BrainPrint: A Discriminative Characterization of Brain Morphology

Christian Wachinger^{*a,b**}, Polina Golland^{*a*}, William Kremen^{*c,d*}, Bruce Fischl^{*a,b*}, Martin Reuter^{*a,b*}, for the Alzheimer's Disease Neuroimaging Initiative^{**}

^a Computer Science and Artificial Intelligence Lab, MIT
^b Massachusetts General Hospital, Harvard Medical School
^c University of California, San Diego
^d VA San Diego, Center of Excellence for Stress and Mental Health

VA buil Diego, Center of Excellence for biress and Mental

Abstract

We introduce *BrainPrint*, a compact and discriminative representation of brain morphology. *BrainPrint* captures shape information of an ensemble of cortical and subcortical structures by solving the eigenvalue problem of the 2D and 3D Laplace-Beltrami operator on triangular (boundary) and tetrahedral (volumetric) meshes. This discriminative characterization enables new ways to study the similarity between brains; the focus can either be on a specific brain structure of interest or on the overall brain similarity. We highlight four applications for *BrainPrint* in this article: (i) subject identification, (ii) age and sex prediction, (iii) brain asymmetry analysis, and (iv) potential genetic influences on brain morphology. The properties of *BrainPrint* require the derivation of new algorithms to account for the heterogeneous mix of brain structures with varying discriminative power. We conduct experiments on three datasets, including over 3000 MRI scans from the ADNI database, 436 MRI scans from the OASIS dataset, and 236 MRI scans from the VETSA twin study. All processing steps for obtaining the compact representation are fully automated, making this processing framework particularly attractive for handling large datasets.

Keywords: Brain Shape, Large Brain Datasets, Brain Similarity, Subject Identification, Brain Asymmetry, Morphological Heritability

1. Introduction

Is it possible to identify individuals based only on the shape of their brain? Are cortical folding patterns unique to a person, similar to a fingerprint? Are brains really highly symmetric? Is it likely that shape of individual brain structures is heritable? To answer these questions, we need an accurate characterization of brain morphology. A common approach in brain morphometry is to quantify the volumes of a set of brain structures (DeCarli et al., 2005), which is, however, only a crude simplification of the full anatomical information. Shape descriptors offer a more informative representation of brain morphology, adding additional information over volumetric measurements. In this work we introduce *BrainPrint*, a holistic representation of the brain anatomy, containing the

Preprint submitted to NeuroImage

shape information of an ensemble of cortical and subcortical structures. The variety of structures included in the BrainPrint yields an extensive characterization of brain anatomy. Furthermore, working with shape representations, rather than directly with image intensities, has the advantage of remaining more robust with respect to intensity changes that may be caused by different scanner hardware or protocols. We quantify the shape information by calculating the spectrum of the Laplace-Beltrami operator both for boundary representations (defined by triangular surface meshes) and for volumetric representations (via tetrahedral meshes). BrainPrint presents a higher dimensional extension to the description of brain structures by volume measurements and therefore naturally integrates shape information into common ROI-based analysis. In addition to allowing us to investigate the questions stated above, the proposed approach provides new means to pursue many interesting research directions in neuroscience where brain morphometry, i.e., the study of brain structure and change, is of importance.

In this work, we use *BrainPrint* specifically to construct a similarity measure between brain scans. An alternative approach to define pairwise similarities could be based on image registration (Gerber et al., 2010; Hamm et al., 2010). However, this is not an intrinsic measure as the regularization term impacts the similarities. Furthermore,

^{*}Corresponding Author. Address: 32, Vassar St., D462, Cambridge, MA, 02139. Tel: +1 6176427026. Email: wachinger@csail.mit.edu

^{**}Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://adni.loni.usc.edu/wp-content/uploads/how_ to_apply/ADNI_Acknowledgement_List.pdf

many applications require large datasets and the cost of aligning a new scan to *all* scans in the database becomes prohibitively expensive for a large number of scans. *Brain-Print* introduces a new framework that is especially beneficial when working with large datasets. The first step extracts information from the image, based on the segmentation of anatomical structures. The second step transfers this information into a compact and discriminative representation, the *BrainPrint*. Any further processing is conducted on this representation, which requires significantly less memory and enables easier modeling, computation, and comparisons than working with the original scans.

Current studies in shape analysis mainly focus on single structures, e.g., the hippocampus (Golland et al., 2005; Gerardin et al., 2009; Shen et al., 2012) or the ventricles (Gerig et al., 2001a,b; Styner et al., 2005; Terriberry et al., 2005). In studies of cortex, thickness and gyrification measures are used (Batchelor et al., 2002; Luders et al., 2006). In contrast, *BrainPrint* contains the shape information of a large number of cortical and subcortical structures. This holistic representation of brain morphology offers advantages for studying the shape variability within and across populations because additional information is available for the statistical analysis. A second advantage of *BrainPrint* is rooted in the intrinsic shape description. This facilitates the statistical analysis because we can directly compute distances between shape descriptors without the need for establishing direct correspondences. Establishing correspondences is challenging and may involve computationally expensive shape registrations (Ng et al., 2014). The composition of different information across structures and dimensions (surface, volume) contained in *BrainPrint* and the use of intrinsic brain shape descriptors to define a distance function distinguishes this work from previous studies in medical shape analysis. A preliminary version of this work with a focus on subject identification was presented at a conference (Wachinger et al., 2014b). The application of Brain-Print for the prediction of Alzheimer's disease won the second prize at the challenge on Computer-Aided Diagnosis of Dementia (Wachinger et al., 2014a).

1.1. Related Work

A 3D object can be represented by the space that it occupies (3D volume representation, e.g., voxels, tetrahedra meshes) or by representing its boundary (2D surface representation, e.g., triangle meshes). Reuter et al. (2006) introduced the "shapeDNA" and demonstrated that the spectra of 3D solid objects and their 2D boundary surfaces contain complementary information: the spectra of the 2D boundary surface is capable of distinguishing two isospectral 3D solids. Therefore, we propose to combine the information from both the 3D solid and 2D boundary shape representations.

While previous work focused on the analysis of the shapeDNA for single brain structures (Bates et al., 2011; Reuter et al., 2007, 2009), to the best of our knowledge this

is the first study that evaluates its application to cortical and a wide range of subcortical structures. Importantly, we investigate the joint modeling of the ensemble. Additionally, most prior work computes the shapeDNA for triangular surface meshes (Bates et al., 2011; Bernardis et al., 2012; Niethammer et al., 2007), while we also work with tetrahedral volume tessellations. Given that the Laplace spectra are isometry invariant, the 2D boundary representation alone may yield a weaker descriptor, due to the large set of potential (near-) isometric deformations. For example, a closed 2D surface with a protrusion pointing inwards yields the same descriptor as one with the protrusion pointing outwards, while the spectra of the enclosed 3D solids differ.

Spectral methods based on eigenfunctions of the Laplace-Beltrami operator (LBO) have been used to study cortical folding variability (Germanaud et al., 2012) with applications to developmental diseases such as microcephaly (Germanaud et al., 2014). Lai et al. (2009) proposed nodal counts as a possibly more discriminative descriptor than the LBO spectrum for shapes like caudate, hippocampus, and putamen. Seo and Chung (2011) studied the compression power of LBO eigenfunctions and compared it to spherical harmonics. LBO eigenfunctions have been used for matching (Lombaert et al., 2013) and for lobar segmentation (Lefevre et al., 2014).

An alternative to structure-specific or region-specific analysis is a voxel-wise comparison. Voxel-, deformation-, and tensor-based morphometry requires deformable registration of images to align corresponding structures. For voxel-based morphometry the spatial distribution of tissue is analyzed in each voxel (Good et al., 2002). Deformationand tensor-based morphometry build on properties of deformation fields (Davatzikos et al., 1996; Miller et al., 1997). Deformation fields have further been used for measuring similarity between scans (Gerber et al., 2010; Hamm et al., 2010; Aljabar et al., 2010; Wolz et al., 2010). Based on the similarity structure, a nearest neighbor graph is constructed and spectral methods yield a low-dimensional embedding. For large repositories, the pairwise alignment of all scans to define the similarity can be computationally prohibitive. Aljabar et al. (2008) derived image similarities from anatomical segmentation overlaps. Zhu et al. (2011) find nearest neighbors by combining edge extraction with spatial pyramid matching. BrainPrint relates to these approaches because it provides a measure of brain similarity.

1.2. Outline

We introduce *BrainPrint* in Sec. 2 and describe the datasets in Sec. 3. Abstractly, we can compute distances between brains with *BrainPrint*, which is interesting for numerous applications in neuroimaging. We highlight four applications in Secs. 4 - 7, where we derive the necessary methodology for each application of *BrainPrint*. The heterogeneity of the applications also requires a different



Figure 1: Overview of the computation of *BrainPrint*. First, MRI scans are processed with FreeSurfer to obtain segmentations of cortical and subcortical structures. Second, a mesh is created for each brain structure. Third, the shapeDNA is computed for all of the η meshes, constituting the *BrainPrint*.

validation procedure for each application. The first application is the identification of subjects. Given a new scan, we aim to identify the subject based on a database of brain scans. For this purpose, we derive a robust classifier, where each structure in the *BrainPrint* votes independently. In a second application, we investigate the similarity structure that *BrainPrint* imposes on the dataset by predicting associated non-imaging data, e.g., age and sex. The third application studies lateral brain asymmetry, where we compare cortical and subcortical structures across hemispheres. We investigate different types of asymmetry, e.q., directional and fluctuating, and evaluate the dependence of asymmetry on age and sex. Specifically for the striatum, we analyze within-subject changes with a multicohort longitudinal model. In the last application, we study the genetic influence on brain morphology by analyzing the differences between monozygotic and dizygotic twins.

2. BrainPrint

An overview for the computation of the BrainPrint is shown in Fig. 1. First, we segment anatomical structures from brain scans with FreeSurfer (Dale and Sereno, 1993; Dale et al., 1999; Fischl et al., 1999a, b, 2002). Second, we create meshes for all cortical and subcortical structures of interest. Finally, we compute compact shape representations for all structures, constituting the BrainPrint. This representation is lightweight, consuming less than 9 kByte for a scan, which makes all further processing steps efficient in terms of both memory and computation. Brain-Print focuses entirely on the geometric properties of the brain and consequently it is a characterization that is robust to intensity variations between scans. Such variations may be introduced by imaging artifacts, inhomogeneities, and scanning protocols. BrainPrint is therefore well suited to operating on large databases with images acquired at different centers and scanners.

