Multi-modal brain fingerprinting: a manifold approximation based framework

Kuldeep Kumar[a,c,*1], Laurent Chauvin[a], Matthew Toews[a], Olivier Colliot[b,c,d], Christian Desrosiers[a]

[a] *Laboratory for Imagery, Vision and Artificial Intelligence, École de technologie supérieure, 1100 Notre-Dame W., Montreal, QC, Canada, H3C1K3*
[b] *Sorbonne Universités, UPMC Univ Paris 06, Inserm, CNRS, Institut du cerveau et la moelle (ICM) - Hôpital Pitié-Salpêtrière, Boulevard de l'hôpital, F-75013, Paris, France*
[c] *Inria Paris, Aramis project-team, 75013, Paris, France*
[d] *AP-HP, Departments of Neurology and Neuroradiology, Hôpital Pitié-Salpêtrière, 75013, Paris, France*

Abstract

In this work, we propose a simple and computationally efficient framework for multi-modal brain fingerprinting analysis from large datasets. The key idea is to represent each image as a bag of local features, and obtain a subject proximity graph using feature matches (over the set of all the images). This subject proximity graph, which is an approximation of the image appearance manifold, can be used to obtain a compact fingerprint representation using concepts of manifold embedding, and also allows out of sample extensions. The proposed framework can work with any local feature descriptors, and is used to compare and contrast MRI modalities and their combination for the task of identification of genetically related subjects. Experiments using the T1/T2-weighted MRI and diffusion MRI data of 1010 Human Connectome Project subjects demonstrate strong links between the proposed manifold approximation and genetic proximity, with a multi-modal brain fingerprint being more discriminative than individual modalities. Single/multi-modal fingerprints for Mono-zygotic twins are found to be more similar than dizygotic or non-twin siblings in a rank retrieval analysis. This is also observed in terms of number of feature matches between twin/sibling pairs, in major brain structures across both hemispheres. The ro-

[1] kkumar@livia.etsmtl.ca

bustness of the proposed framework for image alignment and scan resolution, with high reproducibility using retest scans, suggests the promising future application of multi-modal brain fingerprinting for characterization of individuals in large cohort analysis.

Keywords: Brain fingerprint, Multi-modal, Bag-of-Features, Manifold, Structural MRI, Diffusion MRI, HCP Twin data, Hemisphere asymmetry

1. Introduction

Despite sharing gross similarities, individual brains show a significant amount of variability [1, 2, 3, 4] in terms of structure [5], function [6, 7, 8], and white matter architecture [9, 10]. Recently, various studies have focused on characterizing this variability using brain *fingerprints*, for instance, based on shape [11], functional connectivity [12, 13], white matter fiber geometry [14] or voxel-wise diffusion density [15]. These studies are motivated by the fact that brain characteristics are largely determined by genetic factors that are often unique to individuals [16]. Moreover, various neurological disorders like Parkinson [17] and autism [18] have been linked to specific brain abnormalities that are difficult to describe at the population level. With the rapid improvements in MRI acquisition hardware and analysis tools, and thanks to large brain-related initiatives like the Human Connectome Project (HCP) [19, 20, 21], researchers are better poised to study individual subjects in addition to group level analysis [22, 7, 23, 24], thus taking a step towards precision medicine [25, 26, 27] or precision psychiatry [28, 29, 30, 24].

The importance of brain fingerprinting is evident from the recent surge in studies on this topic. For example, Yeh et al. [15] built a local connectome fingerprint using dMRI data, and applied this fingerprint to the analysis of genetically-related subjects. In [14], Kumar et al. proposed another dMRI-based fingerprint called *Fiberprint*, which characterizes white matter fiber geometry. In [12], Finn et al. considered the correlation between time courses of atlas-defined nodes to generate a functional connectivity profile, and used

this profile to identify individuals across scan sessions, both for task and rest conditions. Moreover, Miranda et al. [31] proposed a linear model to describe the activity of brain regions in resting-state fMRI as a weighted sum of its functional neighboring regions. Their functional fingerprint, derived from the model's coefficients, has the ability to predict individuals using a limited number of non-sequential frames.

Various morphometry-based fingerprints have also been proposed for structural MRI modalities like T1- or T2-weighted images. In [11, 32], Wachinger et al. quantify the shape of cortical and subcortical structures via the spectrum of the Laplace-Beltrami operator. The resulting representation, called *Brainprint*, is used for subject identification and analyzing potential genetic influences on brain morphology. In [33], Toews et al. represent images as a collection of localized image descriptors, and apply scale-space theory to analyze their distribution at the characteristic scale of underlying anatomical structures. This representation is employed to identify distinctive anatomical patterns of genetically-related individuals or subjects with a known brain disease.

So far, fingerprinting studies in the literature have focused on a single modality. However, each modality captures unique properties of the brain and combining multiple modalities can provide a richer, more discriminative information [34, 35, 36, 37, 38, 39, 40]. Hence, the fusion of multiple modalities has been shown superior than single-modality data to identify diseases like schizophrenia, bipolar disorder, major depressive disorder and obsessive-compulsive disorder [35]. Multi-modal neuroimaging biomarkers have also been proposed to predict cognitive deficits in schizophrenia [41]. In [42], the combination of multiple MRI modalities has led to the improved segmentation of isointense infant brain images. In [36], multimodal imaging data is used to predict the brain-age of subjects and detect cognitive impairments. Detailed reviews on multi-modal methods and investigations for psychopathology can be found in [35, 37, 38, 43, 39].

Due to the challenges of combining multiple modalities in a single framework [35, 38], defining a multi-modal brain fingerprinting remains to this day an elusive task. Morphometry-based approaches, such as Brainprint [11], could

potentially be extended to other modalities like dMRI. However, this requires solving non-trivial problems such as the cross-modality alignment of images with different resolutions, the segmentation and correspondence of neuroanatomical structures, etc. Computational efficiency is another important issue when dealing with large-scale, multi-subject and multi-modal datasets like the Human Connectome Project (HCP) [21] and UK Biobank [20]. In this work, we propose a multi-modal brain fingerprinting that overcomes these challenges using manifold approximation. The key idea is to represent each image as a bag of local features, and derive a subject-level proximity graph using feature matches over the entire set of images [33]. This subject proximity graph provides an approximation of the image appearance subspace (i.e., the manifold), which can be used to obtain a compact fingerprint representation.

Manifold learning has been extensively studied in machine learning [44, 45] with many approaches like Isomap [46], Locally Linear Embedding (LLE) [47], Spectral Embedding [48] and Multi-dimensional Scaling (MDS) [49] proposed over the years. As detailed in [50], such techniques have also been used for various problems of medical imaging like registration, segmentation and classification. For example, in [51], Gerber et al. use manifold learning to perform a population analysis of brain images. Similarly, a deep learning based approach is explored in [52] to learn the manifold of brain MRIs. A key factor in such methods is image representation. For instance, the manifold could be approximated using the Euclidean distance between image pairs, however this would not be robust to translation, rotation or scaling, and would suffer from high computational costs. Representations based on local features, often referred to as bag of features (BoF) representations, have been shown to be superior in terms of robustness and matching distinctiveness [33].

BoF representations play a key role in various problems of computer vision, for example, object recognition [53, 54, 55] and image retrieval [56]. This technique can be seen as a way of compressing full images using a few discriminative local features, which can then be matched in sublinear time, for example, using randomized KD-search trees [57]. With respect to brain imaging, BoFs have

been used for morphometry analysis [33, 58], modeling the development of infant brains [59], and image alignment [60]. While they have shown great potential for computer vision and medical imaging, BoFs have, thus far, not been explored for brain fingerprinting. This is mainly due to the fact that BoF representations can have a large and variable size, which makes comparing two fingerprints non-trivial. In this work, this problem is circumvented by embedding the BoF representations in a low-dimensional manifold.

Building on preliminary work [61, 62], we propose an efficient approach based on BoFs and manifold approximation that combines the information from multiple imaging modalities into a single compact fingerprint. The current study presents a comprehensive analysis of the proposed fingerprint using a large-scale dataset from the Human Connectome Project (HCP), where numerous properties/factors are investigated, including the impact of various fingerprint parameters (e.g., manifold dimensionality and proximity graph connectivity), the contribution of individual modalities and/or their combination to the fingerprint's discriminativeness, the fingerprint's robustness to image alignment and scan resolution, and the reproducibility of results using re-test or corrupted scans. Using the HCP twin dataset, we analyze the proposed fingerprint's ability to identify genetically-related subjects (i.e., monozygotic twins, dizygotic twins and non-twin siblings) from a large cohort, and show our multi-modal fingerprint to outperform single-modality approaches or fingerprints based on raw images. In an analysis of local feature correspondences, we identify for each modality the neuroanatomical regions having the most significant differences across groups of genetically-related subjects, as well as between males and females. Lateral asymmetry is also considered in this analysis by comparing the distribution of features matches across hemispheres. To our knowledge, this study constitutes the most in-depth investigation of a multi-modal brain fingerprint.

The rest of this paper is organized as follows. We first present the proposed multi-modal brain fingerprinting framework, detailing the data pre-processing steps, the BoF representation and proximity graph computation, and the mani-

fold embedding of this graph. In Section 3, we then conduct an extensive experimental validation using the T1-weighted, T2-weighted, and diffusion-weighted MRI data of 1010 subjects from the HCP dataset. Finally, we conclude with a summary of our contributions and a discussion of possible extensions.

2. Materials and methods

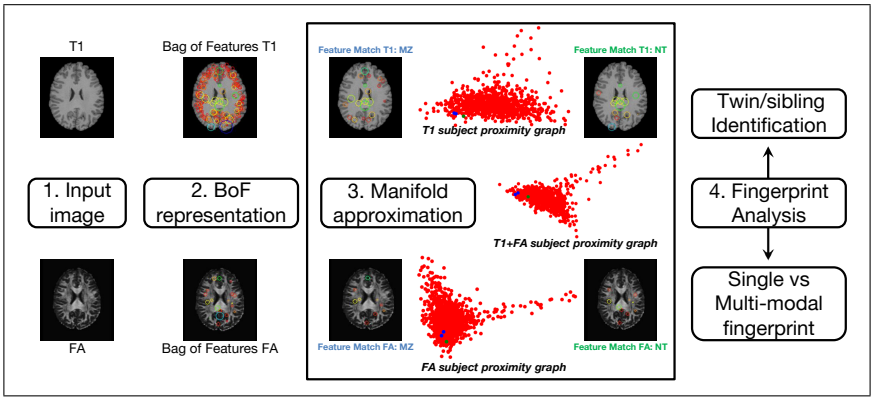

Figure 1: Pipeline of the proposed framework and analysis. For a given input image, we obtain a bag of features representation. Manifold approximation is then obtained in the form of a subject proximity graph using an approximate feature correspondence. Followed by analysis using twin data from Human Connectome Project. (Manifold approximation block also shows pairwise feature matches for a MZ twin pair and MZ-NT sibling pair, represented by blue and green dots in a 2-D spectral embedding representation of subject proximity graph for T1, FA, and T1+FA)

Figure 1 summarizes the pipeline of the proposed multi-modal brain fingerprint framework, which is comprised of four steps. In the first step, we start with pre-processed structural MRI (sMRI) and diffusion MRI (dMRI) data of 1010 subjects from the Human Connectome Project [63, 21]. Four diffusivity measures are obtained from dMRI scans: fractional anisotropy (FA), axial diffusivity (AD), mean diffusivity (MD), and radial diffusivity (RD). The second step then extracts local features from the images of each subject, and encodes subjects as a Bag of Features (BoF). In the third step, the multi-modal fingerprints of subjects are computed using manifold approximation. Towards this goal,

a subject-level proximity graph is first constructed by matching the features of each modality across images, and identifying pairs of subjects with a high number of matches. Fingerprints are then obtained by embedding this graph in a low-dimensional subspace. In the last step, we perform various analyses on the subject fingerprints. The informativeness of individual modalities and their link to genetic proximity are first measured in twin/sibling identification analysis. This analysis is then extended to multi-modal fingerprints, showing the combined effect and complementarity of multiple modalities. Finally, the distribution of feature matches between pairs of subjects are used to identify regions showing significant differences across different sibling types. The following subsections describe each of these steps in greater detail.

2.1. Data and pre-processing

We used the pre-processed structural and diffusion MRI data of 1010 subjects, as well as the retest data of 42 subjects, from the HCP1200 release of the Human Connectome Project [21]. The HCP1200 release provides genetically verified labels for twins and siblings, and is a rich resource to analyze the importance of environmental and genetic influences for traits, phenotypes, and disorders [64, 65, 66, 67, 63]. Table 1 provides the demographic details of the subjects used in this study.

