

Interpretation of Brain Morphology in Association to Alzheimer’s Disease Dementia Classification Using Graph Convolutional Networks on Triangulated Meshes

Emanuel Azcona^{1,4}[0000–0002–1896–4823], Pierre Besson^{2,4}[0000–0002–0591–7542], Yunan Wu^{1,4}[0000–0001–6980–9746], Arjun Punjabi^{1,4}[0000–0003–2770–3112], Adam Martersteck^{3,4}[0000–0001–6997–9390], Amil Dravid^{1,4}[0000–0001–6007–0690], Todd B. Parrish^{3,4}[0000–0002–1184–1572], S. Kathleen Bandt^{2,4}[0000–0002–1525–2775], and Aggelos K. Katsaggelos^{1,4}[0000–0003–4554–0070] for the Alzheimers Disease Neuroimaging Initiative*

¹ Image and Video Processing Laboratory, Department of Electrical and Computer Engineering, Northwestern University, IL, USA, <https://ivpl.northwestern.edu>

² Advanced NeuroImaging and Surgical Epilepsy (ANISE) Lab, Northwestern Memorial Hospital, IL, USA, <https://anise-lab.com/>

³ Neuroimaging Laboratory, Department of Radiology, Northwestern University, IL, USA, neuroimaging.northwestern.edu

⁴ Augmented Intelligence in Medical Imaging, Northwestern University, IL, USA

Abstract. We propose a mesh-based technique to aid in the classification of Alzheimer’s disease dementia (ADD) using mesh representations of the cortex and subcortical structures. Deep learning methods for classification tasks that utilize structural neuroimaging often require extensive learning parameters to optimize. Frequently, these approaches for automated medical diagnosis also lack visual interpretability for areas in the brain involved in making a diagnosis. This work: (a) analyzes brain shape using surface information of the cortex and subcortical structures, (b) proposes a residual learning framework for state-of-the-art graph convolutional networks which offer a significant reduction in learnable parameters, and (c) offers visual interpretability of the network via class-specific gradient information that localizes important regions of interest in our inputs. With our proposed method leveraging the use of cortical and subcortical surface information, we outperform other machine learning methods with a 96.35% testing accuracy for the ADD vs. healthy control problem. We confirm the validity of our model by observing its performance in a 25-trial Monte Carlo cross-validation. The generated visualization maps in our study show correspondences with

* Data used in preparation of this article were obtained from the Alzheimers Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf

current knowledge regarding the structural localization of pathological changes in the brain associated to dementia of the Alzheimer’s type.

Keywords: Graph convolutional networks · Alzheimer’s disease classification · triangulated meshes · neural network interpretability.

1 Introduction

Alzheimer’s disease dementia (ADD) is a clinical syndrome characterized by progressive amnesic multidomain cognitive impairment [27]. The causative underlying pathology is Alzheimers disease (AD), defined as the co-occurrence of neurofibrillary tangles and amyloid-beta plaques. Globally, the number of individuals living with AD is expected to reach 1 out of 85 people by the year 2050 [4]. Automated methods for the computer-aided clinical diagnosis of ADD has been an area of interest in the medical imaging community for the development of assistive tools aiding in the visual inspection of structural information captured by magnetic resonance imaging (MRI).

Previous studies in the neuroanatomical pathologies of AD have demonstrated correlations in cortical folding pattern [5] and different neurodegenerative pathologies. Specific patterns of atrophy in the cortex and subcortical structures have been linked to AD [21, 25]. For example, [5] discusses a potential to focus on high variability in association cortices like the intermediate sulcus of Jensen. As [28] also points out, widespread cortical thinning and a greater rate of atrophy is present in temporal lobe regions, primarily the left parahippocampal gyrus, for subjects with AD. Furthermore, Jong *et al.* [6] discuss irregularities like reduced putamen and thalamus volumes for subjects with AD. In studies such as ADNI, it is common to find bias towards more left-sided atrophy because of the verbal language tests given to assess memory function [8]. For example, if asymmetrical atrophy of the language network is more prominent, subjects may perform worse on verbal tests and be diagnosed with dementia earlier.

