Estimated Age-Guided Cerebral Microbleed Segmentation

Junmo Kwon[1], Jonghun Kim[1], Taehyeon Kim[2], and Hyunjin Park[1](✉)

[1] Department of Electrical and Computer Engineering, Sungkyunkwan University, Suwon 16419, South Korea
{skenfn1231,iproj2,hyunjinp}@skku.edu
[2] School of Applied and Creative Computing, Purdue University, West Lafayette, IN 47907, USA

Abstract. Cerebral microbleed (CMB) segmentation is complicated by extreme class imbalance and data scarcity due to low prevalence. Given the positive correlation between the prevalence of CMBs and aging, integrating age-related biomarkers offers a promising strategy to improve segmentation performance. In this study, we propose AgeCMB-Net, a framework that leverages estimated chronological age derived from anatomical T1-weighted magnetic resonance imaging. A pre-trained Meta-matching model was used to estimate each subject's chronological age. Based on the median age of the training cohort, we implement a two-branch decoder architecture where younger and older adults are processed through different decoders while sharing a common image encoder. Experiments on an in-house dataset and the public MICCAI VALDO 2021 Challenge dataset demonstrate that our approach outperforms conventional segmentation models and existing state-of-the-art methods. Furthermore, our model shows robust generalizability on the ATLAS v2.0 stroke dataset, which has a weaker correlation between estimated age and lesion volume. Our code is available at https://github.com/junmokwon/AgeCMBNet

Keywords: Cerebral microbleeds · Medical image segmentation · Age estimation · Meta-matching.

1 Introduction

Cerebral microbleeds (CMBs) are primary imaging indicators of cerebral small vessel disease [6]. CMBs are characterized by an oval shape and hypointensity on T2*-weighted magnetic resonance imaging (MRI), with a diameter limited to between 2 and 10 mm [6]. CMB detection is complicated by an age-dependent prevalence, as shown in Fig. 1 ($r = 0.36$ on the VALDO 2021 dataset [29]). This suggests that CMBs can be present in healthy older subjects without exhibiting distinctive symptoms [32]. Thus, true negatives (i.e., subjects without symptoms) cannot be filtered during the initial screening stage, causing an enormous annotation burden for radiologists, as this requires a complete full-volume inspection for every subject.

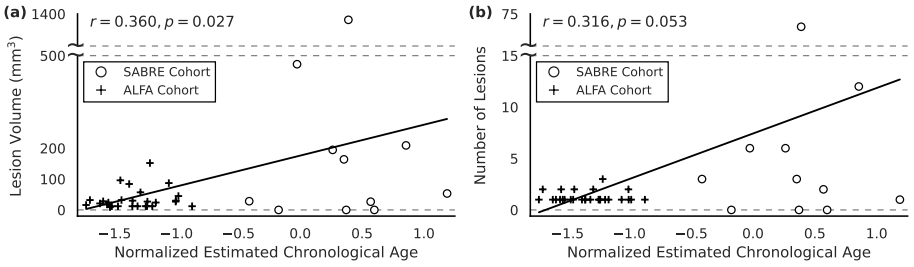

Fig. 1. Distribution of (a) lesion volume and (b) number of lesions across the estimated age on the VALDO 2021 dataset. Our study is motivated by the clinically relevant positive correlation ($r > 0.3$) between estimated age and lesion characteristics.

Existing studies have proposed automated CMB segmentation frameworks to alleviate the significant burden of manual annotation [18,19,20,21]. Based on the standard anatomical definition of the Microbleed Anatomical Rating Scale [7], AG-nnUNet [20] introduced anatomical priors into nnU-Net [13] to guide the prediction of CMB segmentation, surpassing the performance of conventional segmentation baselines. Another study [21] leveraged a pre-training strategy on UNETR++ [28], tackling the data scarcity problem. BP-nnUNet [18] estimated phenotypic information directly from anatomical T1 MRI scans, allowing for the utilization of subject-specific information. Although the concept of applying blood pressure (BP) measurements to segmentation networks was promising, the clinical relevance of BP was limited to deep microbleeds [32].

