

Region-wise stacking ensembles for estimating brain-age using structural MRI

Georgios Antonopoulos^{a,b,*}, Shammi More^{a,b,c}, Simon B. Eickhoff^{a,b}, Federico Raimondo^{a,b}, Kaustubh R. Patil^{a,b}, for the Alzheimer's Disease Neuroimaging Initiative¹

^a Institute of Neuroscience and Medicine (INM-7: Brain and Behaviour), Research Centre, Jülich, Germany

^b Institute of Systems Neuroscience, Heinrich Heine University Düsseldorf, Düsseldorf, Germany

^c Department of Bioinformatics, Fraunhofer Institute for Algorithms and Scientific Computing (SCAI), Sankt Augustin, Germany

ARTICLE INFO

Keywords:

MRI
Age-prediction
Brain-age
Stacking ensemble
Aging

ABSTRACT

Predictive modeling using structural magnetic resonance imaging (MRI) data is a prominent approach to study brain-aging. Machine learning frameworks have been employed to improve predictions and explore healthy and accelerated aging due to diseases. The high-dimensional MRI data pose challenges to building generalizable and interpretable models as well as for data privacy. Common practices are resampling or averaging voxels within predefined parcels which reduces anatomical specificity and biological interpretability. Effectively, naive fusion by averaging can result in information loss and reduced accuracy. We present a conceptually novel two-level stacking ensemble (SE) approach. The first level comprises regional models for predicting individuals' age based on voxel-wise information, fused by a second-level model yielding final predictions. Eight data fusion scenarios were explored using Gray matter volume (GMV) estimates from four large datasets. Performance measured using mean absolute error (MAE), R^2 , correlation and prediction bias, showed that SE outperformed the region-wise averages. The best performance was obtained when first-level regional predictions were obtained as out-of-sample predictions on the application site with second-level models trained on independent and site-specific data (MAE = 4.75 vs baseline regional mean GMV MAE = 5.68). Performance improved as more datasets were used for training. First-level predictions showed improved and more robust aging signal providing new biological insights and enhanced data privacy. Overall, the SE improves accuracy compared to the baseline while preserving or enhancing data privacy. Finally, we show the utility of our SE model on a clinical cohort showing accelerated aging in cognitively impaired and Alzheimer's disease patients.

1. Introduction

The process of aging in humans is complex and it inevitably influences the brain, with negative consequences such as neurodegeneration which can lead to dementia and also is a mortality risk [1, 2]. The structural changes in the brain can be measured non-invasively via Magnetic Resonance Imaging (MRI) scans which capture the brain in volumetric form, comprising smaller volume elements called voxels. Each scan consists of hundreds of thousands of voxels. Following appropriate processing, it is possible to quantify the amount of specific

brain tissues such as gray matter volume (GMV) at each voxel. This in turn permits in-depth study of distinct brain structures in relation to various cognitive, pathological and other physiological processes such as aging. Structural differences in the GMV have been reported between younger and older healthy individuals [3–5] along with continuous, widespread reduction in GMV observed from middle age onwards [6–8]. Additionally, an accelerated loss of both global and local GMV has been documented in neurodegenerative disorders [6,9–13] such as the Alzheimer's disease (AD) and mild cognitive impairment (MCI) compared to normal aging.

* Corresponding author. Applied Machine Learning, Institute of Neuroscience and Medicine, (INM-7: Brain and Behaviour), Research Centre Jülich, 52425, Jülich, Germany.

E-mail address: g.antonopoulos@fz-juelich.de (G. Antonopoulos).

¹ Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf.

<https://doi.org/10.1016/j.combiomed.2025.111182>

Received 27 March 2025; Received in revised form 26 September 2025; Accepted 2 October 2025

Available online 12 October 2025

0010-4825/© 2025 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Leveraging machine learning models to predict chronological age in healthy individuals using GMV aims to estimate the trajectory of healthy brain aging and through it unravel the links to neurodegenerative, psychiatric, and other disorders. Such MRI-derived age-prediction has been shown to be a reliable proxy for overall health [10,11,14]. An elevated brain age stands as an important risk factor for various neurodegenerative and psychiatric disorders [12,13,15–18]. Thus, age prediction models can facilitate early detection of health risk and facilitate effective prevention and treatment strategies. The success of this approach hinges on accurate and biologically meaningful models. Furthermore, sharing MRI data or even its derivatives can lead to privacy issues [19]. Hence it is important that new brain age estimation methods respect individuals' privacy. Furthermore, generalizing across data from multiple scanners remains challenging due to systematic biases. To this end, various machine learning algorithms and GMV-derived features have been tested aiming at providing more accurate predictions as well as insights into healthy brain aging. For instance, Relevance Vector Regression (RVR) outperformed Support Vector Regression (SVR) when trained on downsampled voxels -with and without feature selection, yielding MAE = 5 years in healthy subjects and MAE = 10 years in AD [20]. Gaussian Process Regression using structural MRIs was efficient in estimating mortality via brain age [1] and outperformed RVR, Kernel Ridge Regression and LASSO [1,2]. Additionally, various deep neural network models have been proposed such as convolutional neural networks [12,21–26].

In terms of GMV features, previous studies have used voxel-wise GMV [1,2,27,28], feature reduction methods such as principal component analysis and non-negative matrix factorization [29,30] as well as regional mean values [31,32]. The voxel-based approach is encumbered by the “curse of dimensionality”, high computational demands due to the high number of voxels and high intersubject variability. On the other hand, the use of regional mean GMV is more conducive for machine learning. This simple and efficient type of information fusion can enhance robustness and statistical power mitigating the impact of noise or variability within individual voxels. Additionally, use of a meaningful parcellation scheme increases sensitivity for detecting regional changes, yielding results that are easier to interpret from a neuroanatomical perspective and can provide critical priors for studying brain basis of behavior and disease [33]. Effectively, calculating regional mean of GMV as a form of information fusion might be too naive, as it weighs all the voxels equally regardless of their reliability and relevance to the problem at hand.

Further challenges to real-world use of brain age models are posed by data heterogeneity and data privacy issues. It is desirable that the models should work effectively on test data collected from a new scanner from which no or only limited data is available for training. However, scanner-induced systematic differences make this challenging [34,35]. It is also desirable that data is utilized and shared in a way that preserves privacy, a crucial consideration in healthcare and medicine [36]. Any use of personal and medical data should ensure it's ethical for research and treatment, and comply with legal obligations. For instance, raw MRI data is enough for face identification of the subjects [37] as well as information such as MRI-derived connectomes are unique to individuals like a fingerprint promoting identification [38]. Subjects can be identified based on their regional GMV, even across different preprocessing pipelines, indicating that sharing preprocessed data still poses the risk of privacy violation [34].

To better utilize the information of each voxel, while managing the very high dimensionality of the data, we propose the use of stacking generalization or stacking ensemble (SE) models [39] which effectively mitigate overall bias and variance, and have shown promise in multiple domains [40–42]. The SE framework is based on weighted voting of predictions derived from various models (referred to as generalizers in Ref. [39]). Specifically, we propose a variation of SE in which first level models are trained on the voxel-wise data from each brain region. In contrast to the uniform weighting, as done in conventional models that

use the mean for each region, the first level models weight the contribution of each voxel according to how it captures the aging signal. This approach can effectively handle regions in which the voxels have differing signal to noise ratio. A second level model is then trained using the predictions of the first level models to obtain the final prediction. In addition, the predictions of the first level models are more aligned with the target, in our case age. This allows for better fusion of data across datasets/scanners, in effect mitigating the scanner differences and allowing to combine datasets to train more accurate and generalizable models. In other words, pooling age predictions of first level models is likely to incur lower bias than pooling regional mean GMV. Overall, we expect that this more nuanced and informative fusion of voxel and cross-dataset information contrary to the conventional mean approach will potentially lead to more accurate and biologically meaningful models. The proposed SE framework, to some degree, addresses also the privacy issues. By using the output of the first level, i.e. regional age predictions, for cross-site predictions and sharing our framework promotes interoperability and privacy.

