network, such as the posterior cingulate cortex, precuneus, ventral and dorsal medial prefrontal cortex, show increased levels of activity during resting-state suggesting that this network represents the baseline or default functioning of the human brain. The default mode network has sparked a lot of interest in the rs-fMRI community [31], and several studies have consequently explored disruptions in DMN resting-state connectivity in various neurological and psychiatric disorders, including autism, schizophrenia and Alzheimer's. [32, 33, 34]

Despite the widespread success and popularity of rs-fMRI, the causal origins of ongoing spontaneous fluctuations in the resting brain remain largely unknown. Several studies explored whether resting-state coherent fluctuations have a neuronal origin, or are just manifestations of aliasing or physiological artifacts introduced by the cardiac or respiratory cycle. Over time, evidence in support for a neuronal basis of BOLD-based resting state functional connectivity has accumulated from multiple complementary sources. This includes (a) observed reproducibility of RSFC patterns across independent subject cohorts [5, 4], (b) its persistence in the absence of aliasing and distinct separability from noise components [5, 35], (c) its similarity to known functional networks [1, 2, 11] and (d) consistency with anatomy [36, 37], (e) its correlation with cortical activity studied using other modalities [38, 39, 40] and finally, (f) its systematic alterations in disease [23, 24, 25].

1.2. Application of Machine Learning in rs-fMRI

A vast majority of literature on machine learning for rs-fMRI is devoted to unsupervised learning approaches. Unlike task-driven studies, modelling resting-state activity is not straightforward since there is no controlled stimuli driving these fluctuations. Hence, analysis methods used for characterizing the spatio-temporal patterns observed in task-based fMRI are generally not suited for rs-fMRI. Given the high dimensional nature of fMRI data,

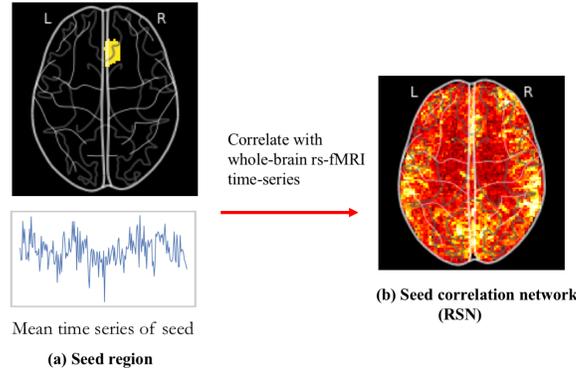


Figure 1: Traditional seed based analysis approach

it is unsurprising that early analytic approaches focused on decomposition or clustering techniques to gain a better characterization of data in spatial and temporal domains. Unsupervised learning approaches like ICA catalyzed the discovery of the so-called resting-state networks or RSNs. Subsequently, the field of resting-state brain mapping expanded with the primary goal of creating brain parcellations, i.e., optimal groupings of voxels (or vertices in the case of surface representation) that describe functionally coherent spatial compartments within the brain. These parcellations aid in the understanding of human functional organization by providing a reference map of areas for exploring the brain’s connectivity and function. Additionally, they serve as a popular data reduction technique for statistical analysis or supervised machine learning.

More recently, departing from the stationary representation of brain networks, studies have shown that RSFC exhibits meaningful variations during the course of a typical rs-fMRI scan [41, 42]. Since brain activity during resting-state is largely uncontrolled, this makes network dynamics even more interesting. Using unsupervised pattern discovery methods, resting-state patterns have been shown to transition between discrete recurring functional connectivity “states”, representing diverse mental processes [42, 43, 44]. In the simplest and most common scenario, dynamic functional connectivity is expressed using sliding-window correlations. In this approach, functional connectivity is estimated in a temporal window of fixed length, which is subsequently shifted by different time steps to yield a sequence of correlation matrices. Recurring correlation patterns can then be identified from this sequence through decomposition or clustering. This dynamic nature of func-

tional connectivity opens new avenues for understanding the flexibility of different connections within the brain as they relate to behavioral dynamics, with potential clinical utility [45].

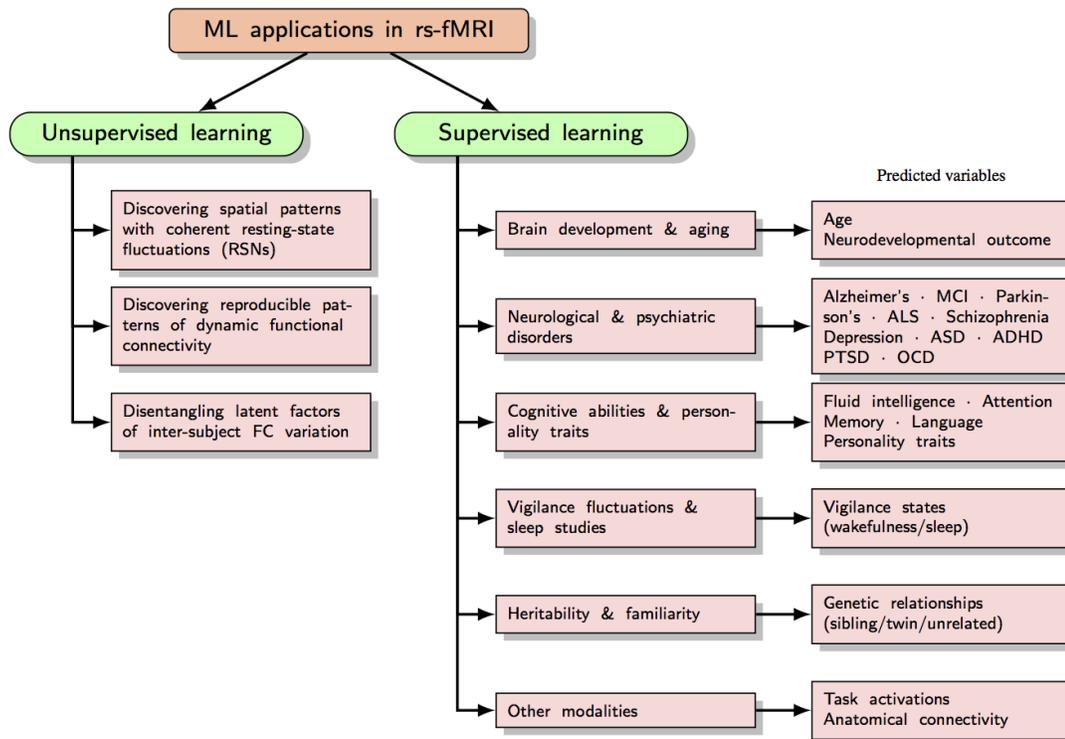


Figure 2: Applications of machine learning methods in resting-state fMRI

Another, perhaps clinically more promising application of machine learning in rs-fMRI expanded in the late 2000s. This new class of applications leveraged supervised machine learning for individual level predictions. The covariance structure of resting-state activity, more popularly known as the "connectome", has garnered significant interest in the field of neuroscience as a sensitive biomarker of disease. Studies have further shown that an individual's connectome is unique and reliable, akin to a fingerprint [46]. Machine learning can exploit these neuroimaging based biomarkers to build diagnostic or prognostic tools. Visualization and interpretation of these models can complement statistical analysis to provide novel insights into the dysfunction

of resting-state patterns in brain disorders. Given the prominence of deep learning in today’s era, several novel neural-network based approaches have also emerged for the analysis of rs-fMRI data. A majority of these approaches target connectomic feature extraction for single-subject level predictions.

In order to organize the work in this rapidly growing field, we sub-divide the machine learning approaches into different classes by methods and application focus. We first differentiate among unsupervised learning approaches based on whether their main focus is to discover (a) the underlying spatial organization that is reflected in coherent fluctuations, (b) the structure in temporal dynamics of resting-state connectivity, or (c) population-level structure for inter-subject comparisons. Next, we move on to discuss supervised learning. We organize this section by discussing the relevant rs-fMRI features employed in these models, followed by commonly used training algorithms, and finally the various application areas where rs-fMRI has shown promise in performing predictions.

2. Unsupervised learning methods

The primary objective of unsupervised learning is to discover latent representations and disentangle the explanatory factors for variation in rich, unlabelled data. These learning methods do not receive any kind of supervision in the form of target outputs (or labels) to guide the learning process. Instead, they focus on learning structure in the data in order to extract relevant signal from noise. Below, we review some important unsupervised learning methods that have advanced rs-fMRI analysis.

2.1. Clustering

Given data points $\{X_1, \dots, X_n\}$, the goal of clustering is to partition the data into K disjoint groups $\{C_1, \dots, C_K\}$. Different clustering algorithms differ in terms of their clustering objective, which is to maximize some notion of within-cluster similarity and/or between-cluster dissimilarity.

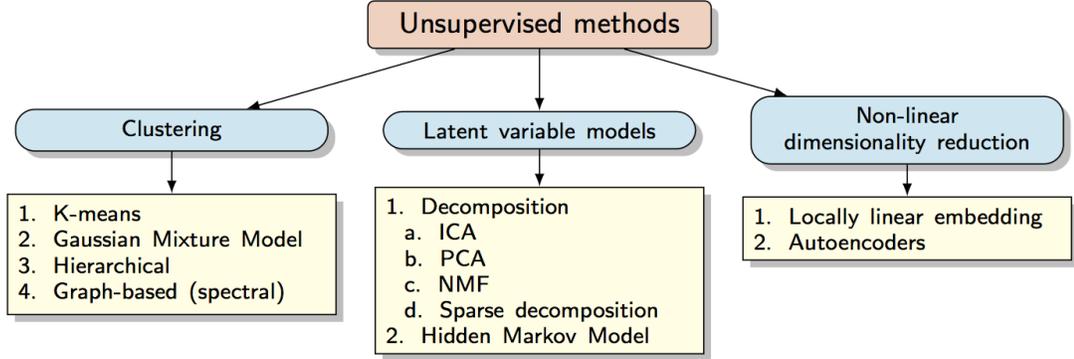


Figure 3: A taxonomy of unsupervised learning methods used for rs-fMRI analysis

K-means. K-means clustering is thus far the most popular learning algorithm for partitioning data. The algorithm aims at minimizing the within-cluster variance. Formally, this corresponds to the following clustering objective,

$$\min \sum_{j=1}^K \sum_{i \in C_j} \left\| X_i - \frac{1}{n_j} \sum_{t \in C_j} X_t \right\|^2$$

where n_j denotes the cardinality of set C_j . This optimization problem is solved using an iterative algorithm, known as the Lloyd’s algorithm. The algorithm begins with initial estimates of cluster centroids and iteratively refines them by (a) assigning each datum to its closest cluster, and (b) updating cluster centroids based on these new assignments.

Gaussian mixture models. Mixture models are often used to represent probability densities of complex multimodal data with hidden components. These models are constructed as mixtures of arbitrary unimodal distributions, each representing a distinct cluster. In the case of Gaussian mixture models, each X_i is assumed to be generated by a two-step process: (a) First, a latent component $z_i \in \{1, \dots, K\}$ is sampled, $z_i \sim \text{Multinomial}(\phi)$ where $\phi_k = P(z_i = k)$; then (b) a random sample is drawn from one of K multivariate gaussians conditional on z_i , i.e. $X_i | z_i = k \sim \mathcal{N}(\mu_k, \Sigma_k)$ where μ_k and Σ_k denote the mean and covariance of the k-th gaussian respectively. Each gaussian distribution thus denotes a unique cluster. The model parameters

$\{\phi, \mu, \Sigma\}$ are obtained by maximizing the complete data likelihood,

$$\begin{aligned} \{\phi_{opt}, \mu_{opt}, \Sigma_{opt}\} &= \arg \max_{\phi, \mu, \Sigma} \sum_{i=1}^n \log P(X_i | \phi, \mu, \Sigma) \\ &= \arg \max_{\phi, \mu, \Sigma} \sum_{i=1}^n \log P(X_i | z_i, \mu, \Sigma) P(z_i | \phi). \end{aligned}$$

Maximum likelihood estimates of GMMs are usually obtained using the Expectation-Maximization (EM) algorithm.

Hierarchical clustering. Hierarchical clustering methods group the data into a set of nested partitions. This multi-resolution structure is often represented with a cluster tree, or dendrogram. Hierarchical clustering is divided into agglomerative or divisive methods, based on whether the clusters are identified in a bottom-up or top-down fashion respectively. Hierarchical agglomerative clustering (HAC), the more dominant approach, initially treats each data point as a singleton cluster and then successively merges them according to pre-specific distance metric until a single cluster containing all observations is formed. Many distance metrics, referred to as linkage criterion, have been proposed in literature that optimize different goals of hierarchical clustering. These include: (a) single-link, where distance between clusters C_1 and C_2 is defined as the distance between their closest points, i.e., $d(C_1, C_2) = \min_{x_i \in C_1, x_j \in C_2} d(x_i, x_j)$, (b) Complete linkage, where this distance is measured between the farthest points, $C(d_1, d_2) = \max_{x_i \in C_1, x_j \in C_2} d(x_i, x_j)$, (c) Average linkage which measures the average distance between members $d(C_1, C_2) = \frac{1}{|C_1||C_2|} \sum_{x_i \in C_1} \sum_{x_j \in C_2} d(x_i, x_j)$ etc. Here, d represents dissimilarity

between observations. Alternate methods for merging have also been proposed, the most popular being Ward's criterion. Ward's method measures how much the within-cluster variance will increase when merging two partitions and minimizes this merging cost. A major drawback is computational complexity, which render HAC methods impractical in applications with large observational data.

Graph-based clustering. Graph based clustering forms another class of similarity-based partitioning methods for data that can be represented using a graph.

Given a weighted undirected graph $G = \{V, E\}$ with vertex set V and edge set E , most graph-partitioning methods optimize a dissociation measure, such as the normalized cut (Ncut). The edge weights $w(i, j)$ represent a function of similarity between vertices i and j . Ncut computes the total edge weights connecting two partitions and normalizes this by their weighted connections to all nodes within the graph. A two-way normalized-cut criteria divides G into disjoint partitions A and B ($A \cup B = V, A \cap B = \phi$) by simultaneously minimizing between-cluster similarity while maximizing within-cluster similarity. This objective criterion is expressed as,

$$\text{Ncut}(A, B) = \frac{\sum_{i \in A, j \in B} w(i, j)}{\sum_{i \in A, j \in V} w(i, j)} + \frac{\sum_{i \in A, j \in B} w(i, j)}{\sum_{i \in V, j \in B} w(i, j)}$$

However, minimizing this objective directly is an NP-hard problem. Spectral clustering algorithms typically solve a relaxation of this problem. This approach can be further extended to obtain a K -way partitioning of the graph. Graph-based clustering approach is often more resilient to outliers, compared to k -means or hierarchical clustering.

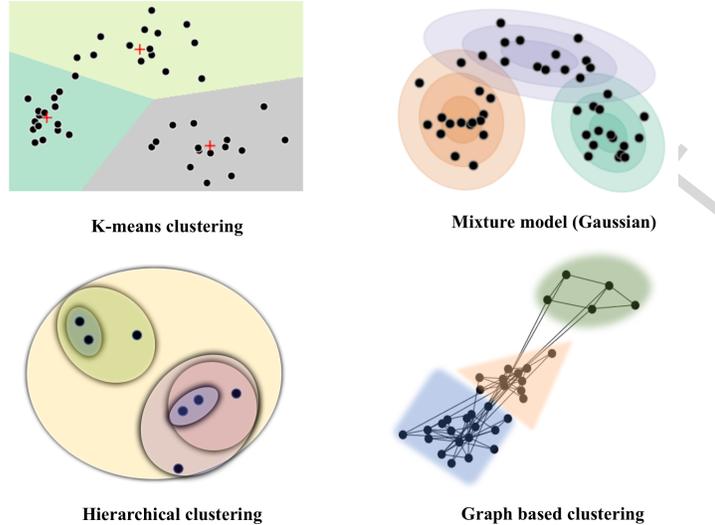


Figure 4: Illustrations of popular clustering algorithms: K-means clustering partitions the data space into Voronoi cells, where each observation is assigned to the cluster with the nearest centroid (marked red in the figure). GMMs assume that each cluster is sampled from a multivariate Gaussian distribution and estimates these probability densities to generate probabilistic assignment of observations to different clusters. Hierarchical (agglomerative) clustering generates nested partitions, where partitions are merged iteratively based on a linkage criteria. Graph-based clustering partitions the graph representation of data so that, for example, number of edges connecting distinct clusters are minimal.

2.2. Latent variable models

2.2.1. Decomposition

Decomposition or factorization based approaches assume that the observed data can be decomposed as a product of simpler matrices, often imposing a specific structure and/or sparsity on these individual matrices. Formally, given data points $X = [x_1, \dots, x_n]$ with $x_i \in R^D$, linear decomposition techniques seek a basis set $W = [w_1, \dots, w_K]$ such that the linear space spanned by W closely reconstructs X .

$$x_i = \sum_{k=1}^K w_k z_i(k)$$

Here, each data point x_i is characterized by unique coefficients $z_i \in R^K$ for the basis set W . Typically, $K < D$ so that decomposition amounts to a dimensionality reduction. In matrix notation, the goal is to find W and Z

such that $X \approx WZ$, where $Z = [z_1, \dots, z_n]$. This ill-posed problem is generally solved by constraining the structure of W and/or Z .

Principal component analysis (PCA). PCA is a linear projection based technique widely used for dimensionality reduction. The goal of PCA is to find an orthonormal basis W that maximizes the variance captured by projected data $Z = W^T X$. This is equivalent to minimizing the reconstruction error of the data points based on the low-dimensional representation Z . Mathematically, this amounts to solving the following optimization problem,

$$W_{opt} = \arg \min_W \|X - WW^T X\|_F^2 \quad \text{subject to } W \in \mathcal{O}^{D \times K}$$

where F denotes the Frobenius norm and $\mathcal{O}^{D \times K}$ denotes the set of $D \times K$ dimensional orthonormal matrices.

Independent component analysis (ICA). Independent Component Analysis (ICA) is a popular method for decomposing data as a linear combination of statistically independent components. In the ICA terminology, W is often known as the mixing matrix whereas Z comprises the source signals. In the above formalism, ICA assumes that the sources, i.e., the rows of Z , are statistically independent. The source signals are recovered using a "whitening" or "unmixing" matrix U , where $U = W^{-1}$. Since $X = WZ$, we obtain $Z = UX$. Popular algorithms thus recover the sources by estimating U such that the components of UX are statistically independent. Common ICA algorithms emulate independence by either minimizing the mutual information between sources (InfoMax) or by maximizing their non-gaussianity (FastICA). ICA usually employs a full-rank matrix factorization and is often preceded with PCA for dimensionality reduction.

Sparse dictionary learning. Sparse dictionary learning is formulated as a linear decomposition problem, similar to ICA/PCA, but with sparsity constraints on the components Z . This results in a non-convex optimization problem of the following form:

$$\{W_{opt}, Z_{opt}\} = \arg \min_{W, Z} \|X - WZ\|_F^2 + C \|Z\|_0$$

In most practical applications, this optimization problem is relaxed by replacing the L0-norm with L1-norm.

Non-negative matrix factorization (NMF). NMF is another dimensionality reduction technique that seeks a low-rank decomposition of the data matrix X with non-negativity constraints on the components W and Z . Typically, this corresponds to solving the following optimization,

$$\{W_{opt}, Z_{opt}\} = \arg \min_{W, Z} \|X - WZ\|_F^2 \quad \text{subject to } W \geq 0, Z \geq 0$$

2.2.2. Hidden Markov Models

Hidden Markov Models (HMMs) are a class of unsupervised learning methods for sequential data. They are used to model a Markov process where the sequence of observations $\{x_1, \dots, x_T\}$ are assumed to be generated from a sequence of underlying hidden states $\{s_1, \dots, s_T\}$, which can be discrete. In a HMM with K states, it is assumed that s_i can take discrete values in $\{1, \dots, K\}$. The parameters of the HMM are learned by maximizing the complete data likelihood,

$$\begin{aligned} \theta_{opt} &= \arg \max_{\theta} P(x_1, \dots, x_T, s_1, \dots, s_T | \theta) \\ &= \arg \max_{\theta} \prod_{t=1}^T P(s_t | s_{t-1}, \theta) P(x_t | s_t, \theta) \end{aligned}$$

Here, $P(s_1 | s_0)$ denotes the initial state distribution π . The state transition probabilities are defined by a transition matrix T with elements $T_{i,j} = P(s_t = j | s_{t-1} = i)$. The conditionals $P(x_t | s_t = k, \theta)$ are captured by an emission probability table $E[k, x_t]$. The parameters θ of this probabilistic model are thus $\{\pi, T, E\}$. This maximum likelihood estimation problem is efficiently solved using a special case of the Expectation-Maximization algorithm, known as the Baum-Welch algorithm.

2.3. Non-linear embeddings

Locally linear embeddings. LLE projects data to a reduced dimensional space while preserving local distances between data points and their neighborhood. LLE algorithm proceeds in two steps. First, each input $X_i, i \in \{1, \dots, n\}$ is approximated as a linear combination of its K closest neighbors. The linear subspace W is obtained by minimizing the reconstruction error, i.e.,

$$\begin{aligned} W_{opt} &= \arg \min_W \sum_i |X_i - \sum_j W_{ij} X_j|^2 \\ &\text{subject to } \sum_j W_{ij} = 1 \end{aligned}$$

Here, $W_{ij} = 0$ if X_j is not one of the K -nearest neighbors of X_i . In the second step, the low-dimensional embeddings Y_i are obtained by minimizing the embedding cost function,

$$Y_{opt} = \arg \min_Y |Y_i - \sum_j W_{ij} Y_j|^2$$

In the latter optimization, W is kept fixed at W_{opt} , while Y_i 's are optimized.

Autoencoders. The autoencoder is an unsupervised neural-network based approach for learning latent representations of high-dimensional data. It encodes the input X into a lower dimensional representation $Z = f_\theta(X)$, known as the bottleneck, which is then decoded to reconstruct the input $\hat{X} = g_\phi(Z)$. Both the encoder f_θ and decoder g_ϕ are neural networks. The autoencoder is trained to minimize the reconstruction error on a set of examples, often measured with an L2 loss, i.e., $\|X - \hat{X}\|^2$. The autoencoder can thus be seen as a non-linear extension of PCA since f_θ and g_ϕ are in general non-linear functions.

