

# SynthBA: Reliable Brain Age Estimation Across Multiple MRI Sequences and Resolutions

1<sup>st</sup> Lemuel Puglisi

*Dept. of Math and Computer Science*  
*University of Catania*  
Catania, Italy  
lemuel.puglisi@phd.unict.it

2<sup>nd</sup> Alessia Rondinella

*Dept. of Math and Computer Science*  
*University of Catania*  
Catania, Italy  
alessia.rondinella@phd.unict.it

3<sup>rd</sup> Linda De Meo

*Department of Neuroinflammation*  
*University College London*  
London, UK  
e.demeo@ucl.ac.uk

4<sup>th</sup> Francesco Guarnera

*Dept. of Math and Computer Science*  
*University of Catania*  
Catania, Italy  
francesco.guarnera@unict.it

5<sup>th</sup> Sebastiano Battiato

*Dept. of Math and Computer Science*  
*University of Catania*  
Catania, Italy  
battiato@dmi.unict.it

6<sup>th</sup> Daniele Ravi

*School of Physics, Engineering and CS*  
*University of Hertfordshire*  
Hatfield, UK  
d.ravi@herts.ac.uk

**Abstract**—Brain age is a critical measure that reflects the biological ageing process of the brain. The gap between brain age and chronological age, referred to as brain PAD (Predicted Age Difference), has been utilized to investigate neurodegenerative conditions. Brain age can be predicted using MRIs and machine learning techniques. However, existing methods are often sensitive to acquisition-related variabilities, such as differences in acquisition protocols, scanners, MRI sequences, and resolutions, significantly limiting their application in highly heterogeneous clinical settings. In this study, we introduce Synthetic Brain Age (SynthBA), a robust deep-learning model designed for predicting brain age. SynthBA utilizes an advanced domain randomization technique, ensuring effective operation across a wide array of acquisition-related variabilities. To assess the effectiveness and robustness of SynthBA, we evaluate its predictive capabilities on internal and external datasets, encompassing various MRI sequences and resolutions, and compare it with state-of-the-art techniques. Additionally, we calculate the brain PAD in a large cohort of subjects with Alzheimer’s Disease (AD), demonstrating a significant correlation with AD-related measures of cognitive dysfunction. SynthBA holds the potential to facilitate the broader adoption of brain age prediction in clinical settings, where re-training or fine-tuning is often unfeasible. The SynthBA source code and pre-trained models are publicly available at <https://github.com/LemuelPuglisi/SynthBA>.

**Index Terms**—brain age, deep learning, MRI, domain randomization

## I. INTRODUCTION

The biological ageing process of the brain can be accelerated by various health conditions, such as neurodegenerative diseases, and by environmental factors. Recent studies have demonstrated that disentangling the biological age of the brain (referred to as brain age) from the chronological age of the individual aids in predicting disease risk [1]. The primary assumption is that the brain age of healthy individuals closely matches their chronological age. Building on this assumption, researchers have developed machine learning techniques to predict the chronological age of healthy individuals from brain MRI data as a means

to estimate their brain age [2]. When these models are applied to subjects with neurological diseases, an increased discrepancy is observed between their predicted brain age and chronological age [3]. This widening gap, referred to as brain PAD, highlights an accelerated ageing process of the brain in diseased individuals [4].

Deep learning has brought significant advancements to brain age modeling, resulting in models that surpass traditional methods in terms of accuracy. However, their sensitivity to variations in acquisition protocols, scanners, MRI sequences, and resolutions limits their practical use across diverse clinical environments. Despite extensive research done in the medical imaging analysis field to address this issue, such as using domain-specific augmentation techniques [5] or domain randomization [6], only a few studies have been conducted to tackle this challenge in the context of brain age modeling. For example, the authors of [7] develop five brain age models, each trained on a large dataset of a different MRI sequence: T1-weighted (T1w), T2-weighted (T2w), Fluid-Attenuated Inversion Recovery (FLAIR), Gradient-Recalled Echo (GRE), and Diffusion-Weighted Imaging (DWI). These models are intended to serve as foundational models that can be fine-tuned on a target dataset to improve predictive accuracy compared to training a model ex novo. While the availability of brain age models for various MRI sequences facilitates their adoption, they still require fine-tuning to achieve low errors. Another relevant work is proposed in [8], which introduces a contrastive regression loss that enhances robustness against site differences [9], including variations in scanner vendor, head coil, and acquisition parameters. Although this method achieves state-of-the-art performance in the multi-site challenge [10], it does not account for higher-level variations, such as different MRI sequences and resolutions. Lastly, [11] proposes a dual-transfer learning strategy to obtain an unbiased, protocol-agnostic brain

