



Published in final edited form as:

Proc IEEE Int Symp Biomed Imaging. 2018 April ; 2018: 527–530. doi:10.1109/ISBI.2018.8363631.

NON-EUCLIDEAN, CONVOLUTIONAL LEARNING ON CORTICAL BRAIN SURFACES

Mahmoud Mostapha^{*}, SunHyung Kim[†], Guorong Wu[‡], Leo Zsembik[†], Stephen Pizer^{*,‡}, and Martin Styner^{*,†}

^{*}Department of Computer Science, University of North Carolina at Chapel Hill, USA

[†]Department of Psychiatry, University of North Carolina at Chapel Hill, USA

[‡]Department of Radiology, University of North Carolina at Chapel Hill, USA

Abstract

In recent years there have been many studies indicating that multiple cortical features, extracted at each surface vertex, are promising in the detection of various neurodevelopmental and neurodegenerative diseases. However, with limited datasets, it is challenging to train stable classifiers with such high-dimensional surface data. This necessitates a feature reduction that is commonly accomplished via regional volumetric morphometry from standard brain atlases. However, current regional summaries are not specific to the given age or pathology that is studied, which runs the risk of losing relevant information that can be critical in the classification process. To solve this issue, this paper proposes a novel data-driven approach by extending convolutional neural networks (CNN) for use on non-Euclidean manifolds such as cortical surfaces. The proposed network learns the most powerful features and brain regions from the extracted large dimensional feature space; thus creating a new feature space in which the dimensionality is reduced and feature distributions are better separated. We demonstrate the usability of the proposed surface-CNN framework in an example study classifying Alzheimers disease patients versus normal controls. The high performance in the cross-validation diagnostic results shows the potential of our proposed prediction system.

Index Terms

MRI; Cortical Surfaces; Alzheimer's Disease; Deep Learning; Convolutional Neural Network

1. INTRODUCTION

The development of sophisticated, noninvasive 3D medical imaging technologies such as magnetic resonance imaging (MRI) provides a means to extract biomarkers to assist in the diagnosis of neurodevelopmental and neurodegenerative diseases. Neuroanatomical abnormalities in the cerebral cortex are widely investigated by examining cortical brain morphometric measures such as cortical thickness and cortical surface area [1]. Such cortical measures are usually extracted from highly-sampled 2D cortical surfaces, forming a high-dimensional feature list (commonly close to 100,000 features or more). Based on the extracted features, machine learning classifiers can be used to predict diagnostic outcomes. Such classifiers undergo a learning process in which they learn the true category labels for

each set of features. However, it is hard to train stable classifiers using such a high-dimensional feature space without over-fitting. One way to solve this is to subdivide the human cortex into a mosaic of anatomically and functionally distinct, spatially contiguous areas using standard parcellation atlases [2]. However, current parcellation atlases are generic and not optimized to a given study. For example, these parcellations do not take into account age or pathology of the studied subjects. Alternatively, dimensionality reduction techniques can be used to solve this problem in a supervised or unsupervised manner [3]. However, unsupervised methods run the risk of losing relevant information while supervised methods tend to be more biased and therefore harder to generalize. Therefore, there is a need to develop data-driven classifiers that can learn directly from such high-dimensional features without the need for a separate feature reduction step.

In recent years, convolutional neural networks (CNNs) have become the predominant machine learning tool in computer vision and speech recognition. The main strength of CNNs is their ability to extract mid-level and high-level abstractions directly from raw data with little need of prior knowledge. Another benefit of CNNs is that they are easier to train and have fewer parameters than fully connected networks with the same number of hidden units. However, extending the use of CNNs in applications where the input data has irregular non-Euclidean structure is still challenging. This is because of the missing notion of a grid on a non-Euclidean surface and the additional need for localized kernel design.