3. Applications of unsupervised learning in rs-fMRI

Unsupervised machine learning methods have proven promising for the analysis of high-dimensional data with complex structures, making it ever more relevant to rs-fMRI. Many unsupervised learning approaches in rs-fMRI aim to parcellate the brain into discrete functional sub-units, akin to atlases. These segmentations are driven by functional data, unlike those approaches that use cytoarchitecture as in the Brodmann atlas, or macroscopic anatomical features, as in the Automated Anatomical Labelling (AAL) atlas [47]. A second class of applications delve into the exploration of brain network dynamics. Unsupervised learning has recently been applied to interrogate the dynamic functional connectome with promising results [43, 42, 44, 48, 49]. Finally, the third application of unsupervised learning focuses on learning latent low-dimensional representations of RSFC to conduct analyses across a population of subjects. We discuss the methods under each of these challenging application areas below.

3.1. Discovering spatial patterns with coherent fluctuations

Mapping the boundaries of functionally distinct neuroanatomical structures, or identifying clusters of functionally coupled regions in the brain is a major objective in neuroscience. Rs-fMRI and machine learning methods provide a promising combination with which to achieve this lofty goal.

In the case of rs-fMRI, the typical approach is to decompose the 4D fMRI data into a linear superposition of distinct spatial modes that show coherent temporal dynamics using techniques like ICA. Clustering is an alternative unsupervised learning approach for analysis of rs-fMRI data. Unlike ICA or dictionary learning, clustering is used to partition the brain surface (or volume) into *disjoint* functional networks. It is important to draw a distinction at this stage between two slightly different applications of clustering since they sometimes warrant different constraints; one direction is focused on identifying functional networks which are often spatially distributed, whereas the other is used to parcellate brain regions. The latter application aims to construct atlases that reflect local areas that constitute the functional neuroanatomy, much like how standard atlases such as the Automated Anatomical Labelling (AAL) [47] delineate macroscopic anatomical regions. One important design decision in the application of clustering is the distance function used to measure dissimilarity between different voxels (or vertices). In the case of rs-fMRI, this distance function is either computed on raw time-series at voxels or between their connectivity profiles. While these two distances are motivated by the same idea of functional coherence, certain differences have been found in parcellations optimized using either criteria [50].

An important requirement for almost all of these methods is the *a priori* selection of the number of clusters/components. These are often determined through cross-validation or through statistics that reflect the quality, stability or reproducibility of decomposition/partitions at different scales.

3.1.1. ICA

ICA has been one of the earliest and most widely used analytic tools for rs-fMRI, driving several pivotal neuroscientific insights into intrinsic brain networks. When applied to rs-fMRI, brain activity is expressed as a linear superposition of distinct spatial patterns or maps, with each map following its own characteristic time course. These spatial maps can reflect a coherent functional system or noise, and several criteria can be used to automatically

I. Discovering spatial maps with coherent temporal dynamics

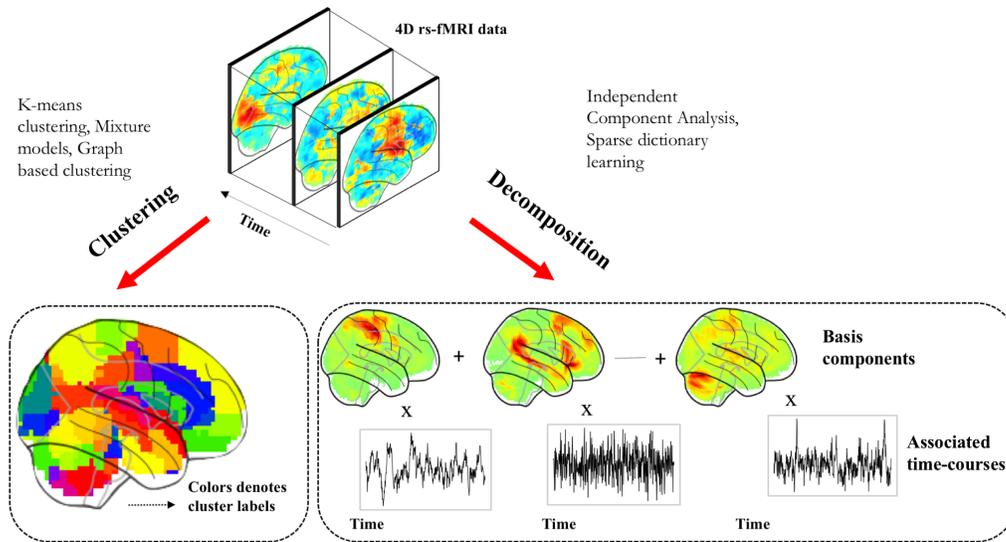


Figure 5: Schematic of application 3.1: In decomposition, the original fMRI data is expressed as a linear combination of spatial patterns and their associated time series - in ICA, the independence of spatial maps is optimized whereas in sparse dictionary learning, the sparsity of maps is encouraged. In clustering, time series or connectivity fingerprints of voxels are clustered to assign voxels to distinct functional networks.

differentiate them. This capability to isolate noise sources makes ICA particularly attractive. In the early days of rs-fMRI, several studies demonstrated marked resemblance between the ICA spatial maps and cortical functional networks known from task-activation studies [2, 4]. While typical ICA models are noise-free and assume that the only stochasticity is in the sources themselves, several variants of ICA have been proposed to model additive noise in the observed signals. Beckmann et al. [2] introduced a probabilistic ICA (PICA) model to extract the connectivity structure of rs-fMRI data. PICA models a linear instantaneous mixing process under additive noise corruption and statistical independence between sources. De Luca et al. [5] showed that PICA can reliably distinguish RSNs from artifactual patterns. Both these works showed high consistency in resting-state patterns across multiple subjects. While there is no standard criteria for validating the ICA patterns, or any clustering algorithm for that matter, reproducibility or reliability is often used for quantitative assessment. More recently, Khorshidi et al. proposed an automated denoising strategy for fMRI based on ICA, known as FIX "FMRIB's ICA-based-X-noiseifier". The authors trained a classifier using manual annotations to label artifactual components based on distinct spatial/ temporal features. These components could represent a variety of structured noise sources and once identified, they can be either subtracted or regressed out of the data to yield clean signals.

ICA can also be extended to make group inferences in population studies. Group ICA is thus far the most widely used strategy, where multi-subject fMRI data are concatenated along the temporal dimension before implementing ICA [51]. Individual-level ICA maps can then be obtained from this group decomposition by back-projecting the group mixing matrix [51], or using a dual regression approach [52]. More recently, Du et al. [53] introduced a group information guided ICA to preserve statistical independence of individual ICs, where group ICs are used to constrain the corresponding subject-level ICs. Varoquaux et al. [54] proposed a robust group-level ICA model to facilitate between-group comparisons of ICs. They introduce a generative framework to model two levels of variance in the ICA patterns, at the group level and at a subject-level, akin to a multivariate version of mixed-effect models. The IC estimation procedure, termed Canonical ICA, employs Canonical Correlation Analysis to identify a joint subspace of common IC patterns across subjects and yields ICs that are well representative of the group.

Alternatively, it is also possible to compute individual-specific ICA maps

and then establish correspondences across them [53] for generating group inferences; however, this approach has been limited because source separations can be very different across subjects, for example, due to fragmentation.

While ICA and its extensions have been used broadly by the rs-fMRI community, it is important to acknowledge its limitations. ICA models linear representations of non-Gaussian data. Whether a linear transformation can adequately capture the relationship between independent latent sources and the observed high-dimensional fMRI data is uncertain and likely unrealistic. Unlike the popular Principal Component Analysis (PCA), ICA does not provide the ordering or the energies of its components, which makes it impossible to distinguish strong and weak sources. This also complicates replicability analysis since known sources i.e., spatial maps could be expressed in any arbitrary order. Extracting meaningful ICs also sometimes necessitates manual selection procedures, which can be inefficient or subjective. In the ideal scenario, each individual component represents either a physiologically meaningful activation pattern or noise. However, this might be an unrealistic assumption for rs-fMRI. Additionally, since ICA assumes non-Gaussianity of sources, Gaussian physiological noise can contaminate the extracted components. Further, due to the high-dimensionality of fMRI, analysis often proceeds with PCA based dimensionality reduction before application of ICA. PCA computes uncorrelated linear transformations of highest variance (thus explaining greatest variability within the data) from the top eigenvectors of the data covariance matrix. While this step is useful to remove observation noise, it could also result in loss of signal information that might be crucial for subsequent analysis. Although ICA optimizes for independence, it does not guarantee independence. Based on studies of functional integration within the brain, assumptions of independence between functional units could themselves be questioned from a neuroscientific point of view. Several papers have suggested that ICA is especially effective when spatial patterns are sparse, with negligible or little overlap. This hints to the possibility that success of ICA is driven by sparsity of the components rather than their independence. Along these lines, Daubechies and colleagues claim that fMRI representations that optimize for sparsity in spatial patterns are more effective than fMRI representations that optimize independence [55].

3.1.2. Learning sparse spatial maps

Sparse dictionary learning is another popular framework for constructing succinct representations of observed data. Varoquaux et al. [56] adopt a

dictionary learning framework for segmenting functional regions from resting-state fMRI time series. Their approach accounts for inter-subject variability in functional boundaries by allowing the subject-specific spatial maps to differ from the population-level atlas. Concretely, they optimize a loss function comprising a residual term that measures the approximation error between data and its factorization, a cost term penalizing large deviations of individual subject spatial maps from group level latent maps, and a regularization term promoting sparsity. In addition to sparsity, they also impose a smoothness constraint so that the dominant patterns in each dictionary are spatially contiguous to construct a well-defined parcellation. In order to prevent blurred edges caused due to the smoothness constraint, Abraham et al. [57] propose a total variation regularization within this multi-subject dictionary learning framework. This approach is shown to yield more structured parcellations that outperform competing methods like ICA and clustering in explaining test data. Similarly, Lv et al. [58] propose a strategy to learn sparse representations of whole-brain fMRI signals in individual subjects by factorizing the time-series into a basis dictionary and its corresponding sparse coefficients. Here, dictionaries represent the co-activation patterns of functional networks and coefficients represent the associated spatial maps. Experiments revealed a high degree of spatial overlap in the extracted functional networks in contrast to ICA that is known to yield spatially non-overlapping components in practice.

3.1.3. *K-means clustering and mixture models*

K-means clustering or mixture models are frequently used for spatial segmentation of fMRI data [59, 37, 60, 61]. Similarity between voxels can be defined by correlating their raw time-series [59] or connectivity profiles [61]. Euclidean distance metrics have also been used on spectral features of time series [37].

K-means clustering has provided several novel insights into functional organization of the human brain. It has revealed the natural division of cortex into two complementary systems, the internally-driven "intrinsic" system and the stimuli-driven "extrinsic" system [59, 60]; provided evidence for a hierarchical organization of RSNs [60]; and exposed the anatomical contributions to co-varying resting-state fluctuations [37].

Golland et al. [62] proposed a Gaussian mixture model for clustering fMRI signals. Here, the signal at each voxel is modelled as a weighted sum of N Gaussian densities, with N determining the number of hypothesized

functional networks and weights reflecting the probability of assignment to different networks. Large-scale systems were explored at several resolutions, revealing an intrinsic hierarchy in functional organization. Yeo et al. [63] used rs-fMRI measurements on 1000 subjects to estimate the organization of large-scale distributed cortical networks. They employed a mixture model to identify clusters of voxels with similar corticocortical connectivity profiles. Number of clusters were chosen from stability analysis and parcellations at both a coarse resolution of 7 networks and a finer scale of 17 networks were identified. A high degree of replicability was attained across data samples, suggesting that these networks can serve as reliable reference maps for functional characterization.

3.1.4. Identifying hierarchical spatial organization

Several studies have provided evidence for a hierarchical organization of functional networks in the brain [62, 60]. Hierarchical agglomerative clustering (HAC) thus provides a natural tool to partition rs-fMRI data and explore this latent hierarchical structure. Earliest applications of clustering to rs-fMRI were based on HAC [64, 36]. This technique thus largely demonstrated the feasibility of clustering for extracting RSNs from rs-fMRI data. Recent applications of HAC have focused on defining whole-brain parcellations for downstream analysis [65, 66, 67]. Spatial continuity can be enforced in parcels, for example, by considering only local neighborhoods as potential candidates for merging [65].

An advantage of hierarchical clustering is that unlike k-means clustering, it does not require knowledge of the number of clusters and is completely deterministic. However, once the cluster tree is formed, the dendrogram must be split at a level that best characterizes the "natural" clusters. This can be determined based on a linkage inconsistency criterion [64], consistency across subjects [36], or advance empirical knowledge [68].

While a promising approach for rs-fMRI analysis, hierarchical clustering has some inherent limitations. It often relies on prior dimensionality reduction, for example by using an anatomical template [36], which can bias the resulting parcellation. It is a greedy strategy and erroneous partitions at an early stage cannot be rectified in subsequent iterations. Single-linkage criterion may not work well in practice since it merges partitions based on the nearest neighbor distance, and hence is not inherently robust to noisy resting-state signals. Further, different metrics usually optimize divergent attributes of clusters. For example, single-link clustering encourages extended

clusters whereas complete-link clustering promotes compactness. This makes the *a priori* choice of distance metric somewhat arbitrary.

3.1.5. Graph based clustering

Functional MRI data can be naturally represented in the form of graphs. Here, nodes represent voxels and edges represent connection strength, typically measured by a correlation coefficient between voxel time series or between connectivity maps [50, 69]. Often, thresholding is applied on edges to limit graph complexity. Graph segmentation approaches, such as those based on Ncut criteria, have been widely used to derive whole-brain parcellations [50, 70, 71]. Population-level parcellations are usually derived in a two stage procedure: First, individual graphs are clustered to extract functionally-linked regions, followed by a second stage where a group-level graph characterizing the consistency of individual cluster maps is clustered [69, 50]. Spatial contiguity can be easily enforced by constraining the connectivity graph to local neighborhoods [50], or through the use of shape priors [71]. Departing from this protocol, Shen et al. [70] propose a groupwise clustering approach that jointly optimizes individual and group parcellations in a single stage and yields spatially smooth group parcellations in the absence of any explicit constraints.

A disadvantage of the Ncut criteria for fMRI is its bias towards creating uniformly sized clusters, whereas in reality functional regions show large size variations. Graph construction itself involves arbitrary decisions which can affect clustering performance [72] e.g., selecting a threshold to limit graph edges, or choosing the neighborhood to enforce spatial connectedness.

**Table 1: Key papers for application 3.1
Discovering spatial patterns with coherent resting-state fluctuations (RSNs)**

Approach a: Decomposition
<p><i>Investigations into resting-state connectivity using independent component analysis (Beckmann et al., 2005)[2]</i> <i>Consistent resting-state networks across healthy subjects (Damoiseaux et al., 2006)[4]</i> <i>Method: ICA, Contribution: Early works demonstrating the striking similarity between ICA spatial maps and cortical functional networks</i></p> <p><i>Group comparison of resting-state fMRI using multi-subject ICA and dual regression (Beckmann et al., 2009) [52]</i> <i>A group model for stable multi-subject ICA on fMRI datasets (Varoquaux et al., 2010)[54]</i> <i>Group information guided ICA for fMRI data analysis (Du et al., 2013) [53]</i> <i>Method: ICA (group-level), Contribution: Influential works discussing analytical approaches for multi-subject ICA in resting-state</i></p> <p><i>Multi-subject dictionary learning to segment an atlas of brain spontaneous activity (Varoquaux et al., 2011) [56]</i> <i>Method: Sparse dictionary learning, Contribution: A multi-subject dictionary learning framework for learning sparse spatial maps</i></p>
Approach b: Clustering
<p><i>Hierarchical clustering to measure connectivity in fMRI resting-state data, (Cordes et al., 2002)[64]</i> <i>Neurophysiological Architecture of Functional Magnetic Resonance Images of Human Brain (Salvador et al., 2005)[36]</i> <i>Method: Hierarchical clustering, Contribution: Earliest applications of clustering to rs-fMRI; highlighted hierarchical organization of functional networks</i></p> <p><i>The organization of the human cerebral cortex estimated by intrinsic functional connectivity, (Yeo et al., 2011)[63]</i> <i>Method: Mixture models, Contribution: Influential large-scale study investigating brain’s functional organization</i></p> <p><i>A whole brain fMRI atlas generated via spatially constrained spectral clustering, (Craddock et al., 2012)[50]</i> <i>Groupwise whole-brain parcellation from resting-state fMRI data for network node identification, (Shen et al., 2013)[70]</i> <i>Method: Graph based clustering, Contribution: Released consistent whole-brain functional atlas for fMRI at varying spatial resolutions based on rs-fMRI data</i></p>

3.1.6. Comments

I. A comment on alternate connectivity-based parcellations. Several papers make a distinction between clustering / decomposition and boundary detection based approaches for network segmentation. In the rs-fMRI literature, several non-learning based parcellations have been proposed, that exploit traditional image segmentation algorithms to identify functional areas based on abrupt RSFC transitions [73, 74]. Clustering algorithms do not mandate spatial contiguity, whereas boundary based methods implicitly do. In contrast, boundary based approaches fail to represent long-range functional associa-

tions, and may not yield parcels that are as connectionally homogeneous as unsupervised learning approaches. A hybrid of these approaches can yield better models of brain network organization. This direction was recently explored by Schaefer et al. [75] with a Markov Random Field model. The resulting parcels showed superior homogeneity compared with several alternate gradient and learning-based schemes. Further, complementing RSFC with other modalities can yield corroborative and perhaps complementary information for delineating areal boundaries. Recently, Glasser et al. approached this problem by developing a multi-modal approach for generating brain parcellations[74]. The authors propose a semi-automated approach that combines supervised machine learning with manual annotations for parcellating regions based on their multi-modal fingerprints (architecture, function, connectivity and topography). Such an approach can be instrumental towards the goal of precise human brain functional mapping.

II. Subject versus population level parcellations. Significant effort in rs-fMRI literature is dedicated to identifying population-average parcellations. The underlying assumption is that functional connectivity graphs exhibit similar patterns across subjects, and these global parcellations reflect common organizational principles. Yet, individual-level parcellations can potentially yield more sensitive connectivity features for investigating networks in health and disease. A central challenge in this effort is to match the individual-level spatial maps to a population template in order to establish correspondences across subjects. Common approaches to obtain subject-specific networks with group correspondence often incorporate back-projection and dual regression [52, 51], or hierarchical priors within unsupervised learning [56, 76]. While a number of studies have developed subject-specific parcellations, the significance of this inter-subject variability for network analysis has only recently been discussed. Kong et al. [76] developed high quality subject-specific parcellations using a multi-session hierarchical Bayesian model, and showed that subject-specific variability in functional topography can predict behavioral measures. Recently, using a novel parcellation scheme based on K-medoids clustering, Salehi et al. [77] showed that individual-level parcellation alone can predict the sex of the individual. These studies suggest the intriguing idea that subject-level network organization, i.e. voxel-to-network assignments, can capture concepts intrinsic to individuals, just like connectivity strength.

III. Is there a universal 'gold-standard' atlas? . When considering the family of different methods, algorithms or modalities , there exist a plethora of diverse brain parcellations at varying levels of granularity. Thus far, there is no unified framework for reasoning about these brain parcellations. Several taxonomic classifications can be used to describe the generation of these parcellations, such as machine learning or boundary detection, decomposition or clustering, multi-modal or unimodal. Even within the large class of clustering approaches, it is impossible to find a single algorithm that is consistently superior for a collection of simple, desired properties of partitioning [78]. Several evaluation criteria have emerged for comparing different parcellations, exposing the inherent trade-offs at work. Arslan et al. [79] performed an extensive comparison of several parcellations across diverse methods on resting-data from the Human Connectome Project (HCP). Through independent evaluations, they concluded that no single parcellation is consistently superior across all evaluation metrics. Recently, Salehi et al. [80] showed that different functional conditions, such as task or rest, generate reproducibly distinct parcellations thus questioning the very existence of an optimal parcellation, even at an individual-level. These novel studies necessitate rethinking about the final goals of brain mapping. Several studies have reflected the view that there is no optimal functional division of the brain, rather just an array of meaningful brain parcellations [65]. Perhaps, brain mapping should not aim to identify functional sub-units in a universal sense, like Brodmann areas. Rather, the goal of human brain mapping should be reformulated as revealing consistent functional delineations that enable reliable and meaningful investigations into brain networks.

IV. A comparison between decomposition and clustering. A high degree of convergence has been observed in the functionally coherent patterns extracted using decomposition and clustering. Decomposition techniques allow soft partitioning of the data, and can thus yield spatially overlapping networks. These models may be more natural representations of brain networks where, for example, highly integrated regions such as network 'hubs' can simultaneously subservise multiple functional systems. Although it is possible to threshold and relabel the generated maps to produce spatially contiguous brain parcellations, these techniques are not naturally designed to generate disjoint partitions. In contrast, clustering techniques automatically yield hard assignments of voxels to different brain networks. Spatial constraints can be easily incorporated within different clustering algorithms to

yield contiguous parcels. Decomposition models can adapt to varying data distributions, whereas clustering solutions allow much less flexibility owing to rigid clustering objectives. For example, k-means clustering function looks to capture spherical clusters. While a thorough comparison between these approaches is still lacking, some studies have identified the trade-offs between choosing either technique for parcellation. Abraham et al. [57] compared clustering approaches with group-ICA and dictionary learning on two evaluation metrics: stability as reflected by reproducibility in voxel assignments on independent data, and data fidelity captured by the explained variance on independent data. They observed a stability-fidelity trade-off: while clustering models yield stable regions but do not explain test data as well, linear decomposition models explain the test data reasonably well but at the expense of reduced stability.

3.2. Discovering patterns of dynamic functional connectivity

Unsupervised learning has also been applied to study patterns of temporal organization or dynamic reconfigurations in resting-state networks. These studies are often based on two alternate hypothesis that (a) dynamic (windowed) functional connectivity cycles between discrete "connectivity states", or (b) functional connectivity at any time can be expressed as a combination of latent "connectivity states". The first hypothesis is examined using clustering-based approaches or generative models like HMMs, while the second is modelled using decomposition techniques. Once stable states are determined across population, the former approach allows us to estimate the fraction of time spent in each state by all subjects. This quantity, known as dwell time or occupancy of the state, shows meaningful variation across individuals [43, 42, 81, 82]. It is important to note than in all these approaches, the RSNs or the spatial patterns are assumed to be stationary over time and it is the temporal coherence that changes with time.

