

EXPLAINABLE BRAIN AGE GAP PREDICTION IN NEURODEGENERATIVE CONDITIONS USING COVARIANCE NEURAL NETWORKS

Saurabh Sihag[†]

Gonzalo Mateos^{*}

Alejandro Ribeiro[‡]

[†] University of Albany, SUNY

^{*} University of Rochester

[‡] University of Pennsylvania

ABSTRACT

Brain age is the estimate of biological age derived from neuroimaging datasets using machine learning algorithms. Increasing *brain age gap* characterized by an elevated brain age relative to the chronological age can reflect increased vulnerability to neurodegeneration and cognitive decline. Hence, brain age gap is a promising biomarker for monitoring brain health. However, black-box machine learning approaches to brain age gap prediction have limited practical utility. Recent studies on coVariance neural networks (VNN) have proposed a relatively transparent deep learning pipeline for neuroimaging data analyses, which possesses two key features: (i) inherent *anatomically interpretability* of derived biomarkers; and (ii) a methodologically interpretable perspective based on *linkage with eigenvectors of anatomic covariance matrix*. In this paper, we apply the VNN-based approach to study brain age gap using cortical thickness features for various prevalent neurodegenerative conditions. Our results reveal distinct anatomic patterns for brain age gap in Alzheimer's disease, frontotemporal dementia, and atypical Parkinsonian disorders. Furthermore, we demonstrate that the distinct anatomic patterns of brain age gap are linked with the differences in how VNN leverages the eigenspectrum of the anatomic covariance matrix, thus lending explainability to the reported results.

1. INTRODUCTION

Aging is a complicated biological process that manifests itself in the form of various progressive physiological and cognitive changes [1]. Recent years have seen an exponential increase in the study of brain aging using machine learning algorithms [2]. A common objective of brain age prediction strategies is to derive an estimate of brain age from neuroimaging data and compare it with chronological age (time since birth) via the brain age gap, i.e., the difference between brain age and chronological age. A large age gap in the brain has been shown to be indicative of accelerated aging in various neurological diseases, implying a greater burden of the disease and the risk of mortality [3]. Subsequently, we use the notation Δ -Age to refer to brain age gap.

The methodologies for inferring Δ -Age in prior works follow the common template of (i) training a regression model to predict chronological age for a healthy population; and then (ii) application to a cohort of interest (often characterized by neurodegenerative condition), under the hypothesis that the trained model can detect 'accelerated aging' [4,5]. Existing studies in this domain have focused exclusively on (i) and there exist numerous machine learning models that have achieved impressive accuracy in predicting chronological age across heterogeneous populations of healthy individuals [6–8]. However, it is noteworthy that the residuals of the regression models that inform the Δ -Age estimates are of primary interest in this

application, as the expectation is that they will drift in a specific direction when deployed to predict chronological age for individuals with adverse health conditions. Empirical evidence from several existing studies hints at decoupling the assessment of Δ -Age from the prediction accuracy of chronological age on healthy population, as better performance in this task may not necessarily lead to a more informative Δ -Age prediction [6, 9, 10]. In this context, we build upon the recently proposed *explanation-driven* pipeline for Δ -Age prediction [10] based on VNN models [11] to provide an inherently transparent and *anatomically interpretable* perspective to evaluation of Δ -Age.

1.1. coVariance Neural Networks

VNNs have recently been studied as graph neural networks (GNNs) operating on the sample covariance matrix as the graph [10–15]. VNNs belong to the convolutional family of deep learning models as they rely on *linear-shift-and-sum* operators [16] on the covariance matrix. Specifically, given a sample covariance matrix $\mathbf{C} \in \mathbb{R}^{m \times m}$, the convolution operation in a VNN is modeled by a coVariance filter, given by $\mathbf{H}(\mathbf{C}) \triangleq \sum_{k=0}^K h_k \mathbf{C}^k$, where scalar parameters $\{h_k\}_{k=0}^K$ are referred to as filter taps that are learned from the data. The application of coVariance filter $\mathbf{H}(\mathbf{C})$ on an input \mathbf{x} translates to combining information across different sized neighborhoods. *VNNs have been studied separately from GNNs* because they draw a fundamental equivalence with principal component analysis (PCA) when \mathbf{C} is the sample covariance matrix [11, 13, 14]. Advantages of VNNs over traditional PCA-based approaches include improved performance, stability [11, 14], and, notably, transferability across multiscale datasets [15]. In this context, VNNs are deep learning models that *bridge the gap* between traditional PCA-based statistical approaches and graph neural networks [10–12, 14, 15].

