

BrainPrint: Identifying Subjects by their Brain

Christian Wachinger^{1,2}, Polina Golland¹, Martin Reuter^{1,2}

¹Computer Science and Artificial Intelligence Lab, MIT

²Massachusetts General Hospital, Harvard Medical School

Abstract. Introducing *BrainPrint*, a compact and discriminative representation of anatomical structures in the brain. *BrainPrint* captures shape information of an ensemble of cortical and subcortical structures by solving the 2D and 3D Laplace-Beltrami operator on triangular (boundary) and tetrahedral (volumetric) meshes. We derive a robust classifier for this representation that identifies the subject in a new scan, based on a database of brain scans. In an example dataset containing over 3000 MRI scans, we show that *BrainPrint* captures unique information about the subject’s anatomy and permits to correctly classify a scan with an accuracy of over 99.8%. All processing steps for obtaining the compact representation are fully automated making this processing framework particularly attractive for handling large datasets.

1 Introduction

Is it possible to identify an individual based on their brain? Are cortical folding patterns unique to a person, similar to a fingerprint? While the unique complexity of the brain may indicate that an unambiguous identification should be possible, there is currently little empirical research that can speak to these questions. 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 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 permit a robust identification.

Here, we introduce *BrainPrint*, a holistic representation of the brain anatomy, containing the shape information of an ensemble of cortical and subcortical structures. The inclusion of only shape information has the advantage to remain independent from the local intensity values in the scan. Moreover, the variety of the different structures included in the *BrainPrint* yields an extensive characterization of the brain anatomy. We quantify the shape information by calculating the spectrum of the Laplace-Beltrami operator (LBO) on both triangular meshes that represent boundary surfaces, e.g., the white matter surface, and tetrahedral

meshes for volumetric representations of individual structures. We then derive a classifier that identifies a subject from an MRI scan based on its *BrainPrint*. We achieve robustness in the identification 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 also determines whether a query brain belongs to an unknown subject, not yet represented in the existing database.

An alternative approach to calculate the similarity between scans could be based on image registration [3, 5]. However, real applications of such identification methods require large datasets and the cost for aligning a new scan to all scans in the database becomes prohibitive for a large number of scans. *BrainPrint* introduces a new framework that is especially beneficial when working with large datasets widely available today. 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 takes less memory and permits easier calculations and comparisons than the original scan.

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. [10] 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 was capable of distinguishing two isospectral 3D solids (GWW-prisms). Therefore, we propose to combine the information from both the 3D solid and 2D boundary shape representations.

While there has been previous work analyzing the shapeDNA for single brain structures [1, 9, 11], to the best of our knowledge this is the first study that evaluates its application to cortical structures 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 [1, 8], 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.

2 Shape Descriptor

We segment anatomical structures from brain scans with FreeSurfer [2]. Next, we compute a compact shape representation that captures important shape information and facilitates the further processing. Since image intensity varies across

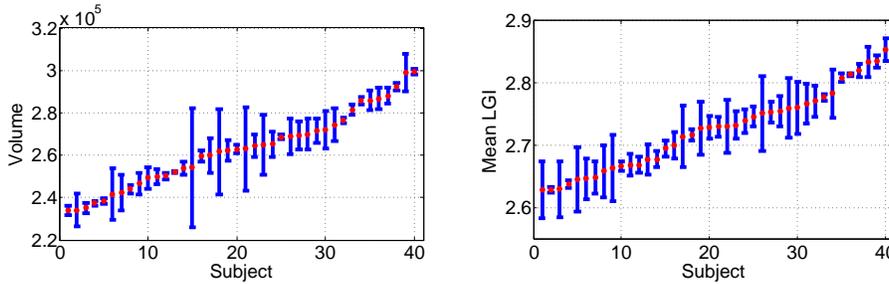


Fig. 1: Mean and standard deviation of the volume (left) and of the mean local gyrification index (right) of the cortex for 40 subjects. Statistics are calculated over several longitudinal scans per subject.

scans, we focus on geometrical properties. Example representations are the volume and the local gyrification index (LGI) of a structure. 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. The LGI was used previously to identify gyral abnormalities [12]. We transform this local measure into a global shape descriptor by computing the mean LGI over the surface. Fig. 1 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.

