Spectral Learning of Surface Data: Ideas from Medical Imaging

Herve Lombaert
ETS Montreal, Canada
Inria Sophia-Antipolis, France

Antonio Criminisi
Microsoft Research
Cambridge, UK

Nicholas Ayache
Inria Sophia-Antipolis, France

Abstract

This method [1] exploits spectral representations of shapes to classify surface data. This learning approach is relevant for problems involving data living on surfaces, such as in brain surface parcellation. Current neuroimaging studies relies on slow mesh deformation towards pre-labeled atlases, requiring hours of computation. Learning techniques offer an attractive computational advantage, however, their representation of spatial information, typically defined in a Euclidean domain, may be inadequate for analyzing surface data, such as cortical parcellation. Indeed, cortical data resides on surfaces that are highly variable in space and shape. Consequently, Euclidean representations of surface data may be inconsistent across individuals. We propose to fundamentally change the spatial representation of surface data, by exploiting spectral coordinates derived, for instance, from the Laplacian eigenfunctions of shapes. They have the advantage to be *geometry aware*. This change of paradigm, from Euclidean to Spectral representations, enables a classifier to be applied *directly* on surface data via spectral coordinates. As an illustration, we extend one learning method, for instance, the Random Decision Forests. Our spectral learning method is shown to significantly improve the accuracy of cortical parcellation over its standard version.

1 Introduction

Efficient algorithms for surface processing and analysis are often sought. In particular, the accurate segmentation of cortical surfaces into sulcal areas is fundamental to neuroscience. Two strategies exist for cortical parcellation. They are either template based [2, 3, 4, 5, 6], via iterative deformations of a pre-labeled atlas, or subject based, via costly processing of sulcal data [7, 8, 9, 10, 11]. Unfortunately, machine learning techniques in cortical analysis [12] has been limited due to the high geometrical variability of the folding pattern across individuals. They raise the fundamental question on how to learn data *directly* on surfaces. Neighborhood structures, often exploited in image segmentation [13], may be ambiguous on surfaces, and more challenging to interpret on highly convoluted cortical surfaces [14]. Neighboring points in 3D may not necessarily lie on a surface, and be even several folds away on the cortex. Consequently, standard learning approaches that use features defined in the Euclidean domain, may not be adequate for analyzing surface data [15]. We propose to represent instead spatial information with geometry-aware features. The spectral decomposition of shapes [16, 17] provides means to efficiently parameterize cortical surfaces with few spectral coordinates [18]. These spectral coordinates constitute, in fact, an explicit parameterization of surfaces. A learning technique could thus exploit such spectral coordinates to learn data directly on surfaces. In this paper, we improve, for instance, the Random Decision Forests (RF) [19, 20], to process surface data via spectral representations of shapes, and name our method **Spectral Forests (SF)**.

2 Method

Spectral Coordinates – Let us build the graph $\mathscr{G} = \{\mathscr{V}, \mathscr{E}\}$ from the set of vertices with position x, and edges of a surface model S. Its graph Laplacian operator is defined [17] as the matrix $\mathcal{L} = D^{-1}(D - W)$, with D, the node degree matrix, and W, the affinity matrix, e.g., $W_{ij} = \|$

31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA.

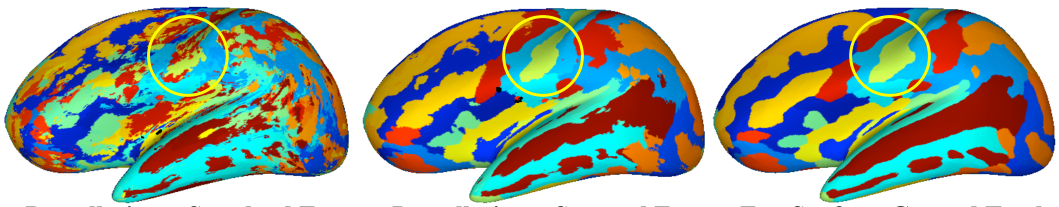

Figure 1: **Spectral Learning of Brain Parcels** – (*Left*) Best cortical parcellation using RF (31.0%), (*Middle*) Parcellation using SF (77.6%, our method), (*Right*) using FS, considered as gold standard. Inflated surfaces show 77 color-coded parcels.

