Local Brain-Age: A U-Net Model

Sebastian G. Popescu [1,2], Ben Glocker [1], David J. Sharp [2,3] and James H. Cole [4,5*]

[1] Biomedical Image Analysis Group, Imperial College London, London, United Kingdom, [2] Computational, Cognitive & Clinical Neuroimaging Laboratory, Imperial College London, London, United Kingdom, [3] Care Research & Technology Centre, UK Dementia Research Institute, London, United Kingdom, [4] Centre for Medical Image Computing, Department of Computer Science, University College London, London, United Kingdom, [5] Dementia Research Centre, Queen Square Institute of Neurology, University College London, London, United Kingdom

We propose a new framework for estimating neuroimaging-derived "brain-age" at a local level within the brain, using deep learning. The local approach, contrary to existing global methods, provides spatial information on anatomical patterns of brain ageing. We trained a U-Net model using brain MRI scans from $n = 3,463$ healthy people (aged 18–90 years) to produce individualised 3D maps of brain-predicted age. When testing on $n = 692$ healthy people, we found a median (across participant) mean absolute error (within participant) of 9.5 years. Performance was more accurate (MAE around 7 years) in the prefrontal cortex and periventricular areas. We also introduce a new voxelwise method to reduce the age-bias when predicting local brain-age "gaps." To validate local brain-age predictions, we tested the model in people with mild cognitive impairment or dementia using data from OASIS3 ($n = 267$). Different local brain-age patterns were evident between healthy controls and people with mild cognitive impairment or dementia, particularly in subcortical regions such as the accumbens, putamen, pallidum, hippocampus, and amygdala. Comparing groups based on mean local brain-age over regions-of-interest resulted in large effects sizes, with Cohen's d values >1.5, for example when comparing people with stable and progressive mild cognitive impairment. Our local brain-age framework has the potential to provide spatial information leading to a more mechanistic understanding of individual differences in patterns of brain ageing in health and disease.

Keywords: brain age, deep learning, dementia, U-net, voxelwise

1. INTRODUCTION

Brain ageing is associated with cognitive decline and an increased risk of neurodegenerative disease, though these effects vary greatly between individuals. Brain atrophy, often measured using structural MRI, is commonly seen in many neurological diseases (Lorenzetti et al., 2009; Chaudhuri, 2013), but also in the normal ageing process. Even hippocampal atrophy, which is often thought to be characteristic of Alzheimer's disease, can be seen in many other neurological and psychiatric conditions, and in normal ageing (Laakso et al., 1996). Evidently, both normal ageing and dementia can affect the same brain regions (Lockhart and DeCarli, 2014). This fact complicates research into the earliest stages of age-related neurodegenerative diseases. Determining whether changes are "normal" and or pathological is challenging. The brain-age paradigm can offer information on whether an individual's brain is changing as expected for their age. The difference between chronological age and "brain-predicted age" obtained from neuroimaging data has been provided insights into the relationship between brain ageing and disease, and may be a useful biomarker

for predicting clinical outcomes (Cole et al., 2018, 2020; Wang et al., 2019; Biondo et al., 2021). For example, in Alzheimer's Disease (AD), patients have previously been shown to have older-appearing brains, and that individuals with mild cognitive impairment (MCI) who had an older-appearing brain were more likely to progress to dementia (Franke et al., 2010, 2012; Gaser et al., 2013; Popescu et al., 2020). However, despite the growing literature employing the brain-age paradigm (Cole et al., 2019; Franke and Gaser, 2019), current approaches tend to generate brain-age predictions at a global level, with a single value per brain image. While some efforts have been made to derive patterns of "feature importance" or similar from brain-age models (Varikuti et al., 2018; Dinsdale et al., 2020; Erramuzpe et al., 2020; Kolbeinsson et al., 2020; Levakov et al., 2020), these patterns are at population-level, and do not apply to the individual.

1.1. Localised Brain Predicted Age

Obtaining a finer-grained picture of brain-ageing patterns for a given brain disease is likely to provide several benefits. Firstly, neuroanatomical patterns should enable inferences to be made about mechanisms underlying the clinical manifestation of the disease. Secondly, better predictive discrimination between clinical groups should be possible, as different groups are likely to be associated with different spatial patterns of age-related brain changes, even in the case where "global" brain-age differences are similar. Thirdly, the local individualised maps should enable fine-grain characterisation of brain changes over time, as the disease progresses or in response to treatment. Finally, spatial patterns of brain-age could be used to discover clinically-relevant subgroups in a data-driven manner, for example using clustering techniques such as principal component analysis, Gaussian mixture models or variational autoencoders.

1.2. Related Work

Limited prior work on local predictions of brain-age are available. Of note, is the early work of Cherubini et al. (2016), who used linear regression models with voxel-level features derived from voxel-based morphometry and diffusion-tensor imaging to demonstrate reasonable prediction results in a small sample of healthy people ($n = 140$). This approach of using a separate linear regression model for each voxel is limited as it does not incorporate contextual information from neighbouring voxels, and is insensitive to non-linear relationships. Other studies have provided local or regional information by training separate models per region e.g., Kaufmann et al., 2019, though again this precludes the incorporation of contextual and global information in the local predictions and is limited to the specific anatomical atlas used to define the brain regions.

Some studies have gone further and extracted "patch" level information on brain-age, subsequently averaging predictions across brain regions to arrive at a global-level prediction (Beheshti et al., 2019; Pawlowski and Glocker, 2019; Bintsi et al., 2020; Gupta et al., 2021). In Bintsi et al. (2020), the authors use a ResNet (He et al., 2016) for each 64^3 3D block, reporting MAE values between 2.16 and 4.19 depending on block origin. While these approaches are promising, the relatively large size of the patch limits spatial resolution which results in less insightful inference in clinical settings. For example, semantic dementia is associated with a relatively localised spatial pattern of atrophy, often the left anterior and middle temporal lobe (Harper et al., 2014; Landin-Romero et al., 2016), which could be overlooked by brain-age prediction models that lack spatial resolution. Alternatively, in Beheshti et al. (2019), the authors introduce a model based on kernel methods introduced in Coupé et al. (2012), whereby they predict the grading at 7^3 voxel patches. However, the authors use Support Vector Regression to aggregate the patch-level results to arrive at a global level prediction and do not provide patch-level results in the cortical regions. Similarly, Gupta et al. (2021) proposed a slice-level MRI encoding network, followed by an aggregation method to obtain global-level predictions. Likewise, the authors do not provide results at finer grained scales.

1.3. Contributions

The goal of this work was to develop a model to accurately predict chronological age at the local level in healthy people, by incorporating voxelwise information using deep learning. U-Nets (Ronneberger et al., 2015), which are typically used for tumor (Feng et al., 2020) or organ (Vesal et al., 2019) segmentation, provide an excellent framework for voxelwise predictions, as their specific architecture enables the inclusion of contextual spatial information into individual predictions. Here, we introduce a deep learning algorithm that is trained to predict localised brain-age, producing high-resolution maps of brain-predicted age differences (brain-PAD maps) covering the entire brain. We hypothesised that brain-PAD in healthy people would be centred on zero and would smoothly vary across regions of the brain. We further hypothesised that people with MCI and dementia patients would see higher brain-PAD values in regions previously reported to dementia-related atrophy. We provide an in-depth analysis of the structural differences seen in people with MCI and AD patients. We provide a means to reduce the so-called "age-bias" in brain-PAD maps and examine the reliability of local brain-age predictions, both within and between scanners.

2. METHODS

2.1. Participants

To train, test and validate our local brain-age model, we collated multiple datasets comprising T1-weighted MRI brain scans. All included datasets were from studies that had been reviewed and approved by the local ethics committees and all participants provided informed consent. All participants were included, notwithstanding exclusions due to failure during quality control after pre-processing. All data were from publicly accessible databases. The **Supplementary Material** includes all links to access the respective databases, alongside chronological age histograms for each datasets (see **Supplementary Figure 1**).

2.1.1. Brain-Age Healthy Controls (BAHC)

This dataset comprises 2001 3D T1-weighted MRI scans from healthy individuals with a male/female ratio of 1,016/985, with a mean age of 36.95 ± 18.12, aged 18–90 years. These data

are an amalgam of 14 separate publicly-available datasets, as used in our previous brain-age research (Cole et al., 2017) (see **Supplementary Table 1** for full details).

2.1.2. Dallas Lifespan Brain Study (DLBS)

This is a major effort designed to understand the antecedents of preservation and decline of cognitive function at different stages of the adult lifespan, with a particular interest in the early stages of a healthy brain's march toward Alzheimer's Disease. For our purpose we have selected solely the T1-weighted MRI scans, totalling $n = 315$ healthy participants aged 18–89 years, with a mean age of 54.61 ± 20.09 and male/female ratio of 117/198. All participants were scanned on a single 3T Philips Achieva scanner equipped with an 8-channel head coil. High-resolution anatomical images were collected with a sagittal T1-weighted 3D MPRAGE sequence (TR = 8.1 ms, TE = 3.7 ms, flip angle = 12°, FOV = 204 × 256, slices = 160, voxel size = 1 mm isotropic. More information can be found at http://fcon_1000. projects.nitrc.org/indi/retro/dlbs.html.

2.1.3. Cambridge Centre for Ageing and Neuroscience (Cam-CAN)

This dataset is part of larger project which is trying to use epidemiological, behavioural and neuroimaging data to understand how individuals can best retain cognitive abilities into old age. All participants were scanned with 3 T Siemens TIM Trio scanner with a 32-channel head coil. The dataset consists of $n = 652$ T1-weighted MRI 3D MPRAGE (TR = 2,250 ms, TE = 2.99 ms, TI = 900 ms, flip angle = 9°, FOV = 256 × 240, slices = 192, voxel size = 1 mm isotropic, scan duration = 1 h) from participants aged 18–88 years, with a mean age of 54.29 ± 18.59 and a male/female ratio of 322/330. More information can be found at https://www.cam-can.org/.

2.1.4. Southwest University Adult Lifespan Dataset (SALD)

This comprises a large cross-sectional sample ($n = 494$; age range = 19–80 years; mean age 45.18 ± 17.44; male/female ratio of 187/307) undergoing a multi-modal (structural MRI, resting state fMRI, and behavioural) neuroimaging. Only T1-weighted 3D MPRAGE (TR = 1,900 ms, TE = 2.52 ms, TI = 900 ms, flip angle = 90°, FOV = 256 × 256, slices = 176, voxel size = 1 mm isotropic) were used here. The goals of the SALD are to give researchers the opportunity to map the structural and functional changes the human brain undergoes throughout adulthood and to replicate previous findings. More information can be found at http://fcon_1000.projects.nitrc.org/indi/retro/sald.html.

2.1.5. Wayne State

The Wayne State longitudinal data set for the Brain Aging in Detroit Longitudinal Study, comprises 200 healthy individuals, with $n = 302$ total anatomical scans across two waves of data collection and mean age of 53.94 ± 15.58, with a male/female ratio of 37/77. All the participants were screened by the local research centres to be free from neurological or psychiatric disorders according to established protocols. All of the neuroimaging data were acquired using a 4T Varian Scanner

(Bruker Biospin, Ettlingen, Germany) with a 3D T1-weighted MPRAGE sequence (TR = 1,600 ms, TE = 4.83 ms, TI = 800 ms, flip angle = 8°, FOV = 256 × 256, voxel size = 0.7 × 0.7 × 1.34 mm). More information can be found at http://fcon_1000. projects.nitrc.org/indi/retro/wayne_11.html.

2.1.6. Within-Scanner Reliability Dataset

Here we used data from the Imperial College London project, STudy Of Reliability of MRI (STORM). The study comprises of 20 participants with a male-female ratio of 12/8, with a mean age at the first scan undertaken of 34.05 ± 8.71. The participants were scanned for the second time at an average distance of 28.35 ± 1.09 days. All participants were free from any neurological or psychiatric disorders. MPRAGE were acquired using a Siemens Verio 3T scanner (TR = 2,300 ms, TE = 2.98 ms, TI = 900 ms, flip angle = 9°, FOV = 256 × 256, slices = 160, voxel size = 1 mm isotropic).

2.1.7. Scanner Calibration Dataset

This study included 11 participants scanned in two different centres, mean age at first scan of 30.88 ± 6.16 and with a male/female ration of 7/4. The two scanning sites were at Imperial College London, where a Siemens Verio 3T scanner was used to acquire MPRAGE (TR = 2,300 ms, TE = 2.98 ms, TI = 900 ms, flip angle = 9°, FOV = 256 × 256, slices = 160, voxel size = 1 mm isotropic), whereas a Philips Ingenia 3T scanner was used at the Academic Medical Center Amsterdam to acquire sagittal Turbo Field Echo (TR = 6.6 ms, TE = 3.1 ms, flip angle = 9°, FOV = 270 × 270, slices = 170, voxel size = 1.1 × 1.1 × 1.2 mm). The mean interval between scans was 68.17 ± 92.23 days.

