Fast Mesh Data Augmentation via Chebyshev Polynomial of Spectral filtering

Shih-Gu Huang, Moo K. Chung, Anqi Qiu and
Alzheimer's Disease Neuroimaging Initiative[*‡]

October 19, 2021

Abstract

Deep neural networks have recently been recognized as one of the powerful learning techniques in computer vision and medical image analysis. Trained deep neural networks need to be generalizable to new data that was not seen before. In practice, there is often insufficient training data available and augmentation is used to expand the dataset. Even though graph convolutional neural network (graph-CNN) has been widely used in deep learning, there is a lack of augmentation methods to generate data on graphs or surfaces. This study proposes two unbiased augmentation methods, Laplace-Beltrami eigenfunction Data Augmentation (LB-eigDA) and Chebyshev polynomial Data Augmentation (C-pDA), to generate new data on surfaces, whose mean is the same as that of real data. LB-eigDA augments data via the resampling of the LB coefficients. In parallel with LB-eigDA, we introduce a fast augmentation approach, C-pDA, that employs a polynomial approximation of LB spectral filters on surfaces. We design LB spectral bandpass filters by Chebyshev polynomial approximation and resample signals filtered via these filters to generate new data on surfaces. We first validate LB-eigDA and C-pDA via simulated data and demonstrate their use for improving classification accuracy. We then employ the brain images of Alzheimer's Disease Neuroimaging Initiative (ADNI) and extract cortical thickness that is represented on the cortical surface to illustrate the use of the two augmentation

---

[*]S.-G. Huang is with Department of Biomedical Engineering, National University of Singapore, Singapore (e-mail: shihgu@gmail.com).

[†]M.K. Chung is with the Department of Biostatistics and Medical Informatics, University of Wisconsin, Madison, WI 53706, USA (e-mail: mkchung@wisc.edu).

[‡]A. Qiu is with Department of Biomedical Engineering and the N.1 Institute for Health, National University of Singapore, Singapore, 117583 (e-mail: bieqa@nus.edu.sg).

1

methods. We demonstrate that augmented cortical thickness has a similar pattern to real data. Second, we show that C-pDA is much faster than LB-eigDA. Last, we show that C-pDA can improve the AD classification accuracy of graph-CNN.

1   Introduction

Deep neural networks have recently been recognized as one of the powerful learning techniques in computer vision and medical image analysis [1, 2]. Training deep neural networks requires a large dataset so that they are generalizable to data that have never been seen before. This is challenging especially in the field of medical image analysis. Building big medical image datasets is expensive and labor-intensive to collect, and is related to patient privacy, and the requirement of medical experts for labeling. Not having enough data could overfit training data so that network models are not generalized to new data. Moreover, studies on rare diseases or medical screening also face the problem of class imbalance with a skewed ratio of majority to minority samples [3, 4]. These obstacles have led to many studies on image data augmentation (see review in [5]). Data augmentation assumes that additional information can be extracted from an original dataset. It is a very powerful approach for overcoming overfitting in deep learning.

Image augmentation inflates the size of training data via either image transformation or oversampling. New images can be generated by warping existing images via geometric (rotation, flipping) and color transformations [6], random erasing [7], and adversarial training [8, 9] such that their labels are preserved. In contrast, oversampling augmentation creates synthetic data by mixing existing images, auto encoder-decoder [10, 11], and generative adversarial networks (GANs) [12, 13]. Even though GANs are powerful, their computation is more expensive compared to image warping methods.

Among existing image augmentation methods [6, 7, 10, 11, 12, 13], image data are defined on an equi-spaced grid in the Euclidean space. However, medical images in the Euclidean space may not fully characterize the geometry of human organs that encompass their intrinsic and complex anatomy, as well as physiological functions. For example, the cerebral cortex is composed of ridges (gyri) and valleys (sulci). Due to the way gyri and sulci are curved, the cortex is thicker in gyri but thinner in sulci. Hence, it is preferred to represent brain images in a way that the underlying geometrical information is encoded. One can express the cerebral cortex as a surface embedded in the 3D Euclidean space. Existing literature has demonstrated that such representation incorporates useful geometry information of the brain into machine

learning for disease diagnosis [14, 15, 16, 17]. Recently, a number of deep neural networks, such as diffusion-convolutional neural networks (DCNNs) [18], PATCHY-SAN [19, 20], gated graph sequential neural networks [21], DeepWalk [22], and spectral graph convolutional neural networks (graph-CNN) [23, 24, 25, 26, 27, 28, 29] can take data on surfaces for classification. The core challenge for implementing CNN on surfaces lies in defining the convolution on surfaces. These existing neural network approaches focus on how to process vertices whose neighborhood has different sizes and connections for the convolution in the spatial domain. Alternately, the convolution can be defined as a multiplication involving a diagonal matrix in the graph Fourier transform derived from a normalized graph Laplacian in the spectral domain. Hence, existing image warping augmentations on equi-spaced grids (e.g., flipping, rotation, shifting) may not directly apply to data on surfaces since the points on surfaces are not on the equi-spaced grid of the Euclidean space. Nevertheless, there is a lack of augmentation approaches to generate data on surfaces.

This study proposes two unbiased augmentation methods, Laplace-Beltrami eigenfunction Data Augmentation (LB-eigDA) and Chebyshev polynomial Data Augmentation (C-pDA), to generate new data on surfaces. These two approaches preserve the mean of real data in each class, which is crucial for classification problems. These two approaches are motivated by the Fourier representation of signals in equi-spaced Euclidean grids. A signal in equi-spaced Euclidean grids can be created as a linear combination of Fourier bases, where the corresponding Fourier coefficients can be generated via the resampling of the Fourier coefficients of existing signals [30, 31, 32]. We adopt this idea and compute the eigenfunctions of the Laplace-Beltrami (LB) operator on a surface. New data on the surface can be constructed via the resampling of the LB coefficients among real data on the surface.

In parallel with LB-eigDA, we introduce a fast augmentation approach, C-pDA, that employs a polynomial approximation of LB spectral filters on surfaces. C-pDA is designed to be in line with graph-CNN [24, 29], where spectral filters are implemented via Chebychev polynomial approximation such that the resulting convolution can be written as a polynomial of the adjacency matrix of a graph. This avoids the cost of calculating the eigenfunctions of a large-scale graph Laplacian. In [24, 29], it is shown that the k-th order Chebyshev polynomial formation of the graph Laplacian is equivalent to k-ring filtering. In C-pDA, we design LB spectral bandpass filters by Chebyshev polynomial approximation and resample filtered real data to generate new data. Due to the recurrence relation of Chebyshev polynomials, the computation of the C-pDA method can be efficient. We

validate LB-eigDA and C-pDA using simulated data with the ground truth of class labels. We further employ the methods to the cortical surface data in Alzheimer's Disease Neuroimaging Initiative (ADNI). We first demonstrate that augmented cortical thickness data have a similar pattern to real data. Second, we show that C-pDA is much faster than LB-eigDA . Last, we illustrate the use of C-pDA to improve the AD classification of the graph-CNN [24].

The main contributions of this study are as follows.

- We introduce two augmentation methods to generate new data on surfaces using the LB eigenfunctions and LB spectral filters.

- We show that C-pDA is computationally more efficient than LB-eigDA.

- We demonstrate that C-pDA improves the graph-CNN performance on the classification of AD patients.

