Graph Convolutions on Spectral Embeddings: Learning of Cortical Surface Data

Karthik Gopinath [*]
ETS Montreal, Canada

Christian Desrosiers
ETS Montreal, Canada

Herve Lombaert
ETS Montreal, Canada

Abstract

This paper presents a novel approach for learning and exploiting surface data directly across surface domains. Direct learning of surface data via graph convolutions provides a new family of fast algorithms for processing brain surfaces. The current limitation of existing state-of-the-art approaches is their inability to compare surface data across different surface domains. Surface bases are indeed incompatible across geometries. This paper leverages recent advances in spectral graph matching to transfer surface data across aligned spectral domains. This enables direct learning of surface data across compatible surface bases. It exploits spectral filters over intrinsic representations of surfaces. We illustrate our approach with an application to brain parcellation and validate the algorithm over 101 manually labeled brain surfaces. The results show a significant improvement in labeling accuracy over recent Euclidean approaches and gaining a drastic speed improvement over conventional methods.

1 Introduction

Neuroimage analysis consists of studying functional and anatomical information over the brain geometry. The thin outer layer of the brain cerebrum is of particular interest due to its key role in cognition, vision and perception. Statistical frameworks on surfaces are, therefore, highly sought for studying various aspects of the brain. Conventional approaches rely on geometrical simplifications, such as spherical inflation and slow mesh deformations [1], often a costly process. For instance, the widely used FreeSurfer [2] takes around 3 hours to parcellate brain surfaces by slowly deforming brain models towards labeled atlases. State-of-the-art learning approaches [3, 4] have the potential to offer a drastic speed advantage over traditional surface-based methods, but operate on image or volumetric spaces. Geometric deep learning [5, 6, 7] recently proposes to use convolutional filters on irregular graphs. The main concern of [8, 9, 10, 11, 12] is their inability to compare surface data across different surface domains. One approach is to map local graph information onto geodesic patches and use conventional spatial convolution via template matching [13, 14, 6]. However, fundamentally, spatial representations of surface data remain defined in Euclidean spaces, for instance, using polar representations of pixels or mesh vertices.

This paper leverages recent advances in spectral graph matching to transfer surface data across aligned spectral domains [15]. This spectral alignment strategy was exploited to learn surface data [16], but was limited to pointwise information, ignoring local patterns within surface neighborhoods. Our novel approach enables a direct learning of surface data across compatible surface bases by exploiting spectral filters over intrinsic representations of surface neighborhoods. We illustrate the learning capabilities of this approach with an application to brain parcellation. The validation over 101 manually labeled brain surfaces [17] shows a significant improvement of spectral graph convolutions over Euclidean approaches, from a Dice score of 51% to 87%. This performance is superior with

[*] `karthik.gopinath.1@etsmtl.net`

Medical Imaging meets NeurIPS (Med-NeurIPS) Workshop, 32nd Conference on Neural Information Processing Systems (NeurIPS 2018), Montréal, Canada.