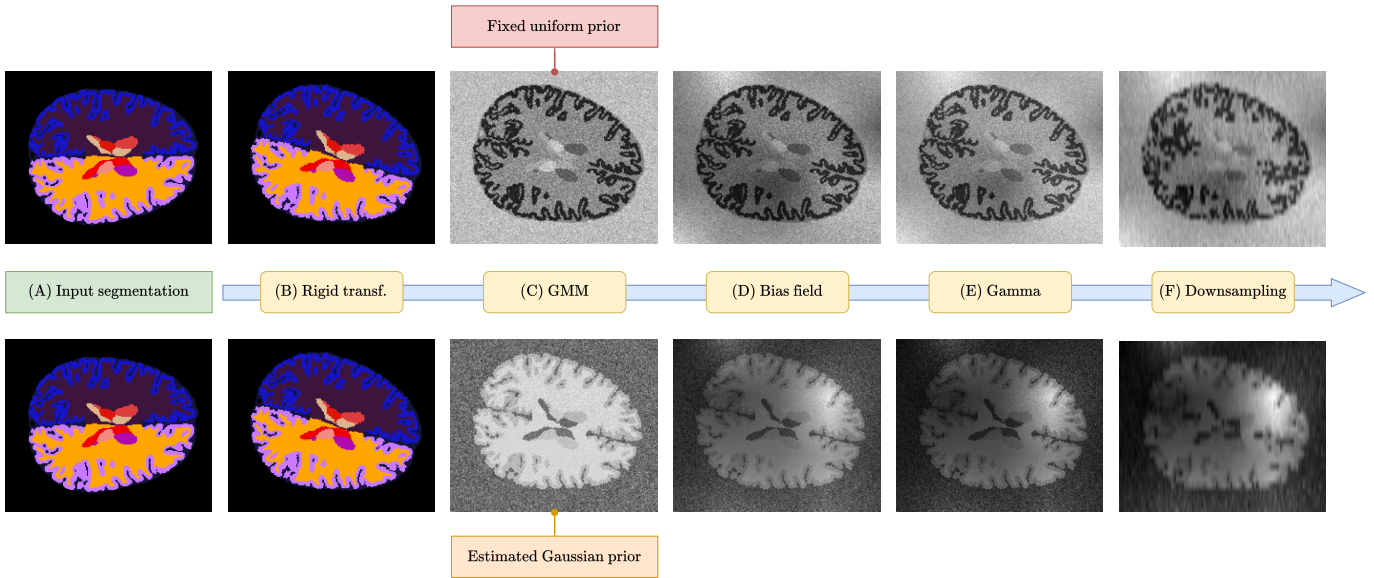


Fig. 1. All the steps from the generative process  $x \sim \mathcal{G}(s)$ . The upper branch illustrates the generative process when the GMM employs a uniform prior, while the lower branch shows the process when the GMM uses a Gaussian prior estimated from real MRIs.

age model. Nevertheless, their experimental results reveal an unfavorable trade-off between predictive accuracy and generalizability.

Our work stems from the practical necessity of computing brain age using clinical-grade MRIs, regardless of their acquisition settings, without the possibility of re-training an existing model. This scenario is typical in clinical trials, where the selected subjects are often impaired, leaving no healthy cohort available to train or fine-tune existing brain age models. This limitation restricts the applicability and, consequently, the large-scale adoption of brain age assessments.

The ability to generalize across acquisition-related variability is essential in medical imaging applications. Domain randomization [6] has recently been employed to address this issue in various MRI-related tasks. This approach involves training a deep learning model on synthetic data generated by a simulator with fully randomized parameters. Although this technique can produce unrealistic data, such as MRIs with atypical contrasts, it improves the model’s ability to generalize to previously unseen data distributions [6]. For example, the authors of [12] developed a generative model for brain MRIs and trained a brain tissue segmentation model, called SynthSeg, by fully randomizing the generator’s parameters. As a result, SynthSeg can adapt to varying contrasts and resolutions. Building on this concept, several robust models have been developed for various MRI-related tasks: SynthStrip [13] for skull-stripping, SynthSR [14] for super-resolution, and SynthMorph [15] for image registration.

