



HAL
open science

Fine-grain atlases of functional modes for fMRI analysis

Kamalaker Dadi, Gaël Varoquaux, Antonia Machlouzarides-Shalit, Krzysztof J. Gorgolewski, Demian Wassermann, Bertrand Thirion, Arthur Mensch

► **To cite this version:**

Kamalaker Dadi, Gaël Varoquaux, Antonia Machlouzarides-Shalit, Krzysztof J. Gorgolewski, Demian Wassermann, et al.. Fine-grain atlases of functional modes for fMRI analysis. *NeuroImage*, 2020, 221, pp.117126. hal-02904869

HAL Id: hal-02904869

<https://inria.hal.science/hal-02904869>

Submitted on 22 Jul 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fine-grain atlases of functional modes for fMRI analysis

Kamalaker Dadi^a, Gaël Varoquaux^a, Antonia Machlouzarides-Shalit^a, Krzysztof J. Gorgolewski^c, Demian Wassermann^a, Bertrand Thirion^a, Arthur Mensch^{a,b}

^aInria, CEA, Université Paris-Saclay, Palaiseau, 91120, France

^bENS, DMA, 45 rue d'Ulm, 75005 Paris

^cDepartment of Psychology, Stanford University, California, USA

Abstract

Population imaging markedly increased the size of functional-imaging datasets, shedding new light on the neural basis of inter-individual differences. Analyzing these large data entails new scalability challenges, computational and statistical. For this reason, brain images are typically summarized in a few signals, for instance reducing voxel-level measures with brain atlases or functional modes. A good choice of the corresponding brain networks is important, as most data analyses start from these reduced signals. We contribute finely-resolved atlases of functional modes, comprising from 64 to 1024 networks. These dictionaries of functional modes (DiFuMo) are trained on millions of fMRI functional brain volumes of total size 2.4TB, spanned over 27 studies and many research groups. We demonstrate the benefits of extracting reduced signals on our fine-grain atlases for many classic functional data analysis pipelines: stimuli decoding from 12,334 brain responses, standard GLM analysis of fMRI across sessions and individuals, extraction of resting-state functional-connectomes biomarkers for 2,500 individuals, data compression and meta-analysis over more than 15,000 statistical maps. In each of these analysis scenarii, we compare the performance of our functional atlases with that of other popular references, and to a simple voxel-level analysis. Results highlight the importance of using high-dimensional “soft” functional atlases, to represent and analyse brain activity while capturing its functional gradients. Analyses on high-dimensional modes achieve similar statistical performance as at the voxel level, but with much reduced computational cost and higher interpretability. In addition to making them available, we provide meaningful names for these modes, based on their anatomical location. It will facilitate reporting of results.

Keywords: Brain imaging atlases; Functional networks; Functional parcellations; Multi-resolution;

1. Introduction

Population imaging has been bringing in terabytes of high-resolution functional brain images, uncovering the neural basis of individual differences (Elliott et al., 2008). While these great volumes of data enable fitting richer statistical models, they also entail massive data storage (Poldrack et al., 2013; Gorgolewski et al., 2017) and challenging high-dimensional data analysis. A popular approach to facilitate data handling is to work with image-derived phenotypes (IDPs), i.e. low-dimensional signals that summarize the information in the images while keeping meaningful representations of the brain (Miller et al., 2016).

While brain atlases originated in characterizing the brain’s microstructure (Brodman, 1909), today they are widely used to study functional connectomes (Sporns et al., 2005; Varoquaux and Craddock, 2013) and for data reduction in functional imaging (Thirion et al., 2006; Craddock et al., 2012). For these applications, the choice of brain regions conditions the signal captured in the data analysis. To define regions well suited to brain-imaging endeavors, there is great progress in building atlases from the neuroimaging data itself (Eickhoff et al., 2018). Yet, most functional atlases describe the brain as parcellations,

locally-uniform functional units, and thus do not represent well functional gradients (Huntenburg et al., 2018).

For functional imaging, brain structures delineated by an atlas should capture the main features of the functional signal, e.g. the functional networks (Smith et al., 2011). In a nutshell, there are two approaches to define well-suited structures. These can strive to select *homogenous* neural populations, typically via clustering approaches (Goutte et al., 1999; Bellec et al., 2010; Craddock et al., 2012; Thirion et al., 2014; Schaefer et al., 2017). They can also be defined via continuous *modes* that map intrinsic brain functional networks (Damoiseaux et al., 2006; Varoquaux et al., 2011; Harrison et al., 2015).

These functional modes have been shown to capture well functional connectivity, with techniques such as Independent Component Analysis (Kiviniemi et al., 2009; Pervaiz et al., 2019) or sparse dictionary learning (Mensch et al., 2016b; Dadi et al., 2019).

High-resolution atlases can give a fine-grained division of the brain and capture more functionally-specific regions and rich descriptions of brain activity (Schaefer et al., 2017). Yet, there is to date no highly-resolved set of “soft” functional modes available, presumably because increas-

ing the dimensionality raises significant computational and statistical challenges (Mensch et al., 2016a; Pervaiz et al., 2019). “Soft” regions take continuous non-negative values, in contrast with hard parcellation atlases. In this paper, we address this need with high-order dictionaries of functional modes (DiFuMo) extracted at a large scale both in terms of data size (3 million volumes of total data size 2.4TB) and resolution (up to 1024 modes). For this, we leverage the wealth of openly-available functional images (Poldrack et al., 2013) and efficient dictionary-learning algorithms to fit on large data. This is unlike ICA which is hard to use for a high number of modes (Pervaiz et al., 2019).

Contributions. We provide Dictionaries of Functional Modes¹ “*DiFuMo*” that can serve as atlases to extract functional signals, e.g. provide IDPs, with different dimensionalities (64, 128, 256, 512, and 1024). These modes are optimized to represent BOLD data well, over a wide range of experimental conditions. They are more finely-resolved than existing brain decompositions with continuous networks. By providing validated fine functional atlases, our goal is to streamline fMRI analysis with reduced representations, to facilitate large-cohort and inter-studies work. Through thorough benchmarking over classic data analysis tasks, we show that these modes gives IDPs that ground better analysis of functional images. Finally, we provide a meaningful label to each mode, summarizing its anatomical location, to facilitate reporting of results.

2. Methods: data-driven fine-grain functional modes

We describe in this section the models and methods underlying our definition of brain structures to extract IDPs.

2.1. Context: Image Derived Phenotypes

While analysis of brain images has been pioneered at the voxel level (Friston et al., 1995), image-derived phenotypes (IDP) are increasingly used in the context of population imaging. Trading voxel-level signals for IDPs has several motivations. First and foremost, it greatly facilitates the analysis on large cohorts: the data are smaller, easier to share, requiring less disk storage, computer memory, and computing power to analyze. It can also come with statistical benefits. For instance, in standard analysis of task responses, e.g. in mass-univariate brain mapping, the statistical power of hypothesis test at the voxel level is limited by multiple comparisons (Friston et al., 1995), while working at the level of IDPs mitigates this problem (Thirion et al., 2006). For predictive modeling, e.g. in multi-variate decoding (Mourão-Miranda et al., 2005), the high-dimensionality of the signals is a challenge to learning models that generalize well—a phenomenon known as the

curse of dimensionality in machine learning (Hastie et al., 2009). Finally, for functional connectomes, working at voxel-level is computationally and statistically intractable as it entails modeling billions of connections. The standard approach is therefore to average signals on regions or networks (Varoquaux and Craddock, 2013).

Functional neuroimaging is currently largely dependent on neuroanatomy for mapping function to structure (Destrieux et al., 2010; Devlin and Poldrack, 2007). Some anatomical structures support well a direct mapping to specific functions (Brett et al., 2002; Rademacher et al., 1993), e.g. the primary visual areas. Yet other functional units are not simply defined from anatomical features, for instance in high-level regions such as the default mode, which is defined from functional data (Leech et al., 2011; Greicius et al., 2003).

2.2. Methods for data-driven functional atlases

Compared to anatomical atlases, defining regions from the functional signal can lead to a better explanation of behavioral outcomes (Dadi et al., 2019), as they capture the functional structure of the brain. Clustering of fMRI timeseries has been heavily used to define brain parcellations (Goutte et al., 1999), or for data reduction in predictive models (Michel et al., 2012). Reference functional brain parcellations have been defined with various clustering algorithms on resting-state fMRI (Bellec et al., 2010; Yeo et al., 2011; Craddock et al., 2012). Another class of approaches seeks modes of brain activity, decomposing the signal as a product of spatial maps and corresponding time-series (Figure 1). The most popular model in neuroimaging is independent component analysis (ICA, Hyvärinen and Oja, 2000), which optimizes spatial independence between extracted maps. It has been extensively used to define resting-state networks (Kiviniemi et al., 2003; Beckmann et al., 2005; Calhoun et al., 2001) and implicitly outlines soft parcellations of the brain at high order (Kiviniemi et al., 2009; Varoquaux et al., 2010b). ICA-defined networks are used to extract the official IDPs of UK BioBank, the largest brain-imaging cohort to date; these have been shown to relate to behavior (Miller et al., 2016).




Figure 1: **Linear decomposition model of fMRI time-series for estimating brain networks:** The fMRI time series \mathbf{X} are factorized into a product of two matrices, \mathbf{D} which contain spatial modes and \mathbf{A} temporal loadings of each mode. p - number of features, n - number of volumes in fMRI image, k - number of dictionaries.

¹<https://parietal-inria.github.io/DiFuMo>

We rely on another decomposition model, dictionary learning (Olshausen and Field, 1997), that enforces sparsity and non-negativity instead of independence on the spatial maps. While less popular than ICA in neuroimaging, sparsity brings the benefit of segmenting well functional regions on a zeroed-out background (Lee et al., 2010; Varoquaux et al., 2011). For our purposes, an important aspect of sparse models is that they have computationally-scalable formulations even with high model order and on large datasets (Mensch et al., 2016a, 2018). Functional modes defined from sparse dictionary learning have been used to predict Autism Spectrum Disorder (Abraham et al., 2017), or mental processes (Mensch et al., 2017).

Rest and task fMRI. Most functional brain atlases have been extracted from rest fMRI (Bellec et al., 2010; Power et al., 2011; Craddock et al., 2012; Yeo et al., 2011; Miller et al., 2016; Schaefer et al., 2017). Brain networks can also be extracted from task fMRI data (Calhoun et al., 2008; Lee et al., 2010), and segment a similar intrinsic large-scale structure (Smith et al., 2009). In our work, we build functional modes from datasets with different experimental conditions, including task and rest. Our goal is to be as general as possible and capture information from different protocols. Indeed, defining networks on task fMRI can help representing these brain images and predicting the corresponding psychological conditions (Duff et al., 2012). We note that appending rest data to the already large task fMRI corpus gives marginal improvement (Figure A9).

2.3. DiFuMo extraction: model and data

We consider BOLD time-series from fMRI volumes, resampled and registered to the MNI template. After temporal concatenation, those form a large matrix $\mathbf{X} \in \mathbb{R}^{p \times n}$, where p is the number of voxels of the images (around $2 \cdot 10^5$), and n is the number of brain images, of the order of 10^6 in the following. To extract DiFuMos, each brain volume is modeled as the linear combination of k spatial functional networks, assembled in a dictionary matrix $\mathbf{D} \in \mathbb{R}^{p \times k}$. We thus assume that \mathbf{X} approximately factorizes as $\mathbf{D}\mathbf{A}$, where the matrix $\mathbf{A} \in \mathbb{R}^{k \times n}$ holds in every column the loadings α_i necessary to reconstruct the brain image \mathbf{x}_i from the networks \mathbf{D} . The *dictionary* \mathbf{D} is to be *learned* from data. For this, we rely on *Stochastic Online Matrix Factorization*² (Mensch et al., 2018, SOMF), that is computationally tractable for matrices large in both directions, as with high-resolution large-scale fMRI data. SOMF solves the constrained ℓ_2 reconstruction problem

$$\min_{\substack{\mathbf{D} \in \mathbb{R}^{p \times k}, \mathbf{A} \in \mathbb{R}^{k \times n} \\ \mathbf{D} \geq 0, \forall j \in [k], \|\mathbf{d}_j\|_1 \leq 1}} \|\mathbf{X} - \mathbf{D}\mathbf{A}\|_F^2 + \lambda \|\mathbf{A}\|_F^2, \quad (1)$$

where λ is a regularization parameter that controls the sparsity of the dictionary \mathbf{D} , via the ℓ_1 and positivity con-

straints.³ Encouraging sparsity in spatial maps is key to obtaining well-localized maps that outline few brain regions. Yet these may still have some overlap, to capture distributed networks with shared anatomical territories. The parameter λ is chosen so that the union of all maps approximately covers the whole brain with minimum overlap between maps. We provide an exhaustive description of the methodological choices made for extracting DiFuMos in Appendix H. In particular, we provide more details on selecting the optimum λ , and on the brain coverage of the DiFuMo atlases (see Table A4). We note that we do not enforce any spatial prior on brain regions, in contrast with existing work based on matrix factorization (Abraham et al., 2014a; Dohmatob et al., 2016). This allows to retain a reasonable computational cost. Dohmatob et al. (2016) have shown the reduced importance of spatial priors when extracting atlases from large datasets.