2 Methods

2.1 Augmentation based on the Laplace-Beltrami representation of signals on a surface mesh

We introduce a data augmentation method based on the Laplace-Beltrami representation of signals on a surface mesh. We denote the surface as $\mathcal{M}$ with the Laplace-Beltrami (LB) operator Δ on $\mathcal{M}$. Let ψ_j be the j^{th} eigenfunction of the LB-operator with eigenvalue λ_j

$$\Delta\psi_j = \lambda_j\psi_j \; , \tag{1}$$

where $0 = \lambda_0 \leq \lambda_1 \leq \lambda_2 \leq \cdots$. A signal $f(x)$ on the surface $\mathcal{M}$ can be represented as a linear combination of the LB eigenfunctions

$$f(x) = \sum_{j=0}^{\infty} c_j\psi_j(x) \; , \tag{2}$$

where c_j is the j^{th} coefficient associated with the eigenfunction $\psi_j(x)$. For n observations, $f_1(x), \cdots, f_n(x)$, $f_i(x)$ can be represented as

$$f_i(x) = \sum_{j=0}^{\infty} c_j^{(i)}\psi_j(x) \; ,$$

4

where $c_j^{(i)}$ is the j^{th} coefficient associated with the j^{th} LB eigenfunction for the i^{th} observation. We like to generate new data based on the frequency resampling of these n observations. This is similar to creating new samples via permuting Fourier coefficients [32]. Let S_n be the permutation group of order n [33] and $\tau \in S_n$ be an element of permutation given by

$$\tau = \left(\begin{array}{cccc} 1 & 2 & \cdots & n \\ \tau(1) & \tau(2) & \cdots & \tau(n) \end{array} \right). \tag{3}$$

$\tau(i)$ indicates element i is permuted to $\tau(i)$. We resample the LB coefficients to obtain new data representation $f_{i'}(x)$:

$$f_{i'}(x) = \sum_{j=0}^{\infty} c_j^{\tau_j(i)} \psi_j(x) , \tag{4}$$

where $\tau_j(\cdot)$ is the permutation on the j^{th} LB coefficients among the n observations. We will refer this approach as *LB eigenfunction Data Augmentation (LB-eigDA)*.

Based on Eq. (4), one can show that the mean of $f_{i'}(x)$ over every possible permutation is the same as that of $f_i(x)$ since the permutation function $\tau(\cdot)$ does not change the mean of the LB coefficients.

2.2 Augmentation via Chebyshev polynomials

Previous research suggests that the augmentation strategy of Gaussian filters leads to the best validation accuracy in medical imaging classification tasks [34]. We now introduce the second data augmentation approach, *Chebyshev polynomial Data Augmentation (C-pDA)*. The idea of C-pDA is similar to the augmentation strategy of Gaussian filters in equi-spaced grids of the Euclidean space by designing LB spectral filters on surfaces. We design LB spectral filters that are similar to spectral filter banks [35]. We can then approximate real data on surfaces using these LB spectral filters and resample the LB spectral filtered signals of real data in order to generate new data on surfaces. To avoid the direct computation of the LB eigenfunctions, we will employ the Chebyshev polynomial approximation of LB spectral filters, which is computationally efficient. In the following, we first describe the Chebyshev polynomial approximation of an LB spectral filter and then design LB spectral bandpass filters for the C-pDA approach.

2.2.1 Chebychev polynomial approximation of LB spectral filters

Consider an LB spectral filter g on the surface $\mathcal{M}$ with spectrum $g(\lambda)$ as

$$g(x, y) = \sum_{j=0}^{\infty} g(\lambda_j)\psi_j(x)\psi_j(y). \tag{5}$$

Based on Eq. (2), the convolution of a signal f with the filter g can be written as

$$h(x) = g * f(x) = \sum_{j=0}^{\infty} g(\lambda_j)c_j\psi_j(x). \tag{6}$$

As suggested in [24, 36, 37, 38, 39, 35], the filter spectrum $g(\lambda)$ in Eq. (6) can be represented as the expansion of Chebyshev polynomials, $T_k, k = 0, 1, 2, \ldots, \infty$, such that

$$g(\lambda) = \sum_{k=0}^{\infty} \theta_k T_k(\lambda) \ . \tag{7}$$

θ_k is the k^{th} expansion coefficient associated with the k^{th} Chebyshev polynomial. T_k is the Chebyshev polynomial of the form $T_k(\lambda) = \cos(k\cos^{-1}\lambda)$ with recurrence

$$T_{k+1}(\lambda) = (2 - \delta_{k0})\lambda \ T_k(\lambda) - T_{k-1}(\lambda),$$

where δ_{k0} is Kronecker delta. The convolution in Eq. (6) can be rewritten as

$$h(x) = g * f(x) = \sum_{k=0}^{\infty} \theta_k T_k(\Delta)f(x). \tag{8}$$

This Chebyshev polynomial approximation of the spectral filter has previously used in diffusion wavelet transform [38, 37, 39, 40], graph convolutional neural network [24, 36], spectral wavelet transform [35], and heat diffusion [41] on graphs. The polynomial method avoids the direct computation of the LB eigenfunctions through the recursive computation of $T_k(\Delta)f(x)$ and preserves local geometric structure of the surface [24, 41].

2.2.2 C-pDA

We design a series of LB spectral bandpass filters, g_l, $l = 1, 2, \ldots, L$, based on Eq. (7) such that

$$g_l(\lambda) = \sum_{k=0}^{\infty} \theta_{lk} T_k(\lambda)$$

where θ_{lk} is the k^{th} Chebyshev expansion coefficient of the l^{th} bandpass filter. The frequency band of the l^{th} bandpass filter is $\lambda \in [\epsilon_l \ \epsilon_{l+1}]$. Now, a signal $f(x)$ on surface $\mathcal{M}$ can be approximated using these filters such that

$$f(x) \approx h_0 + \sum_{l=1}^{L} g_l(\Delta) f(x), \tag{9}$$

where h_0 is the mean of $f(x)$ over the surface. If g_l, $l = 1, 2, \ldots, L$, together span the entire spectrum of $f(x)$, then the spectral information of $f(x)$ is retained.

We develop the C-pDA approach in a way similar to the LB-eigDA approach in Eq. (4) such that

$$f_{i'}(x) = h_0^{\tau_0(i)} + \sum_{l=1}^{L} \Big(g_l(\Delta) f_i(x) \Big)^{\tau_l(i)}, \tag{10}$$

where $\tau_l(\cdot)$ is the permutation on the l^{th} filtered signal among the n observations f_1, f_2,..., f_n such that the i^{th} observation is permuted to the $\tau_l(i)^{th}$ observation. Hence, C-pDA generates new data via resampling the l^{th} filtered outputs among the n observations and summing the resampled signals across L filters. Again, we can show that the mean of $f_{i'}(x)$ over every possible permutation is the same as that of $f_i(x)$ since the permutation function $\tau(\cdot)$ does not change the mean of the filtered signals.

With the Chebyshev polynomial approximation, we can rewrite Eq. (10) as

$$f_{i'}(x) = h_0^{\tau_0(i)} + \sum_{l=1}^{L} \Big(\sum_{k=0}^{\infty} \theta_{lk} T_k(\Delta) f_i(x) \Big)^{\tau_l(i)}. \tag{11}$$

2.3   LB-eigDA and C-pDA numerical implementation

For the implementation of the LB-eigDA in Eq. (4), we adopt the discretization scheme of the LB operator in [35], where surface $\mathcal{M}$ is represented by a triangulated mesh with a set of triangles and vertices v_i. The ij^{th} element of the LB-operator on $\mathcal{M}$ can be computed as

$$\Delta_{ij} = C_{ij}/A_i, \tag{12}$$

where A_i is the Voronoi area of vertex v_i if the triangles containing v_i are nonobtuse [42] and Heron's area if the triangles containing v_i are obtuse [35, 42]. The off-diagonal entries are defined as $C_{ij} = -(\cot \theta_{ij} + \cot \phi_{ij})/2$
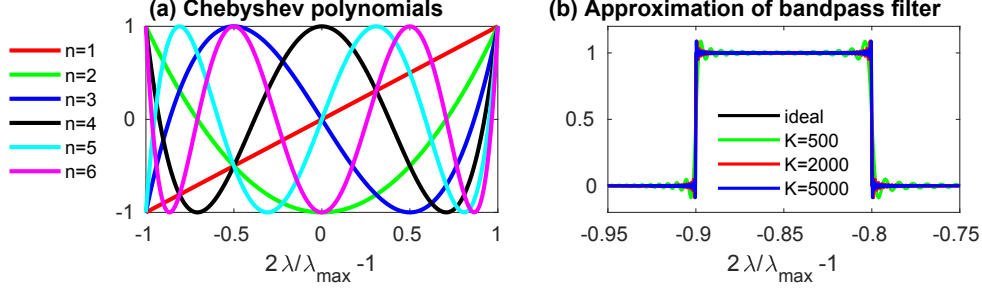
7

Figure 1: (a) Chebyshev polynomials of order 1 to 6. (b) An ideal rectangular bandpass filter with range $\lambda \in [0.05\lambda_{\max}, 0.1\lambda_{\max}]$ and its approximation of Chebyshev polynomials of order up to $K = 500, 2000$, and 5000.

if v_i and v_j form an edge, otherwise $C_{ij} = 0$. The diagonal entries C_{ii} are computed as $C_{ii} = -\sum_j C_{ij}$. Other cotan discretizations of the LB operator are discussed in [43, 44, 45]. When the number of vertices on $\mathcal{M}$ is large, the computation of the LB eigenfunctions can be costly [46].