2.1.8. Open Access Series of Imaging Studies (OASIS3)

This is a retrospective compilation of data for >1,000 participants that were collected across several ongoing projects through the WUSTL Knight ADRC over the course of 30 years. Participants include $n = 609$ cognitively normal adults and n=489 individuals at with MCI or dementia ranging in age from 42 to 95 years. Using Clinical Dementia Rating scale (CDR) scores, we classified participants as healthy control (HC), stable MCI, progressive MCI or AD, as detailed in **Table 1**. Follow-up CDR scores used to define MCI status were from at least 3 years after baseline assessments. We excluded scans which did not pass quality standards after pre-processing pipeline. MPRAGE was collected on Siemens TIM Trio 3T (TR = 2,400 ms, TE = 3.08 ms, TI = 1, flip angle = 8°, FOV = 256 × 256, voxel size = 1 mm isotropic). Further information can be found at https://www.oasis-brains. org/.

2.1.9. Australian Imaging Biomarkers and Lifestyle Study of Aging (AIBL)

This study is a study to discover which biomarkers, cognitive characteristics, and health and lifestyle factors determine subsequent development of symptomatic Alzheimer's Disease (AD) (https://aibl.csiro.au/). The dataset contained $n = 198$ participants with Clinical Dementia Rating scale scores, detailed in **Table 2**. MPRAGE was collected on 1.5 T Siemens Avanto (TR

Characteristics	HC	sMCI	pMCI	AD
	(n = 128)	(n = 29)	(n = 29)	(n = 78)
Males/females, n	70/58	15/14	18/11	33/45
Age, mean (SD) years	68.14 (9.40)	76.44 (6.81)	75.72 (7.68)	75.02 (8.90)
Age, range years	42.66–97.11	59.2–94.44	49.38–93.93	50.35–95.58
Baseline CDR	0.0	0.5	0.5	≥ 1.0
Follow-up CDR	0.0	0.5	≥ 1.0	-

CDR, Clinical Dementia Rating scale; HC, Healthy Controls; pMCI, progressive MCI; sMCI, stable MCI; AD, Alzheimer's Disease.

Characteristics	HC	sMCI	pMCI	AD
	(n = 83)	(n = 64)	(n = 20)	(n = 31)
Males/Females, n	38/45	36/26	11/9	17/14
Age, mean (SD) years	67.28 (9.70)	75.83 (8.85)	71.86 (9.21)	74.23 (10.21)
Age, range years	60.11–86.45	62.54–87.86	54.66–81.24	61.75–87.89
Baseline CDR	0.0	0.5	0.5	≥ 1.0
Follow-up CDR	0.0	0.5	≥ 1.0	-

CDR, Clinical Dementia Rating scale; HC, Healthy Controls; pMCI, progressive MCI; sMCI, stable MCI; AD, Alzheimer's Disease.

= 1,900 ms, TE = 2.13 ms, TI = 900 ms, flip angle = 9°, FOV = 240 × 256, voxel size = 1 × 1 × 1.2 mm isotropic).

2.2. Data Pre-processing

All T1-weighted brain MRI scans were pre-processed using the Statistical Parametric Mapping (SPM12) software package (https://www.fil.ion.ucl.ac.uk/spm/software/spm12/). This entailed tissue segmentation into grey matter (GM) and white matter (WM), followed by a non-linear registration procedure using the DARTEL algorithm (Ashburner, 2007) to the Montreal Neurological Institute 152 (MNI152) space, subsequently followed by resampling to 1.5 mm³ with a 4 mm smoothing kernel.

2.3. Statistical Analysis
2.3.1. Inferential Statistics

Welch's t-test was used to compare groups based on voxel, regional and global brain-PAD values. Welch's t-test is an alternative to the standard student's t-test when the two populations to be compared have uneven variance and optionally also uneven sample size. The t statistics to test whether the populations means is given by:

$$t = \frac{\bar{X}_1 - \bar{X}_2}{s_{\bar{\Delta}}} \quad (1)$$

where $s_{\bar{\Delta}} = \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}$ and s_i^2 represents the unbiased estimator of the variance of a respective sample with n_i participants. To use the test statistics for significance testing, the degrees of

freedom of the associated Student's t-distribution is given by the Welch-Satterthwaite equation:

$$df = \frac{\left(\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}\right)^2}{\frac{(s_1^2/n_1)^2}{n_1-1} + \frac{(s_2^2/n_2)^2}{n_2-1}} \quad (2)$$

2.3.2. Effect Size Estimates

To quantify effect sizes when comparing different disease groups we used the standardised effect size Cohen's d:

$$d = \frac{m_1 - m_2}{\sqrt{\frac{(n_1-1)*s_1+(n_2-1)*s_2}{n_1+n_2-2}}} \quad (3)$$

where m_k is the mean, s_k represents the variance, whereas n_k defines the number of participants within group k. The purpose of this method is to quantify the size of the difference, allowing us to decide if the difference is meaningful.

2.3.3. Intraclass Correlation Coefficient

The intraclass correlation coefficient (ICC) is used to test the reproducibility of a certain quantitative measurement made by a specified number of observers which rate the same participant. The original formula is given as follows:

$$r = \frac{1}{Ns^2} \sum_{n=1}^{N} (x_{n,1} - \tilde{x})(x_{n,2} - \tilde{x}) \quad (4)$$

where $\tilde{x} = \frac{1}{2N} \sum_{n=1}^{N}(x_{n,1} + x_{n,2})$ and $s_2 = \frac{1}{2N}\{\sum_{n=1}^{N}(x_{n,1} - \tilde{x})^2 + \sum_{n=1}^{N}(x_{n,2} - \tilde{x})^2\}$.

Here, we used ICC[2,1] as defined by Shrout and Fleiss (1979). The interval of values ranges from $[-1, 1]$ with values closer to 1 denoting that the observers (e.g., MRI scans or scanners) agree with each other.

2.4. Study Design

In this subsection, we summarise the design of our experiments and which datasets are used in each step.

- BAHC, CamCan, Dallas and SALD were used as training and validation sets (standard 80/20 split) for training the local brain-age U-net model.
- All Wayne State participants and healthy participants from OASIS3 and AIBL were used at testing time. We calculated mean absolute error (MAE) values globally (by averaging across voxel-level brain-predicted age for each participant) and at voxel level, alongside the Pearson's correlation coefficient between chronological age brain-predicted age at both global and voxel level.
- Within-scanner reliability and Scanner calibration datasets were used as test sets to compute voxel-level ICC values to assess the reliability of local brain-age when the same participant is scanned in two different scanners, respectively one the same scanner with short time interval.
- Using subcortical and cortical ROIs from the Harvard-Oxford structural brain atlas, we obtained brain tissues volumes
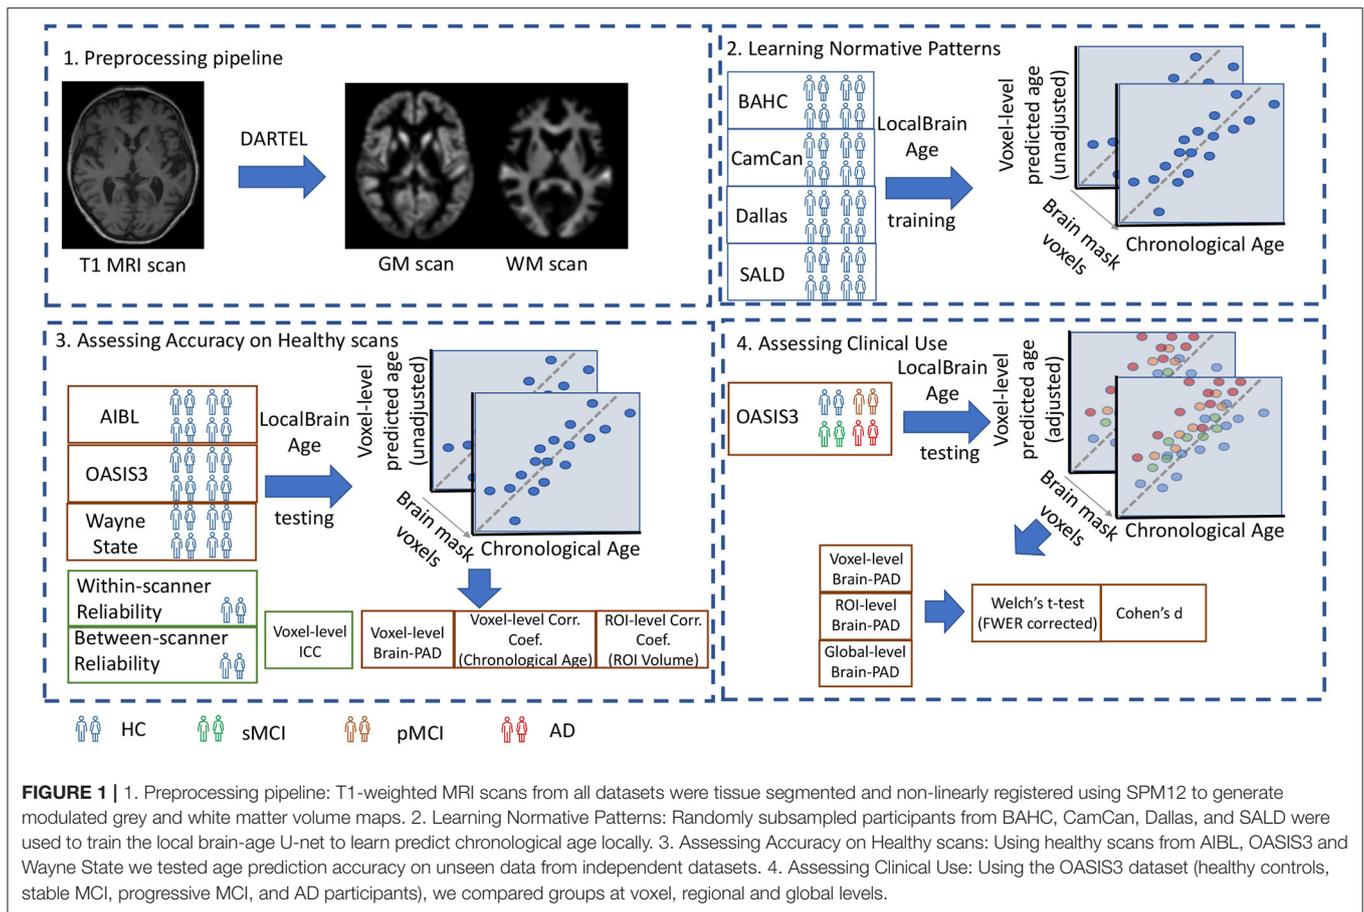
FIGURE 1 | 1. Preprocessing pipeline: T1-weighted MRI scans from all datasets were tissue segmented and non-linearly registered using SPM12 to generate modulated grey and white matter volume maps. 2. Learning Normative Patterns: Randomly subsampled participants from BAHC, CamCan, Dallas, and SALD were used to train the local brain-age U-net to learn predict chronological age locally. 3. Assessing Accuracy on Healthy scans: Using healthy scans from AIBL, OASIS3 and Wayne State we tested age prediction accuracy on unseen data from independent datasets. 4. Assessing Clinical Use: Using the OASIS3 dataset (healthy controls, stable MCI, progressive MCI, and AD participants), we compared groups at voxel, regional and global levels.

(mm^3) for each ROI, alongside ROI-level brain-PAD values which were computed by first averaging voxel-level "brain-predicted age" inside an ROI, then subtracting the participant's chronological age. We calculated a Pearson's correlation coefficient for each ROI.