In this work we use the shapeDNA (Reuter et al., 2006) as a shape descriptor, which performed among the best in a comparison of methods for non-rigid 3D shape re-

trieval (Lian et al., 2012). The ShapeDNA is computed from the intrinsic geometry of an object by calculating the Laplace-Beltrami spectrum. Considering the Laplace-Beltrami operator Δ , we obtain the spectrum by solving the Laplacian eigenvalue problem (Helmholtz equation)

$$\Delta f = -\lambda f. \tag{1}$$

The solution consists of eigenvalue $\lambda_i \in \mathbb{R}$ and eigenfunction f_i pairs, sorted by eigenvalues, $0 \le \lambda_1 \le \lambda_2 \le \ldots \uparrow \infty$. In this work we use "spectrum" and "eigenvalues" synonymously. This partial differential equation can be solved analytically only for a few special domains, e.g. for the sphere, where the eigenfunctions are the spherical harmonics. Here we numerically estimate both the eigenvalues and the corresponding eigenfunctions for arbitrary shapes via the finite element method (Reuter, 2006). To solve the resulting generalized eigenvalue problem, we use the iterative Lanczos algorithm from the ARPACK package (Lehoucq et al., 1998). Fig. 2 illustrates the first seven eigenfunctions of the left white matter surface. Note, that throughout this work we omit the first eigenfunction for closed surfaces, which is constant with eigenvalue zero. The eigenfunctions show natural vibrations of the shape when oscillating at a frequency specified by the square root of the eigenvalue. The complex cortical folding pattern complicates the understanding of the characteristics of the eigenfunctions. We therefore also show the same eigenfunctions mapped on the inflated surface in Fig. 2. Note, that the eigenfunctions are computed on the white matter surface and only visualized on the inflated surface. The first *m* non-zero eigenvalues $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_m)$ form the shapeDNA. Uniform scaling of an object's geometry by a factor s results in scaling of the eigenvalues by a factor $1/s^2$, which holds for manifolds of any dimension (Reuter et al., 2006). For a D-dimensional manifold with Riemannian volume vol (i.e., the area for 2D surfaces), one needs to scale the geometry by $s = \text{vol}^{-1/D}$ to obtain unit volume. Thus, to be independent of the object's scale, we consider normalized eigenvalues

$$\lambda' = \operatorname{vol}^{2/D} \lambda. \tag{2}$$



Figure 2: First two rows: Left white matter surface and first seven non-constant eigenfunctions of the Laplace-Beltrami operator (sorted left to right, top to bottom) calculated on the surface. Increasing positive values of the eigenfunctions are shown in the color gradient from red to yellow and decreasing negative values are shown from dark blue to light blue. Last two rows: Visualization of the eigenfunctions on the inflated white matter surface with level sets. These are the same eigenfunctions as in top rows computed on the original white matter surface; they are visualized on the inflated surface to better understand their pattern of variation. The main directions of variation of the first three eigenfunctions are anterior-posterior, superior-inferior, and lateral-medial, respectively. The following eigenfunctions show higher frequency variations.

The importance of an objects' scale varies with the application, so that we evaluate normalized and un-normalized eigenvalues in most of our experiments.

The eigenvalues are isometry invariant with respect to the Riemannian manifold, meaning that length-preserving deformations do not change the spectrum. While isometric non-congruent surfaces exist (e.g., bending a sheet of paper), two solid bodies embedded in \mathbb{R}^3 are isometric if and only if they are congruent (translated, rotated and mirrored). Some shapes, such as the disc (Kac, 1966), are spectrally determined; meaning that the shape is known, when the spectrum is known. It has also been shown that isospectral but non-isometric shapes exist even on the plane (Gordon et al., 1992), however, so far, only synthetic pairs have been constructed, that share a variety of different geometric features. For example, area/volume, boundary length, number of holes, Euler characteristic, and several curvature integrals are spectrally determined (Minakshisundaram and Pleijel, 1967), and thus must agree for isospectral shapes. Furthermore, our good brain classification results below indicate that such isospectral (or near isospectral) objects are rare in practice. Another

important property of the spectrum is that the eigenvalues change continuously with topology-preserving deformations of the object's geometry. These properties make the shapeDNA well suited for comparing shapes, as initial alignment of the shapes can be completely avoided and minor shape changes cause only small changes in the shapeDNA.

We compute the spectra for all cortical and subcortical structures on the 2D boundary surfaces (triangle meshes) and additionally on the full 3D solid (tetrahedra meshes) for the cortical structures, forming the *BrainPrint*

$$\Lambda = (\lambda_1, \dots, \lambda_\eta), \tag{3}$$

where η is the number of meshes. In this study, we work with 36 subcortical structures and 8 descriptors for cortical structures (left/right, white/gray matter, 2D/3D), yielding $\eta = 44$. Triangle meshes of the cortical surfaces are obtained automatically for each hemisphere using FreeSurfer. Surface meshes of subcortical structures are constructed via marching cubes from the FreeSurfer subcortical segmentation. To construct tetrahedral meshes, we remove handles from the surface meshes using ReMESH (Attene and Falcidieno, 2006), uniformly resample the output to 60K vertices, and create the volumetric mesh with the gmsh package (Geuzaine and Remacle, 2009). We use the linear finite element method (Reuter et al., 2006) with Neumann boundary condition (zero normal derivative) to compute the spectra of the tetrahedral meshes.

Several brain structures that FreeSurfer identifies are lateralized, *e.g.*, left and right hippocampus, so that the shapeDNA for both hemispheres is present in the *Brain-Print*. We compute shape differences between these lateralized structures to express intra-subject relationships. Since shapeDNA is invariant to mirroring, we can directly compute

$$\boldsymbol{\lambda}_{l}^{\text{diff}} = \boldsymbol{\lambda}_{l}^{\text{left}} - \boldsymbol{\lambda}_{l}^{\text{right}}, \qquad (4)$$

for any lateralized structure l. The difference quantifies lateral asymmetries in the brain, which we discuss in details later in the article.

2.1. Distance

The definition of a distance function to quantify shape differences is fundamental for the computation of statistics on shapes. For distance calculations on BrainPrint, it is essential to take the characteristics of the eigenvalue sequence into account, which exhibits an asymptotically linear growth for 2D manifolds (Weyl, 1911, 1912). The average over a set of eigenvalue sequences shown in Fig. 12 in the appendix confirms the linear growth of eigenvalues in practice. The linear growth itself is not problematic for the distance calculation, however, as Fig. 12 shows, the variance grows quadratically. The computation of the Euclidean distance (or any p-norm) therefore causes the higher eigenvalues to dominate the distance, although they only represent a part of an object's geometry. Reuter et al. (2006) propose, among other options, linear re-weighting of eigenvalues

$$\hat{\lambda}_i = \frac{\lambda_i}{i} \tag{5}$$

to reduce the impact of higher eigenvalues on the distance. Fig. 12 shows that the mean and variance on re-weighted eigenvalues no longer exhibit the original growth pattern and therefore yield a more balanced contribution of lower and higher eigenvalues on the distance computation.

An alternative to the re-weighting is to employ the Mahalanobis distance to account for the covariance pattern in the data and to support an equal contribution of all eigenvalues in the sequence. Depending on the dimensionality of shapeDNA, many samples may be required to estimate a positive definite covariance matrix. In such cases, we can restrict the computation of the covariance matrix to the variances on the diagonal. We employ the linear re-weighting and the Mahalanobis distance in our computations. Konukoglu et al. (2013) proposed the weighted spectral distance, which is similar to a division by the squared eigenvalue number and therefore functions as a low-pass filter. To summarize, in the following applications of *BrainPrint* we can either work with the original eigenvalues λ or the volume normalized λ' , where both can be subject to linear re-weighting $\hat{\lambda}$. To simplify the notation, we will refer to them as λ , where the type of processing applied will be clear from the context.

3. Data

In this work we use data from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu), the Open-Access Series of Imaging Studies (OASIS, oasis-brains.org), and the Vietnam Era Twin Study of Aging (VETSA) (Kremen et al., 2006). Results on ADNI are presented in Secs. 4 - 6. Due to its wider age range, we work with OASIS for the age prediction in Sec. 5. The twin data from the VETSA study is used in Sec. 7 to evaluate potential genetic influences on brain morphometry. The ADNI was launched in 2003 by the National Institute on Aging (NIA), the National Institute of Biomedical Imaging and Bioengineering (NIBIB), the Food and Drug Administration (FDA), private pharmaceutical companies and non-profit organizations, as a \$60 million, 5-year public-private-partnership. The primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimer's disease (AD). Determination of sensitive and specific markers of very early AD progression is intended to aid researchers and clinicians to develop new treatments and monitor their effectiveness, as well as lessen the time and cost of clinical trials. The Principal Investigator of this initiative is Michael W. Weiner, MD, VA Medical Center and University of California - San Francisco. ADNI is the result of efforts of many coinvestigators from a broad range of academic institutions and private corporations, and subjects have been recruited from over 50 sites across the U.S. and Canada. The follow up duration of each group is specified in the protocol for ADNI. For up-to-date information, see www.adni-info.org.

The OASIS dataset consists of a series of cross-sectional MR scans from 436 subjects aged 18 to 96 years. One hundred of the included subjects older than 60 years have been clinically diagnosed with very mild to moderate Alzheimer's disease. The subjects are all right-handed and include both men and women.

VETSA is a longitudinal study of cognitive changes and brain aging with baseline in midlife (Kremen et al., 2006), where we work with twin data obtained from participants in the first wave. Participants in the VETSA were drawn from the larger Vietnam Era Twin (VET) Registry, a nationally distributed sample of male-male twin pairs who served in the United States military at some point between 1965 and 1975. VETSA participants are all military veterans; however, nearly 80% did not experience combat situations during their military careers. In comparison to U.S. census data, participants in the VETSA demonstrate similar health and lifestyle characteristics compared to American men in the same age range (Panizzon et al., 2012). To be eligible for the primary VETSA project both members of a twin pair had to agree to participate and be between the ages of 51 and 59 at the time of recruitment. Our dataset includes 67 pairs of monozygotic (MZ) twins and 51 pairs of dizygotic (DZ) twins. The average age is 56.0 years (SD=2.8).

MR T1-weighted scans from all datasets are processed with the cross-sectional pipeline in FreeSurfer. For ADNI, we additionally run the longitudinal processing stream (Reuter et al., 2010, 2012), which is used in the cohort model in Sec. 6.2.3. We construct meshes and compute the *BrainPrint* from the cortical and subcortical segmentations from FreeSurfer.

4. Subject identification

In the first application, we investigate if it is possible to identify an individual based on their brain. While the unique complexity of the brain may suggest that an unambiguous identification should be possible, there is currently little empirical research that supports this hypothesis. One difficulty for identifying the subject of a given brain is that longitudinal changes caused by aging or disease may significantly alter the brain morphometry. Additionally, scanning artifacts, inhomogeneities, and different imaging protocols can cause changes in intensity values in magnetic resonance scans, further complicating the identification. Therefore, a useful subject-specific brain signature must be both stable across time and insensitive to imaging artifacts. Moreover, it needs to provide a holistic representation of the brain to ensure subject identification even if certain parts change. Finally, small changes in the brain should map to small changes in the representation to enable a robust identification.

The characteristics of *BrainPrint* make it well-suited for biometric identification, where we pose the identification task as classification task. We derive a robust classifier by letting each brain structure vote independently for the subject's identity. Not only does our classifier identify previously encountered subjects with high accuracy, but it can also determine whether a query brain belongs to an unknown subject, not yet represented in the existing database.

