

BrainRotViT: Transformer-ResNet Hybrid for Explainable Modeling of Brain Aging from 3D sMRI

Wasif Jalal¹, Md Nafiu Rahman¹, Atif Hasan Rahman^{1,2}, and M. Sohel Rahman^{1,2}

¹Department of Computer Science and Engineering, Bangladesh University of Engineering and Technology

²Corresponding authors: atif@cse.buet.ac.bd, msrahman@cse.buet.ac.bd

Abstract

Accurate brain age estimation from structural MRI is a valuable biomarker for studying aging and neurodegeneration. Traditional regression and CNN-based methods face limitations such as manual feature engineering, limited receptive fields, and overfitting on heterogeneous data. Pure transformer models, while effective, require large datasets and high computational cost. We propose **Brain ResNet over trained Vision Transformer (BrainRotViT)**, a hybrid architecture that combines the global context modeling of vision transformers (ViT) with the local refinement of residual CNNs. A ViT encoder is first trained on an auxiliary age and sex classification task to learn slice-level features. The frozen encoder is then applied to all sagittal slices to generate a 2D matrix of embedding vectors, which is fed into a residual CNN regressor that incorporates subject sex at the final fully-connected layer to estimate continuous brain age. Our method achieves an MAE of 3.34 years (Pearson $r = 0.98$, Spearman $\rho = 0.97$, $R^2 = 0.95$) on validation across 11 MRI datasets encompassing more than 130 acquisition sites, outperforming baseline and state-of-the-art models. It also generalizes well across 4 independent cohorts with MAEs between 3.77 and 5.04 years. Analyses on the brain age gap (BAG), i.e. the difference between the predicted age and actual age, show that aging patterns are associated with Alzheimer’s disease, cognitive impairment, and autism spectrum disorder. Model attention maps highlight aging-associated regions of the brain, notably the cerebellar vermis, precentral and postcentral gyri, temporal lobes, and medial superior frontal gyrus. Our results demonstrate that this method provides an efficient, interpretable, and generalizable framework for brain-age prediction, bridging the gap between CNN- and transformer-based approaches while opening new avenues for aging and neurodegeneration research.

Key words— Deep learning, Biomedical image analysis, Brain aging, MRI, Vision transformer, ResNet, Convolutional neural network, Hybrid model, Neurodegeneration

1 Introduction

The human brain undergoes continuous transformations across the lifespan, representing a natural component of aging that does not inherently signal pathological conditions [1]. Neurodegenerative disorders such as dementia can compromise the brain structure and accelerate aging processes. Understanding and characterizing healthy brain aging patterns therefore becomes essential for distinguishing normal aging from pathological neurodegeneration, potentially enabling earlier detection of neurodegenerative diseases. The Brain Age-Gap (BAG), i.e. the discrepancy between predicted brain age and chronological age, has emerged as a robust biomarker that captures pathological brain processes and offers insights into the rate at which an individual’s brain ages in comparison to others in the population [2, 3]. It is not only associated with various neurological disorders, such as Alzheimer’s disease, cognitive impairment, and Autism Spectrum

Disorder, but also serves as an indicator of all-cause mortality [4, 5, 6, 7, 8]. Brain age estimation has been approached through both conventional and machine learning techniques, analyzing either the whole brain, specific regions, or localized patches [9, 10, 11]. One particular study presented a method using T1-weighted MRI to predict age through region-level and voxel-level metrics [12]. Regression-based machine learning has shown promise for the brain age prediction, with kernel regression applied to whole-brain MRI across diverse age ranges [13]. Various algorithms including Support Vector Regression and Binary Decision Trees have been compared for their brain age prediction capabilities [14]. Additional regression techniques such as Relevance Vector Regression, Twin Support Vector Regression, and Gaussian Process Regression have been explored across different imaging modalities for age estimation and mortality prediction [11, 15, 16, 17]. However, while these conventional machine learning approaches achieve reasonable performance in brain age estimation, they are

limited by the need for manual feature engineering. Deep learning offers automatic feature extraction capabilities, making it particularly suited for handling unstructured data and addressing complex analytical challenges.

Recent advances have focused on developing deep learning approaches for brain age estimation from T1-weighted structural MRI [18, 19, 20]. Research has demonstrated that traditional methods such as Lasso regression and Support Vector Regression fail to capture complex brain structural relationships, necessitating deep learning models [21]. Various architectures have been explored: Convolutional Neural Networks (CNNs) have been employed on both raw and preprocessed MRI data to capture fine-grained structural patterns of the brain [11]. Extensions to three-dimensional CNN architectures further enhanced spatial context modeling compared to their 2D counterparts [22], while Fully Convolutional Networks (SFCNs) inspired by VGG Net achieved notable success on large-scale datasets [23]. Beyond CNNs, attention-based architectures including graph transformers utilizing regions of interest [24] and global-local transformers [25] have demonstrated exceptional performance through their attention mechanisms.

Nevertheless, despite these developments, significant research gaps persist. Pure transformer-based and global-local transformer models (e.g., [25]) typically fuse global and local information across entire volumes or heavy patch-based pipelines, which increases computational cost, memory footprint, and data requirements, thus hindering scalability across moderate-sized, heterogeneous multi-site MRI cohorts. Graph-transformer and ROI-based methods effectively model inter-regional relationships but often rely on parcellation (i.e. predefined ROIs), potentially discarding voxel- or slice-level signals that are crucial for fine-grained brain age prediction [26, 27]. Recent multi-oriented and hybrid designs (e.g., permuted transformer-convolution hybrids [28] and multi-view ViT fusions, such as Triamese-ViT [29]) improve 3D context modeling but do not implement a pipeline that trains on 3D voxel data of whole-brain images in a domain-agnostic approach across a diverse set of independent multi-site cohorts. To the best of our knowledge, no prior transformer-based or hybrid CNN-ViT study on brain age estimation has simultaneously (a) reduced data demands of end-to-end ViT regression while preserving volumetric continuity in a compact, learnable representation, (b) generalized across as many independent cohorts as our study, (c) enabled efficient regression without reliance on ROI graphs or full 3D attention, (d) provided biologically interpretable explanations consistent with existing literature, and (e) analyzed the association of the brain age-gap (BAG) with neurological disorders.

To address these gaps, we introduce a hybrid architecture (Figure 1) that integrates Vision Transformer (ViT)

and Convolutional Neural Network (CNN) components. Our method leverages the global attention capabilities of the transformer, which enhance prediction accuracy by encoding meaningful features, and combines them with the local feature extraction efficiency of a lightweight CNN, which improves computational efficiency compared to an entirely transformer-based setup. Training across diverse multi-site cohorts further strengthens generalizability and clinical robustness. This balance of accuracy, efficiency, and clinical applicability forms the foundation of our contributions outlined below.

- We propose a novel hybrid modeling framework for 3D brain MRI analysis that combines a Vision Transformer (ViT) encoder for learning slice-level embeddings with a lightweight 2D convolutional regression head for aggregating volumetric information. By applying convolutions over transformer embeddings, our model captures both global contextual information and local spatial patterns, making it the first approach of its kind in 3D medical image analysis to achieve competitive performance (MAE = 3.34 years; Table 2) while remaining computationally efficient (Table 5).
- We adopt a domain-agnostic whole-image learning paradigm trained on a large and diverse collection of brain MRI scans from 11 independent multi-site cohorts. This design ensures robustness to site-specific variations and consistent performance across heterogeneous datasets (Tables 3, 4).
- We provide explainable interpretations of the model’s predictions by identifying brain regions most affected by aging, showing a strong alignment with known neuroanatomical evidences and offering new insights for empirical aging research (Section 3.6).
- We further analyze associations between the brain age-gap (BAG) and neurological disorders, namely, Alzheimer’s disease (AD), mild cognitive impairment (MCI), and autism spectrum disorder (ASD), demonstrating the clinical relevance and translational potential of our approach (Sections 3.4, 3.5).

The code used for experimentation in this study is made available through a GitHub repository at <https://github.com/wjalal/BrainRotViT/>.

2 Methodology

Our proposed framework combines transformer-based and convolutional approaches to achieve accurate and interpretable brain age estimation from 3D MRI scans. It covers the entire process, from data collection and preparation to feature learning, prediction, and interpretation. The workflow starts with gathering multiple MRI datasets and applying a consistent preprocessing




Figure 1: **Overview of the proposed framework for 3D brain MRI analysis.** The model integrates a Vision Transformer (ViT) encoder to extract slice-level embeddings with a lightweight 2D convolutional regression head to aggregate volumetric information. By applying convolutions over transformer-derived embeddings, the framework jointly models global contextual and local spatial patterns, enabling subsequent analyses on brain age gaps and thus linking aging with neurological conditions such as Alzheimer’s disease, mild cognitive impairment, and autism spectrum disorder.

pipeline. Then, we use a Vision Transformer (ViT) to learn informative features from MRI slices and employ a residual CNN module to predict brain age. Finally, we generate attention and saliency maps to understand how the model makes its predictions. The following sections describe each of these stages in detail.

2.1 Dataset

In this study, we utilize fifteen publicly available structural MRI datasets. These datasets, including prominent collections such as ADNI, IXI, Cam-CAN, and ABIDE, provide a diverse pool of subjects across various age ranges, cognitive statuses, and acquisition protocols. This diversity enables robust training and evaluation of our model. Table 1 summarizes the demographic information of the 6,060 total samples used in this study. A brief description of each dataset is provided below.

The Autism Brain Imaging Data Exchange II (ABIDE II) [30] is a multisite, open-data initiative established to advance discovery science on the brain connectome in Autism Spectrum Disorder (ASD). It aggregates 1,114 MRI datasets from 521 individuals with ASD and 593 typically developing controls, spanning ages of 5 to 64 years. In this work, we use a curated

subset of 750 high-quality MRI scans from the ABIDE II dataset.

The Alzheimer’s Disease Neuroimaging Initiative (ADNI) [31] is a longitudinal multicenter study aimed at identifying biomarkers of Alzheimer’s disease progression. The ADNI was launched in 2003 as a public-private partnership, led by Principal Investigator Michael W. Weiner, MD. It provides access to multi-modal neuroimaging data, as well as clinical, genetic, and cognitive assessments. In our study, we utilize 900 T1-weighted MRI scans from this cohort.

The **AgeRisk** [32] dataset provides cross-sectional MRI data from healthy participants aged 16 to 81 years, collected as part of a study on risk preference and impulsivity. The protocol included a T1-weighted structural scan and functional scans. We utilize 189 skull-stripped T1-weighted MRI scans from this study.

The dataset from the **BOLD variability study** by Rieck et al. (2021) [33] includes MRI data from 158 healthy, cognitively normal adults (aged 20 to 86 years), recruited from the greater Toronto area. The study investigated age-related differences in cognitive control

(inhibition, task switching, working memory) using MRI. We incorporate the structural MRI scans from all participants in this study.

The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) [34] [35] project is a comprehensive study aimed at understanding the neural mechanisms underlying healthy cognitive ageing. The dataset includes data from a population-based cohort of adults aged 18-90 years. For this study, we utilize T1-weighted MRI volumes from 653 adults.

The Center for Biomedical Research Excellence (COBRE) dataset [36] comprises anatomical and resting-state functional MRI data from 72 patients with schizophrenia and 75 healthy controls, aged 18–65. In this study, we include 146 samples from COBRE.