For the numerical implementation of the C-pDA method in Eq. (11), we need to first determine the order of Chebyshev polynomials while $g_l(\lambda)$ have less overlap for C-pDA. One can quantify the overlap among the filters g_l via training the spectral band between the passband and stopband [47]. A higher-order filter has a narrower transition band than a lower-order filter. Fig. 1 shows the transition bandwidth over order K for Chebyshev polynomials when the filter band is $\lambda \in [0.05\lambda_{\max}, 0.1\lambda_{\max}]$, where $\lambda_{\max}$ is the maximum eigenvalue of the LB operator. In this study, we empirically determined the order of Chebyshev polynomials as $K = 5000$ for C-pDA, which achieves the transition bandwidth as small as 3.5×10^{-4} as illustrated in Fig. 1. L depends on the spectral distribution of the observations and thus is application specific. This study empirically determines L in the below applications.

We take the advantage of the recurrence relation of the Chebyshev polynomials and compute C-pDA recursively. We now describe steps for the numerical implementation of Eq. (11).

1. discretize the surface $\mathcal{M}$ using a triangulated mesh;

2. compute Δ based on Eq. (12) for the surface mesh $\mathcal{M}$;

3. compute the maximum eigenvalue $\lambda_{\max}$ of Δ. For the standardization across surface meshes, we normalize Δ as $\widetilde{\Delta} = \frac{2\Delta}{\lambda_{\max}} - I$, where I is an

8

identity matrix;

4. for the signal f_i of the i^{th} subject, compute $T_k(\widetilde{\Delta})f_i(x)$ recursively by

$$T_{k+1}(\widetilde{\Delta})f_i(x) = (2 - \delta_{k0})\widetilde{\Delta}\ T_k(\widetilde{\Delta})f_i(x) - T_{k-1}(\widetilde{\Delta})f_i(x)$$

with initial conditions

$$T_{-1}(\widetilde{\Delta})f_i(x) = 0$$

and

$$T_0(\widetilde{\Delta})f_i(x) = f_i(x).$$

5. compute each augmented signal f_i' recursively as

$$f_{i'}^k(x) = f_{i'}^{k-1}(x) + \sum_{l=0}^{L}\left(\theta_{lk}T_k(\widetilde{\Delta})f_i(x)\right)^{\tau_l(i)},$$

where

$$\theta_{lk} = \frac{2 - \delta_{k0}}{\pi}\int_{\epsilon_k}^{\epsilon_{k+1}} T_k(\lambda)\frac{d\lambda}{\sqrt{1 - \lambda^2}},$$

where $[\epsilon_k\ \epsilon_{k+1}]$ is the frequency band of g_l. Steps 4 and 5 are repeated from $k = 0$ till $k = K - 1$. In step 5, there is no need to explicitly compute each filtered signal, which saves computational time and memory, especially when a large number of filters are used.

3 Simulation Experiments

A majority of medical applications often face two challenges, limited sample sizes and potential uncertainty of diagnosis [48, 49]. We designed simulation experiments with the ground truth of group labels to illustrate the use of LB-eigDA and C-pDA in the sample size estimation and diagnosis classification.

We performed simulation experiments using a hippocampus surface mesh with 1184 vertices and 2364 triangles. We generated two groups of simulated data on this surface mesh: n samples in Group 0 and m samples in Group 1. We first generated $n + m$ measurements by a normal distribution with mean 0 and variance σ^2, i.e., $\mathcal{N}(0, \sigma^2)$, at each vertex of the hippocampus surface. The first n measurements were considered as samples in Group 0, while the rest of m measurements were added signal 1 in a small patch on the hippocampus (see the red region in Fig. 2 (a)) and were considered as Group
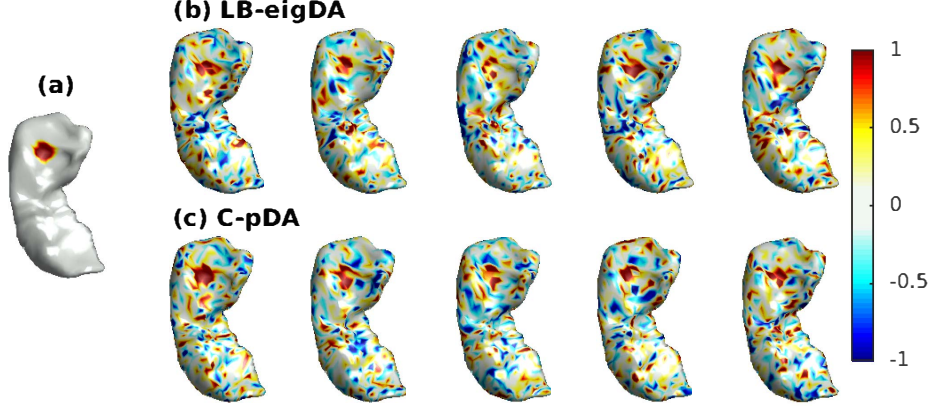
9

Figure 2: Simulated and augmented data in Group 1. (a) Averaged signal over 500 data that were simulated via the distribution $\mathcal{N}(1, \sigma^2)$ in the small patch (red region) of the hippocampus and the distribution of $\mathcal{N}(0, \sigma^2)$ at each vertex on the rest of the hippocampus. (b) five augmented data for Group 1 via the LB-eigDA method ; and (c) five augmented data for Group 1 via the C-pDA method.

1. Thus, Group 0 had the distribution $\mathcal{N}(0, \sigma^2)$ at each vertex, while Group 1 had the distribution $\mathcal{N}(1, \sigma^2)$ in the small patch of the hippocampus and the distribution of $\mathcal{N}(0, \sigma^2)$ at each vertex on the rest of the hippocampus. Fig. 2 (a) shows the signal averaged over 500 samples in Group 1.

To generate augmented data, we computed all the 1184 eigenfunctions for LB-eigDA. The hippocampal surface mesh had the spectrum over [0, 10.9]. For C-pDA, we used 109 bandpass filters whose bandwidth was 0.1 and a mean filter that computed the average value of a signal over the hippocampal surface. Each filter was approximated by Chebyshev polynomials of order 5000. Fig. 2(b) and (c) show 5 augmented data generated by LB-eigDA and C-pDA for Group 1, respectively.