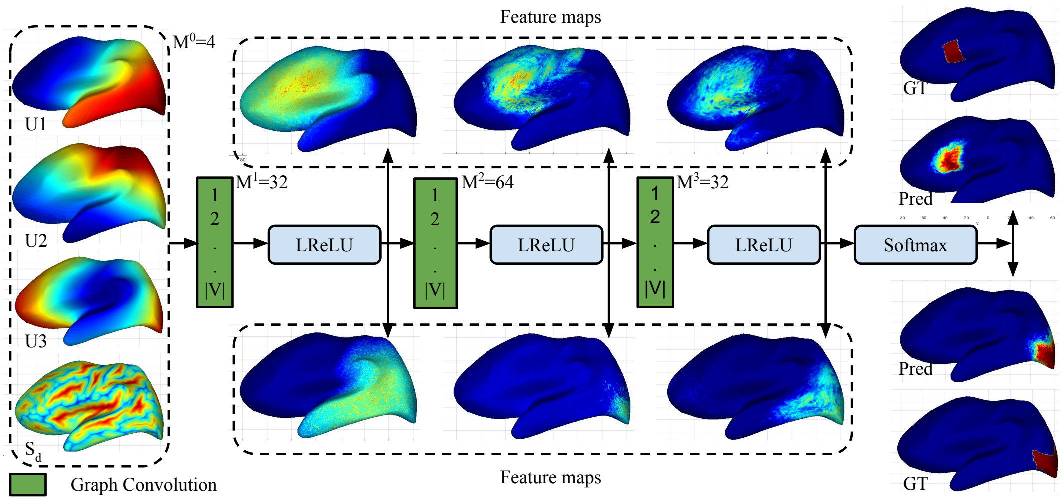

Figure 1: **Overview of the algorithm** – On the left are inputs: Sulcal depth S_d, and corresponding spectral coordinates $(\widehat{\mathbf{u}}_1, \widehat{\mathbf{u}}_2, \widehat{\mathbf{u}}_3)$. In the middle are: Learned convolution layers, M_l, where sample filter responses are shown with coarse to fine geometric features. On the right are: Predicted parcel probabilities (Pred) with ground truth (GT) for two parcels. Brain surfaces are inflated for visualization.

the well established FreeSurfer algorithm [2], which scores 83% [17], while gaining a drastic speed improvement, in the order of seconds. Our contributions are multifold. The transfer of spectral bases across domains enables the design of spectral filters in graph convolutional approaches. Our adaptive spectral filters can consequently learn cortical surface data across multiple geometries, as well as exploit local patterns of data within surface neighborhoods. The next section details the fundamentals of our spectral approach, followed by experiments evaluating the impact of our spectral strategy over standard Euclidean approaches for graph convolutions.

2 Method

An overview of the proposed method is shown in Fig. 1. Firstly, cortical surfaces are modeled as a brain graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, such that $|\mathcal{V}| = N$, and edge set $\mathcal{E}$. Each node i has a feature vector $\mathbf{x}_i \in \mathbb{R}^4$ representing its 3D coordinates and sulcal depth. We map $\mathcal{G}$ to a low-dimension spectral manifold using the normalized graph Laplacian operator $\mathbf{L}$. The eigendecomposition $\mathbf{L}$ is given by $\mathbf{L} = \mathbf{U}\boldsymbol{\Lambda}\mathbf{U}^\top$, with the normalized spectral coordinates of nodes as $\widehat{\mathbf{U}} = \boldsymbol{\Lambda}^{\frac{1}{2}}\mathbf{U}$. The spectral embedding of different brain surfaces are then aligned in the manifold to a reference $\widehat{\mathbf{U}}_{\text{ref}}$ using the Iterative Closest Point (ICP) algorithm. The optimal transformation between matched nodes is then obtained by iterating until convergence. Finally, a geometric convolutional neural network (CNN) is used to map input features, corresponding to the spectral coordinates and sulcal depth of brain graph nodes, to a labeled graph. A generalized convolution operation on a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, with $\mathcal{N}_i = \{j \,|\, (i,j) \in \mathcal{E}\}$, as the neighbors of node $i \in \mathcal{V}$, is defined as:

$$z_{ip}^{(l)} \;=\; \sum_{j \in \mathcal{N}_i} \sum_{q=1}^{M_l} \sum_{k=1}^{K_l} w_{pqk}^{(l)} \cdot y_{jq}^{(l)} \cdot \varphi(\widehat{\mathbf{u}}_i, \widehat{\mathbf{u}}_j;\, \boldsymbol{\Theta}_k^{(l)}) \;+\; b_p^{(l)}, \tag{1}$$

where $\varphi(\widehat{\mathbf{u}}_i, \widehat{\mathbf{u}}_j; \boldsymbol{\Theta}_k)$ is a symmetric kernel in the embedding space with parameter $\boldsymbol{\Theta}_k$. In this work, we follow [6] and use a Gaussian kernel: $\varphi(\widehat{\mathbf{u}}_i, \widehat{\mathbf{u}}_j; \boldsymbol{\mu}_k, \sigma_k) \;=\; \exp\big(-\sigma_k \,\|(\widehat{\mathbf{u}}_j - \widehat{\mathbf{u}}_i) - \boldsymbol{\mu}_k\|^2\big)$.