The Consortium for Reliability and Reproducibility (CORR) [37] dataset is a large-scale aggregation of resting-state MRI and anatomical data from multiple international sites. This collection is designed to assess and improve the reliability of brain imaging findings. In our study, we include 377 anatomical scans from this consortium, further enhancing the protocol variability and robustness of our model.

The Dallas Lifespan Brain Study (DLBS) [38] is designed to investigate age-related changes in brain structure and function across the adult lifespan (20-89 years). In this study, we use T1-weighted MRI volumes from 315 participants.

The 1000 Functional Connectomes Project (FCON1000) [39] includes data from multiple independent imaging studies, primarily focused on resting-state functional connectivity. In this work, we use MRI volumes from 763 adults pooled from different sites, adding variability valuable for robust machine learning.

The IXI (Information eXtraction from Images) dataset [40] is a multi-center, open-access collection of structural magnetic resonance images from cognitively healthy adults, scanned at three different London hospitals. In this study, we take 565 samples with an age range between 18 and 88 years.

The National Institute of Mental Health (NIMH) Healthy Research Volunteer (RV) dataset [41] is a comprehensive collection characterizing healthy adults with sMRI, fMRI, DTI, and extensive clinical and cognitive data. Participants with a history of mental illness are excluded. We utilize 181 samples from this dataset.

The Open Access Series of Imaging Studies (OASIS) includes cross-sectional (OASIS-1, n=416) MRI dataset [42], covering an age range of 18–96 years.

In this study, we use 300 individuals from the OASIS-1 cross-sectional collection.

The Southwest University Adult Lifespan Dataset (SALD) [43] is a large, cross-sectional, multimodal neuroimaging collection comprising 494 healthy adults aged 19–80. The dataset is designed to map structural and functional changes across adulthood. In this study, we include all 494 participants from SALD to enhance adult lifespan coverage.

The SUDMEX-CONN dataset [44] is a case-control study of cocaine use disorder patients from the National Institute of Psychiatry in Mexico City. The protocol included T1-weighted, 10-minute resting-state fMRI, and HARDI-DWI multishell sequences. We included 136 subjects from this collection to increase the diversity of our training set.

The **TrueCrime** dataset [45] provides structural and resting-state MRI scans for 133 participants (N=133 for MRI) from a larger study at the University of Graz on the psychology of true crime consumption. No tasks were performed during the MRI. We leverage the T1-weighted structural MRI scans from this cohort.

2.2 Data Acquisition and Preprocessing Pipeline

To harmonize heterogeneous MRI datasets, we applied a preprocessing pipeline that standardizes both anatomy and image quality. Each scan underwent skull stripping, bias correction, and affine registration to a common template, followed by cropping, resampling into 160 sagittal slices, and intensity normalization. This process ensured consistent spatial alignment and contrast, producing high-quality inputs for transformer-based feature extraction and CNN regression.

2.2.1 Input Data Specifications

We use T1-weighted structural MRI scans acquired with the three-dimensional Magnetization-Prepared Rapid Gradient-Echo (3D MPRAGE) sequence, which provides high-resolution anatomical images with excellent gray–white matter contrast essential for accurate brain morphometry [46].

2.2.2 Skull Stripping and Brain Extraction

Non-brain tissues (skull, scalp, eyes, etc.) are removed using *DeepBrain*¹, a lightweight, U-Net-based [47] skull stripping tool trained on diverse multi-site data. Unlike traditional intensity- or atlas-based methods, it employs convolutional neural networks to accurately delineate brain boundaries across varied acquisition protocols and field strengths.

¹<https://github.com/iitczco/deepbrain>

Table 1: Demographic information of the datasets used in this study.

Dataset	# Total Samples	# Samples Used	Age (Mean $\pm$ SD)
ADNI	5881	900	63.04 $\pm$ 3.54
IXI	565	565	48.69 $\pm$ 16.46
ABIDE	1114	750	17.99 $\pm$ 9.72
DLBS	315	315	54.62 $\pm$ 20.09
COBRE	147	146	36.98 $\pm$ 12.78
FCON1000	1355	763	27.86 $\pm$ 13.74
SALD	494	494	45.18 $\pm$ 17.44
CORR	3357	377	52.95 $\pm$ 13.22
OASIS-1	416	300	64.17 $\pm$ 19.12
CamCAN	655	653	54.31 $\pm$ 18.58
NIMH	189	181	34.62 $\pm$ 11.52
BOLD variability	158	158	49.28 $\pm$ 19.01
TrueCrime	133	133	26.98 $\pm$ 8.80
SUDMEX-CONN	136	136	30.81 $\pm$ 7.70
AgeRisk	189	189	45.05 $\pm$ 19.27
Total		6060	-

2.2.3 Intensity Inhomogeneity Correction

To correct spatial intensity variations from RF field and coil sensitivity artifacts, we apply the N4 bias field correction algorithm [48]. N4 iteratively estimates a smooth multiplicative bias field via B-spline approximations, yielding spatially consistent tissue intensities while preserving true anatomical contrast.

2.2.4 Spatial Normalization and Alignment

Bias-corrected volumes undergo affine registration to a standard template for anatomical correspondence. This 12-degree-of-freedom transformation (translations, rotations, scaling, shearing) aligns brain size, shape, and orientation across subjects, ensuring consistent spatial locations of structures for subsequent feature extraction.

2.2.5 Volume Resampling and Slice Generation

Aligned brains are tightly cropped to intracranial boundaries to remove background and reduce computation, then resampled into 160 sagittal slices using trilinear interpolation. Sagittal orientation leverages bilateral brain symmetry and captures key structures (corpus callosum, hippocampus, cortex) relevant to age-related changes, thus facilitating model interpretability.

2.2.6 Intensity Normalization

Finally, each slice undergoes z-score normalization (zero mean, unit variance) to reduce global intensity differences across scans and sites. This slice-wise normalization stabilizes input distributions and mitigates site-specific intensity biases, improving model robustness across heterogeneous data sources.

2.3 Stage 1: Vision Transformer-Based Feature Learning

This initial stage makes use of a Vision Transformer (ViT) to learn discriminative features from individual

2D sagittal brain slices. By training the ViT on an auxiliary task of classifying combined age-decade and sex categories, the model is encouraged to extract representations sensitive to anatomical variations related to aging and sex. The output of this stage is a 2D feature map, where each column represents a slice and each row represents a learned feature, effectively creating a compact representation of the entire 3D brain volume.

2.3.1 Classification Task Formulation for Representation Learning

To train the Vision Transformer for discriminative feature extraction, we formulate an auxiliary classification task that encourages the learning of age- and sex-sensitive representations. Rather than directly regressing to continuous age values, which could lead to overfitting, we discretize the age variable into decade-based intervals and combine them with biological sex information. Specifically, chronological ages are grouped into bins spanning ten-year ranges (20-29, 30-39, 40-49, etc.), and each bin is further subdivided by biological sex (male or female). This discretization strategy results in composite classes such as “50-59 years, male” or “30-39 years, female,” creating approximately 16-20 distinct classes depending on the age distribution of the dataset.

The formulation of age-sex composite classes serves multiple purposes. First, it transforms the continuous age prediction problem into a classification task, which typically yields more robust feature representations during initial training. Second, it explicitly encourages the model to learn sex-specific aging patterns, acknowledging the well-documented differences in brain aging trajectories between males and females. Third, the decade-based grouping provides sufficient granularity to capture major developmental and aging milestones while maintaining adequate sample sizes within each class for stable training.




Figure 2: **Implementation details of methodology.** (a): Chronological age distribution of the 4086 training samples used in the study. (b): Chronological age distribution of the 1022 validation-set samples used in the study. (c): Pre-processing routine to convert NIfTI-format 3D T1w MRI images to 2D sagittal slices that are used as inputs to our modeling framework. (d): Heatmap of 160 pixel $\times$ 768 pixel 2D feature map (ViT embedding vectors of 160 slices concatenated in order; log-scaled, min-max scaled, and rotated for better visualization). This is essentially a representation of the input that the residual CNN portion of the network receives for one of the samples. Slices change along the vertical axis, while the horizontal axis is the dimension of the embedding vectors. There is an appearance of vertical symmetry across slices, while adjacent slices appear to have similar embeddings, thus creating the appearance of vertical lines of similarly intense positions in the embedding vectors. (e): Heatmap of 160 $\times$ 160 cosine similarity matrix between each slice of 3D samples; averaged across all samples in the study. The pixel grid from left to right, and from top to bottom, both represent the indices of slices 1 to 160. The principal diagonal (top left to bottom right) represents the cosine similarity of each slice with itself, which is 1, as the most intense red. The intensity spreading from the principal diagonal indicates that the vision transformer produces similar embeddings for neighboring slices, while the intensity along the other diagonal suggests that the embeddings of symmetrically antipodal slices are very similar, i.e., the transformer captures the sagittal symmetry of the brain effectively.

2.3.2 Vision Transformer Architecture and Training Protocol

The Vision Transformer processes input slices through a sophisticated attention-based architecture. Each input slice is first divided into non-overlapping patches of fixed size, 16×16 pixels depending on the input resolution. These patches are then linearly projected into a high-dimensional embedding space of dimension $d = 768$, creating a sequence of patch embeddings. Learnable positional encodings are added to these embeddings to preserve spatial information about the original patch locations within the slice.

The sequence of patch embeddings undergoes processing through multiple transformer encoder layers, each consisting of multi-head self-attention (MSA) and position-wise feed-forward network (FFN) sub-layers:

$$\mathbf{z}^{(l+1)} = \mathbf{z}^{(l)} + \text{MSA}(\text{LayerNorm}(\mathbf{z}^{(l)})) \quad (1)$$

$$\mathbf{z}^{(l+1)} = \mathbf{z}^{(l+1)} + \text{FFN}(\text{LayerNorm}(\mathbf{z}^{(l+1)})) \quad (2)$$

where $\mathbf{z}^{(l)}$ represents the feature representations at layer l , and LayerNorm denotes layer normalization for training stability.

During training, we implement a strategic sampling approach where 32 evenly spaced sagittal slices are selected from each subject's complete set of 160 slices.

Each sampled slice is processed independently through the ViT encoder:

$$f_{\text{ViT}} : S'_i \rightarrow \mathbb{R}^{768} \quad (3)$$

The encoder generates a 768-dimensional feature vector for each input slice, capturing both local anatomical patterns and global contextual information through the self-attention mechanism.

The classification head attached to the ViT consists of a single linear layer followed by softmax activation, producing probability distributions over the predefined age-sex classes. The model is optimized using categorical cross-entropy loss:

$$L_{\text{CE}} = - \sum_{c=1}^C y_c \log(\hat{y}_c) \quad (4)$$

where y_c represents the one-hot encoded ground truth and $\hat{y}_c$ denotes the predicted probability for class c .

2.3.3 Dense Feature Map Generation

Upon training of the Vision Transformer on the classification task, the model transitions to serve as a feature extractor for the downstream regression task. The classification head is discarded, and the encoder weights are

frozen to prevent further updates. This frozen encoder is then applied systematically to all 160 sagittal slices from each brain volume, generating a comprehensive feature representation of the entire brain.