In this work we use the shapeDNA [10] as a shape descriptor, which performed among the best in a recent comparison of methods for non-rigid 3D shape retrieval [6]. 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) $\Delta f = -\lambda f$ using the finite element method. The solution consists of eigenvalue $\lambda_i \in \mathbb{R}$ and eigenfunction f_i pairs (sorted by eigenvalues, $0 \leq \lambda_1 \leq \lambda_2 \leq \dots$). To be independent of the objects' scale, we normalize the eigenvalues $\lambda' = \text{vol}^{\frac{2}{D}} \lambda$, where vol is the Riemannian volume of the D -dimensional manifold (i.e., the area for 2D surfaces) [10]. The first l non-zero eigenvalues form the shapeDNA: $\boldsymbol{\lambda} = (\lambda'_1, \dots, \lambda'_l)$.

The eigenvalues are isometry invariant with respect to the Riemannian manifold, meaning that length-preserving deformations will not change the spectrum. This important property permits the comparison of subjects by directly comparing the shapeDNA, without the need for alignment. 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). A second property is that the spectrum continuously changes with topology-preserving deformations of the object boundary. Fig. 2 illustrates the eigenfunctions of the cerebral cortex boundary. The eigenfunctions show natural

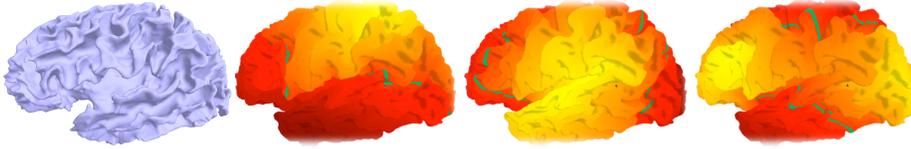


Fig. 2: Left cerebral cortex and first eigenfunctions of the LBO calculated on the surface (yellow – positive, red – negative, and green – zero).

vibrations of the shape when oscillating at a frequency specified by the square root of the eigenvalue.

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 (white and pial surfaces in both hemispheres), forming the *BrainPrint* $\Lambda = (\lambda_1, \dots, \lambda_\eta)$. 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, uniformly resample the output to 60K vertices, and create the volumetric mesh with the gmsh package [4]. We use the linear finite element method [10] with Neumann boundary condition (zero normal derivative) to compute the spectra of the tetrahedral meshes.

3 Classifier

We derive a classifier to assign a new scan to one of the subjects in the database. Since the segmentation or tessellation of specific structures may fail in certain cases, we propose a robust classifier that handles missing information. We build a classifier by combining the results from weak classifiers operating on specific brain structures.

Assuming n subjects $\mathcal{C}_1, \dots, \mathcal{C}_n$ and N scans in a database ($N \geq n$, for repeated scans of subjects). Each scan has its associated *BrainPrint* $\Lambda_1, \dots, \Lambda_N$. Let $S_k \subset \{1, \dots, N\}$ denote scans for subject \mathcal{C}_k . The probability that a the new scan with *BrainPrint* Λ shows subject \mathcal{C}_k is

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

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 . Since we only have a few samples per class, we estimate a global diagonal covariance matrix Σ_s across all scans for each structure. Weighting distances by the variance helps to prevent the domination by higher eigenvalues

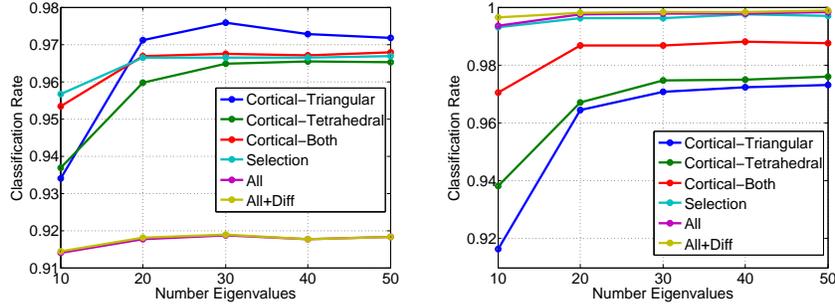


Fig. 3: Classification results for product classifier (left) and voting classifier (right) under variable number of eigenvalues and feature sets.