4.1. Classifier

We derive a classifier for subject identification that assigns a new scan to one of the subjects in the database. Since the segmentation or tessellation may fail for certain brain structures, we seek a robust classifier that handles missing information. We achieve the robustness by combining the results from several weak classifiers operating on single brain structures. We assume n subjects C_1, \ldots, C_n and N scans in a database ($N \geq n$, for repeated scans of



Figure 3: Schematic illustration of the classification of a new scan (red dot). Scans for the same subject (dots of the same color) are represented by the mean (cross). The probability $p(\boldsymbol{\lambda}_s | \mathcal{C}_k)$ that the new scan belongs to class \mathcal{C}_k monotonically decreases with the distance to the mean μ_s^k (black double arrow).

subjects), where each scan has its associated *BrainPrint* $\Lambda_1, \ldots, \Lambda_N$. Let $S_k \subset \{1, \ldots, N\}$ denote scans for subject C_k . The probability that a new scan with *BrainPrint* Λ is an image of subject C_k is

$$p(\mathcal{C}_k|\Lambda) = \frac{p(\Lambda|\mathcal{C}_k) \cdot p(\mathcal{C}_k)}{\sum_{\nu} p(\Lambda|\mathcal{C}_{\nu}) \cdot p(\mathcal{C}_{\nu})} \propto \prod_{s=1,\dots,\eta} p(\boldsymbol{\lambda}_s|\mathcal{C}_k), \quad (6)$$

where we assume a uniform class probability $p(\mathcal{C}_k) \propto 1$ and the conditional independence of structures given the subject. The likelihood is multivariate normal distributed $p(\lambda_s|\mathcal{C}_k) \sim \mathcal{N}(\lambda_s; \mu_s^k, \Sigma_s)$ with the subject mean $\mu_s^k = \frac{1}{|S_k|} \sum_{i \in S_k} \lambda_s^i$ for structure s. Fig. 3 offers a schematic illustration of the computation of $p(\lambda_s|\mathcal{C}_k)$. Since we only have a few samples per class, we estimate a global diagonal covariance matrix Σ_s across all scans for each structure. The weighting by variances helps to prevent the domination by higher eigenvalues that exhibit higher variation, as discussed in Sec. 2.1, thus we do not apply the linear re-weighting of Eq. (5). In order to improve numerical stability, we work with log probabilities

$$\log p(\mathcal{C}_k|\Lambda) = \sum_{s=1,\dots,\eta} \log p(\boldsymbol{\lambda}_s|\mathcal{C}_k) + \text{const.}$$
(7)

$$= \sum_{s=1,\dots,\eta} \log\left((2\pi)^{-\frac{l}{2}} |\Sigma_s|^{-\frac{1}{2}} \right) -$$
(8)

$$\frac{1}{2}(\boldsymbol{\lambda}_s - \boldsymbol{\mu}_s^k)^\top \boldsymbol{\Sigma}_s^{-1}(\boldsymbol{\lambda}_s - \boldsymbol{\mu}_s^k) + \text{const.}$$

Due to positive monotonicity of the log function, all followup computations are performed without change on the log probabilities. The subject identity with the highest probability is assigned to the scan

$$k^* = \arg\max_k p(\mathcal{C}_k|\Lambda). \tag{9}$$

The posterior probability of this classifier is the product of the spectrum likelihoods across all structures, cf. Eq. (6), which may be problematic for structures with low discriminative power. Many subcortical structures do not carry much distinctive shape information and may therefore negatively influence the overall probability. Consequently, we propose a second classifier that is specifically adapted to working with structures that are not very discriminative. Increased robustness is achieved by voting for each structure independently

$$k_s^* = \arg\max_{\mathbf{\lambda}} p(\boldsymbol{\lambda}_s | \mathcal{C}_k), \quad \forall s \in \{1, \dots, \eta\},$$
(10)

with the final vote set to the mode of the vote distribution, $k^* = \text{mode}[k_1^*, \dots, k_{\eta}^*].$

4.2. Results

For this experiment, we work with over 3000 scans from almost 700 subjects from the ADNI dataset, where each subject has between three and six longitudinal scans. We use the cross-sectional processing pipeline from FreeSurfer in this experiment because the longitudinal processing pipeline relies on knowing the subject identity and would therefore not lead to a realistic setup. In addition to the 44 shape descriptors in *BrainPrint*, we include the lateral differences, cf. Eq. (4), between left and right cortical structures to quantify asymmetry, resulting in 4 additional descriptors (white/gray matter, 2D/3D). We perform leave-one-out experiments by removing one scan from the dataset and by aiming to recover the correct identity. Fig. 4 reports the classification results for the product classifier in Eq.(9) and the structure-specific voting in Eq.(10) for normalized and un-normalized eigenvalues. We report classification results as a function of the number of eigenvalues used to represent the shape. Additionally, we vary the set of brain structures in *BrainPrint*: cortical structures with triangular meshes (4), cortical structures with tetrahedral meshes (4), cortical structures for both mesh types (8), a selection of structures with the highest individual performances $(15)^1$, all structures (44), and all structures with the lateral differences of cortical structures (48). The number of structures is shown in parentheses. The results demonstrate a clear difference between the two classifiers. The product classifier achieves the best performance when working with a selection of individually best performing structures. Including more structures reduces the classification results in most cases. We observe a different behavior when working with the structure-specific voting. While both classifiers show the worst performance for sets with cortical structures alone, the addition of more structures leads to a clear improvement for the structure-specific voting. The combination of cortical and subcortical structures with lateral differences yields the best performance.

To further study this behavior, we examine the voting of each structure in more detail, shown in Fig. 5. Each column corresponds to one scan and each row to one structure. The color indicates the subject number. Scans were sorted by subject's index in the database; a perfect feature should show a color gradient from blue to red. The first 8 rows correspond to cortical structures, which exhibit the best performance. The remaining 36 rows show subcortical structures that perform worse than cortical structures and vary in their discriminative power. This explains the poor performance of the product classifier for the whole feature set, as weak features can overwhelm the response of the good features. In contrast, weak features do not degrade the performance of the voting classifier as long as weak features show no bias for a specific subject. Subcortical structures that show a reasonable performance in Fig. 5 correspond to the previously described selection of structures with the highest individual performance in footnote 1. The best retrieval performance of 99.9% is achieved for 50 eigenvalues on all features with the additional difference features on normalized eigenvalues. The best performance for un-normalized eigenvalues of 99.7% is achieved for 30 eigenvalues². Fig. 13 in the appendix shows exemplar scans for which BrainPrint does not correctly identify the subject identity. These scans show imaging artifacts, resulting in skull strip and segmentation errors. Manual correction in FreeSurfer or reacquisition to avoid motion artifacts can therefore be expected to improve the above results.

We compare our approach to two geometric representations of structures, the volume and the local gyrification index (LGI). While volume will be affected by brain atrophy, quantifying the gyrification may be more robust to longitudinal changes, assuming that the folding patterns of the brain remain stable. Schaer et al. (2008) used the LGI to identify gyral abnormalities. We transform this local measure into a global shape descriptor by computing the mean LGI over the surface. Fig. 14 in the appendix shows the mean and standard deviation of these measures calculated from several longitudinal scans per subject. The large variance and overlap across subjects indicates that such representations are not well suited for identifying subjects. For comparison, the classification accuracy for the mean LGI on both hemispheres is 1.0% for the product and 3.9% for the voting classifier. The classification accuracy for the volume, calculated from all cortical and subcortical structures, is 0.03% for the product and 0.6% for the voting classifier, confirming results from Fig. 14. In a further comparison, we compute the mean LGI for cortical regions of interests (ROIs) instead of the entire cortex. We use the Desikan-Killiany parcellation in FreeSurfer (Desikan et al., 2006) to obtain 34 ROIs per hemisphere. This extended

 $^{^1{\}rm The}$ selected structures are: cortical structures, corpus callosum, cerebellum, left/right lateral ventricle, 3rd ventricle, temporal horn of right lateral ventricle, right putamen.

²In addition to classification accuracy, we also computed sensitivity and specificity, but these statistics do not yield additional information for differentiation of feature sets because both measures are very close to one.



Figure 4: Classification results for different feature sets as a function of the number of eigenvalues m. We compare the performance of the product classifier (left column) and voting classifier (right column), as well as normalized eigenvalues (first row) and un-normalized eigenvalues (second row).

LGI-based descriptor leads to a classification rate of 86.4% for the product and 81.4% for the voting classifier.

As an additional experiment, we evaluate the possibility of determining whether a subject is not contained in the database. We study the number of votes the winning subject receives in Fig. 5, when the subject of the scan is included in the database and when the subject is excluded. If the subject in the current scan is in the database, the scan receives about 15 votes for the winning subject class. If the subject is not contained in the database, the number of votes for the winner does not surpass 4. Setting 4 votes as our decision boundary results in only a 0.49% error (false negative) of concluding incorrectly that a subject is not in the database. The false positive rate is zero.

4.3. Discussion

The high classification accuracy of *BrainPrint* suggests that brain structures are unique to individuals and that BrainPrint can potentially serve as biometric identifier. Since our study only includes data on subjects followed over a period of up to 36 months, we cannot currently assess how the accuracy of BrainPrint changes across the entire lifespan of a subject. Unfortunately, such data sets are not yet available. However, since subjects with Alzheimer's disease in our dataset demonstrate pronounced neurodegeneration in a relatively short time, we are optimistic that BrainPrint will remain robust for comparison across longer time periods. Our results demonstrate that cortical structures perform better for subject discrimination than subcortical structures, where the best performance is achieved for a combination of both types of structures. Further, the results show a better classifica-



Figure 5: Left: Subject (color) voted for by each structure (row) for each scan (column). Cortical structures in first 8 rows, subcortical features below. Optimal feature response would show a color gradient from blue to red, since scans are sorted by subject index. Right: Number of votes for the winning subject identity when the correct subject is included (blue) in the database and when it is excluded (green). Decision boundary at 4 votes (red) yields a 0.49% false negative rate.

tion accuracy for normalized eigenvalues, which indicates the importance of shape for improved subject discrimination while the impact of volume is detrimental. We investigated the range from 10 to 50 eigenvalues in this experiment, where the range is set heuristically. Germanaud et al. (2012) studied the relation between eigenvalues (spatial frequencies) and the magnitude of the cortical features, which could provide an alternative for selecting the number of eigenvalues for cortical structures.

The high identification accuracy may raise concerns about privacy issues when publicly distributing de-faced or skull-stripped brain scans together with diagnosis and other sensitive information. However, we do not believe that *BrainPrint* interferes with annonymization because at least a second scan with knowledge of the identity needs to be available to connect to the private information. Identifying similar images in an efficient way can provide the launchpad for a more detailed follow-up analysis, e.q., prediction of localized growth and shrinkage patterns. Since most of our retrieval errors are related to incorrect segmentations, our approach could also be employed for automatic quality control. Furthermore, *BrainPrint* can help identify annonymization errors (mismatch of subject identities), which are difficult to detect and can impede longitudinal studies.

5. Prediction of Non-Imaging Data

In the previous section, we have seen that in a dataset with several scans per subject, the most similar scans are the ones from the same subject. In this section, we further investigate the similarity structure that *BrainPrint* imposes on the dataset by keeping only the baseline scan for each subject. We evaluate the characteristics that are shared among the most similar subjects. To this end, we utilize non-imaging data that is collected alongside the MRI. Such data includes demographics, medical history, neuropsychological scores, blood analysis, diagnosis, and genetics. Studying the relationship of these variables with BrainPrint provides further insights into their correlation with brain morphology. Here, we focus on the relationship of BrainPrint with age and sex, which are the major factors for brain variability. For this evaluation, we vote among the k nearest neighbors to a given scan. These nearest neighbors correspond to subjects that have similar brain morphology, as measured by BrainPrint. The voting results in a predicted value, e.g., for age, which we compare to the actual age of the subject to quantify the correlation between BrainPrint and age. We study the relationship separately for each brain structure and the combined, brain-wide relationship.

5.1. Information propagation with BrainPrint

Given a collection of scans in the database with Brain-Prints $\Lambda_1, \ldots, \Lambda_N$ and the associated non-imaging data A_1, \ldots, A_N , we predict the non-imaging data for a new scan with BrainPrint Λ . We compute the distances from the test scan to all scans in the database and identify the k nearest neighbors for each structure. We employ the linear re-weighting of eigenvalues in combination with the Euclidean distance. We use $\mathcal{I}_s = [i_1^s, \ldots, i_k^s]$ to denote the indices of the nearest neighbors for structure s. Each brain structure can identify different subjects as nearest neighbors. Depending on the type of non-image data, categorical or numerical, two different prediction methods are used. For numerical data, we compute the weighted mean of non-imaging data

$$A_{s} = \frac{1}{\sum_{\nu=1}^{k} w_{i_{\nu}^{s}}} \sum_{\nu=1}^{k} A_{i_{\nu}^{s}} \cdot w_{i_{\nu}^{s}}, \qquad (11)$$

where the weight $w_{i_{\nu}^{s}} = \exp(-\|\boldsymbol{\lambda}_{s} - \boldsymbol{\lambda}_{s,i_{\nu}^{s}}\|^{2}/\tau^{2})$ with variance τ^{2} emphasizes the contribution of closer subjects. In our experiments, we have not observed much difference between the weighted and the unweighted version ($w \propto 1$). We therefore continue with the unweighted version. For categorical data, we perform majority voting and select the label that occurs most frequently

$$A_s = \text{mode}[A_{i_1^s}, \dots, A_{i_k^s}]. \tag{12}$$

A weighted version of the majority voting is also possible by making the contribution of each scan to the outcome distribution monotonically decrease with distance between corresponding *BrainPrints*.