While SE models have been previously used in age-prediction for combining the predictions of various modalities [11,43], they have not been used for region-wise structural MRI models. Popescu et al. [28] utilized structural MRI features to predict age for each brain region, however, they did not combine regional predictions in a meta model. In this study we use a SE framework to enhance brain-age predictions and provide an alternative and more robust biological representation of the contribution of brain regions in healthy aging. We employed structural MRI scans of healthy individuals from four large open datasets covering the adult lifespan. We systematically explored various ways to fuse data from different sites during training and testing. Specifically, we performed cross-site predictions, by keeping a test site separate from the training sets, while ensuring consistent training sample sizes across all set ups. Furthermore, we also explored different methods for conducting out-of-sample (OOS) predictions with and without site-specific models. We compared the conventional baseline utilizing the mean GMV for each region against using the SE framework in all those set ups. Additionally, we assessed the impact of the number of datasets available for training and examined their stability on data coming from different sites. Finally, the two best-performing set-ups were tested on clinical data consisting of cognitively normal (CN), mild cognitive impairment (MCI), and Alzheimer's disease (AD) groups. At the current stage we did not perform any comparison between SE models and the state-of-the-art (SOTA) methods. Our goal is to demonstrate the theoretical and practical promise of SE models in this specific context rather than to immediately outperform highly tuned, task-specific models. Our systematic analysis provides valuable insights regarding the efficacy of our proposed stacking ensemble framework in the domain of brain age prediction, paving the path for clinical applications.

2. Material and methods

2.1. Datasets and preprocessing

We used T1-weighted MRI scans of healthy individuals (total $N = 2926$, covering the whole adult lifespan (18–88 years), coming from 4 open datasets: the Cambridge Center for Ageing and Neuroscience (CamCAN, $N = 650$, mean age = 54 ± 18.6 , min-max = 18–88) [44], Information eXtraction from Images (IXI, $N = 562$, mean age = 48.7 ± 16.45 , min-max = 20–86) (<https://brain-development.org/ixi-dataset/>), the enhanced Nathan Kline Institute-Rockland Sample (eNKI, $N = 597$, mean age = 48.2 ± 18.5 , min-max = 18–85) [45] and the 1000 brains study (1000Brains, $N = 1117$, mean age = 61.8 ± 12.4 , min-max = 21–85) [46]. Additionally, we used 3T T1w images from the Alzheimer's Disease Neuroimaging Initiative 3 (ADNI3 $N = 404$, mean age = 79.9 ± 7.45 , min-max = 50.4–90.6; <https://adni.loni.usc.edu/>) [47, 48] database to evaluate the utility of our model in neurodegenerative disorders. From the available cognitive normal (CN, $N = 372$), we

selected 66 age- and sex-matched to our available Alzheimer's disease (AD, $N = 66$) subjects. The remaining 306 CN scans were used exclusively for model training. Similarly, we used 66, matched for sex and age, subjects with mild cognitive impairment (MCI, $N = 66$) for testing. We used only subjects diagnosed at the baseline.

All datasets underwent the exact same preprocessing. Specifically, we processed the T1 scans to extract modulated GMV in the MNI space for each subject, using Computational Anatomy Toolbox (CAT) version 12.8 [49]. To ensure accurate normalization and segmentation, initial affine registration of T1w images was done with higher than default accuracy ($\text{accstr} = 0.8$). After bias field correction and tissue class segmentation, optimized Geodesic Shooting [50] was used for normalization ($\text{regstr} = 1$). We used 1.5 mm Geodesic Shooting templates and outputted 1.5 mm isotropic images. The normalized GM segments were then modulated for linear and non-linear transformations yielding estimated GMV for each voxel ($N = 399184$ voxels in total).

2.2. The SE model training

In this study we implemented SE with two levels, denoted as L0 and L1. For both levels, we utilized GLMnet (elastic net) [51] due its ability to handle multicollinearity in data, such as the voxels of structural MRI, efficiently. This choice was made to demonstrate the effectiveness of SE, though other models may yield a better performance and will be evaluated in the future. Elastic net regression combines ridge regression (L2 regularization) and LASSO (L1 regularization), which facilitates dealing with multicollinearity by penalizing large coefficients and promoting simpler models. We used the glmnet package in R (version R-4.1.0) that

incorporates an internal process for hyperparameter tuning. By default, it optimizes the regularization parameter λ along with the mixing parameter α using 'adaptive resampling'. Features with near-zero variance were identified and removed. The remaining features underwent centering -removing the mean-, and scaling -dividing with the standard deviation.

At the first level L0, GM voxels were grouped into 873 regions encompassing cortical, subcortical and cerebellar regions, using a parcellation atlas [32] with 800 cortical regions from the Schaefer atlas [52], 36 subcortical regions from the Brainnetome Atlas [53] and 37 cerebellar regions [54]. We chose this specific granularity to retain anatomical specificity. Our choice of the 800-region Schaefer atlas was based on prior evidence demonstrating similar performance of 800 and 1000-parcel resolutions [2], while offering a higher computational efficiency. Nevertheless, other options could be also explored. We used a stratified 3-fold cross-validation scheme as it employs a smaller training sets for each fold and maintains a good trade-off between data coverage and computational efficiency to generate voxel-based and out-of-sample L0 model predictions for each region. Additionally, to establish robustness, we also trained the L0 models using 10-fold CV scheme. L0 predictions (873 per subject) were then used as inputs to train the second level L1 model. For application on new unseen samples, a final L0 model (GLMnet) per region was trained on the complete training set. The L1 GLMnet model was trained on the L0 models' OOS predictions from the CV and provides the final age prediction. A full overview of the training process and the application of the model is illustrated in Fig. 1.

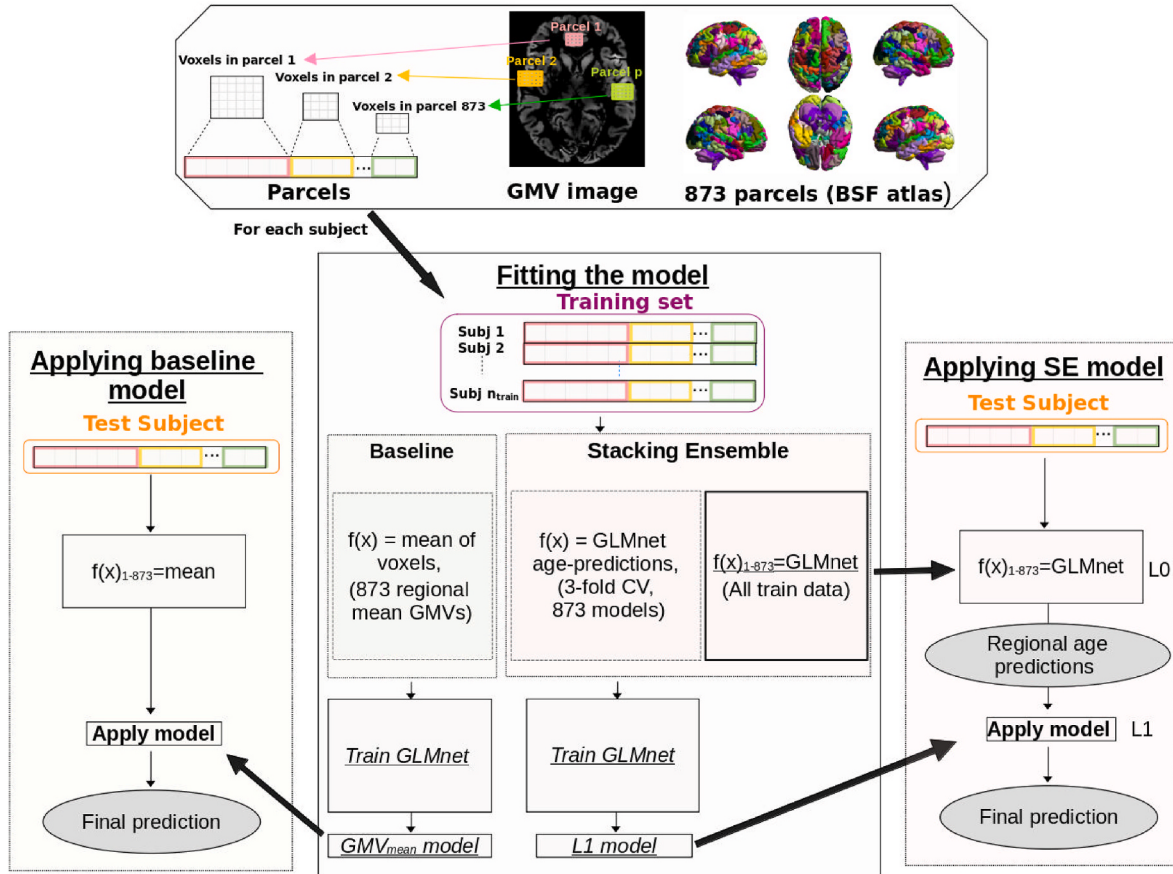


Fig. 1. Illustration of standard age-prediction process. Left: Current standards perform a regional mean of GMV in level-0 ($f(V_i) = \text{mean}(\text{voxels in region } i)$) and with those they train a model in Level-1. In SE models voxels in each region are used to train a model to predict age, in a K-fold scheme, and the predictions are used to train the Level-1 model, which is used to estimate the final prediction. Right: In L0 the trained models (or local averaging) are applied and their output is provided as inputs for the L1 model, which makes the final prediction.