The goal of this study is to utilize clinical information that is generally applicable across all CMB subtypes to improve CMB segmentation performance. To achieve this, we focus on the strong positive correlation between CMB prevalence and estimated age [32]. A primary challenge was the absence of chronological age in all datasets. We addressed this limitation by incorporating estimated ages derived from a pre-trained Meta-matching model [9,4,34]. We propose AgeCMB-Net, an age-stratified nnU-Net segmentation framework that utilizes common encoder components and separate decoding paths for younger and older adults. This design minimizes additional parameter overhead while maximizing the sharing of cross-age imaging features. We collected 850 cases from the Samsung Medical Center and evaluated our model on the MICCAI VALDO 2021 Challenge dataset. To evaluate generalizability across different brain lesions, we extended our experiments to the ATLAS v2.0 dataset for stroke lesion segmentation. Our approach outperformed the most competitive baseline, highlighting its robustness even in scenarios with a weaker correlation between estimated age and lesion volume ($r = 0.125, p = 0.006$). Our study emphasizes the importance of integrating an aging prior into the segmentation network.

Contributions. The main contributions of this study are threefold. **(1)** We utilize estimated age as an important biomarker to assist in CMB segmentation. **(2)** We introduce an age-aware design for CMB segmentation, based on the

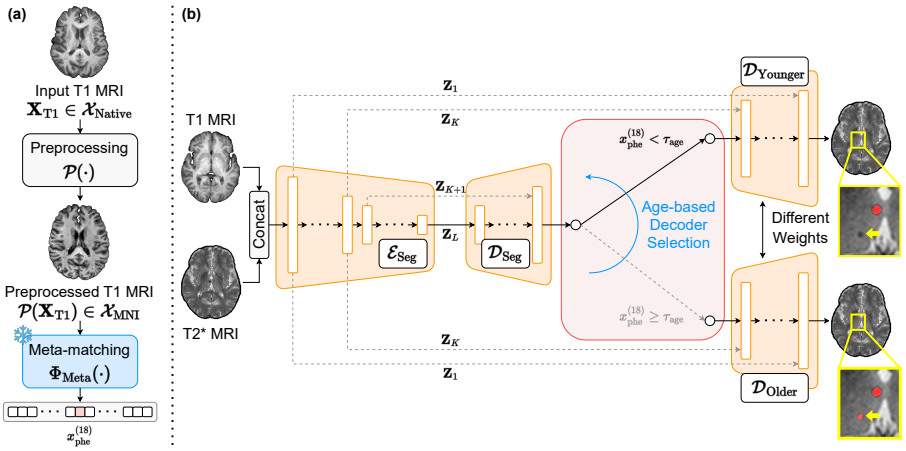

Fig. 2. Overview of our AgeCMBNet framework. (a) Chronological age estimation: A pre-trained Meta-matching model $\Phi_{\text{Meta}}(\cdot)$ extracts a 67-dimensional phenotype vector from preprocessed T1 MRI, where the 18th element $x_{\text{phe}}^{(18)}$ represents the estimated age. **(b) Age-guided segmentation:** Our network utilizes a shared encoder $\mathcal{E}_{\text{Seg}}$ and decoder $\mathcal{D}_{\text{Seg}}$. Based on the estimated age $x_{\text{phe}}^{(18)}$ relative to a data-driven median age threshold τ_{age}, our model dynamically selects a specific decoder for final segmentation. Compared to the younger decoder $\mathcal{D}_{\text{Younger}}$, the older decoder $\mathcal{D}_{\text{Older}}$ is more sensitive to cerebral microbleed (CMB) lesions, as indicated by yellow arrows.

positive correlation between CMB prevalence and chronological age. **(3)** We demonstrate that our framework can be applied to public datasets with missing age information.

