2022 44th Annual International Conference of
the IEEE Engineering in Medicine & Biology Society (EMBC)
Scottish Event Campus, Glasgow, UK, July 11-15, 2022

2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC) | 978-1-7281-2782-8/22/$31.00 ©2022 IEEE | DOI: 10.1109/EMBC48229.2022.9871338

Probing the link between the APOE-ε4 allele and whole-brain gray matter using deep learning

Anees Abrol[1], Ihab Hajjar[2] and Vince Calhoun[1]

[1]Center for Translational Research in Neuroimaging and Data Science (TReNDS): Georgia State University, Georgia Institute of Technology, and Emory University, Atlanta, USA

[2]Department of Neurology, Emory University School of Medicine, Atlanta, USA

***Abstract*— The APOE-ε4 allele is a known genetic risk for Alzheimer's disease (AD). Thus, it can be reasoned that the APOE-ε4 allele would also impact neurodegeneration-associated structural brain changes. Here we probe if the APOE-ε4 genotype directly modulates the human brain's gray matter using a neural network trained on the whole-brain gray matter images from the cognitively normally aging (CN) and AD individuals. To investigate the linkage between the APOE-ε4 allele and whole-brain (voxel-wise) gray matter, we systematically profile our investigation in multiple classification tasks, including diagnostic classification and APOE-ε4 classification conjointly as well as independently. Results suggest that although the MRI data can reliably track and reflect neurodegenerative changes in the brain cross-sectionally, the APOE-ε4 status may not be distinguishable correspondingly. The nonexistence of a direct and convincing modulative effect of APOE-ε4 on the whole-brain gray matter indicates that the gray matter changes may be independent of the APOE-ε4 status, and instead characterize a non-APOE, comorbid mechanism in AD.**

I. Introduction

The ε4 allele of the Apolipoprotein E gene (APOE-ε4) is a strong genetic risk factor for Alzheimer's disease (AD). Thus, it can be reasoned that the APOE-ε4 allele may uniquely impact this neurodegenerative brain condition's morphological expression. While there is abundant literature demonstrating that the APOE-ε4 status is correlated with AD risk and affects AD progression likewise, it is not yet explored if the APOE-ε4 genotype modulates the human brain's gray matter directly.

Recently, deep learning (DL) approaches have shown great promise in diverse medical imaging applications [1-4]. DL approaches exploit the wealth of information available from raw or minimally preprocessed input images to facilitate the automatic and adaptive discovery of task-discriminative representations at multiple hierarchy levels as an integral part of the training procedure. DL approaches are increasingly used to study structural and functional brain imaging modalities with the ultimate objective of understanding mental health non-invasively [5-10]. Our recent work [7] demonstrates using structural MRI (sMRI) data that if trained following prevalent DL practices, DL methods have the potential to scale particularly well and consistently outperform standard machine learning (SML) methods in neuroimaging classification and regression tasks. Additionally, our prior DL study on predicting progression to AD, using baseline sMRI data alone, resulted in state-of-the-art performance even considering multimodal studies. Importantly, in both studies, we demonstrate how DL approaches can provide backward mapping to the input image space through methodical interpretations, therefore facilitating inferences on task-specific brain mechanisms, such as delineating the most influential data features in predicting an attempted task.

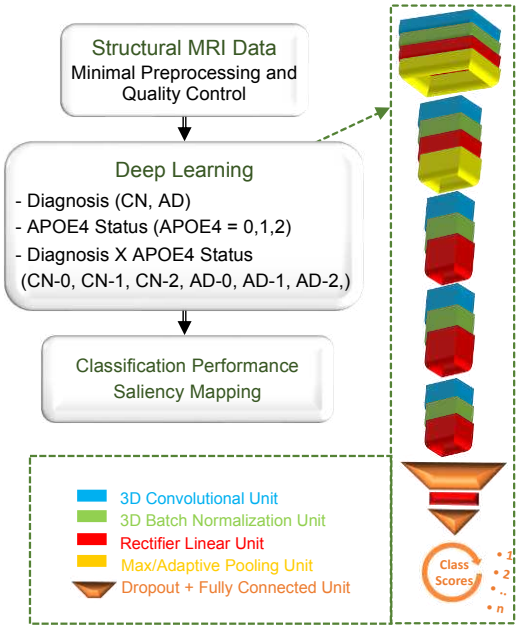

Figure 1: Study Workflow: This figure provides an overview of the experimental design. The ADNI dataset