1.2. VNNs for neuroimaging data analysis

VNNs have been successfully used to extract biomarkers of neurodegeneration in Alzheimer's disease from brain morphometric features [10, 15, 17, 18]. VNNs offer two key features in the context of deriving biomarkers from neuroimaging data: (i) Representations learned by VNNs provide anatomic signatures of biomarkers [10]; and (ii) learning outcomes of VNNs are generalizable to neuroimaging data curated according to other brain atlases under certain regularity conditions [15]. In particular, VNN-based data analysis pipeline has yielded a novel *explanation-driven* perspective to brain age gap evaluation [10], which is distinct from the prevalent machine learning works in this domain that leverage black-box models and focus primarily on the performance of chronological age prediction in healthy individuals for gauging their quality.

1.3. Contributions

In this paper, we leverage the brain age gap prediction pipeline based on VNNs to study Δ -Age from cortical thickness features and its explainability for Alzheimer's disease (AD), frontotemporal degeneration (FTD), as well as corticobasal syndrome and progressive supranuclear palsy (together referred as the cohort of atypical Parkinsonian disorders (APD) in this paper); the data for all cohorts were downloaded from <https://ida.loni.usc.edu>. Our contributions are summarized as follows:

- **Anatomical characterization of Δ -age for different neurodegenerative conditions.** Our results demonstrated significantly elevated Δ -Age for AD, FTD, and APD conditions relative to a healthy population, along with distinct and biologically plausible anatomic patterns associated with Δ -Age for these conditions.
- **Explainability of Δ -Age.** Our experiments demonstrated that a pre-trained VNN model exploited the eigenvectors of the anatomical covariance matrix differently for the aforementioned neurodegenerative conditions, thus rendering explainability to the distinct anatomic patterns associated with Δ -Age.

Previous work on VNN-based Δ -Age prediction has focused primarily on Alzheimer's disease [10, 15, 18]. This paper provides a more comprehensive evaluation than prior works, as it demonstrates the applicability of VNNs for the prediction of Δ -Age for several neurodegenerative conditions and elucidates the interpretability and explainability of Δ -Age across them.

2. COVARIANCE NEURAL NETWORKS

We start by providing a brief overview of the architecture of a VNN model. A single layer of VNN is formed by concatenating a coVariance filter with a pointwise non-linear activation function $\sigma(\cdot)$ (e.g., ReLU, tanh) that satisfies $\sigma(\mathbf{u}) = [\sigma(u_1), \dots, \sigma(u_m)]$ for $\mathbf{u} = [u_1, \dots, u_m]$. Therefore, the output of a single layer VNN with input \mathbf{x} is given by $\mathbf{z} = \sigma(\mathbf{H}(\mathbf{C})\mathbf{x})$ (Fig. 1). The construction of a multi-layer VNN is formalized next.

Remark 1 (Multi-layer VNN). *For an L -layer VNN, denote the coVariance filter in layer ℓ of the VNN by $\mathbf{H}_\ell(\mathbf{C})$ and its corresponding set of filter taps by \mathcal{H}_ℓ . Given a pointwise nonlinear activation function $\sigma(\cdot)$, the relationship between the input $\mathbf{x}_{\ell-1}$ and the output \mathbf{x}_ℓ for the ℓ -th layer is $\mathbf{x}_\ell = \sigma(\mathbf{H}_\ell(\mathbf{C})\mathbf{x}_{\ell-1})$ for $\ell \in \{1, \dots, L\}$, where \mathbf{x}_0 is the input \mathbf{x} .*