2 Methods

Chronological age estimation. A fundamental problem with utilizing clinical information in a segmentation model is that clinical demographics are often missing or anonymized due to subject privacy constraints. To bridge this information gap and avoid potential privacy risks, we leverage the Meta-matching model [9,34], which was pre-trained on the large UK Biobank dataset [1] of 26,848 training subjects (Fig. 2). To utilize the Meta-matching framework, all T1 MRI scans underwent HCP PreFreeSurfer preprocessing including linear registration to MNI space as a prerequisite for chronological age estimation [5]. Let $\mathcal{P} : \mathcal{X}_{\text{Native}} \to \mathcal{X}_{\text{MNI}}$ be the preprocessing function and $\Phi_{\text{Meta}} : \mathcal{X}_{\text{MNI}} \to \mathbb{R}^{67}$ denote the phenotype mapping from registered T1 MRI scans. Each input T1 MRI $\mathbf{X}_{\text{T1}} \in \mathcal{X}_{\text{Native}}$ is mapped to a phenotype vector $\mathbf{x}_{\text{phe}}$:

$$\mathbf{x}_{\text{phe}} = \Phi_{\text{Meta}}(\mathcal{P}(\mathbf{X}_{\text{T1}})) = [x_{\text{phe}}^{(1)}, x_{\text{phe}}^{(2)}, \ldots, x_{\text{phe}}^{(67)}]^\top \in \mathbb{R}^{67}, \tag{1}$$

where $x_{\text{phe}}^{(18)}$ corresponds to the normalized estimated chronological age [9].

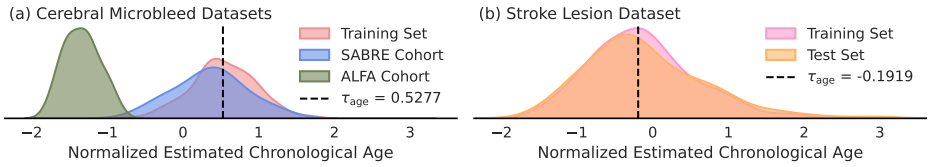

Fig. 3. Distribution of estimated age for (a) CMB datasets including in-house training and external validation cohorts (SABRE, ALFA), and (b) stroke lesion dataset with training and test splits. Dashed lines indicate age thresholds (τ_{age}) for decoder selection.

Definition of younger and older groups. We define the age threshold τ_{age} as the median age of the training cohort. This approach ensures a balanced sample distribution between the younger and older groups, mitigating potential model bias toward a specific group. Selected thresholds across different datasets are illustrated in Fig. 3.

Segmentation architecture. To capture age-dependent lesion prevalence [32], we propose a two-branch decoding strategy. As illustrated in Fig. 2, we adopt the nnU-Net framework with a residual encoder [13,12] as our backbone. Given co-registered 3D T1 and T2* MRI inputs, the shared encoder $\mathcal{E}_{\text{Seg}}$ generates L hierarchical latent feature maps $\{\mathbf{Z}_l\}_{l=1}^{L}$, where $\mathbf{Z}_L$ denotes the features from the deepest bottleneck layer. The corresponding decoder consists of $L-1$ blocks. To investigate the optimal decoder branching strategy, we introduce a split depth $K \in \{0, 1, \ldots, L-1\}$, where high-level features $\{\mathbf{Z}_l\}_{l=K+1}^{L-1}$ are processed by a shared decoder $\mathcal{D}_{\text{Seg}}$. The remaining features $\{\mathbf{Z}_l\}_{l=1}^{K}$ are then passed to either $\mathcal{D}_{\text{Younger}}$ or $\mathcal{D}_{\text{Older}}$, conditioned on the estimated age $x_{\text{phe}}^{(18)}$ relative to the threshold τ_{age}. We consider two boundary cases: (1) $K = 0$, where only the final segmentation heads are separated, and (2) $K = L-1$, where the entire decoder is bifurcated. Our architectural design allows the model to learn age-agnostic features in the upstream layers while predicting age-sensitive lesion characteristics in the downstream branches.

3 Experiments

Datasets. We utilized three datasets to evaluate our framework, with age distributions visualized in Fig. 3. Since age estimation using Meta-matching requires 3T MRI inputs, we excluded all 1.5T MRI scans from public cohorts.

(1) In-house dataset: We collected an in-house dataset of 850 cases for binary CMB segmentation. All cases were acquired with a Philips Achieva 3T scanner and underwent CMB diagnosis based on the recommended criteria [6]. All cases had at least one CMB lesion. We manually split this dataset into 556 training, 123 internal validation, and 171 testing cases.

(2) MICCAI VALDO 2021 Challenge dataset: For external validation, we used a subset of the MICCAI VALDO 2021 Challenge dataset for binary CMB segmentation [29]. From 72 available cases, we excluded 34 cases acquired

using a 1.5T GE MRI scanner. We utilized 38 cases from SABRE and ALFA cohorts [30,26]. These cohorts remained unseen during training.