To solve this problem, Masci et al. [4] introduced a generalization of the CNN paradigm to non-Euclidean manifolds based on a local geodesic system of polar coordinates to extract patches, which are then passed through a cascade of filters and linear and nonlinear operators. The coefficients of the filters and linear combination weights are optimization variables that are learned to minimize a task-specific cost function. Although this implementation is promising, it usually fails if the surface mesh is very irregular or if the radius of the geodesic patches is large compared to a curvature radius of the shape. To solve these drawbacks, we would need to generate highly uniform surface representations as well as to decrease the geodesic path radius, which in turn might be a problem in terms of expected computational complexity. In another attempt, Boscaini et al. [5] presented another generalization based on localized frequency analysis (a generalization of the windowed Fourier transform to manifolds) that is used to extract the local behavior of some dense intrinsic descriptor, roughly acting as an analogy to patches in images. The resulting local frequency representations are then passed through a bank of filters whose coefficients are determined by a learning procedure minimizing a task-specific cost. In this implementation, the authors addressed some of the limitations described in the previous paper; however, this work was only designed to capture shape descriptors on the surface and is not suitable to be used with other surface measures.

To overcome the above problems, we propose a novel CNN architecture that utilizes geodesic-based kernels in learning the optimal features and brain regions in a data-driven way (see Fig. 1). The contributions of this paper are three-fold:

- A general definition of kernels on non-Euclidean cortical surfaces using a locally constructed geodesic grid.

- A general framework for learning on cortical surfaces using novel surface convolution, pooling, and resampling layers.
- An accurate classification in a high-dimensional feature space without the need for a separate dimensionality reduction step.

Details of the proposed CNN extension will be provided next.

2. METHOD

As shown in Fig. 1(a), a novel extension to the classical CNN framework is proposed, which includes several layers that are applied subsequently. Details of these proposed layers are given next.

Convolution on Cortical Surfaces

Convolution of data living on cortical surfaces is defined as a correlation with a surface kernel that is used to extract corresponding patches on the manifold. A localized grid is established at each surface point in a way that considers the intrinsic shape of the underlying manifold. To achieve that, we rely on the geodesic (shortest path) distances computed at each surface point to reconstruct a local system of geodesic coordinates. The geodesic distance to a collection of points satisfies a non-linear differential equation. This so-called the Eikonal equation gives the viscosity $\phi(x, y)$ according to

$$\|\nabla \phi\| = F. \quad (1)$$

This ϕ can be interpreted as a weighted distance map from an initial seed, where the weights are given by the function $F(x, y)$ which is a scalar positive function [6]. It follows that, the curve, $C(t)$ giving the level set of distance, defined as points on the front of the function ϕ at time t propagates following the evolution equation

$$\frac{d}{dt}C(t) = \frac{\vec{n}_{xy}}{P(x, y)}, \quad (2)$$

where \vec{n}_{xy} is the exterior unit normal vector to the curve at the point $C(t) = (x, y)$. The function $F(x, y) = 1/P(x, y)$ is the propagation speed of the front, $C(t)$.

To compute geodesic distances on a triangular manifold $M: (V, E, F)$, the fast marching method (FMM) is used to provide an approximation of the true geodesic distance field for each vertex $v \in V$ on the manifold. FMM relies on an upwind finite difference approximation to the gradient and a resulting causality relationship, which results in a Dijkstra-like algorithm but with an update step for triangles instead of edges [7, 8].

The radial coordinates relative to a vertex v on the cortical surface are defined as the level curves of the geodesic distance function $\phi = r_i$, where r_i is the radius of the i^{th} geodesic ring (Fig. 2(a)). A grid is now formed at v (Fig. 2(b) and (c)). Angular partitions are then created

by splitting the curve of the geodesic distance function $\phi = r_j$ into segments of equal length and propagating directions from these points inward along the gradient $\nabla_M \phi$. This geodesic grid at v is formed from partitions or regions P_k with a similar surface area, and the order of these partitions is established using the spherical parameterization of the surface (Fig. 2(b)). Finally, the feature values over the points in each partition P_k are summarized using a specified number of sampled points m_{kj} (Fig. 2(c)).

In the training phase, N geodesic kernels are used to produce the corresponding feature maps f_1, \dots, f_N . A pooling layer is then used to produce a fused feature map f by selecting the maximum feature response for each v on the surface from f_1, \dots, f_N .