Furthermore, similar to other deep learning models, sufficient expressive power can be facilitated in the VNN architecture by incorporating multiple input multiple output (MIMO) processing at every layer. Formally, consider a VNN layer ℓ that can process $F_{\ell-1}$ number of m -dimensional inputs and outputs F_ℓ number of m -dimensional outputs via $F_{\ell-1} \times F_\ell$ number of filter banks [19]. In this scenario, the input is specified as $\mathbf{X}_{\text{in}} = [\mathbf{x}_{\text{in}}[1], \dots, \mathbf{x}_{\text{in}}[F_{\text{in}}]]$, and the output is specified as $\mathbf{X}_{\text{out}} = [\mathbf{x}_{\text{out}}[1], \dots, \mathbf{x}_{\text{out}}[F_{\text{out}}]]$. The relationship between the f -th output $\mathbf{x}_{\text{out}}[f]$ and the input \mathbf{x}_{in} is given by $\mathbf{x}_{\text{out}}[f] = \sigma\left(\sum_{g=1}^{F_{\text{in}}} \mathbf{H}_{fg}(\mathbf{C})\mathbf{x}_{\text{in}}[g]\right)$, where $\mathbf{H}_{fg}(\mathbf{C})$ is the coVariance filter that processes $\mathbf{x}_{\text{in}}[g]$. Without loss of generality, we assume that $F_\ell = F, \forall \ell \in \{1, \dots, L\}$. In this case, the set of all filter taps is given by $\mathcal{H} = \{\mathcal{H}_{fg}^\ell\}, \forall f, g \in \{1, \dots, F\}, \ell \in \{1, \dots, L\}$, where $\mathcal{H}_{fg}^\ell = \{h_{fg}^\ell[k]\}_{k=0}^K$ and $h_{fg}^\ell[k]$ is the k -th filter tap for filter $\mathbf{H}_{fg}^\ell(\mathbf{C})$. Thus, we can compactly represent a multi-layer VNN architecture capable of MIMO processing via the notation $\Phi(\mathbf{x}; \mathbf{C}, \mathcal{H})$, where the set of filter taps \mathcal{H} captures the full span of its architecture.




Fig. 1. One layer of a VNN model. The covariance filter implicitly manipulates input data according to the eigenvectors of the covariance matrix \mathbf{C} , thus tying the output \mathbf{y} to specific eigenvectors of \mathbf{C} .

Remark 2 (Statistical inference using VNNs). *Theorem 1 in [11] established the equivalence between processing data samples with principal component analysis (PCA) transform and processing data samples with a coVariance filter $\mathbf{H}(\mathbf{C})$. Hence, it can be concluded that input data is processed with VNNs, at least in part, by exploiting the eigenvectors of \mathbf{C} . Unlike simpler PCA-based inference models, VNNs offer stability [11] and transferability guarantees [18], which ensure reproducibility of the inference outcomes by VNNs with high confidence.*

3. METHODS AND MATERIALS

3.1. Datasets for neurodegenerative conditions

In this paper, we leverage cortical thickness measures derived from structural MRI and curated according to Desikan-Killiany brain atlas (68 cortical regions) for various neurodegenerative conditions. All datasets are publicly available at <https://ida.loni.usc.edu/>.

ADNI. This dataset comprised of 118 individuals diagnosed with Alzheimer's disease dementia (AD; age = 73.84 ± 7.56 years, 56 females) and 206 healthy individuals (HC; age = 73.87 ± 6.39 years, 110 females). The 68 cortical thickness features across the cortex for all individuals were downloaded from <https://adni.loni.usc.edu/>. The cortical thickness features had been derived from T1w MRI images using Freesurfer 5.1. This dataset was collected as part of the Alzheimer's disease neuroimaging initiative (ADNI) [20].

NIFD. This dataset spans 119 individuals diagnosed with different forms of FTD (FTD; age = 64.72 ± 6.78 years, 47 females). In the FTD group, 52 individuals were diagnosed with behavioral variant FTD (BV; age = 63.07 ± 5.77 years, 16 females); 36 individuals were diagnosed with semantic variant of primary progressive aphasia (SV; age = 63.68 ± 6.17 years, 15 females); and 31 individuals with non-fluent variant of primary progressive aphasia (PNFA; age = 68.7 ± 7.61 years, 16 females). Additionally, this dataset consisted of 114 healthy individuals (HC; age = 64.51 ± 6.5 years, 65 females). The MRI images were collected as part of the frontotemporal lobar degeneration neuroimaging initiative (FTLDNI). Cortical thickness features (curated according to Desikan-Killiany atlas) were derived using the open-access CAT12 pipeline [21] using their default options. We refer the reader to <https://neuro-jena.github.io/cat12-help/> for detailed processing steps. All outputs were quality checked visually for errors in grey matter segmentation.