(3) ATLAS v2.0 dataset: To evaluate generalizability under a weak age-lesion correlation ($r = 0.125, p = 0.006$), we used the ATLAS v2.0 stroke lesion dataset [25,24]. We excluded 52 cases obtained from 1.5T scanners and utilized 603 cases. We manually split this dataset into 323 training, 120 internal validation, and 120 testing cases.

MRI preprocessing. The pre-trained Meta-matching [9,4,34] requires T1 MRI preprocessing via the HCP PreFreeSurfer pipeline (as shown in Fig. 2) [5,14]. For T2* MRI preprocessing, we performed skull-stripping [11] and bias field correction [31]. For CMB segmentation, T1 MRI scans were rigidly co-registered to their corresponding native T2* space using ANTs [2]. Both T1 and T2* MRI were used as inputs. For stroke lesion segmentation, we used raw T1 MRI scans.

Implementation details. All models in Table 1 and Table 2 were trained using the nnU-Net framework with 3D configurations, except for 2D CMB-UNETR++ [21]. For U-Net variants, we followed the default nnU-Net hyperparameter recommendations and network designs. For transformer-based models, we adopted original hyperparameters from brain tumor segmentation studies. For CMB segmentation, input dimensions were set to $448 \times 384 \times 24$ for U-Net variants, and $448 \times 384 \times 32$ for transformer architectures, with a median spacing of $0.4286^2 \times 6.5$ mm^3. For stroke lesion segmentation, the input size was set to $224 \times 192 \times 160$ with a median spacing of 1.0 mm^3, except for four architectures: TransUNet [3], 3D UX-Net [23], UNETR++ [28], and TransBTS [33], which utilized a 128^3 input size to align with their original configurations for brain tumor segmentation.

Evaluation metrics. We used the Dice score (DSC) and Normalized Surface Dice (NSD) to evaluate segmentation performance. To evaluate detection performance, we utilized the F_1 score based on task-specific criteria. For CMB segmentation, a prediction is considered a true positive if the Euclidean distance between its centroid and the nearest ground-truth centroid is less than 5 mm [29]. We set the NSD tolerance to 5 mm. We did not report the Hausdorff distance (HD) because the HD is undefined for healthy subjects (see Fig. 1). For stroke lesion segmentation, a true positive is defined by the presence of at least one overlapping voxel between the predicted region and the ground truth label [25,24]. The tolerance for NSD was set to 2 mm, following [12]. Unlike CMB segmentation, we reported the 95th percentile Hausdorff distance (HD95), as participants are guaranteed to have at least one lesion.

4 Results

AgeCMBNet outperforms state-of-the-art CMB segmentation methods. Table 1 presents quantitative results comparing our proposed **AgeCMB-Net**, against eight general segmentation models (nnU-Net `ResEncL` [13,12], SwinUNETR [8], SwinUNETR-V2 [10], UNETR++ [28], 3D UX-Net [23], CoTr [35],

Table 1. Quantitative results of CMB segmentation. Values are reported as percentages. The best performance is highlighted in **bold** and the second-best is underlined.

	In-house Dataset			SABRE Cohort			ALFA Cohort		
Method	DSC↑	NSD↑	F_1↑	DSC↑	NSD↑	F_1↑	DSC↑	NSD↑	F_1↑
nnU-Net `ResEncL` [12]	39.13	57.99	61.71	54.57	61.42	59.96	45.04	56.01	53.70
SwinUNETR [8]	40.35	57.98	57.96	39.35	45.83	36.79	26.32	31.56	25.01
SwinUNETR-V2 [10]	41.58	58.75	59.53	51.04	56.14	55.35	43.09	49.44	43.83
UNETR++ [28]	40.05	61.07	61.60	42.47	54.63	55.15	36.51	43.42	48.15
3D UX-Net [23]	42.09	61.91	64.98	52.02	58.12	59.60	27.92	33.38	29.88
TransUNet [3]	46.81	65.17	65.41	42.80	49.36	49.73	33.38	40.42	37.02
CoTr [35]	47.76	65.17	65.67	53.04	60.68	59.12	42.26	51.21	47.60
UniSeg [36]	51.58	69.51	**71.12**	43.36	49.82	48.14	26.41	34.28	23.06
AG-nnUNet [20]	46.19	65.48	68.09	52.19	59.32	58.23	43.24	51.18	54.20
CMB-UNETR++ [21]	49.60	69.47	70.45	43.27	48.36	51.45	43.67	54.74	47.16
BP-nnUNet [18]	46.35	65.28	70.13	57.49	63.22	61.15	46.39	54.74	54.94
AgeCMBNet (Reversed)	46.92	66.16	68.89	58.81	64.03	61.59	40.48	47.22	51.11
AgeCMBNet (Ours)	**51.88**	**70.96**	70.83	**62.52**	**68.65**	**68.34**	**48.14**	**58.07**	**59.14**