2.3. Data split and training setups

In our experiments, we estimated models' performance using Leave-One-Site-Out (LOSO) validation. The LOSO set up mimics the scenario when the models are applied to data from a new scanner not available during training, thus it provides a more accurate estimation of model's ability to generalize across different scanners or sites. When using multiple datasets for training, the training data or models can be combined in various ways. Data from different sites can be either pooled and used as if they are coming from one source for training L0 and L1 models. Alternatively, data from each site can be used separately to train both in L0 and L1 models creating site-specific SE models and the predictions of L1 models can be then averaged to obtain the final prediction. In-between set ups are also possible, for instance L0 predictions can be pooled or averaged prior to training a L1 model. We implemented 8 such set-ups as presented in Table 1. We also examined the impact of training with one or multiple different sites, always testing in an unseen site.

Additionally, beyond the aforementioned set-ups, we tested the scenario where a center or a clinic cannot share raw data from patients. Thus, L0 regional predictions could be estimated directly in the test site (from the test set), using a k-fold CV scheme to get out-of-sample (OOS) predictions and pass these predictions to L1 models, trained either in pooled or site-specific data. Note that the L1 models did not include data from the test site for training. We tested this scenario, using a 3-fold CV to get L0 predictions from the test set.

To benchmark against the conventional standard, we computed the mean of GMV for each region and trained a GLMnet model to predict age. This is equivalent to replacing each of the 873 first level models with an averaging function. We implemented both pooling as well as different set ups of data/model fusion for the GMV mean data. In order

Table 1

All the setups that we compared, given the possible combinations that can occur when creating site-specific models or pooling data from all sites in the two levels of SE models. The subscripts *s* and *p* represent "site" and "pooled," respectively, indicating whether training was conducted on a per-site basis or using pooled data.

Train datasets	L0 operation	L0 train data/output	L1 operation	L1 train data/output	Notation
per site	$f_{1-873}(x) = \text{mean}$	per site	$f_{L1}(x) = \text{GLMnet per site, mean sites}$	three outputs Mean	GMV_sL1_s
	$f_{1-873}(x) = \text{GLMnet}$	per site	$f_{L1}(x) = \text{GLMnet per site, mean sites}$	three outputs Mean	PredL0_sL1_s
	OOS on test set	per site	$f_{L1}(x) = \text{GLMnet per site, mean sites}$	three outputs Mean	OOSPred_sL1_s
	$f_{1-873}(x) = \text{mean}$	pool preds	$f_{L1}(x) = \text{GLMnet on pooled L0 preds}$	one output	GMV_pL1_p^a
	$f_{1-873}(x) = \text{GLMnet}$	per site	$f_{L1}(x) = \text{GLMnet on pooled L0}$	one output	PredL0_sL1_p
	OOS on test set	per site	$f_{L1}(x) = \text{GLMnet on pooled L0}$	one output	OOSPred_sL1_p
pooled	$f_{1-873}(x) = \text{mean}$	pool preds	$f_{L1}(x) = \text{GLMnet}$	one output	GMV_pL1_p^a
	$f_{1-873}(x) = \text{GLMnet}$	pool preds	$f_{L1}(x) = \text{GLMnet}$	one output	PredL0_pL1_p
	OOS on test set	pool preds	$f_{L1}(x) = \text{GLMnet}$	one output	OOSPred_pL1_p

^a These models are the exact same, as the L0 is local averaging and has nothing to do with the train data.

to have a fair comparison with the set-up where we use the test set in a 3-fold scheme in L0, we also implemented an approach for the mean GMV models, where we trained models by pooling the training data sets with the two training folds from the test site.

In all the scenarios, the train and test datasets were exclusive to avoid data leakage.

The utility of using additional datasets was evaluated by comparing the performances of the different setups when using a different number of training data sets. For this, we examined the cases where one, two and three datasets were used to train the models in all possible combinations occurring from our four datasets.

For performance evaluation, we calculated the mean absolute error (MAE), coefficient of determination (R^2), Pearson's correlation between real age and predicted age (r), and prediction bias. The prediction bias is calculated as the correlation between the actual ages and the prediction errors (real age - predicted age) and it expresses the tendency of a model to provide predictions closer to the mean, meaning that younger subjects are predicted as older and older are predicted as younger [35,55]. Although it is not a standard performance metric, it is important for brain-age prediction problems. These metrics are widely used and well described in MRI-based age-prediction problems [56].

2.4. Biological insights

To understand biological aging it is important to identify regions that are strongly and robustly correlated with the real age. Although correlation doesn't imply causation, such insights can shed light on the brain-aging process. To this end, we calculated Pearson's correlation of real age with L0 predictions and mean regional GMV across subjects. These values were compared to identify the more robust approach providing biological insights. Projecting those correlations back to the brain space reveals regions strongly and robustly associated with aging. We ensured the stability of those results by evaluating all four datasets to show that the insights gained are not dataset dependent. This comprehensive analysis was aimed to provide a holistic perspective on our method's robustness and stability.

2.5. Data privacy

To assess how SE compares to conventional methods in data anonymity and privacy, we conducted a multiclass classification task to predict the dataset of origin. Identifying the datasets is a step closer to identifying the subject itself and thus this task serves as a proxy for gauging the potential for privacy violation. For this task, we employed L0 regional age predictions as features and juxtaposed the dataset prediction performance with that achieved using regional mean gray matter volume (GMV). We performed model selection with hyperparameter tuning for both sets of features. The tested models included logistic regression, linear Support Vector Machine (SVM), SVM with radial basis function kernel, Random Forest (RF). The exact hyperparameter space is presented in the Supplementary Table 1. The accuracy of the optimal model serves as a metric, where higher accuracy indicates better preservation of dataset-specific information in the corresponding feature space. Conversely, lower accuracy in the dataset prediction task suggests a higher level of privacy. This analysis permits to directly compare how SE preserves data anonymity in contrast to regional GMV.

3. Results

3.1. Performance

The performance of different set ups was evaluated using leave-one-site-out (LOSO) scheme. We calculated MAE, R^2 , Pearson's correlation between predicted and real values and prediction bias (Fig. 2). A comparative analysis of 3-fold and 10-fold CV yielded comparable results (Supplementary Fig. 2). We chose the 3-fold approach as it is




Fig. 2. Performance of SE models and GMV approaches for chronological age prediction, across four metrics. Performance in terms of MAE (upper left), Pearson's correlation (upper right), R^2 (bottom left) and bias (bottom right) for all setups. Results were extracted by using three datasets for training and one for testing in all possible combinations. Setups that perform OOS prediction on the test set yield better performance in all metrics except for bias.

computationally more efficient. Overall, SE frameworks showed better performance. Specifically, the highest performances were observed for the setups where L0 predictions were obtained from the test site: **OOSPredSL1_p**² (average MAE = 4.75), closely followed by **OOSPredSL1_s** (MAE = 4.9). Setups using pooled L0 predictions to train a single L1 model, independently of how L0 was trained, i.e. **PredLOpL1_p** and **PredLOsL1_p**, showed MAE = 5.1. When using mean region-wise GMV, the model trained using the training data sets together with the 2 training folds from the test site, **GMVpL1_pext**, showed MAE = 5.7. Performance was worse for the other two region-wise mean GMV models trained in three datasets, with the L1p setup (pooled predictions from L0 to train L1) being slightly better compared to L1s (train L1 models from different sites), MAE = 6.2 and MAE = 6.7 respectively. The performance ranking was the same for R^2 , slightly changed in Pearson's r , where the best model was the one where both L0 and L1 models were trained on pooled data. The performance of the models regarding age bias was rather different. The lowest bias was found for **GMVpL1_p** followed by **PredLOpL1_p** and **OOSPredSL1_p** ($b = -0.41$, $b = -0.43$ and $b = -0.44$, respectively). Detailed results can be found in [Supplementary Material Table 2](#).