Data were acquired on a Siemens Skyra 3T scanner [68, 69] and pre-processed as described in [70]. The structural acquisitions include high resolution T1-weighted (T1w) and T2-weighted (T2w) images (0.7 mm isotropic, FOV = 224 mm, matrix = 320, 256 sagittal slices in a single slab), and the diffusion acquisition used following parameters: sequence = Spin-echo EPI; repetition time (TR) = 5520 ms; echo time (TE) = 89.5 ms; resolution = $1.25 \times 1.25 \times 1.25$ mm^3 voxels. Further details can be obtained from HCP1200 data release manual[2]. We used the `hcp2blocks.m` script (described in HCP1200 release) to generate a FamilyID based matrix, only considering subjects having dMRI data

[2] `https://www.humanconnectome.org/documentation/S1200/`

and for which the `HasGT` field is true. Using this selection criterion, we obtained a total of 134 mono-zygotic (MZ) twin pairs, 72 di-zygotic (DZ) twin pairs, and 502 non-twin (NT) sibling pairs. No subject had more than five siblings.

Table 1: Demographics. We have used HCP1200 release subjects with both diffusion MRI and structural MRI data, and HasGT (genetically verified data) field to be true. The family structure and links are obtained from the output of hcp2blocks.m script listed in data release manual. Note: each NT subject has $1 \leq$ Num siblings ≤ 5.

Type	**Total**	**Gender**		**Age**	**Handedness**	
		F	M	Range (median)	L	R
All	1010	542	468	22-36 (29)	95	915
MZ	268	158	110	22-36 (29)	23	245
DZ	144	84	60	22-35 (29)	15	129
NT	502	251	251	22-36 (29)	44	458

For structural MRI we considered T1-weighted (0.7 mm) and T2-weighted (0.7 mm). In the case of dMRI data, signal reconstruction was performed with the freely available DSI Studio toolbox [71] using the DTI reconstruction option. Four widely used diffusivity measures were extracted to characterize white matter micro-structure: fractional anisotropy (FA) [72, 73], axial diffusivity (AD), mean diffusivity (MD) and radial diffusivity (RD). The interpretation of these measures are discussed in [34, 74]. For analyzing the impact of alignment, we also used the MNI space aligned data for T1-weighted (0.7 mm) and T2-weighted (0.7 mm) provided by HCP 1200 release. In addition, to combine structural modalities with dMRI, and to analyze impact of scan resolution, we re-sampled T1- and T2-weighted images to a 1.25 mm resolution, using the linear registration (FLIRT) tool of FSL [75].

2.2. Multi-modal brain fingerprint

Generating brain fingerprints of subjects, based on their multi-modal data, involves multiple steps: extracting local descriptors in images to build a BoF

representation of subjects, building a subject proximity graph by comparing their BoF representations, and embedding this graph in a low-dimensional manifold. Additionally, once the manifold has been constructed, an out-of-sample extension strategy is required to compute the fingerprint of new subjects.

2.2.1. Bag of feature (BoF) representation of subjects

In the first step, a set of local descriptors [76, 54] is obtained from each available image (3D scan). Various local feature extraction and representation approaches [77] can be used, for example, Scale Invariant Feature Transfrom (SIFT) [53], Generalized Robust Invariant Feature (G-RIF) [78] or Speeded UP Robust Features (SURF) [79, 80]. In this work, we use 3-D SIFT descriptors, as they have been well studied for neuro-image analysis [58, 33, 60, 81] and can be computed efficiently.

3D keypoints are located in the scans of each subject by finding the local extrema (i.e., maxima or minima) of the difference of Gaussians (DoG) occurring at multiple scales. Keypoints with a low contrast or corresponding to edge response are discarded, and remaining ones are encoded into a feature vector (i.e, the descriptor) using the histogram of oriented gradients (HOG) within a small neighborhood. Note that these descriptors are robust to changes in illumination, scale and rotation, and are thus efficient for comparing images acquired using different scanners or imaging parameters. Each subject is then represented as an orderless bag of features (BoF), containing all the descriptors found in this subject's scans. This representation provides a simple, robust and extensible way of incorporating data from multiple modalities.

2.2.2. Subject proximity graph

Because the BoFs of two subjects may contain different numbers of descriptors, they cannot be directly compared. To circumvent this problem, we construct an intrinsic manifold of subject appearance using a feature-to-feature nearest-neighbor (NN) graph. In this graph, each descriptor is represented by a node and is connected to its K most similar descriptors based on Euclidean dis-

tance. This feature-to-feature graph is then converted to a subject-to-subject (i.e., subject proximity) graph by considering, for each pair of subjects, the number descriptors in their BoF that are linked in the feature-to-feature graph.

Let $\mathcal{B}_i^m$ and $\mathcal{B}_j^m$ be the BoFs (i.e., set of descriptors) of subjects i and j for modality $m \in \mathcal{M}$, where $\mathcal{M}$ is the set of available modalities. The similarity between these subjects is evaluated as

$$s_{ij} \;=\; \frac{\sum_{m\in\mathcal{M}} |\mathcal{B}_i^m \cap \mathcal{B}_j^m|}{\sum_{m\in\mathcal{M}} |\mathcal{B}_i^m \cup \mathcal{B}_j^m|} \;=\; \frac{\sum_{m\in\mathcal{M}} |\mathcal{B}_i^m \cap \mathcal{B}_j^m|}{\sum_{m\in\mathcal{M}} \left(|\mathcal{B}_i^m| \,+\, |\mathcal{B}_j^m| - |\mathcal{B}_i^m \cap \mathcal{B}_j^m|\right)}, \qquad (1)$$

where $|\mathcal{B}_i^m \cap \mathcal{B}_j^m|$ is the number of edges in the feature-to-feature graph between nodes in $\mathcal{B}_i^m$ and those in $\mathcal{B}_j^m$. When using a single modality, this measure corresponds to the well-known Jaccard similarity. Here, we extend it to a multi-modal setting by comparing the descriptors of each modality separately. We note that the Jaccard distance, defined as one minus the Jaccard similarity, is a *metric* and thus well-suited for constructing the manifold.

When defining the feature-to-feature graph, K determines the number of nearest-neighbor connections for each descriptor. In manifold learning approaches, this parameters controls the locality of the manifold approximation at each point [44]. Its value should be large enough to capture the manifold's local structure, but also restricted so that distances to nearest-neighbors are close to the geodesic. In our experiments, we tested $K \in \{10, 20, \ldots, 50\}$ and found similar results for these values. In what follows, we report results obtained with $K = 20$.

2.2.3. Manifold embedding

A manifold embedding technique is used to obtain compact brain fingerprints from the subject proximity graph. While various approaches could be employed for this task, for instance Isomap [46], locally linear embedding (LLE) [47] and multi-dimensional scaling (MDS) [49], we performed the embedding using Laplacian eigenmaps [48]. This technique, which is connected to the well-known Laplace-Beltrami operator, has the advantage of being efficient and allowing out-of-sample extensions.

In Laplacian eigenmaps, each subject i is mapped to a coordinate vector $\mathbf{x}_i \in \mathbb{R}^k$ of the manifold, whose dimension k is a user parameter. The embedding of subjects in the manifold is made such that two subjects i and j with a high similarity s_{ij} will be close to one another. Let $\mathbf{S} \in \mathbb{R}^{n \times n}$ be the adjacency matrix of the subject proximity graph, as defined in Eq (1), and denote as $\mathbf{L} = \mathbf{D} - \mathbf{S}$ the Laplacian of $\mathbf{S}$, where $\mathbf{D}$ is a diagonal matrix containing the row sums of $\mathbf{S}$. The embedding is accomplished by solving the following problem:

$$\arg\min_{\mathbf{X}} \sum_{i=1}^{n}\sum_{j=1}^{n} s_{ij}\|\mathbf{x}_i - \mathbf{x}_j\|_2^2 = \mathrm{tr}(\mathbf{X}^\intercal \mathbf{L}\mathbf{X}), \quad \text{s.t. } \mathbf{X}^\intercal\mathbf{D}\mathbf{X} = \mathbf{I}. \qquad (2)$$

The constrain on $\mathbf{X}$ removes arbitrary scaling factor in the embedding. As described in [48], the solution to this problem is given by the leading k eigenvectors of the normalized adjacency matrix $\overline{\mathbf{S}} = \mathbf{D}^{-\frac{1}{2}}\mathbf{S}\mathbf{D}^{-\frac{1}{2}}$, starting from the second one[3]. Once computed, the rows of matrix $\mathbf{X}$ correspond to the n subject fingerprints of size k.

2.2.4. Out-of-sample extension

The manifold embedding technique described above computes the fingerprint of all subjects at once. If new subjects are added, this process must be repeated over again, which is inefficient and changes the fingerprint of previous subjects. To alleviate these problems, we use an out-of-sample extension of Laplacian eigenmaps, based on the Nystrom method [82, 83, 84].

Suppose we want to compute the manifold embedding of m new subjects. The first step is to update the nearest-neighbor feature graph with the local descriptors of these new subjects, leaving unchanged the nearest-neighbors of the n base subjects. We then evaluate the pairwise similarities between new subjects and the base ones. Let $\mathbf{P} \in \mathbb{R}^{n \times m}$ be the matrix containing these similarities, the adjacency matrix of the extended subject proximity graph

[3] The first eigenvector contains constant values.

$\mathbf{S}_{\text{ext}} \in \mathbb{R}^{(n+m)\times(n+m)}$ is given by

$$\mathbf{S}_{\text{ext}} = \begin{bmatrix} \mathbf{S} & \mathbf{P} \\ \mathbf{P}^\intercal & \mathbf{Q} \end{bmatrix}. \tag{3}$$

Using the formula in [48], the matrix $\mathbf{Q}$ of similarities between new subjects can be approximated as $\mathbf{P}^\intercal\mathbf{S}^{-1}\mathbf{P}$.

To normalize $\mathbf{S}_{\text{ext}}$, we compute the vector of row sums

$$\mathbf{d}_{\text{ext}} = \begin{bmatrix} \mathbf{s}_r + \mathbf{p}_r \\ \mathbf{p}_c + \mathbf{P}^\intercal\mathbf{S}^{-1}\mathbf{p}_r \end{bmatrix}, \tag{4}$$

where $\mathbf{s}_r, \mathbf{p}_r \in \mathbb{R}^n$ contain the row sums of $\mathbf{S}$ and $\mathbf{P}$, respectively, and $\mathbf{p}_c \in \mathbb{R}^m$ contains the column sum of $\mathbf{P}$. In the case where m is small compared to n, we have that $\mathbf{s}_r \approx \mathbf{s}_r + \mathbf{p}_r$, and thus $\mathbf{d}_{\text{ext}}$ can be approximated as

$$\mathbf{d}_{\text{ext}} \approx \begin{bmatrix} \mathbf{s}_r \\ \mathbf{p}_c + \mathbf{P}^\intercal\mathbf{S}^{-1}\mathbf{p}_r \end{bmatrix}. \tag{5}$$

This strategy, used in [84] for white matter fiber segmentation, allows preserving the previous embedding of base subjects. Let $\mathbf{D}_{\text{ext}}$ be the diagonal matrix with entries corresponding to $\mathbf{d}_{\text{ext}}$, the normalized adjacency matrix of the extended subject graph is calculated as $\overline{\mathbf{S}}_{\text{ext}} = \mathbf{D}_{\text{ext}}^{-\frac{1}{2}}\,\mathbf{S}_{\text{ext}}\,\mathbf{D}_{\text{ext}}^{-\frac{1}{2}}$. The extended embedding is then obtained following Nystrom's method as

$$\mathbf{X}_{\text{ext}} = \begin{bmatrix} \mathbf{U} \\ \overline{\mathbf{P}}^\intercal\mathbf{U}\Lambda^{-1} \end{bmatrix}, \tag{6}$$

where $\mathbf{U}\Lambda\mathbf{U}^\intercal$ is the eigendecomposition of $\overline{\mathbf{S}}$, and $\overline{\mathbf{P}}$ is the normalized submatrix in $\overline{\mathbf{S}}_{\text{ext}}$ corresponding to $\mathbf{P}$. Hence, the embedding of base subjects is the same as in Section 2.2.3, and new subjects are embedded as $\overline{\mathbf{P}}^\intercal\mathbf{U}\Lambda^{-1}$. Once more, a fingerprint of size k is obtained by considering only the k leading eigenvectors in matrix $\mathbf{U}$, ignoring the constant eigenvector.

2.3. Computational efficiency

Computational and memory requirements are key factors when performing large scale analyses. In this section, we evaluate these requirements for the main

steps of the proposed framework. To highlight the efficiency of encoding images with local descriptors, we also compare our framework to a simple fingerprint using full images as features. Other aspects like scan resolution and image alignment requirements are discussed in Section 3. All experiments were performed on a 3.6 GHz processor with 32 GB RAM.