Building on the intuition of the seminal work SynthSeg [12],

we propose **SynthBA**, a deep learning-based brain age model robust to acquisition-related variability. We adapted the generative model from SynthSeg [12] to produce synthetic brain MRIs specifically tailored for the downstream task of brain age prediction. Subsequently, we utilize this generative model to sample MRI scans with random contrasts and resolutions from our cohort of healthy subjects, and we train our brain age model on these synthetic images. The high variability provided by the generative model enables SynthBA to operate with different MRI sequences and resolutions without the need for re-training or fine-tuning while maintaining reasonable accuracy. We evaluate SynthBA using an external dataset comprising MRIs from three distinct MRI sequences at varying resolutions, demonstrating its robustness across various acquisition settings compared to state-of-the-art brain age models. Ultimately, we analyse how the brain PAD obtained using SynthBA from a cohort of subjects with AD correlates with measures of cognitive dysfunction.

## II. METHOD

This section describes our methodology, which involves: (i) two adjustments applied to the generative model introduced in SynthSeg [12], and (ii) the development of a neural network to estimate brain age, trained on synthetic brain MRIs produced by the modified generative model. Additionally, we describe the preprocessing steps applied to the brain MRIs used in our study.

### A. Adjustments to the generative model from SynthSeg.

The generative model presented in SynthSeg [12], referred to as  $\mathcal{G}$ , consists of a series of steps aimed at generating synthetic brain MRIs from brain tissue segmentations. This process is illustrated in Figure 1. A brain tissue segmentation, denoted as  $s$  (input A in Figure 1), is a 3D map that assigns

each voxel at coordinates  $(i, j, k)$  a label  $s_{ijk} = l$ , identifying a specific brain region.

Starting with a tissue segmentation  $s$ , the generative process applies a random spatial transformation  $\phi$  to produce a deformed segmentation  $s' = \phi(s)$  (step B from Figure 1). Originally, in SynthSeg [12],  $\phi$  includes both affine transformations and elastic deformations. The **first adjustment** we applied to the generative model is to restrict the spatial deformation to only a random rigid transformation, explicitly excluding shearing, scaling, and non-linear components. In particular, the exclusion of scaling and elastic deformations aims to prevent transformations that could mimic atrophy patterns associated with ageing, thus avoiding a potential mismatch between the subject’s chronological age and the appearance of the transformed segmentation  $s'$ .

The next step involves generating a 3D image  $x$  from the labels  $s'$  using a Gaussian Mixture Model (GMM). The GMM assigns a Gaussian distribution  $\mathcal{N}(\mu_l, \sigma_l)$  to each segmentation label  $l$  indicating a different brain region. The parameters  $(\mu_l, \sigma_l)$  are randomly selected either from a fixed uniform prior distribution (option a) or an estimated Gaussian prior distribution (option b):

$$\begin{array}{ll} \text{(option a)} & \text{(option b)} \\ \mu_l \sim U[\mu_a, \mu_b] & \mu_l \sim \mathcal{N}(\mu_{\mu,l}^k, \sigma_{\mu,l}^k) \\ \sigma_l \sim U[\sigma_a, \sigma_b] & \sigma_l \sim \mathcal{N}(\mu_{\sigma,l}^k, \sigma_{\sigma,l}^k) \end{array}$$

In SynthSeg [12], the authors provide a method to estimate the prior parameters in (option b) using real images from a specific MRI sequence. Our **second adjustment** involves expanding this feature to enable the estimation of  $K$  sets of prior parameters from  $K$  different MRI sequences. Specifically, the prior parameters  $(\mu_{\mu,l}^k, \sigma_{\mu,l}^k, \mu_{\sigma,l}^k, \sigma_{\sigma,l}^k)$  for each label  $l$  and sequence  $k = 1, \dots, K$  are estimated from brain MRIs reflecting the  $k$ -th MRI sequence. During generation, one set of parameters is randomly selected from the  $K$  available sets. As illustrated in Figure 1, utilizing the uniform prior (option a) results in unrealistic random contrasts (upper branch from Figure 1). Conversely, employing the estimated Gaussian priors (option b) constrains the training domain, yielding more realistic, albeit randomized, MRI contrasts (lower branch from Figure 1). Finally, to generate the synthetic image  $x$ , each voxel  $x_{ijk}$  with label  $s_{ijk} = l$  is sampled from  $x_{ijk} \sim \mathcal{N}(\mu_l, \sigma_l)$  (step C in Figure 1).