The success mentioned above, and the flexibility of DL approaches motivate its application to the neuroimaging objective explored in this work. To investigate the association between the APOE-ε4 allele and whole-brain (voxel-wise) gray matter, here we perform a structural magnetic resonance imaging (sMRI) study to examine if the APOE-ε4 status directly modulates the whole-brain (voxel-wise) gray matter. We probe that using a neural network trained on the whole-brain gray matter images from the cognitively normally aging (CN) and AD individuals from the Alzheimer's Disease Neuroimaging Initiative (ADNI) data repository, systematically profiling our investigation in multiple

classification tasks, including diagnosis and APOE-ε4 status classification independently as well as conjointly (as detailed in the coming sections).

II. Materials and Methods

A. Data

The ADNI study procedures were approved by the institutional review boards of all participating centers as detailed in this document—https://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf. Written, informed consent was obtained from all subjects participating in the study according to the Declaration of Helsinki, and the study was approved by the institutional review board at each participating site.

This study worked with all cognitively normally aging (CN: *n=425*), and AD (CN: *n=293*) individuals with structural MRI scans available in the ADNI 1/2/GO/3 phases (as of September 7, 2021) that passed our image preprocessing pipeline quality check. The above-stated numbers are post-QC sample sizes (see next section for QC details). Notably, we used only one image per subject in all conducted research, retaining the earliest scan passing QC for all subjects. Detailed scanning parameters are available at http://adni.loni.usc.edu/methods/documents/mri-protocols/.

B. Data pre-processing

The ADNI sMRI data preprocessing pipeline featured segmentation into tissue probability maps for gray matter, white matter, and cerebral spinal fluid using the Statistical Parametric Mapping toolbox (SPM12). Using the same toolbox, the segmented gray matter maps were warped to standard Montreal Neurological Institute (MNI) space, modulated, and smoothed using a Gaussian kernel with a full width at half maximum (FWHM) of 6 mm. Quality control (QC) of the preprocessed sMRI datasets included discarding images that exhibited low correlation with individual and/or group level masks.

C. Classification Tasks

We undertook several classification tasks based on subject diagnosis and APOE-ε4 status. The following classification tasks were performed - (1) 2-way diagnostic classification (CN vs. AD), (2) 3-way APOE-ε4 status classification for the CN and AD groups (APOE non-carriers vs. heterozygous vs. homozygous, or APOE4=0 vs. APOE4=1 vs. APOE4=2), and (3) 6-way combined diagnosis by APOE-ε4 status classification tasks (CN-0 vs. CN-1 vs. CN-2 vs. AD-0 vs. AD-1 vs. AD-2).

D. Deep Learning Model and Pipeline

We trained the preprocessed gray matter volume images using our recent deep learning model and pipeline as implemented in our previous work [7]. This pipeline employed an end-to-end trained convolutional neural network (CNN) model to learn the differences in the morphological patterns revealed in the implemented classification tasks (Figure 1), details of which are provided in the following sub-sections.

- *DL training:* Our CNN model's training and testing routines were implemented on the NVIDIA CUDA parallel computing platform on the TReNDS slurm-managed cluster using the PyTorch Lightning research framework. We used the *SGD* (stochastic gradient descent) optimizer with learning rate and weight decay parameters of 0.001 and momentum value of 0.9 in this work. We used a batch size of 8 for all classification tasks in this work and employed the *StepLR* learning rate scheduler callback to decay the learning rate of each parameter by a factor of 0.3 every 50 epochs.
- *Cross-validation procedure:* The ADNI datasets (*n*=718) were stratified into non-overlapping training (*n*=512), validation (*n*=103) and test (*n*=103) partitions, and a stratified Monte-Carlo (i.e., repeated random sub-sampling) cross-validation procedure was employed. Each repetition sampled the data exactly once to ensure a consistent and valid statistical distribution of the evaluated test metrics.
- *Hyperparameter validation and model validation:* Hyperparameter tuning was employed using the training and validation partitions to tune the optimizer and learning rate. The *SGD* optimizer and a learning rate of 0.001 were validated for this experiment. We validated the best-performing model per the epoch with the highest validation balanced accuracy metric in this work. Using the validated model, we evaluated the performance on unseen, held-out test data samples to estimate test metrics for each cross-validation repetition and classification task.
- *Saliency Estimation:* Saliency was estimated using gradient-based sensitivity analysis [11]. More specifically, we computed the gradients of the class probabilities with respect to the input for determining each pixel's relevance with respect to the classification decision. Subsequently, the subject level maps were filtered using a three-dimensional Gaussian filter (with a standard deviation of 2 with truncation at 1.75 standard deviations, equivalent to a smoothing kernel size of 9x9x9) and scaled to a standard [0-1] range for further group-level analysis.