In addition to the separate voting on each structure, we perform a combined prediction with all structures. Based on our experience for subject identification in Sec. 4, we avoid computing one global similarity but rather combine the votes for each structure to limit the impact of less discriminative brain structures. We combine the indices of nearest neighbors that are retrieved for each structure $\mathcal{I}_1, \ldots, \mathcal{I}_\eta$ and select only those scans that appear at least twice. This means that at least two brain structures have to independently vote for a scan to let it participate in the final voting, which is then analogous to the structurespecific voting in Eqs. (11) and (12). This voting procedure is different than the one used for subject identification, where we assign each scan a class label and subsequently search for the most frequent class. The reason for using another voting procedure is that the previous one is difficult to adapt to continuous variables. Further, nondiscriminative structures could have a detrimental effect on the prediction of categorical data with few labels, e.g., gender.

5.2. Results

We evaluate the performance for the prediction of age and sex with *BrainPrint*, where we are particularly interested in identifying the brain structures with the highest prediction accuracy.

5.2.1. Age prediction

We perform age prediction on the ADNI and on the OASIS dataset, which has a wider age range. Next to an improved evaluation, this also makes our results more comparable to other studies on age prediction that usually work on datasets with a wide age range. OASIS contains scans from N = 436 subjects aged 18 to 96 with a mean age of 51.4 years (SD=25.3). Baseline scans from the entire ADNI sample (N = 819) range from 54.4 to 90.9 years with a mean age of 75.2 years (SD=6.8); healthy controls (N = 229) have a mean age of 75.9 years (SD=5.0) and range from 59.9 to 89.6 years. For the prediction of age, we work with 50 eigenvalues and k = 20 nearest neighbors³. Table 1 reports the mean absolute prediction error for several structures with high prediction accuracy. The table also shows the result for the combined prediction on all brain structures. All results are reported with and without normalization. The volume loss in gray matter and white matter with aging correlates with an increase in ventricular spaces. We therefore use the volume of the lateral ventricles as a reference for age prediction, which is one of the brain structures that shows the largest effects of



Figure 6: Scatter plot of true versus estimated ages for the combined prediction with normalized eigenvalues on the OASIS dataset. Least squares regression line with correlation coefficient (Pearson's r = 0.90).

aging (Jernigan et al., 2001; Walhovd et al., 2005; Raz and Rodrigue, 2006). We show a scatter plot of the true and estimated age for the combined prediction on the OASIS dataset in Fig. 6. The plot also contains the least squares regression line with Pearson's r = 0.90.

The prediction error on all of the reported shapes is lower than using the ventricle volume. This indicates that the shape information captured with *BrainPrint* contains additional information about subject similarity that is important for predicting age. The lowest prediction errors are achieved for the combined prediction based on all brain structures. On ADNI, we observe a consistent decrease in prediction accuracy when comparing the prediction on healthy controls with that for the entire dataset. The prediction results for the normalized eigenvalues (nEV) and the un-normalized eigenvalues (EV) are similar. Results of the Wilcoxon signed-rank test on the OASIS dataset indicate that all structures listed in the table yield a significant improvement over the volume prediction on the entire dataset (all *p*-values reported in this article are 2sided). For ADNI the *p*-values are lower, but still most of the listed structures yield a significant improvement. The combined prediction with normalized eigenvalues was most significant $(p < 10^{-49} \text{ on OASIS}).$

5.2.2. Sex prediction

In the second experiment, we evaluate the prediction of sex based on *BrainPrint*. We sample from the ADNI dataset to obtain a subset that has the same number of male and female subjects in each of the three disease groups, yielding N = 684 subjects in this analysis (342)

³We set k = 20 so that k roughly corresponds to the square-root of the number of samples.

Data	VV	Norm	Comb	$\mathbf{C}\mathbf{C}$	Striatum	Putamen	Caudate	3rd Vent	Lat Vent	Hippoc	Amygdala
OASIS	27.12	nEV EV	$8.96 \\ 9.37$	$\begin{array}{c} 13.34\\ 13.21 \end{array}$	$\begin{array}{c} 11.10\\ 13.34 \end{array}$	$\begin{array}{c} 14.23\\ 13.05 \end{array}$	$14.59 \\ 15.62$	$\begin{array}{c} 11.04 \\ 13.68 \end{array}$	$\begin{array}{c} 13.56 \\ 14.10 \end{array}$	$16.57 \\ 17.22$	$\begin{array}{c} 18.28\\ 16.57 \end{array}$
ADNI Controls	4.31	nEV EV	$3.73 \\ 3.77$	$3.70 \\ 3.82$	$3.82 \\ 3.83$	$3.94 \\ 4.01$	$3.94 \\ 3.89$	$3.94 \\ 3.75$	$3.84 \\ 3.87$	$3.86 \\ 3.88$	$3.73 \\ 3.99$
ADNI Entire	5.66	nEV EV	$4.92 \\ 5.03$	$5.26 \\ 5.15$	$5.15 \\ 5.31$	$5.40 \\ 5.55$	$5.29 \\ 5.34$	$5.27 \\ 5.25$	$5.29 \\ 5.32$	$5.34 \\ 5.41$	$5.39 \\ 5.51$

Table 1: Mean absolute age difference between estimated and true age on the OASIS and ADNI (controls and entire) dataset. Shown are results for the combined prediction of all brain structures (Comb) and the separate prediction for a selection of well performing structures. For lateral structures, we show results for the right hemisphere, where results across hemispheres are very similar. The prediction error for using the ventricle volume (VV) is listed as reference. We compare the prediction on the entire dataset and on healthy controls, as well as for normalized (nEV) and un-normalized (EV) eigenvalues.

male and female, CN=220, MCI=282, AD=182, average age 75.1 years (SD=6.64)). Fig. 7 reports the prediction results for the nine best performing structures and the combined prediction. All results are for un-normalized eigenvalues, which perform better than the normalized version due to differences in head size between men and women. The best performance is achieved for the combined prediction, followed by cortical structures. The best performing subcortical structure is the brainstem. The improvement of the cortical structures over the brainstem is significant. This difference between cortical and subcortical structures only holds for the un-normalized eigenvalues, for the normalized eigenvalues they perform similarly. We use the head size as a volume-based reference measure for the prediction of sex. Buckner et al. (2004) suggested that normalization with the intracranial volume (ICV), which is about 12% larger for men than women, corrects for gender-based head size differences. Additionally, Walhovd et al. (2005) reported a significant difference in mean ICV between men and women. ICV is therefore a well-suited reference for the prediction and yields a 0.69 prediction accuracy in our experiments. The combined prediction with *BrainPrint* performs significantly better than the prediction with ICV (p < 0.005). To disentangle the impact of the size and the shape on the sex prediction, we also computed the combined prediction for volume only (0.74), for normalized eigenvalues (0.68), and for unnormalized eigenvalues (0.78). The shape alone performs worse than the volume. The combination of shape and size in the un-normalized eigenvalues yields the best performance.

5.3. Discussion

These experiments provide an interesting insight into the versatility of *BrainPrint*. While subcortical structures achieve best results for the age prediction, cortical structures perform significantly better for the sex prediction. Furthermore, the normalization of the eigenvalues helps to adapt the notion of similarity to the specific application. Experiments indicate increased prediction accuracy for age and sex prediction in *BrainPrint* compared to volumetric measurements, most likely due to the more detailed characterization of brain morphology in *BrainPrint*.

For the prediction of age, the error based on each of the reported shapes is lower than when using the ventricle volume. The lowest prediction errors are achieved for the combined prediction with all brain structures. Moreover, we observe a consistent decrease in prediction accuracy on the entire ADNI dataset when compared to prediction on only healthy controls. This decrease may indicate a change in the pattern of aging for patients with Alzheimer's disease or mild cognitive impairment. The pronounced atrophy in dementia patients may cause the brain morphology of younger subjects with disease to be similar to older subjects without disease, thus reducing the accuracy in age prediction. Our predicted age can be seen as an estimate of biological age, where the discrepancy between predicted and chronological age can serve as a biomarker for neurodegenerative disease (see also Gaser et al. (2013) for a similar approach).

For the prediction of sex, un-normalized eigenvalues yield higher accuracy than normalized eigenvalues. This highlights the importance of the size of brain structures when seeking subjects with matching sex. Our results also show that the volume of structures alone yields better results than the shape alone. The highest accuracy for un-normalized eigenvalues suggests that the shape information captured with *BrainPrint* contains additional information about subject similarity that is important for predicting sex, with a significant improvement from subcortical to cortical structures.

In comparison to other studies on age and sex prediction, Ashburner (2007) used diffeomorphic registration and support vector classification. The reported accuracy for sex prediction was about 87%. The root-mean-square error for age prediction was about 7 years and the best correlation was 0.86. Duchesnay et al. (2007) proposed cortical sulci descriptors for the sex prediction and the best classification pipeline achieved 85% classification accuracy. Our accuracy for sex prediction of 78% is lower, but this may in part be due to different datasets and study designs. In our experiments, we paid special attention to sample



Figure 7: Prediction accuracy of sex for the combined prediction on all brain structures and the nine best performing single structures. Cortical structures perform best, with each one offering a significant improvement in prediction accuracy over the best subcortical structure (brainstem). Bars correspond to the mean prediction accuracy and error bars display standard error. The color of the bars is used for visualization purposes only. *, **, and *** indicate statistical significance levels at 0.05, 0.01, and 0.001, respectively.

sex-matched subsets from ADNI to avoid bias. Concerning age prediction, the dataset used by Ashburner (2007) contains subjects aged 17 to 79 with a mean age of 31.8 years. The age range is similar to the OASIS dataset but the mean age differs. The correlation we measured on OA-SIS is higher than the one reported in Ashburner (2007)but also the prediction error is higher. Since the results are dependent on the dataset and the study design, we can only draw limited conclusion from such a comparison, but overall our results are in a similar range than alternative methods. Note, that we do not train a classifier but average among the nearest neighbors identified by Brain-*Print*. For increasingly larger datasets, we obtain a denser sampling of the brain manifold with the expectation that neighbors will be more similar which could further improve prediction accuracy.

6. Lateral Shape Symmetry

In this section, we study the lateral shape symmetry of the brain. The study of symmetry in neuroscience investigates similarity of the two hemispheres in the brain, where the symmetry of the human body along the vertical body axis is pronounced in the brain (Hugdahl, 2005). Symmetry can be studied from a functional or anatomical perspective. The functional layout of the brain is organized asymmetrically, with hemispheric specializations for key aspects of language and motor function (Geschwind and Galaburda, 1985; Toga and Thompson, 2003). We focus on anatomical brain symmetry, where previous studies reported asymmetries based on voxel-based morphometry (Pepe et al., 2014) as well as the analysis of sulci and other brain features. Three distinctive patterns of population-level brain asymmetries are possible (Gómez-Robles et al., 2013): (i) directional asymmetry occurs when the two sides of the brain are systematically different; (ii) anti-symmetry is the consistent difference between sides, but the direction of this difference varies; and (iii) fluctuating asymmetry is the non-directional departure from bilateral symmetry.