The processing of all slices, rather than just the sampled subset used during training, aims to ensure that no anatomical information is lost in the feature extraction phase. Each slice yields a 768-dimensional embedding vector that encapsulates the learned representations from the ViT training. These embeddings are arranged sequentially according to their spatial position along the sagittal axis, forming a two-dimensional feature map:

$$\mathbf{Z} = [\mathbf{z}_1; \mathbf{z}_2; \dots; \mathbf{z}_{160}]^T \in \mathbb{R}^{160 \times 768} \quad (5)$$

This feature map can be interpreted as a compact, learned representation of the three-dimensional brain volume, where each row corresponds to a sagittal slice and each column represents a learned feature dimension. The spatial arrangement preserves the anatomical continuity along the sagittal axis, enabling the subsequent CNN to exploit spatial relationships between adjacent slices.

2.4 Stage 2: Residual Convolutional Neural Network for Brain Age Regression

In the second stage, the 160×768 feature map generated by the frozen ViT encoder is repurposed as a “pseudo-image.” This novel representation is then processed by a custom residual Convolutional Neural Network (CNN). The CNN is specifically designed to learn spatial relationships between the slice-level features along the sagittal axis. This stage integrates these learned spatial patterns with biological sex information to perform the final regression, predicting a single scalar value for the subject’s brain age.

2.4.1 Feature Map as Pseudo-Image Representation

The feature map generated from the Vision Transformer stage is reconceptualized as a single-channel pseudo-image with dimensions 160×768 . This novel representation treats the rows as spatial positions along the sagittal axis and columns as feature channels, enabling the application of two-dimensional convolutional operations. This approach allows the CNN to learn spatial filters that capture relationships both across adjacent brain slices and between different feature dimensions extracted by the ViT. The convolutional operations can thus identify patterns such as gradual anatomical transitions along the sagittal axis or correlations between specific feature combinations that are indicative of brain age.

2.4.2 Residual Convolutional Block Architecture

The CNN architecture comprises three hierarchical residual convolutional blocks designed to progressively extract

and refine features while maintaining gradient flow via skip connections inspired by *ResNet*. Each block uses a dual-pathway design: the main path applies a convolution, $\text{SiLU}(x) = x \cdot \text{sigmoid}(x)$ activation, and max pooling (kernel size 2, stride 1) to capture features with local translation invariance, while the residual path aligns dimensions using 1×1 convolutions (for channel mismatches) and nearest-neighbor interpolation (for spatial mismatches). Outputs from both paths are merged by element-wise addition to preserve information and support gradient propagation.

ConvBlock 1 uses 8 filters of size (10, 60) to capture broad spatial patterns across slices. ConvBlock 2 refines these into 4 channels with (5, 15) kernels, emphasizing localized features, and ConvBlock 3 compresses them into 1 channel using (2, 6) kernels, producing a compact age-relevant representation.

2.4.3 Fully Connected Layers and Feature Integration

After convolutional feature extraction, the output tensor is flattened into a 1D vector. Since its size depends on the post-pooling spatial dimensions, the first fully connected layer (FC1) is initialized dynamically using a dummy forward pass to determine the exact dimensions.

The flattened features pass through two fully connected layers. FC1 maps them to 512 dimensions, followed by layer normalization (more stable than batch normalization for small batch sizes [49]), SiLU activation, and dropout ($p = 0.3$) for regularization. FC2 reduces this to 128 dimensions, again followed by layer normalization, SiLU, and dropout. This progressive reduction enforces increasingly abstract representations of age-related patterns.

At this stage, biological sex (binary: 0 for male, 1 for female) is concatenated to the 128-dimensional vector, forming a 129-dimensional representation. This late fusion lets earlier layers learn sex-independent features, while the final layer incorporates sex-specific differences in brain aging to improve prediction accuracy.

2.4.4 Final Regression Layer

The final layer of the network consists of a linear transformation that maps the 129-dimensional combined feature vector directly to a single scalar output representing the predicted brain age:

$$\hat{y} = \mathbf{W}_3 \mathbf{h}_{\text{combined}} + b_3 \quad (6)$$

where $\hat{y}$ represents the predicted chronological age in years, $\mathbf{W}_3$ is a learned weight matrix of dimension 1×129 , $\mathbf{h}_{\text{combined}}$ is the concatenated feature vector including sex information, and b_3 is a scalar bias term.

2.4.5 Loss Function

The model is primarily trained using Mean Squared Error (MSE) loss, which penalizes predictions based on the

squared difference between predicted and actual ages. This was adopted as the primary optimization objective due to its superior deterministic regression accuracy on the validation set.

$$L_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^N (a_i - \hat{y}_i)^2 \quad (7)$$

, where a_i denotes the ground-truth chronological age for subject i , $\hat{y}_i$ is the predicted age, and N is the batch size.

For comparison, a Negative Log-Likelihood (NLL) loss [50] was also evaluated, particularly to analyze probabilistic uncertainty modeling and generalization. This approach assumes the model outputs the parameters of a probability distribution (e.g., a Gaussian distribution) rather than a single point estimate. The NLL loss for a Gaussian distribution, where the model predicts both a mean ($\hat{\mu}_i$) and a variance ($\hat{\sigma}_i^2$) for each subject i , is given by:

$$L_{\text{NLL}} = \frac{1}{N} \sum_{i=1}^N \left(\frac{(a_i - \hat{\mu}_i)^2}{2\hat{\sigma}_i^2} + \frac{1}{2} \log(2\pi\hat{\sigma}_i^2) \right) \quad (8)$$

The performance of both loss functions was analyzed on the validation set and in cross-cohort testing, with the MSE loss being adopted as the primary optimization objective for all downstream analyses on brain age gaps and neurological conditions.

2.4.6 Optimization and Regularization

During training, the gradient of the MSE loss with respect to model parameters guides the optimization process. The Adam optimizer is employed with a learning rate of 5×10^{-4} for the CNN and 1×10^{-4} for the ViT backbone, chosen to balance convergence speed with training stability. The adaptive learning rate mechanisms of Adam help navigate the complex loss landscape created by the deep architecture, automatically adjusting the effective learning rate for each parameter based on historical gradient information.

Regularization is achieved through multiple mechanisms to prevent overfitting. **Dropout** layers with probability 0.3 are applied after each fully connected layer, randomly deactivating neurons during training to encourage robust feature learning. **Layer normalization** stabilizes the training process and acts as an implicit regularizer by normalizing the distribution of activations. The **residual connections** in the convolutional blocks also serve a regularizing role by providing alternative gradient paths and preventing the learning of overly complex transformations.

2.5 Model Interpretability via Guided Backpropagation and ViT Patch Mapping

To interpret the age-relevant features learned by our two-stage model, we implement a 3D interpretability pipeline

that combines guided backpropagation of the CNN regression model with patch-wise relevance mapping of the pre-trained Vision Transformer (ViT). This approach allows us to generate slice-level saliency maps, which are subsequently projected onto the ViT patch embeddings and aggregated across subjects to produce a volumetric attention representation of the brain.

2.5.1 Guided Backpropagation of CNN Features

We begin by generating guided gradients with respect to the CNN input feature map. The CNN, previously trained on ViT-derived feature embeddings, is set to evaluation mode. Guided backpropagation is implemented by registering hooks on all ReLU and SiLU activations within the network. Forward hooks record the activation values, while backward hooks allow only positive gradients to propagate through neurons that were active during the forward pass. Given an input tensor representing the sagittal slice feature map (1, 160, 768) and the subject’s binary sex encoding, the guided backpropagation module produces a normalized gradient map for each slice. This map highlights input dimensions that positively influence the predicted brain age.

2.5.2 Projection onto ViT Patch Tokens

To relate the CNN saliency to anatomical locations within the original MRI slices, we map the guided gradients onto ViT patch embeddings. Each preprocessed MRI slice is passed through the ViT backbone to obtain a sequence of patch embeddings. The guided CNN vector is projected onto these embeddings via a dot product, yielding a patch-wise importance score. Additionally, the attention weights from the CLS token to each patch in the final transformer layer are extracted. Both patch scores and CLS attention are normalized to the range $[0, 1]$ for consistent scaling.

2.5.3 Slice-Level Heatmap Generation

Patch scores and CLS attention maps are multiplicatively fused to form a slice-level heatmap. The fused patch grid is then upsampled to the original slice resolution using bilinear interpolation and normalized. This process produces a per-slice saliency image that reflects the combined importance from the CNN and ViT representations.

2.5.4 Aggregation into 3D Attention Volume

To construct a volumetric representation, slice-level heatmaps are averaged across multiple subjects in the validation set for each sagittal slice. The resulting per-slice maps are stacked sequentially along the sagittal axis, producing a 3D attention volume (S, H, W), where S is the number of slices, and H and W are the slice height and width. The volume is finally normalized, thus

Algorithm 1: Proposed BrainRotViT Framework

Input: MRI volume $\mathcal{V}$, chronological age a , sex $s \in \{0, 1\}$

Output: Predicted brain age $\hat{y}$, Brain-Age-Gap BAG = $\hat{y} - a$

1 Stage 1: Vision Transformer (ViT) Representation Learning

2 Sample 32 evenly spaced sagittal slices $\{S_i\}_{i=1}^{32}$ from $\mathcal{V}$

3 for $i \leftarrow 1$ **to** 32 **do**

4 | Divide slice S_i into non-overlapping patches, project them into embeddings, and add positional encodings

5 | Forward the embeddings through L Transformer encoder layers consisting of multi-head self-attention and feed-forward sublayers with residual connections

6 | Apply a linear classification head with softmax to predict age-sex composite class for slice S_i

7 end

8 Compute cross-entropy loss:

$$\mathcal{L}_{\text{CE}} = -\frac{1}{B} \sum_{i=1}^B \sum_{c=1}^C y_{i,c} \log \hat{y}_{i,c}$$

Train until convergence, then discard classification head and freeze Enc_{ViT}

9 Stage 2: Feature Map Construction

10 for $i \leftarrow 1$ **to** 160 **do**

11 | Extract 768-dimensional feature embedding z_i from slice S_i using the frozen ViT encoder

12 end

13 Stack embeddings in spatial order to form feature tensor:

$$Z = [z_1, z_2, \dots, z_{160}]^T \in \mathbb{R}^{160 \times 768}$$

14 Stage 3: Residual CNN Regression

15 Input feature map Z into a residual CNN with three convolutional blocks

16 Each block applies convolution, SiLU activation, and normalization, with residual paths to preserve information

17 Flatten the final feature map F_3 into vector h

18 Pass h through two fully connected layers with layer normalization, SiLU activation, and dropout, producing h_2

19 Concatenate sex information: $h_c = [h_2; s]$

20 Predict brain age with a linear layer:

$$\hat{y} = W_3 h_c + b_3$$

Compute mean squared error:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^N (a_i - \hat{y}_i)^2$$

21 Output: Predicted age $\hat{y}$ and Brain-Age-Gap BAG = $\hat{y} - a$

Algorithm 2: Model Interpretability Pipeline

Input: Trained ViT Encoder Enc_{ViT} , Trained CNN Regressor CNN_{reg}
Input: Validation set subjects $\{(\mathcal{V}_j, s_j)\}_{j=1}^{N_{\text{val}}}$, AAL3 Atlas $\mathcal{A}$
Output: Aggregated 3D Attention Volume $\mathbf{A}_{\text{agg}}$, ROI Importance Scores