4RTNI. This dataset was collected as part of the 4-Repeat Tauopathy Neuroimaging Initiative (4RTNI) and used similar MRI acquisition and clinical assessments as the NIFD dataset. This dataset constituted of 59 individuals diagnosed with progressive supranuclear palsy (PSP; age = 70.79 ± 7.65 years, 32 females) and 45 individuals diagnosed with corticobasal syndrome (CBS; age = 66.71 ± 6.75 years, 24 females). CBS and PSP disorders are among the most common disorders within the broader family of APD [22]. In this paper,

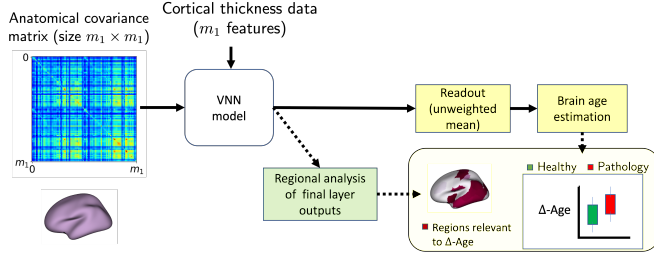


Fig. 2. A VNN-based pipeline for Δ -Age prediction.

we denote the combined cohort of CBS and PSP as APD group. Cortical thickness features curated according to the Desikan-Killiany atlas were derived from T1w MRI using similar pre-processing steps as that for the NIFD dataset. Furthermore, the HC group from the NIFD dataset was considered as the healthy control group for 4RTNI dataset due to the similarity in the MRI acquisition methods of NIFD and 4RTNI datasets.

3.2. Δ -Age prediction

VNN architecture and training. The VNN model was pre-trained on a healthy population to glean information about healthy aging. To facilitate Δ -Age that is transparent and methodologically interpretable, we used a multi-layer VNN model that yielded representations from the input cortical thickness features at the final layer, such that the unweighted mean of these representations formed the estimate for chronological age. Details on how this choice of architecture leads to anatomically interpretable Δ -Age are discussed subsequently. The VNN model consisted of two layers and yielded a representation from the input cortical thickness data via transformation that was dictated by the anatomical covariance matrix. For training this VNN model, we leveraged the cortical thickness features from the healthy control population in the publicly available OASIS-3 dataset [23]. The healthy population in the OASIS-3 dataset consisted of 631 individuals (age = 67.71 ± 8.37 years, 367 females).

For training, we randomly split the available cortical thickness data into a training set of 568 individuals and a test set of 63 individuals. The anatomical covariance matrix was estimated from the cortical thickness data of the training set. The training set was further split into a subset of 498 individuals and a validation set of 70 individuals. The VNN was trained to predict chronological age on the subset of 498 individuals with mean squared error loss optimized using stochastic gradient descent with Adam optimizer for up to 100 epochs. The configuration with the best performance on the validation set of 70 individuals was selected.

The first layer of VNN consisted of 2 filter taps and the second layer consisted of 6 filter taps, with width 61. Thus, in total, VNN model consisted of 22,570 learnable parameters. The batch size used for training was 10 and the learning rate was 0.15. The hyperparameters for the VNN architecture and training were decided during a hyperoptimization procedure based on Optuna [24]. Using this strategy, we trained 10 distinct VNN models with different permutations of the training set. These models achieved a prediction performance of 7.25 ± 0.51 years on the test set and 6.33 years on the complete dataset, with a Pearson's correlation of 0.44 ± 0.014 . Thus, the statistical evidence suggested that VNNs learned information about healthy aging, even though they were weak predictors of chronological age. The results reported in this paper are derived

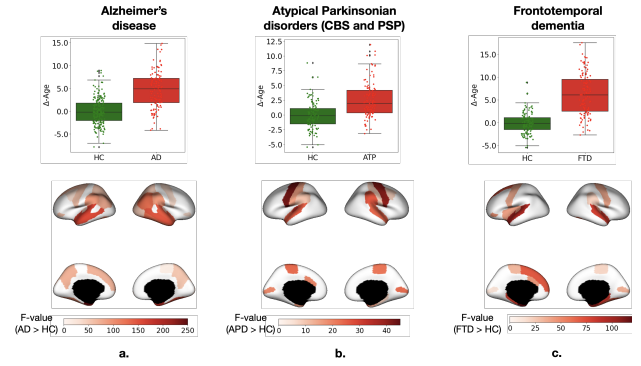


Fig. 3. Brain age estimates and associated anatomic characterizations for (a) AD, (b) APD, and (c) FTD.

from one pre-trained VNN model among the 10 that were pre-trained using the above procedure.