Surface Resampling

To reduce the number of training parameters and to learn an abstracted form of the representation, the feature dimensionality is reduced via surface subsampling. This is done by reducing the number of triangles in the input triangle mesh, forming a good approximation to the original geometry. Based on repeated edge collapses, a surface simplification algorithm is employed [9, 10]. Edges are placed in a priority queue based on a quadric error measure Q that is associated with each vertex $v: (p, s)$ of the surface, where p is the geometric position and s a set of attribute scalars. Let $v': (p', s')$ be the projection of v onto the associated affine subspace, and let Q be given by

$$Q(v) = \|p - p'\|^2 + \lambda \|s - s'\|^2, \quad (3)$$

the weighted sum of the geometric distance error $\|p - p'\|^2$ and the attribute deviation error $\|s - s'\|^2$. As edges are deleted, the quadric error measures associated with the two endpoints of the edge are summed and an optimal collapse point is computed. This edge collapse process is repeated until the desired surface resolution level is reached or topological constraints are violated. In this paper, Gaussian and mean curvatures along with the geodesic distances are used as attributes in the quadric error measure to ensure the uniformity and integrity of the resampled surfaces. Fig. 3 shows a resampled surface using the proposed surface simplification layer.

Finally, these newly defined convolution and resampling techniques will enable the development of a novel network architecture that can extend classical CNN to the non-Euclidean domain.

3. EXPERIMENTAL RESULTS

Data

We applied our surface-CNN to a subset of the preprocessed structural MRI data from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database [11], specifically the first scan available for Alzheimer's disease (AD) patients and normal control patients (CN). This resulted in a dataset of 86 subjects of two classes: 39 AD patients (25 Males, 14 Females; 75.1 ± 8.2 years) and 47 CN patients (23 Males, 24 Females; 73.2 ± 5.7 years).

Surface Preprocessing

The FreeSurfer [12] pipeline was used to segment the MRI images and produce registered triangulated surfaces of the cortical white matter and gray matter. Cortical thickness (CT) and cortical surface area (SA) were then measured. For the purpose of comparison, we summarized CT and SA using the Desikan-Killiany parcellation atlas (DK 70 ROIs) [13]. The resampled cortical surfaces consisted of 81,920 high-resolution triangles (40,962 vertices) per hemisphere.

Network Architecture

As shown in Fig. 4, geodesic convolution was applied first on features that live on a high-resolution average surface using 64 geodesic kernels (2 mm step size for each geodesic ring) with trainable weights and non-linearity in the form of the hyperbolic tangent (\tanh) activation function. The newly generated features were then summarized into a single feature using a maximum pooling layer. This was followed by a surface resampling with a factor of 4, where feature values are mapped to the lower resolution surface by averaging the closest 5 points' values. This process was repeated for the new lower resolution surface but with proportionally increased geodesic kernel ring step size. After reaching a reasonable number of features, fully connected layers were applied with rectified linear (relu) neurons (total of 124,994 units), and to overcome the over-fitting problem associated with deep neural nets with many parameters, dropout techniques were used after each hidden layer (dropout rate of 0.2). Max-norm regularization was also used to constraint kernel weights, which has been shown to be useful in deep neural networks.

Performance Evaluation

The proposed techniques were tested by applying them to each feature extracted (CT and SA) to classify AD subjects from CN subjects. For training, both sets of features were standardized using Z-score normalization that rescales the features so that they will have zero mean and unit variance. To handle the classes' imbalance in the training phase, a method called SMOTE [14] was used to over-sample the minority class to match the number of samples in the majority class with 5 nearest neighbors used to construct these synthetic samples. The proposed network was compiled using the Adam optimization algorithm and the binary cross-entropy loss function. The network was trained using the back-propagation algorithm for 500 epochs and a batch size of 1.

The network performance was evaluated using 10-fold cross-validation via the mean area under the ROC curve (AUC), sensitivity (SEN), specificity (SPC), positive predictive value (PPV), and negative predictive value (NPV). Only the non-synthetic minority class samples were used in the testing phase. The performance of the proposed CNN extension was assessed by a comparison against a similar network with only fully connected layers that was trained using features summaries provided by DK ROIs. The classification results for each method are provided in Table 1.