Input fMRI data. We build the input data matrix \mathbf{X} with BOLD time-series from 25 different task-based fMRI studies and 2 resting state studies, adding up to 2192 functional MRI recording sessions. We gather data from OpenNeuro (Gorgolewski et al., 2017) –Table A5 lists the corresponding studies while Table A6 gives their data-acquisition parameters.

We use *fMRIprep* (Esteban et al., 2019) for minimal preprocessing: brain extraction giving as a reference to correct for head-motion (Jenkinson et al., 2002), and co-registration to anatomy (Greve and Fischl, 2009). All the fMRI images are transformed to MNI template space. We then use *MRIQC* (Esteban et al., 2017) for quality control.

Multi-dimensional DiFuMo atlases. We estimate dictionaries of dimensionality $k \in \{64, 128, 256, 512, 1024\}$. This is useful as the optimal dimensionality for extracting IDPs often depends on the downstream data analysis task. The obtained functional modes segment well-localized regions, as illustrated in Figure 2.

2.4. Extracting signal on functional modes

The functional modes take continuous values (we refer to them as *soft*) and can have some overlap –though in practice this overlap is small. As a consequence, signal extraction calls for more than averaging on regions. The natural formulation is that the extracted signals (the IDPs) should best approximate the brain image $\mathbf{x} \in \mathbb{R}^p$ as a linear combination $\alpha \in \mathbb{R}^k$ of the set of modes in the dictionary $\mathbf{D} \in \mathbb{R}^{p \times k}$. This is solved by linear regression:

$$\alpha = \operatorname{argmin}_{\alpha \in \mathbb{R}^k} \|\mathbf{x} - \mathbf{D}\alpha\|_2^2, \quad \text{i.e.} \quad \alpha = \mathbf{D}^\dagger \mathbf{x}, \quad (2)$$

³Eq. (1) admits multiple local solutions: the obtained atlases may slightly vary across training runs, although similar regions are recovered, and all atlas perform similarly. Mensch et al. (2016b) further discuss reproducibility of matrix factorization across runs.

²Available at: <https://arthurmensch.github.io/mod1/>




Figure 2: **Schema of DiFuMo atlases and their usage in typical fMRI analyses.** DiFuMo atlases are extracted from a massive concatenation of BOLD time-series across fMRI studies, using a sparsity inducing matrix factorization algorithm. We compute the DiFuMo atlases at different resolutions, up to 1024 components. We assess our atlases in 4 benchmarks that measure suitability to classic fMRI analyses. Those are performed on reduced and non-reduced data, with different atlas sizes and a comparison between atlases. The easiest way to view and download DiFuMo atlases is via the online interactive visualizations: parietal-inria.github.io/DiFuMo.

where $D^\dagger = (D^T D)^{-1} D^T \in \mathbb{R}^{k \times p}$ is the pseudo-inverse of D . For atlases composed of non-overlapping regions, such as classic brain parcellations—e.g. BASC (Bellec et al., 2010) or normalized cuts (Craddock et al., 2012)—linear regression simply amounts to averaging the images values in every cluster of D . For overlapping modes as the ones of DiFuMo or the ICA maps used in UKBB (Miller et al., 2016), the linear regression formulation caters for the overlap and softness of the regions.

2.5. Region names: relation to anatomical structures

Relating IDPs to known brain structures facilitates interpretation and discussion of results. Though the DiFuMo atlases are defined from functional signal, we choose to reference their regions by their anatomical location, as it is a common and well-accepted terminology in neuroscience. For each resolution, we match the modes with regions in references of brain structure: the Harvard-Oxford atlas (Desikan et al., 2006), Destrieux atlas (Destrieux et al., 2010), the MIST atlas (Urchs et al., 2019), Johns Hopkins University (JHU) atlas (Hua et al., 2008), and the Dierdrichsen cerebellum atlas (Dierdrichsen et al., 2009).

We name each mode from the anatomical structure that it most overlaps with. When the overlap was weak, a trained neuroanatomist (AMS) looked up the structure in standard classic anatomy references (Henri, 1999; Schmahmann et al., 1999; Rademacher et al., 1992; Ono et al., 1990; Catani and de Schotten, 2012). Appendix F gives more details on the naming of the brain areas.

3. Brain-image analysis on functional modes

We use the reduced representations (IDPs) introduced above for various functional-imaging analytic tasks: standard mass-univariate analysis of brain responses (§3.2); decoding mental processes from brain activity (§3.3); predicting phenotypes from functional connectomes (§3.4); finally, we measure the quality of signal reconstruction after dimension reduction, illustrated on meta-analyses (§3.5).

3.1. Benchmarking several functional atlases

To gauge the usefulness of the extracted IDPs, we compare each analysis pipeline across several functional atlases: DiFuMo and reference atlases are used to compute

Name	Dimensionality	# subj.	Data type	Soft	Extraction method	Reference
BASC	64, 122, 197, 325, 444	200	rest	No	Hierarchical clustering	Bellec et al. (2010)
MIST ATOM ^a	1095	200	rest	No	Region growing	Urchs et al. (2019)
Craddock	200, 400	41	rest	No	Spectral clustering	Craddock et al. (2012)
FIND ^b	90, 499	15	rest	Yes	ICA; Ward clustering	Shirer et al. (2012); Altmann et al. (2015)
Gordon	333	120	rest	No	Local-gradient approach	Gordon et al. (2014)
UKBB ICA	21, 55	4100	rest	Yes	Selected ICA components ^c	Miller et al. (2016)
Schaefer	100, 200, 300, 400, 500, 600, 800, 1000	1489	rest	No	Gradient-weighted Markov Random Field (gwMRF)	Schaefer et al. (2017)
DiFuMo^d	64, 128, 256, 512, 1024	2192	task, rest	Yes	Sparse dictionary learning	This paper

^a https://figshare.com/articles/_/5633638

^b https://findlab.stanford.edu/functional_ROIs.html

^c <https://www.fmrib.ox.ac.uk/ukbiobank/>

^d <https://parietal-inria.github.io/DiFuMo>

Table 1: Functional atlases that we benchmark; they define IDPs for brain-images analysis

functional IDPs. We use the same signal-extraction function (2), but vary the spatial components \mathbf{D} . As a baseline, we also perform the voxel-level analyses, though it entail significantly larger computational costs.

We consider other functional atlases that are multi-resolutions, accessible to download, and volumetric (Table 1): *ICA maps* with $k \in \{21, 55\}$ components, extracted on large-scale rs-fMRI from UKBB (Miller et al., 2016); *bootstrap analysis of stable clusters (BASC)* built with hierarchical clustering on rs-fMRI, with various number of clusters (Bellec et al., 2010); *Multiresolution Intrinsic Segmentation Template (MIST)* functional atoms, identified using the region growing procedure from (Bellec et al., 2006), with $k = 1095$ regions (Urchs et al., 2019); spatially-constrained clustering on rs-fMRI, with $k \in \{200, 400\}$ clusters (Craddock et al., 2012); $k = 333$ cortical areas derived from rs-fMRI using a local gradient approach (Gordon et al., 2014); $k \in \{90, 499\}$ functional regions covering cortical and subcortical gray matter with ICA and Ward clustering (Shirer et al. (2012), Altmann et al. (2015)); and brain parcellations derived with gradient-weighted Markov Random Field, with resolutions similar to ours (Schaefer et al., 2017, k up to 1000).

3.2. Mapping brain response: standard task-fMRI analysis

Standard analysis in task fMRI relates psychological manipulations to brain activity separately for each voxel or region. It models the BOLD signal as a linear combination of experimental conditions—the General Linear Model (GLM, Friston et al., 1995). The BOLD signal forms a matrix $\mathbf{Y} \in \mathbb{R}^{n \times p}$, where p is the number of voxels. With data reduction, we use as input the reduced signal $\mathbf{Y}_{\text{red}} = \mathbf{Y}_{\text{voxel}}(\mathbf{D}^\dagger)^\top \in \mathbb{R}^{n \times k}$ (Equation 2). The GLM models \mathbf{Y} or \mathbf{Y}_{red} as $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$ where $\mathbf{X} \in \mathbb{R}^{n \times q}$ is the design matrix formed by q temporal regressors of interest or nuisance and $\boldsymbol{\varepsilon}$ is noise (Friston et al., 1998). In our experiments, we use the *Nistats* library⁴, and compute a common brain mask from data covering brain areas with ICBM grey matter mask.

⁴<https://nistats.github.io/>

With reduced input \mathbf{Y}_{red} , we obtain one signal per region, as $\boldsymbol{\beta} \in \mathbb{R}^{q \times k}$. The full β -maps can then be reconstructed by setting $\boldsymbol{\beta}_{\text{rec}} = \boldsymbol{\beta}\mathbf{D}^\top \in \mathbb{R}^{q \times p}$. We transform the reconstructed β -maps into z -maps $\mathbf{z} \in \mathbb{R}^{q \times p}$ using base contrasts, before thresholding them with Benjamini and Hochberg (1995) FDR correction for multiple comparisons. We then compare the z -maps obtained using voxels as input, and z -maps using reduced input and reconstructed β -maps, using the Dice (1945) similarity coefficient. We also perform an *intra-subject* analysis detailed in Appendix D.

Data. We consider the Rapid-Serial-Visual-Presentation (RSVP) language task of Individual Brain Charting (IBC) (see Pinho et al., 2018, for experimental protocol and pre-processing). We model six experimental conditions: complex meaningful sentences, simple meaningful sentences, jabberwocky, list of words, lists of pseudowords, consonant strings. β -maps are estimated for each subject using a fixed-effect model over 3 out of the 6 subject’s sessions. We randomly select 3 sessions 10 times to estimate the variance of the Dice index across sessions. As a baseline, we evaluate the mean and variance of the Dice index across z -maps when varying the sessions used in voxel-level GLM.

3.3. Decoding experimental stimuli from brain responses

Decoding predicts psychological conditions from task-related z -maps (Haynes and Rees, 2006). The validity of a decoding model is evaluated on left-out data (following Varoquaux et al., 2017), e.g. left-out subjects for inter-subject decoding (Poldrack et al., 2009). We use linear decoding models: ridge regression for continuous target and Support Vector Machine (SVC, Hastie et al., 2009) for classification. For each study, we separate sessions (for intra-subject decoding) or subjects (for inter-subject decoding) into randomly-chosen train and test folds (20 folds with 30% test size), and measure the test accuracy. We compare the performance of predictive models using voxel-level z -maps or data reduced with functional atlases.

Data. We use 6 open-access task fMRI studies. We perform *inter-subject* decoding in the **emotional** and sensitivity to **pain** experiences from Chang et al. (2015), and

in three studies from HCP900 (Van Essen et al., 2012): **working memory**, **gambling** (Delgado et al., 2000), and **relational processing** (Smith et al., 2007). We perform *intra-subject* decoding using the several sessions of **left** and **right** button press responses in IBC (ARCHI protocol, Pinel et al., 2007). The unthresholded z -maps used in the decoding pipeline are either obtained from Neurovault (Gorgolewski et al., 2015), or computed with the GLM following §3.2. Details are reported in Appendix B.1.

3.4. Predicting phenotype from functional connectomes

Resting-state fMRI can be used to predict phenotypic traits (Richiardi et al., 2010). For this, each subject is represented by a functional connectivity matrix capturing the correlation between brain signals at various locations. Our functional-connectome prediction pipeline comprises three steps: **1)** we extract a reduced representation of the BOLD signal⁵, projecting voxel-level data onto a functional atlas as in §3.2; **2)** we compute a *functional connectome* from the reduced BOLD signals; **3)** we use it as input to a linear model. We compute a connectome from activations with the Ledoit and Wolf (2004) covariance estimator as Varoquaux and Craddock (2013); Brier et al. (2015). We then derive single-subject features from covariance matrices using their tangent space parametrization (Varoquaux et al., 2010a; Barachant et al., 2013; Pervais et al., 2019), as advocated by the benchmarks of Dadi et al. (2019). Those are used to fit an ℓ_2 -penalized logistic regression for classification and a ridge regression for continuous targets. We assess predictive performance with 20 folds, random splits of subjects in train and test sets, with 25% test size.