$x_i - x_j \|^{-1}$ if $\exists e_{ij} \in \mathscr{E}$, 0 otherwise. Its spectral decomposition, $\mathcal{L} = U \Lambda U^{-1}$, provides the eigenvalues $\Lambda = \mathrm{diag}(\lambda_0, ... \lambda_{|V|})$ and eigenfunctions $U = (u^{(0)}, ..., u^{(|V|)})$, where $u^{(\cdot)}$ depicts a vibration mode of shape S [16]. The spectral coordinates of points $p \in \mathscr{V}$ are defined as their normalized eigenfunction values [21], $\texttt{spectral}(p) = \{\lambda_0^{-\frac{1}{2}} u^{(0)}(p), ..., \lambda_{|V|}^{-\frac{1}{2}} u^{(|V|)}(p)\}$. Since these coordinates are defined on surfaces, navigating with them would move us over the surface, whereas an increase in Euclidean coordinates may bring us away from the surface. We correct for slight perturbations in shape isometry using a realignment to an arbitrary reference, $\Lambda^{-\frac{1}{2}} U T_{i \mapsto \mathrm{ref}}$, where the transformation T is found, for instance, with [22]. In practice [18, 22], only the first $k = 5$ spectral coordinates are sufficient to capture the main geometrical properties. The spectral coordinates $\texttt{spectral}(p)$ for a point p of subject i, are the first k elements of the p^{th} row of matrix $\Lambda^{-\frac{1}{2}} U T_{i \mapsto \mathrm{ref}}$.

Spectral Learning – We extend one learning technique beyond the Euclidean domain to classify surface data. As an example, we use Random Forests (RF) [19, 20]. We name the method Spectral Forests (SF). To do so, spatial features are represented using spectral coordinates rather than Euclidean coordinates. These surface basis functions are *geometry aware* and have the property to be invariant to shape isometry. Spectral representations effectively capture *intrinsic* shape information. Location and neighborhoods are defined explicitly on surfaces, which contrasts with the implicit representation of surfaces with Euclidean coordinates. This change of paradigm enables Spectral Forests to process data directly on surfaces.

Cortical Parcellation – The labeling of cortical points into major sulci and gyri, is an application where learning should be performed on surfaces. Our Spectral Forests algorithm represents spatial information with spectral coordinates, which naturally parameterize surfaces in an *intrinsic* spectral domain rather than in an *extrinsic* Euclidean space. The simplest form of feature representation could be, for instance, $f_p = (\texttt{depth}(p), \texttt{spectral}(p))$, which includes data information, such as the sulcal depth at each point, and spatial information, where standard (x, y, z) point values are replaced with k spectral coordinates. This change of paradigm enables a standard RF classifier to be applied on the spectral representation f_p for learning and inferring the major parcels over the brain surface.

3 Results & Conclusion

We segmented 77 cortical parcels on 16 brain surfaces, and validated using leave-one-out. **Running standard RF** produces an average overlap, for all 77 parcels on 16 surfaces, of **27.9%** (± 17.0, min/max parcels = 4.9/65.9), with 21 seconds for training and 66 seconds for testing. **Running Spectral Forests (SF)** produces an average overlap of **74.3%** (± 8.32, min/max parcels = 38.9/94.9), with 17 seconds for training and 23 seconds for testing. Fig. 1 shows the best scoring parcellation of RF, with an average overlap of 31.0% (± 15.5), which contrasts with the SF parcellation on the same subject of 77.6% (± 11.41). **In comparison, FreeSurfer (FS)**, which is considered here as a gold standard, performs with an average overlap of **74.4%** (± 9.7, min/max parcels = 41.2/96.6) among all possible transfers of parcellation maps from all subjects onto all possible reference subjects.

The change from *extrinsic* to *intrinsic* shape representations enables learning techniques to process data directly on surfaces. We implemented this strategy using Random Forests. We found that revisiting the fundamentals of spatial representations, from Euclidean to spectral-based features, improves the parcellation accuracy from 27.9% to 74.3%, which is comparable to the present state-of-the-art, but with a clear speed advantage (23 seconds vs. hours).