3.2.1. Clustering

Several studies have discovered recurring dynamic functional connectivity patterns, known as "states", through k-means clustering of windowed correlation matrices [42, 81, 82, 83, 84]. FC associated with these repeating states shows marked departure from static FC, suggesting that network dynamics provide novel signatures of the resting brain [42]. Notable differences have been observed in the dwell times of multiple states between healthy

II. Discovering reproducible patterns of dynamic functional connectivity

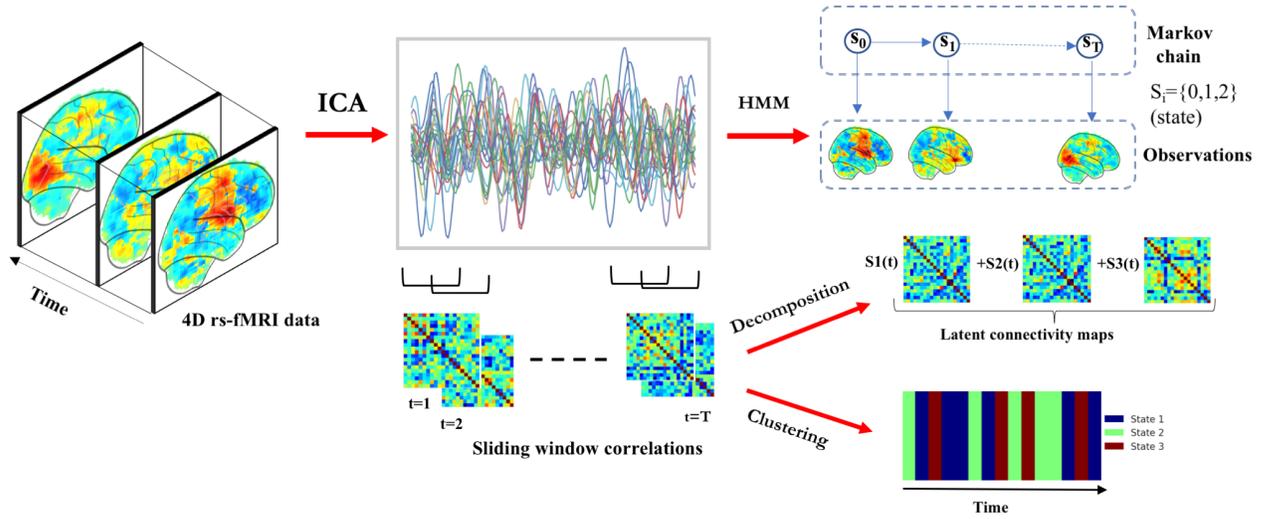


Figure 6: Schematic of application 3.2. Three connectivity states are assumed in the data for illustration purposes

controls and patient populations across schizophrenia, bipolar disorder and psychotic-like experience domains [81, 82, 83].

Abrol et al. [84] performed a large-scale study to characterize the replicability of brain states using standard k-means as well as a more flexible, soft k-means algorithm for state estimation. Experiments indicated reproducibility of most states, as well as their summary measures, such as mean dwell times and transition probabilities etc. across independent population samples. While these studies establish the existence of recurring FC states, behavioral associations of these states is still unknown. In an interesting piece of work, Wang et al. [85] identified two stable dynamic FC states using k-means clustering that showed correspondence with internal states of high- and low-arousal respectively. This suggests that RSFC fluctuations are behavioral state-dependent, and presents one explanation to account for the heterogeneity and dynamic nature of RSFC.

3.2.2. Markov modelling of state transition dynamics

HMMs are another valuable tool to interrogate recurring functional connectivity patterns [44, 43, 86]. The notion of states remains similar to the

”FC states” described above for clustering; however, the characterization and estimation is drastically different. Unlike clustering where sliding windows are used to compute dynamic FC patterns, HMMs model the rs-fMRI time-series directly. Hence, they offer a promising alternative to overcome statistical limitations of sliding-windows in characterizing FC changes.

Several interesting results have emerged through the adoption of HMMs. Vidaurre et al. [43] find that relative occupancy of different states is a subject-specific measure linked with behavioral traits and heredity. Through Markov modelling, transitions between states have been revealed to occur as a non-random sequence [42, 43], that is itself hierarchically organized [43]. Recently, network dynamics modelled using HMMs were shown to distinguish MCI patients from controls [86], thereby indicating their utility in clinical domains.

3.2.3. Finding latent connectivity patterns across time-points

Decomposition techniques for understanding RSFC dynamics have the same flavor as the ones described in section 2.2.1: of explaining data through latent factors; however, the variation of interest is across time in this case. Adoption of matrix decomposition techniques exposes a basis set of FC patterns from windowed correlation matrices. Dynamic FC has been characterized using varied decomposition approaches, including PCA[48], Singular Value Decomposition (SVD)[49], non-negative matrix factorization[87] and sparse dictionary learning[88].

Decomposition approaches, here, diverge from clustering or HMMs as they associate each dFC matrix with multiple latent factors instead of a single component. To compare these alternate approaches, Leonardi et al. [49] implemented a generalized matrix decomposition, termed k-SVD. This factorization generalizes both k-means clustering and PCA subject to variable constraints. Reproducibility analysis in this study indicated that dFC is better characterized by multiple overlapping FC patterns.

Decomposition of dFC has revealed novel alterations in network dynamics between healthy controls and patients suffering from PTSD [88] or multiple sclerosis [48], as well as between childhood and young adulthood [87].

Table 2: Key papers for application 3.2
Discovering reproducible patterns of dynamic functional connectivity

Approach a: Decomposition
<p><i>Principal components of functional connectivity: a new approach to study dynamic brain connectivity during rest (Leonardi et al.,2013)[48]</i> <i>Method:PCA , Contribution: Early work characterizing dFC using latent connectivity patterns and suggesting altered connectivity dynamics in disease</i></p>
Approach b: Clustering
<p><i>Tracking whole-brain connectivity dynamics in the resting state, (Allen et al.,2014)[42]</i> <i>Method: K-means, Contribution: Provided evidence for recurring FC states and suggested marked departure of dynamic connectivity patterns from static FC</i></p> <p><i>Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia (Damaraju 2014)[81]</i> <i>Method: K-means, Contribution: Revealed strong statistical differences in dwell times of multiple FC states between controls and a disease group</i></p>
Approach c: Markov models
<p><i>Unsupervised learning of functional network dynamics in resting state fMRI (Eavani 2013)[44]</i> <i>Method: HMM, Contribution: Earliest application of HMMs to study resting-state functional network dynamics</i></p> <p><i>Brain network dynamics are hierarchically organized in time, (Vidaurre et al.,2017)[43]</i> <i>Method: HMM, Contribution: Demonstrated that transitions between FC states occur in a non-random hierarchically organized fashion and revealed that dwell times of FC states are linked with behavioral traits and heredity.</i></p>

3.3. Disentangling latent factors of inter-subject FC variation

Unsupervised learning can also disentangle latent explanatory factors for FC variation across population. We find two applications here: (i) learning low dimensional embeddings of FC matrices for subsequent supervised learning and (ii) learning population groupings to differentiate phenotypes based solely on FC.

3.3.1. Dimensionality reduction

Rs-fMRI analysis is plagued by the curse of dimensionality, i.e., the phenomenon of increasing data sparsity in higher dimensions. Commonly used data features such as FC between pairs of regions, increase as $\mathcal{O}(n^2)$ with the number of parcellated regions. Further, sample size in typical fMRI studies is typically of the order of tens or hundreds, making it harder to learn generalizable patterns from original high dimensional data. To overcome this,

III. Disentangling latent factors of inter-subject FC variation

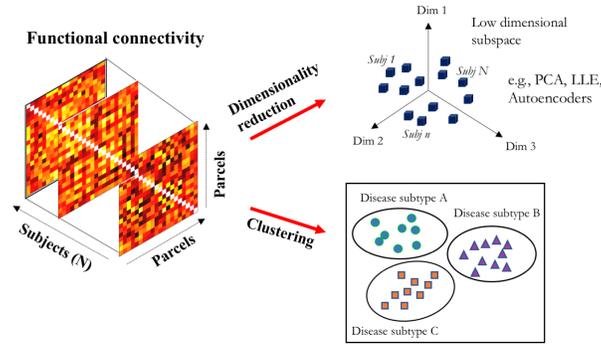


Figure 7: Schematic of application 3.3. Dimensionality reduction of high-dimensional connectomes into 3 latent components is shown for illustration.

linear decomposition methods like PCA or sparse dictionary learning have been widely used for dimensionality reduction of functional connectivity data [89, 90, 91, 92].

Several non-linear embedding methods like Locally linear embedding (LLE) or Autoencoders (AEs) have also garnered attention. LLE embeddings have been employed in rs-fMRI studies, for example, to improve predictions in supervised age regression [93], or for low-dimensional clustering to distinguish Schizophrenia patients from controls [94]. AEs are a neural network based alternative for generating reduced feature sets through nonlinear input transformations. They have been used for feature reduction of RSFC in several studies [86, 95]. AEs can also be used in a pre-training stage for supervised neural network training, in order to direct the learning towards parameter spaces that support generalization [96]. This technique was shown, for example, to improve classification performance of Autism and Schizophrenia using RSFC [97, 98].

3.3.2. Clustering heterogeneous diseases

Clustering can expose sub-groups within a population that show similar FC. Using unsupervised maximum margin clustering [99], Zeng et al. [100] demonstrated that clusters can be associated with disease category (depressed v/s control) to yield high classification accuracy. Recently, Drysdale et al. [101] discovered novel neurophysiological subtypes of depression based on RSFC. Using an agglomerative hierarchical procedure, they identified clus-

tered patterns of dysfunctional connectivity, where clusters showed associations with distinct clinical symptom profiles despite no external supervision. Several psychiatric disorders, like depression, schizophrenia, and autism spectrum disorder, are believed to be highly heterogeneous with widely varying clinical presentations. Instead of labelling them as a unitary syndrome, differential characterization based on disease sub-types can build better diagnostic, prognostic or therapy selection systems. Unsupervised clustering could aid in the identification of these disease subtypes based on their rs-fMRI manifestations.

Table 3: Key papers for application 3.3
Disentangling latent factors of inter-subject RSFC variation

Approach a: Decomposition

Identifying Sparse Connectivity Patterns in the brain using resting-state fMRI (Eavani et al., 2015)[91]
Method: Sparse dictionary learning, *Contribution:* One of the early works explaining inter-subject RSFC variability in terms of sparse connectivity patterns

Approach b: Non-linear embeddings

Discriminative analysis of resting-state functional connectivity patterns of schizophrenia using low dimensional embedding of fMRI (Shen et al., 2010)[94]
Method: LLE, *Contribution:* Proposed an unsupervised learning approach for discriminating Schizophrenia patients from controls with impressive accuracy

Identification of autism spectrum disorder using deep learning and the ABIDE dataset (Heinsfeld et al., 2018)[97]
Deep neural network with weight sparsity control and pre-training extracts hierarchical features and enhances classification performance: Evidence from whole-brain resting-state functional connectivity patterns of schizophrenia (Kim et al., 2016)[98]
Method: Autoencoders, *Contribution:* More recent works demonstrating the advantages of autoencoder based dimensionality reduction/pre-training for downstream classification

Approach c: Clustering

Unsupervised classification of major depression using functional connectivity MRI (Zeng et al., 2014)[100]
Resting-state connectivity biomarkers define neurophysiological subtypes of depression (Drysdale et al., 2017) [101]
Method: Maximum margin clustering/HAC, *Contribution:* Demonstrated the power of clustering approaches for diagnosing depression and identifying its subtypes based on rs-fMRI manifestations

4. Supervised Learning

Supervised learning denotes the class of problems where the learning system is provided input features of the data and corresponding target predictions (or labels). The goal is to learn the mapping between input and label,

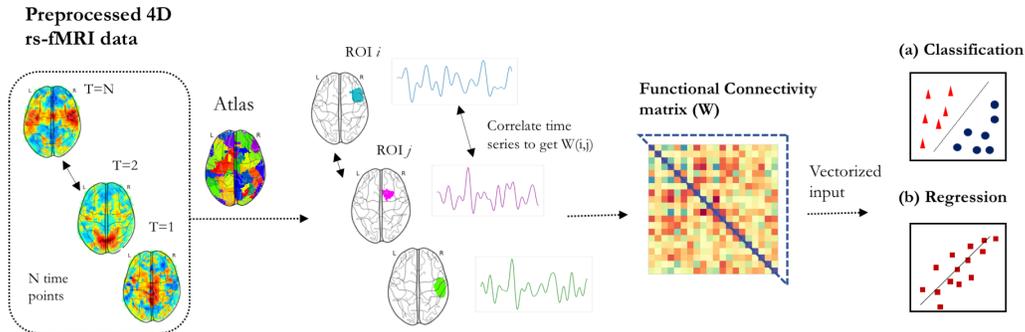


Figure 8: A common classification/regression pipeline for connectomes

so that the system can compute predictions for previously unseen input data points. Prediction of autism from rs-fMRI correlations is an example problem. Since intrinsic FC reflects interactions between cognitively associated functional networks, it is hypothesized that systematic alterations in resting-state patterns can be associated with pathology or cognitive traits. Promising diagnostic accuracy attained by supervised algorithms using rs-fMRI constitute strong evidence for this hypothesis.

In this section, we separate the discussion of rs-fMRI feature extraction from the classification algorithms and application domains.

4.1. Deriving connectomic features

To render supervised learning effective, the most critical factor is feature extraction. Capturing relevant neurophenotypes from rs-fMRI depends on various design choices. Almost all supervised prediction models use brain networks or "connectomes" extracted from rs-fMRI time-series as input features for the learning algorithm. The prototypical prediction pipeline is shown in Figure 8. Here, we discuss critical aspects of common choices for brain network representations in supervised learning.

The first step in the prototypical pipeline is region definition and corresponding time-series extraction. Dense connectomes derived from voxel-level correlations are rarely used in practice for supervised prediction due to their high dimensionality. Both functional and anatomical atlases have been extensively used for this dimensionality reduction. Atlases delineate ROIs within the brain that are often used to study RSFC at a supervoxel scale. Each ROI is represented with a distinct time-course, often computed as the average signal from all voxels within the ROI. Consequently, the data is represented as

an $N \times T$ matrix, where N denotes the number of ROIs and T represents the time-points in the signal. A drawback of using pre-defined atlases is that they may not explain the rs-fMRI dataset very well since they are not optimized for the data at hand. Several studies employ data-driven techniques to define regions within the brain, using unsupervised models such as K-means clustering, Ward clustering, ICA or dictionary learning etc [66, 102]. It is important to note that since we use pairs of ROIs to define whole-brain RSFC, the features grow as $\mathcal{O}(N^2)$ with the number of ROIs. Therefore, in most studies, the network granularity is often limited to the range of 10-400 ROIs.

The second step in this pipeline involves defining connectivity strength for extracting the connectome matrix. Functional connectivity between pairs of ROIs is the most common feature representation of rs-fMRI in supervised learning. In order to extract connectivity matrix, first the covariance matrix needs to be estimated. Sample covariance matrices are subject to a significant amount of estimation error due to the limited number of time-points. This ill-posed problem can be partially resolved through the use of shrinkage transformations [103]. Connectivity strength can then be estimated from the covariance matrix in multiple ways. Pearson’s correlation coefficient is a commonly used metric for estimating functional connectivity. Partial correlation is another metric that has been shown to yield better estimates of network connections in simulated rs-fMRI data [104]. It measures the normalized correlation between two time-series, after removing the effect of all other time-series in the data. Alternatively, one can use a tangent-based reparametrization of the covariance matrix to obtain functional connectivity matrices that respect the Riemannian manifold of covariance matrices [105]. These connectivity coefficients can boost the sensitivity for comparing diseased versus patient populations [66, 105]. It is also possible to define frequency-specific connectivity strength by decomposing the original time-series into multiple frequency sub-bands and correlating signals separately within these sub-bands [106].

A few studies depart from this routine. In graph-theoretic analysis, it is common to represent parcellated brain regions as graph nodes and functional connectivity between nodes as edge weights. This graph based representation of functional connectivity, the human ”connectome”, has been used to infer various topological characteristics of brain networks, such as modularity, clustering, small-worldedness etc. Some discriminative models have exploited these graph-based measures for individual-level predictions [107, 108, 13], although they are more commonly used for comparing groups. While limited

in number, a few studies have also explored rs-fMRI features beyond RSFC. Amplitude of low-frequency fluctuations (ALFF) and local synchronization of rs-fMRI signals or Regional Homogeneity (ReHo) are two alternate measures for studying spontaneous brain activity that have shown discriminative ability [109, 110]. More recently, several studies have also begun to explore the predictive capacity of dynamic FC in supervised models [111, 112].

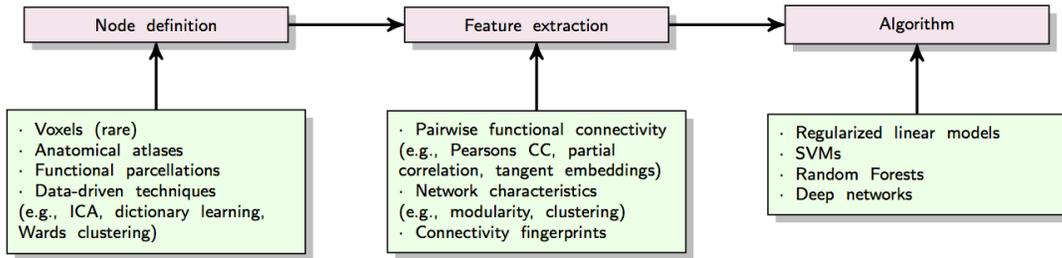


Figure 9: A summary of design choices for supervised learning with rs-fMRI

4.2. Feature selection

The goal of feature selection is to remove noisy, redundant or irrelevant features from the data while minimizing the information loss. Feature selection can often be an advantageous pre-processing step for training supervised learning algorithms, especially in the low sample size regime. In the absence of adequate regularization, large number of features can result in a loss of generalization power. Selecting a subset of features with highest relevance can thus help in building better generalizable models while reducing computational complexity.

Feature selection can be performed in a supervised or unsupervised fashion. Supervised or semi-supervised feature selection techniques choose a subset of features based on their ability to distinguish samples from different classes. These methods thus rely on class labels and can be further classified into filter, wrapper or embedded type models. Filter models first rank features by their importance/relevance for the classification task based on a statistical measure (e.g. t-test) and then select the top-ranked features. Wrapper models select feature subsets based on their predictive accuracy and thus need a pre-determined classification algorithm. Wrapper models thus perform better as they take into account the prediction accuracy estimates during feature selection. Due to the repeated learning and cross-validation,

however, these models are computationally prohibitive. Embedded models combine the advantages of the two by integrating feature selection into the learning algorithm. Regression models such as LASSO belong to this category as they implicitly select features by encouraging sparsity. These feature selection methods are discussed in depth in a detailed review by Tang et al. [113].

An alternative for feature selection is input dimensionality reduction. Methods like PCA or LLE belong to the category of unsupervised feature selection techniques and have been used to reduce the feature set to a manageable size in several studies. However, as pointed out in [114], these are not at all guaranteed to improve classification performance since they are oblivious to class labels.

Further, whether or not feature selection is necessary also depends on the downstream learning algorithm. Support vector machines, in general, deal well with high-dimensional data because of an implicit regularization. In the context of SVMs, Vapnik et al. [115] have shown that an upper bound on generalization error is independent of the number of features. Regularized models, in general, are capable of handling large feature sets. A drawback is that these models necessitate cross-validation to tune hyper-parameters such as the weight of the regularization penalty. This can reduce the effective sample size available for training and/or independent testing.

In some situations, it might be beneficial to exploit domain knowledge to guide feature selection. For example, if certain anatomical regions are known to have altered functional connectivity in disease based on prior studies, it might be advantageous to use this prior knowledge for constructing a focused feature set.

4.3. Methods

The majority of supervised learning methods applied to rs-fMRI are discriminant-based, i.e., they discriminate between classes without any prior assumptions about the generative process. The focus is on correctly estimating the boundaries between classes of interest. Learning algorithms for the same discriminant function (e.g., linear) can be based on different objective functions, giving rise to distinct models. We describe common models below.

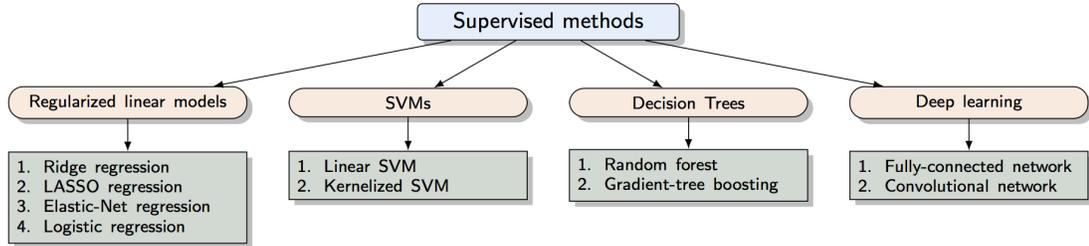


Figure 10: A taxonomy of supervised learning methods used for rs-fMRI analysis

4.3.1. Regularized linear models

A large class of supervised learning algorithms are based on regularized linear models. The goal is to predict a target variable Y given input features X . Without loss of generality and for notational convenience, let us assume that the feature vector contains a single constant entry equal to 1, which allows us to account for a bias term. These algorithms differ in the choice of their likelihood model, $P(Y|X, w)$ and/or prior $P(w)$, where w denotes the parameters of the model. These methods yield optimization problems that are based on a conditional likelihood estimation or a maximum *a posteriori* estimation (MAP) framework.

$$w_{opt} = \arg \max_w P(Y|X, w) : \text{Conditional likelihood}$$

$$w_{opt} = \arg \max_w P(w|X, Y) : \text{MAP}$$

Ridge regression. Ridge regression is another widely used supervised learning algorithm belonging to the class of regularized linear models. The goal is to predict a real-valued output Y given input features X . The conditional likelihood in this algorithm is specified as a multivariate normal distribution where the mean parameter is modelled as a linear combination of input features, i.e., $Y|X \sim \mathcal{N}(w^T x, \sigma^2 I)$. The prior on weight parameters is often modelled a zero-mean gaussian with a diagonal covariance matrix, i.e., $w \sim \mathcal{N}(0, \tau^2 I)$. The optimal weight parameters w are thus optimized within a maximum *a posteriori* estimation (MAP) framework according to,

$$w_{opt} = \arg \min_w \sum_{i=1}^n \frac{1}{2\sigma^2} (y_i - w^T x_i)^2 + \frac{1}{2\tau^2} \|w\|_2^2$$

The MAP estimation problem above is convex and admits an elegant analytical solution.