The generative process continues as detailed in SynthSeg [12]. Specifically, we follow the implementation in [16] to generate a random bias field, which we then apply to the synthetic image  $x$  to simulate the intensity inhomogeneity artefact (step D in Figure 1). The image intensities are then rescaled to the range  $[0, 1]$ , and a random gamma transform is applied to further enhance the heterogeneity of the contrast (step E in Figure 1). Finally, the 3D image  $x$  is resampled to a random lower resolution, which may be isotropic or

anisotropic (step F in Figure 1). Further technical details of each generation step and the parameter ranges for the random transformations are available in [12]. We use the notation  $x \sim \mathcal{G}(s)$  to summarize the process of generating a synthetic scan  $x$  from a brain tissue segmentation  $s$ .

### B. Training the Proposed Brain Age Model

Let  $s$  be a brain tissue segmentation from an MRI of a healthy individual. Using the modified generative model, we sample a synthetic brain MRI  $x \sim \mathcal{G}(s)$ . As mentioned in Section I the brain age  $y_b$  and the chronological age  $y_c$  are equivalent ( $y_b = y_c$ ) for healthy subjects. Therefore, we train a neural network  $f_\theta$ , with parameters  $\theta$ , to predict the brain age  $f_\theta(x) \approx y_b$  from the synthetic brain MRI  $x$ . Specifically, our training process aims to find the parameters  $\theta$  that minimize the  $L_1$  loss  $|y_c - f_\theta(x)|$ , which is equal to minimizing  $|y_b - f_\theta(x)|$ . The training pipeline is summarized in Figure 2. We employ a 3D DenseNet [17] as our parametric model. As the contrast of the synthetic brain MRI is randomly determined at each training iteration, we expect the neural network to estimate the brain age mainly from the structural features of the brain, potentially developing contrast-agnostic properties.

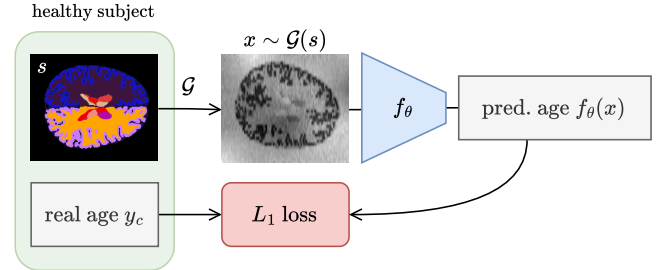


Fig. 2. SynthBA training pipeline.

### C. MRI Preprocessing

Before training, we aligned all brain MRIs to the standard  $1\text{mm}^3$  MNI space to remove variability arising from different MRI orientations. Subsequently, we used SynthSeg v.2 [12] to extract brain tissue segmentations from the brain MRIs, to be used within the generative model to sample synthetic scans. The segmentation identifies 32 different brain regions, plus the background. During the training phase, the synthetic scans are downsampled to a resolution of  $1.4\text{mm}^3$ , resized to a dimension of  $130 \times 130 \times 130$  through padding or cropping, and the intensities are rescaled in the range  $[0, 1]$ .

## III. EXPERIMENTS AND RESULTS

In this section, we first describe the datasets and training settings used in this study. We then compare SynthBA to state-of-the-art brain age models under various acquisition settings. Finally, we provide a brief clinical evaluation of the brain PAD calculated using SynthBA.

TABLE I

NUMBER OF HEALTHY SUBJECTS AND SCANS IN EACH INTERNAL (INT) AND EXTERNAL (EXT) DATASET FOR EACH MRI SEQUENCE AND THEIR RESOLUTIONS. RESOLUTIONS (R) ARE CLASSIFIED AS RESEARCH GRADE (RG) OR CLINICAL GRADE (CG).

Dataset	Type	#subj.	T1w		T2w		FLAIR	
			#scans	R.	#scans	R.	#scans	R.
ADNI	INT	784	3345	RG	-	-	-	-
AIBL	INT	192	596	RG	-	-	-	-
IXI	INT	563	563	RG	584	RG	-	-
HCP	INT	1105	1105	RG	1105	RG	-	-
CoRR	INT	1461	1461	RG	-	-	-	-
OASIS 3	EXT	811	1241	RG	2389	CG	921	CG