- OASIS3 was used at testing time to assess the sensitivity of local brain-age to differences in brain structure between groups. For each participant we computed an mean across voxels global brain-PAD (adjusted for age bias using method described in section 2.6), which we use then to perform Welch's t-test between disease groups, correcting for multiple comparisons using the Bonferroni method. The same method was then applied to local-level brain-PAD values, which were pooled to create a "population" of brain-PAD values at voxel-level. To assess effect sizes of between-group differences we calculated Cohen's d coefficient at global and voxel levels. Lastly, to assess regional differences between disease groups, we used the Harvard-Oxford cortical and subcortical structural atlas, which contains 48 cortical and 21 subcortical structural ROI. We then calculated differences in mean local brain-PAD per ROI between groups using Welch's t-test (Bonferroni corrected), again computing Cohen's d effect sizes. For all experiments in this part we have selected subjects above 60 years old from the healthy controls so as to have a similar
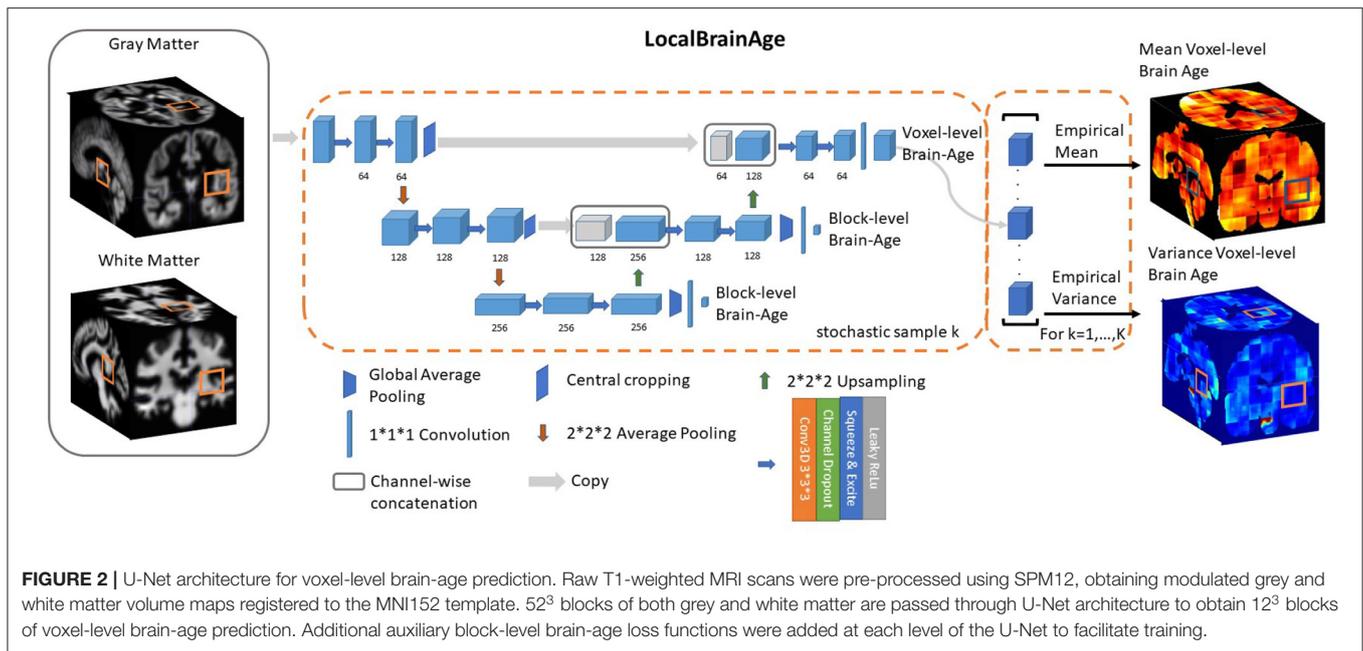
chronological age distribution in relation to the groups with varying degrees of cognitive impairment.

A visual overview of the study design is portrayed in **Figure 1**.

2.5. Local "Brain-Age" Prediction

We used a fully convolutional neural network (CNN) inspired by the U-Net architecture introduced in Ronneberger et al. (2015). Our network architecture is illustrated in **Figure 2**. Input images were the output from SPM12 pre-processing, representing voxelwise maps of GM and WM volume. These images were split into overlapping 3-dimensional blocks of size 52^3 voxels. The convolutional layers in our network used an isotropic 3x3x3 filter, convolved over the input image after which element-wise multiplication with the filter weights and subsequent summation was performed at each location. Subsequently, to allow for non-linear modelling, we passed the obtained values through an "activation function"; we used leaky rectified linear units with alpha = 0.2, more precisely $max(x, 0) + min(x * \alpha, 0)$, thus allowing a small, non-zero gradient when the unit is not active.

The convolution operation is also controlled by its stride, which is how many pixels/voxels are skipped after every element-wise weight multiplication and summation. We set the stride equal to 1.

FIGURE 2 | U-Net architecture for voxel-level brain-age prediction. Raw T1-weighted MRI scans were pre-processed using SPM12, obtaining modulated grey and white matter volume maps registered to the MNI152 template. 52^3 blocks of both grey and white matter are passed through U-Net architecture to obtain 12^3 blocks of voxel-level brain-age prediction. Additional auxiliary block-level brain-age loss functions were added at each level of the U-Net to facilitate training.

Downsampling increases the effective field of view or "receptive field" of layers higher in the hierarchy. For the downsampling part of the U-Net we used at each scale two consecutive 3D $3 \times 3 \times 3$ filter kernels with an initial number of channels = 64, which get multiplied by 2 as we progress down the downsampling path. For downsampling we used $2 \times 2 \times 2$ average pooling.

For the upsampling part of the U-Net we inverted the downsampling architecture, with the downsampling layers being replaced by $2 \times 2 \times 2$ upsampling layers. At each convolution we used a squeeze-and-excite unit. Squeeze & Excite networks were introduced in Hu et al. (2018) and can be viewed as computationally less intensive method of performing attention over the channels of a given feature block. Finally, at the end of the network we obtain predictions over 12^3 voxels blocks.

Besides the voxel-level mean absolute error cost function on the output layer we introduced two additional cost functions at the two other scales of the architecture. We applied global average pooling followed by a dense layer to predict brain-age at block-level. The loss function can be expressed as follows:

$$L = \sum_{i=1}^{M} | y_i - \tilde{y}_i^{voxel} | + \sum_{b=1}^{2} \alpha_b \sum_{i=1}^{M} | y_i - \tilde{y_{i,b}} | \quad (5)$$

where $\tilde{y_{i,b}}$ represent the block-level brain-age prediction of the b-th block. During training, we observed that the addition of these auxiliary loss functions helped stabilise the learning process. During training, α_1 and α_2 are progressively decreased so that the gradients will exclusively flow from the voxel-level predictions after 50,000 training iterations. We used Adam (Kingma and Ba, 2014) for optimising our loss function with a learning rate set to 0.0001. We trained our model for 500,000 iterations, with a minibatch size of 32 (gradient averaging over four splits). We split our healthy participant datasets into training (80%) and validation (20%) sets and the stopping criteria was set based on a visual inspection of the validation loss reaching a plateau. The model was implemented in Tensorflow (Abadi et al., 2016).

2.6. Removing Bias in Predictions

Subtracting chronological age from estimated brain age provides a measure of the difference between an individual's predicted age and chronological age, also known as the brain-age "gap," brain-predicted age difference (brain-PAD) or brain-age "delta." A so-called "regression dilution" has been commonly observed in brain-age prediction algorithms, caused by noise in the neuroimaging features leading to a greater under- or over-estimate of age, the further away a sample is from the training set mean age. In other words, this effect results in the systematic under-estimation of brain-predicted age for older participants and over-estimation for younger participants, which increases as model performance decreases.
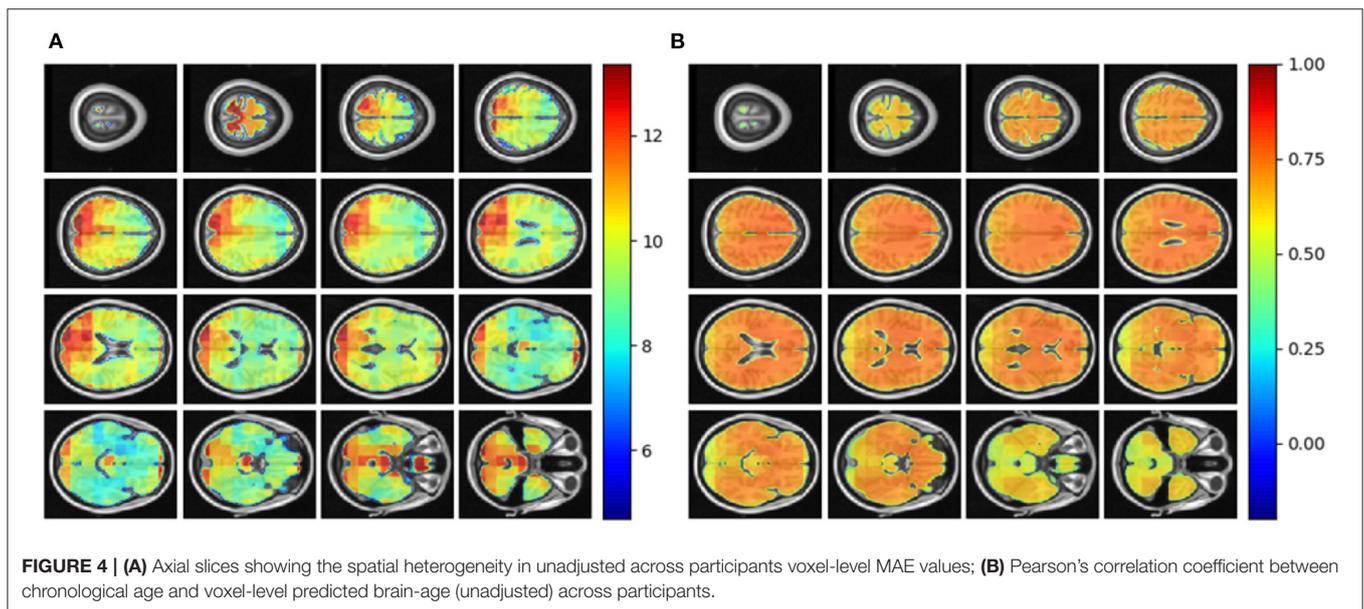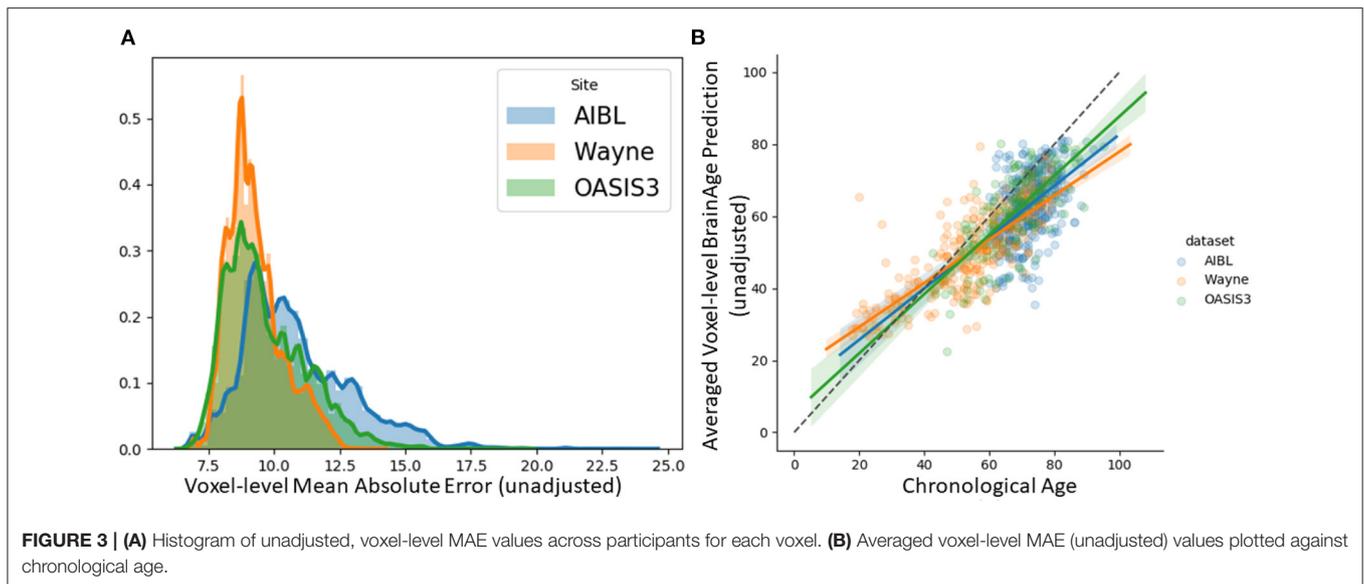
2.6.1. Global-Level

Broadly speaking, two approaches to account for this effect have been reported:

$$\Delta = \alpha * Age + \beta \quad (6)$$

where Δ is the brain-age delta of a group of participants from an external dataset that is used specifically for adjusting the bias. α and β are the parameters of a linear regression with the covariate *Age* representing chronological age.

Then, to obtain the bias-adjusted age we have the following equation:

$$\tilde{\Delta} = \Delta - \alpha * Age + \beta \quad (7)$$

FIGURE 3 | (A) Histogram of unadjusted, voxel-level MAE values across participants for each voxel. (B) Averaged voxel-level MAE (unadjusted) values plotted against chronological age.

FIGURE 4 | (A) Axial slices showing the spatial heterogeneity in unadjusted across participants voxel-level MAE values; (B) Pearson's correlation coefficient between chronological age and voxel-level predicted brain-age (unadjusted) across participants.

Another approach involves using the brain-predicted age in the linear regression, more specifically:

$$\tilde{\Delta} = \Delta - \alpha * \tilde{y} + \beta \tag{8}$$

where $\tilde{y}$ denotes the brain-predicted age. de Lange and Cole (2020) showed that using either formulation results in the same statistical outcome in comparing different disease groups. The authors also argue against using bias adjusted predictions at testing time to assess overall accuracy of the model. However, this standard method used for global-level brain-age prediction did not succeed in de-biasing our predictions at voxel-level. Additional results using this approach are shown in the **Supplementary Material** (see **Supplementary Figure 2**).

2.6.2. Voxel-Level

Here, we used a separate small batch ($n = 200$) of participants randomly selected from the healthy participant datasets (BAHC, CamCan, Dallas, SALD), who were not included in the training or validation set. We obtained testing time predictions for these participants and calculated their voxel-level brain-age delta $\Delta_{i,v}$, where i indicates the i-th participant and v the v-th voxel. We then binned these participants based on their chronological age (5-year intervals, expect the first being between 18 and 25 years). Then, for each bin b we calculated the average voxel-level brain-age delta for that respective bin, which we denote as $\Delta_{b,v}$. This value will represents the average brain-age delta for that voxel given the chronological age interval. Subsequently, to de-bias the voxel-level brain-age delta for

a new participant (e.g., from testing set), $\Delta_{j,v}$ we used the following formula:

$$\tilde{\Delta_{j,v}} = \Delta_{j,v} - \Delta_{b,v} \qquad (9)$$

This method was used in subsequent analysis where indicated.