Anatomically interpretable and explainable Δ -Age prediction. The pre-trained VNN model facilitated the prediction of Δ -Age and associated anatomic interpretability and explainability for the cohorts associated with neurodegenerative conditions in Section 3.1. Figure 2 provides an overview of VNN-based pipeline for Δ -Age prediction.

- *Evaluating Δ -Age.* Δ -Age is evaluated as the difference between brain age and chronological age. Brain age was evaluated from the chronological age estimates formed by the VNN with application of a standard linear regression-based approach [25]; the weights of linear regression model learned from the HC populations in the respective datasets in Section 3.1.
- *Assigning anatomical interpretability to Δ -Age.* Elevated Δ -Age in disease groups could be attributed to the statistical patterns in the representations formed by the pre-trained VNN in its final layer. Since the convolution operations in VNN preserve the original dimensionality of the input data, the representations at the final layer could be mapped to individual brain regions. Hence, we evaluated a set of regional residuals r_i for each brain region i , which were defined as the difference between the chronological estimate and the output in the final layer of the VNN corresponding to that brain region. The elevations in the regional residuals directly contribute to elevated Δ -Age [10, Section 3.3].
- *Explainability of Δ -Age.* The representations are learned by the VNN, *in part*, by transforming the input data according to the eigenspectrum of the anatomical covariance matrix [11]. By leveraging this fact, we characterize the explainability of Δ -Age by evaluating the inner products between the regional residuals derived from representations learned by the VNN and the eigenvectors of the anatomical covariance matrix. We anticipate to observe significant differences in terms of these inner product metrics for disease groups and healthy populations. These experiments will elucidate how VNN processed the cortical thickness information from disease groups differently relative to the healthy population, thus lending *explainability* to the evaluation of the downstream statistic of Δ -Age in different cohorts.

4. RESULTS

For each disease dataset, we used the anatomical covariance matrix estimated only from the respective HC group in the pre-trained VNN

model. VNN was oblivious to the identity or any information about the disease. Figure 3 illustrates that elevated Δ -Age was observed for AD, ATP, and FTD relative to their respective HC groups. The **Δ -Age for AD group** was 4.67 ± 4.04 years, which was elevated relative to the HC group in this dataset (Δ -Age for HC: 0 ± 2.91 years). Figure 3a also illustrates the anatomic characterization of elevated Δ -Age in AD group. The anatomic characterization was derived by plotting the F -values for ANOVA between individual elements of the representations generated by the VNN model for AD and HC groups on the brain surface (only for group differences that had AD > HC and survived Bonferroni correction for multiple comparisons with p -value < 0.05). The directionality AD > HC is relevant to Δ -Age as the rise in the individual elements of the representations learned by VNN for a disease group contributed to the elevated Δ -Age estimate relative to healthy population [10]. The anatomic characterization for Δ -Age in AD spanned bilateral regions in the medial temporal lobe, entorhinal, and temporo-parietal junction, which are relevant regions for AD pathology [26]. Consistent results were observed in the prior work on VNN-based Δ -Age prediction [10].

Similar analysis procedures were followed for 4RTNI, NIFD, and PPMI datasets, and the associated results for APD and FTD disease groups are illustrated in Fig. 3b and Fig. 3c, respectively. The HC group for 4RTNI and NIFD datasets was the same and had Δ -Age = 0 ± 2.33 years. The Δ -Age for FTD disease group was 6.17 ± 4.55 years and that for the APD group was 2.49 ± 3.09 years, suggesting that FTD group exhibited more significant accelerated aging. The **Δ -Age in FTD group** was characterized by superior frontal region in the left hemisphere, bilateral entorhinal, parahippocampal, and superior temporal regions, and relatively less prominent contributions from bilateral precentral regions, which are part of the motor cortex. FTD is characterized by neurodegeneration in the regions in frontal and temporal lobes [27], and hence, the anatomic characterization in Fig. 3c captures their impact in the form of elevated Δ -Age. The **Δ -Age in APD group** was characterized by brain regions comprising bilateral superior temporal, precentral, and occipital lobes. The CBS and PSP pathologies in APD group exhibit symptoms similar to PD in terms of motor dysfunction, but are also characterized by cognitive dysfunction and typically rapid decline in function relative to PD [28]. Hence, the implication of regions in the motor cortex and occipital lobe in Fig. 3b is relevant to the disease characteristic in APD.