- 1 **Initialize:** Aggregated Volume $\mathbf{A}_{\text{agg}} \leftarrow \mathbf{0} \in \mathbb{R}^{S \times H \times W}$ (where $S = 160$)
- 2 **Iterate through each subject in the validation set**
- 3 **for** $j \leftarrow 1$ **to** N_{val} **do**
- 4 **1. Guided Backpropagation of CNN Features**
- 5 Generate feature map $\mathbf{Z}_j \in \mathbb{R}^{160 \times 768}$ from $\mathcal{V}_j$ using frozen Enc_{ViT}
- 6 Perform forward pass: $\hat{y}_j = \text{CNN}_{\text{reg}}(\mathbf{Z}_j, s_j)$
- 7 Compute guided gradient $\mathbf{G}_j = \nabla_{\mathbf{Z}_j} \hat{y}_j$, propagating only positive gradients through active (Si)LU units
- 8 Normalize $\mathbf{G}_j \in \mathbb{R}^{160 \times 768}$
- 9 **2. Projection, Fusion, and Slice-Level Heatmap Generation**
- 10 **for** $i \leftarrow 1$ **to** 160 (each sagittal slice S_i in $\mathcal{V}_j$) **do**
- 11 Pass slice S_i through Enc_{ViT} to get patch embeddings $\mathbf{P}_i \in \mathbb{R}^{N_p \times 768}$ and CLS attention $\alpha_{\text{CLS},i} \in \mathbb{R}^{N_p}$
- 12 Extract slice gradient vector $\mathbf{g}_{j,i} \in \mathbb{R}^{768}$ from $\mathbf{G}_j[i, :]$
- 13 Compute patch-wise importance scores: $\mathbf{s}_{j,i} = \mathbf{P}_i \cdot \mathbf{g}_{j,i}^T \in \mathbb{R}^{N_p}$
- 14 Normalize scores: $\mathbf{s}_{j,i}^{\text{norm}} \leftarrow \text{Normalize}(\mathbf{s}_{j,i}, [0, 1])$
- 15 Normalize attention: $\alpha_{j,i}^{\text{norm}} \leftarrow \text{Normalize}(\alpha_{\text{CLS},i}, [0, 1])$
- 16 Fuse maps: $\mathbf{f}_{j,i} = \mathbf{s}_{j,i}^{\text{norm}} \odot \alpha_{j,i}^{\text{norm}}$ (element-wise multiplication)
- 17 Reshape $\mathbf{f}_{j,i}$ to patch grid and upsample to slice resolution (H, W) using bilinear interpolation $\rightarrow \mathbf{H}_{j,i}$
- 18 $\mathbf{A}_{\text{agg}}[i, :, :] \leftarrow \mathbf{A}_{\text{agg}}[i, :, :] + \text{Normalize}(\mathbf{H}_{j,i})$
- 19 **end**
- 20 **end**
- 21 **3. 3D Volume Aggregation and ROI Identification**
- 22 Average across subjects: $\mathbf{A}_{\text{agg}} \leftarrow \mathbf{A}_{\text{agg}} / N_{\text{val}}$
- 23 Normalize the final 3D volume $\mathbf{A}_{\text{agg}}$
- 24 Align $\mathbf{A}_{\text{agg}}$ with AAL3 atlas $\mathcal{A}$ via cropping and affine registration
- 25 Initialize empty list ROI Scores
- 26 **for** each region R in $\mathcal{A}$ **do**
- 27 Get voxel indices $\mathcal{V}_R$ for region R
- 28 Calculate mean intensity: $I_R = \frac{1}{|\mathcal{V}_R|} \sum_{v \in \mathcal{V}_R} \mathbf{A}_{\text{agg}}[v]$
- 29 Store $(R, I_R, |\mathcal{V}_R|)$ in ROI Scores
- 30 **end**
- 31 Group ROIs (e.g., all cerebellum parts) and compute weighted average intensity based on voxel count $|\mathcal{V}_R|$
- 32 **return** $\mathbf{A}_{\text{agg}}$ and Grouped ROI Scores

achieving an aggregation that preserves anatomical continuity while highlighting regions consistently associated with brain aging across the dataset.

2.5.5 Identification of Regions of Interest

To identify regions of interest (ROI) that receive the most attention from our aging model, the 3D attention volume is cropped, centered, and aligned with affine registration to fit a bounding box of the same size and shape as the Automated Anatomical Labelling atlas 3 (AAL3), so as to get a complete overlap between the 3D attention volume and the atlas. We iterate over all the regions in the AAL atlas to calculate the average intensity of the corresponding voxels in the attention map, and then group the regions and their intensities weighed by voxel count into larger regions (eg. various ROIs of the cerebellum

in the atlas are grouped into a single ROI to represent the cerebellum).

2.6 Experimental Setup

All experiments were carried out on two Debian 12 workstations equipped with NVIDIA GeForce RTX 3060 (12 GB VRAM) GPUs and AMD Ryzen 7 processors (3700X and 5700G), paired with 16 GB of system memory. The implementation was primarily done using the PyTorch framework [51] with CUDA acceleration, and the Vision Transformer (ViT) component was adapted from the official Hugging Face Transformers [52] implementation. Brain MRI preprocessing, including image loading, cropping, resampling, and normalization, was performed using NiBabel [53] and SciPy [54]. The ViT encoder was trained for 5 epochs

on the auxiliary age–sex classification task to learn slice-level representations, while the residual CNN regression head was trained for up to 200 epochs for brain-age prediction. Training followed the same optimization settings described earlier, using the Adam optimizer and mean squared error (MSE) loss. Early stopping based on the validation mean absolute error (MAE) was applied to prevent overfitting.

Overall, the training pipeline remained stable throughout, with the ViT stage converging within roughly six hours and the CNN stage completing in about fourteen hours on the described hardware configuration.

3 Results

This section presents the experimental outcomes of the proposed framework. We first report the quantitative performance on the validation set and compare it with several state-of-the-art (SOTA) approaches. We then evaluate the generalization capability of the model through independent cohort testing. Next, we present the relevance of

3.1 Validation Performance and Comparison with State-of-the-Art

The proposed hybrid Vision Transformer–CNN framework achieved strong predictive performance on the validation set, reaching a mean absolute error (MAE) of **3.34 years**, and Pearson’s correlation coefficient (r) of **0.98** with chronological age. The coefficient of determination (R^2) of **0.95** further indicates high predictive fidelity and stability of the model across subjects.

Table 2 summarizes the comparative results against representative state-of-the-art methods from prior studies, along with baseline models that have been used in multiple studies. For fair comparison, all methods were evaluated on the same distribution of samples from datasets of healthy adults using structural MRI volumes that underwent the same pre-processing routine as input.

Table 2: Comparison of the proposed BrainRotViT model with existing methods on validation data.

Method	MAE	Pearson r	Spearman ρ	R^2
3D-ResNet (baseline)	5.41	0.96	0.91	0.88
3D-ViT (baseline)	6.20	0.82	0.87	0.79
Two-Stage Age Network (TSAN) [55]	3.62	0.93	0.97	0.91
TSAN 1 st -stage	4.12	0.94	0.97	0.93
Global-Local Transformer [25]	4.98	0.95	0.92	0.92
Triamese-ViT [29]	5.80	0.89	0.92	0.80
BrainRotViT (ours)	3.34	0.98	0.97	0.95

3.2 Independent Cohort Testing

To assess the generalization capability of the proposed BrainRotViT, we performed tests on four independent cohorts. Table 3 summarizes the results of cross-cohort evaluation. The proposed model maintained consistently low MAE values in both cohorts, underscoring its strong cross-dataset robustness.

Table 3: Independent cohort evaluation of BrainRotViT. Each dataset was excluded from training and used only for testing.

Dataset	MAE	R^2	Pearson r	Spearman ρ
SALD	4.81	0.86	0.93	0.91
Truecrime	3.75	0.63	0.83	0.55
AgeRisk	5.04	0.81	0.86	0.88
SUDMEX	4.93	0.83	0.76	0.68

To compare the generalizability of our method against state-of-the-art approaches, we chose the TSAN first-stage network since it was among the best in terms of prediction performance.

Table 4: Independent cohort performance comparison with SOTA method.

Dataset	MAE (ours)	MAE (TSAN 1 st -stage)
SALD	4.81	5.79
Truecrime	3.75	5.64
AgeRisk	5.04	7.55
SUDMEX	4.93	5.25

3.3 Training Time and Computational Efficiency

To evaluate computational efficiency, we compared the end-to-end training durations of our proposed BrainRotViT framework with several baseline and state-of-the-art (SOTA) architectures.

Table 5: Comparison of approximate training times across different models. All models were trained on identical hardware settings (RTX 3060, 16 GB RAM).

Model	Training Time (hours)
BrainRotViT (ours)	22
3D-ResNet	24
3D-ViT	30
TSAN	36
Triamese-ViT	52
Global-Local Transformer	60

As shown in Table 5, BrainRotViT achieves an optimal balance between accuracy and computational cost. Despite its hybrid design, it completes training in significantly less time than heavier transformer-based models such as Triamese-ViT and Global-Local Transformer,




Figure 3: **(a)**: Scatter plot of predicted age of validation samples from 11 datasets (MAE = 3.34 years, Pearson $r = 0.98$, Spearman $\rho = 0.97$, $R^2 = 0.95$). **(b)**: Scatter plot of predicted age of independent test samples from 4 datasets (MAE = 4.72 years, Pearson $r = 0.88$, Spearman $\rho = 0.90$, $R^2 = 0.83$). **(c)**: Scatter plot of validation on 5,581 ADNI subjects only. **(d)**: Scatter plot of validation on 1,141 ABIDE-II subjects only. **(e)**: Proportion of AD and AD/MCI subjects with extreme positive aging ($BAG > 1std.dev.$) and extreme negative aging ($BAG < -1std.dev.$) visualized. **(f)** Proportion of AD/MCI subjects with extreme positive aging ($BAG > 1std.dev.$) visualized against other AD/MCI subjects in ADNI. **(g)** Proportion of ASD subjects with extreme positive aging ($BAG > 1std.dev.$) visualized against other ASD subjects in ABIDE-II.

while outperforming both pure CNN and pure transformer baselines in predictive accuracy. This efficiency underscores the scalability of our approach for large, multi-site MRI datasets.

3.4 Impact of Brain Aging on Alzheimer’s Disease and Neurodegeneration

Training and validating our model on 5,581 samples of the ADNI dataset which were grouped into AD (Alzheimer’s Disease), MCI (Mild Cognitive Impairment), and CN (Cognitively Normal), we achieved an MAE of 1.89 years and R^2 of 0.87 (Figure 3c). Plotting the brain age-gaps of each sample against their cognitive status group (Figures 3f, 3e), we found that subjects who exhibited extreme aging (subjects with a positive brain age-gap greater than one standard deviation) were 1.73 times as likely to belong to the AD/MCI group rather than the CN group, compared to those who did not (odds ratio = 1.731, 95% CI [1.192–2.512]; relative risk = 1.319, 95% CI [1.114–1.562]; $p = 0.0048$). Moreover, comparing extreme positive agers with extreme negative agers ($BAG < -1$ std. dev.), we found that the former were 2.18 times as likely to belong to the AD/MCI group rather than CN (odds ratio = 2.180, 95% CI [1.309–3.630]; relative risk = 1.515, 95% CI [1.148–2.001]; $p = 0.0039$). These findings suggest that individuals whose brains undergo accelerated aging are more likely to develop Alzheimer’s disease or mild cognitive impairment than individuals with normally-aged or relatively slow-aged brains, which is widely supported by the existing body of research [2, 4, 56, 5].