Data. We use 7 openly-accessible datasets with diverse phenotypic targets, as summarized in Table A3. We predict diagnostic status for Alzheimer’s disease on **ADNI** (Mueller et al., 2005), PTSD on **ADNIDOD**; Autism Spectrum Disorder on **ABIDE** (Di Martino et al., 2014) and schizophrenia on **COBRE** (Calhoun et al., 2012); drug consumption on **ACPI**; IQ measures on **HCP** (Van Essen et al., 2013); and age (with a regression model) in normal aging with **CamCAN** (Taylor et al., 2017).

3.5. Quality of image reconstruction

The signals extracted on a brain atlas can be seen as a compression, or simplification, of the original signal. Indeed, a full image can be reconstructed from these signals. We quantify the signal loss incurred by this reduction. For this, we project a brain map \mathbf{x} (obtained with the ICBM whole brain mask) onto an atlas (solving Eq. (2)), and compute the best reconstruction of \mathbf{x} from the loadings $\boldsymbol{\alpha}$,

⁵We handle site differences by casting all data in the MNI template space. Abraham et al. (2017) shows the robustness to site difference of predictors based on functional connectomes.

namely $\hat{\mathbf{x}} = \mathbf{D}\boldsymbol{\alpha} \in \mathbb{R}^p$. We compare original and reconstructed images through the R^2 coefficient,

$$R^2(\mathbf{x}, \hat{\mathbf{x}}) = 1 - \frac{\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2}{\|\mathbf{x} - \bar{\mathbf{x}}\|_2^2}, \quad (3)$$

where $\bar{\mathbf{x}} \in \mathbb{R}$ is the spatial mean of map \mathbf{x} . The R^2 coefficient is averaged across all images. Higher R^2 coefficients means that the reduced signals (IDPs) explain more variance of the original images, where $R^2 = 1$ corresponds to no signal loss. The larger the number of signals used, the easier it is to explain variance; it is therefore interesting to compare this measure across atlases with similar number of components. For a fixed number of component in the DiFuMo atlases, R^2 increases with brain coverage.

Data. We use **NeuroVault** (Gorgolewski et al., 2015), the largest public database of statistical maps. To avoid circularity, we exclude maps derived from the studies used to extract the DiFuMo atlases, along with maps that fail semi-automated quality inspection (filtering out thresholded or non-statistical maps), resulting in **15,542** maps.

Meta-analysis of contrasts maps. Ideally, the extracted IDPs should allow to compute meta-analytical summaries of brain activity maps. In this setting, a single map, corresponding to a certain cognitive concept, is computed from many z -maps across different studies, associated to conditions that involve this cognitive concept. We compare the summaries obtained at voxel-level, i.e. averaging the maps $\{\mathbf{x}\}$, with the ones obtained using reconstructed images, i.e. averaging the maps $\{\hat{\mathbf{x}}\}$ used in Eq. (3). We use maps from our curated subset of NeuroVault annotated with terms *motor*, *language* and *face recognition*.

4. Results: comparing atlases for analyses

We report benchmarking results on the analytic tasks listed in the previous section.

4.1. Brain mapping: standard task-fMRI analysis

Figure 3 reports the results of standard analysis of task fMRI (GLM), comparing analysis at the voxel-level with analyses on signals extracted from functional atlases. Best correspondence is obtained at highest dimensionality, as the regions are finer. Notably, analysis with DiFuMo of dimensionality 1024 is markedly closer to voxel-level analysis than using the best-performing alternatives, i.e. the 1000-dimensional Schaefer parcellation and the 1095 MIST atoms. In addition, the Dice index relative to the voxel-level gold standard is comparable to the Dice index between runs of voxel-level GLM estimated across folds. We note that using soft functional modes from only 55 ICA components shows excellent results, comparable to those obtained using the 1000 components Schaefer atlas. This stresses the benefit of *continuous functional modes* for the analysis of task responses. Overall, standard task-fMRI




Figure 3: **Overlap between GLM maps obtained with functional atlases and voxel-level analysis.** **Top:** The overlap is measured with the Dice similarity coefficient. The black line gives a baseline the mean overlap between voxel-level contrast maps over several random selections of sessions per subject. The figure gives Dice similarity scores between the GLM maps computed with signals extracted on functional atlases and at the voxel-level, after reconstruction of full z-maps and voxel-level thresholding with FDR control. The best similarity is achieved for highest dimensionality, though 1024-dimensional DiFuMo atlas largely dominates 1000-dimensional Schaefer parcellation and 1095 MIST. Each point is the mean and the error bar denotes the standard deviation over contrast maps. **Bottom:** The activity maps encoded on 1024-dimensional space capture the same information as voxel-level analysis, while being smoother. Figure A7 shows activity maps for all atlases.

analysis on signals derived from 512 or 1024-dimensional DiFuMo gives results close to the voxel-level gold standard (Figure 3 shows that the maps are also qualitatively similar).⁶ Figure A6 shows similar trends while comparing intra-subject explained-variance maps, both qualitatively and quantitatively. Dimension reduction has the additional benefit of alleviating the burden of correcting for multiple comparisons.

4.2. Decoding mental state from brain responses

Figure 4 shows the impact on decoding performance of reducing signals with various functional atlases. It reports the performance relative to the median across methods for

⁶We note that using too coarse atlases may fail at detecting statistically significant voxel activations, yielding a Dice index of 0.




Figure 4: **Impact of the choice of atlas on decoding performance.** Each point gives the relative prediction score, over 6 different task-fMRI experiments. The thick lines give the median relative score per atlas. The baseline (black) is the relative score. High-order resolutions increase prediction accuracy. Using high-order DiFuMo ($k = 1024$) and Schaefer parcellations ($k = 1000$) gives the best performance and, on average, outperforms voxel-level prediction. Appendix B.2 reports absolute prediction scores for each task.




Figure 5: **Decoding maps of the working memory task, face versus rest,** showed for Voxel-level analysis, DiFuMo, and Schaefer. The maps are highly interpretable with high-dimensional soft modes (DiFuMo 1024) compared to voxel-level analysis. Brain areas important in the visual working memory task—fusiform gyrus and lateral occipital cortex—are clearly visible. Figure A4 gives a full view of decoding weights across atlases and resolutions.

each of the 6 tasks. These results clearly show the importance of high-dimensional functional modes for decoding. Indeed, the higher the atlas resolution, the better the predictions. Using DiFuMo $k = 1024$ or Schaefer $k = 1000$ gives the best performance. In addition, as these functional atlases segment sufficiently-fine regions, prediction from the corresponding signals tends to outperform voxel-level prediction. Indeed, applying multivariate models to a larger number of signals with a limited amount of data is more prone to overfitting—data reduction acts here as a welcome regularization. Qualitatively, brain maps containing decoding weights can be reconstructed. With high-dimensional atlases, they are interpretable and capture information similar to voxel-level analysis (Figure 5).




Figure 6: **Impact of the choice of atlas for predictions based on functional connectomes.** Each data point gives the prediction accuracy relative to the median for one of the 7 phenotypic prediction targets, i.e. each point a dataset. The thick line shows the median over the datasets. While the results are noisy, the optimal dimensionality seems to lie around 300 nodes, and the best-performing atlas is DiFuMo $k = 256$, followed by Craddock $k = 400$ and BASC $k = 444$. Figure A5 report absolute results for each prediction problem.

4.3. Predicting traits from functional connectomes

Figure 6 shows the impact of the choice of functional atlas when predicting phenotypes from functional connectomes. We report the relative prediction accuracy for 7 different prediction problems (each composed of a dataset and a target phenotype); the lines give the median across the prediction problems. Here, we do not report a voxel-level baseline, as it requires to compute covariance matrices of dimensions around $100,000 \times 100,000$ and is therefore computationally and statistically intractable. In contrast with previous results, high-resolution atlases do not provide the best performance, likely because the complexity of the statistical models increases with the square of the number of nodes. The best prediction overall is achieved using DiFuMo $k = 256$, followed by Craddock $k = 400$ and BASC $k = 444$ atlases. Different outcomes have different optimal dimensionality, consistently across atlases (Figure A5): $k \sim 150$ for age prediction; $k \sim 300$ for Autism Spectrum Disorder, PTSD, or IQ prediction; and $k \sim 50$ for Alzheimer’s Disease and drug use prediction.

4.4. Fraction of the original signal captured

Figure 7 (left) displays the R^2 scores summarizing the loss of information when data are reduced on an atlas and reconstructed back to full images. Unsurprisingly, reducing the images with lower-order dimensions (atlases with fewer regions) yields a high loss of information across all methods. DiFuMo $k = 1024$ captures 70% of the original voxel-level signal. Qualitatively, the benefits of functional modes can be seen by comparing the meta-analytic maps

related to motor tasks (Figure 7 right)—Figure A8 shows additional meta-analysis on other topics and atlases. The DiFuMo have clear visual benefits over brain discrete parcellations, such as BASC, as they better capture gradients.

5. Discussion

This paper introduces brain-wide soft functional modes, named DiFuMos and made of a few hundreds to a thousand of brain sub-divisions. They are derived from BOLD time-series across many studies to capture well functional images with a small number of signals. In the context of population imaging, these signals are known as image-derived phenotypes (IDP, Miller et al., 2016) and are crucial to easily scale statistical analysis, building a science of inter-individual differences by relating brain signals to behavioral traits (Dubois and Adolphs, 2016). Reducing the dimensionality of the signals not only come with a $1000\times$ gain in storage, but also with $100\times$ computational speed-up for the analysis (Table A1). Even small-scale studies may need functional nodes, e.g. for computing functional connectomes (Zalesky et al., 2010; Varoquaux and Craddock, 2013). There already exist many functional brain atlases; yet DiFuMos have the unique advantage of being both soft and highly resolved. These features are important to capture gradients of functional information.

Grounding better image-derived phenotypes. Signals extracted from a functional atlas should enable good statistical analysis of brain function. We considered quantitative measures for typical neuroimaging analytic scenarios and compared the fitness of extracting signal on DiFuMo with using existing functional brain atlases. The biggest gains in analysis come from increasing the dimensionality of brain sub-divisions, aside for functional connectome studies where an optimal is found around 200 nodes. Choosing the number of nodes then becomes a tradeoff between complexity of the representation and analytic performance. Importantly, the gains in analytic performance continue way beyond the dimensionality typically used for IDPs (e.g. 55 components from Miller et al., 2016). These results extend prior literature emphasizing the importance of high-dimensional parcellations for fMRI (Abou Elseoud et al., 2011; Thirion et al., 2014; Arslan et al., 2017; Sala-Llonch et al., 2019). To foster good analysis, the second most important aspect of a parcellation appears that it be soft, i.e. continuously-valued. For a given dimensionality, soft modes tend to outperform hard parcellations, whether they are derived with ICA or dictionary learning.

Modes well-adapted to the EPI signal. The functional modes are optimized to fit well a large number of EPI images: 2,192 sessions across 27 studies. As a result, they form a division of the brain well adapted to the signal. For instance, they define regions larger in the white matter and in the CSF than in the grey matter (Figure A1).

Figure 7: **Image reconstruction quality.** **Left:** Quantitative comparison on 15542 statistical images. The R^2 loss between the true and recovered images after compression with brain atlases of multiple resolutions. In general, higher-order atlases capture more signal. **Right:** Meta-analysis summaries for the motor task. High R^2 score (left) correspond to better capturing fine structures of images, as visible on the qualitative images. DiFuMo atlases better capture the gradients and smooth aspects of the original images than hard parcellations, as BASC.





Figure 8: **Modes around the putamen,** for DiFuMo dimensionality 64, 256, and 512. As dimensionality increases: sub-divisions are more refined, modes are split into right and left hemisphere and antero-posterior direction. Each color represents a single mode. [Figure A13](#) details more this breakdown.



A large dataset is needed to capture such implicit regularities of the signal with high-dimensional spatial decompositions. Indeed, running the same model on less data extracts modes with less spatial regularity (Figure A2). The combination of high dimensionality and large dataset leads to significant computational demands. The extraction of DiFuMos was possible thanks to fast algorithms for huge matrix factorization (Mensch et al., 2018), and gathering data representative of a wide variety of scanning protocols via openfMRI (Poldrack et al., 2013).