All setups exhibited significant performance improvement with the inclusion of more datasets in the training process ([Fig. 3](#)). Higher improvements in performance were observed when transitioning from one to two datasets, compared to the increase from two to three. This trend underscores the positive impact of creating the training data with multiple datasets, particularly in the initial stages of model development.

3.2. Biological insights

Correlation values between mean regional GMVs and chronological age across subjects is commonly used for identifying aging-related brain regions. Similar insights from SE models can be derived using the regional age-predictions provided by the L0 models. The results showed that regional predictions of the L0 models exhibit a stronger alignment to age ([Fig. 4](#)). Correlation coefficients' mean (mean of absolute values) for L0 predictions-age was $r_{\text{mean}} = 0.6$ and $r_{\text{mean}} = 0.32$ respectively for GMV-age, yielding significant difference ($p < 0.01e^{13}$). Frontal and cerebellar areas were highlighted in both methods, notably stronger though in SE models. Remarkably, subcortical areas which exhibited positive mean GMV-age correlation, demonstrated a more pronounced

and positive association in SE models. Specifically, four regions had positive GMV-age correlation for all datasets, with the mean across datasets ranging from $r = 0.12$ to $r = 0.29$. In comparison, the corresponding L0 predicted-age to real-age correlations were notably higher, ranging from $r = 0.62$ to $r = 0.8$.

Next, we sought to compare the stability of the biological insights provided by the two methods across different datasets. To this end, we calculated the correlation values between age and L0 regional age OOS-predictions across datasets and compared it to the correlation values calculated between age and regional mean GMV. Correlations between age and regional L0 age predictions demonstrated higher consistency across all datasets compared to the correlations of age and mean GMV ([Fig. 5](#)). Cortical regions exhibited an average L0 predictions-age correlation of $r = 0.54$, while the GMV-age correlation was negative at $r = -0.4$. In subcortical regions, the mean L0 predictions-age correlation was $r = 0.75$, with the mean GMV-age correlation (absolute values) at $r = 0.35$. Interestingly, some regions showed positive GMV-real age correlations; for instance, some thalamic regions had correlations of $r = 0.28$ and $r = 0.29$ when L0 predictions and real age correlations in these regions were $r = 0.8$ and $r = 0.62$, respectively. For the cerebellum, the mean L0 predictions-age correlation was $r = 0.60$, while the GMV-age mean was $r = -0.38$.

3.2.1. Stability across datasets

To mitigate the possibility that the associations between subjects' age and L0 predictions or GMV, are driven by datasets' idiosyncratic properties, we computed the correlation values separately for each dataset. The results ([Fig. 6](#)) consistently exhibited the same pattern across datasets, affirming the overall stronger association of L0 (mean = 0.86 ranging from 0.79 to 0.91) regional age predictions with age compared to mean regional GMV (mean = 0.81 ranging from 0.75 to 0.9).

3.2.2. Application to clinical data

To assess clinical applicability, we evaluated our two best-performing model set-ups, **OOSPredSL1_p** and **OOSPredSL1_s**, in an independent clinical cohort of CN, MCI and AD subjects. Consistent with clinical data common practice, a bias correction model was applied whose parameters were derived from the CN training subjects excluded from the age- and sex-matched test set ([57], [Suppl. Figs. 3-4](#)). For both set-ups, the brain-age gap (BAG, the difference between predicted age and real age) mirrored expected cognitive decline trends, with statistically significant differences observed among all diagnostic groups (CN vs. MCI vs. AD, [Fig. 7](#)).

² The subscripts *s* and *p* represent "site" and "pooled," respectively, indicating whether training was conducted on a per-site basis or using pooled data.




Fig. 3. Comparison of setups' performances across different numbers of train datasets. The impact of the number of sites in the training process in terms of MAE (upper left), Pearson's correlation between predicted and real values (upper right), R^2 (bottom left) and bias (bottom right). Increase in the number of training sets yields better performance for MAE and R^2 . Pearson's correlation demonstrates small improvement and bias seems to be dependent on the setup.




Fig. 4. Regional GMV-age and L0 predictions-age Pearson's correlations, presented in violin plots and projected in the brain. The left panel illustrates Pearson's correlations, across subjects, between age and negated mean GMV per region (green), and that with L0 voxelwise predictions per region (orange). L0 predictions were more correlated to age, but more importantly all regions are positively correlated to age. The right panel demonstrates the same correlation values in the brain space. Positive correlation for GMV is found in the subcortical areas (upper-right), where the correlation of age-L0 predictions appears to be very high (bottom-right).

3.3. Privacy

For estimating whether SE can enhance the data privacy compared to the baseline of using mean GMV, we performed a multiclass classification task. The objective here was to predict the dataset of origin for each subject using either the regional L0 age predictions - out-of-sample estimations within each dataset-, or regional mean GMV as features. We performed a nested 5-fold cross validation with model selection and hyperparameter tuning happening within the inner CV. For both feature spaces linear SVM was chosen resulting in balanced accuracy $ACC_{bal} = 0.63$ for SE L0 predictions and $ACC_{bal} = 0.87$ for mean GMV. The GMV model showed an overall higher confusion while the L0 predictions model was biased towards two sites (Fig. 8).

4. Discussion

Understanding and modeling healthy brain aging is crucial for

developing individualized precision methods for various neurodegenerative and pathological brain disorders. Previous work has shown that advanced MRI-derived brain-age, i.e. a higher predicted age than actual age, is associated with neurodegenerative, psychiatric and other diseases [12,13,15,16,18]. Nevertheless, several important aspects need to be refined prior to integrating brain-age prediction into clinical practice: 1) improved performance and robustness, 2) interpretability and biological insights from prediction models, and 3) improved data privacy. In this context, we propose a two-level stacking model for MRI-based age-prediction. This novel approach performs voxel-based age predictions for each predefined brain region in its first level, which are then fused by a second-level model. Compared to the oftused baseline approach of averaging regional voxels' GMV, our approach offers a more sophisticated data fusion. Contrary to models using regional mean GMV, SE models make better use of the voxel-wise information. Our numerical experiments using large datasets with a wide age range suggest that the proposed SE model provides more reliable and accurate results.




Fig. 5. Regional GMV-age and L0 predictions-age Pearson's correlations across datasets. Pearson correlation across subjects was estimated between regional predicted age and real age, as well as between mean regional GMV and real age. The four panels show these correlations plotted against each other for each region for four datasets. Cortical regions are depicted in yellow, subcortical in pink and cerebellar in blue. We observed that SE L0 predictions to be more aligned with the real age indicating that it better captures the aging process. Here we demonstrate only the cases when one dataset was used for training.

Furthermore, datasets can be fused in different ways forming variations of SE models. Therefore, we tested setups differing in the way that training datasets were combined at both levels of the SE. Specifically, using a LOSO scheme, we evaluated three setups: pooling data from different sites, training per site and averaging results or performing OOS predictions in the first level using the test dataset. First we compared the performance of all SE setups and models to the standard approach of using regional mean GMV, in terms of MAE, R^2 , Pearson's correlation and age bias. Overall, SE setups outperformed the regional mean GMV models in terms of MAE, R^2 and Pearson's correlation. Notably, all SE setups using a part of the test site data for L0 models to obtain out-of-sample predictions performed better than other setups in terms of MAE and R^2 (Suppl. Table 2), with those using regional GMV being the worst. In terms of Pearson's correlation, all SE setups showed $r = 0.93$ (rounded in the second decimal digit) except for PredL0PL1P (pooling data for both L0 and L1) that had a slightly better performance ($r = 0.94$), and all outperformed GMV setups ($r = 0.89$ – 0.9). This way of obtaining OOS L0 predictions effectively models the idiosyncrasies of the specific dataset. Additionally, it also offers the advantage that centers/clinics do not need to share their raw data but only L0 predictions. In this scenario, the test sites can benefit from SE models (specifically L1) trained on publicly available data while only using their own data for L0 models. In terms of the age bias, the results were more mixed. The best and worst bias were both setups using GMV, GMVPL1Pext $b = -0.49$ and GMVSL1S $b = -0.56$. However, the high bias can be addressed by using bias correction, a common practice in age-prediction which can benefit all models [55]. Although a comparison with other models would provide some important insights regarding the performance of SE, such a comparison would be unfair and potentially misleading. While SOTA models demonstrate comparable performance (see e.g. [2], Coles et al., 2018), they typically undergo rigorous

optimization through model selection and feature engineering processes which we did not do here. Our best performing model achieved MAE = 4.75 years, which compares favorably with established classical machine learning approaches for brain-age prediction from structural MRI. Similar studies report MAE = 5.004 years using SVR with GMV features [58] and MAE = 4.3–4.5 years for multi-cohort SVR ensembles [32]. While more complex pipelines using whole-brain VBM, PCA, and intensive hyperparameter tuning can achieve slightly better performance (MAE = 3.66–4.69 years [31]), our simpler GMV-only approach with default set up performs competitively within this range.