For the BoF representation of images, we extracted 3D SIFT features using a publicly available tool[4]. Computing these features took about 3 seconds per image, and approximately 60 minutes for all 1010 images, when processed sequentially. This runtime could however be reduced significantly by processing images in parallel. The feature matching routine [57], for generating the subject proximity graph from the BoF representations of all images, required around 5 minutes to complete. In comparison, calculating the sum of squared distances (SSD) between full images took 1.7 seconds on average for a single pair, and 870,000 seconds for all $(1010 \times 1009)/2 = 509{,}545$ pairs (with parallel computations). In terms of memory, each BoF file is approximately 400 KB in size, compared to 84 MB on average for a NIfTI volume file. In summary, the proposed framework is highly efficient in terms of computational and memory requirements compared to a baseline fingerprint using full images. Moreover, since computing the subject proximity graph has a complexity in $O(N \log N)$ where N is the number of images, and because extending the manifold embedding can be done efficiently using the Nystrom method, our framework is scalable to large datasets.

2.4. Evaluation measures

To measure the link between fingerprint similarity and genetic proximity, we performed a rank retrieval analysis using the sibling information provided in the HCP dataset. In this analysis, we try to identify the twins/siblings of a given subject by comparing its fingerprint with that of all other subjects in the group. Another goal of this analysis is to provide a common platform

[4] `http://www.matthewtoews.com/`

for the quantitative comparison of individual modalities and their combination. Two standard performance metrics for rank retrieval are used to evaluate the fingerprints: mean recall@k and mean average precision (MAP) [85, 86].

Given a subject i, we rank all other subjects by the similarity (i.e., inverse of Euclidean distance) of their fingerprint to that of subject i. Denote as $\mathcal{T}_i$ the set of target siblings/twins of subject i. For instance, if the target group is non-twin sinblings (NT), then $\mathcal{T}_i$ contains the siblings of subject i that are not his/her twin. Moreover, let $\mathcal{S}_i^k$ be the set containing the k subjects with fingerprints most similar to that of i (i.e., the k nearest neighbors of i). For a given value of k, we evaluate the retrieval performance using the measures of recall@k and precision@k:

$$(\text{recall@}k)_i = \frac{|\mathcal{T}_i \cap \mathcal{S}_i^k|}{|\mathcal{T}_i|}, \quad (\text{precision@}k)_i = \frac{|\mathcal{T}_i \cap \mathcal{S}_i^k|}{k}. \tag{7}$$

When analyzing the rank performance for a particular sibling type (i.e., monozygotic twin, dizygotic twin or non-twin sibling), we average values over the set of subjects which have at least one sibling of this type, i.e. the set $\{i,\ \text{s.t.}\ |\mathcal{T}_i| > 0\}$.

We also evaluate performance with the average precision, which extends the above metrics by considering the rank of nearest neighbors:

$$\text{AveP}_i = \frac{1}{|\mathcal{T}_i|} \sum_{k=1}^{n} (\text{precision@}k)_i \times \text{rel}_i(k), \tag{8}$$

where $\text{rel}_i(k)$ is an indicator function with value equal to 1 if the k-th nearest neighbor of i is relevant (i.e., is in $\mathcal{T}_i$), and zero otherwise. The MAP is obtained by averaging AveP values over all subjects having at least one sibling of the target type.

Finally, we use the d-prime sensitivity index [87] to obtain a quantitative measure of separability between the distribution of fingerprint distances corresponding to different sibling types. Let μ_1, μ_2 and σ_1, σ_2 be the means and standard deviations of compared distance distributions (e.g., distance between monozygotic twins versus between dizygotic twins). The d-prime index is com-

puted as

$$\text{d-prime} = \frac{\mu_1 - \mu_2}{\sqrt{\frac{1}{2}\left(\sigma_1^2 + \sigma_2^2\right)}}. \quad (9)$$

In our experiments, we report absolute values of d-prime, higher values indicating better separability.

3. Experiments and results

A comprehensive set of experiments was conducted to analyze the proposed fingerprint and evaluate its usefulness in various applications. In the first experiment, we analyze the manifold embedding of subjects and measure the impact of manifold dimensionality on the fingerprint's ability to capture genetic proximity. We then perform a detailed rank retrieval analysis, in which fingerprints obtained from a single modality or combinations of multiple modalities are used to identify three types of genetically-related subject: monozygotic twins (MZ), dizygotic twins (DZ) and non-twin siblings (NT). The driving hypothesis of this experiment is that individual modalities capture distinct properties of brain tissues, which can be effectively encoded in the fingerprint, and that combining different modalities can help describe the uniqueness of individual brains. Another goal of this experiment is to measure the relationship between the similarity of fingerprints, for different modality combinations, and genetic proximity.

In another experiment, we assess the impact of following factors on the proposed fingerprint: image alignment, scan resolution, inclusion of skull, and subject age. This is followed by a reproducibility analysis, performed with the restest scans of 42 subjects, and a comparison with a baseline fingerprint using full images as features. The objective of these experiments is to demonstrate the robustness and performance of the proposed fingerprint, compared to a full image scan-based fingerprint.

We also present applications of the proposed framework for identifying retest scans, duplicate corrupt scans, and incorrectly-reported zygosity labels. In addition, we use the segmentation masks provided with the HCP data to identify cortical and subcortical brain regions where the distribution of feature matches

between monozygotic twins is significantly different from dizygotic twins. In this analysis, we want to find brain regions which are more influenced by genetic promixity. Finally, we conduct a hemisphere asymmetry analysis using the feature matches for different types of siblings.

3.1. Manifold approximation analysis

To analyze the manifold approximation, we generated fingerprints by projecting the subject proximity graph onto a varying number of spectral components (i.e., leading eigenvectors of the normalized adjacency or Laplacian matrix). Figure 2 (top row) shows a representative 2-dimensional spectral embedding of subject proximity graphs obtained using T1, FA, or both modalities (T1+FA). As described in Section 2.2.2, modalities are combined by aggregating the feature matches in each modality when computing the pairwise subject similarities. In these plots, the position of each red circle corresponds to the 2D fingerprint of a subject. Additionally, in each plot, a single pair of MZ twins is highlighted using two blue circles and their NT sibling highlighted using a green circle.

It can be seen that the distribution of subject embeddings in the manifold varies from T1 to FA, showing that these modalities encode different properties in the fingerprint. Differences between the distributions of FA and T1+FA fingerprints are in part explained by the fact that spectral embeddings are defined up to a rotation or axis flipping. Moreover, we observe for all three modality combinations that genetically-related subjects are near to each other in the manifold, and that MZ twins are closer than their NT sibling.

In the second row of Figure 2, we measure the impact of manifold dimensionality on the fingerprint obtained with T1, FA or T1+FA modalities. The left plot shows the eigenvalues (sorted by decreasing magnitude) of the subject proximity graph's normalized adjacency matrix, which reflect the amount of connectivity information captured by their corresponding eigenvector. This plot indicates that most information is encoded in the first leading eigenvectors and, thus, that a compact fingerprint is possible.

This hypothesis is further confirmed in the middle and right plots of the same row, which evaluate for an increasing number of spectral components (i.e., fingerprint size) how the distributions of distances between MZ fingerprints and between DZ fingerprints differ. The separability between these two distributions of fingerprint distances is measured in terms of d-prime (middle plot) and unpaired t-test p-values (in -$\log_{10}$ scale). In both measures, higher values correspond to a greater separability. For all three modality combinations, a peak separability is observed around 10 eigenvectors, suggesting this value to be suitable for the fingerprint size. The decrease in separability for larger manifold dimensions is due to the fact that the added eigenvectors encode small variations of brain geometry which are not related to genetic proximity. Nevertheless, the difference between fingerprint distances in MZ pairs and in DZ pairs is significant with p-value < 0.01, for all tested manifold sizes and modality combinations.

Comparing the three modality combinations, the diffusion-based fingerprint using FA images provides a higher separability than the fingerprint generated from T1, for all manifold sizes. However, the separability is increased further when combining both modalities in the fingerprint, in line with our hypothesis that multi-modal fingerprints are more discriminative than those based on a single modality. The impact of modalities on the fingerprint is analyzed in greater details.

Finally, the last two rows of Figure 2 give the count histograms and probability density curves (fitted) of Euclidean distances between fingerprints of different sibling types. To generate these results, and in all following experiments, we used a fingerprint of 10 features (i.e., leading eigenvectors of the normalized adjacency matrix). It can be seen that the fingerprints of MZ twins, which share the most genetic material, are significantly more similar than those of DZ twins or NT siblings. This follows the results of various twin studies [88, 64, 89, 16], highlighting the relationship between genetic proximity and anatomical similarity.

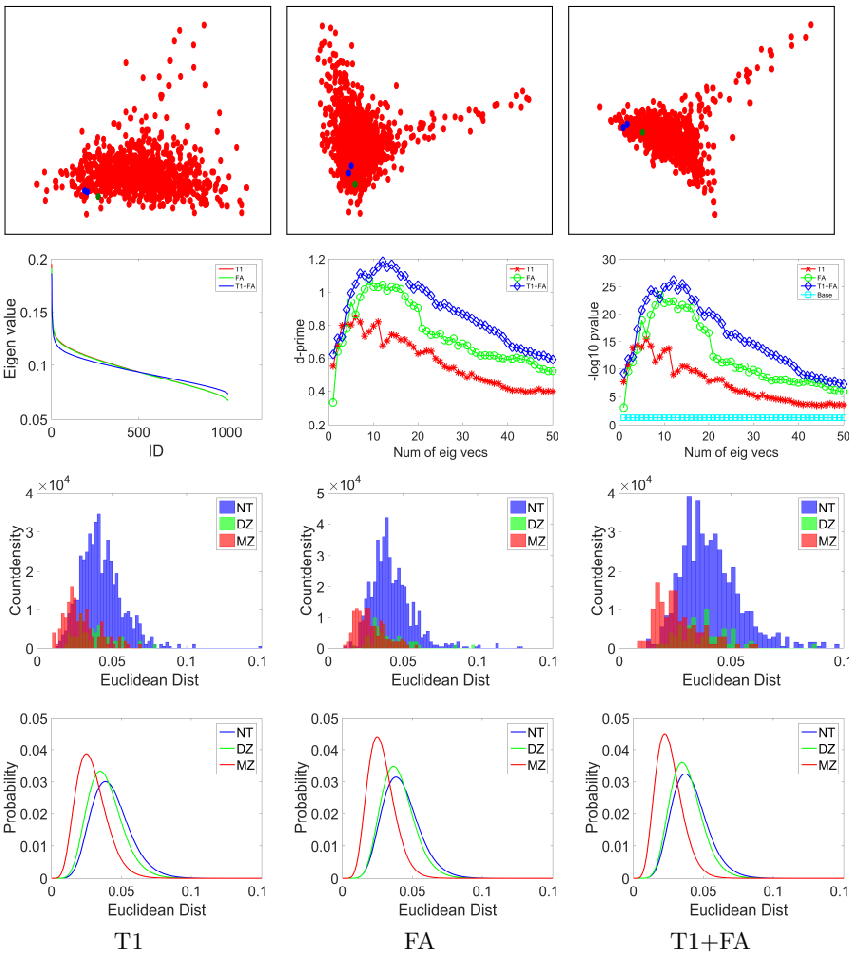

Figure 2: Compact fingerprint analysis. Top row: Representative Spectral embedding visualization using 2-D embedding, blue dots show one pair of MZ twins and green dot shows their NT sibling; Second row: Plots of eigenvalues (excluding the first), Absolute d-prime and -log10 p-value (unpaired t-test, Base p-val = 0.05) for Euclidean distances between MZ pair vs DZ pair fingerprints with increasing number of eigenvectors; Third row: Count-density histograms; and last row: Probability-normalized curves of Euclidean distances between sibling pair fingerprints using 10 eigenvectors (gamma histogram fitting).

3.2. Identification of genetically-related subjects

In this section, we use genetically verified labels of the HCP dataset to determine whether fingerprints generated using different modality combinations

Table 2: Mean average precision (MAP) table comparing different modalities for the task of genetically related subject identification.