We did not limit the DiFuMo modes to gray matter, as measures outside gray matter can be useful in subsequent analysis, for instance to remove the global signal (Murphy and Fox, 2017). In addition, distributed modes extracted from full-brain EPI can separate out noise –such as movement artifacts– and help rejecting it in a later analysis (Perlbarg et al., 2007; Griffanti et al., 2014; Pruim et al., 2015). Some DiFuMo modes indeed segment ventricles or interfaces. Depending on the application, practitioners can choose to restrict signal extraction to a grey-matter mask.

The functional modes are sharp and anatomically relevant. To extract structures defined by brain anatomy or microstructure, atlasing efforts have used anatomical or multimodal imaging (Mori et al., 2005; Desikan et al., 2006; Eickhoff et al., 2007; Glasser et al., 2016). The DiFuMo atlases capture a different signal: brain activity. Yet, thanks to the sparsity and non-negativity constraint, they are made of localized modes which often have a natural anatomical interpretation. Consequently, we have labeled the modes with a unique name based on the most relevant anatomical structure, following Urchs et al. (2019) who also give anatomical labels to functional regions. Indeed, using a common vocabulary of brain structures is impor-

tant for communication across the neuroimaging community. As visible on Figure 8 and Figure A14, the modes are well anchored on anatomical structures such as the putamen and the thalamus. They are however not constrained to contain only one connected region. Smaller dimension DiFuMos indeed capture distributed networks, often comprising bilateral regions. As the dimensionality increases, the networks progressively separate in smaller networks which eventually form single regions. For instance, the left and right putamen appear in the same mode at dimension 64, and are first sub-divided along the antero-posterior direction, and later the left and right putamen are separated (Figure 8). Dimension choice is data driven: it should best explain the functional signal.

6. Conclusion

We provide multidimensional atlases of functional modes that can be used to extract functional signals: parietal-inria.github.io/DiFuMo. They give excellent performance for a wide variety of analytic tasks: GLM-based analysis, mental-process decoding or functional-connectivity analysis. Their availability reduces computational burdens: practitioners can readily perform analyses on a reduced signal, without a costly ROI-definition step. In addition, working on common functional modes across studies facilitates comparison and interpretations of results. To help communication, we have labeled every functional mode to reflect the neuroanatomical structures that it contains. To date, these are the only high-dimensional soft functional modes available. As they have been extracted from a variety of data (more than 2,000 sessions across 27 studies, 2.4TB in size) and improve many

analytic tasks, the rich descriptions of neural activity that they capture is well suited for a broad set of fMRI studies.

7. Acknowledgments

This project has received funding from the European Union’s Horizon 2020 Research and Innovation Programme under Grant Agreement No. 785907 (HBP SGA2) and No 826421 (VirtualBrainCloud). The work of Arthur Mensch has been supported by the European Research Council (ERC Grant Noria). This work acknowledges the support of ANR NeuroRef and ERC-StG NeuroLang. Gaël Varoquaux was partially funded by the Canada First Research Excellence Fund, awarded to McGill University for the Healthy Brains for Healthy Lives initiative.

We also thank Pierre Bellec and Vincent Frouin for their helpful discussions on the experimental work, the neuroimaging community for giving access to fMRI datasets, and open-source contributors on the packages we build upon (including *nilearn*, *fMRIprep*, and *MRIQC*).

Data collection and sharing for this project was funded by the Alzheimer’s Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: AbbVie, Alzheimer’s Association; Alzheimer’s Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research and Development, LLC.; Johnson and Johnson Pharmaceutical Research and Development LLC.; Lumosity; Lundbeck; Merck and Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer’s Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.

References

Abou Elseoud, A., Littow, H., Remes, J., Starck, T., Nikkinen, J., Nissilä, J., Tervonen, O., Timonen, M., Kiviniemi, V., 2011.

Group-ica model order highlights patterns of functional brain connectivity. *Frontiers in Systems Neuroscience* 5, 37.

Abraham, A., Dohmatob, E., Thirion, B., Samaras, D., Varoquaux, G., 2014a. Region segmentation for sparse decompositions: better brain parcellations from rest fMRI. *Frontiers in neuroinformatics* 8.

Abraham, A., Milham, M.P., Di Martino, A., Craddock, R.C., Samaras, D., Thirion, B., Varoquaux, G., 2017. Deriving reproducible biomarkers from multi-site resting-state data: An autism-based example. *NeuroImage* 147, 736–745.

Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Mueller, A., Kossaifi, J., Gramfort, A., Thirion, B., Varoquaux, G., 2014b. Machine learning for neuroimaging with scikit-learn. *Frontiers in neuroinformatics* 8.

Altmann, A., Ng, B., Landau, S.M., Jagust, W.J., Greicius, M.D., 2015. Regional brain hypometabolism is unrelated to regional amyloid plaque burden. *Brain* 138, 3734–3746.

Alvarez, R., Poldrack, R., 2011. Cross-language repetition priming. *Stanford Digital Repository* .

Aron, A.R., Behrens, T.E., Smith, S., Frank, M.J., Poldrack, R.A., 2007. Triangulating a cognitive control network using diffusion-weighted magnetic resonance imaging (mri) and functional mri. *Journal of Neuroscience* 27, 3743–3752.

Aron, A.R., Gluck, M.A., Poldrack, R.A., 2006. Long-term test-retest reliability of functional mri in a classification learning task. *NeuroImage* 29, 1000 – 1006.

Arslan, S., Ktena, S.I., Makropoulos, A., Robinson, E.C., Rueckert, D., Parisot, S., 2017. Human brain mapping: A systematic comparison of parcellation methods for the human cerebral cortex. *NeuroImage* .

Barachant, A., Bonnet, S., Congedo, M., Jutten, C., 2013. Classification of covariance matrices using a riemannian-based kernel for bci applications. *Neurocomputing* 112, 172 – 178.

Beckmann, C., DeLuca, M., Devlin, J., Smith, S., 2005. Investigations into resting-state connectivity using independent component analysis. *Philos Trans R Soc Lond B* 360, 1001.

Behzadi, Y., Restom, K., Liau, J., Liu, T., 2007. A component based noise correction method (compcor) for BOLD and perfusion based fMRI. *Neuroimage* 37, 90.

Bellec, P., Perlberg, V., Jbabdi, S., Pelegrini-Issac, M., Anton, J.L., Doyon, J., Benali, H., 2006. Identification of large-scale networks in the brain using fMRI. *Neuroimage* 29, 1231.

Bellec, P., Rosa-Neto, P., Lyttelton, O., Benali, H., Evans, A., 2010. Multi-level bootstrap analysis of stable clusters in resting-state fMRI. *NeuroImage* 51, 1126.

Benjamini, Y., Hochberg, Y., 1995. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J R STAT SOC B (Methodological)* 57, 289.

Brett, M., Johnsrude, I.S., Owen, A.M., 2002. The problem of functional localization in the human brain. *Nat Rev Neurosci* 3, 243.

Brier, M.R., Mitra, A., McCarthy, J.E., Ances, B.M., Snyder, A.Z., 2015. Partial covariance based functional connectivity computation using ledoit-wolf covariance regularization. *NeuroImage* 121, 29.

Brodman, K., 1909. *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Barth.

Bzdok, D., Varoquaux, G., Grisel, O., Eickenberg, M., Poupon, C., Thirion, B., 2016. Formal models of the network co-occurrence underlying mental operations. *PLOS Computational Biology* 12, 1–31.

Calhoun, V., Sui, J., Kiehl, K., Turner, J., Allen, E., Pearlson, G., 2012. Exploring the psychosis functional connectome: Aberrant intrinsic networks in schizophrenia and bipolar disorder. *Frontiers in Psychiatry* .

- Calhoun, V.D., Adali, T., Pearlson, G.D., Pekar, J.J., 2001. A method for making group inferences from fMRI data using independent component analysis. *Hum Brain Mapp* 14, 140.
- Calhoun, V.D., Kiehl, K.A., Pearlson, G.D., 2008. Modulation of temporally coherent brain networks estimated using ICA at rest and during cognitive tasks. *Hum Brain Map* 29, 828.
- Catani, M., de Schotten, M.T., 2012. *Atlas of Human Brain Connections*. Oxford University Press.
- Cera, N., Tartaro, A., Sensi, S.L., 2014. Modafinil alters intrinsic functional connectivity of the right posterior insula: A pharmacological resting state fmri study. *PLOS ONE* 9, 1–12.
- Chang, L.J., Gianaros, P.J., Manuck, S.B., Krishnan, A., Wager, T.D., 2015. A sensitive and specific neural signature for picture-induced negative affect. *PLOS Biology* 13, 1–28.
- Craddock, R.C., James, G.A., Holtzheimer, P.E., Hu, X.P., Mayberg, H.S., 2012. A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Hum brain map* 33, 1914.
- Dadi, K., Rahim, M., Abraham, A., Chyzyk, D., Milham, M., Thirion, B., Varoquaux, G., 2019. Benchmarking functional connectome-based predictive models for resting-state fMRI. *NeuroImage* 192, 115 – 134.
- Dalcin, L., Bradshaw, R., Smith, K., Citro, C., Behnel, S., Seljebotn, D., 2011. Cython: The best of both worlds. *Computing in Science & Engineering* 13, 31–39.
- Damoiseaux, J.S., Rombouts, S.A.R.B., Barkhof, F., Scheltens, P., Stam, C.J., Smith, S.M., Beckmann, C.F., 2006. Consistent resting-state networks across healthy subjects. *Proc Natl Acad Sci* 103, 13848.
- Delgado, M.R., Nystrom, L.E., Fissell, C., Noll, D.C., Fiez, J.A., 2000. Tracking the hemodynamic responses to reward and punishment in the striatum. *Journal of Neurophysiology* 84, 3072–3077.
- Desikan, R., S., Ségonne, F., Fischl, B., Quinn, B., T., Dickerson, B., C., Blacker, D., Buckner, R., L., Dale, A., M., Maguire, R., P., Hyman, B., T., Albert, M., S., Killiany, R., J., 2006. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage* 31, 968.
- Destrieux, C., Fischl, B., Dale, A., Halgren, E., 2010. Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. *NeuroImage* 53, 1 – 15.
- Devlin, J.T., Poldrack, R.A., 2007. In praise of tedious anatomy. *NeuroImage* 37, 1033 – 1041.
- Di Martino, A., Yan, C.G., Li, Q., Denio, E., Castellanos, F.X., Alaerts, K., Anderson, J.S., Assaf, M., Bookheimer, S.Y., Dapretto, M., et al., 2014. The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Molecular psychiatry* 19, 659–667.
- Dice, L.R., 1945. Measures of the amount of ecologic association between species. *Ecology* 26, 297–302.
- Diedrichsen, J., Balsters, J.H., Flavell, J., Cussans, E., Ramnani, N., 2009. A probabilistic mr atlas of the human cerebellum. *NeuroImage* 46, 39 – 46.
- Dohmatob, E., Mensch, A., Varoquaux, G., Thirion, B., 2016. Learning brain regions via large-scale online structured sparse dictionary-learning, in: *Advances in Neural Information Processing Systems*.
- Dubois, J., Adolphs, R., 2016. Building a science of individual differences from fmri. *Trends in cognitive sciences* 20, 425–443.
- Duff, E.P., Trachtenberg, A.J., Mackay, C.E., Howard, M.A., Wilson, F., Smith, S.M., Woolrich, M.W., 2012. Task-driven ica feature generation for accurate and interpretable prediction using fmri. *NeuroImage* 60, 189 – 203.
- Duncan, K.J., Pattamadilok, C., Knierim, I., Devlin, J.T., 2009. Consistency and variability in functional localisers. *Neuroimage* 46, 1018.
- Eickhoff, S.B., Paus, T., Caspers, S., Grosbras, M.H., Evans, A.C., Zilles, K., Amunts, K., 2007. Assignment of functional activations to probabilistic cytoarchitectonic areas revisited. *Neuroimage* 36, 511.
- Eickhoff, S.B., Yeo, B.T.T., Genon, S., 2018. Imaging-based parcellations of the human brain. *Nat Rev Neurosci* 19, 672.
- Elliott, P., Peakman, T.C., et al., 2008. The UK biobank sample handling and storage protocol for the collection, processing and archiving of human blood and urine. *Int J Epidemiology* 37, 234.
- Esteban, O., Birman, D., Schaer, M., Koyejo, O.O., Poldrack, R.A., Gorgolewski, K.J., 2017. MRIQC: Advancing the automatic prediction of image quality in MRI from unseen sites. *PLOS ONE* 12, 1.
- Esteban, O., Markiewicz, C.J., Blair, R.W., Moodie, C.A., Isik, A.I., Erramuzpe, A., Kent, J.D., Goncalves, M., DuPre, E., Snyder, M., et al., 2019. fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nature methods* 16, 111.
- Foerde, K., Knowlton, B.J., Poldrack, R.A., 2006. Modulation of competing memory systems by distraction. *Proc Natl Acad Sci* 103, 11778.
- Friedman, J., Hastie, T., Höfling, H., Tibshirani, R., 2007. Pathwise coordinate optimization. Technical Report. *Annals of Applied Statistics*.
- Friston, K., Fletcher, P., Josephs, O., Holmes, A., Rugg, M., Turner, R., 1998. Event-related fmri: Characterizing differential responses. *NeuroImage* 7, 30 – 40.
- Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.B., Frith, C., Frackowiak, R., 1995. Statistical parametric maps in functional imaging: A general linear approach. *Hum Brain Mapp* , 189.
- Gabito, E., Manor, D., Karni, A., 2015. Patterns of modulation in the activity and connectivity of motor cortex during the repeated generation of movement sequences. *J Cog Neurosci* 27, 736.
- Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., Smith, S.M., Essen, D.C.V., 2016. A multi-modal parcellation of human cerebral cortex. *Nature* 536, 171–178.
- Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Essen, D.C.V., Jenkinson, M., 2013. The minimal preprocessing pipelines for the human connectome project. *NeuroImage* 80, 105 – 124.
- Gordon, E.M., Laumann, T.O., Adeyemo, B., Huckins, J.F., Kelley, W.M., Petersen, S.E., 2014. Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations. *Cerebral Cortex* 26, 288–303.
- Gorgolewski, K., Esteban, O., Gunnar, S., Brain, W., Poldrack, R., 2017. Openneuro – a free online platform for sharing and analysis of neuroimaging data., in: *23rd Annual Meeting of the Organization for Human Brain Mapping*, p. 1677.
- Gorgolewski, K.J., Storkey, A., Bastin, M.E., Whittle, I.R., Wardlaw, J.M., Pernet, C.R., 2013. A test-retest fmri dataset for motor, language and spatial attention functions. *GigaScience* 2, 2047–217X–2–6.
- Gorgolewski, K.J., Varoquaux, G., Rivera, G., Schwarz, Y., Ghosh, S.S., Maumet, C., Sochat, V.V., Nichols, T.E., Poldrack, R.A., Poline, J.B., Yarkoni, T., Margulies, D.S., 2015. Neurovault.org: a web-based repository for collecting and sharing unthresholded statistical maps of the human brain. *Frontiers in Neuroinformatics* 9, 8.
- Goutte, C., Toft, P., Rostrup, E., Nielsen, F.A., Hansen, L.K., 1999. On clustering fMRI time series. *NeuroImage* 9, 298–310.
- Greicius, M., Krasnow, B., Reiss, A., Menon, V., 2003. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proc Natl Acad Sci* 100, 253.