Gómez-Robles et al. (2013) studied morphological brain symmetry by manually identifying landmarks on the cortex. They observed anatomical asymmetries in both human and chimpanzee brains, but human brains were especially asymmetric. The lack of symmetry in human brains may be a sign of plasticity, which is critical for human cognitive evolution. Greve et al. (2013) studied the association of asymmetry in function with anatomy. The results of the surface-based analysis of the cortex indicated that gross morphometric asymmetry is only subtly related to functional language laterality. Chance and Crow (2007) found sex differences in anatomical and functional asymmetry, which are plausibly related to sex differences in verbal ability in human populations without established counterparts in chimpanzees. These sex differences indicate a possible role of sexual selection in human lateralization and are consistent with data suggesting greater anatomical asymmetry in mature males, but relatively better verbal processing in females (Halpern, 2000). Next to sex, genetics is thought to play an important role in directional asymmetry and anti-symmetry (Gómez-Robles et al., 2013). In contrast, fluctuating asymmetry was suggested to be rather related to environmental factors that affect the brain development. Other studies found little evidence for a genetic influence on brain asymmetry (Bishop, 2013; Eyler et al., 2014).

These mixed results may be a consequence of the complexity and variability of brain structures, which makes it difficult to identify asymmetries and distinguish them from random fluctuations. Also, a cursory examination of the gross brain features fails to identify profound left-right differences (Toga and Thompson, 2003). A careful examination with computational methods to compute statistics on larger populations is therefore required to expose asymmetries. As many structures are represented in BrainPrint for both hemispheres, it is well suited to study structural changes by computing shape distances. This analysis permits us to address the following questions: Which patterns of asymmetry (directional, fluctuating) are present? How similar are structures across hemispheres? Does aging have a significant effect on the brain symmetry? Are asymmetries equally strong for men and women?

6.1. Lateral shape distances in BrainPrint

To study symmetry, we compute lateral (left/right) shape differences of brain structures. More precisely, we compute distances for the following 12 structures on triangular meshes: white matter, gray matter, cerebellum white and gray matter, striatum⁴, lateral ventricles, hippocampus, amygdala, thalamus, caudate, putamen, and accumbens. Additionally, we compute distances on tetrahedral meshes for white and gray matter. Due to the isometry invariance of shapeDNA, we directly compare the eigenvalue sequences as mirrored objects yield the same shapeDNA. We use the linear re-weighted eigenvalues in Eq.(5) for the lateral distance computation for lateral structures l

$$d_l = \|\boldsymbol{\lambda}_l^{\text{left}} - \boldsymbol{\lambda}_l^{\text{right}}\|_2.$$
(13)

6.2. Results

All experiments for lateral shape asymmetry presented in the following sections are performed on the ADNI dataset.

6.2.1. Overall symmetry of hemispheres

First, we investigate the overall symmetry of brain hemispheres. To put the lateral shape distances in Eq. (13) into perspective, we do not only compute the distance across hemispheres for the same subject but for all subjects in the dataset. Next, we compute the rank of the intrasubject distance compared to the inter-subject distances. The rank expresses the number of subjects that have a more similar contralateral hemisphere than the subject itself. A low rank indicates a high lateral shape similarity compared to the variations in the population. Since the rank is easier to interpret in relation to the size of

the dataset, we divide the rank by the number of subjects to obtain a relative measure of brain asymmetry. We use m = 50 eigenvalues with linear re-weighting in this experiment. Fig. 8 reports the median of this measure for lateralized structures, where low values indicate high shape symmetry. We measure significant differences in rank between normalized and un-normalized eigenvalues. Fig. 8 illustrates brain structures sorted by decreasing significance, where we use the Wilcoxon signed-rank test to compute *p*-values. We also compute the rank for volume only to disentangle the impact of shape and size on the brain symmetry. The un-normalized eigenvalues have significantly lower rank than the volume for the majority of structures. This shows that the shape with size information is more discriminative in identifying the brain structure in the other hemisphere than size and shape alone. The only exception is the putamen, where we note a decrease in rank with normalization. We work with healthy controls in this experiment as we aim to avoid disease related effects. Thompson et al. (2007) reported that atrophy associated with Alzheimer's disease may be lateralized with stronger atrophy on the left hemisphere.

We observe that cortical structures exhibit the lowest ranks. This is not surprising because complex shapes are likely to show more variability across the population and lateral asymmetries still lead to a higher intra-subject symmetry compared to inter-subject similarity. Asymmetries in simpler, "potato like" subcortical structures yield a faster decrease in rank due to smaller inter-subject variability. Nevertheless, the high median relative rank of more than 10% for all subcortical structures is surprising. This implies that given a brain structure, *e.g.*, the left caudate, the shape distance to the right caudate of 10% of all subjects in the dataset is lower than the right caudate of the same subject.

6.2.2. Pattern of brain asymmetry

In a further experiment, we investigate the pattern of brain asymmetry. To this end, we examine the lateral differences per eigenvalue instead of the distance. A distribution of the difference around zero indicates fluctuating asymmetry (iii). The directional asymmetry (i) is characterized by a distribution significantly biased from zero. A bimodal distribution of differences identifies antisymmetry (ii). We use Hartigan's dip test to evaluate if the distribution of differences is bimodal (Hartigan and Hartigan, 1985). The *p*-values for all eigenvalues across all structures are above 0.05, suggesting this data provides no evidence for significant bimodality. Since there is no evidence for anti-symmetry (ii), we focus the analysis on distinguishing directional (i) and fluctuating (iii) asymmetry.

Fig. 9 shows a box-and-whisker plot of the first 20 normalized eigenvalues of the white matter triangular mesh. We note that the distribution of differences shows strong deviations from zero for a number of eigenvalues. Specifically the difference of the second eigenvalue shows a strong

 $^{^{4}\}mathrm{Striatum}$ is the combination of caudate, putamen, and accumbens.



Figure 8: Analysis of lateral asymmetry of the brain based on shape (un-normalized and normalized eigenvalues) and volume. The bars show the median of ranks relative to the sample size for lateral brain structures. A low relative rank indicates high lateral symmetry. We compute p-values with the Wilcoxon signed-rank test between the results of un-normalized and normalized eigenvalues (light green), as well as un-normalized eigenvalues and volume (yellow). The structures are sorted in decreasing order of p-value between un-normalized and normalized eigenvalues. The normalization of eigenvalues yields a significant increase in rank for all structures except putamen. The un-normalized eigenvalues have a significantly lower rank than the volume for the majority of structures.

directional asymmetry, where the distribution of the differences is significantly different from a distribution with median zero ($p < 10^{-37}$, Wilcoxon signed-rank test). In Sec. 2, the eigenfunction corresponding to the second eigenvalue on the white matter surface is shown in Fig. 2. The eigenfunction shows variations from left to right. In an analysis across structures, we find that the differences for 18.5% of eigenvalues show significant directional asymmetry, at a significance level that takes Bonferroni correction into account.

6.2.3. Lateral shape symmetry across age and sex

The advantage of computing shape distances, in contrast to working with the high-dimensional shapeDNA, is the reduction to a single scalar value for each lateralized structure. This facilitates statistical analysis because we can use a single linear model per lateralized structure with the distance serving as a dependent variable and age and sex treated as independent variables. Table 2 reports the standardized regression coefficients and *p*-values for age and sex. As in the previous asymmetry experiments, we work with healthy controls. We report results for different lateralized structures and for eigenvalues with and without normalization. Across almost all structures we note an increase in asymmetry with age. For a number of structures highlighted in the table this increase is significant. Aging shows the largest impact on the striatum, whose importance in aging has been previously documented (Raz et al., 2003). The only exception to the overall increase is the decrease in asymmetry for the ventricles with un-normalized eigenvalues. This decrease can be explained by the strong increase in ventricle volume with aging, which causes the decrease of the eigenvalues with aging. Since the eigenval-

14

ues of both hemispheres decrease, this results in a smaller distance with aging.

For the impact of sex, we note a consistently more symmetric brain morphology for women based on normalized eigenvalues. These results are consistent with previous findings in (Halpern, 2000; Chance and Crow, 2007). The difference across sexes is significant for cortical structures, striatum, hippocampus, thalamus and caudate. The largest difference with $\beta = -0.25$ and lowest *p*-value $(p < 10^{-4})$ is measured for the hippocampus. The results for the un-normalized eigenvalues are not as consistent. This may mainly be related to the difference in head size between men and women and the corresponding scaling effect it has on the distances. We also evaluated the effect of adding handedness as additional independent variable to the model. We did not find a significant impact of handedness on brain asymmetry for normalized or un-normalized eigenvalues. This observation is consistent with the results of Good et al. (2002), who have also reported no relationship of handedness and asymmetry.

6.2.4. Longitudinal shape symmetry across age and sex

So far, we have looked at cross-sectional data, which can confound within- and between-subject variations (Schaie and Caskie, 2005). To better separate these two types of variation, we integrate longitudinal data from follow-up scans to the analysis. Single-cohort longitudinal designs are considered one of the best designs for investigating within-subject variations. However, drawbacks of such a design are the length of time and the limitation to a single cohort. Participants in ADNI are recruited from a large age range, causing several age cohorts to be present and therefore constitute an unstructured multicohort longitu-

Structure	No	ormalized	eigenva	lues	Un-normalized eigenvalues				
	А	ge	S	Sex		Age		Sex	
	β	p	β	p	β	p	β	p	
White-Tri	0.15	0.020	-0.16	0.015	0.24	0.000	0.05	0.455	
Pial-Tri	0.09	0.144	-0.15	0.020	0.13	0.053	0.10	0.109	
White-Tet	0.11	0.097	-0.14	0.028	0.27	0.000	0.04	0.501	
Pial-Tet	-0.03	0.681	-0.07	0.256	0.05	0.408	0.17	0.011	
Striatum	0.20	0.002	-0.15	0.021	0.06	0.330	-0.06	0.365	
Ventricle	0.09	0.164	-0.09	0.152	-0.18	0.005	0.01	0.821	
Cerebellum WM	0.03	0.696	-0.04	0.498	0.11	0.099	0.03	0.601	
Cerebellum Cortex	0.16	0.011	-0.09	0.161	0.04	0.541	0.01	0.935	
Hippocampus	0.12	0.068	-0.25	0.000	0.14	0.030	-0.05	0.480	
Amygdala	0.10	0.131	-0.08	0.237	0.16	0.014	0.08	0.204	
Thalamus	0.07	0.252	-0.15	0.025	0.14	0.037	-0.05	0.451	
Caudate	0.17	0.009	-0.19	0.003	0.09	0.152	-0.04	0.495	
Putamen	0.17	0.010	-0.02	0.775	0.03	0.611	0.06	0.367	
Accumbens	-0.02	0.815	-0.05	0.426	0.09	0.161	0.05	0.440	

Table 2: Regression model coefficients and *p*-values for the dependence of lateral shape distances on age and sex. Results are shown for normalized and un-normalized eigenvalues. We show standardized regression coefficients β and highlight signifiant dependencies in bold.

dinal design (Thompson et al., 2011). We use the longitudinal processing stream in FreeSurfer (Reuter et al., 2012) for this experiment to avoid processing bias (Reuter and Fischl, 2011). To study cross-sectional and longitudinal effects, we work with a linear mixed-effects model. We use B_i to denote the age at baseline for subject *i*, X_{ij} to be the time from baseline at follow-up scan *j*, and S_i to be the sex of subject *i*. For the lateral shape distance Y_{ij} as a dependent variable, we employ a linear model

$$Y_{ij} = \beta_0 + \beta_1 B_i + \beta_2 X_{ij} + \beta_3 S_i + b_{0i} + b_{1i} X_{ij}, \quad (14)$$

where $\beta_0, \beta_1, \beta_2, \beta_3$ are fixed effects regression coefficients and b_{0i}, b_{1i} are random effects regression coefficients. The random effects enable modeling subject-specific intercept and slope with respect to the time from the baseline. We investigate the striatum asymmetry with this longitudinal model, because it showed the strongest age related effects. We select healthy controls with at least three longitudinal scans, where the highest number of scans is six. The subject age was centered at 69 years. Fig. 10 displays the estimated within-subject and acrosssubject change of the lateral shape distance for five age cohorts (70, 74, 78, 82, and 86 years of age at baseline). The lateral shape asymmetry increases with age cohort $(\beta_1 = 0.04, p < 0.001)$ and men show stronger asymmetry than women ($\beta_3 = -0.28$, p < 0.05), consistent with the previous results in Table 2. However, withinsubject increase in lateral asymmetry has double the rate than the age of the cohort ($\beta_2 = 0.08$, p < 0.005). Consequently, striatum asymmetry increases faster as a function of within-subject change in age than as a function of cohort age. We also evaluated the interaction between age and the time from the baseline, but did not find a significant dependence. The increase in asymmetry with age may be a sign of brain plasticity; the adaptation of brain structure throughout life. This may be correlated with changes in brain function, which are known to be influenced by a

person's life experience.