3.5 Impact of Autism Spectrum Disorder on Brain Aging

The Autism Brain Imaging Data Exchange II (ABIDE II) is a case-control study that consists of roughly equal numbers of autistic and control subjects. Training and validating our model on 1,141 samples from the dataset, we achieved an MAE of 1.13 years and R^2 of 0.93 (Figure 3d). Plotting the brain age-gaps of each sample against their diagnosis status (Figure 3g), we found that ASD subjects were 1.67 times as likely to exhibit accelerated brain aging ($BAG > 1$ std. dev.) compared to control subjects. (odds ratio = 1.6736, 95% CI [1.100–2.546]; $p = 0.0204$). Our findings suggest that individuals with ASD are significantly more likely to exhibit structural patterns of accelerated brain aging, which coincides with studies that have found age-related changes to be more pronounced in autistic subjects [57, 58].

3.6 Attention Map Analysis

To further interpret the spatial focus of the Vision Transformer within the proposed BrainRotViT framework, we visualize the mean attention distribution across subjects

in the validation set. Figure 4a illustrates the aggregated attention heatmap derived from the middle sagittal slice (slice index 80) averaged over all $n = 1022$ validation samples. The 3D attention visualization in Figure 4b reveals several neuroanatomical regions exhibiting strong activation from the model, corresponding to areas known to undergo age-related structural and functional changes.

The lower and medial posterior regions of the cerebellum show high attention concentration, reflecting their involvement in age-associated decline in motor coordination and memory consolidation, as has been supported by several studies [59, 60, 61]. Particularly the cerebellar vermis, whose volume has been strongly associated with aging-related changes [62, 63, 64], appears as a region that our aging model finds interesting invariably across age ranges. Close to the cerebellum, the cuneus is another region exhibiting high intensity in our attention maps. While recent studies using functional MRI data have noted changes in cuneus activation with age [65, 66], our study uses structural data only, thus hinting that with aging the cuneal cortex undergoes structural changes as well.

The precentral cortex, which exhibited high attention concentration in our model, has been linked with age-associated changes quite commonly in the literature. Prior studies have identified thinning of the precentral cortex and the atrophy of the primary motor cortex as quantitative markers of aging [67, 68]. One study compares such changes with the structural changes in the postcentral gyrus [69], which also appears as a high-attention region in our model. Functional differences in the primary motor cortex, which is located within the precentral cortex, have also been attributed to aging [70, 71, 72], thereby further substantiating the biological interpretability of our aging model.

Across all age ranges, the temporal lobes, particularly the middle temporal gyrus, remained as a prominent region of interest for our aging model. This coincides with several prior studies that have confirmed the role of the temporal lobes in aging and neurodegeneration [73, 74, 75]. Moreover, comparing the volume of the region of relative intensity greater than 0.4 near the temporal gyrus between all age groups (Figures 4d–4k), it was observed that the volume of the high-attention region in the temporal lobe increases with age, suggesting that changes in the temporal lobes accelerate with age, which is also supported by Fujita et al. (2023) [76].

While slightly higher than average attention concentration was observed in a part of the frontal lobe, in general the frontal lobe exhibited relatively lower attention weight compared to postcentral and cerebellar regions (Figure 4b). However, it was observed that the intensity of attention applied by the model to the frontal lobe varies significantly by age group (Figures 4d–4k).




Figure 4: **Model Interpretation.** (a): Example of per-slice attention map extracted from vision transformer. (b) Mid-slice section views of fused 3D attention map along axial, coronal, and sagittal planes (from left to right). (c): Saliency map of trained residual CNN extracted through guided back propagation. (d)–(k): 3D attention map previews by age range.

Sampling the intensity in a region in the medial part of the superior frontal gyrus, the greatest intensity was found among samples in the 10-20 year age group. This suggests that the frontal lobe undergoes the most rapid change specifically in adolescence and early adulthood, as is backed by numerous prior studies [77, 78, 79, 80].

4 Ablation Study

To evaluate the contribution of individual architectural components and hyperparameter choices, we conduct a series of ablation experiments on the validation set. Each experiment isolates or modifies one aspect of the proposed method while keeping all other settings constant. The results highlight the complementary roles of the Vision Transformer (ViT) and CNN modules, the stabilizing effect of residual connections, and the influence of learning rate on convergence and predictive accuracy.

Changing the ViT variant in the backbone We replaced the ViT backbone with several popular variants to observe how sensitive the pipeline is to the transformer choice. The results are summarized in Table 6.

Table 6: Effect of different ViT variants used as the frozen encoder.

Backbone	MAE
ViT base (ours)	3.34
DeiT	3.50
RoPE ViT	3.53
MaxViT	3.39
Swin	3.45

The basic ViT performs best overall. MaxViT gives a close result but is heavier and slower. DeiT and Swin show slightly higher error, indicating that the original ViT base provides the best balance between feature richness and efficiency. The DeiT model benefits from data-efficient training but loses some fine-grained spatial encoding needed for volumetric consistency. Swin Transformer performs better than DeiT due to its hierarchical windows but still falls behind. RoPE ViT (ViT with Rotary Position Embeddings) provides rotational invariance through positional encoding but does not improve regression accuracy in this setting. Overall, the basic ViT remains the most reliable and stable choice for large, multi-site MRI cohorts.

Age-sex grouping granularity We examined the effect of changing the bin width for age-sex composite classes during ViT pretraining. Table 7 shows the results.

During ViT pretraining, age-sex categories were treated as discrete labels under a categorical cross-entropy loss. This setup implicitly encourages the model to cluster feature representations according to these group boundaries. When the bins are too narrow (5 years), the model overfits to minor inter-class variations,

Table 7: Effect of age-sex group granularity used during ViT pretraining.

Group Width (years)	MAE
5	4.60
10 (ours)	3.34
20	5.10

as each class has fewer examples, weakening generalization. On the other hand, overly broad bins (20 years) blur distinct morphological aging cues, making the classification objective less informative. The 10-year grouping provides a meaningful middle ground, giving the ViT enough contrast between classes while preserving intra-class consistency, which in turn improves regression performance after fine-tuning.

Number of evenly spaced slices during ViT training We trained the ViT with different numbers of evenly spaced sagittal slices per subject. The performance is reported in Table 8.

Table 8: Effect of the number of evenly spaced slices sampled per volume during ViT pretraining.

Slices per Volume	MAE
16	3.84
32 (ours)	3.34
64	3.63

Using too few slices reduces spatial coverage, limiting the ViT’s ability to capture subtle structural variations across the brain. Conversely, sampling too many slices increases redundancy, as adjacent slices contain highly correlated information, which adds computational cost without substantially improving feature learning. The 32-slice setting provides a practical balance, offering sufficient anatomical coverage to capture meaningful patterns while maintaining efficient training. This choice also helps the model generalize better by avoiding overfitting to redundant information in dense slice sampling.

Architecture variations in front of ViT We compared several alternative designs to replace the CNN regression head. Table 9 summarizes the results.

Table 9: Effect of different architectures placed after the ViT encoder.

Front End	MAE
CNN with residual blocks (ours)	3.34
ViT	6.34
MLP only	7.91
bi-GRU sequence model	3.63

A simple MLP or another transformer performs poorly, as they fail to preserve the spatial continuity between slice embeddings. While the bi-GRU can model inter-slice dependencies, it lacks sensitivity to local 2D spatial

structure. In contrast, the residual CNN effectively integrates ViT-derived features by capturing fine-grained local spatial information and maintaining spatial coherence across slices.

The ablation results demonstrate that each architectural component contributes meaningfully to the overall performance of proposed method. The ViT provides essential global representations, the CNN enhances spatial localization, and the residual structure ensures efficient training and stability. Moreover, tuning the learning rate plays a pivotal role in achieving optimal convergence. Collectively, these analyses confirm the robustness and balanced design of the proposed hybrid model.

Effect of Methodological Variants in ViT and CNN Integration To further examine the influence of different ways of incorporating sex and age conditioning into the hybrid framework, we conducted additional experiments using the same ViT backbone (ViT-Base) and the best-performing CNN regression head. The goal of these tests was to isolate the contribution of age-sex representation strategies and fusion mechanisms to the final performance. Table 10 summarizes the results.

Table 10: Effect of methodological variants using ViT-Base encoder and best CNN regression head. All results are reported on the validation set.

Model Variant (ViT)	MAE
With age-sex composite classes (ours)	3.34
With age classes only (no sex)	4.94
With age classes + sex input via FiLM	5.61
With direct age regression head (no sex)	5.21
With regression head + sex input via FiLM	6.87
With sex-aware encoder but CNN without sex input	3.71

As seen in Table 10, the proposed configuration, where age and sex are jointly encoded through discrete composite classes during ViT pretraining and later fused in the CNN regression head—achieves the lowest MAE of 3.34 years. Removing sex information or regressing age directly in the ViT substantially increases error, underscoring the benefit of structured, categorical representation learning. Although FiLM-based modulation, which condition the ViT embeddings or intermediate CNN activations on the sex variable through learned affine transformations, improves over unconditioned baselines (MAE = 4.61 vs. 4.87), the categorical composite-class strategy still provides more robust features that generalize better across datasets.

Number of convolutional layers We experimented with varying depths for the CNN regression head. The original three-block configuration remains the best performer, as shown in Table 11.

Shallower networks underfit due to limited receptive fields and weaker hierarchical feature extraction. Adding more layers beyond three increases the number of parameters and slightly degrades performance, likely from over-

Table 11: Effect of varying the number of convolutional blocks in the CNN regression head.

Number of Conv Blocks	MAE
1	4.12
2	3.66
3 (ours)	3.34
4	3.37
5	3.42
6	3.49

smoothing and redundant representations. The three-block configuration provides the best trade-off between expressiveness and generalization, keeping both training stable and inference efficient.

CNN kernel sizes We tested multiple kernel configurations in the residual CNN blocks. Results are given in Table 12.

Table 12: Effect of different kernel size configurations in the residual CNN.

Kernel Sizes	MAE
(10,60),(5,15),(2,6) (ours)	3.34
(12,80),(6,20),(3,8)	3.38
(6,40),(4,12),(2,6)	3.45
(8,50),(5,15),(1,4)	3.39
Dilated variant	3.41

Our selected kernel sizes provide the best compromise between cross-slice coverage and fine-grained spatial precision. Larger kernels extend the receptive field but tend to blur subtle anatomical differences, while smaller ones capture limited context and lead to underfitting. The chosen configuration preserves both local structure and inter-slice coherence effectively.

Activation functions We replaced the activation used in the CNN and fully connected layers to test the effect of nonlinearity. Results are shown in Table 13.

Table 13: Effect of different activation functions in the CNN and FC layers.

