Accurate Brain Age Prediction from MRI: Evaluating Kolmogorov-Arnold and Convolutional Networks

Alessandro Giupponi[*1] ALESSANDRO.GIUPPONI@PHD.UNIPD.IT
Davide De Crescenzo[*1] DAVIDE.DECRESCENZO@STUDENTI.UNIPD.IT
Marco Pinamonti[1] MARCO.PINAMONTI.1@PHD.UNIPD.IT
Manuela Moretto[1,2] MANUELA.MORETTO@UNIPD.IT
Alessandra Bertoldo[1,3] ALESSANDRA.BERTOLDO@UNIPD.IT
Mattia Veronese[1,2] MATTIA.VERONESE@UNIPD.IT
Marco Castellaro[1] MARCO.CASTELLARO@UNIPD.IT

[1] *Department of Information Engineering, University of Padova, Padova, Italy*
[2] *Neuroimaging Department, IoPPN, King's College London, London, UK*
[3] *Padova Neuroscience Center, University of Padova, Padova, Italy*

Editors: Under Review for MIDL 2025

Abstract

Brain age prediction using T1-weighted MRI has become a key biomarker for assessing neurological health, with application in studying neurodegeneration (Soumya Kumari and Sundarrajan, 2024; Mishra et al., 2023; Lea et al., 2021) and brain development (Tanveer et al., 2023). While convolutional neural networks (CNNs) remain a standard approach, recent advances suggest that Kolmogorov-Arnold Networks (KANs) may offer superior performance in image-based task (Bodner et al., 2025; Li et al., 2024). In this study, we present the first use of KANs for brain age prediction from 3D MRI scans, comparing their performance against traditional CNNs. Experimental results show that KAN-based models reduce estimation errors, highlighting their potential for improving brain age assessment.

Keywords: Brain aging, Convolutional Neural Networks, Kolmogorov-Arnold Networks, Neurological biomarker

1. Introduction

Brain age prediction serves as a valuable biomarker, offering insights into neurodegenerative disorders, cognitive decline, and the effects of lifestyle on aging (Natalia et al., 2024; Franke and Gaser, 2019; Dias et al., 2025). Deep learning models, particularly CNNs, have been widely applied to this task due to their capacity to extract meaningful features from MRI scans (Peng et al., 2021; Dartora et al., 2024; Dinsdale et al., 2021). However, recent advances in neural architectures, such as Kolmogorov-Arnold Networks (Liu et al., 2025), provide new opportunities to enhance prediction accuracy (Patel et al., 2024).

KANs utilize the Kolmogorov-Arnold representation theorem to approximate complex functions (Schmidt-Hieber, 2020) more efficiently than conventional neural networks (SS et al., 2024; Yeo et al., 2025). They have shown promising results in classification, segmentation, and image generation tasks. This study investigates the application of convolutional KANs and hybrid CNN-KAN models for brain age prediction, comparing their performance to traditional CNNs.

[*] Contributed equally

2. Materials and Methods

Datset. The study utilized T1-weighted MRI scans from three publicly available datasets: the Human Connectome Project (Bookheimer et al., 2019), the Nathan Kline Istitute - Rockland Sample (Nooner et al., 2012), and the Cambridge Centre for Aging and Neuroscience (Taylor et al., 2017). The combined cohort included 2,129 participants (878 males and 1,250 females), ranging in age from 18 to 100 years. To ensure consistent input dimension (193 x 229 x 193) and spatial alignment across datasets, all images were linearly coregistered to the MNI152 2009c standard space. No additional harmonization was performed. To enhance model robustness, data augmentation (DA) techniques, including rotation ($\pm$ 40°) and translation ($\pm$ 10 pixels), were applied (Connor and M., 2019).

The dataset was randomly split into training (64%), validation (16%), and test (20%) subsets, maintaining age and sex distribution. For cross-validation experiments, the training and validation sets were redefined in each fold. Mann-Whitney U tests confirmed no statistical differences in age or sex between training and test sets ($p = 0.901$) nor between training and validation sets across cross-validation folds (lowest $p = 0.840$).

Models Architecture. The following models were tested: a standard CNN inspired by Cole et al. (2017), serving as a baseline (**CNN**), a convolutional KAN with a fully connected linear KAN output layer (**KAN**), and a hybrid CNN with a final fully connected linear KAN layer (**CNN + KAN-Lin**). All models used 3x3x3 convolutional kernels, and stride of 1 and 2 for the first convolutional layer were tested, however KAN were only evaluated at stride 2 due to its very high memory occupancy demands.