Experiment	**Modality**	**Mean Avg Prec**		
		MZ	DZ	NT
sMRI	T1	0.901	0.139	0.143
	T2	0.902	0.163	0.142
dMRI	FA	0.935	0.182	0.181
	AD	0.694	0.088	0.090
	MD	0.811	0.139	0.118
	RD	0.827	0.114	0.135
Modality Combination (1.25mm)	T1+T2	0.968	0.246	0.235
	T1+FA	0.965	0.251	0.230
	FA+MD	0.965	0.224	0.231
	T1+T2+FA	0.990	0.349	0.287
	T1+T2+dMRI	0.986	0.355	0.339
Skull Impact	T1 Skull	0.989	0.335	0.243
	T2 Skull	0.966	0.288	0.208
Alignment Impact	T1 MNI	0.869	0.106	0.120
	T2 MNI	0.871	0.112	0.129
Resolution Impact	T1 1.25mm	0.859	0.111	0.122
	T2 1.25mm	0.876	0.176	0.147
Baseline Comparison	T1	0.634	0.079	0.051
	T2	0.501	0.062	0.034
	FA	0.687	0.075	0.056
Retest set	T1	0.924	0.191	0.146
	T2	0.908	0.198	0.129
	FA	0.929	0.165	0.181
Random	Rand	0.005	0.003	0.007

380 can identify genetically-related individuals within a group of subjects. For combining structural and diffusion modalities, we considered data at 1.25mm resolution. We refer the reader to Section 2.4 for details on the evaluation protocol

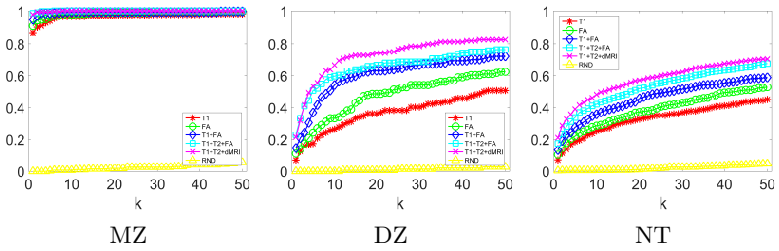

Figure 3: Twin/Sibling identification. Mean recall@k plots for MZ,DZ and NT siblings.

Table 3: Relative informativeness of fingerprints from two modalities. Comparison across modalities or their combination for the task of identification of a given sibling type. Percentage of MZ and DZ twin identification for two modalities. Total number of identification tasks is 268 and 144 for MZ and DZ respectively. We consider 50 nearest neighbors, and if twin is identified within these neighbors, identification is considered a success.

Experiment	**Mod1 vs Mod2**	**Identification % (MZ)**				**Identification % (DZ)**			
		Both	Mod1	Mod2	None	Both	Mod1	Mod2	None
Single Modality	T1 vs T2	97.39	0.75	1.87	0.00	28.47	22.22	23.61	25.69
	T1 vs FA	97.76	0.37	1.87	0.00	34.72	15.97	27.78	21.53
	FA vs MD	98.88	0.75	0.37	0.00	32.64	29.86	14.58	22.92
Modality Combination (1.25mm)	T1 vs T1+FA	98.13	0.00	1.49	0.37	52.08	2.08	20.14	25.69
	FA vs T1+FA	99.63	0.00	0.00	0.37	62.50	0.00	9.72	27.78
	T1 vs T1+T2+dMRI	98.13	0.00	1.87	0.00	52.08	2.08	30.56	15.28
	FA vs T1+T2+dMRI	99.63	0.00	0.37	0.00	61.81	0.69	20.83	16.67

and measures.

Table 2 reports the MAP values obtained in a rank retrieval of three different siblings types (MZ, DZ and NT), using fingerprints generated from various modality combinations. A rich and diverse set of observations can be drawn from this table. In the next section, it is used to analyze the impact of different factors on the proposed fingerprint's ability to identify genetically-related siblings, such as scan resolution, image alignment and skull inclusion. Moreover, results on the difference significance between the distributions of MAP values obtained for different modality combinations and sibling types are also reported

in Table 2 of Supplement material.

In Figure 3, we also report the performance of tested fingerprints using mean recall@k, for $k = 1, \ldots, 50$. This measure, also known as sensitivity, evaluates the proportion of individuals that are genetically-related to a given subject, which are within the k individuals most similar to that subject (in terms of fingerprint distance). To account for chance, we provide the mean recall@k values obtained using a random ranking of subjects.

Comparing modalities, we observe that FA yields the highest MAP and recall@k values among all single-modality fingerprints, with a significant margin. For structure-based fingerprints, T1 and T2 provide similar performances across the different sibling types. Except for FA, other diffusion modalities lead to a lower MAP or recall@k than T1 or T2. Furthermore, higher MAP and recall@k values are obtained when combining multiple modalities, the combination of T1, T2 and all four diffusion measures (dMRI = FA+AD+MD+RD) having the best performance for all sibling types. This applies for combinations within/across structural or diffusion modalities: T1+T2 outperforms T1 and T2, FA+MD performs better than FA and MD, T1+FA outperforms T1 and FA, etc.

With respect to the tested sibling types, we observe a mean recall@k near 100% when using the fingerprint to identify MZ twins, for all modalities and their combinations. This illustrates the high impact of genetic similarity on both structural and diffusion geometry in the brain. Comparing MZ, DZ and NT siblings, we see higher recall@k and MAP values for MZ twins compared to DZ twins or NT siblings, supporting the fact that MZ twins share more genetic information [67]. In contrast, performances obtained for DZ twins and NT siblings are comparable, which reflects the fact that both sibling types have the same genetic proximity. The differences between DZ and NT siblings were found to be not significant in unpaired t-test for individual modalities (Supplement material Table 1; Detailed mean recall plots are also included in Supplement material Figure 1).

To quantify the informativeness of one modality versus another, Table 3 reports the relative percentage of MZ and DZ twins identified by both, a single,

or none of the modalities[5]. The total number of identification tasks is 268 for MZ and 144 for DZ For each task, we consider the $k = 50$ nearest neighbors of a subject in terms of fingerprint distance. The identification is considered a success if the subject's twin is identified within these neighbors. When comparing the success rates of single modalities (top half of the table), we observe that FA identifies more twins uniquely, than when using T1 or MD. This is particularly noticeable for DZ twins, where 27.78% of all pairs were identified by the FA-based fingerprint but not the T1-based ones. Yet, structural modalities still capture brain tissue properties that are not provided by dMRI, as shown by the 15.97% of all DZ pairs that are identified using T1 but not with FA.

As the results in Table 2, we see that combining multiple modalities leads to a more discriminative fingerprint. For example, 1.87% of MZ and 30.56% of DZ twins are identified by fingerprints generated from all modalities (i.e., T1+T2+dMRI) but not from fingerprints based only on T1. Reversely, all MZ twins identified with T1 are also found using T1+T2+dMRI, and only 2.08% of DZ twins are identified uniquely with T1. This last result could be explained by the fact that subjects can have local similarities due to factors not related to genetics.

3.3. Impact of various factors

Factors like image alignment, scan resolution, skull inclusion and subject age, can be confounds when analyzing the proposed fingerprint. In the following subsections, we measure the impact of these factors on the fingerprint's ability to find genetically-related subjects.

3.3.1. Image alignment

Population-level analyses usually require aligning images to a common space or segmenting them into regions of interest, two steps which can be computationally expensive. In this section, we assess whether our framework can be

[5]Results for NT siblings are reported in Table 4 of Supplement material.

applied without the expensive alignment step, and evaluate the robustness of the fingerprint to this process.

Table 2 reports the retrieval performance obtained for fingerprints generated from T1 and T2 images, with or without alignment in MNI space (0.7mm resolution HCP data with affine alignment to MNI template). For all sibling types, MNI space-aligned fingerprints (denoted as MNI in the table) obtained lower MAP values than fingerprints using native space data. This observation, which is consistent across T1/T2 modalities and all sibling types, indicates that image alignment is not required for our fingerprint. Although the difference is not significant in an unpaired t-test with p-value < 0.01 (see Table 2 of Supplement material), bringing subjects to a common space may in fact reduce the discriminativeness of the fingerprint. Note that similar results were obtained using full images as fingerprints (analyzed in the following section), with lower MAP for affine-aligned images.

3.3.2. Scan resolution

Scan resolution is another important factor in multi-modal and multi-subject analyses, for example, sMRI data usually offer higher resolutions compared to dMRI. In this study, we used sMRI data at 0.7mm resolution and dMRI data at 1.25mm resolution, as provided by HCP. However, when comparing or combining sMRI data with dMRI, we used 1.25mm resolution sMRI data obtained with FSL's flirt toolbox [75]. We can thus compare the retrieval performance of fingerprints generated from lower- or higher-resolution images.

Table 2 shows that MAP values for the MZ twin identification task decrease when going from 0.7mm to 1.25mm resolution, for both T1- and T2-based fingerprints. This is due in part to the reduced number of SIFT features extracted from 1.25mm resolution image, compared to 0.7mm resolution ones. However, this is not the case for DZ and NT identification tasks, where contrasting trends are observer for T1 and T2. Moreover, differences in MAP values for the two resolutions are not significant when running an unpaired t-test with p-value < 0.01, for any sibling type (see Supplement material). These results suggest

the robustness of our framework to varying scan resolutions.

3.3.3. Inclusion of skull

Since skull size and shape is strongly influenced by genetics, including skull information in fingerprints could increase their discriminative power. In this experiment, we measure the usefulness of skull tissues for identifying pairs of MZ, DZ and NT subjects.

Table 2 reports the performances of fingerprints based on T1 and T2 image, with or without skull stripping. For both T1 and T2, as well as all sibling types, including the skull in images improves MAP values. These results are significant, with p-value < 0.01, in a one-sided unpaired t-test (see Table 2 of Supplement material). Hence, skull tissues provides additional feature matches which help identify twins and non-twin siblings. However, we should mention that skull stripping is essential to most neuroimaging analyses, and our objective here is only to measure the informativeness of skull tissues and the impact of removing these tissues on the proposed fingerprint.

3.3.4. Subject age

In twin studies, the age of subjects can be a confound when comparing between different sibling types. For instance, DZ twins and NT siblings share the same amount of genetic material, yet DZ twins could be more similar due to their same age. This is particularly important for young subjects, where small age differences can amount to large anatomical variations. The HCP data used in this study was acquired in the age range of 22-36, which corresponds to the plateau/saturation in brain and white matter development [66, 63, 21]. Nevertheless, we analyze whether age differences in non-twin siblings is a contributing factor on performance.

Toward this goal, we divided NT sibling pairs in two groups based on the median age difference of 3 years, and measured the mean performance obtained in each group, for fingerprints generated from T1, T2, and FA. As presented in Table 2, small differences in MAP can be seen between the two groups, although

these differences are not significant with p-value < 0.01 in an unpaired t-test (Supplement material Table 4).

We also evaluated the impact of absolute age on performance. In this case, we divided subjects (not subject pairs) in two groups based on the median subject age of 29 years. Again, no significant differences in MAP are observed across these groups. In summary, using the HCP dataset, we found no significant impact of subject age on the proposed fingerprint.

3.4. Comparison to baseline fingerprint

We compared the performance of our fingerprint to a baseline using full images as features. In this baseline, the similarity of two fingerprints is measured as the sum of squared distances (SSD) between intensities of corresponding voxels.

Table 2 gives the MAP obtained using this baseline, for T1, T2, and FA images in native subject space. Note that we also tested a similar baseline created from MNI aligned images, however this led to lower MAP values. For MZ twin identification, the baseline using FA images performs better than T1 or T2 images, which is consistent with the results of the proposed fingerprint. However, we see that our fingerprint performs consistently better than the baseline, with MAP improvements of 26.7% in T1, 40.1% in T2, and 24.8% in FA, for the task of identifying MZ twins. These improvements are significant in a one-sided unpaired t-test with p-value < 0.05 (see Supplement material).

In summary, while being very compact and efficient see (Section 2.3), our fingerprint based on local features is significantly more informative than a voxel-based representation for identifying genetically-related subjects.

3.5. Results reproducibility

To test the reproducibility of the results, we re-ran the same analysis after replacing the T1, T2 and FA images of 42 subjects with their retest data. Table 2 gives the MAP values obtained following this process. We observe small

differences in MAP, mainly for DZ twins and NT siblings, compared to fingerprints using the original data, however these are not significant (see Supplement material).

We note that the majority of retest subjects available in the HCP data are MZ twins. Since we do not observe significant differences for identifying this type of twins, it shows that the results are reproducible. The small differences in MAP values for DZ twins and NT siblings could be attributed to slight changes in the ordering of retest subjects' nearest neighbors.

3.6. Applications

In this section, we demonstrate the usefulness of our fingerprint on three different applications: 1) the correction of erroneous zygosity labels, 2) the detection of retest and duplicate scans, 3) the visualization and analysis of local feature matches for different modalities, sibling types and neuroanatomical regions.

3.6.1. Zygosity label correction

The Q3 release of the HCP dataset contained self-reported zygosity labels for twins. In the HCP 1200 release, which contains genetically verified zyosity labels, it was found that many self-reported DZ twins were actually MZ twins. In light of this problem, we first evaluate if the proposed framework can be used to identify the twins in large dataset whose self-reported zygosity differs from their true zygosity.