Logistic Regression. Logistic regression employs a Bernoulli distribution to model the conditional probability of an output class Y given the input features X , i.e. $Y|X \sim \text{Bernoulli}(\mu)$. The mean parameter, μ , is specified with a logistic link function $\sigma(\cdot)$ using a linear combination of input features, i.e., $\mu = \sigma(w^T x)$. Given data $\{(x_i, y_i), i = 1, \dots, n\}$, the model parameters w are optimized within a conditional maximum likelihood framework by solving the following convex optimization problem,

$$w_{opt} = \arg \min_w \sum_{i=1}^n \log(1 + \exp(-y_i w^T x_i))$$

The training objective is optimized using iterative methods such as gradient descent or Newton's method. Regularized variants of logistic regression incorporate priors on the weight parameters (e.g., multivariate gaussian) and optimize the MAP estimates instead of the conditional likelihood estimates.

4.3.2. Support Vector Machines (SVMs)

The SVM is the most widely used classification/regression algorithm in rs-fMRI studies. SVMs search for an optimal separating hyperplane between classes that maximizes the margin, i.e., the distance from hyperplane to points closest to it on either side. This results in a classifier of the form $f(x) = \text{sign}(w^T x)$. The model parameters are obtained by solving the following convex optimization problem:

$$w_{opt} = \arg \min_w C \sum_{i=1}^n \max(0, 1 - y_i(w^T x_i)) + \|w_{-b}\|^2.$$

$\|w_{-b}\|^2$ is the L2 norm of the weight vector excluding the bias term. C controls the capacity of the model and determines the margin of the classifier. Tuning C can control overfitting and reduce the generalization error of the model. The resulting classification model is determined by only a subset of training instances that are closest to the boundary, known as the support vectors. SVMs can be extended to seek non-linear separating boundaries via adopting a so-called kernel function. The kernel function, which quantifies the similarity between pairs of points, implicitly maps the input data to higher dimensions. Conceptually, the use of kernel functions allows incorporation of domain-specific measures of similarity. For example, graph-based kernels, such as Weisfeiler-Lehman subtree kernel, can define a distance metric on the graphical representation of functional connectivity data for classification directly in the graph space.

4.3.3. Decision trees and random forests

Decision trees predict the output Y based on a sequence of splits in the input feature space X . The tree is a directed acyclic graph whose nodes represent decision points and edges represent their outcomes. The traversal of this tree in conjunction leads up to a target outcome prediction when a node with no children (leaf node) has been reached. Decision trees are often constructed in a top-down greedy fashion where nodes are split at each step by optimizing a metric that quantifies the consistency between predictions and ground truth. For example, in classification, an often-used information-theoretic metric for quantifying this consistency is Information-Gain, i.e., the reduction in entropy of Y after knowing X . Mathematically, this is expressed as

$$IG(Y, X) = H(Y) - H(Y|X)$$

where H denotes the Shannon entropy. Based on this metric, the first split will use the attribute of X that gives the maximum information gain. Decision trees can offer interpretability, often at the cost of reduced accuracy. Ensembles of decision trees, such as random forests or boosted trees, are thus a more popular choice in most applications since they yield much better prediction performance.

4.3.4. Deep neural networks

An ideal machine learning system should be highly automated, with limited hand-crafting in feature extraction as well as minimal assumptions about the nature of mapping between data and labels. The system should be able to mechanistically learn patterns useful for prediction from observed labelled data. Neural networks are highly promising methods for automated learning. This stems from their capability to approximate arbitrarily complex functions given sufficient labelled data [116].

Deep learning based models or neural networks define a mapping $Y = f(X; \theta)$ and optimize for parameters θ that yield the best functional approximation. The function $f(\cdot)$ is typically composed as a concatenation of simple nonlinear functions, often referred to as layers. A widely-used layer is a fully-connected layer that linearly combines the input variables, and applies a simple elementwise non-linear functions such as a sigmoid. The number of layers determines the depth of the network and controls the complexity of the model. The weights and biases of the layers are optimized via gradient descent based methods to minimize an objective function that quantifies

the empirical risk. Traditionally, the use of neural network algorithms has been limited since neuroimaging is a data-scarce domain, making it difficult to learn a reliable mapping between input and prediction variables. However, with data sharing and open release of large-scale neuroimaging data repositories, neural networks have recently gained adoption in the the rs-fMRI community for supervised prediction tasks. Neural networks with fully connected dense layers have been adopted to learn arbitrary mappings from connectivity features to disease labels [97, 98]. Recently, more advanced neural networks models with local receptive fields, like convolutional neural networks (CNNs), have shown promising classification accuracy using rs-fMRI data [117]. CNNs replace the fully-connected operations by convolutions with a set of learnable filters. Success of this approach stems from its ability to exploit the full-resolution 3D spatial structure of rs-fMRI without having to learn too many model parameters, thanks to the weight sharing in CNNs.

4.3.5. Comments

I. Strengths/weaknesses of diverse approaches. All algorithms have their own strengths and weaknesses and the choice of approach should be driven by several factors such as the prediction task, sample size, and nature of the input features. The training objective in common supervised learning algorithms used for neuroimaging applications, such as regularized linear models or SVMs, is often a combination of two terms: a data loss term that is a measure of the empirical risk or training error and a regularization penalty for the prior that helps combat over-fitting during learning (generalization error). The penalty norm can be critical and is often constrained by our prior knowledge about the data. L1 penalties encourage sparsity in weights whereas L2 penalties can allow kernelization and thus enable non-linear decision functions. L2 penalties lead to dense priors and are useful in learning problems where all features are expected to contribute to the predictive model. L1 penalties are useful when prior belief suggests that only a subset of features will contribute to predictions. Some regression models, e.g., Elastic-Net, employ a linear combination of both these penalties at the expense of an additional hyperparameter for tuning the trade-off between the two. The algorithmic choice is also affected by the end-goal. Models like decision trees or LASSO are often preferred when interpretability is desired over optimal performance whereas high-complexity models like SVMs, Random Forests or Neural Networks are imperative if the goal is to maximize performance.

II. Comments on sample sizes. An important question arises: What is an appropriate sample size for training supervised learning models? Unsurprisingly, research has shown that the sample size needed for learning is dependent on the complexity of the model. Powerful non-linear algorithms typically require more training examples to be effective. In general, one would also expect that the more features in the data, the more training examples would be required to characterize their distribution. Hence, the minimum training size for training a ML algorithm is in general a complex function of input dimensionality, complexity of the chosen model, quality of data, data heterogeneity, separability of classes etc.

Given the significant impact of sample size on classification performance, it is imperative to understand the nature of this relationship. There is significant ongoing research in answering this question using learning curves. These curves model the relationship between sample size and generalization error and can be used to predict the sample size required to train a particular classifier. Several studies have shown that learning curves can be well-characterized with an inverse power-law functional form, with $E(n) \propto n^{-\beta}$, where E denotes the error and n denotes the sample size [118, 119]. Besides empirical justification, many studies have also provided theoretical motivations for the inverse power-law model. The parameters of the learning curve are fitted empirically for a given application domain based on prior classification studies. For traditional algorithms, learning curves are known to plateau, i.e., the performance gains are insignificant beyond a certain sample size. One significant advantage of deep learning methods is that given sufficient capacity, they scale remarkably well with more data. Given the recent surge of interest in single-subject predictions using rs-fMRI, estimating the learning curve for classification of rs-fMRI data could be invaluable for understanding sample size requirements in this domain.

Another critical issue relates to the robustness of the estimated prediction scores. Empirical studies have shown that small sample sizes, typical in neuroimaging studies, result in large error bars on the prediction accuracy. For instance, with a sample size of 100, Varoquaux et al., ballpark the error in estimated prediction accuracy of binary classification tasks to be close to 10%. With 1000 samples, this error reduces down to 3%. Large confidence bounds can potentially invalidate the conclusions of studies based on a small number of samples.

One possible strategy to overcome the limitations of insufficient sample sizes is to exploit unlabelled data in a semi-supervised fashion in order to

increase the effectiveness of supervised learning algorithms. Transfer learning techniques are another promising alternative for enhancing classification performance in the low-data regime. These methods exploit neural networks trained on large datasets or auxiliary tasks by fine-tuning them to a target dataset or classification task. These are relatively unexplored directions in the field of rs-fMRI analysis that hold significant potential to alleviate the sample size limitations.

III. Comments on model evaluation. Cross-validation is a model evaluation technique used to estimate the generalization error of a predictive model. A naive cross-validation strategy is holdout, wherein the data is randomly split into a training and test set and the test score in this single-run is used as an estimate of out-of-sample accuracy. Given the limited sample sizes in most neuroimaging studies, K-fold is the dominant cross-validation choice as it utilizes all data points for both training and validation through repeated holdout, yielding error estimates with much less variance than classic holdout. It first partitions the data into K non-overlapping subsets, $D = \{S_1, \dots, S_K\}$. For each fold i in $\{1, \dots, K\}$, the model is trained on $D - S_i$ and evaluated on S_i . The mean accuracy across all folds is then used to estimate the model performance. While K can be anything, common choices include 5 or 10. When K equals the number of samples in the training set, the resampling procedure is known as leave one-out cross-validation. This can be used with computationally inexpensive models when sample sizes are low, typically less than a hundred.

4.4. Applications of supervised learning in rs-fMRI

Studies harnessing resting-state correlations for supervised prediction tasks are evolving at an unprecedented scale. We describe some interesting applications of supervised machine learning in rs-fMRI below.

4.4.1. Brain development and aging

Machine learning methods have shown promise in investigating the developing connectome. In an early influential work, Dosenbach et al. [120] demonstrated the feasibility of using RSFC to predict brain maturation as measured by chronological age, in adolescents and young adults. Using SVM, they developed a functional maturation index based on predicted brain ages. Later studies showed that brain maturity can be reasonably predicted even

in diverse cohorts distributed across the human lifespan [121, 122]. These works posited rs-fMRI as a valuable tool to predict healthy neurodevelopment and exposed novel age-related dynamics of RSFC, such as major changes in FC of sensorimotor regions [122], or an increasingly distributed functional architecture with age [120]. In addition to characterizing RSFC changes accompanying natural aging, machine learning has also been used to identify atypical neurodevelopment [123].

4.4.2. *Neurological and Psychiatric Disorders*

Machine learning has been extensively deployed to investigate the diagnostic value of rs-fMRI data in various neurological and psychiatric conditions. Neurodegenerative diseases like Alzheimer’s disease [24, 107, 124], its prodromal state Mild cognitive impairment [125, 126, 127, 128], Parkinson’s [129], and Amyotrophic Lateral Sclerosis (ALS) [130] have been classified by ML models with promising accuracy using functional connectivity-based biomarkers. Brain atrophy patterns in neurological disorders like Alzheimer’s or Multiple Sclerosis appear well before behavioral symptoms emerge. Thus, neuroimaging-based biomarkers derived from structural or functional abnormalities are favorable for early diagnosis and subsequent intervention to slow down the degenerative process.

The biological basis of psychiatric disorders has been elusive and the diagnosis of these disorders is currently completely driven by behavioral assessments. rs-fMRI has emerged as a powerful modality to derive imaging-based biomarkers for making diagnostic predictions of psychiatric disorders. Supervised learning algorithms using RSFC have shown promising results for classifying or predicting symptom severity in a variety of psychiatric disorders, including schizophrenia [131, 98, 132, 133], depression [23, 134, 108], autism spectrum disorder [25, 111, 66, 117], attention-deficit hyperactivity disorder [135, 136], social anxiety disorder [137], post-traumatic stress disorder [138] and obsessive compulsive disorder [139]. Several novel network disruption hypotheses have emerged for these disorders as a consequence of these studies. Most of these prediction models are based on standard kernel-based SVMs, and rely on FC between ROI pairs as discriminative features.

4.4.3. *Cognitive abilities and personality traits*

Functional connectivity can also be used to predict individual differences in cognition and behavior [140]. In comparison to task-fMRI studies which capture a single cognitive dimension, the resting state encompasses a wide

repertoire of cognitive states due to its uncontrolled nature. This makes it a rich modality to capture inter-individual variability across multiple behavioral domains. ML models have been shown to predict fluid intelligence [46], sustained attention [141], memory performance [142, 143, 144], language scores [142] from RSFC-based biomarkers in healthy and pathological populations. Recently, the utility of these models was also shown to extend to personality traits such as neuroticism, extraversion, agreeableness and openness [145, 146].

Prediction of behavioral performance is useful in a clinical context to understand how RSFC disruptions in pathology relate to impaired cognitive functioning. Meskaldji et al. [143] used regression models to predict memory impairment in MCI patients from different connectivity measures. Siegel et al. [142] assessed the behavioral significance of network disruptions in stroke patients by training ridge regression models to relate RSFC and structure with performance in multiple domains (memory, language, attention, visual and motor tasks). Among them, memory deficits were better predicted by RSFC, whereas structure was more important for predicting visual and motor impairments. This study highlights how rs-fMRI can complement structural information in studying brain-behavior relationships.

4.4.4. Vigilance fluctuations and sleep studies

A handful of studies have employed machine learning to predict vigilance levels during rs-fMRI scans. Since resting-state studies demand no task-processing, subjects are prone to drifting between wakefulness and sleep. Classification of vigilance states during rs-fMRI is important to remove vigilance confounds and contamination. SVM classifiers trained on cortico-cortical RSFC have been shown to reliably detect periods of sleep within the sca [147, 148]. Tagliazucchi et al. [148] revealed loss of wakefulness in one-third subjects of the experimental cohort, as early as 3 minutes into the scanner. The findings are interesting: While resting state is assumed to capture wakefulness, this may not be entirely true even for very short scan durations. The utility of these studies should not remain limited to classification alone. Through appropriate interpretation and visualization techniques, machine learning can shed new light on the reconfiguration of functional organization as people drift into sleep.

Predicting individual differences in cognitive response after different sleep conditions (e.g. sleep deprivation) using machine learning analysis of rs-fMRI is another interesting research direction. There is significant interest

in examining RSFC alterations following sleep deprivation [149, 150]. While statistical analysis has elucidated the functional reorganization characteristic of sleep deprivation, much remains to be understood about the FC patterns associated with inter-individual differences in vulnerability to sleep deprivation. Yeo et al. [151] trained an SVM classifier on functional connectivity data in the well-rested state to distinguish subjects vulnerable to vigilance decline following sleep deprivation from more resilient subjects, and revealed important network differences between the groups.

4.4.5. Heritability

Understanding the genetic influence on brain structure and function has been a long-standing goal in neuroscience. In a recent study, Ge et al. employed a traditional statistical framework to quantify heritability of whole-brain FC estimates [152]. Investigations into the genetic and environmental underpinnings of RSFC were also pursued within a machine learning framework. Miranda-Dominguez et al. [153] trained an SVM classifier on individual FC signatures to distinguish sibling and twin pairs from unrelated subject pairs. The study unveiled several interesting findings. The ability to successfully predict familial relationships from resting-state fMRI indicates that aspects of functional connectivity are shaped by genetic or unique environmental factors. The fact that predictions remained accurate in young adult pairs suggests that these influences are sustained through development. Further, a higher accuracy of predicting twins compared to non-twin siblings implied that genetics (rather than environment) is likely the stronger predictive force.

4.4.6. Other neuroimaging modalities

Machine learning can also be used to interrogate the correspondence between rs-fMRI and other modalities. The most closely related modality is task-fMRI. Tavor et al. [154] trained multiple regression models to show that resting-state connectivity can predict task-evoked responses in the brain across several behavioral domains. The ability of rs-fMRI, that is a task-free regime, to predict the activation pattern evoked by multiple tasks suggests that resting-state can capture the rich repertoire of cognitive states that is reflected during task-based fMRI. The performance of these regression models was shown to generalize to pathological populations [155], suggesting the clinical utility of this approach to map functional regions in populations incapable of performing certain tasks.

Investigating how structural connections shape functional associations between different brain regions has been the focus of a large number of studies [156]. While neuro-computational models have been promising to achieve this goal, machine learning models are particularly well-equipped to capture inter-individual differences in the structure-function relationship. Deligianni et al. [157] proposed a structured-output multivariate regression model to predict resting-state functional connectivity from DWI-derived structural connectivity, and demonstrated the efficiency of this technique through cross-validation. Venkataraman et al. [158] introduced a novel probabilistic model to examine the relationships between anatomical connectivity measured using DWI tractography and RSFC. Their formulation assumes that the two modalities are generated from a common connectivity template. Estimated latent connectivity estimates were shown to discriminate between control and schizophrenic populations, thereby indicating that joint modelling can also be useful in a clinical context.

Table 4: Key papers for various supervised learning application domains

Brain development and aging
<i>Prediction of individual brain maturity using fMRI (Dosenbach et al.,2010)[120]</i> Method: SVM, Target: Age, Contribution: Early influential work demonstrating the feasibility of using RSFC features for predicting brain maturation.
Neurological and Psychiatric Disorders
<i>Classification of Alzheimer disease, mild cognitive impairment, and normal cognitive status with large-scale network analysis based on resting-state functional MR imaging. (Chen et al.,2011)[24]</i> Method: Fisher LDA, Target: Alzheimer/MCI/controls, Contribution: Early work highlighting the potential of RSFC to diagnose neurological disorders
<i>Deriving reproducible biomarkers from multi-site resting-state data: An Autism-based example(Abraham et al.,2017)[66]</i> Method: Multiple, Target: ASD/controls, Contribution: Extensively evaluated the impact of ROI choice, connectivity metric and classifier on prediction performance in intra-site and inter-site settings
<i>Altered resting state complexity in schizophrenia (Bassett et al.,2012)[132]</i> Method: SVM, Target: Schizophrenia/controls, Contribution: Demonstrated the utility of resting-state network complexity measures in distinguishing patients with schizophrenia
Cognitive abilities and personality traits
<i>Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity (Finn et al.,2015)[46]</i> Method: Linear regression, Target: Fluid intelligence, Contribution: Demonstrated that RSFC can uniquely identify individuals and reliably predict fluid intelligence
<i>Disruptions of network connectivity predict impairment in multiple behavioral domains after stroke (Siegel et al., 2016) [142]</i> Method: Ridge regression, Target: Multiple cognitive measures , Contribution: Demonstrated the ability of ML coupled with RSFC to predict cognitive deficits in clinical populations
Vigilance fluctuations and sleep studies
<i>Automatic sleep staging using fMRI functional connectivity data (Tagliazucchi et al.,2012) [147]</i> <i>Decoding wakefulness levels from typical fMRI resting-state data reveals reliable drifts between wakefulness and sleep (Tagliazucchi et al.,2014) [148]</i> Method: SVM, Target: NREM sleep stages/wakefulness, Contribution: Demonstrated the ability of ML to detect sleep stages from resting-state
Heritability
<i>Heritability of the human connectome: A connectotyping study (Miranda-Dominguez et al.,2018) [153]</i> Method: SVM, Target: Twins/sibling/unrelated, Contribution: Provided evidence for relationship between genetics and RSFC through predictive modelling
Other neuroimaging modalities
<i>Task-free MRI predicts individual differences in brain activity during task performance (Tavor et al.,2016) [154]</i> Method: Multiple regression models, Target: Task-activation map, Contribution: Demonstrated that resting-state can capture a rich repertoire of cognitive states expressed during different behavioral tasks

5. Discussion

5.1. Practical advice for machine learning practitioners

Any machine learning application requires the following: (a) a model that reflects assumed relationships between measurements and other inductive biases, (b) a cost function to quantify how well the model captures our data and finally, (c) an appropriate optimization algorithm to minimize the cost. Successful application of machine learning to rs-fMRI requires a holistic perspective of how these algorithms work, what it means when they fail and most importantly, how to choose an algorithm for a given task or hypothesis. There are three crucial factors that could dictate this choice:

1. What is the research question? What is our prior belief? Unsupervised learning tackles questions about the data-generating process. For example, clustering and decomposition approaches have both been widely used for disentangling the underlying causal sources of rs-fMRI data. However, they represent different prior beliefs and often answer distinct research questions. For example, in the context of discovering RSNs, ICA assumes that the latent components are independent and seeks to recover spatial loci of sources of activation. This decomposition further enables separation of functional activity from noise sources. On the other hand, clustering generally assumes that the activation of each spatial location/region can be explained by exactly one underlying component from a set of clusters. Because this approach results in disjoint functional networks, clustering is the dominant approach for learning spatially contiguous whole-brain parcellations.

When the goal is to make predictions, supervised learning algorithms are the usual choice. The choice of a supervised model again depends on the research question: Is the goal to understand the relationship between labels and features or to build a diagnostic tool? Interpretability is key for the former application whereas highest accuracy can be construed as the primary goal for the latter. Model complexity must thus be chosen in accordance with this end-goal. We recommend that these goals be well-defined before model development.

2. How much data is needed? It is important to assess the quantity of data and whether or not it is feasible to acquire more data. Sample sizes can constrain model complexity. More training examples are required to capture a non-linear relationship between features and labels, than a linear relation-

ship. Data fidelity and regularization must also be weighed in accordance with the sample size. With small sample sizes, regularization becomes even more critical as the model is more likely to overfit on training samples.

3. What is the computational budget? Sometimes, the computational budget can be restrictive. For example, certain algorithms like deep neural networks, have a high computational demand that may not be sustained by available resources. Further, if the number of features is very large, training even low-complexity models can be time consuming. In such cases, models with lower run-timing complexity can take precedence, especially for early investigations. Time, computational budget or space constraints thus must be identified while choosing an appropriate model.

5.2. Limitations and opportunities

Many state-of-the-art techniques for rs-fMRI analysis are rooted in machine learning. Both unsupervised and supervised learning methods have substantially expanded the application domains of rs-fMRI. With large-scale compilation of neuroimaging data and progresses in learning algorithms, an even greater influence is expected in future. Despite the practical successes of machine learning, it is important to understand the challenges encountered in its current application to rs-fMRI. We outline some important limitations and unexplored opportunities below.

One of the biggest challenges associated with unsupervised learning methods is that there is no ground truth for evaluation. There is no *a priori* universal functional map of the brain to base comparisons between parcellation schemes. Further, whole-brain parcellations are often defined at different scales of functional organization, ranging from a few large-scale parcels to several hundreds of regions, making comparisons even more challenging. Although several evaluation criteria have been developed that account for this variability, no single learning algorithm has emerged to be consistently superior in all. Due to the trade-offs among diverse approaches, the choice of which parcellation to use as reference for network analysis is thus largely subjective.

Unsupervised learning approaches for exploring network dynamics are similarly prone to subjectivity. Characterizing dynamic functional connectivity through discrete mental states is difficult, primarily because the repertoire of mental states is possibly infinite. While dFC states are thought to

reflect different cognitive processes, it is challenging to obtain a behavioral correspondence for distinct states since resting-state is not externally probed. This again makes interpretations hard and prone to subjective bias. Machine learning approaches in this direction have thus far relied on cluster statistics to fix the number of FC states. Non-parametric models (e.g. infinite HMMs) provide an unexplored, attractive framework as they adaptively determine the number of states based on the underlying data complexity.