E. Standard Machine Learning Models and Pipelines

Five SML methods, including logistic regression (LR) and support vector machine method with a linear (SVML), polynomial (SVMP), radial-basis function (SVMR), and sigmoidal (SVMS) kernel, were used to estimate a rational baseline to compare the performance of DL models. As feature extraction is vital to boost the performance of SML methods, we explored three dimensionality reduction methods: univariate feature selection (UFS), recursive feature elimination (RFE), and Gaussian random projection (GRP) with hyperparameters tuned in a similar range as in our previous work [7].

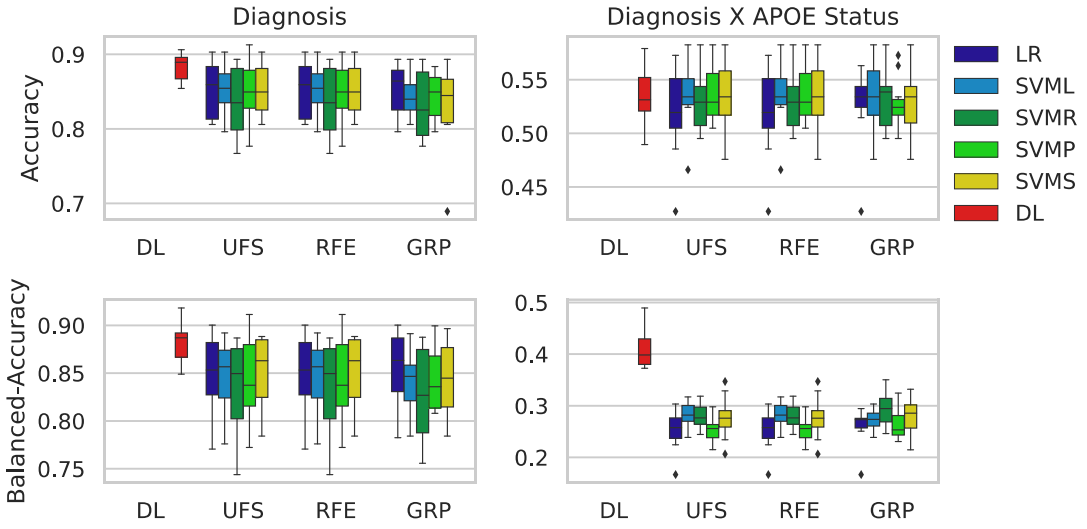

Figure 2: This figure shows the classification performance (accuracy and balanced accuracy metrics) of our deep learning (DL) model and five standard machine learning (SML) models for the (A) 2-way diagnostic classification (left panels) and (B) 6-way diagnosis and APOE status classification tasks (right panels), respectively. The SML methods included logistic regression (LR), and support vector machine methods with a linear (SVML), polynomial (SVMP), radial-basis function (SVMR), and sigmoidal (SVMS) kernel, and the dimensionality reduction methods included univariate feature selection (UFS), recursive feature elimination (RFE), and gaussian random projection (GRP).

III. Results

A. DL Classification Performance

We observed a median classification accuracy and balanced accuracy of 89% for the diagnostic classification (2-way: CN vs. AD) task using our DL model (left top and bottom panels in Figure 2), which is very similar to our previous AD work [6] and a significant improvement (two-tailed paired sample t-test; $p < 0.05$) over the best-performing SML method (86% for LR method on GRP reduced features). Additionally, the distribution of the DL measures has a significantly lower spread than the SML measures for this classification task. This observation suggests that our DL model, when trained on sMRI data, shows improved accuracy to reliably track and visualize neurodegeneration-associated changes in the brain cross-sectionally. However, for the APOE-ε4 status classification task (3-way: APOE4=0 vs. APOE4=1 vs. APOE4=2), accuracy levels were at chance, suggesting that APOE-ε4 status may not be reflected in gray matter maps.