Machine learning (ML) methods have been a growing area of interest in the automated clinical diagnosis for ADD. [2, 24, 38] discuss the use of support vector machines (SVMs) in unimodal and multimodal imaging pipelines for the automated classification of ADD using MRI, PET, and cerebrospinal fluid (CSF). In [23, 30], the use of MRI and PET imaging in multimodal convolutional neural networks (CNNs) for ADD diagnosis is discussed. SVM-based approaches, like those used in [2, 24, 38], have historically been hard to interpret, expensive to train, and often serve as the logical choice only when there is enough domain expertise to construct meaningful kernels. Multimodal volumetric CNNs like [30], often require a lot of memory and frequently are limited to smaller-batch operations or using lower resolution 3D volumes.

Motivated by 3D object detection via surfaces [26], cortical and subcortical irregularities correlated with ADD, our work uses mesh manifolds of the cortex and subcortical structures in the diagnosis of ADD. Our technique leverages a reduction in computational complexity offered by [7]. In [29], Parisot *et al.*

leverage this work from [7] to make similar predictions for Alzheimer’s disease and Autism using graph convolutional networks (GCNs) on ADNI/ABIDE subject population graphs. In [31], their convolutional mesh autoencoder (CoMA) framework uses the same GCN basis from [7] on human face surface meshes to generate new meshes from a learned distribution conditioned on facial expression labels. Their network is also able to reconstruct input meshes from compressed 8-dimensional representations with a 50% reduction in reconstruction error, while using 75% fewer parameters than volumetric models that operate on voxels.

The interpretability of results from ML models has remained an open issue in highlighting regions of interest (ROI) in relation to classification decisions. In this paper we demonstrate that it is possible to (1) extract meaningful surface meshes of the cortex and subcortical structures, (2) achieve accurate predictions for the clinical binary classification of ADD using meshes, (3) extract class-discriminative localization maps for interpretable ROI, and (4) reduce the number of learnable parameters.

2 Methods

Data used in the preparation of this article were obtained from the Alzheimers Disease Neuroimaging Initiative (ADNI) database (<https://adni.loni.usc.edu>). The ADNI was launched in 2003 as a public-private partnership, led by Principal Investigator Michael W. Weiner, MD. The primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimers disease (AD).

2.1 Localized Spectral Filtering on Graphs

Spectral-based graph convolution methods inherit ideas from a graph signal processing (GSP) perspective as described by [37]. Like [7], our work focuses on using undirected graphs defined by a finite set of vertices, $\mathcal{V}$, with $N = |\mathcal{V}|$ vertices, and a corresponding set of edges, $\mathcal{E}$, with scalar edge weights, $e_{ij} = e_{ji} \in \mathcal{E}$, which are stored in the i^{th} rows and j^{th} columns of the adjacency matrix, $\mathbf{A} \in \mathbb{R}^{N \times N}$. A graph’s node attributes are defined using the node feature matrix $\mathbf{X} \in \mathbb{R}^{N \times F}$ where each column, $\mathbf{x}_i \in \mathbb{R}^F$, represents the feature vector for a particular shared feature across each of the vertices, $v_i \in \mathcal{V}$.

A great emphasis in GSP is placed on the normalized graph Laplacian, $\mathbf{L} = \mathbf{I}_N - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$, where $\mathbf{I}_N$ is the identity matrix and $\mathbf{D}_{ii} = \sum_j \mathbf{A}_{ij}$ is the diagonal matrix of node degrees. $\mathbf{L}$ can be factored via the eigendecomposition: $\mathbf{L} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T$, where $\mathbf{U} \in \mathbb{R}^{N \times N}$ is the complete set of orthonormal eigenvectors for $\mathbf{L}$ and $\mathbf{\Lambda} = \text{diag}([\lambda_0, \dots, \lambda_{N-1}]) \in \mathbb{R}^{N \times N}$ is the corresponding set of eigenvalues. Given a spectral filter, g_θ , defined in the graph’s Fourier space [34] as a polynomial of the Laplacian, $\mathbf{L}$, and $\mathbf{U}$ ’s orthonormality, we can

filter $\mathbf{x}$ via multiplication s.t.