### A. Datasets

We collected T1w, T2w, and FLAIR brain MRIs from healthy subjects using six publicly available clinical datasets: ADNI [18], AIBL [19], OASIS 3 [20], IXI [21], HCP [22], and CoRR [23]. Table I reports the number of scans in each dataset for each MRI sequence and their resolutions. We define a research-grade MRI as an isotropic scan with a voxel resolution of approximately  $1\text{mm}^3$  or finer. In contrast, a clinical-grade MRI is defined as an anisotropic scan or a low-resolution isotropic MRI. We aggregate five datasets (ADNI, AIBL, IXI, HCP, CoRR) into a combined internal dataset comprising 4,105 healthy subjects. This dataset includes 7,060 research-grade T1w MRIs and 1,689 research-grade T2w MRIs. The subjects’ ages range from 6 to 95 years, displaying a bimodal distribution with peaks at 25 and 74 years. Females constitute 53% of the subjects. We divided the subjects into training, validation, and test sets in proportions of 85%, 5%, and 10%, respectively. We refer to the T1w and T2w research-grade MRIs from the internal test set as **T1w-INT-RG** and **T2w-INT-RG**, respectively. To evaluate the generalizability of the model, we use OASIS 3 as the external dataset, ensuring that no MRIs from this dataset were used during training or for model selection. The average age in the OASIS 3 dataset is  $68 \pm 9$  and 61% of the subjects are female. We divided OASIS 3 into three separate test sets: research-grade T1w images (**T1w-EXT-RG**), clinical-grade T2w images (**T2w-EXT-CG**), and clinical-grade FLAIR images (**FLR-EXT-CG**). The lowest resolution encountered in clinical-grade MRIs was 6mm along the acquisition direction. Finally, we selected 675 AD subjects from ADNI, considering for each subject only one research-grade T1w MRI and the associated ADAS-Cog-13 cognitive score. We refer to this dataset as **AD-EVAL**.

### B. Experiments settings

We explore two different configurations of SynthBA, depending on which type of prior is used in the GMM step within the generative model. The first configuration, named **SynthBA-u**, utilizes a fixed uniform prior (refer to the upper branch in Figure 1). The second configuration, referred to as **SynthBA-g**, incorporates two Gaussian priors (refer to the lower branch in Figure 1), whose parameters are estimated

from T1w and T2w training scans, respectively. The brain tissue segmentations, derived from the internal training set of T1w images, are utilized by the generative model to dynamically produce synthetic brain MRIs during training. We train SynthBA on these synthetic brain MRIs for 300 epochs. We use the AdamW optimizer [24] with a batch size of 16 and an initial learning rate set to  $10^{-4}$ , which is progressively reduced to  $10^{-5}$  using a cosine decay scheduler [25]. We use a dropout probability of 30%. To mitigate the non-uniform age distribution in the training set, we divide the age range into five-year bins and allocate each training sample to its corresponding bin. During training, we employ uniform sampling across these bins to create a batch, effectively balancing the inclusion of all age groups. We use the validation set to select the model with the best performance. For the 3D DenseNet, we use the implementation available in MONAI [26]. All experiments are conducted on a GeForce RTX 4090 (24GB VRAM).

### C. Comparative study

We conduct an evaluation of both SynthBA-g and SynthBA-u configurations, on the five different test sets detailed in Section III-A (two internal, three external), with varying MRI sequences and resolutions. We compare our models with state-of-the-art baselines: [7] and [8]. For the first baseline [7], we evaluate (i) their five publicly-available pre-trained brain age models, each one specialized on a distinct MRI sequence (GRE, DWI, T1w, T2w, FLAIR), and (ii) their T1w and T2w models fine-tuned respectively on our T1w and T2w training images. For the second baseline [8], there are no publicly-available models; therefore, we evaluate a T1w model and a T2w model trained with their contrastive method on our T1w and T2w training images, respectively. Table II presents the obtained results. Here we report the Mean Absolute Error (MAE) for each method across all internal and external test sets. Importantly, we report the average MAE over all test sets (AVG-ALL) and the average MAE over external test sets only (AVG-EXT). Both average MAEs serve as key metrics for this study, summarizing brain age model performance across various MRI sequences and resolutions.

Table II reveals that some baseline models achieve considerably low MAE on specific MRI sequences and resolutions compared to SynthBA. However, these models fail to generalize across heterogeneous acquisition settings, with their AVG-ALL and AVG-EXT errors ranging from 7 to 19 years and 7 to 20 years, respectively. In contrast, both SynthBA-u and SynthBA-g demonstrate robust performance across various acquisition settings without requiring re-training. In particular, SynthBA-g achieves the lowest average MAEs of  $4.7 \pm 0.61$  and  $4.9 \pm 0.40$  for AVG-ALL and AVG-EXT, respectively. Moreover, SynthBA-g yields the best MAE ( $5.0 \pm 4.75$ ) on FLR-EXT-CG and matches the best MAE ( $5.2 \pm 4.70$ ) on T2w-EXT-CG. We observe that SynthBA-g consistently achieves a lower MAE compared to SynthBA-u and demonstrates

TABLE II  
COMPARISON WITH STATE-OF-THE-ART BRAIN AGE MODELS ON INTERNAL AND EXTERNAL TEST SETS.