As expected, increasing the number of datasets used for training led to prediction performance improvements (Fig. 3). Especially for MAE, when using three datasets, was improved by 0.99 on average for GMV models, 0.42 for models that perform OOS predictions in level 0 and 0.45 for the other models (averaged across different train datasets and setup variations), compared to using one dataset for training. It is worth mentioning, however, that this was not the case for age bias, which demonstrates an unclear pattern in relation to the number of data sets. In fact, differences in bias increased with the number of datasets, likely due to differing age distributions across datasets. Since bias is influenced by the age distribution of the training subjects, its interpretation is challenging in a LOSO setup.

Our analysis of brain age in ADNI participants revealed distinct performance patterns across clinical groups. The brain-age gap of both set-ups differentiated the three groups. This result shows the potential of our proposed methods to capture clinically relevant differences providing an image derived biomarker and aligns with established literature [2,14] linking brain age gaps to neurodegenerative diseases.

A coherent and robust association between chronological age and regional metrics can provide valuable biological insights into the healthy aging mechanism and consequently facilitate identification of




Fig. 6. Dataset based regional GMV-age and L0 predictions-age Pearson's correlations plotted against other datasets. Pearson correlation was estimated between regional predicted age and real age, as well as between mean regional GMV and real age, across subjects. The scatter plots show these correlations against each other for each region. The plots show that SE L0 predictions are more aligned with the real age and can thus provide a better biological insight of the aging process in healthy individuals. Here we demonstrate the predictions obtained via a 3-fold CV scheme for each dataset.




Fig. 7. BAG estimates within CN, MCI, and AD groups from the ADNI dataset. The CN and MCI groups were age- and sex-matched to the AD group. Results are shown for the two best-performing configurations: OOSPedsL1p left (a) and OOSPedsL1s right (b). Significant differences (two-sample *t*-test, after correction for multiple comparisons) were observed between all groups for both configurations.




Fig. 8. Results of dataset prediction when GMV and L0 predictions are provided as features. Confusion matrices for multilabel classification task of dataset of origin prediction with L0 regional predictions (left) and regional GMV (right). A high dataset prediction performance indicates a higher potential for identification of individuals.

brain regions susceptible to neurodegenerative and psychiatric diseases [27,28]. To this end, we conducted a comparative analysis of the correlations between age and two metrics: the regional mean GMV and regional L0 age predictions. The results showed that the correlations were more pronounced for the L0 predictions, suggesting that the associated regional models provide robust biological insights and representation of healthy brain aging (Figs. 4 and 5). Additionally, this result underscores the ability of SE models to capture nuanced aging patterns in brain regions, offering a richer perspective than traditional analyses using mean GMV. Interestingly, subcortical areas showed a weak positive correlation with mean GMV, indicating that their volume increases with age which contradicts the current knowledge [3–8]. This effect could be, however, due to preprocessing artifacts [34]. Specifically, regions in thalamus had a positive GMV-real age correlation of $r = 0.28$ and $r = 0.29$, but showed high positive correlation of L0 predictions and real age, $r = 0.8$ and $r = 0.62$ respectively. Regional predictions from subcortical areas in SE showed a much higher correlation $r = 0.75$, when the mean in subcortical areas for GMV was $r = 0.35$ (mean of absolute values). This suggests that the L0 models were able to extract age-related information from individual voxels whereas their signal was diluted by the GMV averaging. Subcortical regions have been implicated in neurodegeneration processes related to Parkinson's disease and Alzheimer's disease [59,60]. Therefore, appropriate modeling of GMV in those regions regarding healthy aging is essential for clinical application of brain age. In essence, SE effectively uses information from voxels that otherwise is "lost in the crowd" during averaging. Our correlations analysis between datasets suggests that the regional age predictions from SE models maintain a robust and stable association with age across all datasets, further underlining the reliability of the observed correlations (Figs. 5 and 6). This result further supports the notion that the alignment of L0 regional age predictions with chronological age represents biological insights and is not contingent on specific characteristics of datasets. Moreover, with the SE framework individualized brain aging maps could be generated translating brain age predictions into clinically actionable biomarkers. Such maps might reveal distinct spatial patterns and help distinguish between disease subtypes and facilitate monitoring at the individual level. Future work is needed to validate and optimize such results to provide improved clinical interpretations.

Identifying the origin dataset of individuals serves as an indirect measurement of privacy. A feature space that "hides" the dataset of origin also complicates subjects' identification. Dataset identification using L0 predictions of SE proved more challenging compared to using GMV ($ACC_{bal} = 0.63$ and $ACC_{bal} = 0.87$ respectively). For both feature spaces, linear SVM demonstrated the best performance among tested

models (see Supplementary Table 1). Classification analyzes to predict the dataset suggest, to a certain extent, that the efficacy of SE concealing more information essential for identifying the dataset and consequently a subject's identity compared to GMV. This privacy preservation could be further improved by employing more complex models in the first level, such as tree-based models and deep neural networks. However, it is noteworthy that despite the high misclassification between subjects from IXI and eNKI, for CamCAN data and mostly for 1000Brains data the classification was accurate. A potential explanation for that can be the differences in data quality as well as the generally older population of these datasets. Nonetheless, this analysis can serve as a blueprint for other use cases helping the general radiomics community in developing privacy preserving and accurate methods. A potential improvement to the SE, aligned with privacy protection principles, would involve clinics sharing their first-level predictions and/or models with a remote central unit which could then train a comprehensive second-level model. This idea is similar to the approach used in federated learning in spirit and further architectures that combine it with SE could be developed [61].

This study has several limitations that warrant discussion. First, the GLMnet model was chosen due to its interpretability and computational efficiency. We acknowledge that alternative learning algorithms (e.g., random forests, Gaussian process) might yield better performance and additional insights. Second, our analysis was restricted to the Schaefer-800 parcellation, while other atlases or granularities should be evaluated. We used 3-fold CV to generate out-of-sample predictions for the stacking model. Future studies could test other choices of number of folds. These choices were deliberate to prioritize the feasibility testing of our method and future work is needed to identify optimal choices. Taken together, our comprehensive analyses allowed for a thorough understanding of the SE model generalizability and its ability to capture meaningful biological patterns. Future work could consider the application of more, and perhaps more suitable, models through a model selection process together with a more thorough hyperparameter tuning in both stacking ensemble levels. Such refinements will contribute to the overall effectiveness of the ensemble. At first, by improving the diversity and individual performance of base learners, and then by optimizing the combination of their predictions by the meta learner and refining the final prediction, leading to more robust and accurate models which will facilitate clinical application. As performance increased with training data size, including more sites in the training data we can further improve the performance of our SE model. A thorough exploration of model selection, parcellation atlas, and large-scale validation require significantly more computational resources and warrants a separate, carefully designed study. Such carefully designed and well-tuned models will pave the path for broader clinical applications and facilitate

research into specific neurological disorders, including Parkinson's disease and other neurodegenerative conditions. Additionally, the final performance can be further enhanced by applying age-bias-correction in second level outcomes.