that exhibit higher variation. The subject identity with the highest probability is assigned to the scan

$$k^* = \arg \max_k p(\mathcal{C}_k | \mathcal{A}). \quad (2)$$

The posterior probability of this classifier is the product of the posterior probabilities across all structures, cf. Eq. (1), which may be problematic for structures with low discriminative power. Many subcortical structures do not carry much distinctive shape information and can therefore negatively influence the overall probability. We therefore 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

$$k_s^* = \arg \max_k p(\lambda_s | \mathcal{C}_k), \quad \forall s \in \{1, \dots, \eta\}, \quad (3)$$

with the final vote set to the mode of the vote distribution.

4 Results

We perform experiments on data from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) [7]. We work with over 3000 scans from almost 700 subjects, where each subject has between three and six longitudinal scans. Each T1-weighted image from the dataset is processed independently with FreeSurfer. We calculate 36 shape descriptors for subcortical structures and 8 descriptors for cortical structures (left/right, white/gray matter, 2D/3D). Additionally, we calculate the lateral differences of shapeDNA between left and right cortical structures to quantify asymmetries, resulting in 4 additional descriptors.

We perform leave-one-out experiments by removing one scan from the dataset and by aiming to recover the correct identity. Fig. 3 reports the classification results for the product classifier in Eq.(2) and the structure-specific voting in Eq.(3). We report classification results as a function of the number of

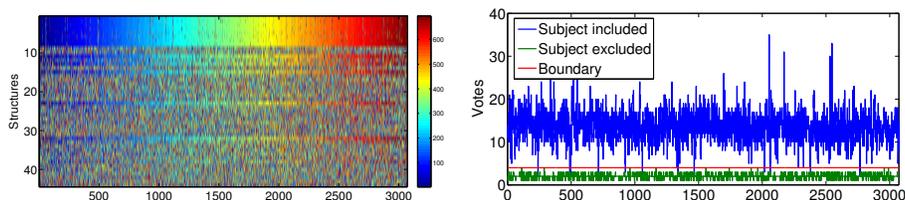


Fig. 4: Left: Subject (color) voted for by each structure (row) for each scan (column). Cortical structures in first 8 rows, subcortical features below. Optimal feature response would show a color gradient from blue to red, since scans are sorted by subject. Right: Number of votes for the winning subject 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.

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), 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 cortical triangular meshes. Adding more features, especially when working with all features, dramatically reduces the classification results. We observe an opposite behavior when working with the structure-specific voting. Subcortical structures alone yield the worst performance in this case. The combination of 3D solid and 2D boundary descriptors leads to a clear improvement. A further improvement is gained by adding subcortical structures.

To further study this behavior, we examine the candidate subject that each structure votes for in Fig. 4. Each column corresponds to one scan and each row to one structure. The color indicates the subject number. Scans were sorted by subject; 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 obscure 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. The best performance of over 99.8% is achieved for 50 eigenvalues on all features with the additional difference features. For comparison, the classification rate for the mean LGI on both hemispheres is 1.0% for the product and 3.9% for the voting classifier. The classification rate 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. 1.

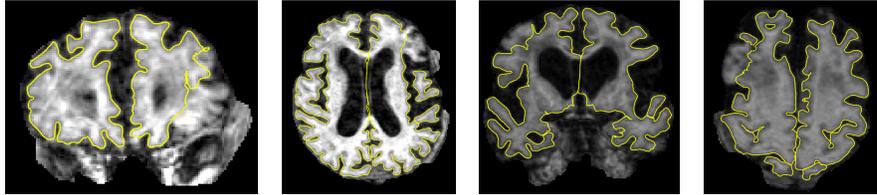


Fig. 5: Coronal and axial slices from two misclassified scans. White matter segmentation is shown in yellow.

Fig. 5 shows the two scans for which *BrainPrint* does not correctly identify the subject identity. These subjects show strong atrophy and imaging artifacts, resulting in pronounced segmentation errors. Manual correction in FreeSurfer or reacquisition to avoid motion artifacts may therefore improve the above results.