Method	Config	T1w-INT-RG	T2w-INT-RG	T1w-EXT-RG	T2w-EXT-CG	FLR-EXT-CG	AVG-ALL	AVG-EXT
[7]	Pre-trained on GRE	22.7 ± 8.97	11.9 ± 6.92	20.4 ± 8.07	20.5 ± 8.25	21.6 ± 8.46	19.4 ± 4.30	20.8 ± 0.63
	Pre-trained on T1w	13.9 ± 7.71	22.5 ± 9.26	11.5 ± 6.28	13.3 ± 8.19	28.4 ± 7.31	17.9 ± 7.24	17.7 ± 9.28
	Pre-trained on T2w	19.7 ± 11.27	4.6 ± 4.72	18.6 ± 7.07	12.2 ± 5.80	18.5 ± 8.60	14.7 ± 6.35	16.4 ± 3.68
	Pre-trained on FLAIR	11.2 ± 6.09	6.9 ± 6.68	8.1 ± 5.09	21.7 ± 9.09	15.9 ± 5.75	12.8 ± 6.07	15.2 ± 6.83
	Pre-trained on DWI	13.5 ± 8.26	10.5 ± 7.44	13.0 ± 6.76	27.0 ± 8.96	11.0 ± 5.58	15.0 ± 6.84	17.0 ± 8.75
	Fine-tuned on T1w	3.7 ± 3.23	7.4 ± 7.73	4.3 ± 3.78	30.8 ± 9.80	7.5 ± 5.05	10.7 ± 11.33	14.2 ± 14.45
	Fine-tuned on T2w	27.9 ± 14.81	<b>3.1 ± 2.54</b>	29.2 ± 6.87	<b>5.2 ± 4.29</b>	15.8 ± 9.46	16.2 ± 12.26	16.7 ± 12.07
[8]	Trained on T1w	<b>3.2 ± 2.75</b>	11.9 ± 10.89	<b>4.1 ± 3.60</b>	11.3 ± 9.75	5.9 ± 4.21	7.3 ± 4.06	7.1 ± 3.76
	Trained on T2w	16.5 ± 11.00	3.4 ± 3.13	13.0 ± 8.46	7.1 ± 5.80	15.3 ± 8.09	11.0 ± 5.61	11.8 ± 4.24
Ours	SynthBA-u	4.5 ± 4.41	4.7 ± 4.90	5.0 ± 3.90	5.9 ± 5.01	6.3 ± 5.38	5.3 ± 0.81	5.7 ± 0.70
	SynthBA-g	3.8 ± 3.78	5.1 ± 4.54	4.5 ± 4.09	<b>5.2 ± 4.70</b>	<b>5.0 ± 4.75</b>	<b>4.7 ± 0.61</b>	<b>4.9 ± 0.40</b>

superior generalization capability when applied to external datasets. This could be attributed to the fact that the synthetic MRIs generated with estimated Gaussian priors from two different MRI sequences (T1w, T2w) more closely resemble real MRIs.

#### D. Clinical evaluation

In this experiment, we conduct a clinical evaluation of SynthBA, examining the relationship between the ADAS-Cog13 score and brain PAD determined by the SynthBA model. The ADAS-Cog13 [27] score serves as a cognitive evaluation tool specifically designed for assessing cognitive function in AD patients. Using the best SynthBA configuration (SynthBA-g), we extract predicted brain age data from subjects within the AD-EVAL dataset and compute the brain PAD. Our analysis reveals a significant Pearson correlation of 0.275 ( $p < 0.001$ ) between brain PAD and the ADAS-Cog-13 score among AD subjects, as depicted in Figure 3.

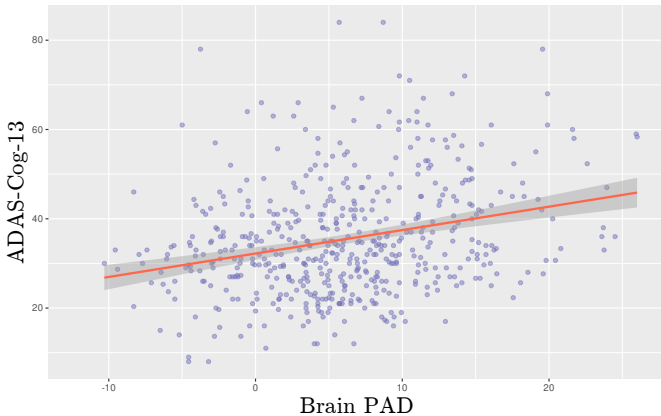


Fig. 3. Association between SynthBA’s brain PAD and ADAS-Cog13.