3. RESULTS

3.1. Local Brain-Age Model Performance in Independent Healthy Test Datasets

We tested the local brain-age model on healthy participants combined from the OASIS3 ($n = 128$), AIBL ($n = 83$) and

TABLE 3 | Pearson's correlation coefficient (r) for different subcortical ROIs from the Harvard-Oxford atlas between ROI-level brain tissue volume and ROI-level brain-PAD.

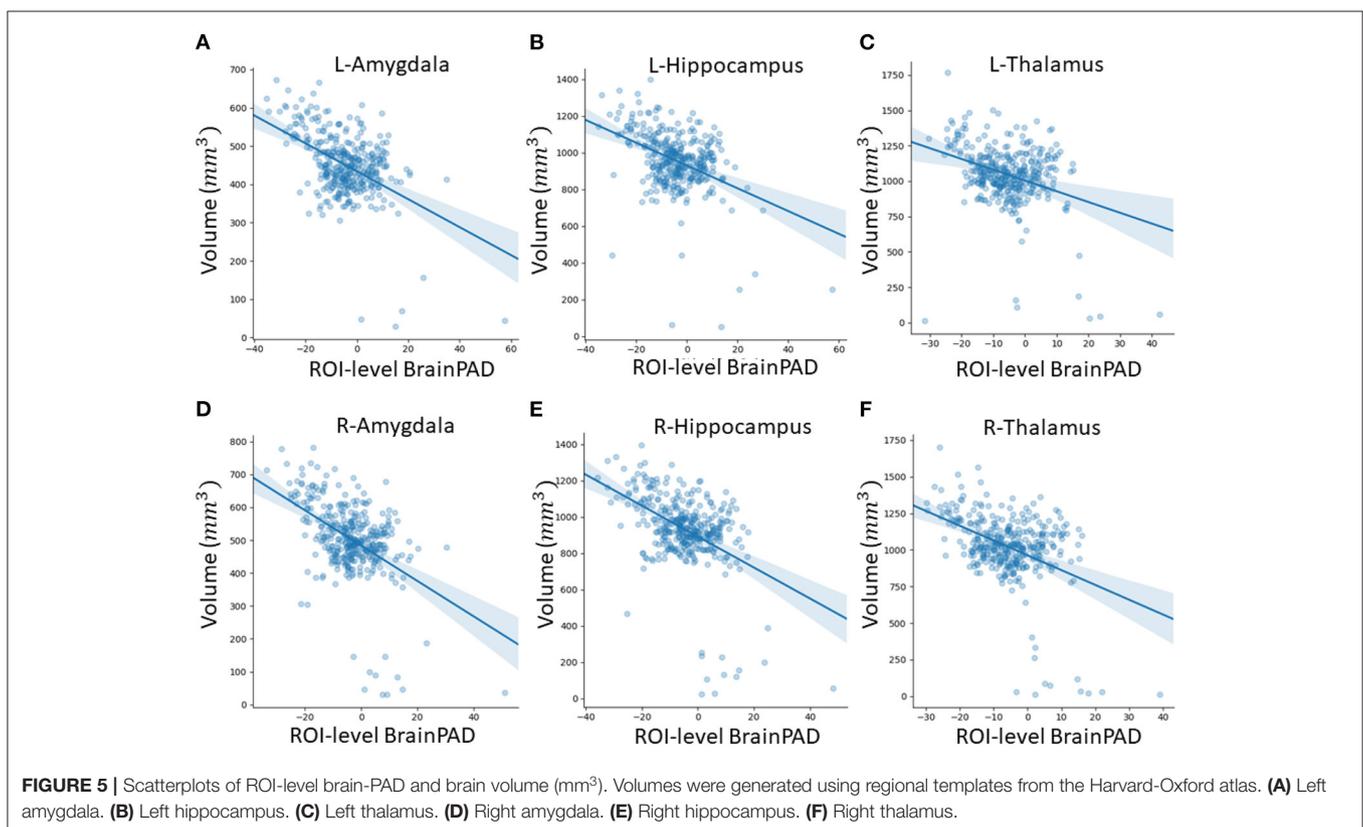| Subcortical ROI name	Left hemisphere	Right hemisphere
Amygdala	−0.48	−0.48
Caudate	−0.13	−0.19
Hippocampus	−0.42	−0.46
Pallidum	−0.19	−0.17
Putamen	−0.32	−0.32
Thalamus	−0.34	−0.43
Accumbens	−0.23	−0.30

Wayne State ($n = 200$) datasets. When looking at voxel-level MAE (unadjusted) values across the brain mask (**Figure 3A**), we mean values for AIBL of 10.84 ± 2.05 (median 10.43) years, for Wayne State 9.28 ± 1.05 (median 9.08) years, and OASIS3 9.70 ± 1.53 (median 9.33) years. The voxel-level MAE (unadjusted) values of the model varied in different brain regions. We observed lower values across the different sites in the prefrontal cortex and subcortical regions and higher MAE in the occipital lobe, cerebellum and brainstem (**Figure 4A**). The correlation coefficient between chronological age and voxel-level predicted brain-age (unadjusted) across participants showed similar patterns, with higher values obtained in the prefrontal cortex and subcortical regions (**Figure 4B**). We obtained a global-level MAE (unadjusted) value by averaging the voxel-level brain-predicted ages across voxels for a given participant. For AIBL, we obtain an average of 10.23 ± 7.08 years (median = 8.86; $r = 0.47$), for Wayne State 8.09 ± 6.08 years (median = 6.92; $r = 0.78$), respectively for OASIS3 we get an average of 8.08 ± 6.40 years (median = 6.26; $r = 0.72$). **Figure 3B** shows the mean voxel-level brain-predicted age for each participant against chronological age.
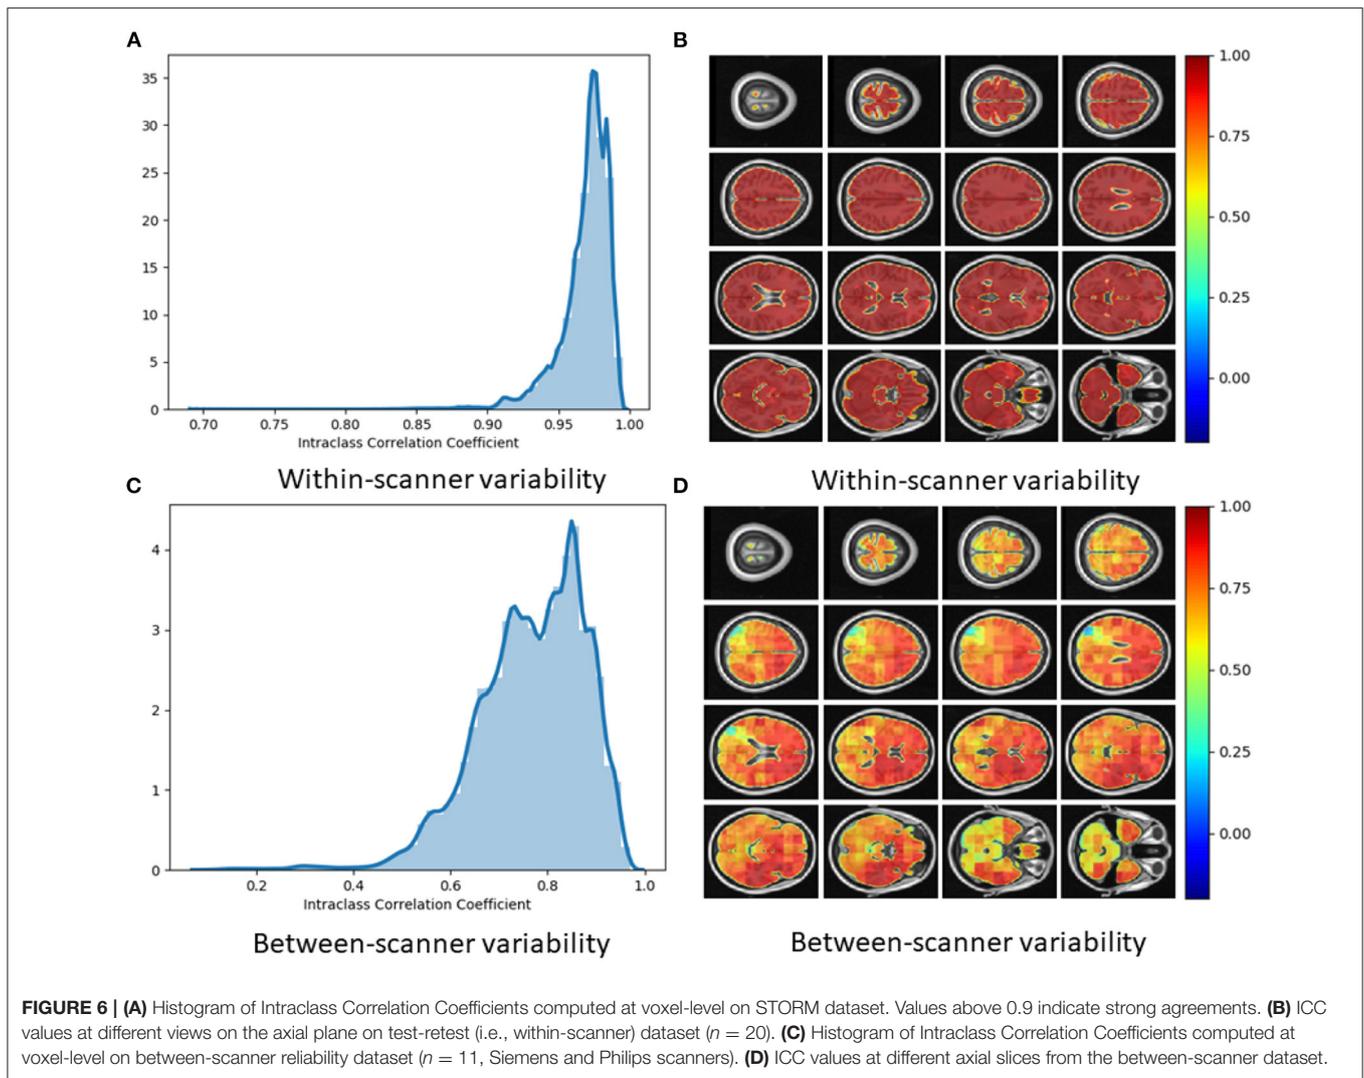
3.2. Regional Brain Volumes and Regional Brain-PAD in Healthy Individuals

In this subsection, we explored ROI-level results, based on the Harvard-Oxford atlas, by effectively averaging voxel-level brain-age predictions within the respective ROI. We include cortical region results in the **Supplementary Material**

FIGURE 5 | Scatterplots of ROI-level brain-PAD and brain volume (mm^3). Volumes were generated using regional templates from the Harvard-Oxford atlas. (A) Left amygdala. (B) Left hippocampus. (C) Left thalamus. (D) Right amygdala. (E) Right hippocampus. (F) Right thalamus.

FIGURE 6 | (A) Histogram of Intraclass Correlation Coefficients computed at voxel-level on STORM dataset. Values above 0.9 indicate strong agreements. **(B)** ICC values at different views on the axial plane on test-retest (i.e., within-scanner) dataset ($n = 20$). **(C)** Histogram of Intraclass Correlation Coefficients computed at voxel-level on between-scanner reliability dataset ($n = 11$, Siemens and Philips scanners). **(D)** ICC values at different axial slices from the between-scanner dataset.

(**Supplementary Table 2**). From **Table 3**, we can observe that the amygdala, hippocampus and thalamus have the strongest negative Pearson's correlation coefficients between ROI-level brain-PAD and ROI-level volumes (**Figure 5**).

3.3. Reliability of Local Brain-Age

Using voxel-level brain-age values for the Within-scanner (test-retest) and Scanner calibration (between-scanner) datasets, ICC was calculated per voxel. Test-retest reliability was very high with the vast majority of voxels having ICC < 0.90 [median ICC = 0.98, (95% confidence intervals 0.92, 0.99)] (**Figure 6A**). This indicated very high reliability of local brain-age predictions within the same scanner. We observed comparatively lower ICC values at the extremities of the brain, see **Figure 6B**. This could be due to residual misregistration or partial volume effects. Between-scanner reliability was lower, with median voxel-level ICC = 0.76 (95% confidence intervals 0.36, 0.93) (**Figure 6C**). Interestingly, the pattern of ICC varied across the brain, with higher values

observed in the prefrontal cortex and lower values in more inferior regions, particularly the brainstem and cerebellum (**Figure 6D**).