- Greve, D.N., Fischl, B., 2009. Accurate and robust brain image alignment using boundary-based registration. *NeuroImage* 48, 63.
- Griffanti, L., Salimi-Khorshidi, G., Beckmann, C.F., Auerbach, E.J., Douaud, G., Sexton, C.E., Zsoldos, E., Ebmeier, K.P., Filippini, N., Mackay, C.E., et al., 2014. ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. *Neuroimage* 95, 232–247.
- Hanson, S.J., Matsuka, T., Haxby, J.V., 2004. Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: is there a “face” area? *NeuroImage* 23, 156 – 166.
- Harrison, S.J., Woolrich, M.W., Robinson, E.C., Glasser, M.F., Beckmann, C.F., Jenkinson, M., Smith, S.M., 2015. Large-scale probabilistic functional modes from resting state fmri. *NeuroImage* 109, 217 – 231.
- Hastie, T., Tibshirani, R., Friedman, J., 2009. *The elements of statistical learning*. Springer.
- Haxby, J.V., Gobbini, I.M., Furey, M.L., et al., 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293, 2425.
- Haynes, J.D., Rees, G., 2006. Decoding mental states from brain activity in humans. *Nat. Rev. Neurosci.* 7, 523.
- Henri, M.D., 1999. *The Human Brain: Surface, Three-dimensional Sectional Anatomy with MRI, and Blood Supply*. Springer.
- Hua, K., Zhang, J., Wakana, S., Jiang, H., Li, X., Reich, D.S., Calabresi, P.A., Pekar, J.J., van Zijl, P.C., Mori, S., 2008. Tract probability maps in stereotaxic spaces: Analyses of white matter anatomy and tract-specific quantification. *NeuroImage* 39, 336.
- Huntenburg, J.M., Bazin, P.L., Margulies, D.S., 2018. Large-scale gradients in human cortical organization. *Trends in cognitive sciences* 22, 21.
- Hyvärinen, A., Oja, E., 2000. Independent component analysis: algorithms and applications. *Neural Networks* 13, 411.
- Iannilli, E., Gasparotti, R., Hummel, T., Zoni, S., Benedetti, C., Fedrighi, C., Tang, C.Y., Van Thriel, C., Lucchini, R.G., 2016. Effects of manganese exposure on olfactory functions in teenagers: A pilot study. *PLOS ONE* 11, 1–9.
- Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage* 17, 825 – 841.
- Jimura, K., Cazalis, F., Stover, E.R.S., Poldrack, R.A., 2014. The neural basis of task switching changes with skill acquisition. *Frontiers in Human Neuroscience* 8, 339.
- Kelly, A.C., Uddin, L.Q., Biswal, B.B., Castellanos, F.X., Milham, M.P., 2008. Competition between functional brain networks mediates behavioral variability. *NeuroImage* 39, 527 – 537.
- Kim, J., Wang, J., Wedell, D.H., Shinkareva, S.V., 2016. Identifying core affect in individuals from fmri responses to dynamic naturalistic audiovisual stimuli. *PLOS ONE* 11, 1–21.
- Kiviniemi, V., Kantola, J., Jauhiainen, J., Hyvärinen, A., Tervonen, O., 2003. Independent component analysis of nondeterministic fmri signal sources. *Neuroimage* 19, 253.
- Kiviniemi, V., Starck, T., Remes, J., Long, X., Nikkinen, J., Haapea, M., Veijola, J., et al., 2009. Functional segmentation of the brain cortex using high model order group PICA. *Hum Brain Map* 30, 3865.
- Ledoit, O., Wolf, M., 2004. A well-conditioned estimator for large-dimensional covariance matrices. *J. Multivar. Anal.* 88, 365.
- Lee, K., Tak, S., Ye, J.C., 2010. A data-driven sparse GLM for fMRI analysis using sparse dictionary learning with MDL criterion. *IEEE Trans Med Imag* 30, 1076.
- Leech, R., Kamourieh, S., Beckmann, C.F., Sharp, D.J., 2011. Fractionating the default mode network: Distinct contributions of the ventral and dorsal posterior cingulate cortex to cognitive control. *J Neurosci* 31, 3217.
- Lepping, R.J., Atchley, R.A., Chryssikou, E., Martin, L.E., Clair, A.A., Ingram, R.E., Simmons, W.K., Savage, C.R., 2016a. Neural processing of emotional musical and nonmusical stimuli in depression. *PLOS ONE* 11, 1–23.
- Lepping, R.J., Atchley, R.A., Savage, C.R., 2016b. Development of a validated emotionally provocative musical stimulus set for research. *Psychology of Music* 44, 1012–1028.
- Mennes, M., Kelly, C., Colcombe, S., Castellanos, F.X., Milham, M.P., 2013. The extrinsic and intrinsic functional architectures of the human brain are not equivalent. *Cerebral Cortex* 23, 223–229.
- Mensch, A., Mairal, J., Bzdok, D., Thirion, B., Varoquaux, G., 2017. Learning Neural Representations of Human Cognition across Many fMRI Studies, in: *Neural Information Processing Systems*, p. 5885.
- Mensch, A., Mairal, J., Thirion, B., Varoquaux, G., 2016a. Dictionary Learning for Massive Matrix Factorization, in: *International Conference on Machine Learning*, pp. 1737–1746.
- Mensch, A., Mairal, J., Thirion, B., Varoquaux, G., 2018. Stochastic Subsampling for Factorizing Huge Matrices. *IEEE Trans Sig Proc* 66, 113.
- Mensch, A., Varoquaux, G., Thirion, B., 2016b. Compressed Online Dictionary Learning for Fast Resting-State fMRI Decomposition, in: *Proc. ISBI*, p. 1282.
- Michel, V., Gramfort, A., Varoquaux, G., Eger, E., Keribin, C., Thirion, B., 2012. A supervised clustering approach for fMRI-based inference of brain states. *Pattern Recognition* 45, 2041.
- Miller, K.L., Alfaro-Almagro, F., et al., 2016. Multimodal population brain imaging in the UK biobank prospective epidemiological study. *Nature Neuroscience* .
- Moran, J.M., Jolly, E., Mitchell, J.P., 2012. Social-cognitive deficits in normal aging. *J Neurosci* 32, 5553.
- Mori, S., Wakana, S., Van Zijl, P.C., Nagae-Poetscher, L., 2005. *MRI atlas of human white matter*. Elsevier.
- Mourão-Miranda, J., Bokde, A.L., Born, C., Hampel, H., Stetter, M., 2005. Classifying brain states and determining the discriminating activation patterns: Support vector machine on functional MRI data. *NeuroImage* 28, 980.
- Mueller, S., Weiner, M., Thal, L., Petersen, R., Jack, C., Jagust, W., Trojanowski, J.Q., Toga, A.W., Beckett, L., 2005. The alzheimer’s disease neuroimaging initiative. *Neuroimaging Clinics of North America* 15, 869.
- Murphy, K., Fox, M.D., 2017. Towards a consensus regarding global signal regression for resting state functional connectivity MRI. *Neuroimage* 154, 169–173.
- Olshausen, B., Field, D., 1997. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision research* 37, 3311.
- Ono, M., Kubik, S., Abernathey, C.D., 1990. *Atlas of the cerebral sulci*. G. Thieme Verlag.
- O’Toole, A.J., Jiang, F., Abdi, H., Haxby, J.V., 2005. Partially distributed representations of objects and faces in ventral temporal cortex. *J Cog Neurosci* 17, 580.
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al., 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825.
- Perlberg, V., Bellec, P., Anton, J.L., Pelegrini-Issac, M., Doyon, J., Benali, H., 2007. CORSICA: correction of structured noise in fMRI by automatic identification of ICA components. *Magn Reson Imaging* 25, 35.
- Pervaz, U., Vidaurre, D., Woolrich, M.W., Smith, S.M., 2019. Optimising network modelling methods for fmri. *bioRxiv* .
- Pinel, P., Thirion, B., Meriaux, S., Jobert, A., Serres, J., Le Bihan, D., Poline, J., Dehaene, S., 2007. Fast reproducible identification and large-scale databasing of individual functional cognitive networks. *BMC neuroscience* 8, 91.