As an additional experiment, we evaluate the possibility to determine whether a subject is not contained in the database. We study the number of votes the winning subject receives in Fig. 4, once when the subject of the scan is included in the database and once when the subject is excluded. If the subject in the current scan exists 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.

5 Discussion and Conclusions

The high classification accuracy of *BrainPrint* suggests that brain structures are unique to individuals and can be used for identification. 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.

The 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. Yet, we currently do not think that *BrainPrint* interferes with anonymization because at least a second scan with knowledge of the identity needs to be available to connect to the private information. In terms of its practical applications, we see *BrainPrint* as an aid when handling large datasets. Identifying similar images in an efficient way can provide the launchpad for a more detailed follow-up analysis, *e.g.*, calculation or prediction of localized growth and shrinkage patterns. Since most of our re-

trieval errors are related to incorrect segmentations, our approach could also be used as an automatic quality control. Furthermore, *BrainPrint* can help identify anonymization errors (mismatch of subject identity), which are difficult to detect and can impede longitudinal studies. Finally, the presented framework of image understanding and compact characterization is relevant for handling large datasets in other fields and not limited to neuroscience.

Acknowledgements: This work was supported in part by the Humboldt foundation, the Martinos Center for Biomedical Imaging (P41-RR014075, P41-EB015896), the National Alliance for Medical Image Computing (U54-EB005149) and the NeuroImaging Analysis Center (P41-EB015902). We thank Anna Rieckmann for revising the manuscript and the Alzheimer’s Disease Neuroimaging Initiative (ADNI) for image data.

References

1. Bates, J., Pafundi, D., Kanel, P., Liu, X., Mio, W.: 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 (2011)
2. 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.: Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron* 33(3), 341–355 (2002)
3. Gerber, S., Tasdizen, T., Fletcher, P.T., Joshi, S., Whitaker, R.: Manifold modeling for brain population analysis. *Medical Image Analysis* 14(5), 643 – 653 (2010)
4. Geuzaine, C., Remacle, J.F.: 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 (2009)
5. Hamm, J., Ye, D.H., Verma, R., Davatzikos, C.: Gram: A framework for geodesic registration on anatomical manifolds. *Med. Image Analysis* 14(5), 633 – 642 (2010)
6. Lian, Z., Godil, A., Bustos, B., Daoudi, M., Hermans, J., Kawamura, S., Kurita, Y., Lavoué, G., Van Nguyen, H., Ohbuchi, R., et al.: A comparison of methods for non-rigid 3D shape retrieval. *Pattern Recognition* 46, 449–461 (2012)
7. Mueller, S.G., Weiner, M.W., Thal, L.J., et al: The alzheimer’s disease neuroimaging initiative. *Neuroimaging Clinics of North America* 15(4), 869–877 (2005)
8. Niethammer, M., Reuter, M., Wolter, F.E., Bouix, S., Peinecke, N., Koo, M.S., Shenton, M.: Global medical shape analysis using the Laplace-Beltrami spectrum. In: Ayache, N., Ourselin, S., Maeder, A.J. (eds.) MICCAI 2007. LNCS, vol. 4791, pp. 850–857. Springer, Heidelberg (2007)
9. Reuter, M., Niethammer, M., Wolter, F.E., Bouix, S., Shenton, M.: Global medical shape analysis using the volumetric Laplace spectrum. In: International Conference on Cyberworlds, NASA-GEM Workshop. pp. 417–426 (2007)
10. Reuter, M., Wolter, F.E., Peinecke, N.: Laplace-Beltrami spectra as ”Shape-DNA” of surfaces and solids. *Computer-Aided Design* 38(4), 342–366 (2006)
11. Reuter, M., Wolter, F.E., Shenton, M., Niethammer, M.: Laplace-Beltrami eigenvalues and topological features of eigenfunctions for statistical shape analysis. *Computer-Aided Design* 41(10), 739–755 (2009)
12. Schaer, M., Cuadra, M.B., Tamarit, L., Lazeyras, F., Eliez, S., Thiran, J.: A surface-based approach to quantify local cortical gyrification. *Medical Imaging, IEEE Transactions on* 27(2), 161–170 (2008)