$$g_\theta *_{\mathcal{G}} \mathbf{x} = g_\theta(\mathbf{L})\mathbf{x} = g_\theta(\mathbf{U}\mathbf{\Lambda}\mathbf{U}^T)\mathbf{x} = \mathbf{U}g_\theta(\mathbf{\Lambda})\mathbf{U}^T\mathbf{x}, \quad (1)$$

where $\theta \in \mathbb{R}^N$ are the parameters of the filter g_θ and $*_{\mathcal{G}}$ is the spectral convolution operator notation borrowed from [7]. Furthermore, $\mathbf{U}^T\mathbf{x}$ is the *graph Fourier transform* (GFT) of the graph signal $\mathbf{x}$, $g_\theta(\mathbf{\Lambda})$ is a filter defined using the spectrum (eigenvalues) of the normalized Laplacian, $\mathbf{L}$, and the left-sided multiplication with $\mathbf{U}$ is the *inverse*-GFT (IGFT). In this context, convolution is implicitly performed by using the duality property of the Fourier transform s.t. a spectral filter is first multiplied with the GFT of a signal, and then the IGFT of their product is determined.

Our approach uses Chebyshev polynomials of the first kind [1, 7] to approximate g_θ using the graph's spectrum s.t.

$$g_\theta(\tilde{\mathbf{L}}) = \sum_{k=0}^{K-1} \theta_k T_k(\tilde{\mathbf{L}}), \quad (2)$$

for the scaled Laplacian $\tilde{\mathbf{L}} = \frac{2\mathbf{L}}{\lambda_{max}} - \mathbf{I}_N$, where λ_{max} is the largest eigenvalue in $\mathbf{\Lambda}$, and K can be interpreted as the kernel size. Chebyshev polynomials of the first kind are defined by the recurrence relation, $T_k(\tilde{\mathbf{L}}) = 2\tilde{\mathbf{L}}T_{k-1}(\tilde{\mathbf{L}}) - T_{k-2}(\tilde{\mathbf{L}})$ where $T_0(\tilde{\mathbf{L}}) = \mathbf{I}$ and $T_1(\tilde{\mathbf{L}}) = \tilde{\mathbf{L}}$ as shown in [7].

2.2 Mesh Extractions of Cortical & Subcortical Structures

Using FreeSurfer v6.0 [10], all MRIs were denoised followed by field inhomogeneity correction, and intensity and spatial normalization. Inner cortical surfaces (interface between gray and white matter) and outer cortical surfaces (CSF/gray matter interface) were extracted and automatically corrected for topological defects. Additionally, seven subcortical structures per hemisphere were segmented (amygdala, nucleus accumbens, caudate, hippocampus, pallidum, putamen, thalamus) and modeled into surface meshes using SPHARM-PDM (<https://www.nitrc.org/projects/spharm-pdm>).

Surfaces were inflated, parameterized to a sphere, and registered to a corresponding spherical surface template using a rigid-body registration to preserve the cortical [10] and subcortical [3] anatomy. Surface templates were converted to meshes using their triangulation schemes. A scalar edge weight, e_{ij} , was assigned to connect vertices v_i and v_j using their geodesic distance, ψ_{ij} , s.t.

$$e_{ij} = e_{ji} = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{\psi_{ij}}{\sigma}\right)^2}. \quad (3)$$

Surface templates were parcellated using a hierarchical bipartite partitioning of their corresponding mesh. Starting with their initial mesh representation of densely triangulated surfaces, spectral clustering was used to define two partitions. These two groups were then each separated yielding four child partitions,

and this process was repeated until the average distance across neighbor partitions was below 2.5 mm. For each partition, the central node was defined as the node whose centrality was highest and the distance across two partitions was defined as the geodesic distance (in mm) across the central vertices. Two partitions were neighbors if at least one node in each partition were connected. Finally, partitions were numbered so that partitions $2i$ and $2i + 1$ at level L had the same parent partition i at level $L - 1$. Therefore, for each level a graph was obtained s.t. the vertices of the graph were the central vertices of the partitions and the edges across neighboring vertices were weighted as in Eq. 3. This serves as an improvement upon [7] to ensure that no singleton is ever produced by pooling operations for the cortex and subcortical structures. At the finest level, meshes had a total of 47,616 vertices: 32,768 vertices for the cortex and 14,848 vertices to represent the subcortical structures.