3.4. Local Brain-Age Differences Between Healthy Controls, People With MCI, and Dementia Patients

We examined patterns of local and global brain-age in people with MCI and dementia patients using cross-sectional data from OASIS3. Firstly, we investigated if the global-level (i.e., averaged within participant) brain-predicted age (adjusted) corresponds to previously reported differences from models that directly predict global brain age. We averaged voxel-level brain-age (adjusted) across voxels per individual to generate an adjusted global-level brain age and then calculate global-level brain-PAD. Global-level brain-PAD (adjusted) mean ($\pm$ standard deviation) values were: -0.65 ± 7.46 (median = 0.95) years for healthy controls, 3.07 ± 4.29 (median = 2.83) years for stable MCI (sMCI), 5.77 ± 5.41
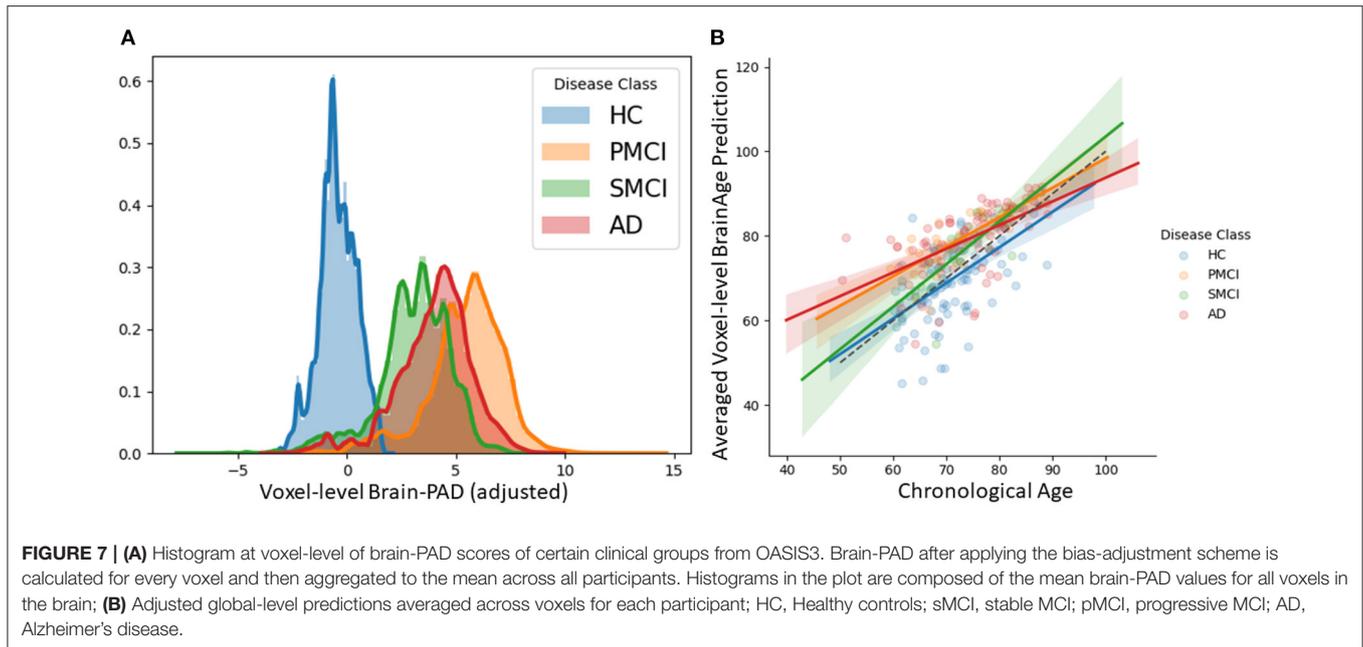
FIGURE 7 | (A) Histogram at voxel-level of brain-PAD scores of certain clinical groups from OASIS3. Brain-PAD after applying the bias-adjustment scheme is calculated for every voxel and then aggregated to the mean across all participants. Histograms in the plot are composed of the mean brain-PAD values for all voxels in the brain; **(B)** Adjusted global-level predictions averaged across voxels for each participant; HC, Healthy controls; sMCI, stable MCI; pMCI, progressive MCI; AD, Alzheimer's disease.

TABLE 4 | Group comparisons of brain-age in OASIS3 participants.

Disease groups	HC	sMCI	pMCI	AD
HC	-	−1284.67 (<0.001)	−2095.58 (<0.001)	−1606.19 (<0.001)
sMCI	−2.18 (0.0369)	-	−684.57 (<0.001)	272.44 (<0.001)
pMCI	−3.67 (0.0007)	-1.44 (0.1611)	-	−411.23 (<0001)
AD	−3.64 (0.0004)	0.83 (0.4141)	−0.90 (0.3714)	-

Upper triangle: Voxel-level brain-age comparisons using paired Welch's t-test results [t statistics value (p-value)] between disease groups. Lower triangle: Global-level brain-age comparisons using independent Welch's t-test results [t statistics value (p-value)] between disease groups.
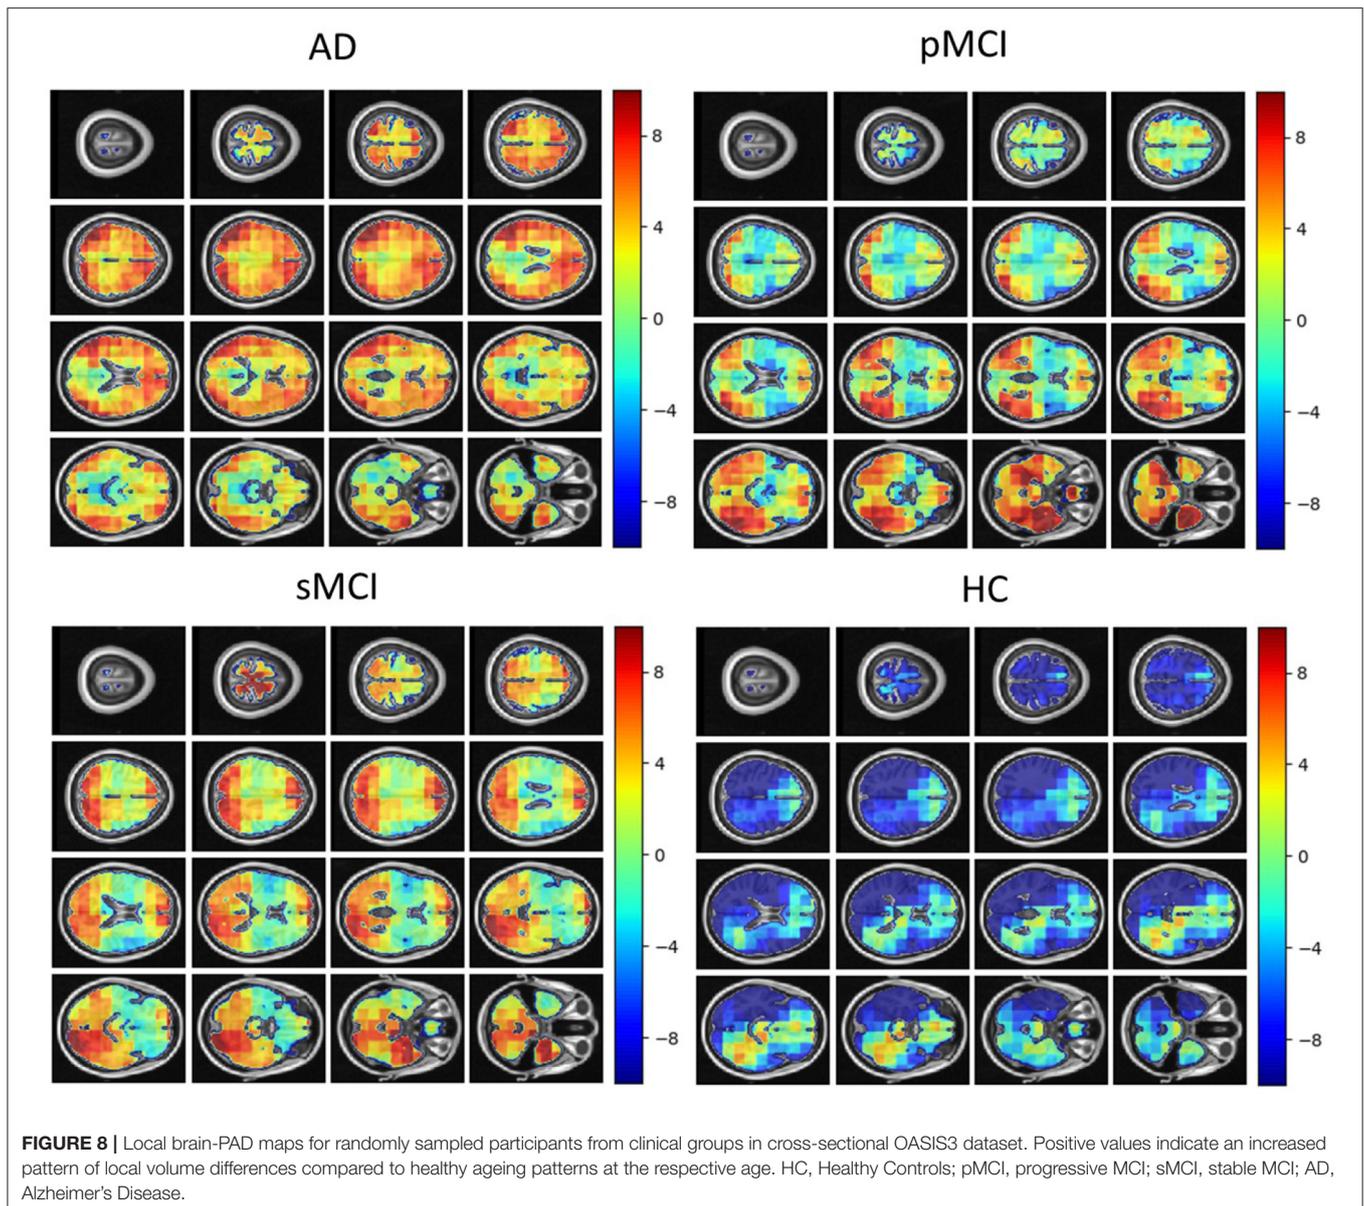
(median = 4.94) years for progressive MCI (pMCI) and 4.34 ± 6.78 (median = 4.63) years for AD patients.

We then assessed the significance of group differences using global-level brain-PAD values by performing independent two-sample Welch's t-tests, finding significant differences between cognitively impaired groups and healthy controls in all cases (HC-AD $t = -3.64$, $p = 0.0004$, df = 88.96, Cohen's $d = -0.70$; HC-sMCI $t = -2.18$, $p = 0.0369$, df = 29.29, Cohen's $d = -0.53$; HC-pMCI $t = -3.67$, $p = 0.0007$, df = 38.66, Cohen's $d = -0.92$). Comparisons between stable and progressive MCI patients and with AD patients were not significant: sMCI-pMCI $p = 0.161$, $t = -1.44$, df = 26.44, Cohen's $d = -0.54$, AD-sMCI $t = 0.83$, $p = 0.414$, df = 20.64, Cohen's $d = 0.19$, AD-pMCI $t = -0.90$, $p = 0.3714$, df = 28.51, Cohen's $d = -0.21$.

Next, we examined local brain-PAD, summarising across all voxels within group. The mean voxel-level brain-PAD (adjusted)

values were: healthy controls $= -0.39 \pm 0.85$ (median $= -0.44$) years, sMCI $= 3.07 \pm 1.67$ (median $= 3.266$) years, pMCI $= 5.45 \pm 1.74$ (median $= 5.663$) years for pMCI, AD patients $= 4.01 \pm 1.71$ (median $= 4.229$) years (**Figure 7**). We then compared groups based on these voxel-level brain-PAD values (adjusted) (**Table 4** upper triangular part) using paired Welch's t-test. Likewise, differences between participants MCI or dementia and healthy controls were significant (HC-sMCI $t = -1284.67$, $p < 0.0001$, df = 723943.40, Cohen's $d = -2.60$; HC-pMCI $t = -2095.58$, $p < 0.0001$, df = 707525.16, Cohen's $d = -4.25$; HC-AD $t = -1606.19$, $p < 0.0001$, df = 971564.04, Cohen's $d = -3.25$). In contrast to the global-level results (lower triangle in **Table 4**), all pairwise differences between groups with MCI or dementia were significant (sMCI-pMCI $t = -684.57$, $p < 0.001$, df = 970331.23, Cohen's $d = -1.38$; sMCI-AD $t = 272.44$, $p < 0.001$, df = 971564.04, Cohen's $d = 0.55$; pMCI-AD $t = -411.23$, $p < 0.001$, df = 971487.0, Cohen's $d = -0.83$).