- Pinho, A.L., Amadon, A., Ruest, T., Fabre, M., Dohmatob, E., Denghien, I., Ginisty, C., Becuwe-Desmidt, S., Roger, S., Laurier, L., Joly-Testault, V., Médiouni-Cloarec, G., Doublé, C., Martins, B., Pinel, P., Eger, E., Varoquaux, G., Pallier, C., Dehaene, S., Hertz-Pannier, L., Thirion, B., 2018. Individual Brain Charting, a high-resolution fMRI dataset for cognitive mapping. *Scientific Data* 5, 180105.
- Poldrack, R.A., Barch, D.M., Mitchell, J.P., et al., 2013. Toward open sharing of task-based fMRI data: the OpenfMRI project. *Frontiers in neuroinformatics* 7.
- Poldrack, R.A., Clark, J., Paré-Blagoev, E.J., Shohamy, D., Creso Moyano, J., Myers, C., Gluck, M.A., 2001. Interactive memory systems in the human brain. *Nature* 414, 546–550.
- Poldrack, R.A., Halchenko, Y.O., Hanson, S.J., 2009. Decoding the large-scale structure of brain function by classifying mental states across individuals. *Psychological Science* 20, 1364.
- Power, J., Cohen, A., Nelson, S., Wig, G., Barnes, K., Church, J., Vogel, A., Laumann, T., Miezin, F., Schlaggar, B., Petersen, S., 2011. Functional network organization of the human brain. *Neuron* 72, 665–678.
- Pruim, R.H., Mennes, M., Buitelaar, J.K., Beckmann, C.F., 2015. Evaluation of ICA-AROMA and alternative strategies for motion artifact removal in resting state fMRI. *Neuroimage* 112, 278–287.
- Rademacher, J., Caviness, V. S., J., Steinmetz, H., Galaburda, A.M., 1993. Topographical Variation of the Human Primary Cortices: Implications for Neuroimaging, Brain Mapping, and Neurobiology. *Cerebral Cortex* 3, 313–329.
- Rademacher, J., Galaburda, A.M., Kennedy, D.N., Filipek, P.A., Caviness, V.S., 1992. Human cerebral cortex: Localization, parcellation, and morphometry with magnetic resonance imaging. *J Cog Neurosci* 4, 352.
- Repovs, G., Barch, D., 2012. Working memory related brain network connectivity in individuals with schizophrenia and their siblings. *Frontiers in Human Neuroscience* 6, 137.
- Richiardi, J., Eryilmaz, H., Schwartz, S., Vuilleumier, P., Van De Ville, D., 2010. Decoding brain states from fMRI connectivity graphs. *NeuroImage* .
- Rizk-Jackson, A., Aron, A., Poldrack, R., 2011. Classification learning and stop-signal (1 year test-retest). *Stanford Digital Repository* .
- Romaniuk, L., Pope, M., Nicol, K., Steele, D., Hall, J., 2016. Neural correlates of fears of abandonment and rejection in borderline personality disorder. *Wellcome Open Research* 1.
- Roy, A., Bernier, R.A., Wang, J., Benson, M., French, Jr., J.J., Good, D.C., Hillary, F.G., 2017. The evolution of cost-efficiency in neural networks during recovery from traumatic brain injury. *PLOS ONE* 12, 1–26.
- Sala-Llonch, R., Smith, S.M., Woolrich, M., Duff, E.P., 2019. Spatial parcellations, spectral filtering, and connectivity measures in fmri: Optimizing for discrimination. *Hum Brain Map* 40, 407.
- Schaefer, A., Kong, R., Gordon, E.M., Laumann, T.O., Zuo, X.N., Holmes, A.J., Eickhoff, S.B., Yeo, B.T.T., 2017. Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cerebral Cortex* 28, 3095–3114.
- Schmahmann, J.D., Doyon, J., McDonald, D., Holmes, C., Lavoie, K., Hurwitz, A.S., Kabani, N., Toga, A., Evans, A., Petrides, M., 1999. Three-dimensional mri atlas of the human cerebellum in proportional stereotaxic space. *NeuroImage* 10, 233 – 260.
- Schonberg, T., Fox, C., Mumford, J., Congdon, E., Trepel, C., Poldrack, R., 2012. Decreasing ventromedial prefrontal cortex activity during sequential risk-taking: An fmri investigation of the balloon analog risk task. *Frontiers in Neuroscience* 6, 80.
- Shirer, W., Ryali, S., Rykhlevskaia, E., Menon, V., Greicius, M., 2012. Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cerebral Cortex* 22, 158.
- Smith, R., Keramatian, K., Christoff, K., 2007. Localizing the rostralateral prefrontal cortex at the individual level. *NeuroImage* 36, 1387 – 1396.
- Smith, S., Fox, P., Miller, K., Glahn, D., Fox, P., Mackay, C., et al., 2009. Correspondence of the brain’s functional architecture during activation and rest. *Proc Natl Acad Sci* 106, 13040.
- Smith, S., Miller, K., Salimi-Khorshidi, G., Webster, M., Beckmann, C., Nichols, T., Ramsey, J., Woolrich, M., 2011. Network modelling methods for fMRI. *Neuroimage* 54, 875.
- Sporns, O., Tononi, G., Kotter, R., 2005. The human connectome: a structural description of the human brain. *PLoS Comput Biol* 1, e42.
- Stephan-Otto, C., Siddi, S., Senior, C., Muñoz-Samons, D., Ochoa, S., Sánchez-Laforga, A.M., Brébion, G., 2017. Visual imagery and false memory for pictures: A functional magnetic resonance imaging study in healthy participants. *PLOS ONE* 12, 1–17.
- Taylor, J.R., Williams, N., Cusack, R., Auer, T., Shafto, M.A., Dixon, M., Tyler, L.K., Cam-CAN, Henson, R.N., 2017. The cambridge centre for ageing and neuroscience (cam-CAN) data repository: Structural and functional MRI, MEG, and cognitive data from a cross-sectional adult lifespan sample. *NeuroImage* 144, 262.
- Thirion, B., Flandin, G., Pinel, P., Roche, A., Ciuciu, P., Poline, J., 2006. Dealing with the shortcomings of spatial normalization: Multi-subject parcellation of fMRI datasets. *Hum brain map* 27, 678.
- Thirion, B., Varoquaux, G., Dohmatob, E., Poline, J., 2014. Which fMRI clustering gives good brain parcellations? *Frontiers in Neuroscience* 8, 167.
- Uncapher, M.R., Hutchinson, J.B., Wagner, A.D., 2011. Dissociable effects of top-down and bottom-up attention during episodic encoding. *Journal of Neuroscience* 31, 12613–12628.
- Urchs, S., Armoza, J., Moreau, C., Benhajali, Y., St-Aubin, J., Urban, P., Bellec, P., 2019. MIST: A multi-resolution parcellation of functional brain networks. *MNI Open Research* 1.
- Van Essen, D., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S., Della Penna, S., Feinberg, D., Glasser, M., Harel, N., Heath, A., Larson-Prior, L., Marcus, D., Michalareas, G., Moeller, S., Oostenveld, R., Petersen, S., Prior, F., Schlaggar, B., Smith, S., Snyder, A., Xu, J., Yacoub, E., 2012. The human connectome project: A data acquisition perspective. *NeuroImage* 62, 2222–2231.
- Van Essen, D.C., Smith, et al., 2013. The wu-minn human connectome project: an overview. *Neuroimage* 80, 62–79.
- Varoquaux, G., Baronnet, F., Kleinschmidt, A., Fillard, P., Thirion, B., 2010a. Detection of brain functional-connectivity difference in post-stroke patients using group-level covariance modeling, in: *MICCAI*.
- Varoquaux, G., Craddock, R.C., 2013. Learning and comparing functional connectomes across subjects. *NeuroImage* 80, 405.
- Varoquaux, G., Gramfort, A., Pedregosa, F., Michel, V., Thirion, B., 2011. Multi-subject dictionary learning to segment an atlas of brain spontaneous activity, in: *Inf Proc Med Imag*, p. 562.
- Varoquaux, G., Raamana, P.R., Engemann, D.A., Hoyos-Idrobo, A., Schwartz, Y., Thirion, B., 2017. Assessing and tuning brain decoders: cross-validation, caveats, and guidelines. *NeuroImage* 145, 166.
- Varoquaux, G., Sadaghiani, S., Pinel, P., Kleinschmidt, A., Poline, J.B., Thirion, B., 2010b. A group model for stable multi-subject ICA on fMRI datasets. *NeuroImage* 51, 288.
- Verstynen, T.D., 2014. The organization and dynamics of corticostriatal pathways link the medial orbitofrontal cortex to future behavioral responses. *J Neurophysio* 112, 2457.




Figure A1: **Region volume (cm^3) of modes on the brain with 1024 dictionary of DiFuMo.** The volume of the modes tends to be larger corresponding to white matter when compared with the cortical gray matter. This justifies the adaptation of DiFuMo atlas to the fMRI signal.




Figure A2: 1024 components trained on two different sizes of the input set of fMRI images. The components trained on the full data have more spatial regularity, while the components trained on 100 volumes have more overlap in some regions of the brain. The additional spatial regularity shows the benefit of large-scale training size in learning a data-driven based functional atlas.

Xue, G., Aron, A.R., Poldrack, R.A., 2008. Common neural substrates for inhibition of spoken and manual responses. *Cerebral Cortex* 18, 1923–1932.

Xue, G., Poldrack, R.A., 2007. The neural substrates of visual perceptual learning of words: Implications for the visual word form area hypothesis. *J Cog Neurosci* 19, 1643.

Yeo, B., Krienen, F., Sepulcre, J., Sabuncu, M., et al., 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J Neurophysio* 106, 1125.

Zalesky, A., Fornito, A., Harding, I.H., Cocchi, L., Yücel, M., Pantelis, C., Bullmore, E.T., 2010. Whole-brain anatomical networks: Does the choice of nodes matter? *NeuroImage* 50, 970 – 983.

Task	# samples	Representation	Time (sec)	Speedup
Emotion	4924	Voxel-level	77.7	46×
		Reduced	1.7	
Pain	84	Voxel-level	1.5	250×
		Reduced	0.006	
Working memory	3140	Voxel-level	874.7	240×
		Reduced	3.7	
Gambling	1574	Voxel-level	298.7	270×
		Reduced	1.12	
Relational	1572	Voxel-level	263.1	405×
		Reduced	0.65	

Table A1: The comparison in computational times while predicting mental state on two set of brain features space: voxel-level $\approx 200,000$ and reduced voxels to DiFuMo 1024. We report the averaged time over 20 cross-validation folds for several task-fMRI conditions. Clearly, there are benefits trading for reduced representations in terms of computation time. On high-resolution brain images like HCP, these are decreased by a factor 200.

Task-fMRI	Prediction task	# maps
NV503: Emotion	Rating: 1, 2, 3, 4, 5	4924
NV504: Pain	Sensitivity: 1, 2, 3	84
HCP: Working mem.	face vs place	3140
HCP: Gambling	loss vs reward	1574
HCP: Relational	relational vs matching	1572
IBC: Archi standard	left vs right hand	1040

Table A2: **Dataset, prediction tasks and dataset size for each of the 6 decoding tasks we consider in §3.3.** z -maps from HCP and IBC were computed using the GLM, while NeuroVault directly provided the β -maps for Emotion and Pain. NV: NeuroVault.

Appendix A. Performance of DiFuMos

As discussed in §5, we report how DiFuMos components are well adapted to the fMRI EPI signal in Figure A1. Figure A2 qualitatively compare components obtained training on the whole data corpus and training on a fraction of it. Better component regularity is obtained with more data. Finally, Table A1 reports the computational speed-ups obtained using DiFuMos IDPs instead of voxel in the decoding experiment. Similar speed-ups are observed in the other validation pipelines.

Appendix B. Details on stimulus decoding

We provide additional details for the decoding pipeline, to complete the description in §3.3.

Appendix B.1. Input data and pre-processing pipelines

The decoding pipeline classifies input unthresholded statistical maps. Table A2 summarizes the task-based studies used to obtain these statistical maps.




Figure A3: **Decoding prediction scores for each brain atlas and target:** Each marker denotes the mean performance of using a certain brain atlas; error bars are the standard deviation of the prediction scores for this atlas. Decoding from high-order dictionaries, and especially from DiFuMos, perform similarly or better than decoding from voxels.

Pre-encoded maps downloaded from Neurovault.org. We download maps related to emotion and pain (Chang et al., 2015) using *Neurovault*, querying the collections 503 and 504. We use the “Rating” & “PainLevel” labels as predictive targets. We predict emotion using ridge regression, and pain-level over 3 classes using Linear SVC. The super-

vised learning pipeline, that includes cross-validation and linear models is implemented with Python based *scikit-learn* (Pedregosa et al., 2011). We use *nilearn* (Abraham et al., 2014b) to download maps from Neurovault.org interface (Gorgolewski et al., 2015). The data acquisition parameters, preprocessing details and estimation of statistical maps are described in Chang et al. (2015).

Statistical maps encoded using the GLM. We compute z -maps from HCP900 (Van Essen et al., 2012) and IBC (Pinho et al., 2018) studies, that comprise high-quality task-fMRI experiments.

HCP. We download fMRI data from the HCP900 release; those are already preprocessed using HCP pipelines (Glasser et al., 2013). We use MNINonLinear-based registered data as input for the GLM, that outputs one z -map per condition per subject. We consider three task-based studies, namely: for **Working Memory**, we consider z -maps based on condition: “0-back faces”, “2-back faces”, “0-back places”, “2-back places”. Similarly, for **Gambling** (Delgado et al., 2000), we consider z -maps for the conditions “loss” and “reward”; finally, on **Relational processing**, we consider z -maps for the conditions “relational processing” and “matching”. For each study, we use Linear SVC on encoded z -maps to predict psychological conditions. The predictive model therefore perform a 2-class or 4-class classification. The experimental protocol and data acquisition parameters are detailed in Van Essen et al. (2012).

IBC. We consider the **Archi Standard** (Pinel et al., 2007) motor task, where subjects are asked to press “left” and “right” button press based on audio and visual instructions. We perform within-subject classification between left and right button press, using z -maps corresponding to each repetition of the instruction. For each of the 13 available subjects, a linear model is trained on the z -maps from all but one session and prediction is performed on the left-out session. Each subject provides 80 encoded z -maps across 4 sessions. We use data preprocessed following the pipelines of Pinho et al. (2018).