Using this formulation, we define a fully-convolutional network composed of 3 graph convolution layers with feature map sizes of $M_1 = 32$, $M_2 = 64$ and $M_3 = 32$, each one having $K_l = 4$ Gaussian kernels. The size of the last layer corresponds to the number of parcels to be segmented (32 in our case). Leaky ReLU is applied after each layer to obtain filter responses and a softmax operation after the last graph convolution layer to obtain the parcel probabilities of each node. Finally, cross-entropy is employed as output loss function for back-propagating the error and updating network parameter $\boldsymbol{\Theta} = \{w_{pqk}^{(l)},\, b_p^{(l)},\, \boldsymbol{\Theta}_k^{(l)}\}$ using standard gradient descent optimization.

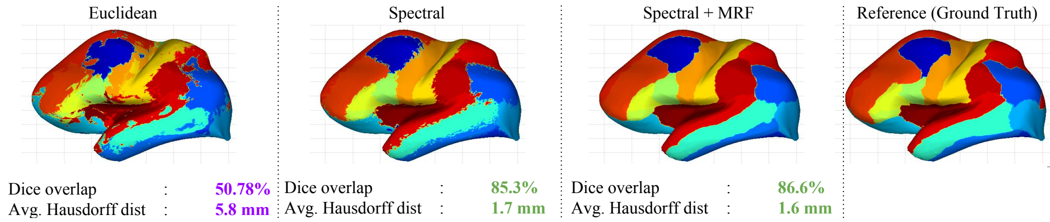

Figure 2: **Cortical Parcellation** – (*First/Left*) Learning with Euclidean coordinates, resulting in low Dice score (50.7%) and inconsistent boundaries (Hausdorff distance of 5.8mm). (*Second*) Learning with Spectral coordinates, improving Dice score (85.3%) and boundary regularity (1.7mm). (*Third*) Spectral domain regularized with MRF, leading to a dice score (86.6%) consistent boundaries (1.6mm) with respect to the reference (*Fourth*). The brain surface is inflated for visualization.

3 Results

We now evaluate the performance of our contributions by highlighting the advantage of moving graph learning frameworks from a conventional Euclidean domain to a Spectral domain. Our validation is performed on Mindboggle [17], the largest publicly available dataset of manually labeled brain MRI. It consists of 101 subjects collected from different sites, with cortical meshes varying from 102K to 185K vertices. Each brain surface contains 32 manually labeled parcels. Results are measured in terms of average Dice overlap and Hausdorff distances [16], and shown in Fig. 2.

We evaluate the improvement of moving the learning operations from the Euclidean domain to a Spectral domain. In our baseline, similarly to the latest approaches of graph convolutions networks [6], we learn from input features in the Euclidean domain. Each cortical point is represented using sulcal depth and its spatial location. This algorithm performs with an average Dice overlap of 50.8% (± 20.5, min/max = 0.0 / 85.5%) across all parcels in our dataset. The Hausdorff distance averaged across all parcels is 5.8mm. Fig. 2 clearly illustrates the current limitation of existing graph convolutions approaches. In a geometry-aware Spectral domain, using the same architecture and data split as before, the average Dice overlap across all parcels improves to 85.3% (± 5.3, min/max = 69.5 / 95.1%). The Hausdorff distance averaged across all parcels is now reduced to 1.7mm. This is a 68% improvement over learning in the conventional Euclidean domain. The qualitative results of Fig. 2 show that our cortical parcellation is almost similar to the manual parcellation. The boundary, however, is irregular and requires further regularization. As an illustration of further refinement, we use Markov random field (MRF) regularization [18] for both Euclidean and Spectral outputs. MRF regularization further improves the overall classification accuracy from 50.8% to 58.1% in the Euclidean domain, and from 85.3% to 86.6% in the Spectral domain. Similar improvement is observed in terms of Hausdorff distance, with a reduction from 5.8mm to 4.9mm in the Euclidean domain, and from 1.7mm to 1.6mm in the Spectral domain.

Results for all parcels are also at par with FreeSurfer's (from 83.2% vs 86.6%), while gaining a significant improvement in computation time, from 3 hours to 18 seconds for processing one subject with our spectral graph convolution network.