Training and Evaluation. Training optimized the Mean Squared Error (MSE) loss between actual and predicted brain age using the Adam optimizer (learning rate: 0.0001) over 1000 epochs, validating every 50 epochs. Five-fold cross-validation was performed for stride-2 models. The best-performing models were selected based on lowest validation loss and evaluated on the test set using Mean Absolute Error (MAE) and Pearson Correlation Coefficient (PCC). In the cross-validation scenario, final performance metrics were obtained via median ensembling of predictions provided by the best models in each fold.

3. Results and Discussion

As shown in Table 1, the KAN model with stride 2 outperformed the CNN, reducing error by 15.16%. The CNN + KAN-Lin hybrid achieved an 11.72% improvement with data augmentation, offering the best balance between accuracy and computational load. Moreover, data augmentation improved model generalizability in all three scenarios, significantly improving the performance on unseen test data.

Due to computational constraints and hybrid's model efficiency, stride-1 evaluation excluded the KAN model. The CNN + KAN-Lin model still outperformed CNN by 5.77%. However, the performance gain was smaller than with stride-2, likely because high-resolution input allowed CNN layer to extract finer features, reducing the added value of the KAN layer.

Despite improvements, both models exhibited age-related bias (Figure 1), overestimating younger ages and underestimating older ones. The hybrid model produced a smoother Predicted Age Difference (PAD) curve and improved accuracy in middle-aged groups, though biases persisted at age extremes (≤ 30 and ≥ 70 years), indicating a need for better age-related feature representation or bias correction methods.

Stride	Method	Without DA		With DA	
		MAE	PCC	MAE	PCC
2	CNN	5.982	0.908	4.588	0.947
	KAN	5.240	0.930	4.561	0.946
	CNN + KAN-Lin	5.286	0.932	4.051	0.959
1	CNN	4.929	0.944	4.158	0.958
	CNN + KAN-Lin	4.994	0.943	3.918	0.962

Table 1: MAE and PCC obtained for the different models with and without the use of data augmentation.

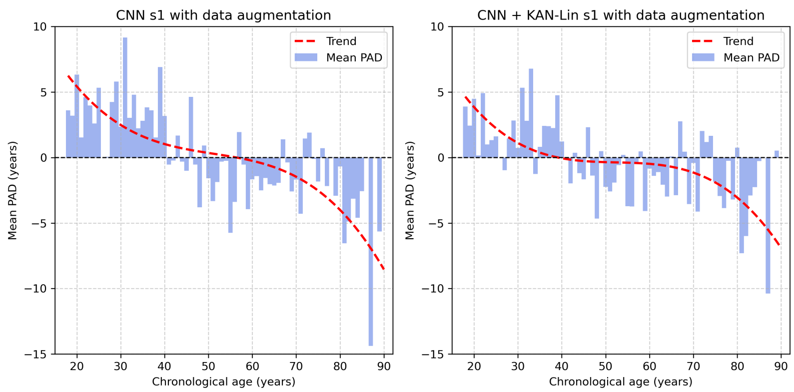

Figure 1: Mean PAD across chronological age bins (5-year intervals) for two models (left: CNN stride-1, right: CNN + KAN-Lin stride-1) trained with data augmentation. The red dashed line represents a third-order polynomial fit of PAD.

4. Conclusion

This study highlights the potential of Kolmogorov-Arnold Networks for brain age prediction using T1-weighted MRI scans. While KAN-based models achieved superior accuracy, the hybrid architecture combining CNN and KAN layers offered the best trade-off between performance and computational efficiency, and demonstrated robust generalizability with data augmentation. Although both models exhibited age-related bias—overestimating younger subjects and underestimating older ones—the CNN + KAN-Lin model produced smoother PAD distributions and higher accuracy in middle-aged groups. This hybrid approach capitalizes on the strengths of both architectures, leveraging CNNs for spatial feature extraction and KANs for complex functional approximations. Nevertheless, further work is needed to mitigate bias at the age extremes, potentially through targeted regularization strategies or debiasing techniques. Overall, this work supports the integration of KANs into neuroimaging pipelines for brain age estimation and opens the door to exploring their application in broader medical imaging tasks.