#### IV. CONCLUSIONS

In this study, we introduce SynthBA, a robust brain age model that maintains consistent performance across different

acquisition settings. We demonstrate its capabilities on both research-grade and clinical-grade MRIs from three different sequences (T1w, T2w, and FLAIR) and compare its performance to state-of-the-art brain age models, showcasing its superiority in heterogeneous acquisition settings. Future directions include enhancing the generative model with various MRI artefacts to increase predictive robustness and incorporating random orientation to avoid alignment to a common template, especially in low-resolution, anisotropic scans, where alignment can be detrimental. Our work encourages the widespread adoption of brain age analysis by extending its application to scenarios where retraining or fine-tuning is impractical. To facilitate this, we publicly release our SynthBA models as a standalone software package<sup>1</sup>.

#### ACKNOWLEDGMENTS

Lemuel Puglisi is enrolled in a PhD program at the University of Catania, fully funded by PNRR (DM 118/2023). Francesco Guarnera is funded by the PNRR MUR project PE0000013-FAIR. Alessia Rondinella is a PhD candidate enrolled in the National PhD in Artificial Intelligence, XXXVII cycle, course on Health and life sciences, organized by Università Campus Bio-Medico di Roma.

#### REFERENCES

- [1] J. H. Cole, S. J. Ritchie, M. E. Bastin, V. Hernández, S. Muñoz Maniega, N. Royle, J. Corley, A. Pattie, S. E. Harris, Q. Zhang *et al.*, “Brain age predicts mortality,” *Molecular psychiatry*, vol. 23, no. 5, pp. 1385–1392, 2018.
- [2] J. H. Cole, R. P. Poudel, D. Tsagkrasoulis, M. W. Caan, C. Steves, T. D. Spector, and G. Montana, “Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker,” *NeuroImage*, vol. 163, pp. 115–124, 2017.
- [3] J. H. Cole, J. Raffel, T. Friede, A. Eshaghi, W. J. Brownlee, D. Chard, N. De Stefano, C. Enzinger, L. Pirpamer, M. Filippi *et al.*, “Longitudinal assessment of multiple sclerosis with the brain-age paradigm,” *Annals of neurology*, vol. 88, no. 1, pp. 93–105, 2020.
- [4] W. Huang, X. Li, H. Li, W. Wang, K. Chen, K. Xu, J. Zhang, Y. Chen, D. Wei, N. Shu *et al.*, “Accelerated brain aging in amnesic mild cognitive impairment: relationships with individual cognitive decline, risk factors for alzheimer disease, and clinical progression,” *Radiology: Artificial Intelligence*, vol. 3, no. 5, p. e200171, 2021.