Explainability of Δ -Age. Next, we analyzed the inner products between the regional residuals derived from the representations learned by VNNs and the eigenvectors of the anatomical covariance matrix. The eigenvectors of the anatomical covariance matrix were organized from 0 to 67, with the eigenvector 0 associated with the largest eigenvalue and the 67-th eigenvector associated with the smallest eigenvalue. The inner product metrics were significantly different (ANOVA, p -value < 0.0001) between the AD group and HC groups for the eigenvectors 0, 1, 2, and 6 of the anatomical covariance matrix (Fig. 4a). Thus, the VNN model processed the cortical thickness features for AD group significantly different relative to the HC group leading to distinct distributions in Δ -Age in Fig. 3a, and these differences were dictated by variations in how the VNN exploited the eigenvectors of the anatomical covariance matrix for AD and HC groups.

Similar results were obtained for the FTD cohort, where the most significant group differences in the inner product measures were observed for eigenvectors 0, 1, 4, and 5 (Fig. 4b). Thus, the variations in how the VNN exploited the eigenvectors of the anatomical covariance matrix for FTD and HC groups *explained* the variations in Δ -Age in Fig. 3c. Notably, the inner product measures

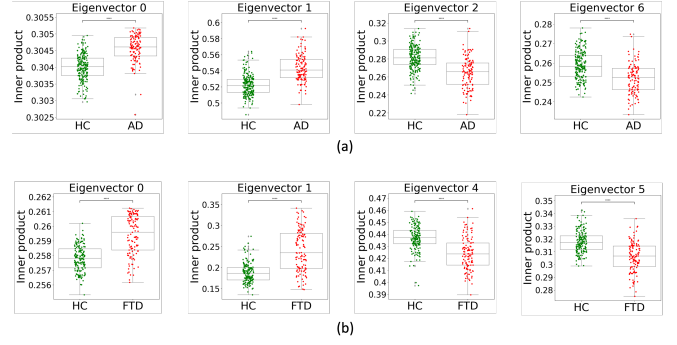


Fig. 4. Explaining Δ -Age in disease groups ((a) AD and (b) FTD) in terms of the group differences between inner products of VNN representations and eigenvectors of anatomical covariance matrices. (**** : p -value $\leq 1 \exp -4$)

were significantly different between the APD and HC groups only for eigenvector 8, which explains the relatively smaller elevation in Δ -Age in the APD group relative to FTD or AD groups in Fig. 3.

5. DISCUSSION

The results in Fig. 3 illustrate that the distributions of Δ -Age estimates can be substantially overlapping across different diseases. Hence, Δ -Age, by itself, is not a sufficient indicator to characterize a disease. In this context, the anatomic characterization of Δ -Age offered by VNN embellishes its informative aspect about neurodegeneration. Notably, the anatomic characterizations of Δ -Age for AD, FTD, and ATP disease groups were unique and characteristic of the respective diseases.

Comparison with existing literature. Existing studies in this domain are focused primarily on brain age prediction, which is to be contrasted with this paper's focus on Δ -Age prediction. Moreover, existing studies utilize the state-of-the-art post-hoc, model-agnostic methods, such as, SHAP, LIME [29], saliency maps [30], and layer-wise relevance propagation [31] to *explain* the brain age predictions. These methods add anatomical interpretability to brain age estimates by assigning some importance to the input features (often associated with specific anatomic regions). Unlike these approaches, we leverage the *inherent explainability* of the VNN model and our results bring into focus the *properties that a VNN gains when it is exposed to the information provided by chronological age of healthy controls and whether and how these properties translate to a meaningful Δ -Age estimate.*

Future work. Future work could entail clinical validation of reported Δ -Age measures by exploring their associations with various clinical markers of a disease. Moreover, longitudinal analyses could provide insights into how Δ -Age progresses with disease burden.