6.3. Discussion

The results on the overall symmetry of hemispheres show that the normalization of eigenvalues decreases the rank. The volume is therefore an important factor in lateral shape symmetry. Focusing only on shape causes hemispheres of other subjects to appear more similar than the contralateral hemisphere in the same subject. For most structures, the volume is more discriminative than the shape alone, but the combination of shape and volume leads in most cases to highly significantly lower ranks. Our analysis of the pattern of asymmetry suggests fluctuating and directional asymmetry. In a follow-up study, it would be interesting to further evaluate the eigenfunctions that correspond to the eigenvalues to investigate further the type of shape changes that are directional. Our results for the influence of age and sex on lateral shape asymmetry showed significant dependency for a number brain structures. The general trend is that asymmetry increases with age and is higher in men than in women. To further differentiate within-subject and across-subject changes, we analyzed the change with age and sex in a longitudinal multicohort study. The results for the striatum show an even more pronounced increase of asymmetry with age within subjects.

7. Genetic Influences on Brain Morphology

In this section, we investigate the potential genetic influence on brain morphology represented by *BrainPrint* by comparing differences between twins in monozygotic (MZ) vs. dizygotic (DZ) cohorts. Analysis of genetic and environmental influences on brain structure and its changes over time is important for the basic understanding of brain



Figure 9: Box-and-whisker plot of differences between normalized eigenvalues on left and right white matter triangular mesh. Red lines indicate the median, the boxes extend to the 25th and 75th percentiles, and the whiskers reach to the most extreme values not considered outliers (red crosses). Outliers are points that are farther than 1.5 times the interquartile range away from either end of the box. Distribution of differences around zero indicate fluctuating asymmetry, whereas strong deviations from zero indicate directional asymmetry. Distributions of differences with non-zero median are marked with a star when significant in Wilcoxon signed -rank test (Bonferroni corrected).

aging (Kremen et al., 2010). Genetically informative designs, such as twin studies, are well suited for differentiating the underlying genetic and family environmental sources of resemblance (Eaves et al., 1978; Neale and Cardon, 1992). Twin studies usually include both MZ twins who share 100% of their genes, and DZ twins who, on average, share 50% of their genes. The reported proportion of phenotypic variance due to genes is generally consistent across large scale twin studies (Schmitt et al., 2007; Peper et al., 2007; Blokland et al., 2012). However, the majority of studies focuses almost entirely on volume, cortical thickness, or surface area rather than shape. Previous shape studies on twin data focused on lateral ventricles (Gerig et al., 2001a,b; Styner et al., 2005; Terriberry et al., 2005) or quantified gyral and sulcal patterns (Schmitt et al., 2007; Im et al., 2011). In contrast, our shape representation includes a large variety of cortical and subcortical brain structures and uses both 2D surface and 3D volume representations. In this study, we compare brain shapes across the MZ and DZ twin groups as an initial test for potential genetic influence. Direct estimation of heritability is outside the scope of this paper, as BrainPrint describes each structure by a vector instead of a single scalar, which makes standard metrics for quantifying heritability not straightforward to compute (Neale and Cardon, 1992).



Figure 10: Longitudinal data for lateral distances of striatum. Lines and line segments show model fitted from linear mixed-effects model. Line segments show withinsubject slope for five age cohorts (70, 74, 78, 82, and 86 years of age at baseline). The longer lines indicate crosssectional slopes at baseline.

A multivariate extension of such metrics will be subject of further research.

7.1. BrainPrint for twin analysis

To test the hypothesis that brain shapes are more similar in MZ twins than in DZ twin pairs, we first compute shape distances across a twin pair for each structure captured by *BrainPrint*. Given a twin pair (a, b) with *BrainPrint* (Λ^a, Λ^b) , we compute the twin shape distance for a brain structure $s \in \{1, \ldots, \eta\}$

$$d_s(\Lambda^a, \Lambda^b) = \|\boldsymbol{\lambda}_s^a - \boldsymbol{\lambda}_s^b\|_2.$$
(15)

For instance, we compute a shape distance between the right caudate of one twin to the right caudate of the other twin in a pair. According to the hypothesis, we expect this distance to be lower for MZ twins than for DZ twin pairs. Based on the results from the subject identification, we use m = 30 eigenvalues for this study and employ the linear re-weighting to balance their impact. In addition to analyzing average shape distances across groups we also compute a rank by relating the distance between twins to the distance to all other subjects in the dataset (236 in total).

7.2. Results

Fig. 11 illustrates the distances and ranks for the 15 most significant brain structures both with and without size normalization of the eigenvalues. We report the mean and standard error, together with Bonferroni corrected *p*-values of the group differences between MZ and DZ twins

using non-parametric testing (Wilcoxon rank-sum test). The cortical structures, and particularly the white matter, show the most significant shape differences between MZ and DZ twins. Furthermore, the shape distances on the volumetric (tetrahedral) meshes are usually more significant than on the surface (triangular) meshes, which could be due to the isometry invariance of our spectral shape descriptor, which yields an increased deformation flexibility on surfaces. While the order of cortical structures remains relatively stable across normalization and between distance or rank computation, the subcortical structures demonstrate larger variations, but overall similar structures are selected as the most significant ones.

In a final experiment, we compare within-pair differences for brain asymmetry in MZ and DZ twins. As a measure of asymmetry we compute within subject lateral shape distances based on *BrainPrint*. We then compute the absolute difference of this asymmetry measure across twins and compare these differences across the two groups. Here, we did not find any significant differences in brain asymmetry between MZ and DZ twins. This is consistent with recent findings (Bishop, 2013; Eyler et al., 2014), where little support for strong genetic influences on asymmetry was found.

7.3. Discussion

The comparison of distances between subjects may depend on additional factors, such as age, sex, and environment. The advantage of studying twin data is that most of these confounding factors are controlled because MZ and DZ twins have the same age, experienced a similar environment, and, in this study, have the same sex. The bar plots in Fig. 11 consistently show significantly lower shape distances and ranks for MZ twins compared to DZ twins independent of eigenvalue normalization. While we limit the listing to the 15 most significant structures, the vast majority of structures in BrainPrint are significant at a level of 0.05. These results suggest that *BrainPrint* may be capable of capturing genetic influences on brain morphology. Generally, the lack of evidence for asymmetry differences across groups may indicate that lateral shape distances are mainly influenced by fluctuating asymmetry, which was suggested by Gómez-Robles et al. (2013) to be related to environmental factors during brain development rather than genetic factors.

Genetic influences play a substantial role in explaining individual differences in brain structure. Studying such influences is an important element for understanding both normal brain development and pathological brain development in genetically-based disorders. Most imaging genetic studies to date, including twin studies, have examined genetic and environmental influences on brain structures, but not on shape. Shape changes may be associated with development, aging, cognition, and other factors. Thus, a complete understanding of the genetic of brain changes should include examination of changes in shape. Our preliminary analysis of MZ and DZ twins suggests that the characterization of brain morphology with *BrainPrint* is well worth pursuing for a more complete genetic analysis. Further work is required to formally evaluate the extent of genetic and environmental influences and to quantify heritability for shape or asymmetry, and to evaluate whether heritabilities for different structures are significantly different from one another. Such work will be an important step toward demonstrating the potential of shape analysis for gene association.

8. Conclusions

We have introduced *BrainPrint*, an extensive characterization of brain morphology. It is computed by solving the eigenvalue problem of the Laplace-Beltrami operator on meshes from cortical and subcortical brain structures. The classification rate of 99.9% for identification of subjects shows that BrainPrint captures discriminative information about a subject's morphology. We continued studying the similarity structure that BrainPrint imposes on the dataset by evaluating its correlation with non-imaging data that is associated with the scan. Our results show that shape information yields a more accurate prediction of sex and age than volume measures. Further, the shape of cortical structures is better suited for sex prediction, while subcortical structures perform better for age prediction. This highlights the richness of the descriptor and the possibility to adapt the notion of similarity to the application by focusing on different brain structures. In our analysis for brain asymmetry, we found large withinsubject differences across hemispheres, when comparing to the population-based distances. For a majority of brain structures, the structure in one hemisphere is more similar to the structure in the other hemisphere for 10% of the dataset than for the same subject. Moreover, we found an increase in asymmetry with age and a further increase from women to men. The strongest increase in asymmetry was found in the striatum, which we analyzed in more detail with a multicohort longitudinal model to differentiate between within-subject and across-subject changes. The results for the twin study indicate genetic influences on brain morphology.

BrainPrint contains shape information from triangular and tetrahedral meshes, where previous applications of shapeDNA mainly focused on surface meshes. In several of our applications, we note a better performance for the Laplace spectra for volumetric meshes than for surface meshes. The volumetric spectra yield a higher classification rate for subject identification, a higher accuracy for sex prediction, and more significant differences in shape distance between MZ and DZ twins. These results show that the theoretically higher discriminative power of volumetric shapeDNA than surface shapeDNA also yields improvements in practice. Laplace spectra for volumetric meshes could be more discriminative because the spectra on objects' surfaces can be identical (or very similar) al-



Figure 11: Distances and ranks between monozygotic (MZ) and dizygotic (DZ) twins for BrainPrint structures, sorted in decreasing order of significance. Bars show the mean of distances and ranks, error bars correspond to two standard error. Numbers show the Bonferroni corrected *p*-value of the Wilcoxon rank-sum test that distances/ranks are samples from distributions with equal medians. We present results with and without normalization of eigenvalues. The total number of scans is 236.

though the objects are different. Since Laplace spectra are isometry invariant (invariant to length preserving transformations), inward and outward pointing protrusions can lead to highly similar spectra as distances on the 2D manifold remain stable.

The processing framework of BrainPrint is advantageous when working with large datasets widely available today. The segmentation of brain structures and the computation of the compact BrainPrint descriptor is performed independently for each scan. The cost for computing BrainPrint only grows linearly with the number of subjects. All further processing steps of the statistical analysis, which may involve the computation of pairwise distances and therefore higher order computational complexity, are performed on the compact representation. It is exactly this quadratic growth in costs that makes the computation of pairwise similarities with image registration prohibitively expensive for large datasets. BrainPrint and appropriate variations are therefore highly relevant for handling large datasets, not necessarily limited to neuroscience.