Table 2. Quantitative results of stroke lesion segmentation using the ATLAS v2.0 dataset. The best performance is highlighted in **bold** and the second-best is underlined.

Metric	Trans-BTS	Swin-UNETR	Trans-UNet	Uni-Seg	SwinUN-ETR-V2	3D UX-Net	UNE-TR++	nnU-Net	Ours
DSC (%)↑	54.96	58.28	58.96	60.67	60.82	61.07	61.36	63.92	**64.11**
NSD (%)↑	59.38	63.58	64.28	65.89	66.98	67.02	66.42	70.48	**70.77**
HD95 (mm)↓	40.48	27.91	31.21	28.46	27.44	26.92	29.31	23.66	**21.08**
F_1 (%)↑	58.77	57.83	45.15	51.92	58.56	62.63	46.58	**68.19**	67.25

TransUNet [3], and UniSeg [36]) and three task-specific models (AG-nnUNet [20], CMB-UNETR++ [21], and BP-nnUNet [18]) across three different cohorts. AgeCMBNet consistently outperformed all baselines in every cohort, with the exception of the F_1 score on the in-house dataset. Notably, our AgeCMBNet achieved a noticeable performance margin over baselines across all metrics on the SABRE cohort. This implies that our AgeCMBNet demonstrates strong generalizability to unseen datasets with similar age distributions. We observed a similar trend when evaluating our AgeCMBNet on the ALFA cohort, which emphasizes the effectiveness of isolating the younger adult population using a younger decoder. Fig. 4 shows qualitative results, where AgeCMBNet consistently performed well on both in-house and unseen cohorts.

AgeCMBNet is generalizable to brain lesions with weak correlation. We conducted experiments using the ATLAS v2.0 dataset [25,24] to evaluate the generalizability of our approach. For this experiment, we excluded three task-specific models and added TransBTS [33] as a competing method. Table 2

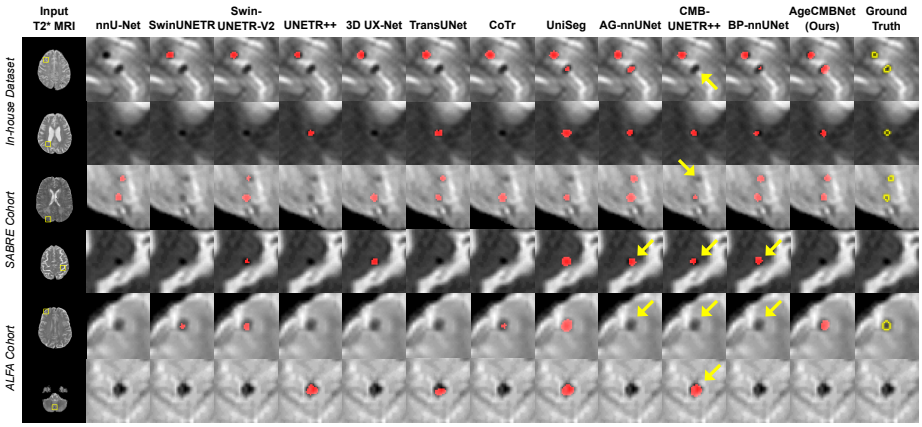

Fig. 4. Qualitative results of CMB segmentation. Predictions and ground-truth annotations are shown in red and yellow outlines, respectively. Yellow arrows indicate false positives or false negatives to highlight performance differences between task-specific models and our AgeCMBNet.