Acknowledgments

We acknowledge the use of data from three public databases. Data collection and sharing for this project were provided by the Cambridge Centre for Ageing and Neuroscience (CamCAN). CamCAN funding was supported by the UK Biotechnology and Biological Sciences Research Council (grant number BB/H008217/1), together with support from the UK Medical Research Council and the University of Cambridge, UK. Additionally, data from the HCP-Aging 2.0 Release were utilized in this research. Research reported in this publication was supported by the National Institute On Aging of the National Institutes of Health under Award Number U01AG052564 and by funds provided by the McDonnell Center for Systems Neuroscience at Washington University in St. Louis. The HCP-Aging 2.0 Release data used in this report came from DOI: 10.15154/1520707. We also acknowledge the use of data from the Enhanced Nathan Kline Institute – Rockland Sample (eNKI-RS).

References

Alexander Dylan Bodner, Antonio Santiago Tepsich, Jack Natan Spolski, and Santiago Pourteau. Convolutional kolmogorov-arnold networks, 2025. URL https://arxiv.org/abs/2406.13155.

Susan Y. Bookheimer, David H. Salat, Melissa Terpstra, Beau M. Ances, Deanna M. Barch, Randy L. Buckner, Gregory C. Burgess, Sandra W. Curtiss, Mirella Diaz-Santos, Jennifer Stine Elam, Bruce Fischl, Douglas N. Greve, Hannah A. Hagy, Michael P. Harms, Olivia M. Hatch, Trey Hedden, Cynthia Hodge, Kevin C. Japardi, Taylor P. Kuhn, Timothy K. Ly, Stephen M. Smith, Leah H. Somerville, Kâmil Uğurbil, Andre van der Kouwe, David Van Essen, Roger P. Woods, and Essa Yacoub. The lifespan human connectome project in aging: An overview. *NeuroImage*, 185:335–348, 2019. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2018.10.009.

James H. Cole, Rudra P.K. Poudel, Dimosthenis Tsagkrasoulis, Matthan W.A. Caan, Claire Steves, Tim D. Spector, and Giovanni Montana. Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker. *NeuroImage*, 163: 115–124, 2017. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2017.07.059. URL https://www.sciencedirect.com/science/article/pii/S1053811917306407.

Shorten Connor and Khoshgoftaar Taghi M. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6, 2019. doi: 10.1186/s40537-019-0197-0.

Caroline Dartora, Anna Marseglia, Gustav Mårtensson, Gull Rukh, Junhua Dang, J-Sebastian Muehlboeck, Lars-Olof Wahlund, Rodrigo Moreno, José Barroso, Daniel Ferreira, Helgi B. Schiöth, Eric Westman, for the Alzheimer's Disease Neuroimaging Initiative , the Australian Imaging Biomarkers , Lifestyle Flagship Study of Ageing, the Japanese Alzheimer's Disease Neuroimaging Initiative , and the AddNeuroMed Consortium . A deep learning model for brain age prediction using minimally preprocessed t1w images as input. *Frontiers in Aging Neuroscience*, 15, 2024. ISSN 1663-4365. doi: 10.3389/fnagi.2023.1303036. URL https://www.frontiersin.org/journals/aging-neuroscience/articles/10.3389/fnagi.2023.1303036.

Maria Fátima Dias, João Valente Duarte, Paulo de Carvalho, and Miguel Castelo-Branco. Unravelling pathological ageing with brain age gap estimation in alzheimer's disease, diabetes and schizophrenia. *Brain Communications*, 7(2):fcaf109, 03 2025. ISSN 2632-1297. doi: 10.1093/braincomms/fcaf109. URL https://doi.org/10.1093/braincomms/fcaf109.

Nicola K. Dinsdale, Emma Bluemke, Stephen M. Smith, Zobair Arya, Diego Vidaurre, Mark Jenkinson, and Ana I.L. Namburete. Learning patterns of the ageing brain in mri using deep convolutional networks. *NeuroImage*, 224:117401, 2021. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2020.117401. URL https://www.sciencedirect.com/science/article/pii/S1053811920308867.

Katja Franke and Christian Gaser. Ten years of brainage as a neuroimaging biomarker of brain aging: What insights have we gained? *Frontiers in Neurology*, 10, 2019. ISSN 1664-2295. doi: 10.3389/fneur.2019.00789. URL https://www.frontiersin.org/journals/neurology/articles/10.3389/fneur.2019.00789.

Baecker Lea, Garcia-Dias Rafael, Vieira Sandra, Scarpazza Cristina, and Mechelli Andrea. Machine learning for brain age prediction: Introduction to methods and clinical applications. *eBioMedicine*, 72, 2021. doi: doi:10.1016/j.ebiom.2021.103600.