Activation	MAE
SiLU (ours)	3.34
ReLU	3.41
Leaky ReLU	3.54
GELU	3.47

SiLU yields the lowest error, likely due to its smooth gradient flow and compatibility with residual layers. ReLU performs reasonably well but lacks the subtle gradient transitions that SiLU provides. Leaky ReLU slightly improves over ReLU for negative activations but offers no clear benefit here. GELU falls behind too, performing worse than the other three. Overall, SiLU provides the most stable convergence and consistent generalization across datasets.

Impact of Late Sex Fusion To evaluate the contribution of biological sex information in the CNN regression stage, we compare the performance of the model with and without late fusion of the binary sex variable at the final fully connected layer. As shown in Table 14, incorporating sex improves model accuracy substantially. The model achieves a validation MAE of **3.34 years** with sex fusion, while the MAE increases to **4.59 years** when sex is excluded. This demonstrates that sex-specific differences in brain structure play a critical role in accurate brain age estimation. The late fusion strategy allows earlier layers to learn generalized, sex-independent representations while leveraging sex information only at the final stage, improving both robustness and biological interpretability.

Table 14: Effect of incorporating biological sex information in the CNN regression head.

Configuration	Validation MAE (years)
With Sex Fusion	3.34
Without Sex Fusion	4.59

Effect of Learning Rate Variation We further study the influence of the learning rate on convergence behavior and final predictive performance. The proposed model was originally trained with a learning rate of 5×10^{-4} , which achieved the best trade-off between stability and accuracy. To explore the sensitivity of the model to this hyperparameter, we vary the learning rate across an order of magnitude and interpolate the results around the optimal configuration.

Table 15: Effect of learning rate on model performance (validation set).

Learning Rate	1×10^{-4}	5×10^{-4} (ours)	1×10^{-3}	5×10^{-3}	1×10^{-2}
MAE	4.52	3.34	3.78	4.96	6.27
RMSE	4.72	3.46	3.92	5.12	6.85
r	0.96	0.98	0.97	0.95	0.91
R ²	0.90	0.93	0.91	0.87	0.80

As shown in Table 15, both excessively low and high learning rates lead to suboptimal results. Lower rates cause slower convergence and local minima entrapment, while higher rates result in unstable updates and overfitting. The learning rate of 5×10^{-4} yields the best overall metrics, indicating a stable optimization regime for the hybrid ViT-CNN architecture.

Effect of Removing Residual Connections in CNN Residual connections are known to facilitate gradient propagation and stabilize deep convolutional architectures. To verify their impact, we train a variant of the CNN regression head where skip connections are removed, forcing the network to learn purely sequential transformations.

Table 16: Impact of removing residual connections in the CNN on validation performance.

Model Variant	MAE	r	R ²
With Residual Connections	3.34	0.98	0.93
Without Residual Connections	3.68	0.97	0.91

The removal of residual pathways leads to a moderate increase in MAE (from 3.34 to 3.68) and slight degradation in correlation metrics. This confirms that residual links effectively stabilize learning and preserve fine-grained spatial representations without causing vanishing gradients.

Comparison of Loss Functions: MSE vs. NLL

We further analyze the impact of different regression loss functions on model performance and generalization. The standard Mean Squared Error (MSE) loss penalizes large deviations more strongly and encourages smooth convergence, while the Negative Log-Likelihood (NLL) loss introduces probabilistic uncertainty modeling. The comparison is summarized in Table 17.

Table 17: Performance comparison between MSE and NLL losses on validation and cross-cohort test (SALD dataset).

Loss Function	Val. MAE	Cross-Cohort MAE
MSE Loss	3.34	5.83
NLL Loss	3.45	4.81

The results indicate that while the NLL loss yields slightly higher MAE values on the validation set, it provides better cross-cohort testing performance due to the added benefit of uncertainty estimation. However, MSE loss remains superior for regression accuracy on the same distribution, enforces a tighter bound on brain age gaps (BAG), and is therefore adopted as the primary optimization objective in downstream analyses on the BAG and its association with neurological disorders.

5 Discussion

In this study, we propose a framework that demonstrates how combining the global contextual modeling capacity of Vision Transformers with the spatial precision of residual CNNs yields a powerful and interpretable architecture for brain-age estimation. Our BrainRotViT model consistently outperformed both CNN-based and transformer-based state-of-the-art approaches while maintaining computational efficiency. Although we strive to make fair comparisons with methods proposed in prior studies, it remains an inherently complicated task. It is common in this domain for studies to compare against others by directly quoting the reported metrics [81, 82, 83]; however, this does not accurately represent the true learning capabilities of those models. Each study uses its own dataset collection with a unique distribution, and to ensure a fair comparison, all




Figure 5: **Ablation study results.** (a) Comparison of ViT backbone architectures; (b) Impact of age-sex group granularity; (c) Effect of slice count per volume; (d) Front-end architecture comparison; (e) Sex information integration methods; (f) Optimal depth of convolutional blocks; (g) Kernel size configurations; (h) Activation function comparison; (i) Sex fusion effectiveness; (j) Learning rate sensitivity; (k) Impact of residual connections; (l) Loss function comparison on validation and cross-cohort test sets.

methods should be trained and validated on the same data distribution. In our work, we made an effort to achieve fairness by re-running the training, validation, and testing pipelines of prior methods on an identical distribution of MRI images, all processed through the same preprocessing routine.

To make our model robust for testing on external datasets, we curated an extensive collection of open-access sMRI images from 15 different global datasets. Such a diverse data distribution is particularly vulnerable to site-specific artifacts, differences in scanner protocols, and demographic variability. Although the heterogeneity of this data somewhat affects the validation MAE, our model demonstrates excellent predictive accuracy on individual datasets (Figures 3c, 3d), exceeding that of existing state-of-the-art methods [25, 84], while also remaining robust across external datasets

(Tables 3, 4) and training faster. These results highlight the model’s resilience to site effects, which is a critical property for translational neuroimaging applications.

Although our model effectively learns the regions of the brain associated with aging, it exhibited relatively low attention to the frontal cortex compared to the higher attention observed in post-frontal cortical regions and the cerebellum. We found that substantial attention to the frontal lobe occurred primarily among subjects aged 10–20 years, suggesting that frontal cortical changes are less prominent in older individuals. This finding contrasts with prior literature suggesting that the frontal cortex is, in general, the region most susceptible to age-related change [85, 86], but nonetheless offers a new perspective on studying brain aging across the human lifespan.

Our analysis of brain age gaps among autistic and control subjects in the case-control ABIDE-II study suggests that autistic subjects are more likely to exhibit higher age gaps, i.e., to show accelerated aging. While prior studies have reported that aging-related structural changes appear more pronounced among autistic individuals, there is no direct evidence in the literature to suggest that their brains age faster than those of controls. In fact, one study reported lower age gaps among ASD subjects compared to controls [7], supporting the notion of delayed development in autism. However, we did not observe any such statistically significant pattern in our brain age-gap data.

The design of our learning architecture is novel in its combination of a transformer encoder with a convolutional regression module that applies convolutions over a 2D matrix of concatenated embeddings. While sagittal slices clearly follow a spatial axis, treating the embedding dimension of transformer outputs as spatial may seem semantically inappropriate. Embedding vectors represent abstract features, and there is no inherent reason to expect a direct correspondence between the positional values of two embeddings produced by the same transformer. However, we argue that neighboring sagittal slices of a 3D MRI are so anatomically similar that a vision transformer will tend to embed them in similar ways, making it plausible that the continuity of positional values between adjacent embeddings carries semantic meaning. To validate this assumption, we calculated the cosine similarity among all slice embeddings for each 3D MRI and found that neighboring slice embeddings are indeed highly similar, as are embeddings of slices that are antipodal in sagittal symmetry (Figure 2e). Nonetheless, we experimented with replacing the CNN in the architecture’s second stage with a Bi-GRU sequential model and another transformer, since these architectures are more semantically aligned with embedding sequences, but they failed to achieve the local refinement provided by the residual CNN (Table 9).

Finally, we acknowledge the limitations posed by the availability of computational resources. Wider access to more powerful GPUs and larger memory would enable the exploration of more sophisticated architectures, such as employing 3D attention or multiple vision transformers coupled to a CNN optimized under a shared loss function. However, these constraints encouraged us to design a model that is not only accurate but also lightweight and fast to train compared to those proposed in previous studies. Beyond improving brain-age prediction, we believe the proposed architecture itself holds promise for broader applications. By modifying the regression or classification heads, it could be adapted to classify cognitive status, predict conversion from mild impairment to Alzheimer’s disease, or diagnose other neurological conditions from structural MRI patterns, given the availability of sufficiently large imaging datasets.

6 Acknowledgments

Data used in the preparation of this article were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database. 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). The grantee organization is the Northern California Institute for Research and Education. In the past, ADNI has also received funding from the National Institute of Biomedical Imaging and Bioengineering, the Canadian Institutes of Health Research, and generous private sector contributions through the Foundation for the National Institutes of Health (FNIH) from several companies and non-profits (See <https://adni.loni.usc.edu> for full up-to-date information). Data collection and sharing for this project was also partly provided by the Cambridge Centre for Ageing and Neuroscience (CamCAN). CamCAN funding was provided by the UK Biotechnology and Biological Sciences Research Council (grant number BB/H008217/1), together with support from the UK Medical Research Council and University of Cambridge, UK. We thank Annette Weekes-Holder, Garry Detzler, and Baycrest Imaging Centre staff for their help with data collection for the BOLD variability dataset. In the context of the DLBS dataset, we thank the former lab members and associates who have made this research possible, and who are listed in the Keys to the Kingdom. In addition, we thank the participants who generously contributed their time and effort to the DLBS project. This study made use of data from the IXI dataset (<https://brain-development.org/ixi-dataset/>), which are made available under the Creative Commons CC BY-SA 3.0 license. Data used in the study were provided in part by OASIS (OASIS-1: Cross-Sectional: Principal Investigators: D. Marcus, R. Buckner, J. Csernansky, J. Morris; P50 AG05681, P01 AG03991, P01 AG026276, R01 AG021910, P20 MH071616, U24 RR021382). For all the other open-access datasets used in our study, we cordially thank the Neuroimaging Tools and Resources Collaboratory (NITRC) and the Open-Neuro (<https://openneuro.org/>) initiative for providing the data through their image repositories in easily accessible formats.

References

- [1] N. C. Fox and J. M. Schott, “Imaging cerebral atrophy: normal ageing to alzheimer’s disease,” *The Lancet*, vol. 363, no. 9406, pp. 392–394, 2004.
- [2] K. Franke and C. Gaser, “Ten years of brainage as a neuroimaging biomarker of brain aging: what insights have we gained? *front. neurol.* 10, 789,” 2019.
- [3] R. Zhang, F. Yi, H. Mao, Z. Huang, K. Wang, and J. Zhang, “Brain age gap as a predictive biomarker that links aging, lifestyle, and neuropsychiatric health,” *Communications Medicine*, vol. 5, no. 1, p. 441, 2025.