<sup>1</sup><https://github.com/LemuelPuglisi/SynthBA>

- [5] L. Puglisi, A. Eshaghi, G. Parker, F. Barkhof, D. C. Alexander, and D. Ravi, "Deepbrainprint: A novel contrastive framework for brain mri re-identification," in *Medical Imaging with Deep Learning*. PMLR, 2024, pp. 716–729.
- [6] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2017, pp. 23–30.
- [7] D. A. Wood, M. Townend, E. Guilhem, S. Kafiabadi, A. Hammam, Y. Wei, A. Al Busaidi, A. Mazumder, P. Sasieni, G. J. Barker *et al.*, "Optimising brain age estimation through transfer learning: A suite of pre-trained foundation models for improved performance and generalisability in a clinical setting," *Human Brain Mapping*, vol. 45, no. 4, p. e26625, 2024.
- [8] C. A. Barbano, B. Dufumier, E. Duchesnay, M. Grangetto, and P. Gori, "Contrastive learning for regression in multi-site brain age prediction," 2023.
- [9] J. M. Bayer, P. M. Thompson, C. R. Ching, M. Liu, A. Chen, A. C. Panzenhagen, N. Jahanshad, A. Marquand, L. Schmaal, and P. G. Sämann, "Site effects how-to and when: An overview of retrospective techniques to accommodate site effects in multi-site neuroimaging analyses," *Frontiers in neurology*, vol. 13, p. 923988, 2022.
- [10] B. Dufumier, A. Grigis, J. Victor, C. Ambroise, V. Frouin, and E. Duchesnay, "Openbbh: a large-scale multi-site brain mri data-set for age prediction and debiasing," *NeuroImage*, vol. 263, p. 119637, 2022.
- [11] P. A. Valdes-Hernandez, C. Laffitte Nodarse, J. A. Peraza, J. H. Cole, and Y. Cruz-Almeida, "Toward mr protocol-agnostic, unbiased brain age predicted from clinical-grade mris," *Scientific Reports*, vol. 13, no. 1, p. 19570, 2023.
- [12] B. Billot, D. N. Greve, O. Puonti, A. Thielscher, K. Van Leemput, B. Fischl, A. V. Dalca, J. E. Iglesias *et al.*, "Synthseg: Segmentation of brain mri scans of any contrast and resolution without retraining," *Medical image analysis*, vol. 86, p. 102789, 2023.
- [13] A. Hoopes, J. S. Mora, A. V. Dalca, B. Fischl, and M. Hoffmann, "Synthstrip: skull-stripping for any brain image," *NeuroImage*, vol. 260, p. 119474, 2022.
- [14] J. E. Iglesias, B. Billot, Y. Balbastre, C. Magdamo, S. E. Arnold, S. Das, B. L. Edlow, D. C. Alexander, P. Golland, and B. Fischl, "Synthsr: A public ai tool to turn heterogeneous clinical brain scans into high-resolution t1-weighted images for 3d morphometry," *Science advances*, vol. 9, no. 5, p. eadd3607, 2023.
- [15] M. Hoffmann, B. Billot, D. N. Greve, J. E. Iglesias, B. Fischl, and A. V. Dalca, "Synthmorph: learning contrast-invariant registration without acquired images," *IEEE transactions on medical imaging*, vol. 41, no. 3, pp. 543–558, 2021.
- [16] D. Ravi, F. Barkhof, D. C. Alexander, L. Puglisi, G. J. Parker, A. Eshaghi, A. D. N. Initiative *et al.*, "An efficient semi-supervised quality control system trained using physics-based mri-artefact generators and adversarial training," *Medical Image Analysis*, vol. 91, p. 103033, 2024.
- [17] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4700–4708.
- [18] 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 *et al.*, "Alzheimer's disease neuroimaging initiative (adni): clinical characterization," *Neurology*, vol. 74, no. 3, pp. 201–209, 2010.
- [19] K. A. Ellis, A. I. Bush, D. Darby, D. De Fazio, J. Foster, P. Hudson, N. T. Lautenschlager, N. Lenzo, R. N. Martins, P. Maruff *et al.*, "The australian imaging, biomarkers and lifestyle (aibl) study of aging: methodology and baseline characteristics of 1112 individuals recruited for a longitudinal study of alzheimer's disease," *International psychogeriatrics*, vol. 21, no. 4, pp. 672–687, 2009.
- [20] P. J. LaMontagne, T. L. Benzinger, J. C. Morris, S. Keefe, R. Hornbeck, C. Xiong, E. Grant, J. Hassensstab, K. Moulder, A. G. Vlassenko *et al.*, "Oasis-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and alzheimer disease," *MedRxiv*, pp. 2019–12, 2019.
- [21] "IXI - Information eXtraction from Images," <http://www.brain-development.org>.
- [22] D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub, K. Ugurbil, W.-M. H. Consortium *et al.*, "The wu-minn human connectome project: an overview," *Neuroimage*, vol. 80, pp. 62–79, 2013.
- [23] X.-N. Zuo, J. S. Anderson, P. Bellec, R. M. Birn, B. B. Biswal, J. Blautzik, J. Breitner, R. L. Buckner, V. D. Calhoun, F. X. Castellanos *et al.*, "An open science resource for establishing reliability and reproducibility in functional connectomics," *Scientific data*, vol. 1, no. 1, pp. 1–13, 2014.
- [24] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [25] —, "Sgdr: Stochastic gradient descent with warm restarts," in *International Conference on Learning Representations*, 2016.
- [26] M. J. Cardoso, W. Li, R. Brown, N. Ma, E. Kerfoot, Y. Wang, B. Murrey, A. Myronenko, C. Zhao, D. Yang *et al.*, "Monai: An open-source framework for deep learning in healthcare," *arXiv preprint arXiv:2211.02701*, 2022.
- [27] R. C. Mohs, D. Knopman, R. C. Petersen, S. H. Ferris, C. Ernesto, M. Grundman, M. Sano, L. Bieliauskas, D. Geldmacher, C. Clark *et al.*, "Development of cognitive instruments for use in clinical trials of antideementia drugs: additions to the alzheimer's disease assessment scale that broaden its scope," *Alzheimer Disease & Associated Disorders*, vol. 11, pp. 13–21, 1997.