5. Conclusion

Here, we introduced a novel two-level SE model for MRI-based brain-age prediction. Our approach demonstrates superior performance and robustness compared to the conventional method of using regional mean GMV. Although all SE configurations outperformed the mean GMV models, the highest accuracy was obtained in configurations where level-0 predictions were derived from the test site, thereby optimizing the utilization of the available data. The SE framework not only achieves competitive predictive accuracy but also offers enhanced biological interpretability by capturing nuanced, voxel-wise aging patterns that are lost during averaging. Furthermore, the model shows potential for application to new data and improved data privacy, as it relies on sharing regional predictions rather than raw data, aligning with principles of federated learning. Crucially, the model's applicability to clinical data was demonstrated in an independent cohort, where the brain-age gap derived from our best-performing setups successfully differentiated between cognitively normal, mild cognitive impairment, and Alzheimer's disease subjects, aligning with expected cognitive decline trends. This demonstrates the potential of our SE model as a clinically relevant biomarker for neurodegenerative progression. Overall, our SE model is powerful for modeling healthy brain aging and holds potential for future clinical application as a biomarker for neurodegenerative and psychiatric diseases.

CRedit authorship contribution statement

Georgios Antonopoulos: Writing – review & editing, Writing – original draft, Visualization, Methodology, Formal analysis, Data curation. **Shammi More:** Writing – review & editing, Data curation. **Simon B. Eickhoff:** Writing – review & editing, Funding acquisition. **Federico Raimondo:** Writing – review & editing, Visualization. **Kaustubh R. Patil:** Writing – review & editing, Methodology, Funding acquisition, Formal analysis, Conceptualization.

Ethics statement

Ethical approval and informed consent were obtained locally for each study covering both participation and subsequent data sharing. The ethics proposals for the use and retrospective analyses of the datasets were approved by the Ethics Committee of the Medical Faculty at the Heinrich-Heine-University Düsseldorf.

Data/code availability statement

The codes used for preprocessing, feature extraction and model training are available at https://github.com/juaml/stacking_ensemble_paper.

Ethics statement

Ethical approval and informed consent were obtained locally for each study covering both participation and subsequent data sharing. The ethics proposals for the use and retrospective analyses of the datasets were approved by the Ethics Committee of the Medical Faculty at the Heinrich-Heine-University Düsseldorf.

Declaration of competing interest

KRP has a pending patent application DE 102022114713A1 filed on June 10, 2022.

Acknowledgements

This study was partly supported by the Helmholtz-AI project BrainAge4AD (ZT-I-PF-5-163), and the Helmholtz Portfolio Theme “Supercomputing and Modeling for the Human Brain”.

The clinical data used in the preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). The ADNI was launched in 2003 as a public-private partnership, led by Principal Investigator Michael W. Weiner, MD. The primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimer's disease (AD). Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: AbbVie, Alzheimer's Association; Alzheimer's Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer's Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.combiomed.2025.111182>.

References

- [1] 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, I.J. Deary, Brain age predicts mortality, *Mol. Psychiatr.* 23 (5) (2018) 5, <https://doi.org/10.1038/mp.2017.62>.
- [2] S. More, G. Antonopoulos, F. Hoffstaedt, J. Caspers, S.B. Eickhoff, K.R. Patil, Brain-age prediction: a systematic comparison of machine learning workflows, *Neuroimage* 270 (2023) 119947, <https://doi.org/10.1016/j.neuroimage.2023.119947>.
- [3] D.V. Callaert, A. Ribbens, F. Maes, S.P. Swinnen, N. Wenderoth, Assessing age-related gray matter decline with voxel-based morphometry depends significantly on segmentation and normalization procedures, *Front. Aging Neurosci.* 6 (2014). <https://www.frontiersin.org/article/10.3389/fnagi.2014.00124>.
- [4] S. Ramanoël, E. Hoyau, L. Kauffmann, F. Renard, C. Pichat, N. Boudiaf, A. Krainik, A. Jaillard, M. Baci, Gray matter volume and cognitive performance during normal aging. A voxel-based morphometry study, *Front. Aging Neurosci.* 10 (2018). <https://www.frontiersin.org/articles/10.3389/fnagi.2018.00235>.
- [5] D. Terribilli, M.S. Schauffelberger, F.L.S. Duran, M.V. Zanetti, P.K. Curiati, P. R. Menezes, M. Scazufca, E. Amaro, C.C. Leite, G.F. Busatto, Age-related gray matter volume changes in the brain during non-elderly adulthood, *Neurobiol. Aging* 32 (2–6) (2011) 354–368, <https://doi.org/10.1016/j.neurobiolaging.2009.02.008>.
- [6] I. Driscoll, C. Davatzikos, Y. An, X. Wu, D. Shen, M. Kraut, S.M. Resnick, Longitudinal pattern of regional brain volume change differentiates normal aging from MCI, *Neurology* 72 (22) (2009) 1906–1913, <https://doi.org/10.1212/WNL.0b013e3181a82634>.