To ascertain the absence of a direct and convincing modulative effect of APOE-ε4 on whole-brain gray matter, we also trained our neural network for prediction of APOE-ε4 status when the two tasks are trained conjointly (6-way: CN-0 vs. CN-1 vs. CN-2 vs. AD-0 vs. AD-1 vs. AD-2). We observed a median accuracy and balanced accuracy of 53% and 40% for this conjoint task, respectively, although the chance prediction level was approximately 17% (right top and bottom panels in Figure 2). For this task, the accuracy of SML methods was very similar to our DL model, but the balanced accuracy dropped significantly, as seen in the bottom right panel of Figure 2.

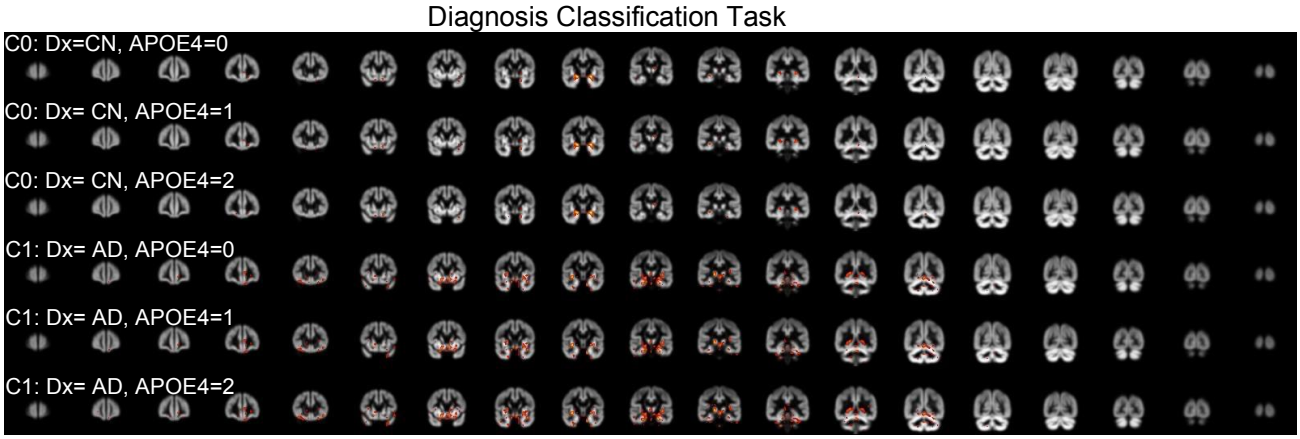

Figure 3: Aggregate (group-level) saliency maps were computed to introspect the MRI correlates of APOE status for the diagnostic classification task.

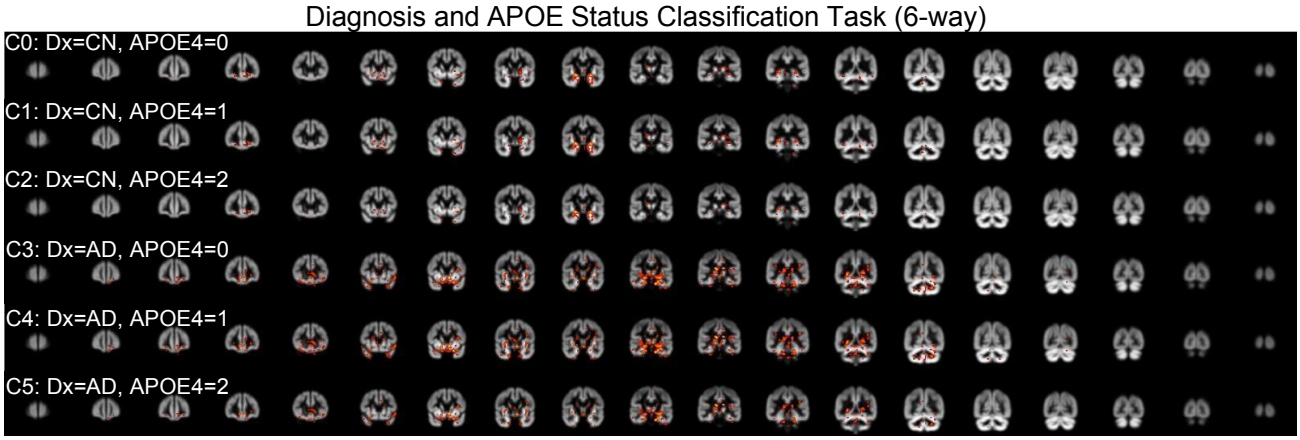

Figure 4: Aggregate (group-level) saliency maps were computed to introspect the MRI correlates of APOE status for the conjoined diagnosis and APOE status classification task.