In earlier experiments, we found that the MZ twin of subjects was always within the 10 nearest neighbors of a subject (i.e., a mean recall@k of 100% was obtained for $k \leq 10$), regardless of the modality combination used for the fingerprint. Conversely, a much lower percentage of DZ twins could be identified in these lists of nearest neighbors. Based on this idea, we find incorrectly reported MZ candidates as the DZ twins which are within the 10 nearest neighbors of a subject.

Table 4 reports the percentage of DZ-to-MZ twins (62 in total) correctly identified by the proposed fingerprint, the baseline using full images, both or none of these methods, for T1, T2 and FA modalities. The results show that our fingerprint can identify most incorrectly self-reported MZ twins, with a detection rate of 90.32% for T1, 93.55% for T2, and 96.77% for FA. For all modalities, over 30% of cases were identified uniquely by our fingerprint. In contrast, no DZ-to-MZ twins were identified uniquely by the baseline fingerprint. In conclusion, the proposed fingerprint can be used effectively to detect misreported zygosity labels.

Table 4: Analysis of self-reported zygosity to genetically verified zygosity detection. Percentage of DZ-to-MZ subject's twin identification within 10 nearest neighbors for proposed framework vs Baseline (Full image based pairwise SSD). Total number of identification tasks is 62. We consider 10 nearest neighbors, and if twin is identified within these neighbors, identification is considered a success.

Modality	**Identification** %			
	Both	Proposed	Base	None
T1	58.06	32.26	0.00	9.70
T2	51.61	41.94	0.00	6.50
FA	61.29	35.48	0.00	3.20

3.6.2. Retest and duplicate scan identification

To analyze our fingerprint's ability to detect repeat scans of the same subjects (acquired after a time gap), we used the data of 1010 subjects + 42 retest subjects, and considered the task of identifying repeats scan in a rank retrieval analysis.

Following the same evaluation protocol as for identifying MZ/DZ/NT siblings, we obtained an MAP value of 1 for fingerprints generated from T1, T2 or FA. Thus, in all cases, the single most similar fingerprint to that of a subject corresponded to this subject's retest data. Moreover, when considering the number of local feature matches in the subject similarity (i.e., $\sum_{m\in\mathcal{M}} |\mathcal{B}_i^m \cap \mathcal{B}_j^m|$

in Eq (1)), we observed more matches for the retest data of a subject than for the subject's MZ twin.

Duplicate scans in a dataset, for example resulting from noise corruption, renaming or other manual errors, can introduce bias in analyses. Therefore, we also assessed if our fingerprint could detect duplicate scans of the same subject, corrupted by noise. For this experiment, we introduced duplicate scans for 42 T1 images, to which was added random noise (uniformly distributed random numbers in the $[-a, a]$ range, where $a \in \{20, 60, 100, 150, 200, 400\}$; the mean and stdev of image intensities are respectively 720 and 185).

Running a rank retrieval analysis using duplicate scans as target, we again obtained an MAP value of 1, for all tested noise levels. As in the retest scan identification task, the number of local feature matches was higher with corrupted duplicates than with images of MZ twins. Compared to retest scans, the number of feature matches was nearly half for corrupted duplicated, suggesting that noise can reduce matches to some extent. Overall, the results of this experiment demonstrate that our fingerprint can preserve brain characteristics over different scans of a subject.

3.6.3. Local feature correspondence analysis

To understand the advantages and limitations of a BoF-based fingerprint compared to voxel-wise or shape-based methods, we perform an in-depth analysis of local feature matches between subjects. In order to compare ours findings with those of related fingerprint studies like Brainprint [11], we limit our analysis to genetically-related subjects from HCP and to structural MRI modalities. Other applications of BoF representations for neuro-image analysis have been well studied in the literature [58, 33, 60, 81].

We start with a qualitative visualization of pairwise feature matches between subjects of different sibling types, for T1, T2 an FA images. The distribution of matches in these modalities in then analyzed using the segmentation maps (WM parcellation) files provided with HCP data. Furthermore, we also report cortical and subcortical regions having significantly different match distributions

across sibling types, these regions have a closer relationship to genetic proximity.

615 Finally, we perform a lateral asymmetry analysis in which the distribution of matches in hemispheres are compared.

Scale-space visualization of features correspondences

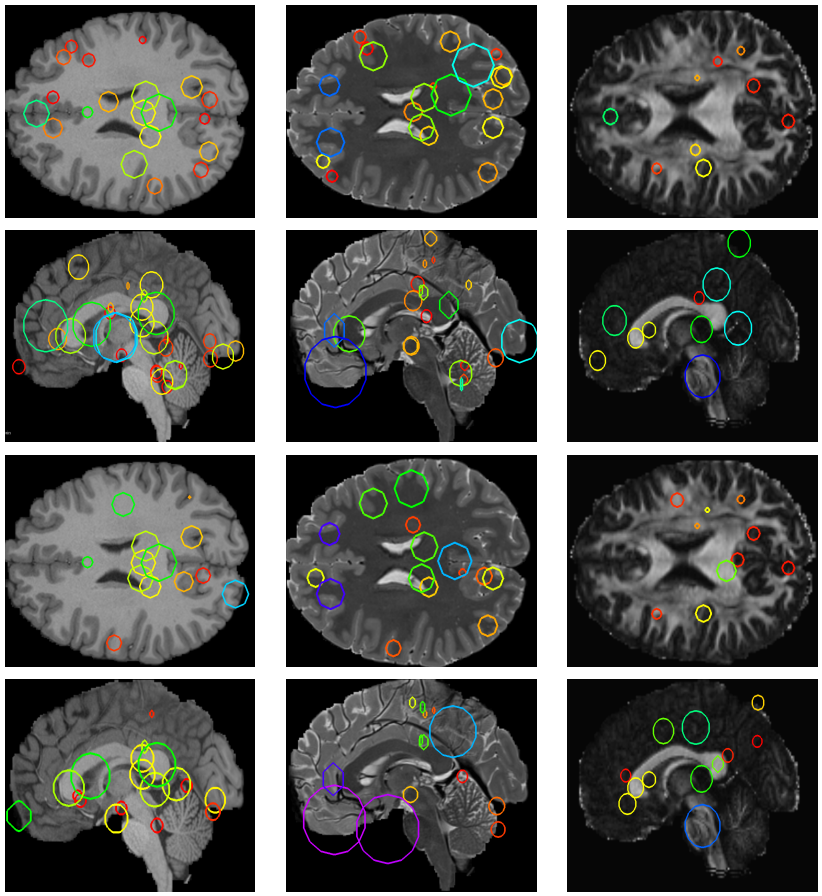

Figure 4: Example of feature correspondences for a MZ twin pair (rows 1-2) and MZ twin subject and it's Full sibling (FS) (rows 3-4). Scale space is represented using circle radius (for the visible slice).

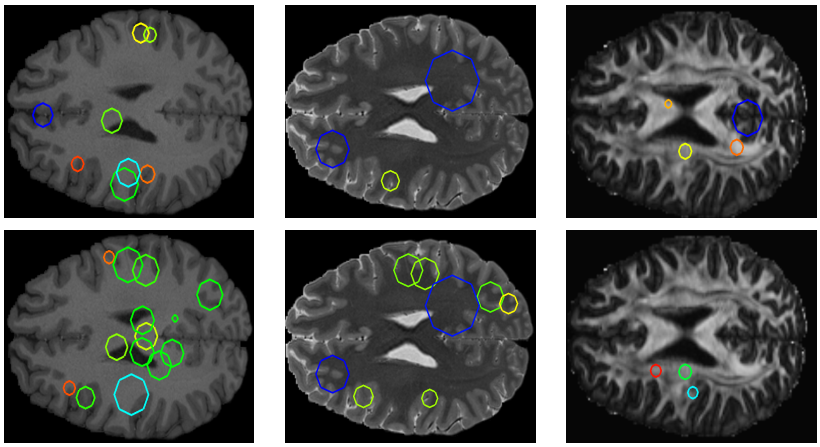

Figure 5: Example of feature correspondences for a DZ twin pair (top row) and DZ twin subject and it's Full sibling (FS) (bottom row). Scale space is represented using circle radius (for the visible slice).

Analyzing local feature matches between sibling pairs provides information in terms of their location as well as scale. In 3D SIFT features, scale corresponds to the variance of a Gaussian blur kernel for which the corresponding voxel in the blurred is a local extrema [54, 53]. It thus coincide, to a certain degree, with the size of structures in which these features are located.

Figure 4 gives a scale-space visualization of features matched with his/her MZ twin and with another non-twin sibling, for T1, T2 and FA images. Likewise, Figure 5 shows features matches for a subject and his/her DZ twin, and with another non-twin sibling. In both figures, scale information is represented using the circles' radius. Note that circles represent the intersection of 3D spheres with the visible slice and, thus, non-intersecting features are hidden in this 2D visualization.

It can be seen that different image modalities generally result in distinct, complementary feature correspondences throughout the brain. In T1 and T2 images, features are mainly located in the frontal lobe, corpus callosum and cerebellum. Smaller-scale features are also visible along various cortical regions,

as well as in subcortical structures near the basal ganglia. Moreover, images based on diffusion measures have less matches than in structural modalities. These matches are located mostly inside or near to white matter: larger-scale features in the corpus-callosum, and smaller-scale ones in the brain stem and along white matter bundles. The distribution of features in prominent brain regions is further analyzed in the next section.

Comparing different sibling types, we see a greater number of matches between MZ twins than between DZ twins or NT siblings. This observation, which is easier to visualize in T1 and T2 images, is consistent with other analyses on twin datasets. In terms of feature location and scale, we observe a slightly higher number of matches in the frontal cortex for MZ twins, however, no obvious pattern can be drawn from one set of representative plots.

Region-wise analysis of feature matches

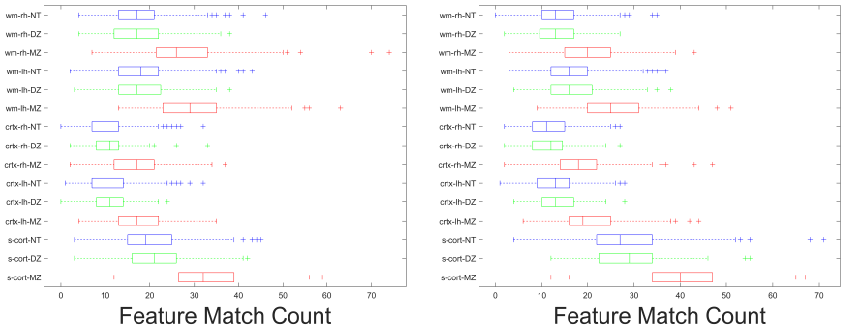

Figure 6: Box plot comparison between MZ, DZ, and NT for pairwise feature match counts for T1 (left) and T2 (right) for major structures. Red, green and blue correspond to MZ, DZ, and NT respectively. Note that the number of region-wise matches is very similar for NT and DZ pairs (blue, green), and significantly higher for MZ pairs (red). This indicates that degree of structural similarity is equivalent for NT and DZ pairs, but significantly higher for MZ pairs.

The previous section provided a qualitative analysis of feature matches for representatives pairs of siblings. Here, we analyze the distribution of feature matches across atlas-defined neuroanatomical regions, measured over the entire

Table 5: Significant parcellations for T1 and T2 for MZ vs DZ and MZ v sNT, along with the HolmBonferroni corrected p-values (-$\log_{10}$ scale) obtained using unpaired t-test (Feature match count)

Label	**T1**		**T2**	
	MZ vs DZ	MZ vs NT	MZ vs DZ	MZ vs NT
subcortical	29.31	50.31	26.06	39.41
Crtx-LH	22.85	35.17	23.37	38.87
Crtx-RH	21.48	39.76	25.38	37.73
WM-LH	37.64	62.83	27.88	47.38
WM-RH	23.38	36.80	21.88	32.62
L-Lat-Vent	5.84	11.34	5.31	7.50
R-Lat-Vent	4.21	10.72	3.99	7.57
L-VentralDC	1.49	6.06	5.16	2.98
R-VentralDC	0.45	0.60	0.00	0.01
R-Cerebellum-WM	3.98	15.40	0.00	0.57
L-Cerebellum-WM	4.82	11.11	2.33	6.55
R-Putamen	0.48	0.51	2.34	1.24
L-Putamen	0.87	0.35	0.06	0.30
L-Cerebellum-Crtx	5.74	6.26	5.58	13.81
L-Thalamus-Proper	2.71	4.03	0.37	0.01
4th-Ventricle	1.49	2.24	1.86	3.91
L-Hippocampus	3.23	3.76	4.61	5.85
CC-Anterior	1.83	0.51	0.40	0.71
R-Cerebellum-Crtx	5.96	11.93	3.38	7.57
3rd-Ventricle	0.45	0.51	0.37	0.43

group of subjects. For each scan, segmentation labels were obtained from the Freesurfer-processed data, using LUT table for label descriptions.