9. Acknowledgements

We thank the reviewers for helpful suggestions to improve the manuscript. Support for this research was provided in part by the Humboldt foundation, the National Cancer Institute (1K25-CA181632-01), the Martinos Center for Biomedical Imaging (P41-RR014075, P41-EB015896), the National Alliance for Medical Image Computing (U54-EB005149), the NeuroImaging Analysis Center (P41-EB015902) the National Center for Research Resources (U24 RR021382), the National Institute for Biomedical Imaging and Bioengineering (5P41EB015896-15, R01EB006758), the National Institute on Aging (AG022381, 5R01AG008122-22, AG018344, AG018386), the National Center for Alternative Medicine (RC1 AT005728-01), the national Institute for Neurological Disorders and Stroke (R01 NS052585-01, 1R21NS072652-01, 1R01NS070963, R01NS083534), and was made possible by the resources provided by Shared Instrumentation Grants 1S10RR023401, 1S10RR019307, and 1S10RR023043. Additional support was provided by The Autism & Dyslexia Project funded by the Ellison Medical Foundation, and by the NIH Blueprint for Neuroscience Research (5U01-MH093765), part of the multi-institutional Human Connectome Project. MR was additionally supported by a grant from the Massachusetts Alzheimer's Disease Research Center (5 P50 AG005134), the MGH Neurology Clinical Trials Unit, and the Harvard NeuroDiscovery Center. In addition, BF has a financial interest in CorticoMetrics, a company whose medical pursuits focus on brain imaging and measurement technologies. BF's interests were reviewed and are managed by Massachusetts General Hospital and Partners HealthCare in accordance with their conflict of interest policies.

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

10. Appendix

In the appendix, we provide additional figures to describe the core functionality of the proposed method and to extend the experimental evaluation. Fig. 12 illustrates the linear growth of eigenvalues and the quadratic growth of the variance. The figure also shows that the mean and variance on re-weighted eigenvalues no longer exhibit the original growth pattern. Fig. 13 shows two exemplar scans for which BrainPrint does not correctly identify the subject identity. These scans exhibit imaging artifacts, resulting in skull strip and segmentation errors. Manual correction in FreeSurfer or reacquisition to avoid motion artifacts can be expected to improve the above results. Fig. 14 reports the mean and standard deviation of the volume and local gyrification index calculated from several longitudinal scans per subject. The large variance and overlap across subjects indicates that such representations are not well suited for identifying subjects.

11. References

Aljabar, P., Rueckert, D., Crum, W. R., 2008. Automated morphological analysis of magnetic resonance brain imaging using spectral analysis. Neuroimage 43 (2), 225–235.



Figure 12: Mean (left) and variance (right) of eigenvalues with and without linear re-weighting.



Figure 13: Coronal and axial slices from two misclassified scans. White matter segmentation is shown in yellow. Imaging artifacts cause errors in skull stripping and segmentation.

- Aljabar, P., Wolz, R., Srinivasan, L., Counsell, S., Boardman, J. P., Murgasova, M., Doria, V., Rutherford, M. A., Edwards, A. D., Hajnal, J. V., et al., 2010. Combining morphological information in a manifold learning framework: application to neonatal mri. In: Medical Image Computing and Computer-Assisted Intervention– MICCAI 2010. Springer, pp. 1–8.
- Ashburner, J., 2007. A fast diffeomorphic image registration algorithm. Neuroimage 38 (1), 95–113.
- Attene, M., Falcidieno, B., 2006. Remesh: An interactive environment to edit and repair triangle meshes. In: Shape Modeling and Applications, 2006. IEEE International Conference on. IEEE, pp. 41–41.
- Batchelor, P. G., Castellano Smith, A. D., Hill, D. L. G., Hawkes, D. J., Cox, T. C. S., Dean, A., 2002. Measures of folding applied to the development of the human fetal brain. Medical Imaging, IEEE Transactions on 21 (8), 953–965.
- Bates, J., Pafundi, D., Kanel, P., Liu, X., Mio, W., 2011. Spectral signatures of point clouds and applications to detection of alzheimer's disease through neuroimaging. In: IEEE International Symposium on Biomedical Imaging. pp. 1851–1854.
- Bernardis, E., Konukoglu, E., Ou, Y., Metaxas, D., Desjardins, B., Pohl, K., 2012. Temporal shape analysis via the spectral signature. In: Medical Image Computing and Computer-Assisted Intervention. pp. 49–56.
- Bishop, D. V., 2013. Cerebral asymmetry and language development: cause, correlate, or consequence? Science 340 (6138).
- Blokland, G. A., de Zubicaray, G. I., McMahon, K. L., Wright, M. J., 2012. Genetic and environmental influences on neuroimaging phenotypes: a meta-analytical perspective on twin imaging studies. Twin Research and Human Genetics 15 (03), 351–371.

Buckner, R. L., Head, D., Parker, J., Fotenos, A. F., Marcus, D.,

Morris, J. C., Snyder, A. Z., 2004. A unified approach for morphometric and functional data analysis in young, old, and demented adults using automated atlas-based head size normalization: reliability and validation against manual measurement of total intracranial volume. Neuroimage 23 (2), 724–738.

- Chance, S. A., Crow, T. J., 2007. Distinctively human: cerebral lateralisation and language in homo sapiens. J Anthropol Sci 85, 83–100.
- Dale, A. M., Fischl, B., Sereno, M. I., 1999. Cortical surface-based analysis: I. segmentation and surface reconstruction. Neuroimage 9 (2), 179–194.
- Dale, A. M., Sereno, M. I., 1993. Improved localizadon of cortical activity by combining eeg and meg with mri cortical surface reconstruction: A linear approach. Journal of cognitive neuroscience 5 (2), 162–176.
- Davatzikos, C., Vaillant, M., Resnick, S. M., Prince, J. L., Letovsky, S., Bryan, R. N., 1996. A computerized approach for morphological analysis of the corpus callosum. Journal of computer assisted tomography 20 (1), 88–97.
- DeCarli, C., Massaro, J., Harvey, D., Hald, J., Tullberg, M., Au, R., Beiser, A., DÕAgostino, R., Wolf, P. A., 2005. Measures of brain morphology and infarction in the framingham heart study: establishing what is normal. Neurobiology of aging 26 (4), 491– 510.
- Desikan, R. S., Ségonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., Buckner, R. L., Dale, A. M., Maguire, R. P., Hyman, B. T., et al., 2006. An automated labeling system for subdividing the human cerebral cortex on mri scans into gyral based regions of interest. Neuroimage 31 (3), 968–980.
- Duchesnay, E., Cachia, A., Roche, A., Rivière, D., Cointepas, Y., Papadopoulos-Orfanos, D., Zilbovicius, M., Martinot, J.-L.,



Figure 14: Mean and standard deviation of the volume (left) and of the mean local gyrification index (right) of the cortex for 40 subjects. Statistics are computed over several longitudinal scans per subject.

Régis, J., Mangin, J.-F., 2007. Classification based on cortical folding patterns. Medical Imaging, IEEE Transactions on 26 (4), 553–565.

- Eaves, L., Last, K., Young, P., Martin, N., 1978. Model-fitting approaches to the analysis of human behaviour. Heredity 41 (3), 249–320.
- Eyler, L. T., Vuoksimaa, E., Panizzon, M. S., Fennema-Notestine, C., Neale, M. C., Chen, C.-H., Jak, A., Franz, C. E., Lyons, M. J., Thompson, W. K., et al., 2014. Conceptual and data-based investigation of genetic influences and brain asymmetry: A twin study of multiple structural phenotypes. Journal of cognitive neuroscience 26 (5), 1100–1117.
- Fischl, B., Salat, D. H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., Montillo, A., Makris, N., Rosen, B., Dale, A. M., 2002. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. Neuron 33 (3), 341–355.
- Fischl, B., Sereno, M. I., Dale, A. M., 1999a. Cortical surface-based analysis: Ii: Inflation, flattening, and a surface-based coordinate system. Neuroimage 9 (2), 195–207.
- Fischl, B., Sereno, M. I., Tootell, R. B., Dale, A. M., et al., 1999b. High-resolution intersubject averaging and a coordinate system for the cortical surface. Human brain mapping 8 (4), 272–284.
- Gaser, C., Franke, K., Klöppel, S., Koutsouleris, N., Sauer, H., Initiative, A. D. N., et al., 2013. Brainage in mild cognitive impaired patients: predicting the conversion to Alzheimer's disease. PloS one 8 (6), e67346.
- Gerardin, E., Chételat, G., Chupin, M., Cuingnet, R., Desgranges, B., Kim, H.-S., Niethammer, M., Dubois, B., Lehéricy, S., Garnero, L., et al., 2009. Multidimensional classification of hippocampal shape features discriminates alzheimer's disease and mild cognitive impairment from normal aging. Neuroimage 47 (4), 1476– 1486.
- Gerber, S., Tasdizen, T., Fletcher, P. T., Joshi, S., Whitaker, R., 2010. Manifold modeling for brain population analysis. Medical Image Analysis 14 (5), 643 – 653.
- Gerig, G., Styner, M., Jones, D., Weinberger, D., Lieberman, J., 2001a. Shape analysis of brain ventricles using spharm. In: Mathematical Methods in Biomedical Image Analysis, 2001. MMBIA 2001. IEEE Workshop on. IEEE, pp. 171–178.
- Gerig, G., Styner, M., Shenton, M. E., Lieberman, J. A., 2001b. Shape versus size: Improved understanding of the morphology of brain structures. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2001. Springer, pp. 24–32.
- Germanaud, D., Lefèvre, J., Fischer, C., Bintner, M., Curie, A., des Portes, V., Eliez, S., Elmaleh-Bergès, M., Lamblin, D., Passemard, S., et al., 2014. Simplified gyral pattern in severe developmental microcephalies? new insights from allometric modeling

for spatial and spectral analysis of gyrification. NeuroImage 102, 317–331.

- Germanaud, D., Lefèvre, J., Toro, R., Fischer, C., Dubois, J., Hertz-Pannier, L., Mangin, J.-F., 2012. Larger is twistier: Spectral analysis of gyrification (spangy) applied to adult brain size polymorphism. Neuroimage 63 (3), 1257–1272.
- Geschwind, N., Galaburda, A. M., 1985. Cerebral lateralization: Biological mechanisms, associations, and pathology. Archives of neurology 42 (5), 428–459.
- Geuzaine, C., Remacle, J.-F., 2009. Gmsh: A 3-d finite element mesh generator with built-in pre-and post-processing facilities. International Journal for Numerical Methods in Engineering 79 (11), 1309–1331.
- Golland, P., Grimson, W. E. L., Shenton, M. E., Kikinis, R., 2005. Detection and analysis of statistical differences in anatomical shape. Medical image analysis 9 (1), 69–86.
- Gómez-Robles, A., Hopkins, W. D., Sherwood, C. C., 2013. Increased morphological asymmetry, evolvability and plasticity in human brain evolution. Proceedings of the Royal Society B: Biological Sciences 280 (1761), 20130575.
- Good, C. D., Johnsrude, I. S., Ashburner, J., Henson, R. N., Fristen, K., Frackowiak, R. S., 2002. A voxel-based morphometric study of ageing in 465 normal adult human brains. In: Biomedical Imaging, 2002. 5th IEEE EMBS International Summer School on. IEEE, pp. 16–pp.
- Gordon, C., Webb, D. L., Wolpert, S., 1992. One cannot hear the shape of a drum. Bulletin of the American Mathematical Society 27 (1), 134–138.
- Greve, D. N., Van der Haegen, L., Cai, Q., Stufflebeam, S., Sabuncu, M. R., Fischl, B., Brysbaert, M., 2013. A surface-based analysis of language lateralization and cortical asymmetry. Journal of cognitive neuroscience 25 (9), 1477–1492.
- Halpern, D. F., 2000. Sex Differences in Cognitive Abilities. Psychology Press.
- Hamm, J., Ye, D. H., Verma, R., Davatzikos, C., 2010. Gram: A framework for geodesic registration on anatomical manifolds. Med. Image Analysis 14 (5), 633 – 642.
- Hartigan, J. A., Hartigan, P., 1985. The dip test of unimodality. The Annals of Statistics, 70–84.
- Hugdahl, K., 2005. Symmetry and asymmetry in the human brain. European Review 13 (S2), 119–133.
- Im, K., Pienaar, R., Lee, J.-M., Seong, J.-K., Choi, Y. Y., Lee, K. H., Grant, P. E., 2011. Quantitative comparison and analysis of sulcal patterns using sulcal graph matching: A twin study. Neuroimage 57 (3), 1077–1086.
- Jernigan, T. L., Archibald, S. L., Fennema-Notestine, C., Gamst, A. C., Stout, J. C., Bonner, J., Hesselink, J. R., 2001. Effects of age on tissues and regions of the cerebrum and cerebellum.