Chenxin Li, Xinyu Liu, Wuyang Li, Cheng Wang, Hengyu Liu, Yifan Liu, Zhen Chen, and Yixuan Yuan. U-kan makes strong backbone for medical image segmentation and generation, 2024. URL https://arxiv.org/abs/2406.02918.

Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljačić, Thomas Y. Hou, and Max Tegmark. Kan: Kolmogorov-arnold networks, 2025. URL https://arxiv.org/abs/2404.19756.

Shiwangi Mishra, Iman Beheshti, and Pritee Khanna. A review of neuroimaging-driven brain age estimation for identification of brain disorders and health conditions. *IEEE Reviews in Biomedical Engineering*, 16:371–385, 2023. doi: 10.1109/RBME.2021.3107372.

De Bonis Maria Luigia Natalia, Fasano Giuseppe, Lombardi Angela, Ardito Carmelo, Ferrara Antonio, Di Sciascio Eugenio, and Di Noia Tommaso. Explainable brain age prediction: a comparative evaluation of morphometric and deep learning pipelines. *Brain Informatics*, 11, 2024. doi: 10.1186/s40708-024-00244-9.

Kate Nooner, Stanley Colcombe, Russell Tobe, Maarten Mennes, Melissa Breland, Alexis Moreno, Laura Panek, Shaquanna Brown, Stephen Zavitz, Qingyang Li, Sharad Sikka, David Gutman, Saroja Bangaru, Rochelle Schlachter, Stephanie Kamiel, Ayesha Anwar, Caitlin Hinz, Michelle Kaplan, Anna Rachlin, and Michael Milham. The nki-rockland sample: A model for accelerating the pace of discovery science in psychiatry. *Frontiers in Neuroscience*, 6:152, 10 2012. doi: 10.3389/fnins.2012.00152.

Salil Patel, Vicky Goh, James FitzGerald, and Chrystalina Antoniades. 2d and 3d deep learning models for mri-based parkinson's disease classification: A comparative analysis of convolutional kolmogorov-arnold networks, convolutional neural networks, and graph convolutional networks. 07 2024. doi: 10.48550/arXiv.2407.17380.

Han Peng, Weikang Gong, Christian F. Beckmann, Andrea Vedaldi, and Stephen M. Smith. Accurate brain age prediction with lightweight deep neural networks. *Medical Image Analysis*, 68:101871, 2021. ISSN 1361-8415. doi: https://doi.org/10.1016/j.media.2020.101871. URL https://www.sciencedirect.com/science/article/pii/S1361841520302358.

Johannes Schmidt-Hieber. The kolmogorov-arnold representation theorem revisited. *CoRR*, abs/2007.15884, 2020. URL https://arxiv.org/abs/2007.15884.

L.K. Soumya Kumari and R. Sundarrajan. A review on brain age prediction models. *Brain Research*, 1823:148668, 2024. ISSN 0006-8993. doi: https://doi.org/10.1016/j.brainres.2023.148668. URL https://www.sciencedirect.com/science/article/pii/S0006899323004390.

Sidharth SS, Keerthana AR, Gokul R, and Anas KP. Chebyshev polynomial-based kolmogorov-arnold networks: An efficient architecture for nonlinear function approximation, 2024. URL https://arxiv.org/abs/2405.07200.

M. Tanveer, M.A. Ganaie, Iman Beheshti, Tripti Goel, Nehal Ahmad, Kuan-Ting Lai, Kaizhu Huang, Yu-Dong Zhang, Javier Del Ser, and Chin-Teng Lin. Deep learning for brain age estimation: A systematic review. *Information Fusion*, 96:130–143, 2023. ISSN 1566-2535. doi: https://doi.org/10.1016/j.inffus.2023.03.007. URL https://www.sciencedirect.com/science/article/pii/S156625352300088X.

Jason R. Taylor, Nitin Williams, Rhodri Cusack, Tibor Auer, Meredith A. Shafto, Marie Dixon, Lorraine K. Tyler, Cam-CAN, and Richard 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: 262–269, 2017. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2015.09.018. Data Sharing Part II.

Sanghyun Yeo, Phuong Anh Nguyen, Anh Ngoc Le, and Satyam Mishra. Kan-pdes: A novel approach to solving partial differential equations using kolmogorov-arnold networks—enhanced accuracy and efficiency. In Akhtar Kalam, Saad Mekhilef, and Sheldon S. Williamson, editors, *Innovations in Electrical and Electronics Engineering*, pages 43–62, Singapore, 2025. Springer Nature Singapore. ISBN 978-981-97-9112-5.