Figure 6 shows the box plot distributions of feature matches between pairs of MZ, DZ and NT siblings, and for T1 and T2 images. Feature match counts are reported for five broad regions: non-white matter subcortex (`s-cort`), left/right cortex (`crtx-lh/rh`) and left/right white matter (`wm-lh/rh`). Note that mapping local features to a finer cortical parcellation is difficult due to the limited

thickness of the cortex. Subcortical regions are further analyzed below.

Comparing across sibling types, we observe a higher number of feature matches for MZ pairs across all five regions and both T1 and T2 modalities. This confirms once again that the local features employed in our fingerprint captures brain characteristics related to genetic proximity. Analyzing the region-wise distribution of feature matches, all five regions are well represented. Since the number of local features in a region is proportional its size, it is not surprising that the cortex has the least matches. Yet, such features are also produced by intensity variations (i.e., edges), thus explaining why many matches are found in the cortex. Finally, when comparing T1 and T2 modalities, we see small differences in the match counts, however these are not statistically significant.

To identify regions showing a strong relationship to genetic proximity, Table 5 gives the p-values (-$\log_{10}$ scale) of an unpaired t-test comparing the mean number of matches between subjects of a given sibling type versus another sibling type (e.g., MZ vs DZ). Significance values are provided for the five major regions described above, as well for 15 prominent subcortical structures matching the analysis by Wachinger et al. [11]. To account for multiple comparisons (i.e., one for each tested region), reported p-values have been corrected using the Holm-Bonferroni procedure [90]. Moreover, to account for age and size bias in this analysis, we selected NT pairs with less than 3 years age difference, and matched the number of NT pairs to MZ pairs using a simple bipartite matching based on age.

From Table 5, we observe significant differences between MZ twins and DZ-twins/NT-siblings (-$\log_{10}$(p-value) > 2), for all five major regions and for both T1 and T2 images. In subcortical structures of T1 images, cerebellum white matter and cortex (left and right), lateral ventricles (left and right), left hippocampus and left thalamus proper have a significantly different number feature matches in MZ twins than in DZ twins or NT subjects. Comparing results obtained with T1 and T2, the same structures are significant across both modalities, differences in significance reflecting the complimentary of these modalities.

Hemisphere asymmetry analysis

In our last experiment, we analyze the symmetry of feature match counts across brain hemispheres, for major structures. Toward this goal, we considered only right-handed (RH) subjects, and limited sibling pairs to subjects with same gender (i.e., a male and his brother, or a female and her sister). For non-twin siblings, we also restricted our analysis to subject pairs with less than 3 years of age difference.

Table 6 gives the results of two-sided unpaired t-tests comparing the feature match counts between cortical or white matter regions (Freesurfer LUT labels) in left- and right- hemispheres. To analyze gender effects, we also report results individually for RH male siblings and RH female siblings. Overall, we observe significant asymmetry in white matter regions (with -$\log_{10}$(p-value) > 2) of MZ twins, the highest significance values obtained for T2 images. No clear pattern is found across sibling types, although hemispherical differences are generally higher in MZ twins than in DZ twins or NT siblings. Likewise, no conclusion can be drawn when comparing results for male and female sibling pairs, with significance values varying across different sibling types and modalities.

Table 6: Hemisphere asymmetry analysis. For a given modality and twin type we compare feature match count differences across hemisphere for major structures. (unpaired t-test + feature match count)

Modality	**Type**	**RH Female**		**RH Male**		**RH Pairs**	
		Crtx	WM	Crtx	WM	Crtx	WM
T1	MZ	0.95	1.20	0.15	**2.57**	0.87	**2.73**
	DZ	1.05	0.78	0.35	0.99	1.12	0.06
	NT	**1.71**	0.39	0.84	0.11	**1.89**	0.09
T2	MZ	**1.95**	**9.13**	**1.52**	**5.77**	**3.00**	**13.97**
	DZ	1.06	**3.60**	1.29	1.23	**1.93**	**4.06**
	NT	1.04	1.22	1.23	**5.35**	**1.90**	**5.84**

4. Discussion

In this section, we summarize the findings of this study and emphasize their link to previous investigations. We also highlight its limitations and discuss additional considerations.

Identification of genetically-related subjects

Our experiments on the task of identifying genetically-related subjects led to various useful observations. We established that the proposed fingerprint, generated from individual modalities or their combination, respects the genetic relationships between siblings, with MZ twins being more similar than DZ twins or NT siblings [88, 64, 89, 16].

Analyzing the manifold approximation, we also showed that a discriminative fingerprint could be obtained with only 10 spectral components (i.e., leading eigenvectors of the normalized adjacency matrix of the subject proximity graph). When compared to a baseline using full images as features, this compact fingerprint yielded significantly better performances, for all modalities and sibling types. This illustrates the high efficiency of our fingerprint and its advantages for comparing large groups of subjects. Moreover, while Laplacian eigenmaps were used to embed the subject proximity graph, the proposed framework is generic and other approaches (e.g., see [44, 45]) can be employed for this task.

The comparison of fingerprints obtained from structural and diffusion MRI highlighted the informativeness and complementarity of these modalities. Our experiments showed fingerprints generated from FA images to be more discriminative than those based on T1 or T2 images. Within structural modalities, T1- and T2-based fingerprints lead to similar performances in terms of MAP. However these modalities offer complementary information, each of them identifying uniquely a significant proportion of sibling pairs. Furthermore, results of this study demonstrate the usefulness of combining multiple modalities in a brain fingerprint. Thus, better performances were obtained with a combined set of modalities than with these modalities alone. Our results are consistent with previous studies underlining the benefit of a multi-modal fusion [35, 40, 61, 62].

Finally, the fingerprint proposed in this work is motivated by the recent increase in multi-modal brain studies. Multi-modal MRI has been shown useful for the analysis of neurodegenrative disorders [35]. For example, linked functional and structural deficits in distributed cortico-striato-thalamic circuits may be closely related to cognitive impairments measured in schizophrenia [41]. Combining multi-modality data also achieves a higher classification performance for identifying subjects with schizophrenia [91]. In [92], the fusion of sMRI and DTI data led to the discovery of linked alterations of gray matter (GM) and white matter (WM) morphology in adults with autism spectrum disorder (ASD).

Robustness to various factors

Our factor impact analysis demonstrated the robustness of the proposed fingerprint to the non-alignment of images. In fact, affine-aligning images in MNI space led to less discriminative fingerprints, suggesting the loss of information during the alignment process. The outstanding performances obtained on non-aligned images shows the advantages of using local features such as SIFT for encoding images [60, 81]. Furthermore, since image alignment is key for most population level analysis [22], by alleviating this requirement, the proposed fingerprint may help save computational costs and avoids errors introduced during alignment.

Experiments have also shown that scan resolution does not have a significant impact on results, although using lower resolution images reduces the number detected features. Data acquired from multiple sites or scanners often need to be brought to same resolution, introducing small errors during interpolation and re-sampling. The proposed fingerprint may thus be of help for multi-site studies, and pave the way to resolution-independent analyses.

Another interesting observation came from analyzing the inclusion/exclusion of skull tissues, considering skull tissues in the fingerprint helping to identify genetically-related siblings by increasing the number of feature matches. In addition, we showed that age was not a confounding factor in the analysis, partially due to the age range of the selected subjects being around the plateau

of brain development [63].

The reproducibility of results was also evaluated using retest scans from the HCP dataset. Using retest scans led to no significant changes in results, further validating the robustness of our fingerprint to image acquisition. However, a detailed longitudinal analysis with longer between-scan times would be required to fully confirm this claim.

Applicability of the proposed fingerprint

We used the proposed fingerprint to find incorrectly reported zygosity labels and identify retest/duplicate scans of the same subjects. Results show the outstanding performance for these tasks, with MAP values near 100% for all tested modalities (i.e., T1, T2 and FA). Hence, our fingerprint could serve as efficient and reliable tool for detecting inconsistent information in large cohorts. Another potential application could be to provide physicians with related cases in clinical settings like MCI diagnostic assistance [93].

While various twin studies have analyzed genetic influences based on volume, cortical thickness, surface area, and morphometry [11], this is the first work to use local features and manifold approximation for this problem. Analyzing the distribution of features matches across brain regions, in images of different modalities, reveals many interesting insights. Results identify various neuroanatomical regions (e.g., cerebellum, lateral ventricles, ventral diencephalon, hippocampus and thalamus proper) having significantly different match counts in MZ twins than DZ twins or NT siblings. These findings relate to those reported in [11], which were obtained on a different dataset (mean subject of age of 56 years, compared to a median of 29 years in the HCP dataset). Another key aspect of our analysis is the size of the subject cohort, larger than that of related studies [89, 94].

The asymmetry of function in the brain, for example the hemispheric specializations of language and motor functions, has been extensively studied [95, 96]. Similarly, studies have analyzed anatomical brain asymmetries based on voxel-based morphometry, sulci and other brain features [97, 11]. Sex-related dif-

ferences in anatomical and functional asymmetry were observed in [98], with greater anatomical asymmetry in mature males and better verbal processing in females [99]. Directional asymmetry and anti-symmetry are thought to be influenced by genetics [100].

The multi-modal and multi-region analysis presented in this work extends previous studies of brain asymmetry in the literature by considering sibling types. Accounting for various confounds, including gender, genetics, handedness and age, this analysis has shown a greater asymmetry in feature matches between MZ twins than DZ twins and NT siblings, mostly found in white matter regions and T2 images. Moreover, differences in asymmetry appear to be directional.

Finding meaningful differences using gross brain features is a difficult task [96, 11]. Differences in results across the two modalities T1 or T2 could be due to complementarity of the two or other confounds also. An in-depth analysis using larger dataset and using both left and right handed subjects will help better understand the hemisphere asymmetry. Another aspect would be to understand what does feature level asymmetry convey for future studies, keeping it an open question.

Additional considerations

In this work, we used a rank retrieval analysis to evaluate the relation between fingerprint similarity and genetic proximity. Mean recall@k and mean average precision (MAP) were employed to measure sensitivity, specificity, and relative informativeness of fingerprints generated from different modality combinations. However, estimating heritability directly, for instance using the approach described in [101], would provide a better quantification of genetic influence on fingerprint features.

Moreover, when building the subject proximity graph, we assumed the independence of feature matches across modalities. However, a deeper analysis could be carried out to investigate false feature matches and correlation between features matches across modality. As mentioned before, other manifold embed-

ding methods like Locally Linear Embedding (LLE) [47] could also be employed for this step.

In this study we limited ourselves to sMRI and dMRI data. However, the proposed framework is generic and could be extended to other modalities like PET-MRI and fMRI. Such extension would potentially capture similarities, in terms of brain function, between subjects at rest or performing a given task. Moreover, our multi-modal fingerprint offers a unique opportunity to measure the respective informativeness and complementarity of functional MRI, with respect to structural and diffusion MRI.

Finally, this study focused on comparing and combining different modalities for identifying genetically-related subjects, misreported zygosity labels and duplicate/restest scans. An interesting extension of this work would to be to assess whether our fingerprint can be used as a biomarker to identify subjects with cognitive or neurological disorders. Publicly available data, for instance from the ADNI dataset [33] or Parkinson's Progression Markers Initiative (PPMI) dataset [102], could be used for this analysis. Likewise, the proposed fingerprint could be used with longitudinal data to analyze the trajectories of aging subjects or subjects suffering from a neurodegenerative disease.

5. Conclusion

We presented a brain fingerprint, based on manifold approximation, for the multi-modal analysis of genetically-related subjects. In this fingerprint, each subject is represented as a bag of local features extracted from multiple image modalities. A subject proximity graph is constructed by counting the number of nearest features matches between pairs of subjects. Fingerprints are finally obtained by embedding this graph in a low-dimension manifold, using Laplacian eigenmaps.

In a rank retrieval analysis, mean recall@k and mean average precision were used to measure the relation between fingerprint similarity and genetic proximity, as well as the contribution/complementarity of information from different

MRI modalities. Results indicated that a compact fingerprint of only 10 features could identify genetically-related subjects better than a baseline using full images as features. Our experiments also showed fingerprints generated from FA images to be more discriminative than those based on T1 or T2 images, and that each modality provides complementary information which can uniquely identify some sibling pairs. Furthermore, we demonstrated the benefit of considering multiple modalities in the fingerprint, combined modalities leading to a better performance than considering these modalities separately.