References

[1] H. Lombaert, A. Criminisi, and N. Ayache. Spectral Forests: Learning of Surface Data, Application to Cortical Parcellation. In *MICCAI*, 2015.

[2] K. J. Behnke, M. E. Rettmann, D. L. Pham, D. Shen, S. M. Resnick, C. Davatzikos, and J. L. Prince. Automatic classification of sulcal regions of the human brain cortex using pattern recognition. *TMI*, 2003.

[3] G. Le Goualher, E. Procyk, D. L. Collins, R. Venugopal, C. Barillot, and A. C. Evans. Automated extraction and variability analysis of sulcal neuroanatomy. *TMI*, 1999.

[4] G. Lohmann and D. Y. von Cramon. Automatic labelling of the human cortical surface using sulcal basins. *Med. Image Anal.*, 2000.

[5] D. Rivière, J.-F. Mangin, D. Papadopoulos-Orfanos, J.-M. Martinez, V. Frouin, and J. Régis. Automatic recognition of cortical sulci of the human brain using a congregation of neural networks. *Med. Image Anal.*, 2002.

[6] B. Fischl, A. van der Kouwe, C. Destrieux, E. Halgren, F. Segonne, D. H. Salat, E. Busa, L. J. Seidman, J. Goldstein, D. Kennedy, V. Caviness, N. Makris, B. Rosen, and A. M. Dale. Automatically parcellating the human cerebral cortex. *Cereb. Cortex*, 2004.

[7] M. E. Rettmann, X. Han, C. Xu, and J. L. Prince. Automated sulcal segmentation using watersheds on the cortical surface. *NeuroImage*, 2002.

[8] F. Yang and F. Kruggel. Automatic segmentation of human brain sulci. *Med. Image Anal.*, 2008.

[9] Y. Shi, Z. Tu, A. L. Reiss, R. A. Dutton, A. D. Lee, A. M. Galaburda, I. Dinov, P. M. Thompson, and A. W. Toga. Joint sulcal detection on cortical surfaces with graphical models and boosted priors. *TMI*, pages 361–73, 2009.

[10] D. W. Shattuck, A. A. Joshi, D. Pantazis, E. Kan, R. A. Dutton, E. R. Sowell, P. M. Thompson, A. W. Toga, and R. M. Leahy. Semi-automated method for delineation of landmarks on models of the cerebral cortex. *Neuroscience*, 2009.

[11] A. Cachia, J.-F. Mangin, D. Rivière, D. Papadopoulos-Orfanos, F. Kherif, I. Bloch, and J. Régis. A generic framework for the parcellation of the cortical surface into gyri using geodesic Voronoï diagrams. *Med. Image Anal.*, 2003.

[12] Z. Tu, S. Zheng, A. L. Yuille, A. L. Reiss, R. A. Dutton, A. D. Lee, A. M. Galaburda, I. Dinov, P. M. Thompson, and A. W. Toga. Automated extraction of the cortical sulci based on a supervised learning approach. *TMI*, 2007.

[13] V. Lempitsky, M. Verhoek, Noble, and A. Blake. Random forest classification for automatic delineation of myocardium in real-time 3D echocardiography. In *FIMH*, 2009.

[14] J. V. Stough, C. Ye, S. H. Ying, and J. L. Prince. Thalamic Parcellation from Multi-modal Data using Random Forests. *ISBI*, 2013.

[15] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning. *IEEE Signal Processing*, 2017.

[16] F. Chung. *Spectral Graph Theory*. AMS, 1996.

[17] L. Grady and J. R. Polimeni. *Discrete Calculus*. Springer, 2010.

[18] H. Lombaert, L. Grady, J. Polimeni, and F. Cheriet. FOCUSR: Feature Oriented Correspondence using Spectral Regularization - A Method for Accurate Surface Matching. *PAMI*, 2012.

[19] L. Breiman. Random forests. *Mach. Learn.*, 45, 2001.

[20] A. Criminisi and J. Shotton. *Decision Forests for Computer Vision and Medical Image Analysis*. Springer, 2013.

[21] R. M. Rustamov. Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Eurographics*, 2007.

[22] H. Lombaert, M. Arcaro, and N. Ayache. Brain Transfer: Spectral Analysis of Cortical Surfaces and Functional Maps. In *IPMI*, 2015.