As documented in Table 1, the proposed network performs notably better than typically used regional summaries, according to every performance metric. To show the ability of the proposed CNN extension to provide an insight on which brain regions are contributing more

in separating the two classes, Fig. 5 shows the average features learned after applying the first block of convolution, pooling, and resampling layers. Fig. 6 show those regions showing statistically significant group differences. This clearly demonstrates the potential of the proposed CNN extension in dynamically learning regional ROIs that is specific to the training dataset.

4. CONCLUSIONS

This paper presents a novel data-driven generalization of CNNs on non-Euclidean cortical surfaces. The proposed CNN extension is applicable to various kinds of clinical applications that involve learning from high-dimensional features living on non-Euclidean manifolds.

The high performance of the proposed techniques is confirmed on features living on cortical surfaces. The obtained results are promising in demonstrating the ability of the proposed techniques in transforming high-dimensional features into a compact high-level representation where classes are better separated and important surface regions are emphasized. However, more experiments need to be conducted with different datasets to find the optimal network architecture and parameters (e.g., determining the optimal kernel size and number of partitions remains an open point of research).

Our future work will include extending the proposed network to be able to learn from a combination of extracted features (e.g., CT and SA together). Also, we plan to design a synthetic experiment to investigate and confirm the usefulness, significance, and reliability of the proposed techniques. Finally, we plan to use the proposed techniques to introduce an extension to deep convolutional generative adversarial networks (GANs) that can produce synthetic features on cortical surfaces, which better deal with the imbalance problem that usually exists in medical datasets.

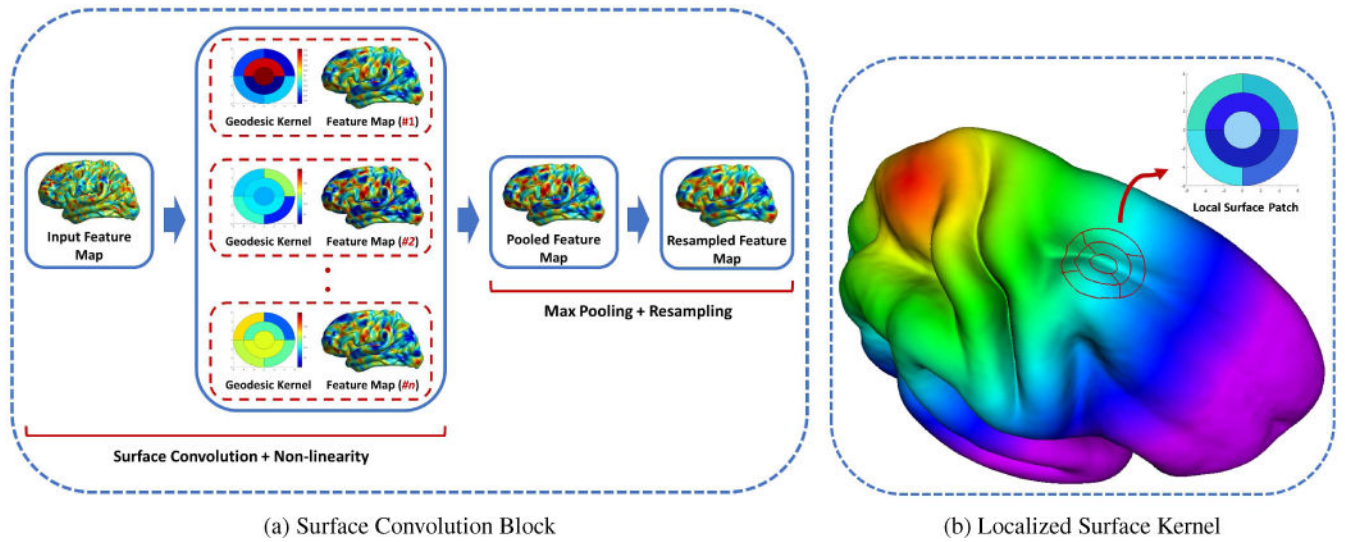
Acknowledgments

Funding was provided by the SALT (Shape Analysis Toolbox for Medical Image Computing Projects) NIH grant (R01EB021391), as well as the following NIH grants: U54HDO79124, R01EB021391, and U01 AG024904.