We employed a convolutional neural network (CNN) that was a modified version of the graph-CNN in [24, 36]. We employed the LB operator instead of the graph Laplacian in the CNN in this study. We called it as an LB-based spectral CNN. Fig. 3) shows the LB-based spectral CNN architecture with two convolutional layers due to the relatively small surface mesh of the hippocampus and one fully connected layer. The two convolutional layers had 8 and 16 filters, respectively. Each filter was characterized by the Chebyshev polynomials of order 7. Moreover, each layer also included

10

hippocampal surface

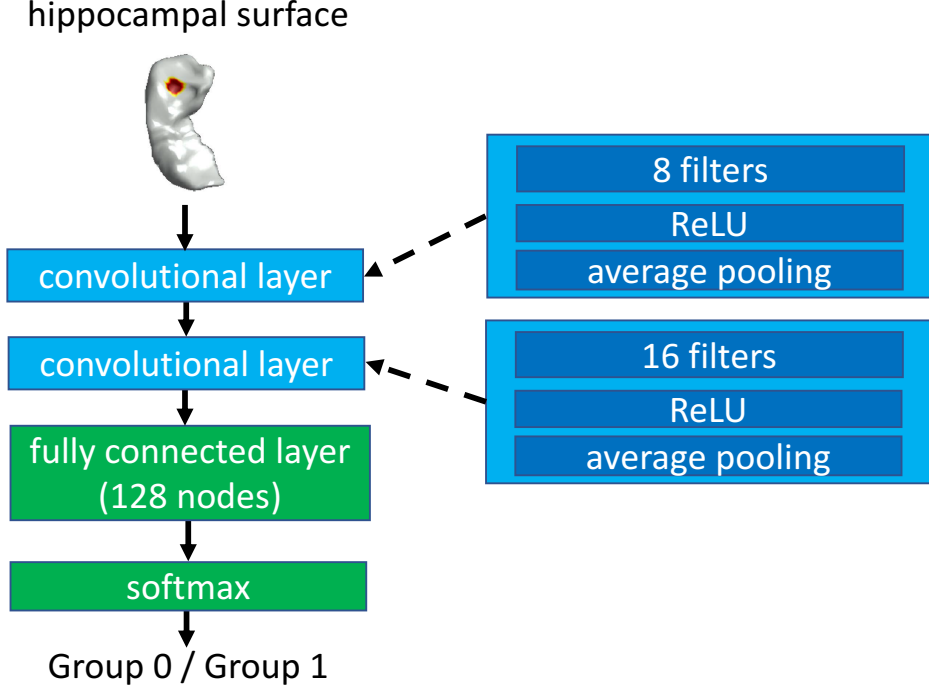|	
convolutional layer	8 filters
	ReLU
	average pooling
convolutional layer	16 filters
	ReLU
	average pooling
fully connected layer (128 nodes)	
softmax	

Group 0 / Group 1

Figure 3: The LB-based spectral CNN with 2 convolutional layers and one fully connected layer. Each convolutional layer is comprised of filters approximated by the Chebyshev polynomials of order 7, a rectified linear unit (ReLU), and average pooling.

a rectified linear unit (ReLU) and average pooling. We trained the network with an initial learning rate of 10^{-3}, and a learning rate decay of 0.05 for every 20 epochs. We applied the ten-fold cross-validation, where one fold was used for testing and the other 9 folds were for training (75%) and validation (25%). Fig. 5 (a) shows the classification accuracy versus total sample size $n + m$ with ratio $n/m = 2$, which was similar to real ADNI data used below in this study. $\sigma = 0.6$ was used. A higher value of σ resulted in a similar curve except that more samples were required to reach the same classification accuracy. The accuracy reached 98.1% when the total sample size was 3000 and then increased slowly as the sample size increased.

To demonstrate the use of the two augmented data in classification, we fixed the total sample size as 3000 ($n = 2000$, $m = 1000$). Among the 3000 samples, 2025, 675, and 300 samples were respectively used as the training,
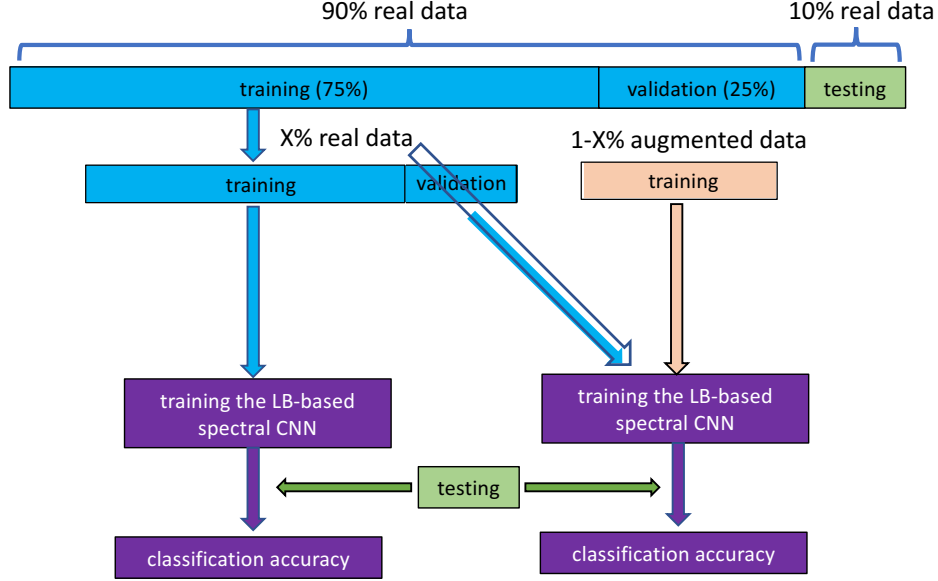
11

Figure 4: Real and augmented data used in the LB-based spectral CNN. $X\%$ indicates that the percentage of the training set is real data and $1 - X\%$ are augmented data.

validation, and testing samples. As illustrated in Fig. 4, when only a smaller fraction of the simulated data in the training, denoted as $X\%$, was available, we applied the augmentation methods to add $1 - X\%$ augmented data to the training set. For instance, if $X = 10$, we only used 203 of the training samples and employed LB-eigDA or C-pDA to generate 1822 augmented data as additional training samples. The augmentation was employed separately for the two groups. The validation (675 samples) and testing (300 samples) sets remained the same. The classification accuracy was respectively 95.5% for LB-eigDA and 92.5% for C-pDA. Without the augmented data, the classification accuracy was 80.3%, more than 10% lower than that obtained using the data augmented by LB-eigDA and C-pDA. Fig. 5 (b) shows that LB-eigDA and C-pDA improved the classification accuracy when compared to that without augmented data. Moreover, the LB-eigDA method performed in general better than the C-pDA method. This is mainly because the C-pDA method employs the polynomial approximation of the LB spectral filters.
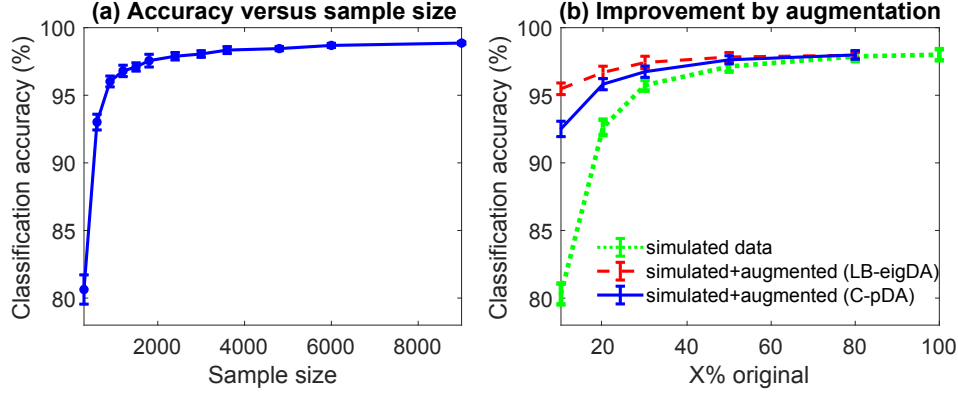
Figure 5: Classification accuracy on simulated and augmented data. (a) The classifications accuracy using simulated data when the sample size increase from 500 to 9000. (b) The green dotted line shows the classification accuracy when only $X\%$ of simulated data was used as the training set. The red dashed and blue solid lines show the classification accuracy when only $X\%$ of the training set were simulated data and $1 - X\%$ of the training set were augmented data by the LB-eigDA and C-pDA methods, respectively.

4 Results

We used MRI data from ADNI. We first illustrate the similarity of augmented data by LB-eigDA and C-pDA to real MRI data. We then compare the computational cost of the LB-eigDA and C-pDA approaches. Finally, we show the use of C-pDA in the LB-based spectral CNN to improve the classification accuracy of Alzheimer's patients.