- [4] P. R. Millar, B. A. Gordon, P. H. Luckett, T. L. Benzing, C. Cruchaga, A. M. Fagan, J. J. Hassenstab, R. J. Perrin, S. E. Schindler, R. F. Allegri, *et al.*, “Multimodal brain age estimates relate to alzheimer disease biomarkers and cognition in early stages: a cross-sectional observational study,” *Elife*, vol. 12, p. e81869, 2023.
- [5] M. M. Wittens, S. Denissen, D. M. Sima, E. Fransen, E. Niemantsverdriet, C. Bastin, F. Benoit, B. Bergmans, J.-C. Bier, P. P. De Deyn, *et al.*, “Brain age as a biomarker for pathological versus healthy ageing—a remember study,” *Alzheimer’s research & therapy*, vol. 16, no. 1, p. 128, 2024.
- [6] C. Torenvliet, A. Groenman, T. Radhoe, J. A. van Rentergem, W. Van der Putten, and H. Geurts, “A longitudinal study on cognitive aging in autism,” *Psychiatry research*, vol. 321, p. 115063, 2023.
- [7] Q. Wang, K. Hu, M. Wang, Y. Zhao, Y. Liu, L. Fan, and B. Liu, “Predicting brain age during typical and atypical development based on structural and functional neuroimaging,” *Human Brain Mapping*, vol. 42, no. 18, pp. 5943–5955, 2021.
- [8] J. H. Cole, “Multimodality neuroimaging brain-age in uk biobank: relationship to biomedical, lifestyle, and cognitive factors,” *Neurobiology of Aging*, vol. 92, pp. 34–42, 2020.
- [9] K. Franke, G. Ziegler, S. Klöppel, and C. Gaser, “Estimating the age of healthy subjects from t1-weighted mri scans using kernel methods: Exploring the influence of various parameters,” *NeuroImage*, vol. 50, no. 3, pp. 883–892, 2010.
- [10] I. Beheshti, O. Potvin, and S. Duchesne, “Patch-wise brain age longitudinal reliability,” *Human Brain Mapping*, vol. 42, no. 3, pp. 690–698, 2021.
- [11] J. H. Cole, R. P. Poudel, D. Tsagkrasoulis, M. W. Caan, C. Steves, T. D. Spector, and G. Montana, “Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker,” *NeuroImage*, vol. 163, pp. 115–124, 2017.
- [12] I. Beheshti, N. Maikusa, and H. Matsuda, “The accuracy of t1-weighted voxel-wise and region-wise metrics for brain age estimation,” *Computer Methods and Programs in Biomedicine*, vol. 214, p. 106585, 2022.
- [13] K. Franke, E. Luders, A. May, M. Wilke, and C. Gaser, “Brain maturation: Predicting individual brainage in children and adolescents using structural mri,” *NeuroImage*, vol. 63, no. 3, pp. 1305–1312, 2012.
- [14] I. Beheshti, M. A. Ganaie, V. Paliwal, A. Rastogi, I. Razzak, and M. Tanveer, “Predicting brain age using machine learning algorithms: A comprehensive evaluation,” *IEEE Journal of Biomedical and Health Informatics*, vol. 26, no. 4, pp. 1432–1440, 2022.
- [15] T. T. Le, R. T. Kuplicki, B. A. McKinney, H.-W. Yeh, W. K. Thompson, M. P. Paulus, T. . I. , R. L. Aupperle, J. Bodurka, Y.-H. Cha, J. S. Feinstein, S. S. Khalsa, J. Savitz, W. K. Simmons, and T. A. Victor, “A non-linear simulation framework supports adjusting for age when analyzing brainage,” *Frontiers in Aging Neuroscience*, vol. Volume 10 - 2018, 2018.
- [16] M. A. Ganaie, M. Tanveer, and I. Beheshti, “Brain age prediction using improved twin svr,” *Neural Computing and Applications*, vol. 36, no. 1, pp. 53–63, 2024.
- [17] J. H. Cole, S. J. Ritchie, M. E. Bastin, M. C. Valdés Hernández, S. Muñoz Maniega, N. Royle, J. Corley, A. Pattie, S. E. Harris, Q. Zhang, N. R. Wray, P. Redmond, R. E. Marioni, J. M. Starr, S. R. Cox, J. M. Wardlaw, D. J. Sharp, and I. J. Deary, “Brain age predicts mortality,” *Molecular Psychiatry*, vol. 23, no. 5, pp. 1385–1392, 2018.
- [18] M. Tanveer, M. Ganaie, I. Beheshti, T. Goel, N. Ahmad, K.-T. Lai, K. Huang, Y.-D. Zhang, J. Del Ser, and C.-T. Lin, “Deep learning for brain age estimation: A systematic review,” *Information Fusion*, vol. 96, pp. 130–143, 2023.
- [19] X. Feng, Z. C. Lipton, J. Yang, S. A. Small, and F. A. Provenzano, “Estimating brain age based on a uniform healthy population with deep learning and structural magnetic resonance imaging,” *Neurobiology of Aging*, vol. 91, pp. 15–25, 2020.
- [20] S. He, Y. Feng, P. E. Grant, and Y. Ou, “Deep relation learning for regression and its application to brain age estimation,” *IEEE Transactions on Medical Imaging*, vol. 41, no. 9, pp. 2304–2317, 2022.
- [21] N. Amoroso, M. La Rocca, L. Bellantuono, D. Diacono, A. Fanizzi, E. Lella, A. Lombardi, T. Maggipinto, A. Monaco, S. Tangaro, and R. Bellotti, “Deep learning and multiplex networks for accurate modeling of brain age,” *Frontiers in Aging Neuroscience*, vol. Volume 11 - 2019, 2019.
- [22] M. Ueda, K. Ito, K. Wu, K. Sato, Y. Taki, H. Fukuda, and T. Aoki, “An age estimation method using 3d-cnn from brain mri images,” in *ISBI 2019 - 2019 IEEE International Symposium on Biomedical Imaging*, Proceedings - International Symposium on Biomedical Imaging, (United States), pp. 380–383, IEEE Computer Society, Apr. 2019. Publisher Copyright: © 2019 IEEE.; 16th IEEE International Symposium on Biomedical Imaging, ISBI 2019 ; Conference date: 08-04-2019 Through 11-04-2019.
- [23] H. Peng, W. Gong, C. F. Beckmann, A. Vedaldi, and S. M. Smith, “Accurate brain age prediction with lightweight deep neural networks,” *Medical Image Analysis*, vol. 68, p. 101871, 2021.
- [24] H. Cai, Y. Gao, and M. Liu, “Graph transformer geometric learning of brain networks using multimodal mr images for brain age estimation,” *IEEE Transactions on Medical Imaging*, vol. 42, no. 2, pp. 456–466, 2023.
- [25] S. He, P. E. Grant, and Y. Ou, “Global-local transformer for brain age estimation,” *IEEE Transactions on Medical Imaging*, vol. 41, no. 1, pp. 213–224, 2022.
- [26] H. Cai, Y. Gao, and M. Liu, “Graph transformer geometric learning of brain networks using multimodal mr images for brain age estimation,” *IEEE Transactions on Medical Imaging*, vol. 42, no. 2, pp. 456–466, 2023.
- [27] I. Beheshti, N. Maikusa, and H. Matsuda, “The accuracy of t1-weighted voxel-wise and region-wise metrics for brain age estimation,” *Computer Methods and Programs in Biomedicine*, vol. 214, p. 106585, Feb. 2022.
- [28] Y. Yang, X. Guo, Z. Chang, C. Ye, Y. Xiang, H. Lv, and T. Ma, “Estimating brain age with global and local dependencies,” in *arXiv preprint arXiv:2209.08933*, 2022.

- [29] Z. Zhang and R. Jiang, "Triamese-vit: A 3d-aware method for robust brain age estimation from mris," *arXiv preprint arXiv:2401.09475*, 2024.
- [30] A. Di Martino, M. P. Milham, *et al.*, "Autism brain imaging data exchange ii (abide ii)." https://fcon_1000.projects.nitrc.org/indi/abide/abide_II.html. Accessed 2025-08-26.
- [31] S. G. Mueller, M. W. Weiner, L. J. Thal, and *et al.*, "The alzheimer's disease neuroimaging initiative," *Neuroimaging Clinics of North America*, vol. 15, no. 4, pp. 869–xii, 2005.
- [32] L. Tisdall and R. Mata, "agerisk", 2023.
- [33] J. R. Rieck, G. Baracchini, B. DeSouza, D. Nichol, E. Howard, and C. L. Grady, "bold variability during cognitive control for an adult lifespan sample", 2024.
- [34] J. R. Taylor, N. Williams, R. Cusack, T. Auer, M. A. Shafto, M. Dixon, L. K. Tyler, CamCAN, and R. N. Henson, "The cambridge centre for ageing and neuroscience (camcan) data repository: Structural and functional mri, meg, and cognitive data from a cross-sectional adult lifespan sample," *NeuroImage*, 2016.
- [35] M. A. Shafto, L. K. Tyler, M. Dixon, J. R. Taylor, J. B. Rowe, R. Cusack, A. J. Calder, W. D. Marslen-Wilson, J. Duncan, T. Dalgleish, R. N. Henson, C. Brayne, CamCAN, and F. E. Matthews, "The cambridge centre for ageing and neuroscience (camcan) study protocol: a cross-sectional, lifespan, multidisciplinary examination of healthy cognitive ageing," *BMC Neurology*, vol. 14, p. 204, 2014.
- [36] Center for Biomedical Research Excellence, "COBRE: Center for Biomedical Research Excellence (schizophrenia) dataset." http://fcon_1000.projects.nitrc.org/indi/retro/cobre.html, 2017. RRID:SCR.010482. Dataset of raw anatomical and functional MRI data from 72 patients with schizophrenia and 75 healthy controls (ages 18–65).
- [37] X.-N. Zuo, J. S. Anderson, P. Bellec, R. M. Birn, B. B. Biswal, J. Blautzik, J. C. Breitner, R. L. Buckner, V. D. Calhoun, F. X. Castellanos, A. Chen, B. Chen, J. Chen, S. J. Colcombe, W. Courtney, R. C. Craddock, A. Di Martino, H.-M. Dong, X. Fu, Q. Gong, and M. P. . . . Milham, "An open science resource for establishing reliability and reproducibility in functional connectomics," *Scientific Data*, vol. 1, p. 140049, 2014.
- [38] D. Park, J. Hennessee, E. T. Smith, M. Chan, C. Katen, G. Wig, K. Rodrigue, and K. Kennedy, "the dallas lifespan brain study", 2024.
- [39] re3data.org, "1000 functional connectomes project." <https://doi.org/10.17616/R3W05R>, 2025. Editing status: 2025-08-22; last accessed: 2025-10-09.
- [40] I. C. L. Biomedical Image Analysis Group, "Ixi (information extraction from images) dataset." https://www.nitrc.org/projects/ixi_dataset, 2015. Accessed 2025-08-26.
- [41] A. C. Nugent, A. G. Thomas, M. Mahoney, A. Gibbons, J. Smith, A. Charles, J. S. Shaw, J. D. Stout, A. M. Namyst, A. Basavaraj, E. Earl, D. Moraczewski, E. Guinee, M. Liu, T. Riddle, J. Snow, S. Japee, M. Andrews, A. Pavletic, S. Sinclair, V. Roopchansingh, P. A. Bandettini, and J. Chung, "the nimh healthy research volunteer dataset", 2025.
- [42] D. S. Marcus, R. L. Buckner, J. G. Csernansky, and J. C. Morris, "Oasis-1 (cross-sectional) and oasis-2 (longitudinal) mri datasets." <https://sites.wustl.edu/oasisbrains/>. Accessed 2025-08-26.
- [43] D. Wei, K. Zhuang, Q. Chen, W. Yang, *et al.*, "Southwest university adult lifespan dataset (sald)." https://fcon_1000.projects.nitrc.org/indi/retro/sald.html. Accessed 2025-08-26.
- [44] E. A. Garza-Villarreal, J. J. G. Olvera, T. Balducci, D. A. Valdez, A. Valencia, and J. Rasgado, "sud-mex.conn: The mexican dataset of cocaine use disorder patients.", 2021.
- [45] C. Perchtold-Stefan, C. Rominger, K. Koschutnig, and A. Fink, "truecrime", 2024.
- [46] M. Brant-Zawadzki, G. D. Gillan, and W. R. Nitz, "Mp rage: a three-dimensional, t1-weighted, gradient-echo sequence-initial experience in the brain," *Radiology*, vol. 182, no. 3, pp. 769–775, 1992.
- [47] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," 2015.
- [48] N. J. Tustison, B. B. Avants, P. A. Cook, Y. Zheng, A. Egan, P. A. Yushkevich, and J. C. Gee, "N4itk: Improved n3 bias correction," *IEEE Transactions on Medical Imaging*, vol. 29, no. 6, pp. 1310–1320, 2010.
- [49] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [50] D. A. Nix and A. S. Weigend, "Estimating the mean and variance of the target probability distribution," in *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)*, vol. 1, pp. 55–60 vol.1, 1994.
- [51] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," 2019.
- [52] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush, "Huggingface's transformers: State-of-the-art natural language processing," 2020.
- [53] M. Brett, C. J. Markiewicz, M. Hanke, M.-A. Côté, B. Cipollini, P. McCarthy, C. Cheng, Y. O. Halchenko, D. S. Marcus, O. Hinds, *et al.*, "Nipype: A flexible, lightweight and extensible neuroimaging data processing framework in python," *Frontiers in Neuroinformatics*, vol. 14, p. 10, 2020.
- [54] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, *et al.*, "Scipy 1.0: Fundamental algorithms for scientific computing in python," *Nature Methods*, vol. 17, pp. 261–272, 2020.
- [55] J. Cheng, Z. Liu, H. Guan, Z. Wu, H. Zhu, J. Jiang, W. Wen, D. Tao, and T. Liu, "Brain age estimation from mri using cascade networks with ranking loss," *IEEE Transactions on Medical Imaging*, vol. 40, no. 12, pp. 3400–3412, 2021.