Neurobiology of aging 22(4), 581-594.

- Kac, M., 1966. Can one hear the shape of a drum? American Mathematical Monthly, 1–23.
- Konukoglu, E., Glocker, B., Criminisi, A., Pohl, K. M., 2013. WESD–Weighted spectral distance for measuring shape dissimilarity. Pattern Analysis and Machine Intelligence, IEEE Transactions on 35 (9), 2284–2297.
- Kremen, W. S., Prom-Wormley, E., Panizzon, M. S., Eyler, L. T., Fischl, B., Neale, M. C., Franz, C. E., Lyons, M. J., Pacheco, J., Perry, M. E., et al., 2010. Genetic and environmental influences on the size of specific brain regions in midlife: the vetsa mri study. Neuroimage 49 (2), 1213–1223.
- Kremen, W. S., Thompson-Brenner, H., Leung, Y.-M. J., Grant, M. D., Franz, C. E., Eisen, S. A., Jacobson, K. C., Boake, C., Lyons, M. J., 2006. Genes, environment, and time: the vietnam era twin study of aging (vetsa). Twin Research and Human Genetics 9 (06), 1009–1022.
- Lai, R., Shi, Y., Dinov, I., Chan, T. F., Toga, A. W., 2009. Laplacebeltrami nodal counts: A new signature for 3d shape analysis. In: Biomedical Imaging: From Nano to Macro, 2009. ISBI'09. IEEE International Symposium on. IEEE, pp. 694–697.
- Lefevre, J., Auzias, G., Germanaud, D., 2014. Brain lobes revealed by spectral clustering. In: International Conference on Pattern Recognition.
- Lehoucq, R. B., Sorensen, D. C., Yang, C., 1998. ARPACK users' guide: solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods. Vol. 6. Siam.
- Lian, Z., Godil, A., Bustos, B., Daoudi, M., Hermans, J., Kawamura, S., Kurita, Y., Lavoué, G., Van Nguyen, H., Ohbuchi, R., et al., 2012. A comparison of methods for non-rigid 3D shape retrieval. Pattern Recognition 46, 449–461.
- Lombaert, H., Grady, L., Polimeni, J. R., Cheriet, F., 2013. Focusr: Feature oriented correspondence using spectral regularization–a method for precise surface matching. Pattern Analysis and Machine Intelligence, IEEE Transactions on 35 (9), 2143–2160.
- Luders, E., Thompson, P., Narr, K., Toga, A., Jancke, L., Gaser, C., 2006. A curvature-based approach to estimate local gyrification on the cortical surface. Neuroimage 29 (4), 1224–1230.
- Miller, M., Banerjee, A., Christensen, G., Joshi, S., Khaneja, N., Grenander, U., Matejic, L., 1997. Statistical methods in computational anatomy. Statistical methods in medical research 6 (3), 267–299.
- Minakshisundaram, S., Pleijel, A., 1967. Some properties of the eigenvalues of the laplacian. J. Difffl. Geom 1, 43–69.
- Neale, M., Cardon, L., 1992. Methodology for genetic studies of twins and families. No. 67. Springer.
- Ng, B., Toews, M., Durrleman, S., Shi, Y., 2014. Shape analysis for brain structures. In: Shape Analysis in Medical Image Analysis. Springer, pp. 3–49.
- Niethammer, M., Reuter, M., Wolter, F.-E., Bouix, S., Peinecke, N., Koo, M.-S., Shenton, M., 2007. Global medical shape analysis using the Laplace-Beltrami spectrum. In: Ayache, N., Ourselin, S., Maeder, A. J. (Eds.), MICCAI 2007. Vol. 4791 of LNCS. Springer, Heidelberg, pp. 850–857.
- Panizzon, M. S., Fennema-Notestine, C., Kubarych, T. S., Chen, C.-H., Eyler, L. T., Fischl, B., Franz, C. E., Grant, M. D., Hamza, S., Jak, A., et al., 2012. Genetic and environmental influences of white and gray matter signal contrast: A new phenotype for imaging genetics? Neuroimage 60 (3), 1686–1695.
- Pepe, A., Dinov, I., Tohka, J., 2014. An automatic framework for quantitative validation of voxel based morphometry measures of anatomical brain asymmetry. NeuroImage 100, 444–459.
- Peper, J. S., Brouwer, R. M., Boomsma, D. I., Kahn, R. S., Pol, H., Hilleke, E., 2007. Genetic influences on human brain structure: a review of brain imaging studies in twins. Human brain mapping 28 (6), 464–473.
- Raz, N., Rodrigue, K. M., 2006. Differential aging of the brain: patterns, cognitive correlates and modifiers. Neuroscience & Biobehavioral Reviews 30 (6), 730–748.
- Raz, N., Rodrigue, K. M., Kennedy, K. M., Head, D., Gunning-Dixon, F., Acker, J. D., 2003. Differential aging of the human

striatum: longitudinal evidence. American Journal of Neuroradiology 24 (9), 1849–1856.

- Reuter, M., 2006. Laplace Spectra for Shape Recognition. Books on Demand GmbH.
- Reuter, M., Fischl, B., 2011. Avoiding asymmetry-induced bias in longitudinal image processing. NeuroImage 57 (1), 19-21. URL http://dx.doi.org/10.1016/j.neuroimage.2011.02.076
- Reuter, M., Niethammer, M., Wolter, F.-E., Bouix, S., Shenton, M., 2007. Global medical shape analysis using the volumetric Laplace spectrum. In: International Conference on Cyberworlds, NASA-GEM Workshop. pp. 417–426.
- Reuter, M., Rosas, H. D., Fischl, B., 2010. Highly accurate inverse consistent registration: A robust approach. NeuroImage 53 (4), 1181–1196.

URL http://dx.doi.org/10.1016/j.neuroimage.2010.07.020

Reuter, M., Schmansky, N. J., Rosas, H. D., Fischl, B., 2012. Withinsubject template estimation for unbiased longitudinal image analysis. NeuroImage 61 (4), 1402–1418.

URL http://dx.doi.org/10.1016/j.neuroimage.2012.02.084

- Reuter, M., Wolter, F.-E., Peinecke, N., 2006. Laplace-Beltrami spectra as "Shape-DNA" of surfaces and solids. Computer-Aided Design 38 (4), 342–366.
- Reuter, M., Wolter, F.-E., Shenton, M., Niethammer, M., 2009. Laplace-Beltrami eigenvalues and topological features of eigenfunctions for statistical shape analysis. Computer-Aided Design 41 (10), 739–755.
- Schaer, M., Cuadra, M. B., Tamarit, L., Lazeyras, F., Eliez, S., Thiran, J., 2008. A surface-based approach to quantify local cortical gyrification. Medical Imaging, IEEE Transactions on 27 (2), 161– 170.
- Schaie, K. W., Caskie, G. I., 2005. Methodological issues in aging research. Handbook of research methods in developmental science, 21–39.
- Schmitt, J. E., Eyler, L. T., Giedd, J. N., Kremen, W. S., Kendler, K. S., Neale, M. C., 2007. Review of twin and family studies on neuroanatomic phenotypes and typical neurodevelopment. Twin Research and Human Genetics 10 (05), 683–694.
- Seo, S., Chung, M. K., 2011. Laplace-beltrami eigenfunction expansion of cortical manifolds. In: Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on. IEEE, pp. 372– 375.
- Shen, K.-k., Fripp, J., Mériaudeau, F., Chételat, G., Salvado, O., Bourgeat, P., 2012. Detecting global and local hippocampal shape changes in alzheimer's disease using statistical shape models. Neuroimage 59 (3), 2155–2166.
- Styner, M., Lieberman, J. A., McClure, R. K., Weinberger, D. R., Jones, D. W., Gerig, G., 2005. Morphometric analysis of lateral ventricles in schizophrenia and healthy controls regarding genetic and disease-specific factors. Proceedings of the National Academy of Sciences of the United States of America 102 (13), 4872–4877.
- Terriberry, T. B., Joshi, S. C., Gerig, G., 2005. Hypothesis testing with nonlinear shape models. In: Information processing in medical imaging. Springer, pp. 15–26.
- Thompson, P. M., Hayashi, K. M., Dutton, R. A., CHIANG, M.-C., Leow, A. D., Sowell, E. R., De Zubicaray, G., Becker, J. T., Lopez, O. L., Aizenstein, H. J., et al., 2007. Tracking alzheimer's disease. Annals of the New York Academy of Sciences 1097 (1), 183–214.
- Thompson, W. K., Hallmayer, J., O'Hara, R., 2011. Design considerations for characterizing psychiatric trajectories across the lifespan: application to effects of apoe-e4 on cerebral cortical thickness in alzheimer's disease. American Journal of Psychiatry 168 (9), 894–903.
- Toga, A. W., Thompson, P. M., 2003. Mapping brain asymmetry. Nature Reviews Neuroscience 4 (1), 37–48.
- Wachinger, C., Batmanghelich, K., Golland, P., Reuter, M., 2014a. Brainprint in the computer-aided diagnosis of alzheimer's disease. In: Challenge on Computer-Aided Diagnosis of Dementia, MIC-CAI.
- Wachinger, C., Golland, P., Reuter, M., 2014b. Brainprint: Identifying subjects by their brain. In: Medical Image Computing and Computer-Assisted Intervention-MICCAI 2014. Springer, pp. 41–

48.

- Walhovd, K. B., Fjell, A. M., Reinvang, I., Lundervold, A., Dale, A. M., Eilertsen, D. E., Quinn, B. T., Salat, D., Makris, N., Fischl, B., 2005. Effects of age on volumes of cortex, white matter and subcortical structures. Neurobiology of aging 26 (9), 1261–1270.
- Weyl, H., 1911. Über die asymptotische verteilung der eigenwerte. Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse 1911, 110–117.
- Weyl, H., 1912. Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen (mit einer Anwendung auf die Theorie der Hohlraumstrahlung). Mathematische Annalen 71 (4), 441–479.
- Wolz, R., Aljabar, P., Hajnal, J., Hammers, A., Rueckert, D., 2010. LEAP: Learning embeddings for atlas propagation. NeuroImage 49 (2), 1316–1325.
- Zhu, P., Awate, S. P., Gerber, S., Whitaker, R., 2011. Fast shapebased nearest-neighbor search for brain MRIs using hierarchical feature matching. In: Medical Image Computing and Computer-Assisted Intervention-MICCAI 2011. Springer, pp. 484–491.