Vertex features were defined as the Cartesian coordinates of the surface vertices in the subjects native space registered to the surface templates. This can create issues if the original scans are not registered to the same template, as was also done by Ranjan *et al.* in [31]. Similar studies, like that of Gutierrez-Becker and Wachinger [15], implement “rotation network” modules as the first few layers of their neural network (NN) architecture to aid in correcting and aligning their samples to a common template. Performing our template registration as an additional preprocessing step reduces the complexity of our NN architecture and eliminates the need of incorporating an “alignment” term to our cost function to optimize later, as was needed in [15].

Cortical vertices were assigned 6 features: the x , y , and z coordinates of both the white matter (WM) and gray matter (GM) vertices in the native space. This was decided because vertices on these surfaces use the same edge weights and therefore the same “faces” with different coordinates for the vertices of the respective triangles. Similar to the cortex, subcortical vertices had 3 features: their corresponding x , y , and z coordinates in the native space as well. To maintain the same number of features for all vertices per scan, the corresponding cortical and subcortical feature matrices were block-diagonalized into a single node feature matrix per scan s.t. $\mathbf{X} \in \mathbb{R}^{47,616 \times 9}$. Sample meshes extracted from a randomly selected HC and one with ADD are demonstrated in Figure 1.

2.3 Residual Network Architecture

Inspired by the work of He *et al.* in [16], we propose an improvement upon ChebNet [7] using residual connections within GCNs, which have been shown in prior work to address the common “vanishing gradient” problem and improve the performance of deep NNs. Typically, these types of residual networks (ResNets) are implemented by using batch normalization (BN) [19] before a ReLU activation function, and followed by convolution as seen in Fig. 2. Using ResBlocks (Fig. 2), max-pooling operations as described by Defferrard *et al.* [7], and a standard fully connected (FC) layer [32], the total architecture used in our study is defined in Fig. 3. An additional ResBlock, which we refer to as a “post-ResBlock,” was




Fig. 1. Cortical meshes from a randomly selected HC subject (blue) and meshes of the subcortical structures from a randomly selected ADD subject (yellow). Presented are lateral views (a-b) of the HC’s left hemisphere (LH) and right hemisphere (RH) cortical meshes respectively. Medial views of the ADD subject’s LH and RH subcortical structure meshes are also presented (c-d).

introduced prior to the FC layer as a linear mapping tool to match the number of FC units.




Fig. 2. Single ResBlock in the GCN architecture used in this study. Linear mapping of F_{in} to F_{out} channels is implemented using a convolutional layer, $*_{\mathcal{G}}$. This is done to match the number of input features to the number of desired feature maps.

2.4 Grad-CAM Mesh Adaptation

Interpretability of CNNs was addressed by [33] via their gradient-weighted class activation map (Grad-CAM) approach. In their work, images are fed to CNNs and gradients for each class score (logits prior to softmax) are extracted at the last convolutional layer. Using these gradients, they perform a global average pooling (GAP) operation for each feature map per class to extract “neuron importance weights,” $\alpha_c^{(k)} \in \mathbb{R}^{c \times k}$, whose formulation we readapt for meshes s.t.

$$\alpha_c^{(k)} = \frac{1}{N} \sum_n \frac{\partial y_c}{\partial A_n^{(k)}}, \quad (4)$$

where y_c corresponds to the class score of class c , and $A_n^{(k)}$ refers to the value at vertex n for the k -th feature map $A^{(k)} \in \mathbb{R}^N$. A set of neuron importance




Fig. 3. Residual GCN used for the binary classification of ADD. In this study, max-pooling operations are used to downsample the vertex dimension by a factor of 2.

weights, $\alpha_c^{(k)}$, is extracted for each k -th feature map, $A^{(k)}$, and projected onto them to get the class activation maps (CAMs) s.t.

$$M_c = \text{ReLU} \left(\sum_k \alpha_c^{(k)} A^{(k)} \right) \in \mathbb{R}^N. \quad (5)$$

As a consequence of pooling, CAMs are upsampled to the same number of nodes as the input mesh for a direct “overlay” using a trivial interpolation by going backward along the hierarchical tree used by the pooling operations.

3 Experimental Design

3.1 Dataset & Preprocessing

T1-weighted MRIs from ADNI [20] were selected with ADD/HC diagnosis labels given up to 2 months after the corresponding scan. This was taken as a precaution to ensure that each diagnosis had clinical justification. The dataset in our study consisted of 1,191 different scans for 435 unique subjects. Section 3.2 outlines our stratified data splitting strategy to ensure no data leakage occurs at the subject level across the training, validation, and testing sets [12].