Moreover, our analysis demonstrated the robustness of the proposed fingerprint to various factors, including image alignment, scan resolution and subject age. The reproducibility of results was also confirmed using retest scans from the HCP dataset, showing our fingerprint to be robust to variability in image acquisition.

The usefulness of our fingerprint was assessed on the tasks of identifying incorrectly reported zygosity and retest/duplicate scans in large dataset. Results of this experiment highlighted the effectiveness of our fingerprint, with MAP values near 100% for all test cases. Moreover, analyzing the distribution of features matches across the brain revealed neuroanatomical regions (e.g., cerebellum, lateral ventricles, ventral diencephalon, hippocampus and thalamus proper) with significantly different match counts in MZ twins compared to DZ twins or NT siblings. In addition, our analysis has shown a greater asymmetry in feature matches between MZ twins than DZ twins and NT siblings, mostly found in white matter regions and T2 images. Some of these findings are consistent with recent twin studies in the literature.

This work could be extended by further investigating the differences, in terms of feature location and similarity, between dizygotic twins and non-twin siblings. A deeper analysis of aging effects could also be performed, for instance, using longitudinal data. Such analysis would help understand the effect of neuroplasticity on individual brain characteristics.

Acknowledgements

Data were provided in part by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

References

[1] E. M. Gordon, T. O. Laumann, B. Adeyemo, S. E. Petersen, Individual variability of the system-level organization of the human brain, Cerebral Cortex 27 (1) (2017) 386–399.

[2] R. H. Grabner, D. Ansari, G. Reishofer, E. Stern, F. Ebner, C. Neuper, Individual differences in mathematical competence predict parietal brain activation during mental calculation, Neuroimage 38 (2) (2007) 346–356.

[3] J. Rademacher, U. Bürgel, S. Geyer, T. Schormann, A. Schleicher, H.-J. Freund, K. Zilles, Variability and asymmetry in the human precentral motor system, Brain 124 (11) (2001) 2232–2258.

[4] M. V. Ruiz-Blondet, Z. Jin, S. Laszlo, CEREBRE: A novel method for very high accuracy event-related potential biometric identification, IEEE Transactions on Information Forensics and Security 11 (7) (2016) 1618–1629.

[5] J.-F. Mangin, D. Riviere, A. Cachia, E. Duchesnay, Y. Cointepas, D. Papadopoulos-Orfanos, P. Scifo, T. Ochiai, F. Brunelle, J. Regis, A framework to study the cortical folding patterns, Neuroimage 23 (2004) S129–S138.

[6] D. M. Barch, G. C. Burgess, M. P. Harms, S. E. Petersen, B. L. Schlaggar, M. Corbetta, M. F. Glasser, S. Curtiss, S. Dixit, C. Feldt, et al., Func-

tion in the human connectome: task-fMRI and individual differences in behavior, Neuroimage 80 (2013) 169–189.

[7] E. M. Gordon, T. O. Laumann, B. Adeyemo, A. W. Gilmore, S. M. Nelson, N. U. Dosenbach, S. E. Petersen, Individual-specific features of brain systems identified with resting state functional correlations, NeuroImage 146 (2017) 918–939.

[8] S. Mueller, D. Wang, M. D. Fox, B. T. Yeo, J. Sepulcre, M. R. Sabuncu, R. Shafee, J. Lu, H. Liu, Individual variability in functional connectivity architecture of the human brain, Neuron 77 (3) (2013) 586–595.

[9] U. Bürgel, K. Amunts, L. Hoemke, H. Mohlberg, J. M. Gilsbach, K. Zilles, White matter fiber tracts of the human brain: three-dimensional mapping at microscopic resolution, topography and intersubject variability, Neuroimage 29 (4) (2006) 1092–1105.

[10] M. T. de Schotten, A. Bizzi, F. Dell'Acqua, M. Allin, M. Walshe, R. Murray, S. C. Williams, D. G. Murphy, M. Catani, et al., Atlasing location, asymmetry and inter-subject variability of white matter tracts in the human brain with MR diffusion tractography, Neuroimage 54 (1) (2011) 49–59.

[11] C. Wachinger, P. Golland, W. Kremen, B. Fischl, M. Reuter, A. D. N. Initiative, et al., Brainprint: a discriminative characterization of brain morphology, NeuroImage 109 (2015) 232–248.

[12] E. S. Finn, X. Shen, D. Scheinost, M. D. Rosenberg, J. Huang, M. M. Chun, X. Papademetris, R. T. Constable, Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity, Nature neuroscience.

[13] T. Xu, A. Opitz, R. C. Craddock, M. J. Wright, X.-N. Zuo, M. P. Milham, Assessing variations in areal organization for the intrinsic brain: from fingerprints to reliability, Cerebral Cortex 26 (11) (2016) 4192–4211.

[14] K. Kumar, C. Desrosiers, K. Siddiqi, O. Colliot, M. Toews, Fiberprint: A subject fingerprint based on sparse code pooling for white matter fiber analysis, NeuroImage 158 (2017) 242 – 259. doi:http://dx.doi.org/10.1016/j.neuroimage.2017.06.083.
URL http://www.sciencedirect.com/science/article/pii/S1053811917305669

[15] F.-C. Yeh, D. Badre, T. Verstynen, Connectometry: A statistical approach harnessing the analytical potential of the local connectome, NeuroImage 125 (2016) 162–171.

[16] P. M. Thompson, T. Ge, D. C. Glahn, N. Jahanshad, T. E. Nichols, Genetics of the connectome, Neuroimage 80 (2013) 475–488.

[17] R. Geevarghese, D. E. Lumsden, N. Hulse, M. Samuel, K. Ashkan, Subcortical structure volumes and correlation to clinical variables in parkinson's disease, Journal of Neuroimaging 25 (2) (2015) 275–280.

[18] S. Goldman, L. M. O'Brien, P. A. Filipek, I. Rapin, M. R. Herbert, Motor stereotypies and volumetric brain alterations in children with autistic disorder, Research in autism spectrum disorders 7 (1) (2013) 82–92.

[19] K. Amunts, C. Ebell, J. Muller, M. Telefont, A. Knoll, T. Lippert, The human brain project: creating a european research infrastructure to decode the human brain, Neuron 92 (3) (2016) 574–581.

[20] C. Sudlow, J. Gallacher, N. Allen, V. Beral, P. Burton, J. Danesh, P. Downey, P. Elliott, J. Green, M. Landray, et al., UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age, PLoS medicine 12 (3) (2015) e1001779.

[21] D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub, K. Ugurbil, W.-M. H. Consortium, et al., The WU-minn human connectome project: an overview, Neuroimage 80 (2013) 62–79.

[22] J. Dubois, R. Adolphs, Building a science of individual differences from fmri, Trends in cognitive sciences 20 (6) (2016) 425–443.

[23] E. M. Gordon, T. O. Laumann, A. W. Gilmore, D. J. Newbold, D. J. Greene, J. J. Berg, M. Ortega, C. Hoyt-Drazen, C. Gratton, H. Sun, et al., Precision functional mapping of individual human brains, Neuron 95 (4) (2017) 791–807.

[24] L. Waller, H. Walter, J. D. Kruschwitz, L. Reuter, S. Mller, S. Erk, I. M. Veer, Evaluating the replicability, specificity, and generalizability of connectome fingerprints, NeuroImage 158 (2017) 371 – 377. doi:https://doi.org/10.1016/j.neuroimage.2017.07.016.
URL http://www.sciencedirect.com/science/article/pii/S1053811917305840

[25] M. Baumann, M. Krause, J. Overgaard, J. Debus, S. M. Bentzen, J. Daartz, C. Richter, D. Zips, T. Bortfeld, Radiation oncology in the era of precision medicine, Nature Reviews Cancer 16 (4) (2016) 234–250.

[26] F. S. Collins, H. Varmus, A new initiative on precision medicine, New England Journal of Medicine 372 (9) (2015) 793–795.

[27] H. Hampel, S. OBryant, S. Durrleman, E. Younesi, K. Rojkova, V. Escott-Price, J.-C. Corvol, K. Broich, B. Dubois, S. Lista, et al., A precision medicine initiative for alzheimers disease: the road ahead to biomarker-guided integrative disease modeling, Climacteric 20 (2) (2017) 107–118.

[28] D. Bzdok, A. Meyer-Lindenberg, Machine learning for precision psychiatry, arXiv preprint arXiv:1705.10553.

[29] E. S. Finn, R. T. Constable, Individual variation in functional brain connectivity: implications for personalized approaches to psychiatric disease, Dialogues in clinical neuroscience 18 (3) (2016) 277.

[30] D. Fraguas, C. Díaz-Caneja, M. O'donovan, R. Gur, C. Arango, et al., Mental disorders of known aetiology and precision medicine in psychiatry:

a promising but neglected alliance, Psychological medicine 47 (2) (2017) 193–197.

[31] O. Miranda-Dominguez, B. D. Mills, S. D. Carpenter, K. A. Grant, C. D. Kroenke, J. T. Nigg, D. A. Fair, Connectotyping: model based fingerprinting of the functional connectome, PloS one 9 (11) (2014) e111048.

[32] C. Wachinger, D. H. Salat, M. Weiner, M. Reuter, A. D. N. Initiative, Whole-brain analysis reveals increased neuroanatomical asymmetries in dementia for hippocampus and amygdala, Brain 139 (12) (2016) 3253–3266.

[33] M. Toews, W. Wells, D. L. Collins, T. Arbel, Feature-based morphometry: Discovering group-related anatomical patterns, NeuroImage 49 (3) (2010) 2318–2327.

[34] A. L. Alexander, J. E. Lee, M. Lazar, A. S. Field, Diffusion tensor imaging of the brain, Neurotherapeutics 4 (3) (2007) 316–329.

[35] V. D. Calhoun, J. Sui, Multimodal fusion of brain imaging data: A key to finding the missing link (s) in complex mental illness, Biological psychiatry: cognitive neuroscience and neuroimaging 1 (3) (2016) 230–244.

[36] F. Liem, G. Varoquaux, J. Kynast, F. Beyer, S. K. Masouleh, J. M. Huntenburg, L. Lampe, M. Rahim, A. Abraham, R. C. Craddock, et al., Predicting brain-age from multimodal imaging data captures cognitive impairment, NeuroImage 148 (2017) 179–188.

[37] S. Liu, W. Cai, S. Liu, F. Zhang, M. Fulham, D. Feng, S. Pujol, R. Kikinis, Multimodal neuroimaging computing: a review of the applications in neuropsychiatric disorders, Brain informatics 2 (3) (2015) 167–180.

[38] S. Liu, W. Cai, S. Liu, F. Zhang, M. Fulham, D. Feng, S. Pujol, R. Kikinis, Multimodal neuroimaging computing: the workflows, methods, and platforms, Brain informatics 2 (3) (2015) 181–195.

[39] J. Sui, T. Adali, Q. Yu, J. Chen, V. D. Calhoun, A review of multivariate methods for multimodal fusion of brain imaging data, Journal of neuroscience methods 204 (1) (2012) 68–81.

[40] Q. Wang, M. Sun, L. Zhan, P. Thompson, S. Ji, J. Zhou, Multi-modality disease modeling via collective deep matrix factorization, in: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '17, ACM, 2017, pp. 1155–1164. doi:10.1145/3097983.3098164.

[41] J. Sui, G. D. Pearlson, Y. Du, Q. Yu, T. R. Jones, J. Chen, T. Jiang, J. Bustillo, V. D. Calhoun, In search of multimodal neuroimaging biomarkers of cognitive deficits in schizophrenia, Biological Psychiatry 78 (11) (2015) 794 – 804, schizophrenia: Glutamatergic Mechanisms of Cognitive Dysfunction and Treatment. doi:http://dx.doi.org/10.1016/j.biopsych.2015.02.017.
URL http://www.sciencedirect.com/science/article/pii/S0006322315001274

[42] W. Zhang, R. Li, H. Deng, L. Wang, W. Lin, S. Ji, D. Shen, Deep convolutional neural networks for multi-modality isointense infant brain image segmentation, NeuroImage 108 (2015) 214–224.

[43] S. Sharma, Translational multimodality neuroimaging, Current Drug Targets 18 (9) (2017) 1039–1050.

[44] Y. Bengio, A. Courville, P. Vincent, Representation learning: A review and new perspectives, IEEE transactions on pattern analysis and machine intelligence 35 (8) (2013) 1798–1828.

[45] X. Huo, X. S. Ni, A. K. Smith, A survey of manifold-based learning methods, Recent advances in data mining of enterprise data (2007) 691–745.

[46] J. B. Tenenbaum, V. De Silva, J. C. Langford, A global geometric frame-

work for nonlinear dimensionality reduction, science 290 (5500) (2000) 2319–2323.