4 Conclusion

This paper presented a novel framework for learning surface data via spectral graph convolutions. The algorithm leverages recent advances in spectral matching to enable the comparison of surface data across different surface domains. Our experiments illustrated the benefits of our approach with an application to cortical surface parcellation. This is a particularly challenging problem where current graph convolution approaches remain limited by the inability to compare surface data across brain geometries. This typically results in spatial irregularities of parcel boundaries as illustrated in Fig. 2. By capturing the geometry of the spectral manifold, the proposed method can improve the parcellation accuracy to a Dice score of 85.3%, from 50.7% with graph convolutions in the Euclidean space. The performance of our method 85.3% is better than the state-of-the-art approach FreeSurfer 83.2% for cortical parcellation, also, it reduces the computation time by an order of magnitude (18 seconds vs hours for FreeSurfer). While the potential of our method was demonstrated on cortical parcellation, it can be applied to other analyses of surface data, potentially leading to new families of geometry-based biomarkers for neurological disorders.

Acknowledgments

This work was supported by a Natural Sciences and Engineering Research Council of Canada (NSERC). We gratefully acknowledge the support of NVIDIA Corporation for the donation of the Titan X Pascal GPU used for this research.

References

[1] N. Tustison et al. Large-scale evaluation of ANTs and FreeSurfer cortical thickness measurements. *NeuroImage*, 2014.

[2] B. Fischl et al. Automatically parcellating the human cortex. *C. Cortex*, 2004.

[3] J. Dolz, C. Desrosiers, and I. B. Ayed. 3D fully convolutional networks for subcortical segmentation in MRI: A large-scale study. *NeuroImage*, 2017.

[4] K. Kamnitsas, C. Ledig, V. F. J. Newcombe, J. P. Simpson, A. D. Kane, David K. Menon, D. Rueckert, and B. Glocker. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *Medical image analysis*, 2017.

[5] M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing*, 2017.

[6] F. Monti, D. Boscaini, J. Masci, E. Rodolà, J. Svoboda, and M. M. Bronstein. Geometric deep learning on graphs using mixture model CNNs. In *CVPR*, 2017.

[7] R. Levie, F. Monti, X. Bresson, and M. M. Bronstein. CayleyNets: Graph convolutional neural networks with complex rational spectral filters. *arxiv*, 2017.

[8] M. Bronstein, K. Glashoff, and T. Loring. Making laplacians commute. *CoRR*, 2013.

[9] A. Kovnatsky, M. M. Bronstein, A. M. Bronstein, K. Glashoff, and R. Kimmel. Coupled quasi-harmonic bases. In *Computer Graphics Forum*, 2013.

[10] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: A flexible representation of maps between shapes. In *SIGGRAPH*, 2012.

[11] D. Eynard, A. Kovnatsky, M. M. Bronstein, K. Glashoff, and A. M. Bronstein. Multimodal manifold analysis by simultaneous diagonalization of Laplacians. *IEEE PAMI*, 2015.

[12] S. Parisot, S. I. Ktena, E. Ferrante, M. Lee, R. G. Moreno, B. Glocker, and D. Rueckert. Spectral graph convolutions for population-based disease prediction. In *MICCAI*, 2017.

[13] J. Masci, D. Boscaini, M. M. Bronstein, and P. Vandergheynst. Geodesic convolutional neural networks on Riemannian manifolds. In *ICCV-3dRR*, 2015.

[14] D. Boscaini, J. Masci, E. Rodolà, and M. M. Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *NIPS*, 2016.

[15] H. Lombaert, M. Arcaro, and N. Ayache. Brain transfer: Spectral analysis of cortical surfaces and functional maps. In *IPMI*, 2015.

[16] H. Lombaert, A. Criminisi, and N. Ayache. Spectral forests: Learning of surface data, application to cortical parcellation. In *MICCAI*, 2015.

[17] A. Klein et al. Mindboggling morphometry of human brains. *PLOS Bio*, 2017.

[18] Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max- flow algorithms for energy minimization in vision. *IEEE PAMI*, 2004.