Individual local brain-age maps from example participants are shown in **Figure 8**. From **Figure 9**, we can observe that local brain-age model is able to detect group differences across the whole brain when comparing healthy controls with AD patients or comparing the pMCI group with the sMCI group (after correction for multiple comparisons). Other group contrasts showed more varied spatial patterns of significant voxels. From **Figure 10**, we can observe that the largest differences are in the temporal lobe and subcortical regions when comparing AD patients to healthy controls. For a more in-depth look at differences between disease groups, we extended the analysis to investigate atlas-based subcortical ROIs. The nucleus accumbens, putamen, pallidum, and hippocampus were the most discriminative ROIs in terms of Cohen's d scores both for separating AD patients from healthy controls and stable from
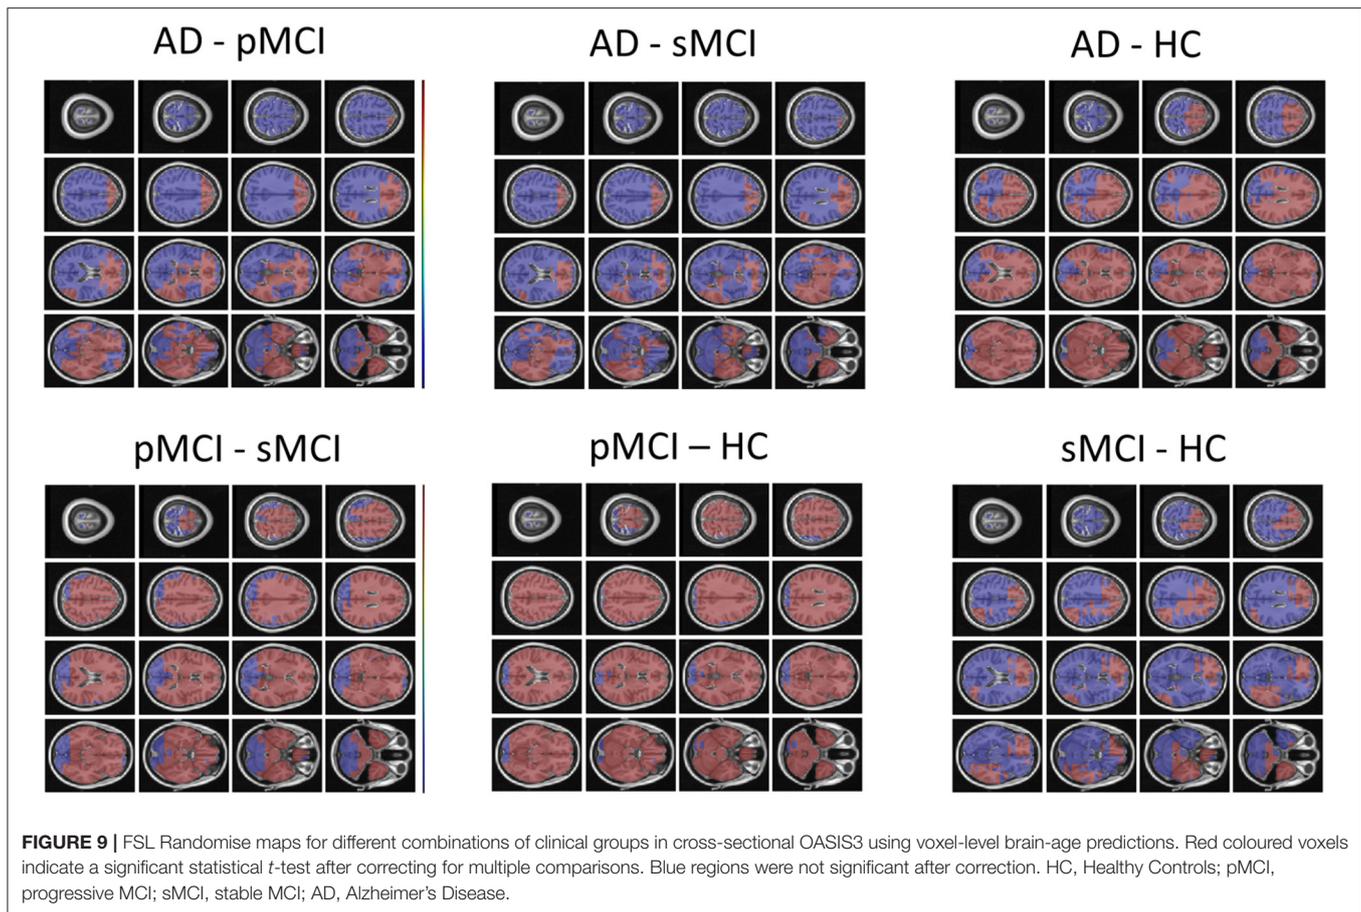
FIGURE 8 | Local brain-PAD maps for randomly sampled participants from clinical groups in cross-sectional OASIS3 dataset. Positive values indicate an increased pattern of local volume differences compared to healthy ageing patterns at the respective age. HC, Healthy Controls; pMCI, progressive MCI; sMCI, stable MCI; AD, Alzheimer's Disease.

progressive MCI (**Table 5**). We also include histograms of the local brain-PAD scores for each disease group per subcortical ROI to visualise the different distributions that drive the report effect sizes (**Figure 11**). For example, the high Cohen's *d* values for the nucleus accumbens may be due to the low variance in brain-PAD values in this small region. We have provided similar graphics for the cortical regions in the **Supplementary Material** (**Supplementary Figure 6**).

4. DISCUSSION

In this paper, we introduced a novel deep-learning framework capable of reliably predicting age from neuroimaging data at a local neuroanatomical level. Training the powerful U-net architecture on n = 3463 healthy people, we present

the first proof-of-concept, to our knowledge, that generating such localised brain-age predictions is feasible. While average performance of our model (MAE = 9.94 ± 1.73 years) is below what has been reported with purely global (MAE ∼ 3 years, Peng et al., 2021), slice-level (MAE between 5 and 7.5 years, Ballester et al., 2021), or patch-level (MAE 2.5–4 years, Bintsi et al., 2020), we show both high reliability and reasonable generalisability to three entirely independent datasets (OASIS3, AIBL, Wayne State). Importantly, we achieved a resolution of 23^3 voxels, substantially more fine-grain than previous patch-level work (64^3 voxels, Bintsi et al., 2020). In fact, we were able to generate voxel-level prediction, though as within-block homogeneity was high, our effective resolution was lower than single voxel. Future improvements to network architecture may be able to improve the effective resolution still further.

FIGURE 9 | FSL Randomise maps for different combinations of clinical groups in cross-sectional OASIS3 using voxel-level brain-age predictions. Red coloured voxels indicate a significant statistical *t*-test after correcting for multiple comparisons. Blue regions were not significant after correction. HC, Healthy Controls; pMCI, progressive MCI; sMCI, stable MCI; AD, Alzheimer's Disease.

Even though the mean global performance of the local brain-age model was relatively poor, the model still demonstrated sensitivity to cognitive impairment and dementia, suggesting that despite the noise at test time, the relevant signal can still be observed. Previous work involving brain-age and dementia have obtained "brain-AGE" scores of −0.2 years for sMCI, 6.2 years for pMCI, and 6.7 years for AD on the ADNI dataset (Franke et al., 2012). Our results are generally in line with these previous findings, though here, we observed "older-appearing" brains on our sMCI group (mean brain-PAD = 2.8 years in contrast to −0.2 years in Franke et al., 2012). Moreover, we were able to generate spatial maps of brain-PAD for each individual, showing how the patterns of brain-ageing may vary across the brain in a single patient. At the local level, we observed widespread patterns of group differences in brain-PAD maps, including when comparing sMCI and pMCI groups, suggesting that brain-ageing is more pronounced in those MCI patients who go on to develop dementia within 3 years.

It has been commonly reported that the early stages of AD involve atrophy in the medial temporal lobe (MTL) including the hippocampus, and the amygdala, entorhinal, and parahippocampal cortices (Braak and Braak, 1991; Jack et al., 2010; Johnson et al., 2012; Klein-Koerkamp et al., 2014). Our voxel-level analysis showed brain-PAD

differences between healthy controls and people with MCI in these key AD-related regions. Furthermore, the ROI-level analysis of brain-PAD show widespread differences with particularly strong effects in the nucleus accumbens, putamen, pallidum, hippocampus, and amygdala as well as cortical and ventricular regions. While further research is required to further improve the model performance and spatial precision, these results suggest that the local brain-age predictions are sensitivity to local patterns of brain atrophy.

The validity of the predictions from the local brain-age model are further supported by the observed significant negative correlations between ROI volumes and ROI-level brain-PAD. In a similar analysis, Levakov et al. (2020) identified the lateral ventricles, inferior lateral ventricles, 3rd ventricles, non-ventricles CSF and left/right choroid plexus as the ROIs (using the FreeSurfer Desikan-Killiany atlas) having the strongest relationships between age normalised volume and brain-age "gap." Here, we also show relationships in GM ROIs (e.g., amygdala, hippocampus, thalamus, parahippocampal gyrus (anterior division), inferior temporal gyrus, temporal-occipital part, intracalcarine cortex). As lower brain volumes are associated with ageing, the observed negative relationships between ROI volume and ROI brain-PAD suggests that indeed the ROI-level brain-PAD captures some age-related variance.
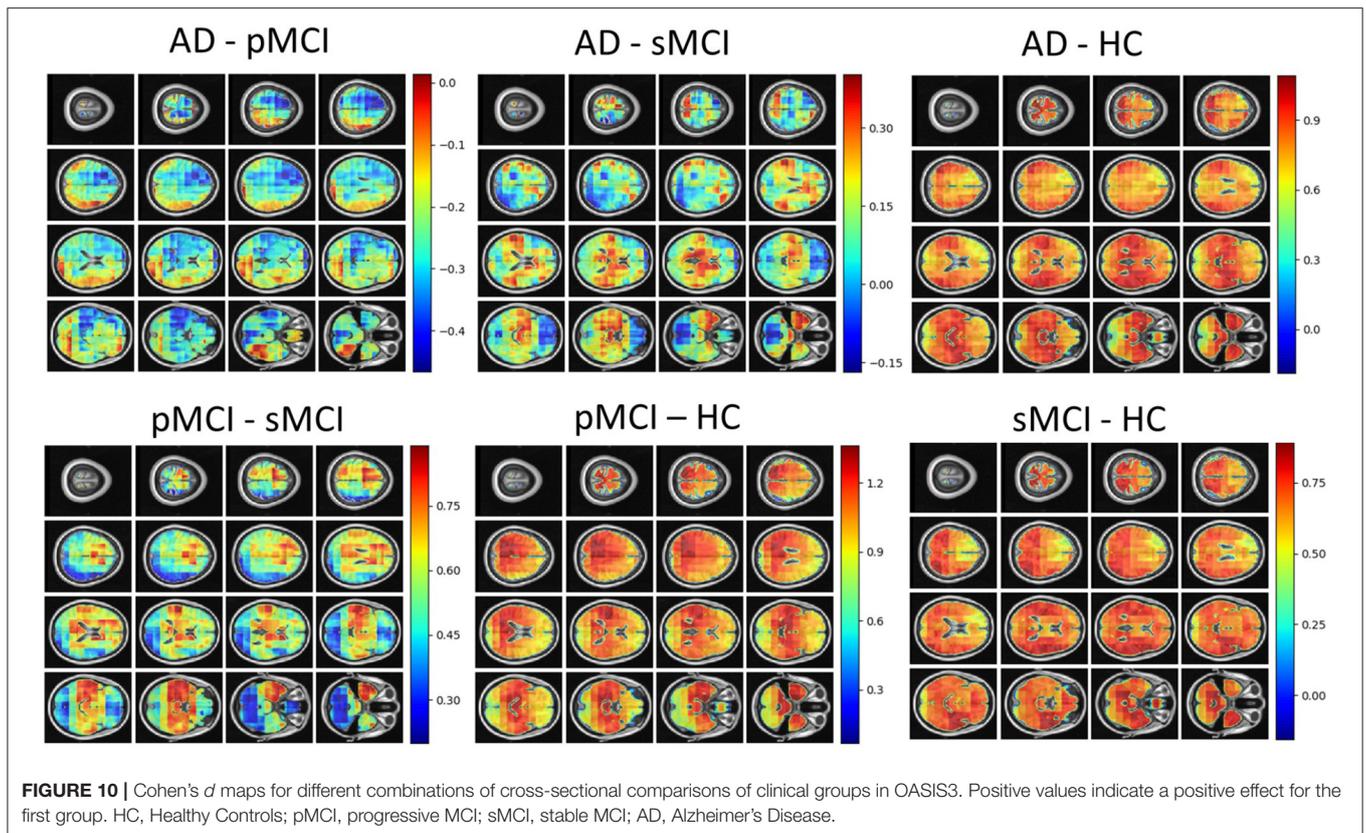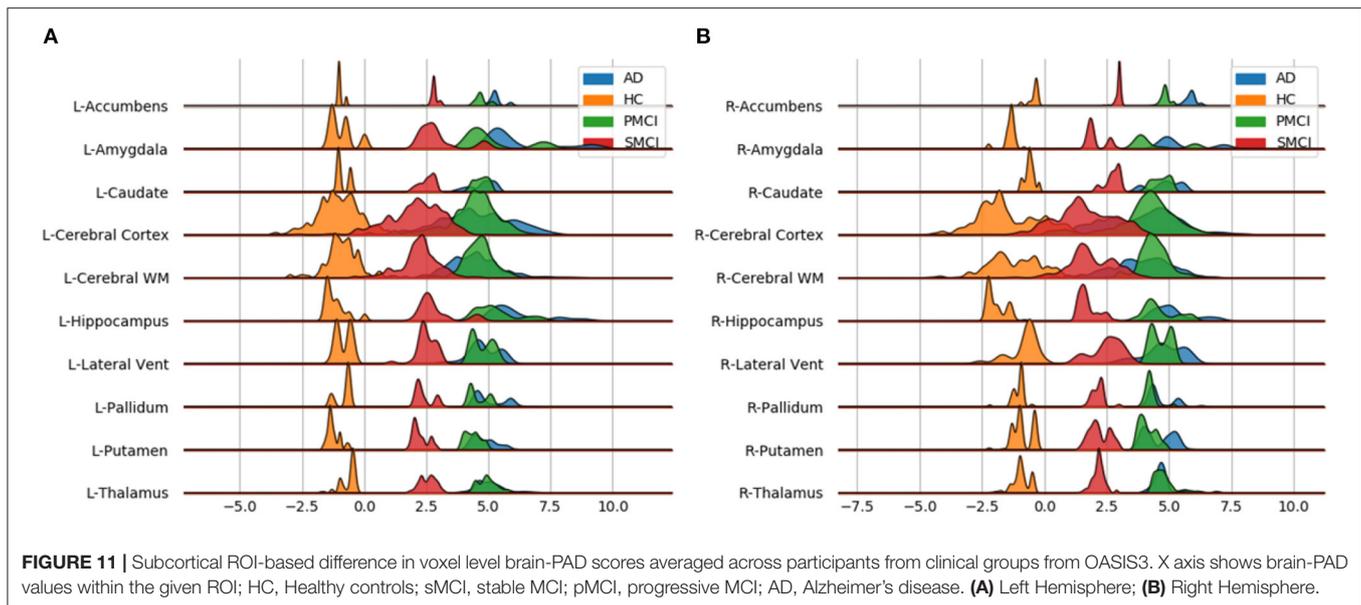
FIGURE 10 | Cohen's *d* maps for different combinations of cross-sectional comparisons of clinical groups in OASIS3. Positive values indicate a positive effect for the first group. HC, Healthy Controls; pMCI, progressive MCI; sMCI, stable MCI; AD, Alzheimer's Disease.