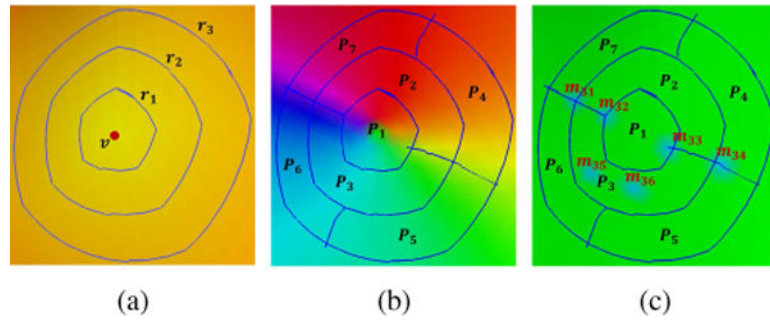
References

1. Hazlett HC, Gu H, Munsell BC, Kim SH, et al. Early brain development in infants at high risk for autism spectrum disorder. *Nature*. 2017; 542(7641):348–351. [PubMed: 28202961]
2. Glasser MF, Coalson TS, Robinson EC, Hacker CD, et al. A multi-modal parcellation of human cerebral cortex. *Nature*. 2016; 536(7615):171–178. [PubMed: 27437579]
3. Van Der Maaten L, Postma E, Van den Herik J. Dimensionality reduction: a comparative. *J Mach Learn Res*. 2009; 10:66–71.
4. Masci J, Boscaini D, Bronstein M, Vandergheynst P. Geodesic convolutional neural networks on riemannian manifolds. *Proceedings of the IEEE international conference on computer vision workshops*. 2015:37–45.
5. Boscaini D, Masci J, Melzi S, Bronstein MM. , et al. *Computer Graphics Forum*. Vol. 34. Wiley Online Library; 2015. Learning class-specific descriptors for deformable shapes using localized spectral convolutional networks; 13–23.
6. Peyre G, Cohen LD. Geodesic remeshing using front propagation. *International Journal of Computer Vision*. 2006:145–156.
7. Sethian JA. Fast marching methods. *SIAM Review*. 1998; 41:199–235.

8. Kimmel R, Sethian J. Computing geodesic paths on manifolds. Proceedings of the National Academy of Sciences. Jul.1998 95
9. Garland M, Heckbert PS. Proceedings of the 24th annual conference on Computer graphics and interactive techniques. ACM Press/Addison-Wesley Publishing Co; 1997. Surface simplification using quadric error metrics; 209–216.
10. Hoppe H. Proceedings of the conference on Visualization' 99: celebrating ten years. IEEE Computer Society Press; 1999. New quadric metric for simplifying meshes with appearance attributes; 59–66.
11. Jack CR, Bernstein MA, Fox NC, Thompson P, et al. The alzheimer's disease neuroimaging initiative (adni): Mri methods. Journal of magnetic resonance imaging. 2008; 27(4):685–691. [PubMed: 18302232]
12. Fischl B. Freesurfer. Neuroimage. 2012; 62(2):774–781. [PubMed: 22248573]
13. Desikan RS, Ségonne F, Fischl B, Quinn BT, et al. An automated labeling system for subdividing the human cerebral cortex on mri scans into gyral based regions of interest. Neuroimage. 2006; 31(3):968–980. [PubMed: 16530430]
14. Chawla NV, Bowyer KW, Hall LO, Kegelmeyer WP. Smote: synthetic minority over-sampling technique. Journal of artificial intelligence research. 2002; 16:321–357.

**Fig. 1.**

The proposed CNN extension on cortical surfaces. (a) A geodesic block consist of surface convolution, pooling, and resampling layers that are applied subsequently. (b) Surface kernel that is reconstructed using a local system of geodesic coordinates.

**Fig. 2.**

Local geodesic kernel partitioned into seven ordered regions. (a) Geodesic distance level curves. (b) Angular partitions created using inward ray shooting and ordered by the surface spherical parameterization. (c) Regional feature summaries are created by averaging six sampled feature values (four measurements at the region corners and two measurements at the straddling of each region). Showing here the measurements locations for P_3 as an example.