Meshes for each MRI were extracted following the process described in Section 2.2. The spatial standard deviation from Eq. 3, σ , was set to 2 ad-hoc. The visual quality for each mesh was assessed manually via a direct overlay over slices of the corresponding MRI. Laplacians for the cortex and each subcortical structure were block-diagonalized to create one overall $\mathbf{L}$ representing a single mesh with multiple connected components. Extracted feature matrices for each sample were min-max normalized per feature to the interval $[-1, 1]$ prior to feeding batches of data into the networks. The added zeros during block-diagonalization (as discussed in Section 2.2) were ignored during each normalization step.

3.2 Network Architecture & Training

Extra care was taken in the shuffling of samples to avoid bias from subject overlap in our cross-validation [12]. A custom dataset splitting function was implemented s.t. the distribution of labels was preserved amongst each set while also ensuring

to avoid subject overlap. 20% of the samples were selected at random for the testing set. Of the remaining 80%, 20% of those were withheld as the validation set, while the remaining belonged to the training set. A 25-trial Monte Carlo cross-validation was performed using this data split scheme.

The architecture in Fig. 3 was implemented using 16 kernels per convolutional layer (not including the post-ResBlock), Chebyshev polynomials of order $K = 3$, and pooling windows of size $p = 2$. Four alternating ResBlock and pooling layers were cascaded as shown in Fig. 3 prior to the post-ResBlock. The number of units at the post-ResBlock and FC layer was 128. Our GCN was optimized by minimizing a standard binary cross-entropy loss function

$$\mathcal{L} = -\frac{1}{N} \sum_{n=1}^N y_n \log(\hat{y}_n) + (1 - y_n) \log(1 - \hat{y}_n), \quad (6)$$

where $\hat{y}_n$ is the predicted class for the n^{th} sample of N total samples and y_n is the ground truth label for the same sample index, n .

Networks were trained using batches of 32 samples per step for 100 epochs in each Monte Carlo trial. The Adam [22] optimizer was used with a learning rate of 5×10^{-4} and a learning rate decay of 0.999. Experiments were implemented in Python 3.6 using Tensorflow 1.13.4 using an NVIDIA GeForce GTX TITAN Z GPU in a Dell Precision Tower 7910 with Linux Mint 19.2.

4 Results & Discussion

4.1 ADD vs. HC Classification

Our cross-validation includes the same multilayer perceptron (MLP) classifier architecture, ridge classifier, and a 100-estimator random forest classifier set up by Parisot *et al.* in [29], where a similar graph approach is also used on the classification of ADD based on population graphs. The MLP designed was synonymous to the design in [29] s.t. the number of hidden layers and parameters was the same as our GCNs. Demonstrated in Figure 4, our GCN outperformed other standard classifiers not limited to graph methods on our dataset split.




Fig. 4. Monte Carlo cross-validation accuracy results for GCN and baseline model architectures from [29] used on brain meshes.

Table 1. Model comparison to classifiers in studies not limited to surface methods.

Study	Data	ADD/HC	Acc. (%)	Sens. (%)	Spec. (%)	AUC (%)
[30]	MRI	–/– (723)	73.76	–	–	–
[30]	MRI+PET _{amyloid}	–/– (723)	92.34	–	–	–
[24]	MRI+PET _{FDG}	51/52	94.37	94.71	94.04	97.24
[36]	MRI+CSF	96/111	91.80	88.50	94.60	95.80
[18]	MRI	228/188	84.13	82.45	85.63	90.00
[2]	MRI	92/94	93.01	89.13	96.80	93.51
[17]	MRI	70/70	97.60	–	–	–
[14]	MRI	200/232	94.74	95.24	94.26	–
Ours	MRI	167/265	96.35	92.37	96.74	96.84

The results in Table 1 highlight comparable metrics of our model versus other studies that operate on voxels from full 3D MRI volumes, including [30]. In their work, Punjabi *et al.* train a multi-modal CNN using both volumetric MRI and FDG-PET imaging for the same task, which we outperform while training and evaluating on a smaller subset of their subject population. Furthermore, volumetric models like those in [30] rely on 3D CNNs with far more learned parameters, e.g. [30]’s 200,194,502 weights ($\times 2$ for fusion model), in comparison to our GCN’s 497,522 learned parameters needed for comparable results. Like [31], we also achieve comparable results with far less learning parameters by working on meshes and focusing on brain *shape* instead of working on raw voxels obtained from MRIs and using voxel-based approaches.