TABLE 5 | Welch's *t*-test statistic (*p*-value/Cohen's *d*) values for different subcortical ROIs from the Harvard-Oxford atlas.

Subcortical ROI name	AD vs. HC left	pMCI vs. sMCI left	AD vs. HC right	pMCI vs. sMCI right
Brain stem	−159.23 (<0.001/1.82)	−111.17 (<0.001/0.53)	−	−
Amygdala	−89.70 (<0.001/4.41)	−83.48 (<0.001/1.43)	−74.61 (<0.001/4.00)	−66.36 (<0.001/1.57)
Caudate	−97.02 (<0.001/4.69)	−98.82 (<0.001/1.69)	−101.28 (<0.001/4.41)	−102.43 (<0.001/1.47)
Cerebral cortex	−576.47 (<0.001/1.98)	−539.03 (<0.001/0.81)	−562.68 (<0.001/1.93)	−528.37 (<0.001/0.89)
Cerebral WM	−1083.13 (<0.001/5.29)	−857.54 (<0.001/2.48)	−1018.97 (<0.001/4.32)	−867.04 (<0.001/2.56)
Hippocampus	−145.32 (<0.001/4.95)	−131.04 (<0.001/1.72)	−146.84 (<0.001/5.77)	−139.80 (<0.001/2.65)
Lateral ventricle	−9.77 (<0.001/0.30)	−9.67 (<0.001/0.11)	−8.57 (<0.001/0.27)	−8.37 (<0.001/0.09)
Pallidum	−191.76 (<0.001/11.80)	−314.02 (<0.001/5.91)	−215.40 (<0.001/13.31)	−175.65 (<0.001/8.31)
Putamen	−562.54 (<0.001/15.59)	−382.47 (<0.001/6.82)	−427.50 (<0.001/10.66)	−306.99 (<0.001/5.13)
Thalamus	−148.15 (<0.001/4.02)	−148.75 (<0.001/1.60)	−165.48 (<0.001/4.80)	−162.59 (<0.001/2.11)
Accumbens	−461.68 (<0.001/30.02)	−205.64 (<0.001/9.60)	−472.20 (<0.001/31.00)	−209.78 (<0.001/9.80)

For Cohen's d, higher values indicate a positive effect size for the first disease group specified.

As biomarker of brain health, brain-age models may have clinical utility, either prognostically or in the context of clinical trials of neuroprotective treatments. While previous studies have reported standardised effect sizes from global brain-age, we used atlas ROIs to summarise regional values of local brain-PAD and generated Cohen's *d* values from pairwise group comparisons. Using conventional hippocampal volumetric measures, Henneman et al. (2009) reported baseline effect size of 0.73 when comparing controls and MCI groups, and 0.33 when comparing people with MCI and AD patients. With

our local brain-age framework, the control-MCI effect size for the hippocampus (average bilaterally) was $d = 5.45$ and the MCI-AD effect size was $d = 0.48$. Using voxel-based morphometry, Risacher et al. (2009) generated Cohen's *d* values for the hippocampus ($d = 0.6$) and amygdala ($d = 0.45$), when comparing stable and progressive MCI patients. Here, our local brain-age framework resulted in $d = 2.18$ for the hippocampus and $d = 1.5$ for the bilateral amygdala in the same context. This suggests that use of the brain-age paradigm to capture local age-related changes, relative to a healthy ageing model, could

FIGURE 11 | Subcortical ROI-based difference in voxel level brain-PAD scores averaged across participants from clinical groups from OASIS3. X axis shows brain-PAD values within the given ROI; HC, Healthy controls; sMCI, stable MCI; pMCI, progressive MCI; AD, Alzheimer's disease. **(A)** Left Hemisphere; **(B)** Right Hemisphere.

increase statistical power in experimental research and clinical trials, relative to conventional volumetric imaging biomarkers. Potentially, the ROI-based brain-PAD values could even be used in a classification framework to distinguish between people with stable or progressive MCI.

Out proposed U-Net local brain-age framework has some strengths and weaknesses. Our model was assessed on a large multi-site testing set with a flat distribution of chronological age across the adult lifespan (18–90 years; **Supplementary Figure 1**), a wider interval than a number of studies that rely on UK Biobank (Bintsi et al., 2020; Gupta et al., 2021; Peng et al., 2021) or other narrower-age range studies. Our model showed excellent test-retest reliability, giving confidence that the model could be applied longitudinally to assess individual patterns of brain-ageing changes. However, the between-scanner reliability was moderate, similar to our previous work using deep learning to predict brain age (Cole et al., 2017). In the latter work, brain-age prediction was performed directly on raw MRI scans, hence the deep learning model may be overfitting to some site or effects. One might expect that image pre-processing may partially ameliorate these site effects. However, this is not uniformly the case, as previous research has demonstrated (Glocker et al., 2019). Consequently, one drawback of the current algorithm is the requirement to have a healthy population from a given clinical site to use as a control group, as site or scanner effects may result in the local brain-PAD distribution not being centred at zero. In **Supplementary Figures 7**, **8**, we show that these scanner effects have only a marginal effect on the statistical comparisons between MCI or dementia groups using the AIBL dataset in reference to our main results from OASIS3. Nevertheless, harmonisation of scanner and site is a key direction for future work as the removal of residual scanner effects is likely to improve model generalisability considerably and is an important prerequisite for the clinical adoption of neuroimaging biomarker pipelines.

We trained our regression U-Net with the ground truth objective at a voxel-level (given by a three-dimensional block

filled with the chronological age), in order to encourage the network to emphasise the context encoded in its lower layers. As the individual voxel location we are aiming to obtain a prediction for is not necessarily related to the imposed ground truth output, the U-Net architecture is biased toward using the context information. Hence, in the worst case scenario where no voxel-level relationship is learned, the true resolution of our voxel-level predictions is actually blocks of 23^3 voxels. The final output field-of-view (FOV) was calculated starting from the first convolutional layer where the FOV is 3^3 voxels, which gets increased by 2^3 voxels per convolutional operation in the downstream part of the U-Net. The average pooling layers increase the FOV by 1^3 voxels, since their stride is set to 1, while the upsampling layers do not increase the FOV as they merely repeat existing information. Lastly, the first convolutional layer in the upstream layers only adds 1^3 voxels (since a $2 * 2 * 2$ block inside the operating field of the $3 * 3 * 3$ filter contains the same repeated information, hence no increase in the FOV) whereas the second adds 3^3 voxels (since a $3 * 3 * 3$ filter will have access to 3 additional voxels stemming from the upsampling layer). While this means that our resolution is not necessarily at the voxel-level, 23^3 voxels is still substantially higher resolution compared to existing models in literature. In the 3D block approach of Bintsi et al. (2020), blocks are much larger, 64^3 voxels. Hence, any block-level age prediction will be biased toward the global-level brain age prediction as the blocks include a substantial portion of the overall brain. Moreover, in splitting the whole brain into blocks, naturally some blocks will include non-brain tissue or empty space, which will naturally reduce the amount of discriminative information present there, reducing the validity of results for regions within the respective block.

One drawback of this study is the categorisation of subjects into sMCI vs. pMCI vs. AD. We attempted to closely follow the ADNI clinical procedure, respectively to classify a subject as AD if they had a mini-mental state examination (MMSE) score between 24 and 26, CDR of 0.5 or 1.0. Subjects with an

MMSE score between 24 and 30, CDR of 0.5 alongside memory complains, objective memory loss measured by education adjusted scores and absence of significant levels of impairment in other cognitive domains are classified as MCI. Unfortunately, for OASIS3 we only had access to CDR scores, hence there is the possibility of misclassifications occurring in our data curation.

We demonstrated how our U-net framework can predict age from neuroimaging data. However, this approach could be trained on any continuous or categorical outcome measure to generate individual maps of how given outcomes vary across the brain. For example, one could generate spatial maps of predicted values of fluid biomarkers (e.g., amyloid, tau), genotype or polygenic risk score, cognitive measures (e.g., MMSE scores). Such an approach could be used as an alternative to techniques like VBM, to provide mechanistic insights into the relationship between local brain regions and individual deviations from healthy/normal levels of a given outcome measure. While VBM is the *de facto* method to quantitatively assess differences between groups at voxel-level (Koikkalainen et al., 2016), we believe the local brain-age framework is complementary to this. In VBM analysis, one assesses the statistical models at a voxel level based on volume or intensity, though local context is only really accounted for at the cluster inference stage. Brain-PAD implicitly measures this deviation of the diseases group from what constitutes a normative pattern of ageing, by placing the participant on a distribution of normative ageing for a given local area (i.e., the voxel and its local context). We leave for further work the comparison between VBM and local brain-age.

One potential direction to take local brain-age further is in disease subtyping. Local brain-PAD maps could be used as input to clustering algorithms with the goal of identifying subgroups of patients that have spatially similar patterns of brain ageing. The putative subgroups may undergoing distinct pathological processes that effect different regions of the brain and may have different trajectories of disease progression or may respond differently to treatments. Such approaches have been applied to volumetric brain maps before (Dong et al., 2015), but the addition of brain-PAD information as a local index of age-adjusted brain health could increase sensitivity, as has been seen in global brain-age research (Franke and Gaser, 2012).

Lastly, an important aspect of voxel-level brain-age prediction is how to properly adjust for age-related biases in predictions. In section 2.6, we have introduced two different methods, one which was successfully used for global-level brain-age prediction before, respectively a new technique for voxel-level predictions. In section 5 in the **Supplementary Material**, we provide additional results and diagnostics for the aforementioned techniques. From our analysis, it results that global-level bias-adjustment techniques fail in the voxel-level case. Further research should focus on why this happens and what are some possible remedies.

5. CONCLUSION

We have introduced a new deep learning framework that is capable of reliably estimating brain-age with high spatial resolution, providing information on spatial patterns of age-related changes to brain volume. We were able to demonstrate the potential clinical relevance of the model by mapping differences in local and regional brain-PAD scores in patients with cognitive impairment and dementia. This work illustrates how the sensitivity of conventional global brain-age analysis can be augmented with individualised spatial maps offering potential mechanistic insights, with the goal of opening the "black box" of the machine learning algorithms that underpin the brain-age paradigm.

DATA AVAILABILITY STATEMENT

Publicly available datasets were analysed in this study. This data can be found here: details of multiple data sources are provided in the **Supplementary Material**.

ETHICS STATEMENT

Ethical review and approval was not required for the study on human participants in accordance with the local legislation and institutional requirements. The patients/participants provided their written informed consent to participate in this study.

AUTHOR CONTRIBUTIONS

SP: conceptualisation, methodology, software, formal analysis, investigation, writing—original draft, writing—review and editing, and visualisation. BG, DS, and JC: conceptualisation, methodology, writing—review and editing, and supervision. All authors contributed to the article and approved the submitted version.

FUNDING

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fnagi.2021.761954/full#supplementary-material

REFERENCES

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., et al. (2016). "Tensorflow: a system for large-scale machine learning," in *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)* (Savannah, GA), 265–283.