A significant challenge in single-subject prediction using rs-fMRI is posed by the fact that rs-fMRI features can be described in multiple ways. There is no recognized gold-standard atlas for time-series extraction, nor is there a consensus on the optimal connectivity metric. Further, even the fMRI pre-processing strategies can vary considerably. Exploration across this space is cumbersome, especially for advanced machine learning models like neural networks that are slow to train. An ideal system should be invariant to these choices. However, this is hardly the case for rs-fMRI where large deviations have been reported in prediction performance in relationship to these factors [66].

Another challenge in training robust prediction systems on large populations stems from the heterogeneity of multi-site rs-fMRI data. Resting-state is easier to standardize across sites compared to task-based protocols since it does not rely on external stimuli. However, differences in acquisition protocols and scanner characteristics across sites still constitute a significant source of heterogeneity. Multi-site studies have shown little to no improvement in prediction accuracy compared to single-site studies, despite the larger sample sizes [25, 159]. While it is possible to normalize out site effects from data, more advanced tools are needed in practice to mitigate this bias.

High diagnostic accuracies achieved by supervised learning methods should be interpreted with caution. Several confounding variables can induce systematic biases in estimates of functional connectivity. For example, head motion is known to affect connectivity patterns in the default mode network and frontoparietal control network [160]. Further, motion profiles also vary systematically between subgroups of interest, e.g., diseased patients often move more than healthy controls. Apart from generating spurious associations, this could affect the interpretability of supervised prediction studies. Independent statistical analysis is critical to rule out the effect of confounding variables on predictions, especially when these variables differ across the groups being explored.

Methodological innovations are needed to improve prediction accuracy

to levels suitable for clinical translation. Several factors make comparison of methods across studies tedious. Cross-validation is the most commonly employed strategy for reporting performance of ML models. However, small sizes (common in rs-fMRI studies) are shown to yield large error bars [161], indicating that data-splits can significantly impact performance. Generalizability and interpretability should remain the key focus while developing predictive models on rs-fMRI data. These are critical attributes to achieve clinical translation of machine learning models. Uncertainty estimation is another challenge in any application of supervised learning; ideally, class assignments by any classification algorithm should be accompanied by an additional measure that reflects the uncertainty in predictions. This is especially important for clinical diagnosis, where it is important to know a reliability measure for individual predictions.

Most existing studies focus on classifying a single disease versus controls. The ability of a diagnostic system to discriminate between multiple psychiatric disorders is much more useful in a clinical setting [162]. Hence, there is a need to assess the efficacy of ML models for differential diagnosis. Integrating rs-fMRI with complementary modalities like diffusion-weighted MRI can possibly yield even better neurophenotypes of disease, and is another challenging yet promising research proposition.

6. Conclusions

We have presented a comprehensive overview of the current state-of-the-art of machine learning in rs-fMRI analysis. We have organized the vast literature on this topic based upon applications and techniques separately to enable researchers from both neuroimaging and machine learning communities to identify gaps in current practice.

Table 5: Key related review papers in the field

Multi-subject Independent Component Analysis of fMRI: A Decade of Intrinsic Networks, Default Mode, and Neurodiagnostic Discovery (Calhoun et al. 2009)[163]

A focused review of group ICA discussing methodologies, discovery of RSNs and their diagnostic potential

Imaging-based parcellations of the human brain (Eickhoff et al.,2018)[164]

A detailed exploration into approaches for deriving imaging based parcellations and lurking challenges in the field

Dynamic functional connectivity: Promise, issues, and interpretations (Hutchison et al.,2013)[165]

An early review on findings, methods and interpretations of dynamical functional connectivity

The Chronnectome: Time-Varying Connectivity Networks as the Next Frontier in fMRI Data Discovery (Calhoun et al.,2014) [166]

A detailed review of methods for dynamic functional connectivity analysis with a focus on decomposition techniques

The dynamic functional connectome: State-of-the-art and perspectives (Prete et al.,2017)[167]

A comprehensive review of analytical approaches for dynamic functional connectivity analysis and future perspectives

On the nature of resting fMRI and time-varying functional connectivity (Lurie et al.,2018)[168]

A discussion of diverse perspectives on time-varying connectivity in rs-fMRI

Clinical Applications of Resting State Functional Connectivity (Fox et al.,2010)[169]

An early short review focused on clinical applications of rs-fMRI

Single Subject Prediction of Brain Disorders in Neuroimaging: Promises and Pitfalls (Arbabshirani et al. 2017)[170]

Extensive survey of studies on single subject prediction of brain disorders, including opinions on promises/limitations

Acknowledgements

This work was supported by NIH R01 grants (R01LM012719 and R01AG053949), the NSF NeuroNex grant 1707312, and NSF CAREER grant (1748377).

- [1] B. Biswal, F. Z. Yetkin, V. M. Haughton, J. S. Hyde, Functional connectivity in the motor cortex of resting human brain using echo-planar MRI, *Magn Reson Med* 34 (4) (1995) 537–541.
- [2] C. F. Beckmann, M. DeLuca, J. T. Devlin, S. M. Smith, Investigations into resting-state connectivity using independent component analysis, *Philosophical Transactions of the Royal Society B: Biological Sciences* 360 (1457) (2005) 1001–1013. doi:10.1098/rstb.2005.1634.
URL <http://rstb.royalsocietypublishing.org/cgi/doi/10.1098/rstb.2005.1634>
- [3] D. Cordes, V. M. Haughton, K. Arfanakis, G. J. Wendt, P. A. Turski, C. H. Moritz, M. A. Quigley, M. E. Meyerand, Mapping functionally related regions of brain with functional connectivity MR imaging, *AJNR Am J Neuroradiol* 21 (9) (2000) 1636–1644.
- [4] J. S. Damoiseaux, S. A. Rombouts, F. Barkhof, P. Scheltens, C. J. Stam, S. M. Smith, C. F. Beckmann, Consistent resting-state networks across healthy subjects, *Proc. Natl. Acad. Sci. U.S.A.* 103 (37) (2006) 13848–13853.
- [5] M. De Luca, C. F. Beckmann, N. De Stefano, P. M. Matthews, S. M. Smith, fMRI resting state networks define distinct modes of long-distance interactions in the human brain, *Neuroimage* 29 (4) (2006) 1359–1367.
- [6] N. U. Dosenbach, D. A. Fair, F. M. Miezin, A. L. Cohen, K. K. Wenger, R. A. Dosenbach, M. D. Fox, A. Z. Snyder, J. L. Vincent, M. E. Raichle, et al., Distinct brain networks for adaptive and stable task control in humans, *Proceedings of the National Academy of Sciences* 104 (26) (2007) 11073–11078.
- [7] M. D. Fox, M. Corbetta, A. Z. Snyder, J. L. Vincent, M. E. Raichle, Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems, *Proceedings of the National Academy of Sciences* 103 (26) (2006) 10046–10051.
- [8] M. Hampson, B. S. Peterson, P. Skudlarski, J. C. Gatenby, J. C. Gore, Detection of functional connectivity using temporal correlations in mr images, *Human brain mapping* 15 (4) (2002) 247–262.

- [9] D. S. Margulies, A. C. Kelly, L. Q. Uddin, B. B. Biswal, F. X. Castellanos, M. P. Milham, Mapping the functional connectivity of anterior cingulate cortex, *Neuroimage* 37 (2) (2007) 579–588.
- [10] W. W. Seeley, V. Menon, A. F. Schatzberg, J. Keller, G. H. Glover, H. Kenna, A. L. Reiss, M. D. Greicius, Dissociable intrinsic connectivity networks for salience processing and executive control, *Journal of Neuroscience* 27 (9) (2007) 2349–2356.
- [11] S. M. Smith, P. T. Fox, K. L. Miller, D. C. Glahn, P. M. Fox, C. E. Mackay, N. Filippini, K. E. Watkins, R. Toro, A. R. Laird, C. F. Beckmann, Correspondence of the brain’s functional architecture during activation and rest, *Proc. Natl. Acad. Sci. U.S.A.* 106 (31) (2009) 13040–13045.
- [12] M. D. Greicius, G. Srivastava, A. L. Reiss, V. Menon, Default-mode network activity distinguishes alzheimer’s disease from healthy aging: evidence from functional mri, *Proceedings of the National Academy of Sciences* 101 (13) (2004) 4637–4642.
- [13] K. Supekar, V. Menon, D. Rubin, M. Musen, M. D. Greicius, Network analysis of intrinsic functional brain connectivity in Alzheimer’s disease, *PLoS Comput. Biol.* 4 (6) (2008) e1000100.
- [14] Y. I. Sheline, M. E. Raichle, Resting state functional connectivity in preclinical alzheimers disease, *Biological psychiatry* 74 (5) (2013) 340–347.
- [15] D. P. Kennedy, E. Courchesne, The intrinsic functional organization of the brain is altered in autism, *Neuroimage* 39 (4) (2008) 1877–1885.
- [16] C. S. Monk, S. J. Peltier, J. L. Wiggins, S.-J. Weng, M. Carrasco, S. Risi, C. Lord, Abnormalities of intrinsic functional connectivity in autism spectrum disorders, *Neuroimage* 47 (2) (2009) 764–772.
- [17] J. V. Hull, Z. J. Jacokes, C. M. Torgerson, A. Irimia, J. D. Van Horn, Resting-state functional connectivity in autism spectrum disorders: A review, *Frontiers in psychiatry* 7 (2017) 205.

- [18] A. Anand, Y. Li, Y. Wang, J. Wu, S. Gao, L. Bukhari, V. P. Mathews, A. Kalnin, M. J. Lowe, Activity and connectivity of brain mood regulating circuit in depression: a functional magnetic resonance study, *Biological psychiatry* 57 (10) (2005) 1079–1088.
- [19] M. D. Greicius, B. H. Flores, V. Menon, G. H. Glover, H. B. Solvason, H. Kenna, A. L. Reiss, A. F. Schatzberg, Resting-state functional connectivity in major depression: abnormally increased contributions from subgenual cingulate cortex and thalamus, *Biological psychiatry* 62 (5) (2007) 429–437.
- [20] P. C. Mulders, P. F. van Eijndhoven, A. H. Schene, C. F. Beckmann, I. Tendolkar, Resting-state functional connectivity in major depressive disorder: a review, *Neuroscience & Biobehavioral Reviews* 56 (2015) 330–344.
- [21] M. Liang, Y. Zhou, T. Jiang, Z. Liu, L. Tian, H. Liu, Y. Hao, Widespread functional disconnectivity in schizophrenia with resting-state functional magnetic resonance imaging, *Neuroreport* 17 (2) (2006) 209–213.
- [22] J. M. Sheffield, D. M. Barch, Cognition and resting-state functional connectivity in schizophrenia, *Neuroscience & Biobehavioral Reviews* 61 (2016) 108–120.
- [23] R. C. Craddock, P. E. Holtzheimer, X. P. Hu, H. S. Mayberg, Disease state prediction from resting state functional connectivity, *Magn Reson Med* 62 (6) (2009) 1619–1628.
- [24] G. Chen, B. D. Ward, C. Xie, W. Li, Z. Wu, J. L. Jones, M. Franczak, P. Antuono, S. J. Li, Classification of Alzheimer disease, mild cognitive impairment, and normal cognitive status with large-scale network analysis based on resting-state functional MR imaging, *Radiology* 259 (1) (2011) 213–221.
- [25] J. A. Nielsen, B. A. Zielinski, P. T. Fletcher, A. L. Alexander, N. Lange, E. D. Bigler, J. E. Lainhart, J. S. Anderson, Multisite functional connectivity MRI classification of autism: ABIDE results, *Front Hum Neurosci* 7 (2013) 599.

- [26] M. D. Fox, M. Greicius, Clinical applications of resting state functional connectivity, *Frontiers in systems neuroscience* 4 (2010) 19.
- [27] M. Greicius, Resting-state functional connectivity in neuropsychiatric disorders, *Current opinion in neurology* 21 (4) (2008) 424–430.
- [28] D. Zhang, M. E. Raichle, Disease and the brain’s dark energy, *Nature Reviews Neurology* 6 (1) (2010) 15.
- [29] D. Cordes, J. Carew, H. Eghbalnid, E. Meyerand, M. Quigley, K. Arfanakis, Resting-State Functional Connectivity Study using Independent Component Analysis, *Proceedings ISMRM* 1706.
URL <https://cds.ismrm.org/ismrm-1999/PDF6/1706.pdf>
- [30] C. F. Beckmann, S. M. Smith, Tensorial extensions of independent component analysis for multisubject fMRI analysis, *Neuroimage* 25 (1) (2005) 294–311.
- [31] M. D. Greicius, B. Krasnow, A. L. Reiss, V. Menon, Functional connectivity in the resting brain: a network analysis of the default mode hypothesis, *Proc. Natl. Acad. Sci. U.S.A.* 100 (1) (2003) 253–258.
- [32] M. Jung, H. Kosaka, D. N. Saito, M. Ishitobi, T. Morita, K. Inohara, M. Asano, S. Arai, T. Munesue, A. Tomoda, Y. Wada, N. Sadato, H. Okazawa, T. Iidaka, Default mode network in young male adults with autism spectrum disorder: relationship with autism spectrum traits, *Mol Autism* 5 (2014) 35.
- [33] D. Ongur, M. Lundy, I. Greenhouse, A. K. Shinn, V. Menon, B. M. Cohen, P. F. Renshaw, Default mode network abnormalities in bipolar disorder and schizophrenia, *Psychiatry Res* 183 (1) (2010) 59–68.
- [34] W. Koch, S. Teipel, S. Mueller, J. Benninghoff, M. Wagner, A. L. Bokde, H. Hampel, U. Coates, M. Reiser, T. Meindl, Diagnostic power of default mode network resting state fMRI in the detection of Alzheimer’s disease, *Neurobiol. Aging* 33 (3) (2012) 466–478.
- [35] D. Cordes, V. M. Haughton, K. Arfanakis, J. D. Carew, P. A. Turski, C. H. Moritz, M. A. Quigley, M. E. Meyerand, Frequencies contributing to functional connectivity in the cerebral cortex in ”resting-state” data, *AJNR Am J Neuroradiol* 22 (7) (2001) 1326–1333.

- [36] R. Salvador, J. Suckling, M. R. Coleman, J. D. Pickard, D. Menon, E. Bullmore, Neurophysiological architecture of functional magnetic resonance images of human brain, *Cerebral Cortex* 15 (9) (2005) 1332–2342. doi:10.1093/cercor/bhi016.
- [37] A. Mezer, Y. Yovel, O. Pasternak, T. Gorfine, Y. Assaf, Cluster analysis of resting-state fMRI time series, *NeuroImage* 45 (4) (2009) 1117–1125. doi:10.1016/j.neuroimage.2008.12.015.
URL <http://linkinghub.elsevier.com/retrieve/pii/S1053811908012706>
- [38] H. Laufs, A. Kleinschmidt, A. Beyerle, E. Eger, A. Salek-haddadi, C. Preibisch, K. Krakow, Eeg-correlated fmri of human alpha activity., *NeuroImage* 19 4 (2003) 1463–76.
- [39] J. S. Damoiseaux, M. D. Greicius, Greater than the sum of its parts: a review of studies combining structural connectivity and resting-state functional connectivity, *Brain Structure and Function* 213 (6) (2009) 525–533. doi:10.1007/s00429-009-0208-6.
URL <https://doi.org/10.1007/s00429-009-0208-6>
- [40] Y. Nir, R. Mukamel, I. Dinstein, E. Privman, M. Harel, L. Fisch, H. Gelbard-Sagiv, S. Kipervasser, F. Andelman, M. Y. Neufeld, U. Kramer, A. Arieli, I. Fried, R. Malach, Interhemispheric correlations of slow spontaneous neuronal fluctuations revealed in human sensory cortex, *Nat. Neurosci.* 11 (9) (2008) 1100–1108.
- [41] C. Chang, G. H. Glover, Time-frequency dynamics of resting-state brain connectivity measured with fMRI, *Neuroimage* 50 (1) (2010) 81–98.
- [42] E. A. Allen, E. Damaraju, S. M. Plis, E. B. Erhardt, T. Eichele, V. D. Calhoun, Tracking whole-brain connectivity dynamics in the resting state, *Cerebral Cortex* 24 (3) (2014) 663–676. doi:10.1093/cercor/bhs352.
- [43] D. Vidaurre, S. M. Smith, M. W. Woolrich, Brain network dynamics are hierarchically organized in time, *Proceedings of the National Academy of Sciences* 114 (48) (2017) 201705120. arXiv:arXiv:1408.1149, doi:10.1073/pnas.1705120114.
URL <http://www.pnas.org/lookup/doi/10.1073/pnas.1705120114>

- [44] H. Eavani, T. D. Satterthwaite, R. E. Gur, R. C. Gur, C. Davatzikos, Unsupervised learning of functional network dynamics in resting state fMRI, *Inf Process Med Imaging* 23 (2013) 426–437.
- [45] J. M. Reinen, O. Y. Chen, R. M. Hutchison, B. T. T. Yeo, K. M. Anderson, M. R. Sabuncu, D. Ongur, J. L. Roffman, J. W. Smoller, J. T. Baker, A. J. Holmes, The human cortex possesses a reconfigurable dynamic network architecture that is disrupted in psychosis, *Nat Commun* 9 (1) (2018) 1157.
- [46] E. S. Finn, X. Shen, D. Scheinost, M. D. Rosenberg, J. Huang, M. M. Chun, X. Papademetris, R. T. Constable, Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity, *Nature Neuroscience* 18 (11) (2015) 1664–1671. arXiv:15334406, doi:10.1038/nn.4135.
URL <http://dx.doi.org/10.1038/nn.4135>
- [47] N. Tzourio-Mazoyer, B. Landeau, D. Papathanassiou, F. Crivello, O. Etard, N. Delcroix, B. Mazoyer, M. Joliot, Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain, *Neuroimage* 15 (1) (2002) 273–289.
- [48] N. Leonardi, J. Richiardi, M. Gschwind, S. Simioni, J. M. Annoni, M. Schluep, P. Vuilleumier, D. Van De Ville, Principal components of functional connectivity: A new approach to study dynamic brain connectivity during rest, *NeuroImage* 83 (2013) 937–950. doi:10.1016/j.neuroimage.2013.07.019.
URL <http://dx.doi.org/10.1016/j.neuroimage.2013.07.019>
- [49] N. Leonardi, W. R. Shirer, M. D. Greicius, D. Van De Ville, Disentangling dynamic networks: Separated and joint expressions of functional connectivity patterns in time, *Human Brain Mapping* 35 (12) (2014) 5984–5995. doi:10.1002/hbm.22599.
- [50] R. C. Craddock, G. A. James, P. E. Holtzheimer, X. P. Hu, H. S. Mayberg, A whole brain fMRI atlas generated via spatially constrained spectral clustering, *Human Brain Mapping* 33 (8) (2012) 1914–1928. arXiv:NIHMS150003, doi:10.1002/hbm.21333.

- [51] V. D. Calhoun, T. Adali, G. D. Pearlson, J. J. Pekar, A method for making group inferences from functional MRI data using independent component analysis, *Hum Brain Mapp* 14 (3) (2001) 140–151.
- [52] B. CF, M. CE, N. Filippini, S. SM, Group comparison of resting-state fmri data using multi-subject ica and dual regression, *Neuroimage* 47. doi:10.1016/S1053-8119(09)71511-3.
- [53] Y. Du, Y. Fan, Group information guided ICA for fMRI data analysis, *NeuroImage* 69 (2013) 157–197. doi:10.1016/j.neuroimage.2012.11.008. URL <http://dx.doi.org/10.1016/j.neuroimage.2012.11.008>
- [54] G. Varoquaux, S. Sadaghiani, P. Pinel, A. Kleinschmidt, J. B. Poline, B. Thirion, A group model for stable multi-subject ICA on fMRI datasets, *NeuroImage* 51 (1) (2010) 288–299. arXiv:1006.2300, doi:10.1016/j.neuroimage.2010.02.010. URL <http://dx.doi.org/10.1016/j.neuroimage.2010.02.010>
- [55] I. Daubechies, E. Roussos, S. Takerkart, M. Benharrosh, C. Golden, K. D’Ardenne, W. Richter, J. D. Cohen, J. Haxby, Independent component analysis for brain fmri does not select for independence, *Proceedings of the National Academy of Sciences* 106 (26) (2009) 10415–10422. arXiv:<http://www.pnas.org/content/106/26/10415.full.pdf>, doi:10.1073/pnas.0903525106. URL <http://www.pnas.org/content/106/26/10415>
- [56] G. Varoquaux, A. Gramfort, F. Pedregosa, V. Michel, B. Thirion, Multi-subject dictionary learning to segment an atlas of brain spontaneous activity, *Inf Process Med Imaging* 22 (2011) 562–573.
- [57] A. Abraham, E. Dohmatob, B. Thirion, D. Samaras, G. Varoquaux, Extracting brain regions from rest fMRI with total-variation constrained dictionary learning, *Med Image Comput Comput Assist Interv* 16 (Pt 2) (2013) 607–615.
- [58] J. Lv, X. Jiang, X. Li, D. Zhu, S. Zhang, S. Zhao, H. Chen, T. Zhang, X. Hu, J. Han, J. Ye, L. Guo, T. Liu, Holistic atlases of functional networks and interactions reveal reciprocal organizational architecture of cortical function, *IEEE Trans Biomed Eng* 62 (4) (2015) 1120–1131.