Nonetheless, since this is a pooled task, it is reasonable to assume that the performance was driven by the easier of the two tasks, which is the diagnosis here. We further interrogate this by exploring the saliency distributions in the different brain regions for these tasks as the next sub-section covers.

B. Brain saliency mapping

Now, if our neural network is indeed learning APOE-ε4 status, we expect those differences to appear on the saliency maps. However, as revealed by Figure 4, the saliences are very similar within the diagnostic groups and like the patterns indicated in the diagnostic classification task (Figure 3).

Notably, the highest saliency values localized in distinct temporal lobe regions for both groups in both tasks align with previous literature. Overall, the saliency results confirm that our deep learning multi-class task classifier primarily recorded diagnostic information rather than learning the APOE-ε4 status.

IV. Discussion

This study sought evidence for APOE-ε4 associated morphological variations detectable directly at the whole-brain voxel level. Results suggest that the sMRI data, when probed with DL, can accurately and reliably predict neurodegenerative diagnostic categories. However, the APOE-ε4 status does not appear to be uniquely reflected in the gray matter. Our observations thus suggest that gray matter changes may be independent of APOE-ε4 status and may characterize a non-APOE, comorbid mechanism in AD.

Although limited by small sample size, this study is distinctly indicative of the trends in cross-examining APOE-ε4 status directly from gray matter. Further exploratory work could involve pooling the available AD data from other studies and additionally exploring other phenotypes and assessments relevant to AD. Crucially, an extension of the undertaken approach to large-scale univariate or multivariate zygosity testing in single nucleotide polymorphisms, for example, through single-task or multi-task learning, could aid in reliable identification of individuals at risk of AD.

Acknowledgments

This study was supported by National Institutes of Health (NIH) grants RF1AG063153, R01EB006841, and R01EB020407 and National Science Foundation grant 2112455 to Dr. Vince D. Calhoun.

References

[1] D. Shen, G. Wu, and H.-I. Suk, "Deep Learning in Medical Image Analysis,", *Annu Rev Biomed Eng,* vol. 19, pp. 221-248, 2017.

[2] G. Litjens *et al.*, "A survey on deep learning in medical image analysis," *Medical image analysis,* vol. 42, pp. 60-88, 2017.

[3] G. Hinton, "Deep learning—a technology with the potential to transform health care," *Jama,* vol. 320, no. 11, pp. 1101-1102, 2018.

[4] H. Greenspan, B. Van Ginneken, and R. M. Summers, "Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique," *IEEE transactions on medical imaging,* vol. 35, no. 5, pp. 1153-1159, 2016.

[5] S. Vieira, W. H. Pinaya, and A. Mechelli, "Using deep learning to investigate the neuroimaging correlates of psychiatric and neurological disorders: Methods and applications," *Neuroscience & Biobehavioral Reviews,* vol. 74, pp. 58-75, 2017.

[6] A. Abrol, M. Bhattarai, A. Fedorov, Y. Du, S. Plis, and V. Calhoun, "Deep residual learning for neuroimaging: An application to predict progression to Alzheimer's disease," *Journal of Neuroscience Methods,* vol. 339, p. 108701, 2020/06/01/ 2020.

[7] A. Abrol *et al.*, "Deep learning encodes robust discriminative neuroimaging representations to outperform standard machine learning," *Nature Communications,* vol. 12, no. 1, p. 353, 2021/01/13 2021, doi: 10.1038/s41467-020-20655-6.

[8] J. H. Cole *et al.*, "Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker," *NeuroImage,* vol. 163, pp. 115-124, 2017.

[9] K. Oh, Y.-C. Chung, K. W. Kim, W.-S. Kim, and I.-S. Oh, "Classification and visualization of Alzheimer's disease using volumetric convolutional neural network and transfer learning," *Scientific Reports,* vol. 9, no. 1, pp. 1-16, 2019.

[10] A. Abrol, R. Hassanzadeh, S. Plis, and V. Calhoun, "Deep learning in resting-state fMRI," in *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 1-5 Nov. 2021 2021, pp. 3965-3969.

[11] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps," *CoRR,* vol. abs/1312.6034, 2014.