4.2 Class Activation Map Visualization

By employing Grad-CAM on our best GCN, an average CAM was generated for true positive (TP) predictions (Fig. 5). We project our CAM onto the cortical template [11] provided by FreeSurfer [10] and the homemade subcortical structure templates detailed in [3]. The color scale highlights areas from least-to-most influential in TP predictions. The patterns in the CAM match previously described distributions of cortical and subcortical atrophy [9,21]. One reason we may observe a mismatch between the CAM and expected atrophy in the inferior parietal lobule could be the degree of variability in highly folded association cortex, e.g., the intermediate sulcus of Jensen is found only in some individuals [5,35]. The slightly more left lateralized pattern in the CAM aligns with previous reports that propose greater pathologic burden and neurodegeneration of the language network which leads to worsening on verbal-based neuropsychological measures of memory resulting in a diagnosis for ADD [8].

5 Conclusion & Future Work

In this work, we demonstrated the effectiveness of using cortical and subcortical surface meshes in the context of binary ADD clinical diagnosis and ROI visualization in TP predictions. Furthermore, we compared the cross-validation results




Fig. 5. Average TP CAMs on the cortical template from [10,11] (top) and subcortical structures from [3] (bottom) including: (a-b, e-f) lateral-medial views of the LH respectively, (c-d, g-h) medial-lateral views of the RH respectively.

of our model for the same ADD vs. HC problem using other ML models on our data. Additionally, our final results were comparable to the results of other studies that use traditional neuroimaging modalities as inputs. When compared to the performance of the multimodal approach used in [30], our model outperforms their approach, thus potentially indicating the reliability of leveraging shape information represented as meshes to perform the same binary classification task.

Natural extensions of this work could be to (1) expand our classification problem to include a third class from ADNI, mild cognitive impairment (MCI), (2) increase the population in our study to include those in ADNI3 [20], (3) work on longitudinal predictions, and (4) compare our model’s performance in using only the cortex, subcortical structures, or both. Additionally, having a 3D-volume-to-mesh dataset offers the potential for developing generative networks, as in [13], for performing the graph extraction preprocessing step described in Section 2.2. This will provide more autonomy and limit the need for the manual quality assessment (QA) of meshes as a part of our pipeline.

6 Acknowledgements

This work was funded in part by the Biomedical Data Driven Discovery Training Grant from the National Library of Medicine (5T32LM012203) through Northwestern University, and the National Institute on Aging. The authors would also like to thank the QUEST High Performance Computing Cluster at Northwestern University for computational resources.

Data collection and sharing for this project was funded by the Alzheimers Disease Neuroimaging Initiative (ADNI). Data collection and sharing for this project was funded by the Alzheimer’s Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: AbbVie, Alzheimers Association; Alzheimers Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (<https://www.fnih.org>). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimers Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.

References

1. Arfken, G.B., Weber, H.J., Harris, F.E.: *Mathematical Methods for Physicists*. Academic Press, third edn. (2013). <https://doi.org/10.1016/C2013-0-10310-8>
2. Beheshti, I., et al.: Classification of alzheimer’s disease and prediction of mild cognitive impairment-to-alzheimer’s conversion from structural magnetic resource imaging using feature ranking and a genetic algorithm. *Computers in biology and medicine* **83**, 109–119 (2017). <https://doi.org/10.1016/j.combiomed.2017.02.011>
3. Besson, P., et al.: Intra-subject reliability of the high-resolution whole-brain structural connectome. *NeuroImage* **102**, 283 – 293 (2014). <https://doi.org/10.1016/j.neuroimage.2014.07.064>
4. Brookmeyer, R., et al.: Forecasting the global burden of Alzheimer’s disease. *Alzheimer’s and Dementia* **3**, 186–191 (2007). <https://doi.org/10.1016/j.jalz.2007.04.381>
5. Bruyn, G.: *Atlas of the cerebral sulci*, vol. 93. G. Thieme Verlag (1991)
6. De Jong, L.W., et al.: Strongly reduced volumes of putamen and thalamus in Alzheimer’s disease: An MRI study. *Brain* **131**(12), 3277–3285 (2008). <https://doi.org/10.1093/brain/awn278>
7. Defferrard, M., Bresson, X., Vandergheynst, P.: Convolutional neural networks on graphs with fast localized spectral filtering. In: *Advances in Neural Information Processing Systems*. pp. 3844–3852 (2016), <http://papers.nips.cc/paper/6081-convolutional-neural-networks-on-graphs-with-fast-localized-spectral-filtering.pdf>