6. ACKNOWLEDGEMENT

Data used in the preparation of this manuscript were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database, and the Frontotemporal Lobar Degeneration Neuroimaging Initiative (FTLDNI) database (all available at <https://ida.loni.usc.edu>). Data collection and sharing for the Alzheimer's Disease Neuroimaging Initiative (ADNI) is funded by the National Institute on Aging (National Institutes of Health Grant U19AG024904). FTLDNI was

funded through the National Institute of Aging, and started in 2010.

7. REFERENCES

- [1] Carlos López-Otín et al., “The hallmarks of aging,” *Cell*, vol. 153, no. 6, pp. 1194–1217, 2013.
- [2] Lea Baecker, Rafael Garcia-Dias, Sandra Vieira, Cristina Scarpazza, and Andrea Mechelli, “Machine learning for brain age prediction: Introduction to methods and clinical applications,” *EBioMedicine*, vol. 72, pp. 103600, 2021.
- [3] Mariona Jové, Manuel Portero-Otín, Alba Naudí, Isidre Ferrer, and Reinald Pamplona, “Metabolomics of human brain aging and age-related neurodegenerative diseases,” *Journal of Neuropathology & Experimental Neurology*, vol. 73, no. 7, pp. 640–657, 2014.
- [4] James H Cole and Katja Franke, “Predicting age using neuroimaging: Innovative brain ageing biomarkers,” *Trends in Neurosciences*, vol. 40, no. 12, pp. 681–690, 2017.
- [5] Katja Franke and Christian Gaser, “Ten years of brainage as a neuroimaging biomarker of brain aging: What insights have we gained?,” *Frontiers in neurology*, p. 789, 2019.
- [6] Vishnu M Bashyam, Guray Erus, Jimit Doshi, Mohamad Habes, Ilya M Nasrallah, Monica Truelove-Hill, Dhivya Srinivasan, Liz Mamourian, Raymond Pomponio, Yong Fan, et al., “MRI signatures of brain age and disease over the lifespan based on a deep brain network and 14 468 individuals worldwide,” *Brain*, vol. 143, no. 7, pp. 2312–2324, 2020.
- [7] Han Peng, Weikang Gong, Christian F Beckmann, Andrea Vedaldi, and Stephen M Smith, “Accurate brain age prediction with lightweight deep neural networks,” *Medical image analysis*, vol. 68, pp. 101871, 2021.
- [8] Xin Niu, Fengqing Zhang, John Kounios, and Hualou Liang, “Improved prediction of brain age using multimodal neuroimaging data,” *Human brain mapping*, vol. 41, no. 6, pp. 1626–1643, 2020.
- [9] Robert J Jirsaraie, Aaron J Gorelik, Martins M Gatavins, Denis A Engemann, Ryan Bogdan, Deanna M Barch, and Aristeidis Sotiras, “A systematic review of multimodal brain age studies: Uncovering a divergence between model accuracy and utility,” *Patterns*, vol. 4, no. 4, 2023.
- [10] Saurabh Sihag, Gonzalo Mateos, Corey McMillan, and Alejandro Ribeiro, “Explainable brain age prediction using covariance neural networks,” in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [11] Saurabh Sihag, Gonzalo Mateos, Corey McMillan, and Alejandro Ribeiro, “coVariance neural networks,” in *Proc. Conference on Neural Information Processing Systems*, Nov. 2022.
- [12] Andrea Cavallo, Madeline Navarro, Santiago Segarra, and Elvin Isufi, “Fair covariance neural networks,” *arXiv preprint arXiv:2409.08558*, 2024.
- [13] Andrea Cavallo, Zhan Gao, and Elvin Isufi, “Sparse covariance neural networks,” *arXiv:2410.01669*, vol. cs.LG, 2024.
- [14] Andrea Cavallo, Mohammad Sabbaqi, and Elvin Isufi, “Spatiotemporal covariance neural networks,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2024, pp. 18–34.
- [15] Saurabh Sihag, Gonzalo Mateos, Corey McMillan, and Alejandro Ribeiro, “Transferability of covariance neural networks,” *IEEE Journal of Selected Topics in Signal Processing*, pp. 1–16, 2024.
- [16] Antonio Ortega, Pascal Frossard, Jelena Kovačević, José MF Moura, and Pierre Vandergheynst, “Graph signal processing: Overview, challenges, and applications,” *Proceedings of the IEEE*, vol. 106, no. 5, pp. 808–828, 2018.