4.1 MRI data acquisition and preprocessing

We used ADNI-2 cohort (`adni.loni.ucla.edu`) acquired from participants aging from 55 to 90 using either 1.5 or 3T scanners. For the typical 1.5T acquisition, repetition time (TR) = 2400 ms, minimum full echo time (TE) and inversion time (TI)= 1000 ms, flip angle= 8°, field-of-view (FOV)= 240×240 mm^2, acquisition matrix= $256 \times 256 \times 170$ in the x-, y-, and z-dimensions, yielding a voxel size of $1.25 \times 1.25 \times 1.2$ mm^3. For the 3T scans, TR= 2300 ms, minimum full TE and TI = 900 ms, flip angle= 8°, FOV= 260×260 mm^2, acquisition matrix = $256 \times 256 \times 170$, yielding a voxel size of $1.0 \times 1.0 \times 1.2$ mm^3.

We utilized the structural T1-weighted MRI from the ADNI-2 dataset.

The number of visits of each subject varied from 1 to 7 (i.e., baseline, 3-, 6-, 12-, 24-, 36-, and 48-month), and at each visit, the subjects were diagnosed with one of the four clinical statuses based on the criteria in the ADNI protocol (`adni.loni.ucla.edu`): healthy control (HC), early mild cognitive impairment (MCI), late MCI, and Alzheimer's disease (AD). In this study, we illustrated the use of the augmentation methods via the HC/AD classification since it has been well studied using T1-weighted image data (e.g., [50, 51, 52, 53, 54, 55, 56, 36]). Hence, this study involved 643 subjects with HC or AD scans (392 subjects had HC scans; 253 subjects had AD scans). There were 8 subjects who fell into both groups due to the conversion from HC to AD. Tables 1 lists the demographic information of the ADNI-2 cohort.

The T1-weighted images were segmented using FreeSurfer (version 5.3.0) [57]. The white and pial cortical surfaces were generated at the boundary between white and gray matter and the boundary of gray matter and CSF, respectively. Cortical thickness was computed as the distance between the white and pial cortical surfaces. It represents the depth of the cortical ribbon. We represented cortical thickness on the mean surface, the average between the white and pial cortical surfaces. We employed large deformation diffeomorphic metric mapping (LDDMM) [58, 59] to align individual cortical surfaces to the atlas and transferred the cortical thickness of each subject to the atlas. The cortical atlas surface was represented as a triangulated mesh with 655,360 triangles and 327,684 vertices. At each surface vertex, a spline regression implemented by piecewise step functions [60] was performed to regress out the effects of age and gender. The residuals from the regression were used in the below LB-based spectral CNN.

Table 1: Demographic information of the ADNI-2 cohort with MRI scans.

	HC	AD
the number of subjects[†]	400	261
the number of scans	1122	587
gender (female/male)	607/515	254/333
age (years; mean±SD)	75.3±6.8	75.3±7.7

[†] There are 8 subjects who fall into both the HC and AD groups due to the conversion from HC to AD. Abbreviations: HC, healthy controls; AD: Alzheimer's disease; SD, standard deviation.

4.2 LB-eigDA and C-pDA augmentation

We extracted cortical thickness data from 500 ADNI brain MRI scans and then used them to generate augmented cortical thickness via LB-eigDA and C-pDA.

C-pDA requires determining the number of filters and the bandwidth of each filter. These parameters are dependent on the spectrum of real data and application specific. First, we analyzed the spectrum of cortical thickness data, which was predominantly in the low-frequency band. More filters with narrow bandwidth were needed in the low frequency, while fewer filters with wide bandwidth were needed in the high frequency. Second, the discrimination of cortical thickness between controls and AD patients lies in the low-frequency band. Hence, we empirically designed more filters in the low-frequency band based on the following procedure.

Let $\lambda_{\max}$ be the maximum eigenvalue of the LB-operator of the cortical surface mesh. We divided the spectral range of $[0, \frac{\lambda_{\max}}{4^{m-1}}]$ into 2^{m+1} equal-width frequency bands, where m is an integer between 1 and 5, and assigned a bandpass filter to each frequency band. This procedure resulted in a total of 109 filters. Fig. 6 illustrates the filters used in this study. Moreover, the order of the Chebyshev polynomials needs to be determined so that the transition of the filters is sharp. As illustrated in Fig. 1, when $K = 5000$, the approximation of the Chebyshev polynomials converges fast and has a small transition bandwidth. For the rest of this study, we employed $K = 5000$ for C-pDA.

On the other hand, only one parameter, the number of LB eigenfunctions, is needed for LB-eigDA. This study used 5000 eigenfunctions for LB-eigDA, which covered the spectral range critical to the discrimination of controls and AD patients.

We employed LB-eigDA and C-pDA and generated 500 augmented cortical thickness data based on 500 randomly selected data from ADNI. Fig. 7(a) illustrates cortical thickness averaged over the 500 real data. Fig. 7(b) and (c) show 5 augmented thickness data that were respectively generated by LB-eigDA and C-pDA. This figure suggests that the pattern of the augmented data from the two methods is similar to the averaged pattern observed in real data.

Moreover, Fig. 8 shows the thickness averaged over the 500 real data (green solid line), the 500 LB-eigDA augmented data (blue dashed line), and the 500 C-pDA data (red dotted line), respectively. Both LB-eigDA and C-pDA preserved the mean of the real thickness data at each vertex of the cortical surface mesh. Empirically, the largest difference between the real
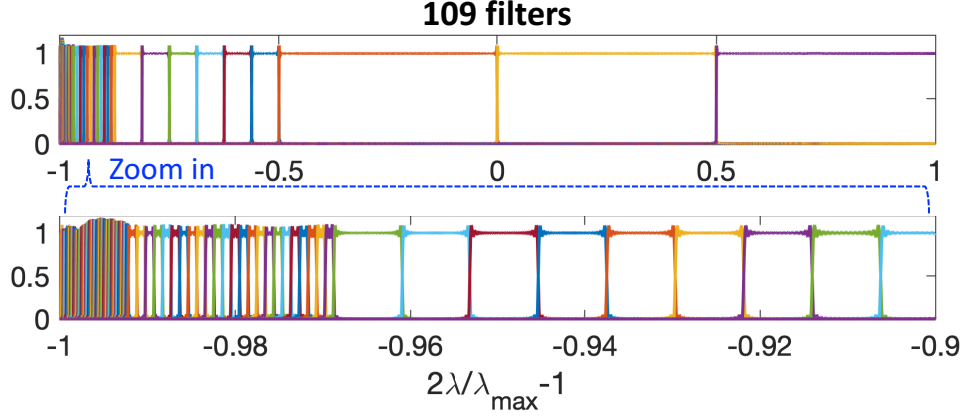
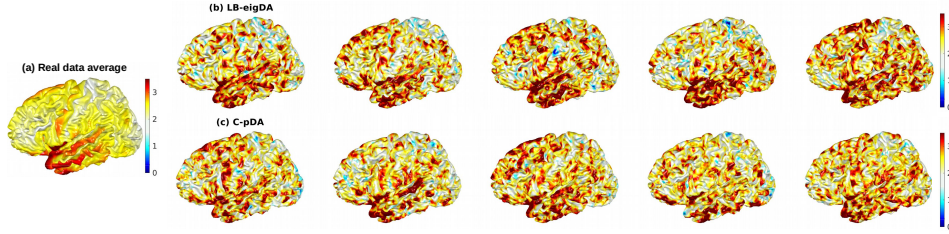Figure 6: A filter bank with 109 bandpass filters used in the C-pDA method.