- [7] A. Giorgio, L. Santelli, V. Tomassini, R. Bosnell, S. Smith, N. De Stefano, H. Johansen-Berg, Age-related changes in grey and white matter structure throughout adulthood, *Neuroimage* 51 (3) (2010) 943–951, <https://doi.org/10.1016/j.neuroimage.2010.03.004>.
- [8] D.H. Salat, J.A. Kaye, J.S. Janowsky, Prefrontal gray and white matter volumes in healthy aging and alzheimer disease, *Arch. Neurol.* 56 (3) (1999) 338–344, <https://doi.org/10.1001/archneur.56.3.338>.
- [9] C. Davatzikos, F. Xu, Y. An, Y. Fan, S.M. Resnick, Longitudinal progression of alzheimer's-like patterns of atrophy in normal older adults: the SPARE-AD index, *Brain* 132 (8) (2009) 2026–2035, <https://doi.org/10.1093/brain/awp091>.
- [10] N. Koutsouleris, C. Davatzikos, S. Borgwardt, C. Gaser, R. Bottlender, T. Frodl, P. Falkai, A. Riecher-Rössler, H.-J. Möller, M. Reiser, C. Pantelis, E. Meisenzahl, Accelerated brain aging in schizophrenia and beyond: a neuroanatomical marker of psychiatric disorders, *Schizophr. Bull.* 40 (5) (2014) 1140–1153, <https://doi.org/10.1093/schbul/sbt142>.
- [11] F. Liem, G. Varoquaux, J. Kynast, F. Beyer, S. Kharabian Masouleh, J. M. Huntenburg, L. Lampe, M. Rahim, A. Abraham, R.C. Craddock, S. Riedel-Heller, T. Luck, M. Loeffler, M.L. Schroeter, A.V. Witte, A. Villringer, D.S. Margulies, Predicting brain-age from multimodal imaging data captures cognitive impairment, *Neuroimage* 148 (2017) 179–188, <https://doi.org/10.1016/j.neuroimage.2016.11.005>.
- [12] J. Wang, M.J. Knol, A. Tiulpin, F. Dubost, M. de Bruijne, M.W. Vernooij, H.H. Adams, M.A. Ikram, W.J. Niessen, G.V. Roshchupkin, Gray matter age prediction as a biomarker for risk of dementia, *Proc. Natl. Acad. Sci.* 116 (42) (2019) 21213–21218, <https://doi.org/10.1073/pnas.1902376116>.
- [13] T. Wyss-Coray, Ageing, neurodegeneration and brain rejuvenation, *Nature* 539 (7628) (2016) 180–186, <https://doi.org/10.1038/nature20411>.
- [14] J.H. Cole, K. Franke, Predicting Age using neuroimaging: innovative brain ageing biomarkers, *Trends Neurosci.* 40 (12) (2017) 681–690, <https://doi.org/10.1016/j.tins.2017.10.001>.
- [15] A.M. Fjell, L. McEvoy, D. Holland, A.M. Dale, K.B. Walhovd, What is normal in normal aging? Effects of aging, amyloid and Alzheimer's disease on the cerebral cortex and the hippocampus, *Prog. Neurobiol.* 117 (2014) 20–40, <https://doi.org/10.1016/j.pneurobio.2014.02.004>.
- [16] C.D. Good, I.S. Johnsrude, J. Ashburner, R.N. Henson, K.J. Friston, R. S. Frackowiak, A voxel-based morphometric study of ageing in 465 normal adult human brains, *Neuroimage* 14 (1 Pt 1) (2001) 21–36, <https://doi.org/10.1006/nimg.2001.0786>.
- [17] T. Hajek, K. Franke, M. Kolenic, J. Capkova, M. Matejka, L. Propper, R. Uher, P. Stopkova, T. Novak, T. Paus, M. Kopecek, F. Spaniel, M. Alda, Brain Age in early stages of bipolar disorders or schizophrenia, *Schizophr. Bull.* 45 (1) (2019) 190–198, <https://doi.org/10.1093/schbul/sbx172>.
- [18] G.B. Karas, P. Scheltens, S.A.R.B. Rombouts, P.J. Visser, R.A. van Schijndel, N. C. Fox, F. Barkhof, Global and local gray matter loss in mild cognitive impairment and Alzheimer's disease, *Neuroimage* 23 (2) (2004) 708–716, <https://doi.org/10.1016/j.neuroimage.2004.07.006>.
- [19] D. Eke, I.E.J. Aasebo, S. Akintoye, W. Knight, A. Karakasidis, E. Mikulan, P. Ochang, G. Ogoh, R. Oostenveld, A. Pigorini, B.C. Stahl, T. White, L. Zehl, Pseudonymisation of neuroimages and data protection: increasing access to data while retaining scientific utility, *Neuroimage: Report 1* (4) (2021) 100053, <https://doi.org/10.1016/j.nyrp.2021.100053>.
- [20] K. Franke, G. Ziegler, S. Klöppel, C. Gaser, Estimating the age of healthy subjects from T1-weighted MRI scans using kernel methods: exploring the influence of various parameters, *Neuroimage* 50 (3) (2010) 883–892, <https://doi.org/10.1016/j.neuroimage.2010.01.005>.
- [21] J.H. Cole, R.P.K. Poudel, D. Tsagkrasoulis, M.W.A. Caan, C. Steves, T.D. Spector, G. Montana, Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker, *Neuroimage* 163 (2017) 115–124, <https://doi.org/10.1016/j.neuroimage.2017.07.059>.
- [22] N.K. Dinsdale, E. Bluemke, S.M. Smith, Z. Arya, D. Vidaurre, M. Jenkinson, A.I. L. Namburete, Learning patterns of the ageing brain in MRI using deep convolutional networks, *Neuroimage* 224 (2021) 117401, <https://doi.org/10.1016/j.neuroimage.2020.117401>.
- [23] B.A. Jonsson, G. Bjornsdottir, T.E. Thorgerisson, L.M. Ellingsen, G.B. Walters, D. F. Gudbjartsson, H. Stefansson, K. Stefansson, M.O. Ulfarsson, Brain age prediction using deep learning uncovers associated sequence variants, *Nat. Commun.* 10 (1) (2019), <https://doi.org/10.1038/s41467-019-13163-9>, 1.
- [24] J. Lee, B.J. Burkett, H.-K. Min, M.L. Senjem, E.S. Lundt, H. Botha, J. Graff-Radford, L.R. Barnard, J.L. Gunter, C.G. Schwarz, K. Kantarci, D.S. Knopman, B.F. Boeve, V. J. Lowe, R.C. Petersen, C.R. Jack, D.T. Jones, Deep learning-based brain age prediction in normal aging and dementia, *Nature Aging* 2 (5) (2022) 5, <https://doi.org/10.1038/s43587-022-00219-7>.
- [25] H. Peng, W. Gong, C.F. Beckmann, A. Vedaldi, S.M. Smith, Accurate brain age prediction with lightweight deep neural networks, *Med. Image Anal.* 68 (2021) 101871, <https://doi.org/10.1016/j.media.2020.101871>.
- [26] A. Lombardi, A. Monaco, G. Donvito, N. Amoroso, R. Bellotti, S. Tangaro, Brain Age prediction with morphological features using deep neural networks: results from predictive analytic competition 2019, *Front. Psychiatr.* 11 (2021), <https://www.frontiersin.org/articles/10.3389/fpsyt.2020.619629>.
- [27] N. Koutsouleris, C. Davatzikos, S. Borgwardt, C. Gaser, R. Bottlender, T. Frodl, P. Falkai, A. Riecher-Rössler, H.-J. Möller, M. Reiser, C. Pantelis, E. Meisenzahl, Accelerated brain aging in schizophrenia and beyond: a neuroanatomical marker of psychiatric disorders, *Schizophr. Bull.* 40 (5) (2014) 1140–1153, <https://doi.org/10.1093/schbul/sbt142>.
- [28] S.G. Popescu, B. Glocker, D.J. Sharp, J.H. Cole, Local Brain-Age: a U-Net model, *Front. Aging Neurosci.* 13 (2021), <https://www.frontiersin.org/articles/10.3389/fnagi.2021.761954>.
- [29] K. Franke, C. Gaser, T.J. Roseboom, M. Schwab, S.R. de Rooij, Premature brain aging in humans exposed to maternal nutrient restriction during early gestation, *Neuroimage* 173 (2018) 460–471, <https://doi.org/10.1016/j.neuroimage.2017.10.047>.
- [30] D.P. Varikuti, S. Genon, A. Sotiras, H. Schwender, F. Hoffstaedter, K.R. Patil, C. Jockwitz, S. Caspers, S. Moebus, K. Amunts, C. Davatzikos, S.B. Eickhoff, Evaluation of non-negative matrix factorization of grey matter in age prediction, *Neuroimage* 173 (2018) 394–410, <https://doi.org/10.1016/j.neuroimage.2018.03.007>.
- [31] L. Baecker, R. Garcia-Dias, S. Vieira, C. Scarpazza, A. Mechelli, Machine learning for brain age prediction: introduction to methods and clinical applications, *EBioMedicine* 72 (2021), <https://doi.org/10.1016/j.ebiom.2021.103600>.
- [32] C.R. Eickhoff, F. Hoffstaedter, J. Caspers, K. Reetz, C. Mathys, I. Dogan, K. Amunts, A. Schnitzler, S.B. Eickhoff, Advanced brain ageing in Parkinson's disease is related to disease duration and individual impairment, *Brain Communications* 3 (3) (2021) fcab191, <https://doi.org/10.1093/braincomms/fcab191>.
- [33] S.B. Eickhoff, B.T.T. Yeo, S. Genon, Imaging-based parcellations of the human brain, *Nat. Rev. Neurosci.* 19 (11) (2018) 11, <https://doi.org/10.1038/s41583-018-0071-7>.
- [34] G. Antonopoulos, S. More, F. Raimondo, S.B. Eickhoff, F. Hoffstaedter, K.R. Patil, A systematic comparison of VBM pipelines and their application to age prediction, *Neuroimage* 279 (2023) 120292, <https://doi.org/10.1016/j.neuroimage.2023.120292>.
- [35] C. Wachinger, A. Rieckmann, S. Pölsterl, Detect and correct bias in multi-site neuroimaging datasets, *Med. Image Anal.* 67 (2021) 101879, <https://doi.org/10.1016/j.media.2020.101879>.
- [36] P. Gooding, T. Kariotis, Ethics and law in research on algorithmic and data-driven technology in mental health care: scoping review, *JMIR Ment. Health* 8 (6) (2021) e24668, <https://doi.org/10.2196/24668>.
- [37] C.G. Schwarz, W.K. Kremers, T.M. Therneau, R.R. Sharp, J.L. Gunter, P. Vemuri, A. Arani, A.J. Spychalla, K. Kantarci, D.S. Knopman, R.C. Petersen, C.R. Jack, Identification of anonymous MRI research participants with face-recognition software, *N. Engl. J. Med.* 381 (17) (2019) 1684–1686, <https://doi.org/10.1056/NEJMc1908881>.
- [38] E.S. Finn, X. Shen, D. Scheinost, M.D. Rosenberg, J. Huang, M.M. Chun, X. Papademetris, R.T. Constable, Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity, *Nat. Neurosci.* 18 (11) (2015) 11, <https://doi.org/10.1038/nn.4135>.
- [39] D.H. Wolpert, Stacked generalization, *Neural Netw.* 5 (2) (1992) 241–259, [https://doi.org/10.1016/S0893-6080\(05\)80023-1](https://doi.org/10.1016/S0893-6080(05)80023-1).
- [40] H. Cao, Y. Gu, J. Fang, Y. Hu, W. Ding, H. He, G. Chen, Application of stacking ensemble learning model in quantitative analysis of biomaterial activity, *Microchem. J.* 183 (2022) 108075, <https://doi.org/10.1016/j.microc.2022.108075>.
- [41] M. Liang, T. Chang, B. An, X. Duan, L. Du, X. Wang, J. Miao, L. Xu, X. Gao, L. Zhang, J. Li, H. Gao, A stacking ensemble learning framework for genomic prediction, *Front. Genet.* 12 (2021) 600040, <https://doi.org/10.3389/fgenet.2021.600040>.
- [42] X. Zhang, C. Tang, S. Wang, W. Liu, W. Yang, D. Wang, Q. Wang, F. Tang, A stacking ensemble model for predicting the occurrence of carotid atherosclerosis, *Front. Endocrinol.* 15 (2024), <https://doi.org/10.3389/fendo.2024.1390352>.
- [43] P. Kal, R. Dahnke, F. Hoffstaedter, C. Gaser, A.D.N. Initiative, BrainAGE: revisited and reframed machine learning workflow, *Hum. Brain Mapp.* 45 (3) (2024) e26632, <https://doi.org/10.1002/hbm.26632>.
- [44] J.R. Taylor, N. Williams, R. Cusack, T. Auer, M.A. Shafto, M. Dixon, L.K. Tyler, Cam-Can, R.N. Henson, The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) data repository: structural and functional MRI, MEG, and cognitive data from a cross-sectional adult lifespan sample, *Neuroimage* 144 (2017) 262–269, <https://doi.org/10.1016/j.neuroimage.2015.09.018>.
- [45] K.B. Nooner, S.J. Colcombe, R.H. Tobe, M. Mennes, M.M. Benedict, A.L. Moreno, L. J. Panek, S. Brown, S.T. Zavitz, Q. Li, S. Sikka, D. Gutman, S. Bangaru, R. T. Schlachter, S.M. Kamel, A.R. Anwar, C.M. Hinz, M.S. Kaplan, A.B. Rachlin, M. P. Milham, The NKI-Rockland sample: a model for accelerating the pace of discovery science in psychiatry, *Front. Neurosci.* 6 (2012) 152, <https://doi.org/10.3389/fnins.2012.00152>.
- [46] S. Caspers, S. Moebus, S. Lux, N. Pundt, H. Schütz, T.W. Mühleisen, V. Gras, S. B. Eickhoff, S. Romanzetti, T. Stöcker, R. Stirnberg, M.E. Kirlangic, M. Minnerop, P. Pieperhoff, U. Mödder, S. Das, A.C. Evans, K.-H. Jöckel, R. Erbel, K. Amunts, Studying variability in human brain aging in a population-based German cohort—Rationale and design of 1000BRAINS, *Front. Aging Neurosci.* 6 (2014), <https://www.frontiersin.org/articles/10.3389/fnagi.2014.00149>.
- [47] C.R. Jack, M.A. Bernstein, N.C. Fox, P. Thompson, G. Alexander, D. Harvey, B. Borowski, P.J. Britson, J.L. Whitwell, C. Ward, A.M. Dale, J.P. Felmlee, J. L. Gunter, D.L.G. Hill, R. Killiany, N. Schuff, S. Fox-Bosetti, C. Lin, C. Studholme, C. S. DeCarli, G. Krueger, H.A. Ward, G.J. Metzger, K.T. Scott, R. Mallozzi, D. Blezek, J. Levy, J.P. Debbins, A.S. Fleisher, M. Albert, R. Green, G. Bartzokis, G. Glover, J. Mugler, M.W. Weiner, The Alzheimer's disease neuroimaging initiative (ADNI): MRI methods, *J. Magn. Reson. Imag.* 27 (4) (2008) 685–691, <https://doi.org/10.1002/jmri.21049>.
- [48] R.C. Petersen, P.S. Aisen, L.A. Beckett, M.C. Donohue, A.C. Gamst, D.J. Harvey, C. R. Jack, W.J. Jagust, L.M. Shaw, A.W. Toga, J.Q. Trojanowski, M.W. Weiner, Alzheimer's Disease Neuroimaging Initiative (ADNI): clinical characterization,