8. Derflinger, S., et al.: Grey-matter atrophy in Alzheimer’s disease is asymmetric but not lateralized. *Journal of Alzheimer’s Disease* **25**(2), 347–357 (2011). <https://doi.org/10.3233/JAD-2011-110041>
9. Dickerson, B.C., et al.: The cortical signature of Alzheimer’s disease: Regionally specific cortical thinning relates to symptom severity in very mild to mild AD dementia and is detectable in asymptomatic amyloid-positive individuals. *Cerebral Cortex* **19**(3), 497–510 (2009). <https://doi.org/10.1093/cercor/bhn113>
10. Fischl, B.: FreeSurfer. *NeuroImage* **62**(2), 774–81 (8 2012). <https://doi.org/10.1016/j.neuroimage.2012.01.021>
11. Fischl, B., et al.: High-resolution intersubject averaging and a coordinate system for the cortical surface. *Human brain mapping* **8**(4), 272–284 (1999). [https://doi.org/10.1002/\(SICI\)1097-0193\(1999\)8:4<272::AID-HBM10>3.0.CO;2-4](https://doi.org/10.1002/(SICI)1097-0193(1999)8:4<272::AID-HBM10>3.0.CO;2-4)
12. Fung, Y.R., et al.: Alzheimer’s Disease Brain MRI Classification: Challenges and Insights. *arXiv preprint arXiv:1906.04231* (6 2019), <https://arxiv.org/abs/1906.04231>
13. Goodfellow, I.J., et al.: Generative adversarial nets. In: *Advances in Neural Information Processing Systems*. vol. 3, pp. 2672–2680 (2014), <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>
14. Gupta, A., Ayhan, M., Maida, A.: Natural image bases to represent neuroimaging data. In: *International conference on machine learning*. pp. 987–994 (2013), <http://proceedings.mlr.press/v28/gupta13b.pdf>
15. Gutiérrez-Becker, B., Wachinger, C.: Learning a conditional generative model for anatomical shape analysis. In: *International Conference on Information Processing in Medical Imaging*. pp. 505–516. Springer (2019). https://doi.org/10.1007/978-3-030-20351-1_39
16. He, K., et al.: Deep residual learning for image recognition. In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. vol. 2016-Decem, pp. 770–778 (2016). <https://doi.org/10.1109/CVPR.2016.90>
17. Hosseini-Asl, E., Keynton, R., El-Baz, A.: Alzheimer’s disease diagnostics by adaptation of 3d convolutional network. In: *2016 IEEE International Conference on Image Processing (ICIP)*. pp. 126–130. IEEE (2016). <https://doi.org/10.1109/ICIP.2016.7532332>
18. Hu, K., et al.: Multi-scale features extraction from baseline structure mri for mci patient classification and ad early diagnosis. *Neurocomputing* **175**, 132 – 145 (2016), [10.1016/j.neucom.2015.10.043](https://doi.org/10.1016/j.neucom.2015.10.043)
19. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: *32nd International Conference on Machine Learning, ICML 2015. Proceedings of Machine Learning Research*, vol. 1, pp. 448–456 (2015), <http://proceedings.mlr.press/v37/ioffe15.pdf>
20. Jack, C.R., et al.: The Alzheimer’s Disease Neuroimaging Initiative (ADNI): MRI methods (2008). <https://doi.org/10.1002/jmri.21049>
21. Kälin, A.M., et al.: Subcortical shape changes, hippocampal atrophy and cortical thinning in future Alzheimer’s disease patients. *Frontiers in Aging Neuroscience* (2017). <https://doi.org/10.3389/fnagi.2017.00038>
22. Kingma, D.P., Ba, J.L.: Adam: A method for stochastic optimization. In: *3rd International Conference on Learning Representations* (2015), <https://arxiv.org/pdf/1412.6980.pdf>
23. Li, R., et al.: Deep learning based imaging data completion for improved brain disease diagnosis. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. vol. 8675 LNCS, pp. 305–312 (2014). https://doi.org/10.1007/978-3-319-10443-0_39