Figure 7: Augmented cortical thickness. (a) Cortical thickness averaged over 500 real datasets; (b) five augmented thickness data via the LB-eigDA method; and (c) five augmented thickness data via the C-pDA method.

and augmented data was smaller than 10^{-8} mm. Moreover, we computed Pearson's correlation of the averaged real data with the 500 augmented data. Fig. 9 shows the distribution of these correlation values for the LB-eigDA and C-pDA methods. The correlation value of the LB-eigDA augmented thickness was in the range of [0.58, 0.68] with mean and standard deviation of 0.64 ± 0.02, while the C-pDA augmented data showed the correlation in the range of [0.53, 0.72] with mean and standard deviation of 0.65 ± 0.03. Overall, both the LB-eigDA and C-pDA methods can generate new data whose pattern is similar to that of real data.

The LB-eigDA computational time was dependent on the number of the LB eigenfunctions, while the C-pDA computational time was related to the order of Chebyshev polynomials. Fig. 10 shows the LB-eigDA computational time as a function of the number of the LB eigenfunctions and
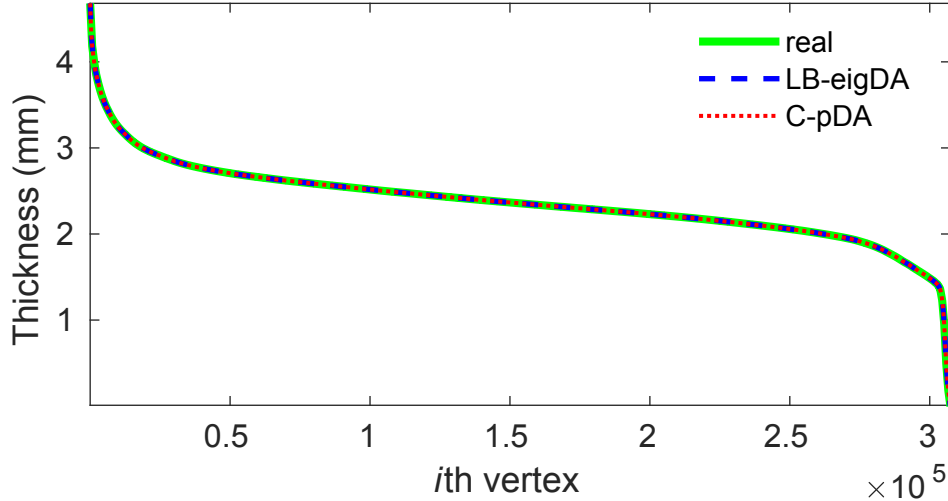
Figure 8: Sorted thickness values at each vertex on the cortical surface mesh. The green solid, blue dashed, and red dotted lines represent the thickness value at a particular vertex averaged over the 500 real data, 500 augmented data via LB-eigDA, and 500 augmented data via C-pDA, respectively. For the purpose of visualization, the thickness averaged over 500 original data is sorted in a descend manner across all the vertices on the cortical surface mesh. The augmented data follow the sorted vertex index.

the C-pDA computational time as the order of Chebyshev polynomials, K. This figure suggests that more LB eigenfunctions used in LB-eigDA allow the augmentation over a wider spectrum but require a high computational cost when the cortical surface mesh is large (the cortical surface mesh with 327,684 vertices). The LB-eigDA computational cost was exponentially increased as the number of the LB eigenfunctions increased. In contrast, the C-pDA computational time was approximately a linear function of the order of Chebyshev polynomials. Compared to C-pDA , LB-eigDA was 70 times slower when $K = 5000$.

4.3 Does classification improve by data augmentation?

We illustrate the use of the C-pDA method to classify healthy controls (HC) and AD patients based on the cortical thickness of the ADNI dataset. Again, we employed the LB-based spectral CNN with the architecture similar to that in Fig. 3, but used five convolutional layers. Each layer involved 8, 16,
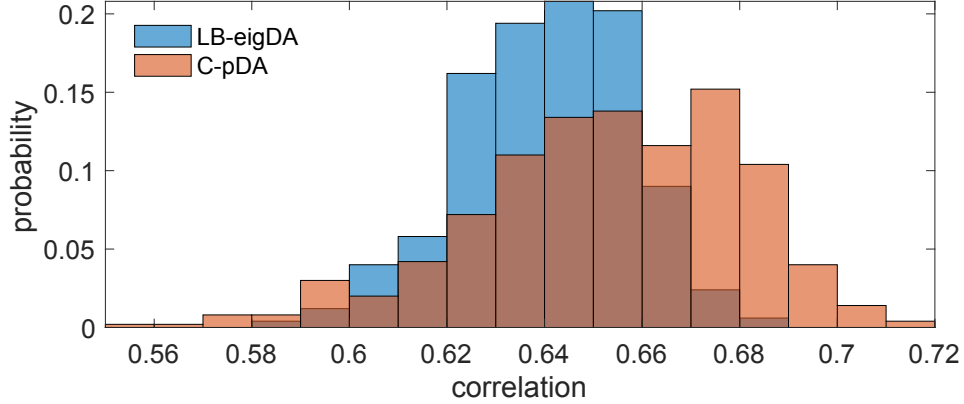
Figure 9: The distribution of the correlation between the 500 augmented thickness data and the thickness averaged over 500 real data (blue bar for the LB-eigDA and orange bar for the C-pDA method.

32, 64, and 128 filters, respectively. The initial learning rate was 10^{-3}, and the learning rate decay was 0.05 for every 20 epochs. In this experiment, the total sample from the ADNI dataset was 1709 (HC: $n = 1122$; AD: $n = 587$). Ten-fold cross-validation was adopted. One fold of real data was left out for testing. The remaining nine folds of data were further separated into training (75%) and validation (25%) sets. When the MRI datasets were separated into the training, validation, and testing sets, we considered subjects instead of MRI scans so that the scans from the same subjects were in the same set to avoid potential data over leakage.

The HC/AD classification accuracy based on the real ADNI data and the LB-based spectral CNN was $90.9 \pm 0.6\%$. However, when only a smaller set of the real data was available ($X\%$ of the training set), that is, the training sample size was reduced, the classification accuracy dropped as illustrated by the red dashed line in Fig.11. When only 10% of the real data was available, the classification accuracy was 75.8% and decreased 15% compared to that using the full ADNI data.

We previously showed that both C-pDA and LB-eigDA have the same results but C-pDA was more computationally efficient than LB-eigDA. Thus, the following experiments only used C-pDA with 109 filters and the Chebyshev polynomials of order $K = 5000$. As illustrated in Fig. 4, the training samples contained $X\%$ of real ADNI data and $1 - X\%$ augmented data, where $X = 10, 20, \cdots, 80$. We added $1 - X\%$ augmented data using C-pDA in the LB-based spectral CNN and computed the network performance us-
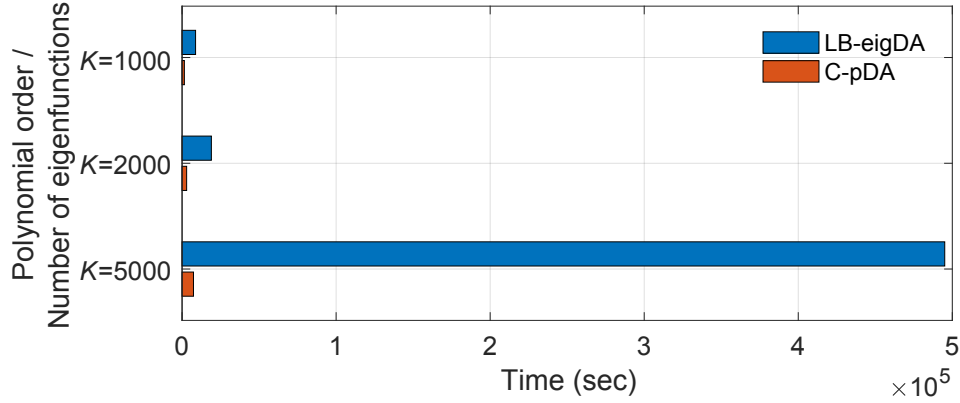
Figure 10: Computational time of the LB-eigDA (blue) and C-pDA (red-orange) methods for generating 500 augmented thickness data from 500 randomly selected subjects in ADNI.

ing the testing real data. The augmentation was done separately for the HC and AD groups. For instance, when 90% of the training samples were augmented data and 10% of the training samples were real data, the classification accuracy was 83.3% and improved by 7.5%. Fig.11 shows that C-pDA can increase the sample size and improve the HC/AD classification accuracy.