Table 3. Ablation results of CMB segmentation on the split depth for seven stages. The best performance is highlighted in **bold** and the second-best is underlined.

Split Depth	Params	SABRE Cohort			ALFA Cohort			Average (VALDO)		
0 1 2 3 4 5 6	(M)	DSC↑	NSD↑	F_1↑	DSC↑	NSD↑	F_1↑	DSC↑	NSD↑	F_1↑
● ○ ○ ○ ○ ○ ○	137.88	62.52	68.65	68.34	**48.14**	**58.07**	**59.14**	**52.31**	**61.13**	**61.80**
○ ● ○ ○ ○ ○ ○	137.91	57.37	64.81	66.22	31.00	42.19	39.75	38.64	48.74	47.41
○ ○ ● ○ ○ ○ ○	138.01	61.13	68.35	68.78	38.80	47.76	44.57	45.26	53.72	51.58
○ ○ ○ ● ○ ○ ○	138.44	**63.55**	**71.72**	**70.62**	37.90	42.72	40.12	45.33	51.12	48.95
○ ○ ○ ○ ● ○ ○	142.63	58.03	63.25	61.00	43.43	48.68	52.96	47.66	52.90	55.29
○ ○ ○ ○ ○ ● ○	148.98	56.75	63.92	66.12	38.55	45.76	51.73	43.82	51.01	55.90
○ ○ ○ ○ ○ ○ ●	154.93	58.57	63.36	60.67	41.79	48.71	45.56	46.65	52.96	49.93

shows that our method achieved the best DSC, NSD, and HD95, with a close second in the F_1 score. This result emphasizes that our approach is effective even in a weak age-correlation scenario, highlighting that our AgeCMBNet is generalizable to other brain lesions.

Ablation study on split depth. Following architectural recommendations of the nnU-Net framework [12], we evaluated all possible split depth choices for a model with $L = 7$ stages, as presented in Table 3. We also reported the total number of model parameters to analyze performance in terms of model complexity. Deeper split points increase the number of layers in younger and older decoders, thus increasing the model parameters [22]. In the VALDO 2021 dataset, the SABRE cohort shares a similar age distribution with our training set, whereas the ALFA cohort is distinct (see Fig. 3). We found that while a split depth of $K = 3$ achieved the highest performance on the SABRE cohort,

it failed to generalize to the ALFA cohort. This suggests that branching with $K = 3$ leads to overfitting toward the specific age distribution of the training data; thus, a smaller K is beneficial in this case. We chose $K = 0$ as it achieved the best average performance on the VALDO 2021 dataset.

Corrupting the decoder selection degrades the generalizability. We intentionally inverted the decoder choice to generate a reversed segmentation output. Two AgeCMBNet models in Table 1 indicate the reversed and intended cases. These results suggest that proper decoder selection is crucial and severely affects the performance of CMB segmentation on unseen datasets. The ALFA cohort presented the largest performance discrepancy among cohorts. This aligns with our expectations given its demographic bias toward younger adults, as shown in Fig. 3 (a). Interestingly, even with the performance degradation, the reversed AgeCMBNet surpassed BP-nnUNet [18] and nnU-Net [12]. This implies that when dealing with training samples obtained from a single site, adjusting the training data distribution with different decoders is more important than the architectural design.

5 Conclusion

We present an effective decoder separation strategy for CMB segmentation, utilizing separate younger and older segmentation decoders based on the strong positive correlation between chronological age and the prevalence of CMB lesions. Experiments on two CMB datasets demonstrate that our framework achieves promising segmentation performance by incorporating estimated age rather than relying on actual chronological age information. Moreover, because our approach uses image-based age estimation, AgeCMBNet can be easily incorporated into existing clinical workflows without requiring additional metadata, emphasizing its practical feasibility. We also highlight that our framework is effective for other datasets, such as ATLAS v2.0 for stroke lesion segmentation, despite a weak correlation between estimated age and lesion volume. Our extensive ablation studies on different branching points indicate that deep decoders for younger and older adults are not beneficial; shallow decoders are not only superior but also alleviate parameter overhead. Additionally, we observe that a comprehensive understanding of the training data distribution is important for improving generalizability in small lesion segmentation. Future research will explore label-free segmentation [16,17] to minimize annotation burden, while leveraging healthy controls [27] and MRI synthesis [15] to mitigate data scarcity and class imbalance.