Fig. 3.

An example of a resampled surface showing preserved shape and geodesic distances at different resolutions.

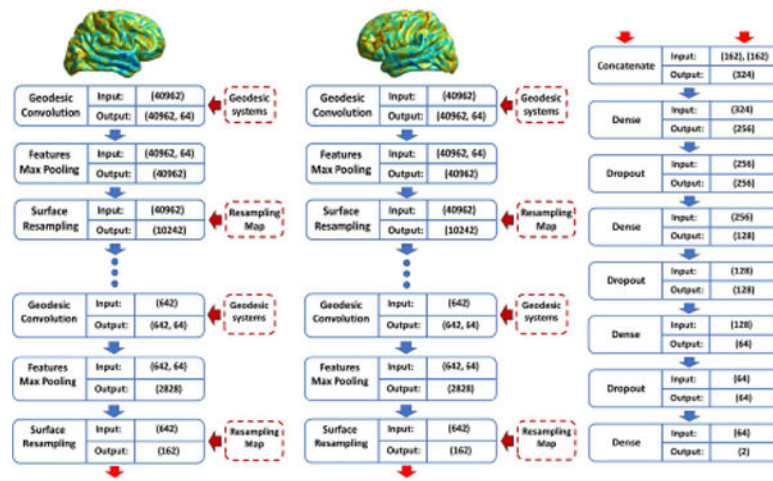


Fig. 4.

The proposed CNN extension network architecture containing four convolutional blocks for each hemisphere followed by fully connected and dropout layers.

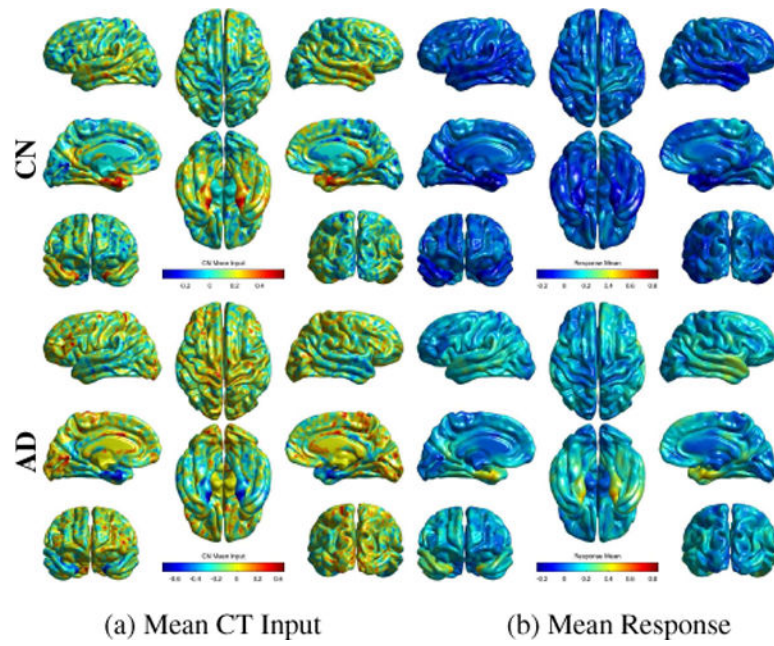


Fig. 5. Visualization of the mean class-specific learned features after applying the first surface convolution block of layers. The proposed network can learn the most powerful features and brain regions that better separate different classes.




Fig. 6.

Uncorrected statistical t -test results on the first geodesic block output indicated that several brain regions (36% of the whole brain showing here in red) showed statistically significant (p -value < 0.05) group differences.

Classification accuracy of the proposed CNN extension compared to a typical regional summary.

Table 1

		AUC	ACC	SEN	SPC	NPV	PPV
CT	Proposed	0.94	91.86	92.31	91.49	93.48	90.00
	Freesurfer ROIs [13]	0.84	81.40	79.49	82.98	82.98	79.49
SA	Proposed	0.87	86.05	84.62	87.23	87.23	84.62
	Freesurfer ROIs [13]	0.67	67.44	58.97	74.47	68.63	65.71