- [17] Saurabh Sihag, Gonzalo Mateos, and Alejandro Ribeiro, “Towards a foundation model for brain age prediction using covariance neural networks,” *arXiv preprint arXiv:2402.07684*, 2024.
- [18] Saurabh Sihag, Gonzalo Mateos, Corey McMillan, and Alejandro Ribeiro, “Predicting brain age using transferable coVariance neural networks,” in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing*, Jun. 2023.
- [19] Fernando Gama, Joan Bruna, and Alejandro Ribeiro, “Stability properties of graph neural networks,” *IEEE Transactions on Signal Processing*, vol. 68, pp. 5680–5695, 2020.
- [20] Clifford R Jack Jr, Matt A Bernstein, Nick C Fox, Paul Thompson, Gene Alexander, Danielle Harvey, Bret Borowski, Paula J Britson, Jennifer L. Whitwell, Chadwick Ward, et al., “The alzheimer’s disease neuroimaging initiative (adni): Mri methods,” *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 27, no. 4, pp. 685–691, 2008.
- [21] Christian Gaser, Robert Dahnke, Paul M Thompson, Florian Kurth, Eileen Luders, and Alzheimer’s Disease Neuroimaging Initiative, “CAT—a computational anatomy toolbox for the analysis of structural mri data,” *bioRxiv*, pp. 2022–06, 2022.
- [22] Brent L Fogel, Mary C Clark, and Daniel H Geschwind, “The neurogenetics of atypical parkinsonian disorders,” in *Seminars in neurology*. Thieme Medical Publishers, 2014, vol. 34, pp. 217–224.
- [23] Pamela J LaMontagne, Tammie LS Benzinger, John C Morris, Sarah Keefe, Russ Hornbeck, Chengjie Xiong, Elizabeth Grant, Jason Hassenstab, Krista Moulder, Andrei G Vlassenko, et al., “OASIS-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and Alzheimer disease,” *MedRxiv*, 2019.
- [24] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama, “Optuna: A next-generation hyperparameter optimization framework,” in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 2623–2631.
- [25] Iman Beheshti, Scott Nugent, Olivier Potvin, and Simon Duchesne, “Bias-adjustment in neuroimaging-based brain age frameworks: A robust scheme,” *NeuroImage: Clinical*, vol. 24, pp. 102063, 2019.
- [26] Rahul S Desikan, Howard J Cabral, Bruce Fischl, Charles RG Guttmann, Deborah Blacker, Bradley T Hyman, Marilyn S Albert, and Ronald J Killiany, “Temporoparietal mr imaging measures of atrophy in subjects with mild cognitive impairment that predict subsequent diagnosis of alzheimer disease,” *American Journal of Neuroradiology*, vol. 30, no. 3, pp. 532–538, 2009.
- [27] Jonathan D Rohrer, “Structural brain imaging in frontotemporal dementia,” *Biochimica et Biophysica Acta (BBA)-Molecular Basis of Disease*, vol. 1822, no. 3, pp. 325–332, 2012.
- [28] Nikolaus R McFarland, “Diagnostic approach to atypical parkinsonian syndromes,” *CONTINUUM: Lifelong Learning in Neurology*, vol. 22, no. 4, pp. 1117–1142, 2016.
- [29] Angela Lombardi, Domenico Diacono, Nicola Amoroso, Alfonso Monaco, João Manuel RS Tavares, Roberto Bellotti, and Sabina Tangaro, “Explainable deep learning for personalized age prediction with brain morphology,” *Frontiers in neuroscience*, vol. 15, pp. 578, 2021.
- [30] Chenzhong Yin, Phoebe Imms, Mingxi Cheng, Anar Amgalan, Nahian F Chowdhury, Roy J Massett, Nikhil N Chaudhari, Xinghe Chen, Paul M Thompson, Paul Bogdan, et al., “Anatomically interpretable deep learning of brain age captures domain-specific cognitive impairment,” *Proceedings of the National Academy of Sciences*, vol. 120, no. 2, pp. e2214634120, 2023.
- [31] Simon M Hofmann, Frauke Beyer, Sebastian Lapuschkin, Ole Goltermann, Markus Loeffler, Klaus-Robert Müller, Arno Villringer, Wojciech Samek, and A Veronica Witte, “Towards the interpretability of deep learning models for multi-modal neuroimaging: Finding structural changes of the ageing brain,” *NeuroImage*, vol. 261, pp. 119504, 2022.