- [59] Y. Golland, P. Golland, S. Bentin, R. Malach, Data-driven clustering reveals a fundamental subdivision of the human cortex into two global systems, *Neuropsychologia* 46 (2) (2008) 540–553. doi:10.1016/j.neuropsychologia.2007.10.003.
- [60] M. H. Lee, C. D. Hacker, A. Z. Snyder, M. Corbetta, D. Zhang, E. C. Leuthardt, J. S. Shimony, Clustering of resting state networks, *PLoS ONE* 7 (7) (2012) e40370.
- [61] J. H. Kim, J. M. Lee, H. J. Jo, S. H. Kim, J. H. Lee, S. T. Kim, S. W. Seo, R. W. Cox, D. L. Na, S. I. Kim, Z. S. Saad, Defining functional SMA and pre-SMA subregions in human MFC using resting state fMRI: functional connectivity-based parcellation method, *Neuroimage* 49 (3) (2010) 2375–2386.
- [62] P. Golland, Y. Golland, R. Malach, Detection of spatial activation patterns as unsupervised segmentation of fmri data, *MICCAI 10 Pt 1* (2007) 110–8.
- [63] B. T. Thomas Yeo, F. M. Krienen, J. Sepulcre, M. R. Sabuncu, D. Lashkari, M. Hollinshead, J. L. Roffman, J. W. Smoller, L. Zöllei, J. R. Polimeni, B. Fischl, H. Liu, R. L. Buckner, The organization of the human cerebral cortex estimated by intrinsic functional connectivity, *Journal of Neurophysiology* 106 (3) (2011) 1125–1165. doi:10.1152/jn.00338.2011. URL <http://www.physiology.org/doi/10.1152/jn.00338.2011>
- [64] D. Cordes, V. Haughton, J. D. Carew, K. Arfanakis, K. Maravilla, Hierarchical clustering to measure connectivity in fMRI resting-state data, *Magnetic Resonance Imaging* 20 (4) (2002) 305–317. doi:10.1016/S0730-725X(02)00503-9.
- [65] T. Blumensath, S. Jbabdi, M. F. Glasser, D. C. Van Essen, K. Ugurbil, T. E. Behrens, S. M. Smith, Spatially constrained hierarchical parcellation of the brain with resting-state fMRI, *Neuroimage* 76 (2013) 313–324.
- [66] A. Abraham, M. P. Milham, A. Di Martino, R. C. Craddock, D. Samaras, B. Thirion, G. Varoquaux, Deriving reproducible biomarkers from

- multi-site resting-state data: An Autism-based example, *Neuroimage* 147 (2017) 736–745.
- [67] B. Thirion, G. Varoquaux, E. Dohmatob, J. B. Poline, Which fMRI clustering gives good brain parcellations?, *Front Neurosci* 8 (2014) 167.
- [68] Y. Wang, T.-Q. Li, Analysis of whole-brain resting-state fmri data using hierarchical clustering approach, *PLOS ONE* 8 (10) (2013) 1–9. doi:10.1371/journal.pone.0076315.
URL <https://doi.org/10.1371/journal.pone.0076315>
- [69] M. van den Heuvel, R. Mandl, H. H. Pol, Normalized cut group clustering of resting-state fMRI data, *PLoS ONE* 3 (4). doi:10.1371/journal.pone.0002001.
- [70] X. Shen, F. Tokoglu, X. Papademetris, R. T. Constable, Groupwise whole-brain parcellation from resting-state fMRI data for network node identification, *NeuroImage* 82 (2013) 403–415. arXiv:NIHMS150003, doi:10.1016/j.neuroimage.2013.05.081.
URL <http://dx.doi.org/10.1016/j.neuroimage.2013.05.081>
- [71] N. Honnorat, H. Eavani, T. D. Satterthwaite, R. E. Gur, R. C. Gur, C. Davatzikos, GraSP: Geodesic Graph-based Segmentation with Shape Priors for the functional parcellation of the cortex, *NeuroImage* 106 (2015) 207–221. arXiv:NIHMS150003, doi:10.1016/j.neuroimage.2014.11.008.
URL <http://dx.doi.org/10.1016/j.neuroimage.2014.11.008>
- [72] M. Maier, U. von Luxburg, M. Hein, How the result of graph clustering methods depends on the construction of the graph, *ArXiv e-prints* arXiv:1102.2075.
- [73] E. M. Gordon, T. O. Laumann, B. Adeyemo, J. F. Huckins, W. M. Kelley, S. E. Petersen, Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations, *Cerebral Cortex* 26 (1) (2016) 288–303. doi:10.1093/cercor/bhu239.
- [74] M. F. Glasser, T. S. Coalson, E. C. Robinson, C. D. Hacker, J. Harwell, E. Yacoub, K. Ugurbil, J. Andersson, C. F. Beckmann, M. Jenkinson, S. M. Smith, D. C. Van Essen, A multi-modal parcellation of human cerebral cortex, *Nature Publishing Group* 536.

doi:10.1038/nature18933.

URL <http://balsa.wustl.edu/WN56>.

- [75] A. Schaefer, R. Kong, E. M. Gordon, T. O. Laumann, X.-N. Zuo, A. J. Holmes, S. B. Eickhoff, B. T. Yeo, Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI, *Cerebral Cortex* (2017) 1–20doi:10.1093/cercor/bhx179.
URL <https://academic.oup.com/cercor/article-lookup/doi/10.1093/cercor/bhx179>
- [76] R. Kong, J. Li, C. Orban, M. R. Sabuncu, H. Liu, A. Schaefer, N. Sun, X. N. Zuo, A. J. Holmes, S. B. Eickhoff, B. T. T. Yeo, Spatial Topography of Individual-Specific Cortical Networks Predicts Human Cognition, Personality, and Emotion, *Cereb. Cortex*.
- [77] M. Salehi, A. Karbasi, X. Shen, D. Scheinost, R. T. Constable, An exemplar-based approach to individualized parcellation reveals the need for sex specific functional networks, *NeuroImage* 170 (2018) 54–67. doi:10.1016/j.neuroimage.2017.08.068.
URL <http://dx.doi.org/10.1016/j.neuroimage.2017.08.068>
- [78] J. Kleinberg, An impossibility theorem for clustering, in: *Proceedings of the 15th International Conference on Neural Information Processing Systems, NIPS'02*, MIT Press, Cambridge, MA, USA, 2002, pp. 463–470.
URL <http://dl.acm.org/citation.cfm?id=2968618.2968676>
- [79] S. Arslan, S. I. Ktena, A. Makropoulos, E. C. Robinson, D. Rueckert, S. Parisot, Human brain mapping: A systematic comparison of parcellation methods for the human cerebral cortex, *Neuroimage* 170 (2018) 5–30.
- [80] M. Salehi, A. S. Greene, A. Karbasi, X. Shen, D. Scheinost, R. T. Constable, There is no single functional atlas even for a single individual: Parcellation of the human brain is state dependent, *bioRxiv*:<https://www.biorxiv.org/content/early/2018/10/01/431833.full.pdf>, doi:10.1101/431833.
URL <https://www.biorxiv.org/content/early/2018/10/01/431833>
- [81] E. Damaraju, E. A. Allen, A. Belger, J. M. Ford, S. McEwen, D. H. Mathalon, B. A. Mueller, G. D. Pearlson, S. G. Potkin, A. Preda, J. A.

- Turner, J. G. Vaidya, T. G. van Erp, V. D. Calhoun, Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia, *Neuroimage Clin* 5 (2014) 298–308.
- [82] B. Rashid, E. Damaraju, G. D. Pearlson, V. D. Calhoun, Dynamic connectivity states estimated from resting fMRI Identify differences among Schizophrenia, bipolar disorder, and healthy control subjects, *Front Hum Neurosci* 8 (2014) 897.
- [83] A. D. Barber, M. A. Lindquist, P. DeRosse, K. H. Karlsgodt, Dynamic Functional Connectivity States Reflecting Psychotic-like Experiences, *Biol Psychiatry Cogn Neurosci Neuroimaging* 3 (5) (2018) 443–453.
- [84] A. Abrol, E. Damaraju, R. L. Miller, J. M. Stephen, E. D. Claus, A. R. Mayer, V. D. Calhoun, Replicability of time-varying connectivity patterns in large resting state fMRI samples, *Neuroimage* 163 (2017) 160–176.
- [85] C. Wang, J. L. Ong, A. Patanaik, J. Zhou, M. W. Chee, Spontaneous eyelid closures link vigilance fluctuation with fMRI dynamic connectivity states, *Proc. Natl. Acad. Sci. U.S.A.* 113 (34) (2016) 9653–9658.
- [86] H. I. Suk, C. Y. Wee, S. W. Lee, D. Shen, State-space model with deep learning for functional dynamics estimation in resting-state fMRI, *Neuroimage* 129 (2016) 292–307.
- [87] L. R. Chai, A. N. Khambhati, R. Ciric, T. M. Moore, R. C. Gur, R. E. Gur, T. D. Satterthwaite, D. S. Bassett, Evolution of brain network dynamics in neurodevelopment, *Network Neuroscience* 1 (1) (2017) 14–30.
- [88] X. Li, D. Zhu, X. Jiang, C. Jin, X. Zhang, L. Guo, J. Zhang, X. Hu, L. Li, T. Liu, Dynamic functional connectomics signatures for characterization and differentiation of PTSD patients, *Hum Brain Mapp* 35 (4) (2014) 1761–1778.
- [89] V. D. Calhoun, J. Sui, K. Kiehl, J. Turner, E. Allen, G. Pearlson, Exploring the psychosis functional connectome: aberrant intrinsic networks in schizophrenia and bipolar disorder, *Front Psychiatry* 2 (2011) 75.

- [90] E. Amico, J. Goni, The quest for identifiability in human functional connectomes, *Sci Rep* 8 (1) (2018) 8254.
- [91] H. Eavani, T. D. Satterthwaite, R. Filipovych, R. E. Gur, R. C. Gur, C. Davatzikos, Identifying Sparse Connectivity Patterns in the brain using resting-state fMRI, *Neuroimage* 105 (2015) 286–299.
- [92] H. Eavani, T. D. Satterthwaite, R. E. Gur, R. C. Gur, C. Davatzikos, Discriminative sparse connectivity patterns for classification of fMRI Data, *Med Image Comput Comput Assist Interv* 17 (Pt 3) (2014) 193–200.
- [93] A. Qiu, A. Lee, M. Tan, M. K. Chung, Manifold learning on brain functional networks in aging, *Medical Image Analysis* 20 (1) (2015) 52–60. doi:10.1016/j.media.2014.10.006.
URL <http://dx.doi.org/10.1016/j.media.2014.10.006>
- [94] H. Shen, L. Wang, Y. Liu, D. Hu, Discriminative analysis of resting-state functional connectivity patterns of schizophrenia using low dimensional embedding of fMRI, *Neuroimage* 49 (4) (2010) 3110–3121.
- [95] X. Guo, K. C. Dominick, A. A. Minai, H. Li, C. A. Erickson, L. J. Lu, Diagnosing Autism Spectrum Disorder from Brain Resting-State Functional Connectivity Patterns Using a Deep Neural Network with a Novel Feature Selection Method, *Front Neurosci* 11 (2017) 460.
- [96] D. Erhan, Y. Bengio, A. Courville, P.-A. Manzagol, P. Vincent, S. Bengio, Why does unsupervised pre-training help deep learning?, *J. Mach. Learn. Res.* 11 (2010) 625–660.
URL <http://dl.acm.org/citation.cfm?id=1756006.1756025>
- [97] A. S. Heinsfeld, A. R. Franco, R. C. Craddock, A. Buchweitz, F. Meneguzzi, Identification of autism spectrum disorder using deep learning and the ABIDE dataset, *Neuroimage Clin* 17 (2018) 16–23.
- [98] J. Kim, V. D. Calhoun, E. Shim, J. H. Lee, Deep neural network with weight sparsity control and pre-training extracts hierarchical features and enhances classification performance: Evidence from whole-brain resting-state functional connectivity patterns of schizophrenia, *Neuroimage* 124 (Pt A) (2016) 127–146.

- [99] L. Xu, J. Neufeld, B. Larson, D. Schuurmans, Maximum margin clustering, in: L. K. Saul, Y. Weiss, L. Bottou (Eds.), *Advances in Neural Information Processing Systems 17*, MIT Press, 2005, pp. 1537–1544.
URL <http://papers.nips.cc/paper/2602-maximum-margin-clustering.pdf>
- [100] L. L. Zeng, H. Shen, L. Liu, D. Hu, Unsupervised classification of major depression using functional connectivity MRI, *Human Brain Mapping* 35 (4) (2014) 1630–1641. doi:10.1002/hbm.22278.
- [101] A. T. Drysdale, L. Grosenick, J. Downar, K. Dunlop, F. Mansouri, Y. Meng, R. N. Fetcho, B. Zebley, D. J. Oathes, A. Etkin, A. F. Schatzberg, K. Sudheimer, J. Keller, H. S. Mayberg, F. M. Gunning, G. S. Alexopoulos, M. D. Fox, A. Pascual-Leone, H. U. Voss, B. J. Casey, M. J. Dubin, C. Liston, Resting-state connectivity biomarkers define neurophysiological subtypes of depression, *Nat. Med.* 23 (1) (2017) 28–38.
- [102] K. Dadi, M. Rahim, A. Abraham, D. Chyzyk, M. Milham, B. Thirion, G. Varoquaux, Benchmarking functional connectome-based predictive models for resting-state fMRI, working paper or preprint (Jun. 2018).
URL <https://hal.inria.fr/hal-01824205>
- [103] G. Varoquaux, A. Gramfort, J.-B. Poline, B. Thirion, Brain covariance selection: Better individual functional connectivity models using population prior, in: *Proceedings of the 23rd International Conference on Neural Information Processing Systems - Volume 2, NIPS'10*, Curran Associates Inc., USA, 2010, pp. 2334–2342.
URL <http://dl.acm.org/citation.cfm?id=2997046.2997156>
- [104] S. M. Smith, K. L. Miller, G. Salimi-Khorshidi, M. Webster, C. F. Beckmann, T. E. Nichols, J. D. Ramsey, M. W. Woolrich, Network modelling methods for FMRI, *Neuroimage* 54 (2) (2011) 875–891.
- [105] G. Varoquaux, F. Baronnet, A. Kleinschmidt, P. Fillard, B. Thirion, Detection of brain functional-connectivity difference in post-stroke patients using group-level covariance modeling, *Med Image Comput Comput Assist Interv* 13 (Pt 1) (2010) 200–208.

- [106] J. Richiardi, H. Eryilmaz, S. Schwartz, P. Vuilleumier, D. Van De Ville, Decoding brain states from fMRI connectivity graphs, *Neuroimage* 56 (2) (2011) 616–626.
- [107] A. Khazaee, A. Ebrahimzadeh, A. Babajani-Feremi, Identifying patients with Alzheimer’s disease using resting-state fMRI and graph theory, *Clin Neurophysiol* 126 (11) (2015) 2132–2141.
- [108] A. Lord, D. Horn, M. Breakspear, M. Walter, Changes in community structure of resting state functional connectivity in unipolar depression, *PLoS ONE* 7 (8) (2012) e41282.
- [109] C. Z. Zhu, Y. F. Zang, M. Liang, L. X. Tian, Y. He, X. B. Li, M. Q. Sui, Y. F. Wang, T. Z. Jiang, Discriminative analysis of brain function at resting-state for attention-deficit/hyperactivity disorder, *Med Image Comput Comput Assist Interv* 8 (Pt 2) (2005) 468–475.
- [110] M. Mennes, X. N. Zuo, C. Kelly, A. Di Martino, Y. F. Zang, B. Biswal, F. X. Castellanos, M. P. Milham, Linking inter-individual differences in neural activation and behavior to intrinsic brain dynamics, *Neuroimage* 54 (4) (2011) 2950–2959.
- [111] T. Price, C. Y. Wee, W. Gao, D. Shen, Multiple-network classification of childhood autism using functional connectivity dynamics, *Med Image Comput Comput Assist Interv* 17 (Pt 3) (2014) 177–184.
- [112] T. M. Madhyastha, M. K. Askren, P. Boord, T. J. Grabowski, Dynamic connectivity at rest predicts attention task performance, *Brain Connect* 5 (1) (2015) 45–59.
- [113] J. Tang, S. Alelyani, H. Liu, Feature selection for classification: A review, in: *Data Classification: Algorithms and Applications*, 2014.
- [114] F. J. S. Pereira, T. M. Mitchell, M. Botvinick, Machine learning classifiers and fmri: A tutorial overview, *NeuroImage* 45 (2009) s199–s209.
- [115] V. Vapnik, O. Chapelle, Bounds on error expectation for support vector machines, *Neural Computation* 12 (2000) 2013–2036.
- [116] K. Hornik, M. Stinchcombe, H. White, Multilayer feedforward networks are universal approximators, *Neural Networks* 2 (5) (1989) 359

– 366. doi:[https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8).

URL <http://www.sciencedirect.com/science/article/pii/0893608089900208>

- [117] M. Khosla, K. Jamison, A. Kuceyeski, M. R. Sabuncu, Ensemble learning with 3D convolutional neural networks for connectome-based prediction, ArXiv e-prints arXiv:1809.06219.
- [118] J. Hestness, S. Narang, N. Ardalani, G. F. Diamos, H. Jun, H. Kianinejad, M. M. A. Patwary, Y. Yang, Y. Zhou, Deep learning scaling is predictable, empirically, CoRR abs/1712.00409.
- [119] S. Mukherjee, P. Tamayo, S. Rogers, R. M. Rifkin, A. Engle, C. Campbell, T. R. Golub, J. P. Mesirov, Estimating dataset size requirements for classifying dna microarray data, *Journal of computational biology : a journal of computational molecular cell biology* 10 2 (2003) 119–42.
- [120] N. U. Dosenbach, B. Nardos, A. L. Cohen, D. A. Fair, J. D. Power, J. A. Church, S. M. Nelson, G. S. Wig, A. C. Vogel, C. N. Lessov-Schlaggar, K. A. Barnes, J. W. Dubis, E. Feczko, R. S. Coalson, J. R. Pruett, D. M. Barch, S. E. Petersen, B. L. Schlaggar, Prediction of individual brain maturity using fMRI, *Science* 329 (5997) (2010) 1358–1361.
- [121] L. Wang, L. Su, H. Shen, D. Hu, Decoding lifespan changes of the human brain using resting-state functional connectivity MRI, *PLoS ONE* 7 (8) (2012) e44530.
- [122] T. B. Meier, A. S. Desphande, S. Vergun, V. A. Nair, J. Song, B. B. Biswal, M. E. Meyerand, R. M. Birn, V. Prabhakaran, Support vector machine classification and characterization of age-related reorganization of functional brain networks, *Neuroimage* 60 (1) (2012) 601–613.
- [123] G. Ball, P. Aljabar, T. Arichi, N. Tsur, D. Cox, N. Merchant, P. Nonnena, J. V. Hajnal, A. D. Edwards, S. J. Counsell, Machine-learning to characterise neonatal functional connectivity in the preterm brain, *Neuroimage* 124 (Pt A) (2016) 267–275.
- [124] E. Challis, P. Hurley, L. Serra, M. Bozzali, S. Oliver, M. Cercignani, Gaussian process classification of Alzheimer’s disease and mild cognitive impairment from resting-state fMRI, *Neuroimage* 112 (2015) 232–243.

- [125] C. Y. Wee, P. T. Yap, D. Zhang, L. Wang, D. Shen, Group-constrained sparse fMRI connectivity modeling for mild cognitive impairment identification, *Brain Struct Funct* 219 (2) (2014) 641–656.
- [126] X. Chen, H. Zhang, Y. Gao, C. Y. Wee, G. Li, D. Shen, High-order resting-state functional connectivity network for MCI classification, *Hum Brain Mapp* 37 (9) (2016) 3282–3296.
- [127] B. Jie, D. Zhang, C. Y. Wee, D. Shen, Topological graph kernel on multiple thresholded functional connectivity networks for mild cognitive impairment classification, *Hum Brain Mapp* 35 (7) (2014) 2876–2897.
- [128] C. Y. Wee, S. Yang, P. T. Yap, D. Shen, Sparse temporally dynamic resting-state functional connectivity networks for early MCI identification, *Brain Imaging Behav* 10 (2) (2016) 342–356.
- [129] D. Long, J. Wang, M. Xuan, Q. Gu, X. Xu, D. Kong, M. Zhang, Automatic classification of early Parkinson’s disease with multi-modal MR imaging, *PLoS ONE* 7 (11) (2012) e47714.
- [130] R. C. Welsh, L. M. Jelson-Swain, B. R. Foerster, The utility of independent component analysis and machine learning in the identification of the amyotrophic lateral sclerosis diseased brain, *Front Hum Neurosci* 7 (2013) 251.
- [131] A. Venkataraman, T. J. Whitford, C. F. Westin, P. Golland, M. Kubicki, Whole brain resting state functional connectivity abnormalities in schizophrenia, *Schizophr. Res.* 139 (1-3) (2012) 7–12.
- [132] D. S. Bassett, B. G. Nelson, B. A. Mueller, J. Camchong, K. O. Lim, Altered resting state complexity in schizophrenia, *Neuroimage* 59 (3) (2012) 2196–2207.
- [133] Y. Fan, Y. Liu, H. Wu, Y. Hao, H. Liu, Z. Liu, T. Jiang, Discriminant analysis of functional connectivity patterns on Grassmann manifold, *Neuroimage* 56 (4) (2011) 2058–2067.
- [134] L. L. Zeng, H. Shen, L. Liu, L. Wang, B. Li, P. Fang, Z. Zhou, Y. Li, D. Hu, Identifying major depression using whole-brain functional connectivity: a multivariate pattern analysis, *Brain* 135 (Pt 5) (2012) 1498–1507.

- [135] A. Eloyan, J. Muschelli, M. B. Nebel, H. Liu, F. Han, T. Zhao, A. D. Barber, S. Joel, J. J. Pekar, S. H. Mostofsky, B. Caffo, Automated diagnoses of attention deficit hyperactive disorder using magnetic resonance imaging, *Front Syst Neurosci* 6 (2012) 61.
- [136] D. A. Fair, J. T. Nigg, S. Iyer, D. Bathula, K. L. Mills, N. U. Dosenbach, B. L. Schlaggar, M. Mennes, D. Gutman, S. Bangaru, J. K. Buitelaar, D. P. Dickstein, A. Di Martino, D. N. Kennedy, C. Kelly, B. Luna, J. B. Schweitzer, K. Velanova, Y. F. Wang, S. Mostofsky, F. X. Castellanos, M. P. Milham, Distinct neural signatures detected for ADHD subtypes after controlling for micro-movements in resting state functional connectivity MRI data, *Front Syst Neurosci* 6 (2012) 80.
- [137] F. Liu, W. Guo, J. P. Fouche, Y. Wang, W. Wang, J. Ding, L. Zeng, C. Qiu, Q. Gong, W. Zhang, H. Chen, Multivariate classification of social anxiety disorder using whole brain functional connectivity, *Brain Struct Funct* 220 (1) (2015) 101–115.
- [138] Q. Gong, L. Li, M. Du, W. Pettersson-Yeo, N. Crossley, X. Yang, J. Li, X. Huang, A. Mechelli, Quantitative prediction of individual psychopathology in trauma survivors using resting-state fMRI, *Neuropsychopharmacology* 39 (3) (2014) 681–687.
- [139] B. J. Harrison, C. Soriano-Mas, J. Pujol, H. Ortiz, M. Lopez-Sola, R. Hernandez-Ribas, J. Deus, P. Alonso, M. Yucel, C. Pantelis, J. M. Menchon, N. Cardoner, Altered corticostriatal functional connectivity in obsessive-compulsive disorder, *Arch. Gen. Psychiatry* 66 (11) (2009) 1189–1200.
- [140] S. Mueller, D. Wang, M. D. Fox, B. T. Yeo, J. Sepulcre, M. R. Sabuncu, R. Shafee, J. Lu, H. Liu, Individual variability in functional connectivity architecture of the human brain, *Neuron* 77 (3) (2013) 586–595.
- [141] M. D. Rosenberg, E. S. Finn, D. Scheinost, X. Papademetris, X. Shen, R. T. Constable, M. M. Chun, A neuromarker of sustained attention from whole-brain functional connectivity, *Nat. Neurosci.* 19 (1) (2016) 165–171.