5   Discussion

This study introduces the LB-eigDA and C-pDA methods to generate augmented data on surfaces. Using the simulation with the ground truth label, we demonstrate that both methods improve the performance of graph-CNN. In particular, LB-eigDA has the potential to outperform C-pDA method since C-pDA approximates the LB spectral filters using Chebyshev polynomials. Nevertheless, when the mesh becomes large, LB-eigDA is computationally intensive while C-pDA is computationally efficient. C-pDA generates augmented thickness data and improves the AD classification accuracy in a real clinical application.

To our best knowledge, this study provides the first unbiased oversampling approaches for data augmentation on surfaces. These methods have a great potential to open new research areas in graph CNN in conjunction with generative adversarial networks (GANs). In particular, the formulation of the C-pDA method is consistent with that the LB-based spectral CNN
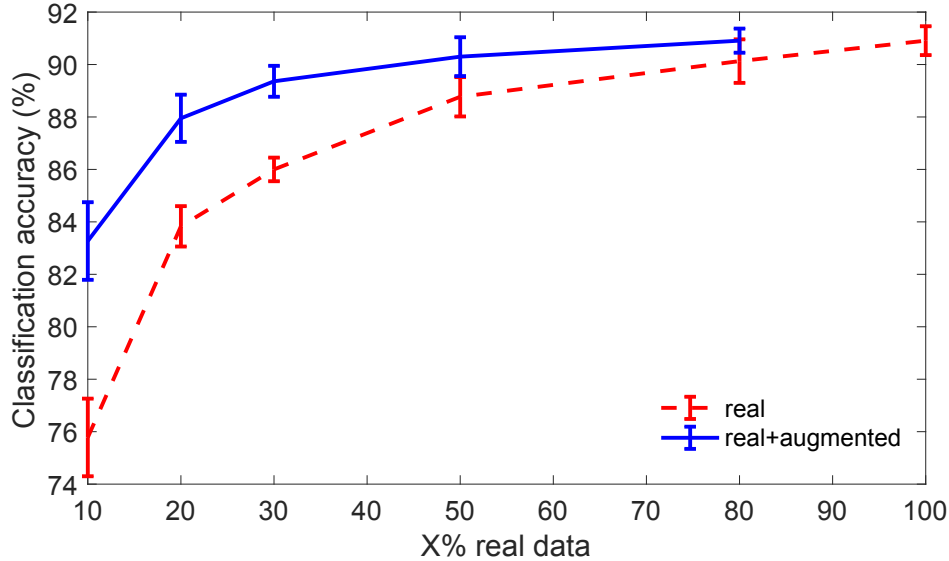
Figure 11: Classification accuracy. The red dashed line shows the classification accuracy when only $X\%$ of the real training samples were used in the training of the LB-based spectral CNN. The blue solid line shows the classification accuracy when only $X\%$ of the real raining set and $1 - X\%$ of the augmented data were used in the training of the LB-based spectral CNN, where the augmented data were generated by the C-pDA method.

[24, 36], which is feasible to adapt the C-pDA and graph network to the GAN framework. Further investigation will be needed.

Acknowledgements

References

[1] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. Van Der Laak, B. Van Ginneken, and C. I. Sánchez, "A survey on deep learning in medical image analysis," *Medical image analysis*, vol. 42, pp. 60–88, 2017.

[2] D. Shen, G. Wu, and H.-I. Suk, "Deep learning in medical image analysis," *Annual Review of Biomedical Engineering*, vol. 19, pp. 221–248, 2017.

[3] M. A. Mazurowski, P. A. Habas, J. M. Zurada, J. Y. Lo, J. A. Baker, and G. D. Tourassi, "Training neural network classifiers for medical decision making: The effects of imbalanced datasets on classification performance," *Neural Networks*, vol. 21, no. 2-3, pp. 427–436, 2008.

[4] J. Ker, L. Wang, J. Rao, and T. Lim, "Deep learning applications in medical image analysis," *Ieee Access*, vol. 6, pp. 9375–9389, 2017.

[5] J. L. Leevy, T. M. Khoshgoftaar, R. A. Bauder, and N. Seliya, "A survey on addressing high-class imbalance in big data," *Journal of Big Data*, vol. 5, no. 1, p. 42, 2018.

[6] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *Journal of Big Data*, vol. 6, no. 1, p. 60, 2019.

[7] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang, "Random erasing data augmentation," *arXiv preprint arXiv:1708.04896*, 2017.

[8] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv preprint arXiv:1412.6572*, 2014.

[9] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2096–2030, 2016.

[10] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.

[11] T. DeVries and G. W. Taylor, "Dataset augmentation in feature space," in *5th International Conference on Learning Representations,*

ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings. OpenReview.net, 2017. [Online]. Available: https://openreview.net/forum?id=HyaF53XYx

[12] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.

[13] X. Yi, E. Walia, and P. Babyn, "Generative adversarial network in medical imaging: A review," *Medical Image Analysis*, vol. 58, p. 101552, 2019.

[14] X. Yang, A. Goh, S. Chen, and A. Qiu, "Evolution of hippocampal shapes across the human lifespan," *Hum Brain Mapp.*, vol. 34, pp. 3075–3085, 2013.

[15] Y. Fan, R. Gur, R. Gur, X. Wu, D. Shen, M. Calkins, and C. Davatzikos, "Unaffected family members and schizophrenia patients share brain structure patterns: a high-dimensional pattern classification study," *Biol Psychiatry*, vol. 63, no. 1, pp. 118–124, 2008.

[16] L. G. Apostolova, I. D. Dinov, R. A. Dutton, K. M. Hayashi, A. W. Toga, J. L. Cummings, and P. M. Thompson, "3D comparison of hippocampal atrophy in amnestic mild cognitive impairment and alzheimer's disease," *Brain*, vol. 129, pp. 2867–2873, 2006.

[17] A. Qiu, C. Fennema-Notestine, A. Dale, M. Miller, and the Alzheimer's Disease Neuroimaging Initiative, "Regional shape abnormalities in mild cognitive impairment and Alzheimer's disease," *Neuroimage*, vol. 45, pp. 656–661, 2009.

[18] J. Atwood and D. Towsley, "Diffusion-convolutional neural networks," *arXiv preprint arXiv:1511.02136*, 2015.

[19] M. Niepert, M. Ahmed, and K. Kutzkov, "Learning convolutional neural networks for graphs," in *Proceeding of the 33rd International Conference on Machine Learning.* ACM, 2016, p. 2014—2023.

[20] D. K. Duvenaud, D. Maclaurin, J. Aguilera-Iparraguirre, R. Gomez-Bombarelli, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams, "Convolutional networks on graphs for learning molecular fingerprints," *arXiv preprint arXiv:1509.09292*, 2015.

[21] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel, "Gated graph sequence neural networks," *arXiv preprint arXiv:1511.05493*, 2015.

[22] B. Perozzi, R. Al-Rfou, and S. Skiena, "Deepwalk: Online learning of social representations," in *Proceedings of the 20th ACM SIGKDD*. ACM, 2014, p. 701–710.

[23] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun, "Spectral networks and locally connected networks on graphs," *arXiv preprint arXiv:1312.6203*, 2013.

[24] M. Defferrard, X. Bresson, and P. Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," in *Proceedings of the 30th International Conference on Neural Information Processing Systems*, ser. NIPS, 2016, pp. 3844–3852.

[25] M. Henaff, J. Bruna, and Y. LeCun, "Deep convolutional networks on graph-structured data," *arXiv preprint arXiv:1506.05163*, 2015.

[26] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *arXiv preprint arXiv:1609.02907*, 2016.

[27] L. Yi, H. Su, X. Guo, and L. Guibas, "Syncspeccnn: Synchronized spectral cnn for 3d shape segmentation," in *Computer Vision and Pattern Recognition (CVPR), Conference on*. IEEE, 2017, pp. 6584–6592.