[47] S. T. Roweis, L. K. Saul, Nonlinear dimensionality reduction by locally linear embedding, science 290 (5500) (2000) 2323–2326.

[48] M. Belkin, P. Niyogi, Laplacian eigenmaps for dimensionality reduction and data representation, Neural computation 15 (6) (2003) 1373–1396.

[49] J. B. Kruskal, M. Wish, Multidimensional scaling, Vol. 11, Sage, 1978.

[50] P. Aljabar, R. Wolz, D. Rueckert, Manifold learning for medical image registration, segmentation, and classification, Machine Learning in Computer-Aided Diagnosis: Medical Imaging Intelligence and Analysis: Medical Imaging Intelligence and Analysis (2012) 351.

[51] S. Gerber, T. Tasdizen, P. T. Fletcher, S. Joshi, R. Whitaker, A. D. N. Initiative, et al., Manifold modeling for brain population analysis, Medical image analysis 14 (5) (2010) 643–653.

[52] T. Brosch, R. Tam, A. D. N. Initiative, et al., Manifold learning of brain mris by deep learning, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2013, pp. 633–640.

[53] D. G. Lowe, Object recognition from local scale-invariant features, in: Computer vision, 1999. The proceedings of the seventh IEEE international conference on, Vol. 2, Ieee, 1999, pp. 1150–1157.

[54] D. G. Lowe, Distinctive image features from scale-invariant keypoints, International journal of computer vision 60 (2) (2004) 91–110.

[55] M. Vidal-Naquet, S. Ullman, Object recognition with informative features and linear classification., in: ICCV, Vol. 3, 2003, p. 281.

[56] C.-F. Tsai, Bag-of-words representation in image annotation: A review, ISRN Artificial Intelligence 2012.

[57] M. Muja, D. G. Lowe, Fast approximate nearest neighbors with automatic algorithm configuration., VISAPP (1) 2 (331-340) (2009) 2.

[58] B. Ng, M. Toews, S. Durrleman, Y. Shi, Shape analysis for brain structures, in: Shape Analysis in Medical Image Analysis, Springer, 2014, pp. 3–49.

[59] M. Toews, W. M. Wells, L. Zöllei, A feature-based developmental model of the infant brain in structural MRI, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2012, pp. 204–211.

[60] M. Toews, W. M. Wells, Efficient and robust model-to-image alignment using 3d scale-invariant features, Medical image analysis 17 (3) (2013) 271–282.

[61] M. Toews, W. M. Wells, How are siblings similar? how similar are siblings? large-scale imaging genetics using local image features, in: 2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI), 2016, pp. 847–850. doi:10.1109/ISBI.2016.7493398.

[62] K. Kumar, L. Chauvin, M. Toews, O. Colliot, C. Desrosiers, Multi-modal analysis of genetically-related subjects using sift descriptors in brain mri, in: Workshop on Computational Diffusion MRI, CDMRI 2017, MICCAI Workshop, 2017.

[63] D. C. Van Essen, K. Ugurbil, E. Auerbach, D. Barch, T. Behrens, R. Bucholz, A. Chang, L. Chen, M. Corbetta, S. W. Curtiss, et al., The human connectome project: a data acquisition perspective, Neuroimage 62 (4) (2012) 2222–2231.

[64] S. C. Kanchibhotla, K. A. Mather, W. Wen, P. R. Schofield, J. B. Kwok, P. S. Sachdev, Genetics of ageing-related changes in brain white matter integrity–a review, Ageing research reviews 12 (1) (2013) 391–401.

[65] P. Kochunov, D. C. Glahn, J. Lancaster, P. M. Thompson, V. Kochunov, B. Rogers, P. Fox, J. Blangero, D. Williamson, Fractional anisotropy of cerebral white matter and thickness of cortical gray matter across the lifespan, Neuroimage 58 (1) (2011) 41–49.

[66] P. Kochunov, N. Jahanshad, D. Marcus, A. Winkler, E. Sprooten, T. E. Nichols, S. N. Wright, L. E. Hong, B. Patel, T. Behrens, et al., Heritability of fractional anisotropy in human white matter: a comparison of human connectome project and ENIGMA-DTI data, Neuroimage 111 (2015) 300–311.

[67] T. J. Polderman, B. Benyamin, C. A. De Leeuw, P. F. Sullivan, A. Van Bochoven, P. M. Visscher, D. Posthuma, Meta-analysis of the heritability of human traits based on fifty years of twin studies, Nature genetics 47 (7) (2015) 702–709.

[68] S. N. Sotiropoulos, S. Jbabdi, J. Xu, J. L. Andersson, S. Moeller, E. J. Auerbach, M. F. Glasser, M. Hernandez, G. Sapiro, M. Jenkinson, et al., Advances in diffusion mri acquisition and processing in the human connectome project, Neuroimage 80 (2013) 125–143.

[69] K. Uğurbil, J. Xu, E. J. Auerbach, S. Moeller, A. T. Vu, J. M. Duarte-Carvajalino, C. Lenglet, X. Wu, S. Schmitter, P. F. Van de Moortele, et al., Pushing spatial and temporal resolution for functional and diffusion mri in the human connectome project, Neuroimage 80 (2013) 80–104.

[70] M. F. Glasser, S. N. Sotiropoulos, J. A. Wilson, T. S. Coalson, B. Fischl, J. L. Andersson, J. Xu, S. Jbabdi, M. Webster, J. R. Polimeni, et al., The minimal preprocessing pipelines for the human connectome project, Neuroimage 80 (2013) 105–124.

[71] F.-C. Yeh, V. J. Wedeen, W.-Y. I. Tseng, Generalized q-sampling imaging, IEEE transactions on medical imaging 29 (9) (2010) 1626–1635.

[72] P. Kochunov, D. Williamson, J. Lancaster, P. Fox, J. Cornell, J. Blangero, D. Glahn, Fractional anisotropy of water diffusion in cerebral white matter across the lifespan, Neurobiology of aging 33 (1) (2012) 9–20.

[73] N. Jahanshad, P. V. Kochunov, E. Sprooten, R. C. Mandl, T. E. Nichols, L. Almasy, J. Blangero, R. M. Brouwer, J. E. Curran, G. I. de Zubicaray, et al., Multi-site genetic analysis of diffusion images and voxelwise heritability analysis: A pilot project of the ENIGMA–DTI working group, Neuroimage 81 (2013) 455–469.

[74] T. M. Nir, N. Jahanshad, J. E. Villalon-Reina, A. W. Toga, C. R. Jack, M. W. Weiner, P. M. Thompson, A. D. N. I. (ADNI, et al., Effectiveness of regional dti measures in distinguishing alzheimer's disease, mci, and normal aging, NeuroImage: clinical 3 (2013) 180–195.

[75] M. Jenkinson, C. F. Beckmann, T. E. Behrens, M. W. Woolrich, S. M. Smith, Fsl, Neuroimage 62 (2) (2012) 782–790.

[76] G. Csurka, C. Dance, L. Fan, J. Willamowski, C. Bray, Visual categorization with bags of keypoints, in: Workshop on statistical learning in computer vision, ECCV, Vol. 1, Prague, 2004, pp. 1–2.

[77] T. Tuytelaars, K. Mikolajczyk, et al., Local invariant feature detectors: a survey, Foundations and trends® in computer graphics and vision 3 (3) (2008) 177–280.

[78] S. Kim, K.-J. Yoon, I. S. Kweon, Object recognition using a generalized robust invariant feature and gestalt's law of proximity and similarity, Pattern Recognition 41 (2) (2008) 726–741.

[79] H. Bay, T. Tuytelaars, L. Van Gool, Surf: Speeded up robust features, Computer vision–ECCV 2006 (2006) 404–417.

[80] H. Bay, A. Ess, T. Tuytelaars, L. Van Gool, Speeded-up robust features (surf), Computer vision and image understanding 110 (3) (2008) 346–359.

[81] M. Toews, C. Wachinger, R. S. J. Estepar, W. M. Wells III, A feature-based approach to big data analysis of medical images, in: International Conference on Information Processing in Medical Imaging, Springer, 2015, pp. 339–350.

[82] Y. Bengio, J.-f. Paiement, P. Vincent, O. Delalleau, N. L. Roux, M. Ouimet, Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering, in: Advances in neural information processing systems, 2004, pp. 177–184.

[83] C. Fowlkes, S. Belongie, F. Chung, J. Malik, Spectral grouping using the nystrom method, IEEE transactions on pattern analysis and machine intelligence 26 (2) (2004) 214–225.

[84] L. J. O'Donnell, C.-F. Westin, Automatic tractography segmentation using a high-dimensional white matter atlas, IEEE transactions on medical imaging 26 (11) (2007) 1562–1575.

[85] C. Buckley, E. M. Voorhees, Evaluating evaluation measure stability, in: Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval, ACM, 2000, pp. 33–40.

[86] A. Turpin, F. Scholer, User performance versus precision measures for simple search tasks, in: Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval, ACM, 2006, pp. 11–18.

[87] S. D. Gale, D. J. Perkel, A basal ganglia pathway drives selective auditory responses in songbird dopaminergic neurons via disinhibition, The Journal of Neuroscience 30 (3) (2010) 1027–1037.

[88] M.-C. Chiang, K. L. McMahon, G. I. de Zubicaray, N. G. Martin, I. Hickie, A. W. Toga, M. J. Wright, P. M. Thompson, Genetics of white matter

development: a DTI study of 705 twins and their siblings aged 12 to 29, Neuroimage 54 (3) (2011) 2308–2317.

[89] J. S. Peper, R. M. Brouwer, D. I. Boomsma, R. S. Kahn, H. Pol, E. Hilleke, Genetic influences on human brain structure: a review of brain imaging studies in twins, Human brain mapping 28 (6) (2007) 464–473.

[90] S. Holm, A simple sequentially rejective multiple test procedure, Scandinavian journal of statistics (1979) 65–70.

[91] J. Sui, E. Castro, H. He, D. Bridwell, Y. Du, G. D. Pearlson, T. Jiang, V. D. Calhoun, Combination of fmri-smri-eeg data improves discrimination of schizophrenia patients by ensemble feature selection, in: Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE, IEEE, 2014, pp. 3889–3892.

[92] T. Itahashi, T. Yamada, M. Nakamura, H. Watanabe, B. Yamagata, D. Jimbo, S. Shioda, M. Kuroda, K. Toriizuka, N. Kato, et al., Linked alterations in gray and white matter morphology in adults with high-functioning autism spectrum disorder: a multimodal brain imaging study, NeuroImage: Clinical 7 (2015) 155–169.

[93] Y. Gao, E. Adeli-M, M. Kim, P. Giannakopoulos, S. Haller, D. Shen, Medical image retrieval using multi-graph learning for mci diagnostic assistance, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2015, pp. 86–93.

[94] J. E. Schmitt, L. T. Eyler, J. N. Giedd, W. S. Kremen, K. S. Kendler, M. C. Neale, Review of twin and family studies on neuroanatomic phenotypes and typical neurodevelopment, Twin Research and Human Genetics 10 (5) (2007) 683–694.

[95] N. Geschwind, A. M. Galaburda, Cerebral lateralization: Biological mechanisms, associations, and pathology: I. a hypothesis and a program for research, Archives of neurology 42 (5) (1985) 428–459.

[96] A. W. Toga, P. M. Thompson, Mapping brain asymmetry, Nature reviews. Neuroscience 4 (1) (2003) 37.

[97] A. Pepe, I. Dinov, J. Tohka, An automatic framework for quantitative validation of voxel based morphometry measures of anatomical brain asymmetry, NeuroImage 100 (2014) 444–459.

[98] S. A. Chance, T. J. Crow, Distinctively human: cerebral lateralisation and language in homo sapiens, J Anthropol Sci 85 (2007) 83–100.

[99] D. F. Halpern, Sex differences in cognitive abilities, Psychology press, 2013.

[100] A. Gómez-Robles, W. D. Hopkins, C. C. Sherwood, Increased morphological asymmetry, evolvability and plasticity in human brain evolution, in: Proc. R. Soc. B, Vol. 280, The Royal Society, 2013, p. 20130575.

[101] T. Ge, M. Reuter, A. M. Winkler, A. J. Holmes, P. H. Lee, L. S. Tirrell, J. L. Roffman, R. L. Buckner, J. W. Smoller, M. R. Sabuncu, Multidimensional heritability analysis of neuroanatomical shape, Nature communications 7 (2016) 13291.

[102] K. Marek, D. Jennings, S. Lasch, A. Siderowf, C. Tanner, T. Simuni, C. Coffey, K. Kieburtz, E. Flagg, S. Chowdhury, et al., The parkinson progression marker initiative (ppmi), Progress in neurobiology 95 (4) (2011) 629–635.