- Neurology 74 (3) (2010) 201–209, <https://doi.org/10.1212/WNL.0b013e3181cb3e25>.
- [49] C. Gaser, R. Dahnke, P.M. Thompson, F. Kurth, E. Luders, the Alzheimer's Disease Neuroimaging Initiative, CAT: a computational anatomy toolbox for the analysis of structural MRI data, *GigaScience* 13 (2024) giae049, <https://doi.org/10.1093/gigascience/giae049>.
- [50] J. Ashburner, K.J. Friston, Diffeomorphic registration using geodesic shooting and Gauss-Newton optimisation, *Neuroimage* 55 (3) (2011) 954–967, <https://doi.org/10.1016/j.neuroimage.2010.12.049>.
- [51] J.H. Friedman, T. Hastie, R. Tibshirani, Regularization paths for generalized Linear models via coordinate descent, *J. Stat. Software* 33 (2010) 1–22, <https://doi.org/10.18637/jss.v033.i01>.
- [52] A. Schaefer, R. Kong, E.M. Gordon, T.O. Laumann, X.-N. Zuo, A.J. Holmes, S. B. Eickhoff, B.T.T. Yeo, Local-Global parcellation of the Human cerebral cortex from intrinsic functional connectivity MRI, *Cerebr. Cortex* 28 (9) (2018) 3095–3114, <https://doi.org/10.1093/cercor/bhx179>.
- [53] L. Fan, H. Li, J. Zhuo, Y. Zhang, J. Wang, L. Chen, Z. Yang, C. Chu, S. Xie, A. R. Laird, P.T. Fox, S.B. Eickhoff, C. Yu, T. Jiang, The human brainnetome Atlas: a new brain Atlas based on connectional Architecture, *Cerebr. Cortex* 26 (8) (2016) 3508–3526, <https://doi.org/10.1093/cercor/bhw157>.
- [54] R.L. Buckner, F.M. Krienen, A. Castellanos, J.C. Diaz, B.T.T. Yeo, The organization of the human cerebellum estimated by intrinsic functional connectivity, *J. Neurophysiol.* 106 (5) (2011) 2322–2345, <https://doi.org/10.1152/jn.00339.2011>.
- [55] I. Beheshti, S. Nugent, O. Potvin, S. Duchesne, Bias-adjustment in neuroimaging-based brain age frameworks: a robust scheme, *Neuroimage: Clinical* 24 (2019) 102063, <https://doi.org/10.1016/j.nicl.2019.102063>.
- [56] A.-M.G. de Lange, M. Anatiürk, J. Rokicki, L.K.M. Han, K. Franke, D. Alnaes, K. P. Ebmeier, B. Draganski, T. Kaufmann, L.T. Westlye, T. Hahn, J.H. Cole, Mind the gap: performance metric evaluation in brain-age prediction, *Hum. Brain Mapp.* 43 (10) (2022) 3113–3129, <https://doi.org/10.1002/hbm.25837>.
- [57] J.H. Cole, Multimodality neuroimaging brain-age in UK biobank: relationship to biomedical, lifestyle, and cognitive factors, *Neurobiol. Aging* 92 (2020) 34–42, <https://doi.org/10.1016/j.neurobiolaging.2020.03.014>.
- [58] P.F. Da Costa, J. Dafflon, W.H.L. Pinaya, Brain-Age prediction using shallow machine learning: predictive analytics competition 2019, *Front. Psychiatr.* 11 (2020) 1664, <https://doi.org/10.3389/fpsy.2020.604478>, 0640.
- [59] Y. Gupta, K.H. Lee, K.Y. Choi, J.J. Lee, B.C. Kim, G.R. Kwon, Early diagnosis of Alzheimer's disease using combined features from voxel-based morphometry and cortical, subcortical, and hippocampus regions of MRI T1 brain images, *PLoS One* 14 (10) (2019) e0222446, <https://doi.org/10.1371/journal.pone.0222446>.
- [60] S. Toshkhujaev, K.H. Lee, K.Y. Choi, J.J. Lee, G.-R. Kwon, Y. Gupta, R.K. Lama, Classification of alzheimer's disease and mild cognitive impairment based on cortical and subcortical features from MRI T1 brain images utilizing four different types of datasets, *J. Healthc. Eng.* (2020) 3743171, <https://doi.org/10.1155/2020/3743171>, 2020.
- [61] B. McMahan, E. Moore, D. Ramage, S. Hampson, B.A.y. Arcas, Communication-Efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 2017, pp. 1273–1282, in: <https://proceedings.mlr.press/v54/mcmahan17a.html>.