- [56] F. Biondo, A. Jewell, M. Pritchard, D. Aarsland, C. J. Steves, C. Mueller, and J. H. Cole, "Brain-age is associated with progression to dementia in memory clinic patients," *NeuroImage: Clinical*, vol. 36, p. 103175, 2022.
- [57] B. B. Braden and C. Riecken, "Thinning faster? age-related cortical thickness differences in adults with autism spectrum disorder," *Research in autism spectrum disorders*, vol. 64, pp. 31–38, 2019.
- [58] A. Dickinson, S. Jeste, and E. Milne, "Electrophysiological signatures of brain aging in autism spectrum disorder," *Cortex*, vol. 148, pp. 139–151, 2022.
- [59] A. Arleo, M. Bareš, J. A. Bernard, H. R. Bogoian, M. M. Bruchhage, P. Bryant, E. S. Carlson, C. C. Chan, L.-K. Chen, C.-P. Chung, *et al.*, "Consensus paper: cerebellum and ageing," *The Cerebellum*, vol. 23, no. 2, pp. 802–832, 2024.
- [60] D. S. Woodruff-Pak, M. R. Foy, G. G. Akopian, K. H. Lee, J. Zach, K. P. T. Nguyen, D. M. Comalli, J. A. Kennard, A. Agelan, and R. F. Thompson, "Differential effects and rates of normal aging in cerebellum and hippocampus," *Proceedings of the National Academy of Sciences*, vol. 107, no. 4, pp. 1624–1629, 2010.
- [61] C. L. Samstag, N. H. Chapman, L. E. Gibbons, J. Geller, N. Loeb, S. Dharap, M. Yagi, D. G. Cook, K. F. Pagulayan, P. K. Crane, *et al.*, "Neuropathological correlates of vulnerability and resilience in the cerebellum in alzheimer's disease," *Alzheimer's & Dementia*, vol. 21, no. 2, p. e14428, 2025.
- [62] V. Koppelmans, Y. Y. Hoogendam, S. Hirsiger, S. Mérillat, L. Jäncke, and R. D. Seidler, "Regional cerebellar volumetric correlates of manual motor and cognitive function," *Brain Structure and Function*, vol. 222, no. 4, pp. 1929–1944, 2017.
- [63] J. A. Bernard and R. D. Seidler, "Moving forward: age effects on the cerebellum underlie cognitive and motor declines," *Neuroscience & Biobehavioral Reviews*, vol. 42, pp. 193–207, 2014.
- [64] N. Raz, I. J. Torres, W. D. Spencer, K. White, and J. D. Acker, "Age-related regional differences in cerebellar vermis observed in vivo," *Archives of Neurology*, vol. 49, no. 4, pp. 412–416, 1992.
- [65] M. K. Heckner, E. C. Cieslik, S. B. Eickhoff, J. A. Camilleri, F. Hoffstaedter, and R. Langner, "The aging brain and executive functions revisited: implications from meta-analytic and functional-connectivity evidence," *Journal of cognitive neuroscience*, vol. 33, no. 9, pp. 1716–1752, 2021.
- [66] E. Knights, R. Henson, A. Morcom, D. Mitchell, and K. Tsvetanov, "Neural evidence of functional compensation for fluid intelligence in healthy ageing. *elife* 13," *RP93327*, 2023.
- [67] B. C. Clark and J. L. Taylor, "Age-related changes in motor cortical properties and voluntary activation of skeletal muscle," *Current aging science*, vol. 4, no. 3, pp. 192–199, 2011.
- [68] D. H. Salat, R. L. Buckner, A. Z. Snyder, D. N. Greve, R. S. Desikan, E. Busa, J. C. Morris, A. M. Dale, and B. Fischl, "Thinning of the cerebral cortex in aging," *Cerebral cortex*, vol. 14, no. 7, pp. 721–730, 2004.
- [69] L. Zhou, N. Tian, Z.-J. Geng, B.-K. Wu, L.-Y. Dong, and M.-R. Wang, "Diffusion tensor imaging study of brain precentral gyrus and postcentral gyrus during normal brain aging process," *Brain and behavior*, vol. 10, no. 10, p. e01758, 2020.
- [70] J. Wen, Z. Liang, C. Li, H. Pang, L. Jiang, J. Li, X. Guan, J. Zhang, H. Rusinek, X. Xu, *et al.*, "Motor-cognitive aging: The role of motor cortex and its pathways," *NeuroImage*, p. 121472, 2025.
- [71] C.-H. Cheng, M.-Y. Lin, and S.-H. Yang, "Age effect on automatic inhibitory function of the somatosensory and motor cortex: an meg study," *Frontiers in Aging Neuroscience*, vol. 10, p. 53, 2018.
- [72] J. Cirillo, C. R. Ward, N. Gant, S. A. Reading, A. Ren, and W. D. Byblow, "Modulation of primary motor cortex after experimentally induced and use-dependent plasticity in young and older adults," *Experimental Brain Research*, vol. 243, no. 7, p. 159, 2025.
- [73] E. D. Bigler, C. V. Andersob, and D. D. Blatter, "Temporal lobe morphology in normal aging and traumatic brain injury," *American Journal of Neuroradiology*, vol. 23, no. 2, pp. 255–266, 2002.
- [74] G. Bartzokis, M. Beckson, P. H. Lu, K. H. Nuechterlein, N. Edwards, and J. Mintz, "Age-related changes in frontal and temporal lobe volumes in men: a magnetic resonance imaging study," *Archives of general psychiatry*, vol. 58, no. 5, pp. 461–465, 2001.
- [75] S. Hrybouski, S. R. Das, L. Xie, L. E. Wisse, M. Kelley, J. Lane, M. Sherin, M. DiCalogero, I. Nasrallah, J. Detre, *et al.*, "Aging and alzheimer's disease have dissociable effects on local and regional medial temporal lobe connectivity," *Brain Communications*, vol. 5, no. 5, p. fcad245, 2023.
- [76] S. Fujita, S. Mori, K. Onda, S. Hanaoka, Y. Nomura, T. Nakao, T. Yoshikawa, H. Takao, N. Hayashi, and O. Abe, "Characterization of brain volume changes in aging individuals with normal cognition using serial magnetic resonance imaging," *JAMA Network Open*, vol. 6, no. 6, pp. e2318153–e2318153, 2023.
- [77] N. Gogtay, J. N. Giedd, L. Lusk, K. M. Hayashi, D. Greenstein, A. C. Vaituzis, T. F. Nugent III, D. H. Herman, L. S. Clasen, A. W. Toga, *et al.*, "Dynamic mapping of human cortical development during childhood through early adulthood," *Proceedings of the national academy of sciences*, vol. 101, no. 21, pp. 8174–8179, 2004.
- [78] J. N. Giedd, J. Blumenthal, N. O. Jeffries, F. X. Castellanos, H. Liu, A. Zijdenbos, T. Paus, A. C. Evans, and J. L. Rapoport, "Brain development during childhood and adolescence: a longitudinal mri study," *Nature neuroscience*, vol. 2, no. 10, pp. 861–863, 1999.
- [79] R. K. Lenroot and J. N. Giedd, "Brain development in children and adolescents: insights from anatomical magnetic resonance imaging," *Neuroscience & biobehavioral reviews*, vol. 30, no. 6, pp. 718–729, 2006.
- [80] S.-J. Blakemore, "Imaging brain development: the adolescent brain," *Neuroimage*, vol. 61, no. 2, pp. 397–406, 2012.
- [81] A. Aghaei, M. E. Moghaddam, and A. D. N. Initiative, "Brain age gap estimation using attention-based resnet

- method for alzheimer’s disease detection,” *Brain Informatics*, vol. 11, no. 1, p. 16, 2024.
- [82] R. Rajabli, M. Soltaninejad, V. S. Fonov, D. Bzdok, and D. L. Collins, “Brain age prediction: Deep models need a hand to generalize,” *Human Brain Mapping*, vol. 46, no. 11, p. e70254, 2025. e70254 HBM-24-0902.R2.
- [83] I. Kianian and H. Sajedi, “Brain age estimation with a greedy dual-stream model for limited datasets,” *Neurocomputing*, vol. 596, p. 127974, 2024.
- [84] H. Kim, S. Park, S. W. Seo, D. L. Na, H. Jang, J. P. Kim, H. J. Kim, S. H. Kang, and K. Kwak, “A novel deep learning-based brain age prediction framework for routine clinical mri scans,” *npj Aging*, vol. 11, no. 1, p. 70, 2025.
- [85] R. Peters, “Ageing and the brain: This article is part of a series on ageing edited by professor chris bulbitt,” *Postgraduate medical journal*, vol. 82, no. 964, pp. 84–88, 2006.
- [86] Y. Blinkouskaya, A. Caçoiló, T. Gollamudi, S. Jalalian, and J. Weickenmeier, “Brain aging mechanisms with mechanical manifestations,” *Mechanisms of ageing and development*, vol. 200, p. 111575, 2021.