GLM specification. For both datasets, the input z -maps are estimated from the raw fMRI data by fitting a GLM. We use *Nistats*⁷, a Python package for the statistical analysis of fMRI data. The temporal regressors of the model are specified according to the timing of stimulus presentations convolved with hemodynamic models ($spm + derivative$). We use polynomial model to capture the low-frequency drifts in the data.

Appendix B.2. Detailed results

To complete the summarizing Figure 4, we report the raw prediction scores separately for each decoding tasks in Figure A3. Prediction accuracy increases with the size of functional atlases. Using 1024 atlases allows to match or

⁷<https://nistats.github.io/>

Decoding working memory: 0BK face versus rest




Figure A4: Decoding classification weight maps for the HCP working memory task (0BK face), obtained with voxel-level decoding and decoding over various functional atlases. Using DiFuMo yield highly interpretable weight maps; it clearly delineates the fusiform gyrus and lateral occipital cortex.

pass the performance of voxel-based prediction. In terms of interpretation, the weights are much smoother and blobs are clearly visible in the weight classification maps obtained using DiFuMo. This is illustrated on Figure A4 for face-versus-place decoding in the working-memory HCP study.

Rest-fMRI	Prediction groups	Samples
HCP900	High IQ vs Low IQ 213/230	443 subjects
ABIDE	Autism vs control 402/464	866 subjects
ACPI	Marijuana use vs control 62/64	126 subjects
ADNI	Alzheimers vs MCI 40/96	136 subjects
ADNIDOD	PTSD vs control 89/78	167 subjects
COBRE	Schizophrenia vs control 65/77	142 subjects
CamCAN	Age 24 – 86	626 subjects

Table A3: Resting-state fMRI datasets used in the pipeline described on §3.4 for predicting phenotypic labels from functional connectomes. In CamCAN, age is predicted using ridge regression. The groups from other datasets are predicted using logistic regression. IQ - Fluid intelligence, PTSD - Post Traumatic Stress Disorder, MCI - Mild Cognitive Impairment.

Appendix C. Details on biomarker prediction

We consider multiple datasets to account for the diversity of prediction targets in biomarker prediction problem. We report datasets, prediction groups and prediction targets in Table A3.

Appendix C.1. Input data and prediction settings

The connectivity features built from functional atlases predict various clinical outcomes (neuro-degenerative and neuro-psychiatric disorders, drug abuse impact) and psychological traits.

Group classification. We use the Alzheimer’s Disease Neuroimaging Initiative⁸ (ADNI) and (ADNIDOD) (Mueller et al., 2005) to predict neuro-degenerative diseases. We discriminate between Alzheimer’s Disease (AD) from Mild Cognitive Impairment (MCI) group on ADNI. We discriminate between post-traumatic stress disorder (PTSD) and healthy individuals on ADNIDOD. We use data from the Center for Biomedical Research Excellence⁹ (COBRE Calhoun et al. (2012) to predict schizophrenia diagnosis of individuals. We classify autism and healthy individuals on Autism Brain Imaging Data Exchange database (ABIDE, Di Martino et al. (2014), Finally, we consider data from Addiction Connectome Preprocessed Initiative¹⁰ (ACPI), where we discriminate Marijuana consumers versus control subjects.

⁸www.adni-info.org

⁹<https://www.mrn.org/research/details/cobre>

¹⁰http://fcon_1000.projects.nitrc.org/indi/ACPI/html/




Figure A5: **Connectome prediction scores for each brain atlas and target:** Each marker denotes the mean performance of using a certain brain atlas; error bars are the standard deviation of the prediction scores for this atlas. BASC and DiFuMo-based atlases give good prediction scores up to $k = 256$ ROIs.

Psychological traits. We first stratify individuals from HCP900 release (Van Essen et al., 2013) into groups of high and low IQ, and perform binary classification on these. The details about the stratification into these groups are described in Dadi et al. (2019).

Age regression. We use Cambridge Center for Ageing and Neuroscience (CamCAN) dataset (Taylor et al., 2017) to study brain ageing. This study comprises wide range of age groups spanning from 24 – 86. We use ridge regression to predict age on this dataset.

Appendix C.2. Data acquisition parameters and pre-processing pipelines

The data acquisition details for ADNI, ADNIDOD, COBRE, ABIDE, ACPI and HCP are described in Dadi et al. (2019); those for CamCAN in Taylor et al. (2017). We pre-process individuals from CamCAN, ADNI, ADNIDOD and COBRE. All rs-fMRI acquisitions are pre-processed with standard steps, described in Dadi et al. (2019). The other considered datasets provide preprocessed data. We report the total number of subjects included in the analysis in Table A3, after excluding for severe scanning artifacts, head movements with amplitude larger than 2 mm and individuals who have more than one clinical diagnosis.

Confound removal and temporal signal pre-processing. The strategy we use for cleaning temporal signals is the same as in Dadi et al. (2019). We briefly outline these steps here. We regress out 10 CompCor (Behzadi et al., 2007) components on the whole brain and six motion related signals which are provided in the ADNI, ADNIDOD, COBRE, CamCAN datasets. We do not perform any additional preprocessing steps on ABIDE, ACPI and HCP. For all datasets, the signal of each region is normalized, detrended and bandpass-filtered between 0.01 and 0.1Hz. All these steps are done with *nilearn* (Abraham et al., 2014b).

Appendix C.3. Detailed results

Figure 6 summarizes the impact of the brain atlases and ROIs in predicting diverse targets on rs-fMRI images. Figure A5 shows the absolute prediction scores for each target separately. DiFuMo-based predictions are on par with those using UKBB ICA components, Craddock et al. (2012) and BASC atlases.

Appendix D. Intra-subject encoding

In §3.2, we compare group-level z -maps computed at the voxel-level and on reduced data using the DICE similarity coefficient. We also performed an intra-subject, across sessions, *standard* analysis. We consider the Rapid-Serial-Visual-Presentation (RSVP) language task of Individual Brain Charting (IBC) (see Pinho et al. (2018) for details on experimental protocol and data pre-processing).

Encoding activations across single-subject sessions




Figure A6: **Intra-subject univariate prediction of brain response in the language task protocol of the IBC dataset.** We compare R^2 -maps obtained using voxel based and functional-atlas based encoding models. Encoding models based on high-order atlases better explain the signal of an unseen session. The comparison is made for a single subject; results are similar across subjects.

Encoding model. In this setting, we fit a GLM on the several acquisition sessions of each subject considered separately. That is, we compute a single β -map per session and condition, forming a set of maps $\beta \in \mathbb{R}^{q \times p}$. β is either computed directly at the voxel-level or using functional atlases, in which case we set $\beta = \beta_{red} D^T$, with $\beta \in \mathbb{R}^{q \times k}$.

We then use a leave-one-session-out cross-validation scheme to compare the observed, single-session, time se-

Group-level z -map: complex sentence for more atlases




Figure A7: **Encoding maps with low dimensional brain representation using various high-dimensional atlases.** This figure is complementary to the reported DICE similarity coefficient Figure 3. The displayed group-level maps across atlases are used for computing these overlapping coefficients. The displayed maps correspond to the single condition “subject presented with complex meaningful sentences”. A more detailed stimulus description is provided in [Pinho et al. \(2018\)](#).

ries $Y \in \mathbb{R}^{n \times p}$ to the reconstructed time-series $\hat{Y} = X\bar{\beta}$, where $\bar{\beta}$ are the average β -maps across the 5 training sessions. We obtain R^2 -maps, where each voxel holds the proportion of variance explained by the model

$$r_i = 1 - \frac{\|y_i - \hat{y}_i\|_2^2}{\|y_i - \bar{y}_i\|_2^2},$$

where y_i is the univariate time-series in \mathbb{R}^n associated to voxel i and \bar{y}_i is its temporal mean. We finally average R^2 scores across leave-one-session-out folds, and thresh-

old non-positive values. The resulting R^2 -maps provides information on how much encoded β -maps are able to predict univariate voxel activation on new sessions. A value close to 1 means that the voxel activation is well predicted by the encoding model, while a 0 value means that the voxel activation cannot be predicted. We compare the R^2 -maps across the various data-reduction methods for estimating β .

Validation. To measure the difference between R^2 maps \mathbf{R} computed from voxels and R^2 maps $\tilde{\mathbf{R}}$ computed from DiFuMos, we report correlation coefficients ρ between \mathbf{R} and $\tilde{\mathbf{R}}$, and the slope s predicting the activations $\tilde{\mathbf{R}}$ from the activations \mathbf{R} . This slope indicates a form of signal loss due to using functional atlases. We expect it to be smaller than 1, in part because projection on functional atlases have a noise reduction effect.

Results. As shown in Figure A6, using higher order DiFuMo atlases leads to a loss of explained variance R^2 of only 6% compared to working directly with voxels, which may imputed to a denoising effect. Qualitatively, the R^2 maps are much comparable. DiFuMo ($k = 1024$) is therefore suitable for intra-subject encoding tasks; they make these much less costly. Using lower-order atlases yield stronger signal loss.

Appendix E. Extra statistical and meta-analysis maps

Figure A7 displays the group statistical maps obtained by fitting a GLM on dimension-reduced data, with various atlases. It is complementary to Figure 3.

Figure A8 shows the meta-analysis summary images for two additional cognitive topics: *language* and *face*. We compare non-reduced images with reduced images using DiFuMo ($k = 1024$) and BASC ($k = 444$). The images reduced with DiFuMos are easier to interpret than the ones reduced with BASC for both topics. Quantitatively, we recall that Figure 7 shows the better performance of DiFuMos for image compression.

Appendix F. DiFuMos naming details

A measure of overlap with a reference anatomical atlas allows to match each DiFuMo component with a specific anatomical region, e.g. “postcentral gyrus”. Where there are more than one component for each anatomical region, the functional atlas region are further characterized by an anatomical spatial descriptions, e.g. “postcentral gyrus inferior”. Finally, we append the localisation of the region in the left or right hemisphere, e.g. “postcentral gyrus inferior RH”. Some of the nodes from DiFuMo atlases overlaps a fraction of several regions in the anatomical atlas—those are named by a trained neuroanatomist.

Appendix G. Extracting DiFuMo from resting-state and task data separately

In this section, we evaluate the impact of learning DiFuMos separately on task and rest fMRI data and compare them to using both data types in the training dataset. We separate the training dataset (Table A5) into task fMRI studies and resting-state studies (adding each up to 1970 and 222 fMRI recording sessions respectively). We train DiFuMo atlases (of all dimensions) on each data corpus separately, and compare the performance of the obtained atlases with the atlases trained on both datasets used jointly.

Results. Figure A9 compares the performance of atlases trained on rest, task and rest+task data, for predicting mental states across 6 task fMRI studies (see Table A2 for the details about the studies) and predicting traits from connectomes across 7 different cohorts (see Table A3 for the list of cohorts). Overall, it reveals that using functional modes learnt from both data types – task and rest – marginally improves brain signal extraction.


We do not observe a significant impact of using task-specific modes for predicting traits from functional connectomes. Previous work have established a strong correspondance between resting-state and task functional networks (Smith et al., 2009; Bzdok et al., 2016). This can explain the little difference observed when appending resting-state data to the training corpus.

Appendix H. DiFuMo extraction: implementation details

To run the Stochastic Online Matrix Factorization algorithm, we relied on the open-source Python package readily available at <https://arthurmensch.github.io/modl/>. The MODL package is implemented using Cython (Dalcin et al., 2011), *scikit-learn* (Pedregosa et al., 2011) and *nilearn* (Abraham et al., 2014b).

We use a coordinate descent algorithm (Friedman et al., 2007) to solve for the inner ℓ_1 regularization problems with positivity constraints. Setting optimum λ for DiFuMo extraction is a challenging task as the DiFuMo extraction is an unsupervised learning problem. On one hand the number of modes extracted should approximately cover the whole brain; on the other hand the overlap should be minimal between the modes. As can be seen from Table A4, higher λ gives a lower brain coverage whereas lower value has 100% coverage starting from low-dimensions. Yet lower λ creates overlap between the modes as shown on Figure A10. We set $\lambda = 0.001$ as it provides a good compromise in between modes overlap and brain coverage, as summarized on Figure A11. It gives full brain coverage for high-order atlases. To disambiguate the role of controlling model parameters, we kept λ fixed for all dimensions.

a. Meta-analysis of "language"



b. Meta-analysis of "face"





Figure A8: Meta-analysis on cognitive topics –language (a.) and face (b.) – from statistical images: We compare images reconstructed with DiFuMo ($k = 1024$) and BASC ($k = 444$) with voxel-level averages (top left). The topic-related activations are better visualized using DiFuMo (top middle) than using BASC (top right). DiFuMo results are closer to voxel-level averages, as the signal loss is minimal when projecting on this atlas.