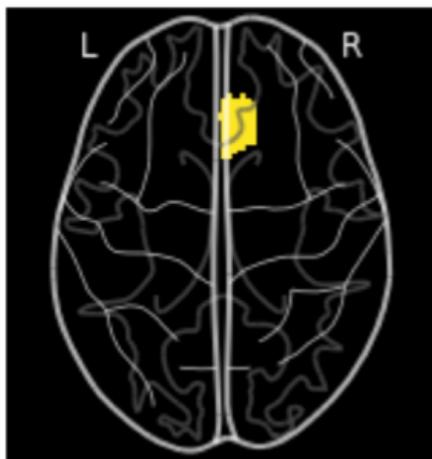
- [142] J. S. Siegel, L. E. Ramsey, A. Z. Snyder, N. V. Metcalf, R. V. Chacko, K. Weinberger, A. Baldassarre, C. D. Hacker, G. L. Shulman, M. Corbetta, Disruptions of network connectivity predict impairment in multiple behavioral domains after stroke, *Proc. Natl. Acad. Sci. U.S.A.* 113 (30) (2016) E4367–4376.
- [143] D. E. Meskaldji, M. G. Preti, T. A. Bolton, M. L. Montandon, C. Rodriguez, S. Morgenthaler, P. Giannakopoulos, S. Haller, D. Van De Ville, Prediction of long-term memory scores in MCI based on resting-state fMRI, *Neuroimage Clin* 12 (2016) 785–795.
- [144] D. C. Jangraw, J. Gonzalez-Castillo, D. A. Handwerker, M. Ghane, M. D. Rosenberg, P. Panwar, P. A. Bandettini, A functional connectivity-based neuromarker of sustained attention generalizes to predict recall in a reading task, *Neuroimage* 166 (2018) 99–109.
- [145] W. T. Hsu, M. D. Rosenberg, D. Scheinost, R. T. Constable, M. M. Chun, Resting-state functional connectivity predicts neuroticism and extraversion in novel individuals, *Soc Cogn Affect Neurosci* 13 (2) (2018) 224–232.
- [146] A. D. Nostro, V. I. Muller, D. P. Varikuti, R. N. Plaschke, F. Hoffstaedter, R. Langner, K. R. Patil, S. B. Eickhoff, Predicting personality from network-based resting-state functional connectivity, *Brain Struct Funct* 223 (6) (2018) 2699–2719.
- [147] E. Tagliazucchi, F. von Wegner, A. Morzelewski, S. Borisov, K. Jahnke, H. Laufs, Automatic sleep staging using fMRI functional connectivity data, *Neuroimage* 63 (1) (2012) 63–72.
- [148] E. Tagliazucchi, H. Laufs, Decoding wakefulness levels from typical fMRI resting-state data reveals reliable drifts between wakefulness and sleep, *Neuron* 82 (3) (2014) 695–708.
- [149] X. J. Dai, C. L. Liu, R. L. Zhou, H. H. Gong, B. Wu, L. Gao, Y. X. Wang, Long-term total sleep deprivation decreases the default spontaneous activity and connectivity pattern in healthy male subjects: a resting-state fMRI study, *Neuropsychiatr Dis Treat* 11 (2015) 761–772.
- [150] Y. Zhu, Z. Feng, J. Xu, C. Fu, J. Sun, X. Yang, D. Shi, W. Qin, Increased interhemispheric resting-state functional connectivity after

- sleep deprivation: a resting-state fMRI study, *Brain Imaging Behav* 10 (3) (2016) 911–919.
- [151] B. T. Yeo, J. Tandi, M. W. Chee, Functional connectivity during rested wakefulness predicts vulnerability to sleep deprivation, *Neuroimage* 111 (2015) 147–158.
- [152] T. Ge, A. J. Holmes, R. L. Buckner, J. W. Smoller, M. R. Sabuncu, Heritability analysis with repeat measurements and its application to resting-state functional connectivity, *Proc. Natl. Acad. Sci. U.S.A.* 114 (21) (2017) 5521–5526.
- [153] O. Miranda-Dominguez, E. Feczko, D. S. Grayson, H. Walum, J. T. Nigg, D. A. Fair, Heritability of the human connectome: A connectotyping study, *Netw Neurosci* 2 (2) (2018) 175–199.
- [154] I. Tavor, O. Parker Jones, R. B. Mars, S. M. Smith, T. E. Behrens, S. Jbabdi, Task-free MRI predicts individual differences in brain activity during task performance, *Science* 352 (6282) (2016) 216–220.
- [155] O. Parker Jones, N. L. Voets, J. E. Adcock, R. Stacey, S. Jbabdi, Resting connectivity predicts task activation in pre-surgical populations, *Neuroimage Clin* 13 (2017) 378–385.
- [156] F. Abdelnour, H. U. Voss, A. Raj, Network diffusion accurately models the relationship between structural and functional brain connectivity networks, *Neuroimage* 90 (2014) 335–347.
- [157] F. Deligianni, G. Varoquaux, B. Thirion, E. Robinson, D. J. Sharp, A. D. Edwards, D. Rueckert, A probabilistic framework to infer brain functional connectivity from anatomical connections, *Inf Process Med Imaging* 22 (2011) 296–307.
- [158] A. Venkataraman, Y. Rathi, M. Kubicki, C. F. Westin, P. Golland, Joint modeling of anatomical and functional connectivity for population studies, *IEEE Trans Med Imaging* 31 (2) (2012) 164–182.
- [159] C. Dansereau, Y. Benhajali, C. Risterucci, E. M. Pich, P. Orban, D. Arnold, P. Bellec, Statistical power and prediction accuracy in multisite resting-state fMRI connectivity, *Neuroimage* 149 (2017) 220–232.

- [160] K. R. Van Dijk, M. R. Sabuncu, R. L. Buckner, The influence of head motion on intrinsic functional connectivity MRI, *Neuroimage* 59 (1) (2012) 431–438.
- [161] G. Varoquaux, Cross-validation failure: Small sample sizes lead to large error bars, *Neuroimage* 180 (Pt A) (2018) 68–77.
- [162] T. Wolfers, J. K. Buitelaar, C. F. Beckmann, B. Franke, A. F. Marquand, From estimating activation locality to predicting disorder: A review of pattern recognition for neuroimaging-based psychiatric diagnostics, *Neurosci Biobehav Rev* 57 (2015) 328–349.
- [163] V. D. Calhoun, T. Adali, Multisubject independent component analysis of fmri: A decade of intrinsic networks, default mode, and neurodiagnostic discovery, *IEEE Reviews in Biomedical Engineering* 5 (2012) 60–73.
- [164] S. B. Eickhoff, B. T. T. Yeo, S. Genon, Imaging-based parcellations of the human brain, *Nature Reviews Neuroscience* 19 (2018) 672–686.
- [165] R. M. Hutchison, T. Womelsdorf, E. A. Allen, P. A. Bandettini, V. D. Calhoun, M. Corbetta, S. D. Penna, J. H. Duyn, G. H. Glover, J. Gonzalez-Castillo, D. A. Handwerker, S. D. Keilholz, V. Kiviniemi, D. A. Leopold, F. de Pasquale, O. Sporns, M. Walter, C. Chang, Dynamic functional connectivity: Promise, issues, and interpretations, *NeuroImage* 80 (2013) 360–378.
- [166] V. D. Calhoun, R. A. Miller, G. D. Pearlson, T. Adal, The chronnectome: Time-varying connectivity networks as the next frontier in fmri data discovery, *Neuron* 84 (2014) 262–274.
- [167] M. G. Preti, T. A. W. Bolton, D. V. D. Ville, The dynamic functional connectome: State-of-the-art and perspectives, *NeuroImage* 160 (2017) 41–54.
- [168] D. J. Lurie, D. Kessler, D. S. Bassett, R. F. Betzel, M. Breakspear, S. Keilholz, e. a. Aaron Kucyi, On the nature of resting fmri and time-varying functional connectivity., *PsyArXivdoi:10.31234/osf.io/xtzre*.
- [169] M. D. Fox, M. D. Greicius, Clinical applications of resting state functional connectivity, in: *Front. Syst. Neurosci.*, 2010.

- [170] M. Arbabshirani, S. M. Plis, J. Sui, V. D. Calhoun, Single subject prediction of brain disorders in neuroimaging: Promises and pitfalls, *NeuroImage* 145 (2017) 137–165.

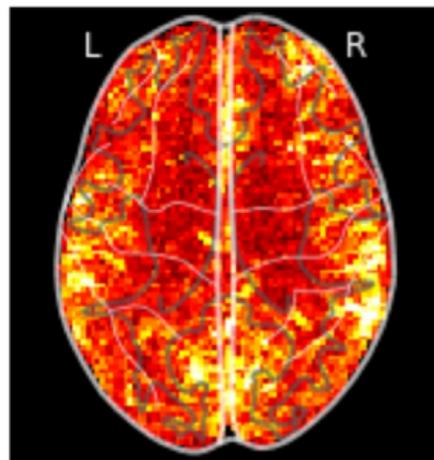
ACCEPTED MANUSCRIPT



Mean time series of seed

(a) Seed region

Correlate with
whole-brain rs-fMRI
time-series



**(b) Seed correlation network
(RSN)**

Figure 1

ML applications in rs-fMRI

Unsupervised learning

Discovering spatial patterns with coherent resting-state fluctuations (RSNs)

Discovering reproducible patterns of dynamic functional connectivity

Disentangling latent factors of inter-subject FC variation

Supervised learning

Brain development & aging

Neurological & psychiatric disorders

Cognitive abilities & personality traits

Vigilance fluctuations & sleep studies

Heritability & familiarity

Other modalities

Predicted variables

Age
Neurodevelopmental outcome

Alzheimer's · MCI · Parkinson's · ALS · Schizophrenia
Depression · ASD · ADHD
PTSD · OCD

Fluid intelligence · Attention
Memory · Language
Personality traits

Vigilance states
(wakefulness/sleep)

Genetic relationships
(sibling/twin/unrelated)

Task activations
Anatomical connectivity

Figure 2

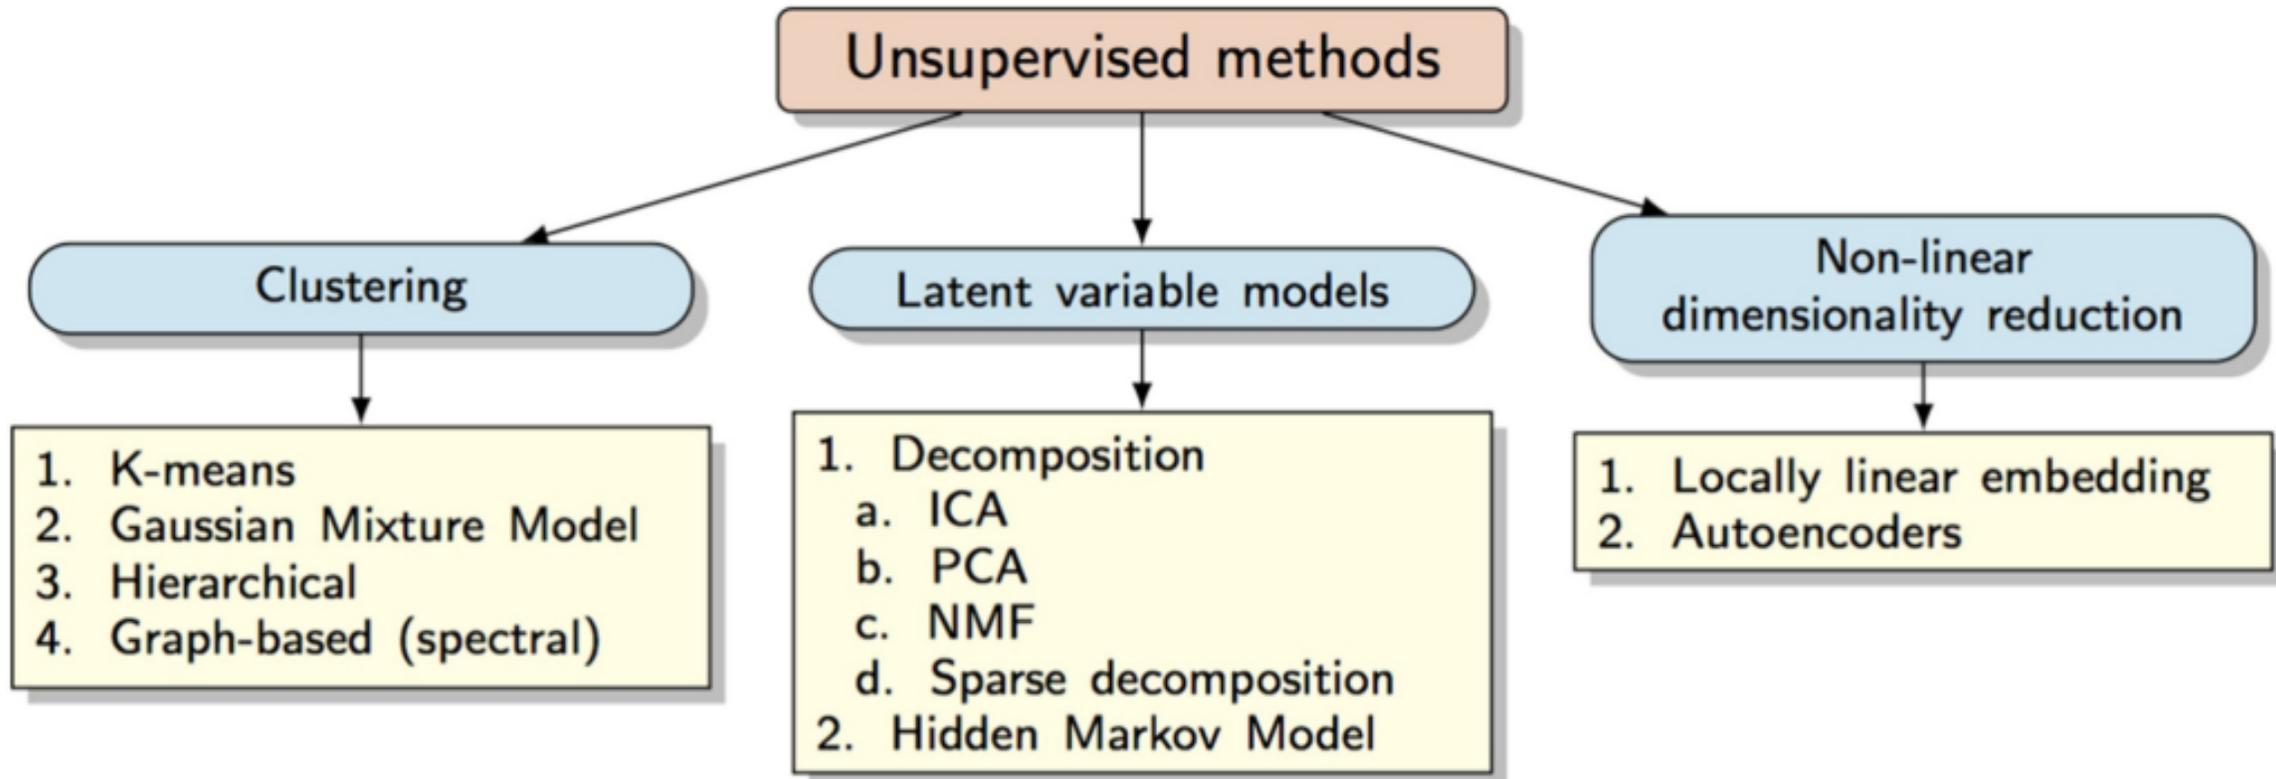
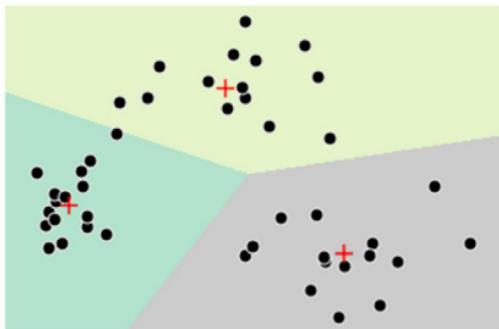
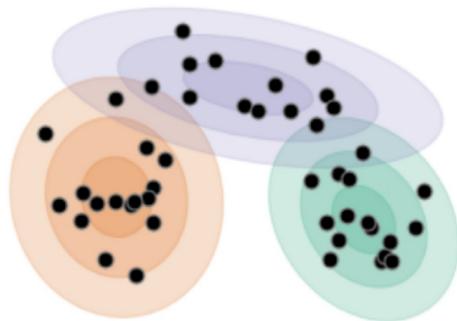


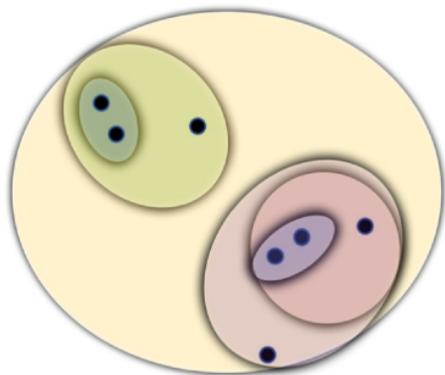
Figure 3



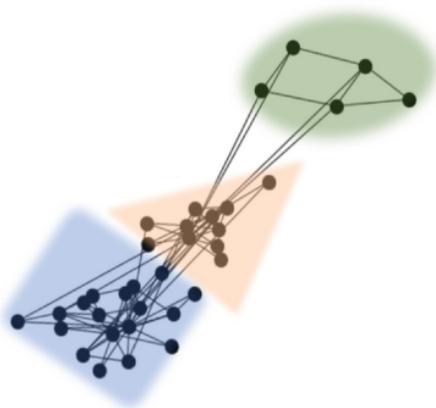
K-means clustering



Mixture model (Gaussian)