Ashburner, J. (2007). A fast diffeomorphic image registration algorithm. *Neuroimage* 38, 95–113. doi: 10.1016/j.neuroimage.2007.07.007

Ballester, P. L., da Silva, L. T., Marcon, M., Esper, N. B., Frey, B. N., Buchweitz, A., et al. (2021). Predicting brain age at slice level: convolutional neural networks and consequences for interpretability. *Front. Psychiatry* 12:118. doi: 10.3389/fpsyt.2021.598518

Beheshti, I., Gravel, P., Potvin, O., Dieumegarde, L., and Duchesne, S. (2019). A novel patch-based procedure for estimating brain age across adulthood. *Neuroimage* 197, 618–624. doi: 10.1016/j.neuroimage.2019.05.025

Bintsi, K.-M., Baltatzis, V., Kolbeinsson, A., Hammers, A., and Rueckert, D. (2020). Patch-based brain age estimation from mr images. *arXiv preprint arXiv:2008.12965*. doi: 10.1007/978-3-030-66843-3_10

Biondo, F., Jewell, A., Pritchard, M., Aarsland, D., Steves, C. J., Mueller, C., et al. (2021). Brain-age predicts subsequent dementia in memory clinic patients. *medRxiv*. 16:e037378. doi: 10.1101/2021.04.03.21254781

Braak, H., and Braak, E. (1991). Neuropathological stageing of alzheimer-related changes. *Acta Neuropathol.* 82, 239–259. doi: 10.1007/BF00308809

Chaudhuri, A. (2013). Multiple sclerosis is primarily a neurodegenerative disease. *J. Neural Trans.* 120, 1463–1466. doi: 10.1007/s00702-013-1080-3

Cherubini, A., Caligiuri, M. E., Péran, P., Sabatini, U., Cosentino, C., and Amato, F. (2016). Importance of multimodal MRI in characterizing brain tissue and its potential application for individual age prediction. *IEEE J. Biomed. Health Inform.* 20, 1232–1239. doi: 10.1109/JBHI.2016.2559938

Cole, J. H., Marioni, R. E., Harris, S. E., and Deary, I. J. (2019). Brain age and other bodily "ages": implications for neuropsychiatry. *Mol. Psychiatry* 24, 266–281. doi: 10.1038/s41380-018-0098-1

Cole, J. H., Poudel, R. P., Tsagkrasoulis, D., Caan, M. W., Steves, C., Spector, T. D., et al. (2017). Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker. *Neuroimage* 163, 115–124. doi: 10.1016/j.neuroimage.2017.07.059

Cole, J. H., Raffel, J., Friede, T., Eshaghi, A., Brownlee, W. J., Chard, D., et al. (2020). Longitudinal assessment of multiple sclerosis with the brain-age paradigm. *Ann. Neurol.* 88, 93–105. doi: 10.1002/ana.25746

Cole, J. H., Ritchie, S. J., Bastin, M. E., Hernández, M. V., Maniega, S. M., Royle, N., et al. (2018). Brain age predicts mortality. *Mol. Psychiatry* 23, 1385–1392. doi: 10.1038/mp.2017.62

Coupé, P., Eskildsen, S. F., Manjón, J. V., Fonov, V. S., Collins, D. L., Initiative, A. D. N., et al. (2012). Simultaneous segmentation and grading of anatomical structures for patient's classification: application to Alzheimer's disease. *Neuroimage* 59, 3736–3747. doi: 10.1016/j.neuroimage.2011.10.080

de Lange, A.-M. G., and Cole, J. H. (2020). Commentary: Correction procedures in brain-age prediction. *Neuroimage: Clin.* 26:102229. doi: 10.1016/j.nicl.2020.102229

Dinsdale, N. K., Bluemke, E., Smith, S. M., Arya, Z., Vidaurre, D., Jenkinson, M., et al. (2020). Learning patterns of the ageing brain in mri using deep convolutional networks. *Neuroimage* 224:117401. doi: 10.1016/j.neuroimage.2020.117401

Dong, A., Honnorat, N., Gaonkar, B., and Davatzikos, C. (2015). Chimera: clustering of heterogeneous disease effects via distribution matching of imaging patterns. *IEEE Trans. Med. Imaging* 35, 612–621. doi: 10.1109/TMI.2015.2487423

Erramuzpe, A., Schurr, R., Yeatman, J., Gotlib, I., Sacchet, M., Travis, K., et al. (2020). A comparison of quantitative r1 and cortical thickness in identifying age, lifespan dynamics, and disease states of the human cortex. *Cereb. Cortex.* 31, 1211–1226. doi: 10.1093/cercor/bhaa288

Feng, X., Tustison, N. J., Patel, S. H., and Meyer, C. H. (2020). Brain tumor segmentation using an ensemble of 3D U-Nets and overall survival prediction using radiomic features. *Front. Comput. Neurosci.* 14:25. doi: 10.3389/fncom.2020.00025

Franke, K., and Gaser, C. (2012). Longitudinal changes in individual brainage in healthy aging, mild cognitive impairment, and Alzheimer's disease. *GeroPsych.* 25, 235–245. doi: 10.1024/1662-9647/a000074

Franke, K., and Gaser, C. (2019). Ten years of brainage as a neuroimaging biomarker of brain aging: what insights have we gained? *Front. Neurol.* 10:789. doi: 10.3389/fneur.2019.00789

Franke, K., Luders, E., May, A., Wilke, M., and Gaser, C. (2012). Brain maturation: predicting individual brainage in children and adolescents using structural MRI. *Neuroimage* 63, 1305–1312. doi: 10.1016/j.neuroimage.2012.08.001

Franke, K., Ziegler, G., Klöppel, S., Gaser, C., Alzheimer's Disease Neuroimaging Initiative, et al. (2010). Estimating the age of healthy subjects from T1-weighted MRI scans using kernel methods: exploring the influence of various parameters. *Neuroimage* 50, 883–892. doi: 10.1016/j.neuroimage.2010.01.005

Gaser, C., Franke, K., Klöppel, S., Koutsouleris, N., Sauer, H., Alzheimer's Disease Neuroimaging Initiative, et al. (2013). Brainage in mild cognitive impaired patients: predicting the conversion to Alzheimer's disease. *PLoS ONE* 8:e67346. doi: 10.1371/journal.pone.0067346

Glocker, B., Robinson, R., Castro, D. C., Dou, Q., and Konukoglu, E. (2019). Machine learning with multi-site imaging data: an empirical study on the impact of scanner effects. *arXiv [preprint] arXiv:1910.04597*.

Gupta, U., Lam, P. K., Ver Steeg, G., and Thompson, P. M. (2021). "Improved brain age estimation with slice-based set networks," in *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)* (Nice: IEEE), 840–844. doi: 10.1109/ISBI48211.2021.9434081

Harper, L., Barkhof, F., Scheltens, P., Schott, J. M., and Fox, N. C. (2014). An algorithmic approach to structural imaging in dementia. *J. Neurol. Neurosurg. Psychiatry* 85, 692–698. doi: 10.1136/jnnp-2013-306285

He, K., Zhang, X., Ren, S., and Sun, J. (2016). "Identity mappings in deep residual networks," in *European Conference on Computer Vision* (Amsterdam: Springer), 630–645. doi: 10.1007/978-3-319-46493-0_38

Henneman, W. J., Sluimer, J. D., Barnes, J., van der Flier, W. M., Sluimer, I. C., Fox, N. C., et al. (2009). Hippocampal atrophy rates in Alzheimer disease: added value over whole brain volume measures. *Neurology* 72, 999–1007. doi: 10.1212/01.wnl.0000344568.09360.31

Hu, J., Shen, L., and Sun, G. (2018). "Squeeze-and-excitation networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (Salt Lake City, UT), 7132–7141. doi: 10.1109/CVPR.2018.00745

Jack, C. R. Jr., Wiste, H. J., Vemuri, P., Weigand, S. D., Senjem, M. L., Zeng, G., et al. (2010). Brain beta-amyloid measures and magnetic resonance imaging atrophy both predict time-to-progression from mild cognitive impairment to Alzheimer's disease. *Brain* 133, 3336–3348. doi: 10.1093/brain/awq277

Johnson, K. A., Fox, N. C., Sperling, R. A., and Klunk, W. E. (2012). Brain imaging in Alzheimer disease. *Cold Spring Harbor Perspect. Med.* 2:a006213. doi: 10.1101/cshperspect.a006213

Kaufmann, T., van der Meer, D., Doan, N. T., Schwarz, E., Lund, M. J., Agartz, I., et al. (2019). Common brain disorders are associated with heritable patterns of apparent aging of the brain. *Nat. Neurosci.* 22, 1617–1623.

Kingma, D. P., and Ba, J. (2014). Adam: a method for stochastic optimization. *arXiv [preprint] arXiv:1412.6980*.

Klein-Koerkamp, Y., A., Heckemann, R., T., Ramdeen, K., Moreaud, O., et al. (2014). Amygdalar atrophy in early Alzheimers disease. *Curr. Alzheimer Res.* 11, 239–252. doi: 10.2174/1567205011666140131123653

Koikkalainen, J., Rhodius-Meester, H., Tolonen, A., Barkhof, F., Tijms, B., Lemstra, A. W., et al. (2016). Differential diagnosis of neurodegenerative diseases using structural mri data. *Neuroimage Clin.* 11, 435–449. doi: 10.1016/j.nicl.2016.02.019

Kolbeinsson, A., Filippi, S., Panagakis, Y., Matthews, P. M., Elliott, P., Dehghan, A., et al. (2020). Accelerated MRI-predicted brain ageing and its associations with cardiometabolic and brain disorders. *Sci. Rep.* 10, 1–9. doi: 10.1038/s41598-020-76518-z

Laakso, M., Partanen, K., Riekkinen, P., Lehtovirta, M., Helkala, E.-L., Hallikainen, M., et al. (1996). Hippocampal volumes in Alzheimer's disease, Parkinson's disease with and without dementia, and in vascular dementia an MRI study. *Neurology* 46, 678–681. doi: 10.1212/WNL.46.3.678

Landin-Romero, R., Tan, R., Hodges, J. R., and Kumfor, F. (2016). An update on semantic dementia: genetics, imaging, and pathology. *Alzheimers Res. Therapy* 8, 1–9. doi: 10.1186/s13195-016-0219-5

Levakov, G., Rosenthal, G., Shelef, I., Raviv, T. R., and Avidan, G. (2020). From a deep learning model back to the brain-identifying regional predictors and their relation to aging. *Hum. Brain Mapp.* 41, 3235–3252. doi: 10.1002/hbm.25011

Lockhart, S. N., and DeCarli, C. (2014). Structural imaging measures of brain aging. *Neuropsychol. Rev*. 24, 271–289. doi: 10.1007/s11065-014-9268-3

Lorenzetti, V., Allen, N. B., Fornito, A., and Yücel, M. (2009). Structural brain abnormalities in major depressive disorder: a selective review of recent mri studies. *J. Affect. Disord*. 117, 1–17. doi: 10.1016/j.jad.2008.11.021

Pawlowski, N., and Glocker, B. (2019). Is texture predictive for age and sex in brain MRI? *arXiv [preprint] arXiv:1907.10961*.

Peng, H., Gong, W., Beckmann, C. F., Vedaldi, A., and Smith, S. M. (2021). Accurate brain age prediction with lightweight deep neural networks. *Med. Image Anal*. 68:101871. doi: 10.1016/j.media.2020.101871

Popescu, S. G., Whittington, A., Gunn, R. N., Matthews, P. M., Glocker, B., Sharp, D. J., et al. (2020). Nonlinear biomarker interactions in conversion from mild cognitive impairment to Alzheimer's disease. *Hum. Brain Map*. 41, 4406–4418. doi: 10.1002/hbm.25133

Risacher, S. L., Saykin, A. J., Wes, J. D., Shen, L., Firpi, H. A., and McDonald, B. C. (2009). Baseline MRI predictors of conversion from MCI to probable AD in the ADNI cohort. *Curr. Alzheimer Res*. 6, 347–361. doi: 10.2174/156720509788929273

Ronneberger, O., Fischer, P., and Brox, T. (2015). "U-Net: convolutional networks for biomedical image segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Munich: Springer), 234–241. doi: 10.1007/978-3-319-24574-4_28

Shrout, P. E., and Fleiss, J. L. (1979). Intraclass correlations: uses in assessing rater reliability. *Psychol. Bull*. 86:420. doi: 10.1037/0033-2909.86.2.420

Varikuti, D. P., Genon, S., Sotiras, A., Schwender, H., Hoffstaedter, F., Patil, K. R., et al. (2018). Evaluation of non-negative matrix factorization of grey matter in age prediction. *Neuroimage* 173, 394–410. doi: 10.1016/j.neuroimage.2018.03.007

Vesal, S., Ravikumar, N., and Maier, A. (2019). A 2D dilated residual u-net for multi-organ segmentation in thoracic ct. *arXiv [preprint] arXiv:1905.07710*.

Wang, J., Knol, M. J., Tiulpin, A., Dubost, F., de Bruijne, M., Vernooij, M. W., et al. (2019). Gray matter age prediction as a biomarker for risk of dementia. *Proc. Natl. Acad. Sci. U.S.A*. 116, 21213–21218. doi: 10.1073/pnas.1902376116