24. Liu, F., Wee, C.Y., Chen, H., Shen, D.: Inter-modality relationship constrained multi-modality multi-task feature selection for alzheimer's disease and mild cognitive impairment identification. *NeuroImage* **84**, 466 – 475 (2014). <https://doi.org/10.1016/j.neuroimage.2013.09.015>
25. Liu, T., et al.: Cortical gyrification and sulcal spans in early stage Alzheimer's disease. *PLoS ONE* (2012). <https://doi.org/10.1371/journal.pone.0031083>
26. Masci, J., Boscaini, D., Bronstein, M.M., Vandergheynst, P.: Geodesic Convolutional Neural Networks on Riemannian Manifolds. In: *Proceedings of the IEEE International Conference on Computer Vision*. vol. 2015-Febru, pp. 832–840 (2015). <https://doi.org/10.1109/ICCVW.2015.112>
27. McKhann, G.M., et al.: The diagnosis of dementia due to Alzheimer's disease: Recommendations from the National Institute on Aging-Alzheimer's Association workgroups on diagnostic guidelines for Alzheimer's disease. *Alzheimer's and Dementia* **7**(3), 263–269 (2011). <https://doi.org/10.1016/j.jalz.2011.03.005>
28. Pacheco, J., et al.: Greater cortical thinning in normal older adults predicts later cognitive impairment. *Neurobiology of Aging* **36**(2), 903–908 (2015). <https://doi.org/10.1016/j.neurobiolaging.2014.08.031>
29. Parisot, S., Ktena, S.I., Ferrante, E., Lee, M., Guerrero, R., Glocker, B., Rueckert, D.: Disease prediction using graph convolutional networks: Application to autism spectrum disorder and alzheimers disease. *Medical image analysis* **48**, 117–130 (2018). <https://doi.org/10.1016/j.media.2018.06.001>
30. Punjabi, A., et al.: Neuroimaging modality fusion in Alzheimer's classification using convolutional neural networks. *PLoS ONE* **14**(12), 1–14 (2019). <https://doi.org/10.1371/journal.pone.0225759>
31. Ranjan, A., Bolkart, T., Sanyal, S., Black, M.J.: Generating 3d faces using convolutional mesh autoencoders. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 704–720 (2018). https://doi.org/10.1007/978-3-030-01219-9_43
32. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning Internal Representations by Error Propagation. In: *Readings in Cognitive Science: A Perspective from Psychology and Artificial Intelligence*, pp. 399–421. MIT Press, Cambridge, MA (2013)
33. Selvaraju, R.R., et al.: Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. In: *International Journal of Computer Vision* (2019). <https://doi.org/10.1109/ICCV.2017.74>
34. Shuman, D.I., et al.: The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine* **30**(3), 83–98 (2013). <https://doi.org/10.1109/MSP.2012.2235192>
35. Thompson, P.M., et al.: Cortical variability and asymmetry in normal aging and Alzheimer's disease. *Cerebral Cortex* **8**(6), 492–509 (1998). <https://doi.org/10.1093/cercor/8.6.492>
36. Westman, E., Muehlboeck, J.S., Simmons, A.: Combining mri and csf measures for classification of alzheimer's disease and prediction of mild cognitive impairment conversion. *NeuroImage* **62**(1), 229 – 238 (2012). <https://doi.org/10.1016/j.neuroimage.2012.04.056>
37. Wu, Z., et al.: A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems* pp. 1–21 (2020). <https://doi.org/10.1109/TNNLS.2020.2978386>
38. Zhang, D., Shen, D.: Multi-modal multi-task learning for joint prediction of multiple regression and classification variables in alzheimer's disease. *NeuroImage* **59**(2), 895 – 907 (2012). <https://doi.org/10.1016/j.neuroimage.2011.09.069>