Acknowledgments. We are grateful to Sang Won Seo for curating the dataset from the Samsung Medical Center. We greatly thank the organizers and participants of the MICCAI VALDO 2021 challenge for the public dataset. This study was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (RS-2025-00518533) and the SMC-SKKU Future Convergence Research Program Grant (SMO125123). This study was also supported by the AI Graduate School Support Program (Sungkyunkwan University) (RS-2019-II190421), the National

Research Foundation of Korea (RS-2024-00408040), the ICT Creative Consilience Program Grant (IITP-2026-RS-2020-II201821), and the AI Computing Infrastructure Enhancement (GPU Rental Support) User Support Program (RQT-25-090077), all funded by the Ministry of Science and ICT (MSIT), Republic of Korea.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Alfaro-Almagro, et al.: Image processing and Quality Control for the first 10,000 brain imaging datasets from UK Biobank. Neuroimage **166**, 400–424 (2018)
2. Avants, B.B., et al.: A reproducible evaluation of ANTs similarity metric performance in brain image registration. Neuroimage **54**(3), 2033–2044 (2011)
3. Chen, J., Mei, J., Li, X., Lu, Y., Yu, Q., Wei, Q., Luo, X., Xie, Y., et al.: TransUNet: Rethinking the U-Net architecture design for medical image segmentation through the lens of transformers. Medical Image Analysis **97**, 103280 (2024)
4. Chen, P., et al.: Multilayer meta-matching: Translating phenotypic prediction models from multiple datasets to small data. Imaging Neuroscience **2**, 1–22 (2024)
5. Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., et al.: The minimal preprocessing pipelines for the Human Connectome Project. Neuroimage **80**, 105–124 (2013)
6. Greenberg, S.M., Vernooij, M.W., Cordonnier, C., et al.: Cerebral microbleeds: a guide to detection and interpretation. The Lancet Neurology **8**(2), 165–174 (2009)
7. Gregoire, S., et al.: The microbleed anatomical rating scale (MARS) reliability of a tool to map brain microbleeds. Neurology **73**(21), 1759–1766 (2009)
8. Hatamizadeh, A., et al.: Swin UNETR: Swin Transformers for Semantic Segmentation of Brain Tumors in MRI Images. In: Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries. vol. LNCS 12962, pp. 272–284 (2022)
9. He, T., et al.: Meta-matching as a simple framework to translate phenotypic predictive models from big to small data. Nature neuroscience **25**(6), 795–804 (2022)
10. He, Y., Nath, V., Yang, D., Tang, Y., Myronenko, A., Xu, D.: SwinUNETR-V2: Stronger Swin Transformers with Stagewise Convolutions for 3D Medical Image Segmentation. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2023. vol. LNCS 14223, pp. 416–426 (2023)
11. Hoopes, A., Mora, J.S., Dalca, A.V., Fischl, B., Hoffmann, M.: SynthStrip: skull-stripping for any brain image. NeuroImage **260**, 119474 (2022)
12. Isensee, F., Wald, T., Ulrich, C., Baumgartner, M., Roy, S., et al.: nnU-Net Revisited: A Call for Rigorous Validation in 3D Medical Image Segmentation . In: proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. vol. LNCS 15009. Springer Nature Switzerland (October 2024)
13. Isensee, F., et al.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nature methods **18**(2), 203–211 (2021)
14. Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith, S.M.: Fsl. Neuroimage **62**(2), 782–790 (2012)
15. Kim, J., Park, H.: Visual prompt tuning for task-flexible medical image synthesis. Computer Methods and Programs in Biomedicine **277**, 109244 (2026)
16. Kim, J., et al.: Enhancing intracranial vessel segmentation using diffusion models without manual annotation for 3D Time-of-Flight Magnetic Resonance Angiography. Computerized Medical Im‘aging and Graphics **125**, 102651 (2025)