[28] S. I. Ktena, S. Parisot, E. Ferrante, M. Rajchl, M. Lee, B. Glocker, and D. Rueckert, "Distance metric learning using graph convolutional networks: Application to functional brain networks," *arXiv preprint arXiv:1703.02161*, 2017.

[29] D. I. Shuman, B. Ricaud, and P. Vandergheynst, "Vertex-frequency analysis on graphs," *Applied and Computational Harmonic Analysis*, vol. 40, no. 2, p. 260–291, 2016.

[30] Y. Tang, "Deep learning using linear support vector machines," *arXiv preprint arXiv:1306.0239*, 2013.

[31] S. Ravanbakhsh, J. Schneider, and B. Poczos, "Deep learning with sets and point clouds," *arXiv preprint arXiv:1611.04500*, 2016.

[32] Y. Wang, H. Ombao, and M. Chung, "Topological data analysis of single-trial electroencephalographic signals," *Annals of Applied Statistics*, vol. 12, pp. 1506–1534, 2018.

[33] M. Chung, L. Xie, S.-G. Huang, Y. Wang, J. Yan, and L. Shen, "Rapid acceleration of the permutation test via transpositions," in *International Workshop on Connectomics in Neuroimaging*, vol. 11848. Springer, 2019, pp. 42–53.

[34] Z. Hussain, F. Gimenez, D. Yi, and D. Rubin, "Differential data augmentation techniques for medical imaging classification tasks," in *Annual Symposium proceedings (AMIA)*, 2017, p. 979–984.

[35] M. Tan and A. Qiu, "Spectral Laplace-Beltrami wavelets with applications in medical images," *IEEE Transactions on Medical Imaging*, vol. 34, pp. 1005–1017, 2015.

[36] C.-Y. Wee, C. Liu, A. Lee, J. S. Poh, H. Ji, A. Qiu, and A. D. N. Initiative, "Cortical graph neural network for AD and MCI diagnosis and transfer learning across populations," *NeuroImage: Clinical*, vol. 23, p. 101929, 2019.

[37] R. R. Coifman and M. Maggioni, "Diffusion wavelets," *Applied and Computational Harmonic Analysis*, vol. 21, no. 1, pp. 53–94, 2006.

[38] D. K. Hammond, P. Vandergheynst, and R. Gribonval, "Wavelets on graphs via spectral graph theory," *Applied and Computational Harmonic Analysis*, vol. 30, no. 2, pp. 129 – 150, 2011.

[39] W. H. Kim, D. Pachauri, C. Hatt, M. K. Chung, S. Johnson, and V. Singh, "Wavelet based multi-scale shape features on arbitrary surfaces for cortical thickness discrimination," in *Advances in Neural Information Processing Systems*, 2012, pp. 1241–1249.

[40] C. Donnat, M. Zitnik, D. Hallac, and J. Leskovec, "Learning structural node embeddings via diffusion wavelets," in *Proc. 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018, pp. 1320–1329.

[41] S.-G. Huang, I. Lyu, A. Qiu, and M. K. Chung, "Fast polynomial approximation of heat kernel convolution on manifolds and its application to brain sulcal and gyral graph pattern analysis," *IEEE Transactions on Medical Imaging*, vol. 39, no. 6, pp. 2201–2212, 2020.

[42] M. Meyer, M. Desbrun, P. Schröder, and A. H. Barr, "Discrete differential-geometry operators for triangulated 2-manifolds," in *Visualization and mathematics III*. Springer, 2003, pp. 35–57.

[43] M. Chung and J. Taylor, "Diffusion smoothing on brain surface via finite element method," in *Proceedings of IEEE International Symposium on Biomedical Imaging (ISBI)*, vol. 1, 2004, pp. 432–435.

[44] A. Qiu, D. Bitouk, and M. Miller, "Smooth functional and structural maps on the neocortex via orthonormal bases of the Laplace-Beltrami operator," *IEEE Transactions on Medical Imaging*, vol. 25, pp. 1296–1396, 2006.

[45] M. Chung, A. Qiu, S. Seo, and H. Vorperian, "Unified heat kernel regression for diffusion, kernel smoothing and wavelets on manifolds and its application to mandible growth modeling in CT images," *Medical Image Analysis*, vol. 22, pp. 63–76, 2015.

[46] S.-G. Huang, I. Lyu, A. Qiu, and M. Chung, "Fast polynomial approximation of heat diffusion on manifolds and its application to brain sulcal and gyral graph pattern analysis," *IEEE Transactions on Medical Imaging*, pp. under 2nd review, arXiv:1911.02 721, 2019.

[47] A. Oppenheim, R. Schafer, and J. Buck, *Discrete-time signal processing*. Upper Saddle River, NJ: Prentice Hall, 1999.

[48] N. A. Ranginwala, L. S. Hynan, M. F. Weiner, and C. L. White III, "Clinical criteria for the diagnosis of alzheimer disease: still good after all these years," *The American Journal of Geriatric Psychiatry*, vol. 16, no. 5, pp. 384–388, 2008.

[49] T. Tong, R. Wolz, Q. Gao, R. Guerrero, J. V. Hajnal, D. Rueckert, A. D. N. Initiative *et al.*, "Multiple instance learning for classification of dementia in brain mri," *Medical image analysis*, vol. 18, no. 5, pp. 808–818, 2014.

[50] R. Cuingnet, E. Gerardin, J. Tessieras, G. Auzias, S. Lehéricy, M.-O. Habert, M. Chupin, H. Benali, and O. Colliot, "Automatic classification of patients with Alzheimer's disease from structural MRI: a comparison of ten methods using the ADNI database," *Neuroimage*, vol. 56, no. 2, pp. 766–781, 2011.

[51] X. Liu, D. Tosun, M. W. Weiner, N. Schuff, A. D. N. Initiative *et al.*, "Locally linear embedding (lle) for mri based alzheimer's disease classification," *Neuroimage*, vol. 83, pp. 148–157, 2013.

[52] E. Hosseini-Asl, R. Keynton, and A. El-Baz, "Alzheimer's disease diagnostics by adaptation of 3d convolutional network," in *2016 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2016, pp. 126–130.

[53] S. Korolev, A. Safiullin, M. Belyaev, and Y. Dodonova, "Residual and plain convolutional neural networks for 3d brain mri classification," in *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*. IEEE, 2017, pp. 835–838.

[54] M. Liu, J. Zhang, E. Adeli, and D. Shen, "Landmark-based deep multi-instance learning for brain disease diagnosis," *Medical image analysis*, vol. 43, pp. 157–168, 2018.

[55] J. Islam and Y. Zhang, "Brain mri analysis for alzheimer's disease diagnosis using an ensemble system of deep convolutional neural networks," *Brain Informatics*, vol. 5, no. 2, p. 2, 2018.

[56] S. Basaia, F. Agosta, L. Wagner, E. Canu, G. Magnani, R. Santangelo, M. Filippi, A. D. N. Initiative *et al.*, "Automated classification of alzheimer's disease and mild cognitive impairment using a single mri and deep neural networks," *NeuroImage: Clinical*, vol. 21, p. 101645, 2019.

[57] B. Fischl, D. H. Salat, E. Busa, M. Albert, M. Dieterich, C. Haselgrove, A. Van Der Kouwe, R. Killiany, D. Kennedy, S. Klaveness *et al.*, "Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain," *Neuron*, vol. 33, no. 3, pp. 341–355, 2002.

[58] J. Zhong, D. Y. L. Phua, and A. Qiu, "Quantitative evaluation of lddmm, freesurfer, and caret for cortical surface mapping," *Neuroimage*, vol. 52, no. 1, pp. 131–141, 2010.

[59] J. Du, L. Younes, and A. Qiu, "Whole brain diffeomorphic metric mapping via integration of sulcal and gyral curves, cortical surfaces, and images," *NeuroImage*, vol. 56, no. 1, pp. 162–173, 2011.

[60] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An introduction to statistical learning*. Springer, 2013, vol. 112.