Following Mensch et al. (2018), we access a random fraction of each record at each iteration to accelerate training. We use a subsampling ratio value $r = 12$, run the algorithm on a single epoch, and use a learning rate $\beta = 0.92$, as empirically proposed by Mensch et al. (2016a). We compute a brain mask that intersect the signal from all the resampled images, using Nilearn (Abraham et al., 2014b). DiFuMo atlases are learned using this mask.

The training scripts are available at https://github.com/KamalakerDadi/DiFuMo_analysis_scripts.

A note on spatial overlap among modes. The variants of sparsity control λ (0.01, 0.001, 0.0001) produce modes that have minimum to strong spatial overlap among them within each resolution. We propose modes ($\lambda = 0.001$) that are spatially localized with minimal spatial overlap. At higher resolution, the overlap increases among

a. Predicting mental tasks



b. Predicting traits from functional connectomes




Figure A9: **Impact of the choice of DiFuMo atlases training data on decoding and functional connectomes.** Each point gives the relative prediction score, over 6 different task-fMRI experiments. The thick lines give the median relative score per type of data in task, rest and combination of both. Using the combination of both is marginally better than using task data alone to learn parcellations. We observe no crucial dominance of task-dependent variance in dictionaries in decoding performance.

modes, resulting in higher mutual correlation, as shown on Figure A12. Additionally, we plot the spatial overlap of maps, intraparietal sulcus at resolution $k = 64$ with $\lambda = 0.001$. This is a compromise between keeping the modes strictly disjoint and having them overlapping strongly, which would challenge their very definition.


Dimension	λ	Non-zero voxels (%)
64	0.01	12%
	0.001	59%
	0.0001	100%
128	0.01	21%
	0.001	77%
	0.0001	100%
256	0.01	34%
	0.001	98%
	0.0001	100%
512	0.01	51%
	0.001	100%
	0.0001	100%
1024	0.01	73%
	0.001	100%
	0.0001	100%

Table A4: **Coverage of the whole brain (%) as a function of dimension and range of λ in DiFuMo model.** While the higher dimension 1024 of DiFuMo maps yields a coverage of the whole brain, other DiFuMo atlases with smaller dimensions do not have such coverage. We recommend the optimum trade-off value as $\lambda = 0.001$ by analysing both the degree of overlap between modes as shown on Figure A10 and non-zero voxels. A lower λ gives full coverage of the whole brain but suffers from strong overlap between modes whereas higher λ has lost the brain coverage. The values are reported based on union of all modes per dimension.

fMRI study	Version	Cognitive task	#Subjects	#Sessions	#Runs	Conditions
(Schonberg et al., 2012)	ds000001_R2.0.4	Balloon Analog Risk-taking	16	-	3	balloon analog risk
(Aron et al., 2006)	ds000002_R2.0.5	Classification learning	17	-	2	deterministic classification mixed event related probe probabilistic classification
(Xue and Poldrack, 2007)	ds000003_R2.0.2	Rhyme judgment	13	-	-	rhyme judgment
(Jimura et al., 2014)	ds000006_R2.0.1	ds000006	14	2	6	living nonliving decision- with plain or mirror reversed text
(Xue et al., 2008)	ds000007_R2.0.1	Stop-signal task with spoken & manual responses	20	-	2	stop manual stop vocal stop word
(Aron et al., 2007)	ds000008_R2.0.0	Stop-signal task with unconditional and conditional stopping	14	-	3	conditional stop signal stop signal
(Foerde et al., 2006)	ds000011_R2.0.1	Classification learning and tone counting	14	-	2	Classification probe without feedback Dual task weather prediction Single task weather prediction Tone counting
(Rizk-Jackson et al., 2011)	ds000017_R2.0.1	Classification learning and stop-signal (1 year test-retest)	8	2	3	probabilistic classification selective stop signal task
(Alvarez and Poldrack, 2011)	ds000051_R2.0.2	Cross-language repetition priming	13	-	8	abstract concrete judgment
(Poldrack et al., 2001)	ds000052_R2.0.0	Classification learning and reversal	14	-	2	weather prediction reversal weather prediction
(Mennes et al., 2013)	ds000101_R2.0.0	Simon task	21	-	2	simon
(Kelly et al., 2008)	ds000102_R2.0.0	Flanker task (event-related)	26	-	2	flanker
(Haxby et al., 2001) (O'Toole et al., 2005) (Hanson et al., 2004)	ds000105_R2.0.2	Visual object recognition	6	-	12	object viewing
(Duncan et al., 2009)	ds000107_R2.0.2	Word and object processing	49	-	2	1-back task
(Moran et al., 2012)	ds000109_R2.0.2	False belief task	36	-	2	theory of mind
(Uncapher et al., 2011)	ds000110_R2.0.1	Incidental encoding task (Posner Cueing Paradigm)	18	-	10	Incidental encoding task
(Gorgolewski et al., 2013)	ds000114_R2.0.1	A test-retest fMRI dataset for motor, language and spatial attention functions	10	2	-	covert verb generation finger footlips line bisection overt verb generation overt word generation
(Repovs and Barch, 2012)	ds000115_R2.0.0	Working memory in healthy and schizophrenic individuals	1	-	-	letter 0-back task letter 1-back task letter 2-back task
(Cera et al., 2014)	ds000133_R1.0.0	Modafinil alters intrinsic functional connectivity of the right posterior insula: a pharmacological resting state fMRI study	26	2	3	rest
(Verstynen, 2014)	ds000164_R1.0.1	Stroop task	28	-	-	stroop
(Gabitov et al., 2015)	ds000170_R1.0.1	Learning and memory: motor skill consolidation and intermanual transfer	15	-	3	Trained Hand Trained Sequence Trained Hand Untrained Sequence Untrained Hand Trained Sequence
(Lepping et al., 2016a) (Lepping et al., 2016b)	ds000171_R1.0.0	Neural Processing of Emotional Musical and Nonmusical Stimuli in Depression	39	-	5	music non music
(Iannilli et al., 2016)	ds000200_R1.0.0	Pre-adolescents Exposure to Manganese	1	-	-	olfactory
(Stephan-Otto et al., 2017)	ds000203_R1.0.2	Visual imagery and false memory for pictures	26	-	2	visual imagery- false memory
(Kim et al., 2016)	ds000205_R1.0.0	Affective Videos	11	-	2	functional localizer view
(Romaniuk et al., 2016)	ds000214_R1.0.0	EUPD Cyberball	40	-	-	Cyberball
(Roy et al., 2017)	ds000220_R1.0.0	Cost Analysis TBI	26	2	-	rest

Table A5: Large-scale fMRI datasets downloaded from OpenNeuro to build our multi-scale functional atlases. Data are pre-processed using *fMRIprep*. The data acquisition parameters of each study are listed on [Table A6](#). The corpus is 2.4TB in total.

Figure A10: **Overlap between maps with $\lambda = 0.0001$ in DiFuMo model**, for DiFuMo dimensionality 64, 256, and 512. As dimensionality increases, subdivisions are strongly overlapped lacking clear delineation across brain regions. [Table A4](#) gives details about the percentage of non-zero voxels covering whole brain for $\lambda = 0.01, 0.001, 0.0001$



Quantification of overlaps for all DiFuMo modes




Figure A11: **Comparison of modes overlap for all proposed DiFuMo atlases.** The y-axis shows how many voxels are at the intersect of exactly n modes (x-axis), without thresholding the modes. On average, at least two modes are shared between voxels. Coverage (i.e. low number of voxel at the intersect of 0 mode) is higher for finer grain atlases, at the cost of larger overlap. Note that the overlaps often comprises voxels with small values, and may thus be weak.

Figure A12: **Spatial overlap and correlation among modes for the proposed DiFuMos.** **Left:** Qualitative visualization of spatial overlap among modes showing Intraparietal sulcus shared as networks with other modes (Left and right hemisphere). The overlap is shown with $k = 64$. **Right:** Pairwise correlation among modes within each resolution of proposed DiFuMo atlases. The overlap gets stronger with the increase in the resolution.





Figure A13: **Interpretation of higher-dimensional modes of DiFuMo:** The putamen segmentation is refined as dimension of DiFuMo increases. A single mode contains the left and right putamen in lower dimension (a), when higher order atlases hold separate components for them. Larger atlases model the detailed organization within the sub-structures, which may be crucial in discriminative tasks.




Figure A14: **Modes around the thalamus**, for DiFuMo dimensionality 64, 256, and 512. As dimensionality increases: sub-divisions are more refined, modes are split into middle, lateral, inferior, superior, anterior, posterior direction and left hemisphere. Each color represents a single mode.

fMRI study	MR scanner	Slice orientation	FoV (mm)	Voxel size (mm)	Matrix size	TR (msec)	TE (msec)	Flip angle (°)	Number of volumes
(Schonberg et al., 2012)	3T Siemens AG Allegra (Erlangen, Germany)	axial	-	4 × 4 × 4	64 × 64	2000	30	90	300
(Aron et al., 2006)	3T Siemens Allegra	-	-	4 × 4 × 4	64 × 64	2000	30	90	180
(Xue and Poldrack, 2007)	3T Siemens Allegra (Iselin, NJ)	-	200	4 × 4 × 4	64 × 64	2000	30	90	160
(Jimura et al., 2014)	3T Siemens Allegra (Erlangen, Germany)	-	200	4 × 4 × 4	64 × 64	2000	30	90	205
(Xue et al., 2008)	3T Siemens Allegra	-	200	4 × 4 × 4	64 × 64	2000	30	90	182
(Aron et al., 2007)	3T Siemens Allegra	-	200	4 × 4 × 4	64 × 64	2000	30	90	176
(Foerde et al., 2006)	3T Siemens Allegra	-	200	4 × 4 × 4	64 × 64	2000	30	-	208
(Poldrack et al., 2001)	3T Siemens Allegra	axial	200	5 × 5 × 5	64 × 64	3000	30	-	225
(Mennes et al., 2013)	3T Siemens Allegra	-	192	3 × 3 × 4	64 × 64	2000	30	80	101
(Kelly et al., 2008)	3T Siemens Allegra	-	192	3 × 3 × 4	64 × 64	2000	30	80	146
(Haxby et al., 2001)	3T GE	sagittal	240	3.5 × 3.5 × 3.5	-	2500	30	90	121
(Duncan et al., 2009)	1.5T Siemens	-	192	3 × 3 × 3	64 × 64	3000	50	-	165
(Moran et al., 2012)	3T Siemens Tim Trio	axial	-	3 × 3 × 3	-	2000	35	-	179
(Uncapher et al., 2011)	3T GE Signa	axial	-	3.44 × 3.44 × 3.44	64 × 64	2000	30	75	186
(Gorgolewski et al., 2013)	1.5T GE Signa	-	256	4 × 4 × 4	64 × 64	2500	50	90	varied
(Repovs and Barch, 2012)	3T Tim Trio	-	256	4 × 4 × 4	64 × 64	2500	27	90	137
(Cera et al., 2014)	3T Philips	transaxial	256	4 × 4 × 4	64 × 64	1671	35	75	145
(Verstynen, 2014)	3T Siemens	-	-	3.2 × 3.2 × 4	-	1500	20	90	370
(Gabitov et al., 2015)	3T GE	axial	220	3.4 × 3.4 × 3.4	64 × 64	3000	35	90	45
(Lepping et al., 2016a)	3T Siemens Skyra (Erlangen, Germany)	axial	220	2.9 × 2.9 × 3	64 × 64	3000	25	90	105
(Iannilli et al., 2016)	1.5T Siemens Aera (Erlangen, Germany)	axial	-	3.3 × 3.3 × 3.3	64 × 64	2500	50	-	120
(Stephan-Otto et al., 2017)	1.5T GE Signa	axial	240	4 × 4 × 4	64 × 64	2000	40	90	267
(Kim et al., 2016)	3T Siemens Trio (Erlangen)	axial	-	3 × 3 × 3	64 × 64	2200	35	90	365
(Romaniuk et al., 2016)	3T Siemens Magnetom Verio	axial	220	3.4 × 3.4 × 5	64 × 64	1560	26	66	341
(Roy et al., 2017)	3T Philips Achieva	axial	240	4 × 4 × 4	80 × 80	2000	30	90	144

Table A6: Data acquisition parameters for each fMRI study that we use for training DiFuMo atlases. Data are downloaded from OpenNeuro.