17. Kim, J., et al.: Weakly-supervised segmentation using sparse single point annotations for lumen and wall of carotid arteries in 3D MRI. Computer Methods and Programs in Biomedicine **269**, 108881 (2025)
18. Kwon, J., Kim, J., Kim, T., Seo, S.W., Cho, H.h., Park, H.: Blood Pressure Assisted Cerebral Microbleed Segmentation via Meta-matching . In: proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2025. vol. LNCS 15960. Springer Nature Switzerland (September 2025)
19. Kwon, J., Seo, S.W., Cho, H.h., Park, H.: Identifying Optimal nnU-Net Configuration for Cerebral Microbleed Segmentation. In: 2026 IEEE 23rd International Symposium on Biomedical Imaging (ISBI). pp. 1–4 (2026). https://doi.org/10.1109/ISBI61048.2026.11515994
20. Kwon, J., Seo, S.W., Park, H.: Anatomically-Guided Segmentation of Cerebral Microbleeds in T1-weighted and T2*-weighted MRI . In: proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. vol. LNCS 15002. Springer Nature Switzerland (October 2024)
21. Kwon, J., Seo, S.W., Park, H.: Enhancing Cerebral Microbleed Segmentation with Pretrained UNETR++. In: 2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM). pp. 3372–3377 (2024)
22. Kwon, J., et al.: Leveraging segmentation-guided spatial feature embedding for overall survival prediction in glioblastoma with multimodal magnetic resonance imaging. Computer Methods and Programs in Biomedicine **255**, 108338 (2024)
23. Lee, H.H., et al.: 3D UX-Net: A Large Kernel Volumetric ConvNet Modernizing Hierarchical Transformer for Medical Image Segmentation. In: The Eleventh International Conference on Learning Representations (ICLR) (2023)
24. Liew, S.L., Anglin, J.M., et al.: A large, open source dataset of stroke anatomical brain images and manual lesion segmentations. Scientific data **5**(1), 1–11 (2018)
25. Liew, S.L., et al.: A large, curated, open-source stroke neuroimaging dataset to improve lesion segmentation algorithms. Scientific data **9**(1), 320 (2022)
26. Molinuevo, J.L., et al.: The ALFA project: a research platform to identify early pathophysiological features of Alzheimer's disease. Alzheimer's & Dementia: Translational Research & Clinical Interventions **2**(2), 82–92 (2016)
27. Nikseresht, G., et al.: Automated detection of cerebral microbleeds on ex-vivo MRI scans of community-based older adults. NeuroImage **335**, 121986 (2026)
28. Shaker, A.M., Maaz, M., Rasheed, H., Khan, S., Yang, M.H., Khan, F.S.: UNETR++: Delving into Efficient and Accurate 3D Medical Image Segmentation. IEEE Transactions on Medical Imaging **43**(9), 3377–3390 (2024)
29. Sudre, C.H., et al.: Where is VALDO? VAscular lesions detection and segmentation challenge at MICCAI 2021. Medical Image Analysis **91**, 103029 (2024)
30. Tillin, T., et al.: The relationship between metabolic risk factors and incident cardiovascular disease in Europeans, South Asians, and African Caribbeans: SABRE (Southall and Brent Revisited)—a prospective population-based study. Journal of the American College of Cardiology **61**(17), 1777–1786 (2013)
31. Tustison, N.J., et al.: N4ITK: improved N3 bias correction. IEEE transactions on medical imaging **29**(6), 1310–1320 (2010)
32. Vernooij, M., van der Lugt, A., et al.: Prevalence and risk factors of cerebral microbleeds: the Rotterdam Scan Study. Neurology **70**(14), 1208–1214 (2008)
33. Wang, W., Chen, C., Ding, M., et al.: TransBTS: Multimodal Brain Tumor Segmentation Using Transformer. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2021. vol. LNCS 12901, pp. 109–119 (2021)
34. Wulan, N., et al.: Translating phenotypic prediction models from big to small anatomical MRI data using meta-matching. Imaging Neuroscience **2**, 1–21 (2024)

35. Xie, Y., Zhang, J., et al.: CoTr: Efficiently Bridging CNN and Transformer for 3D Medical Image Segmentation. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2021. vol. LNCS 12903, pp. 171–180 (2021)
36. Ye, Y., et al.: UniSeg: A Prompt-driven Universal Segmentation Model as well as A Strong Representation Learner. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2023. vol. LNCS 14222, pp. 508–518 (2023)