Hierarchical clustering



Graph based clustering

Figure 4

I. Discovering spatial maps with coherent temporal dynamics

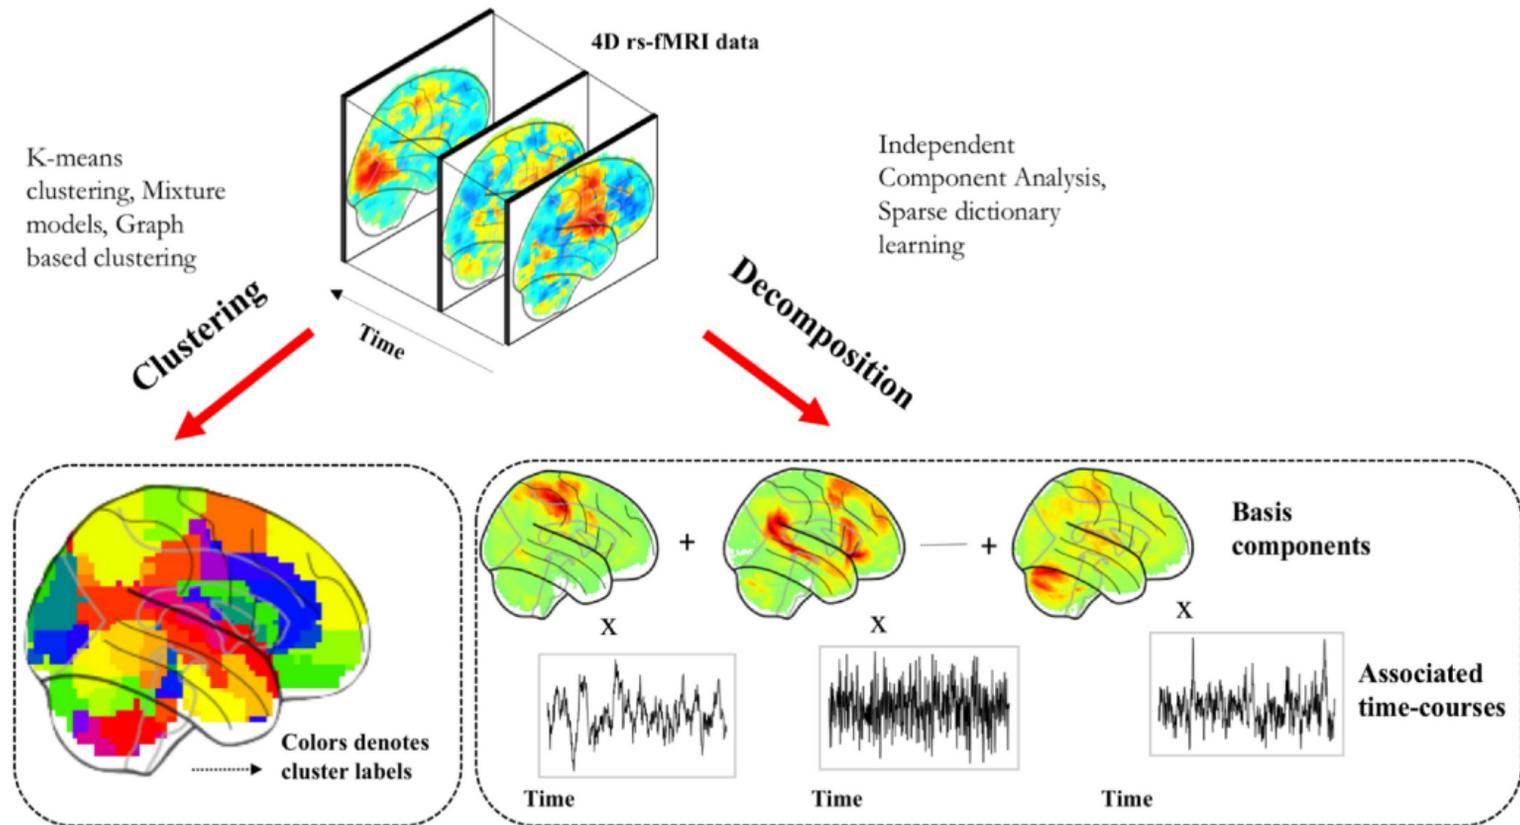


Figure 5

II. Discovering reproducible patterns of dynamic functional connectivity

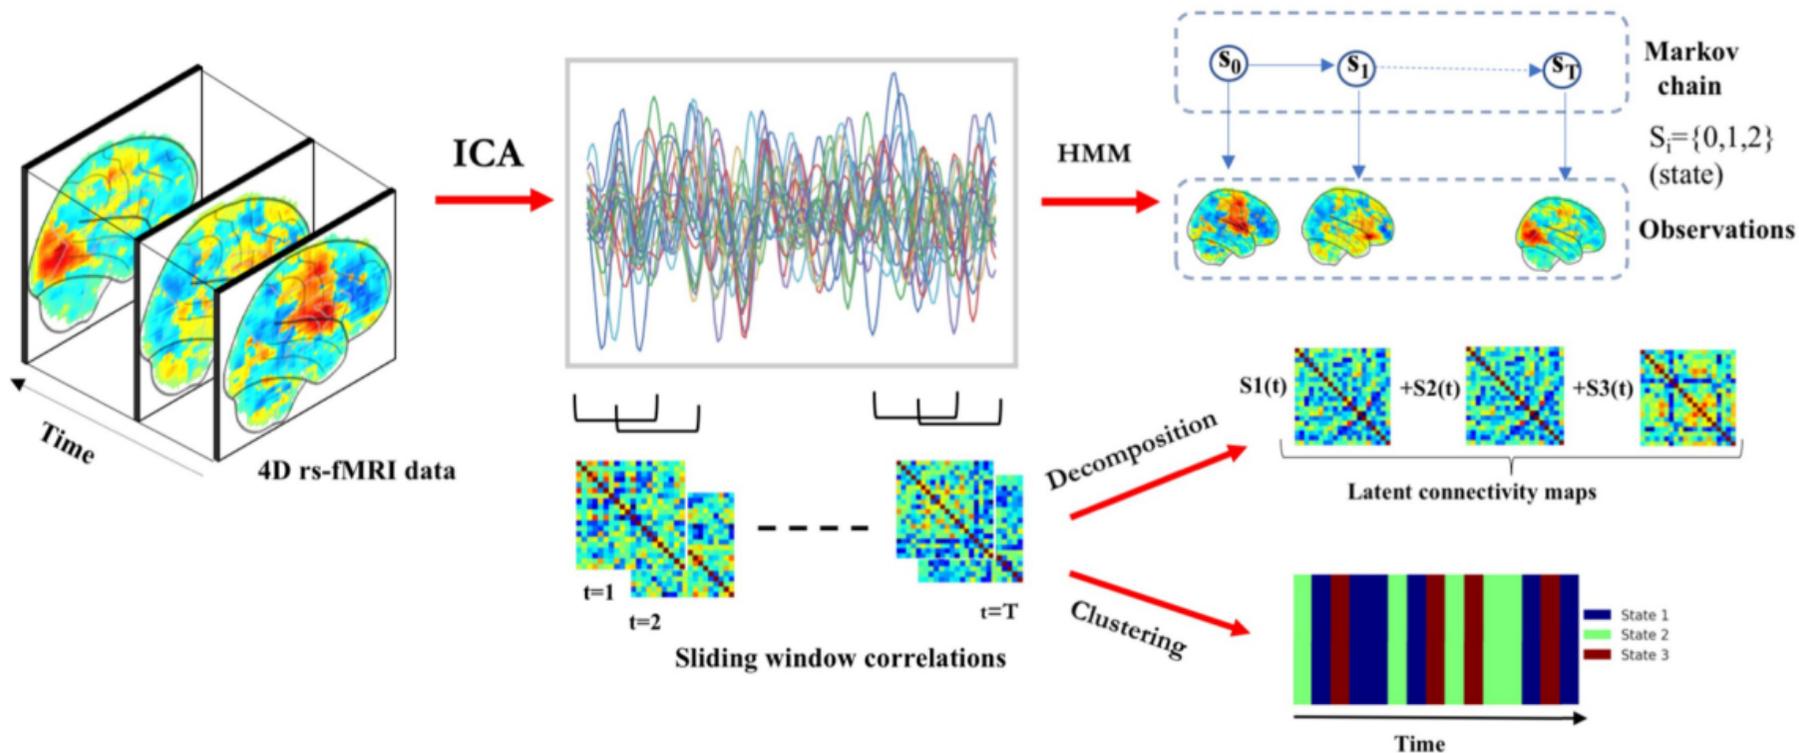


Figure 6

III. Disentangling latent factors of inter-subject FC variation

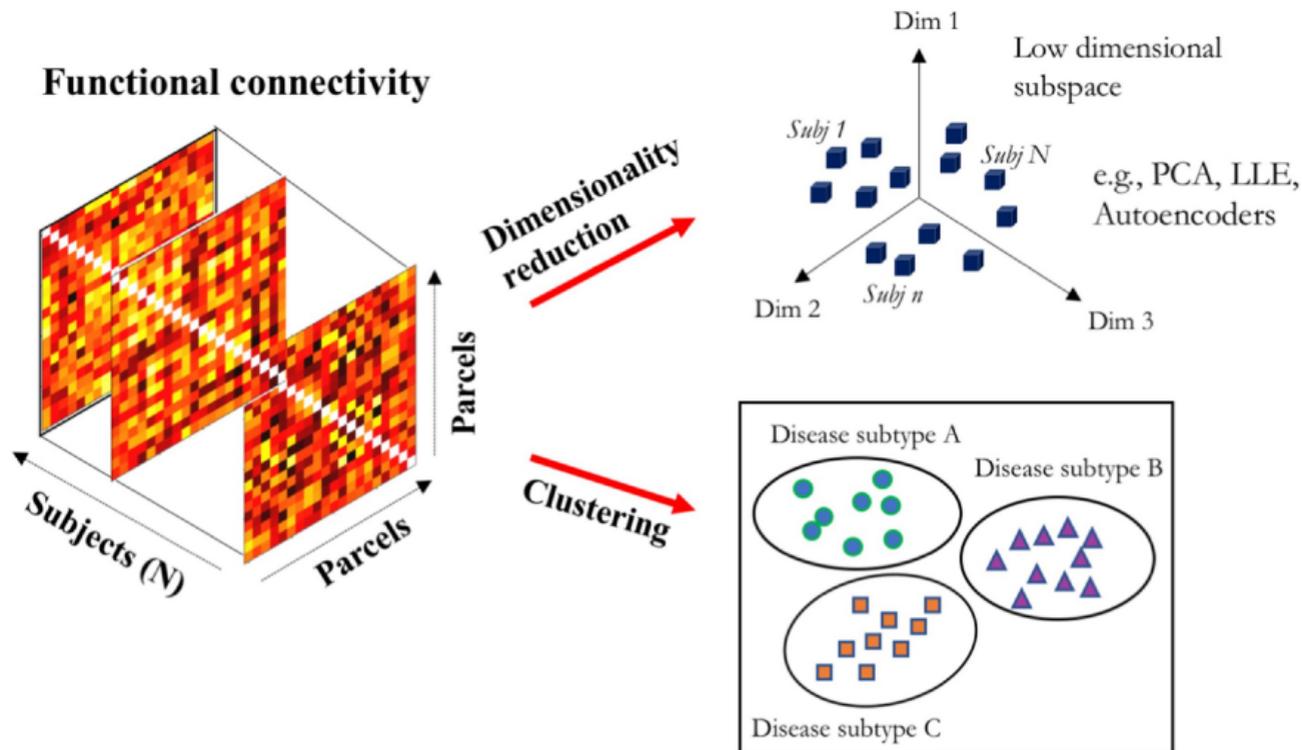


Figure 7

Preprocessed 4D
rs-fMRI data

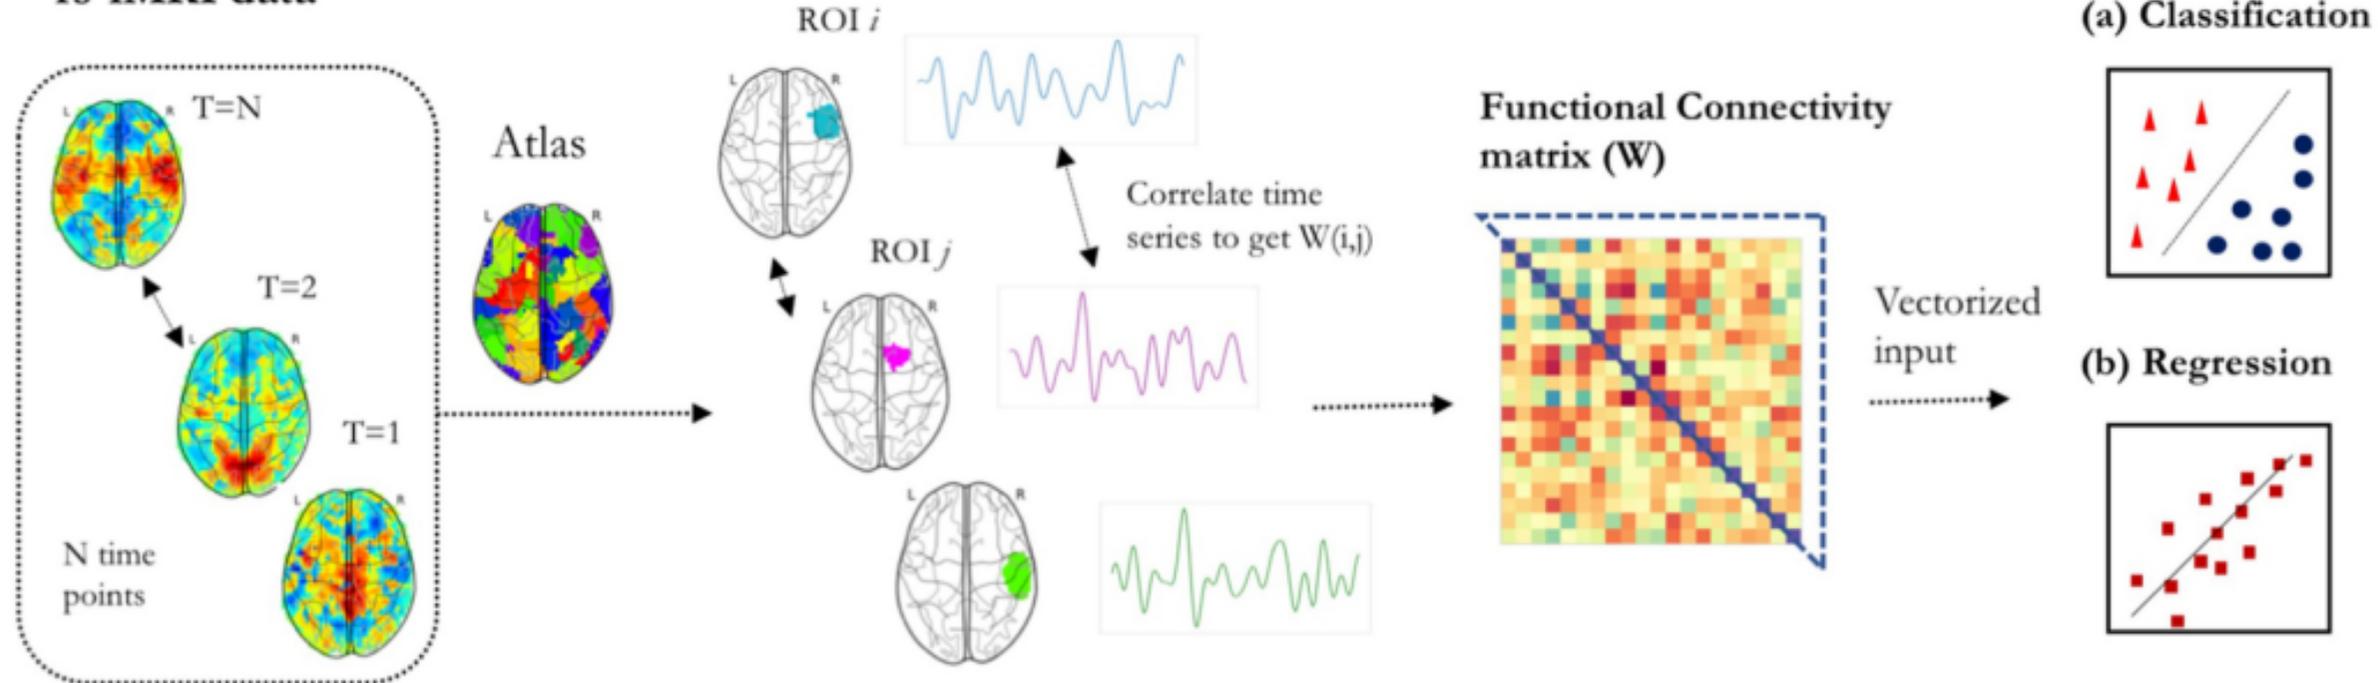


Figure 8

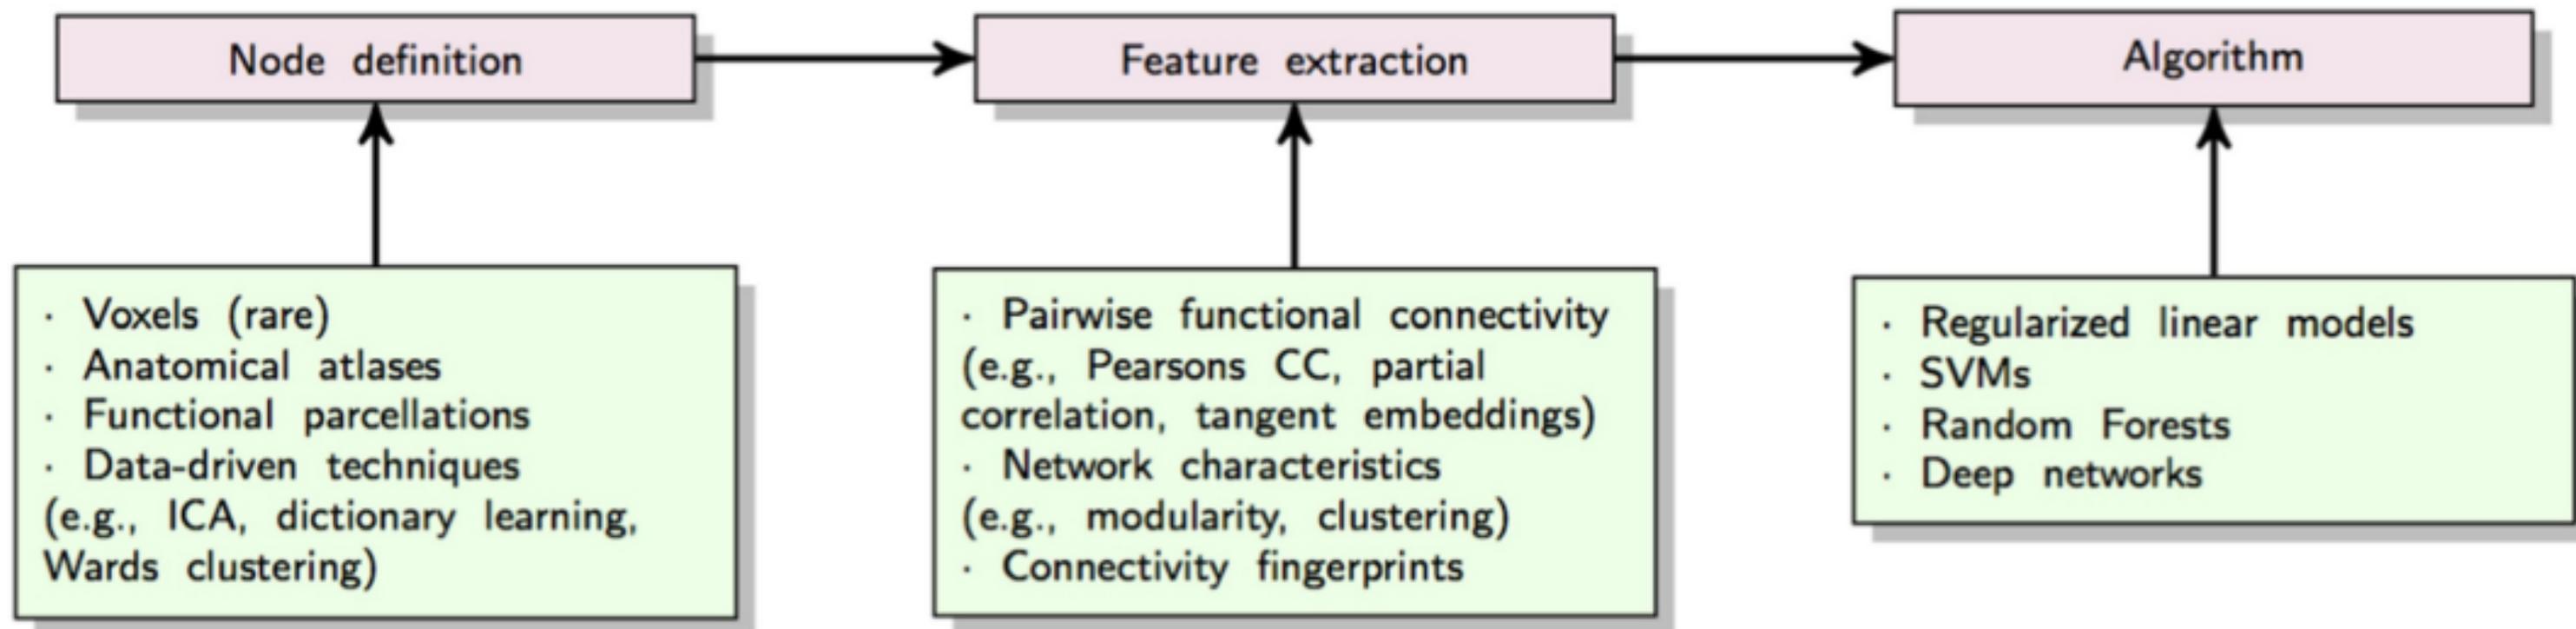


Figure 9

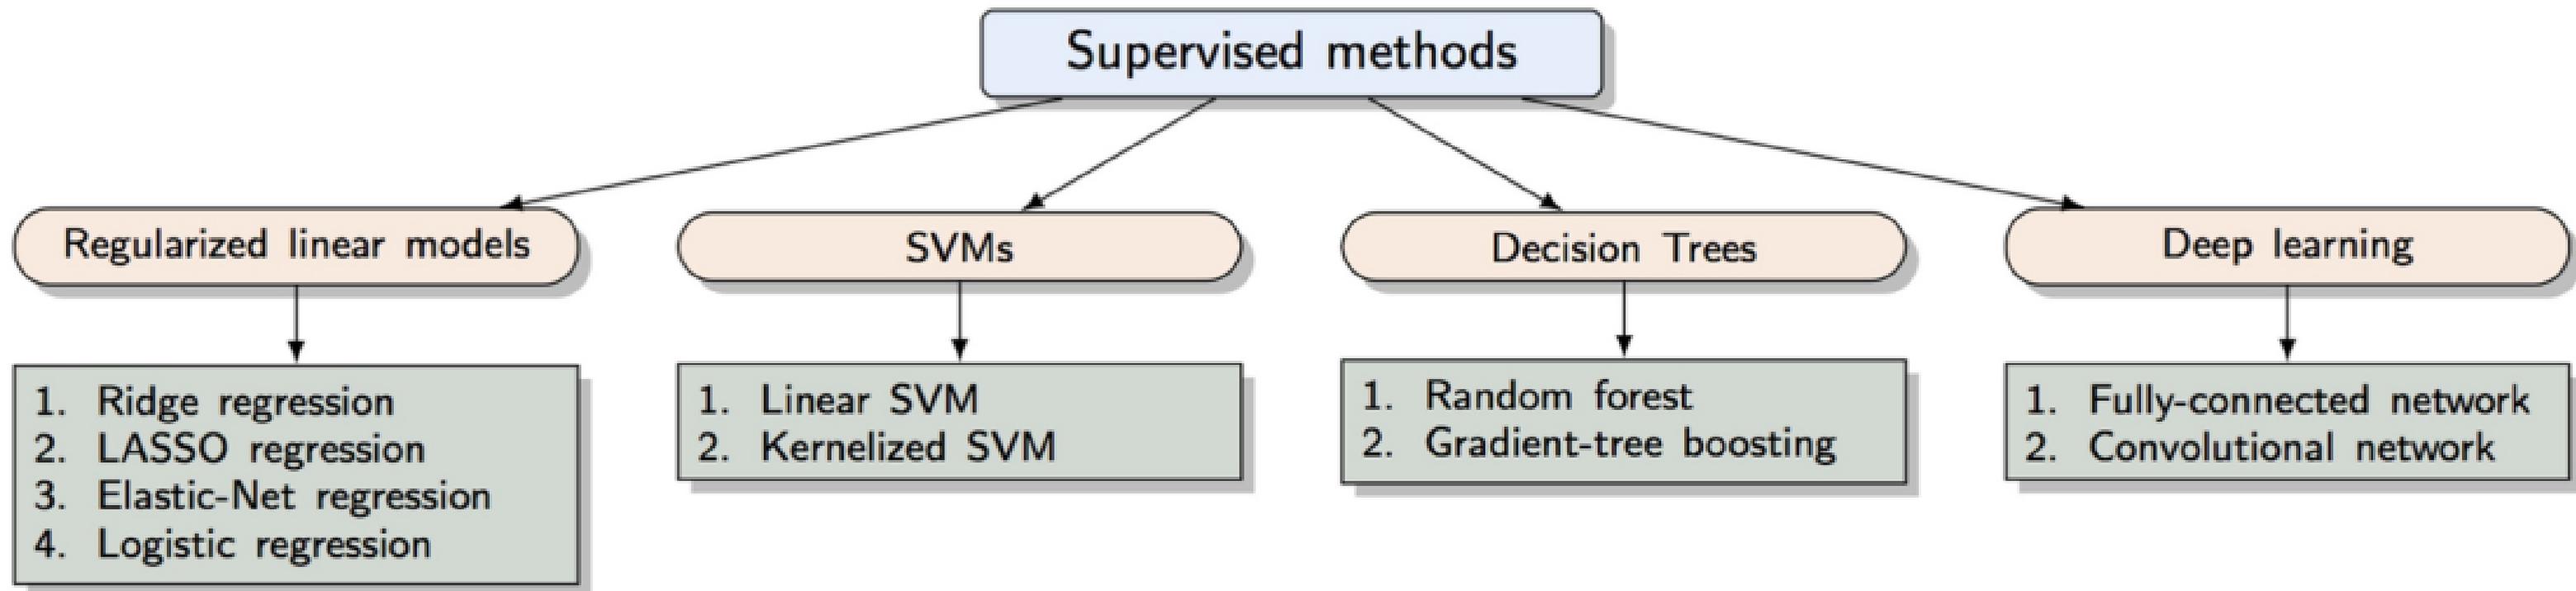


Figure 10