

A Unified Framework for Joint Detection of Lacunes and Enlarged Perivascular Spaces

Lucas He^{1,2*}, Krinos Li³, Hanyuan Zhang¹, Runlong He¹, Silvia Ingala⁴, Luigi Lorenzini⁵, Marleen de Bruijne⁶, Frederik Barkhof^{1,5,7}, Rhodri Davies^{2,8,†}, and Carole H. Sudre^{1,2,9,†}

¹ Hawkes Institute, University College London, UK

² Unit for Lifelong Health and Aging, University College London, UK

³ Bioengineering Department and Imperial-X, Imperial College London, UK

⁴ Department of Diagnostic Radiology, Copenhagen University Hospital, Denmark

⁵ Department of Radiology & Nuclear Medicine, Amsterdam UMC, Vrije Universiteit, The Netherlands

⁶ Department of Radiology and Nuclear Medicine, Erasmus MC, The Netherlands

⁷ Queen Square Institute of Neurology, University College London, UK

⁸ Barts Heart Centre, St Bartholomew's Hospital, London, UK

⁹ Department of Biomedical Computing - School of Biomedical Engineering and Imaging Sciences, King's College London, UK

Lucas.he.23@ucl.ac.uk

Abstract. Cerebral small vessel disease (CSVD) markers, specifically enlarged perivascular spaces (EPVS) and lacunae, present a unique challenge in medical image analysis due to their radiological mimicry. Standard segmentation networks struggle with feature interference and extreme class imbalance when handling these divergent targets simultaneously. To address these issues, we propose a morphology-decoupled framework where Zero-Initialized Gated Cross-Task Attention exploits dense EPVS context to guide sparse lacune detection. Furthermore, biological and topological consistency are enforced via a mixed-supervision strategy integrating Mutual Exclusion and Centerline Dice losses. Finally, we introduce an Anatomically-Informed Inference Calibration mechanism to dynamically suppress false positives based on tissue semantics. Extensive 5-fold cross-validation on the VALDO 2021 dataset ($N = 40$) demonstrates state-of-the-art performance, notably surpassing task winners in lacunae detection precision (71.1%, $p = 0.01$) and F1-score (62.6%, $p = 0.03$). Furthermore, evaluation on the external EPAD cohort ($N = 1762$) confirms the model's robustness for large-scale population studies. Code: <https://github.com/lucas-ucl/Lacune-EPVS-MTL>.

Keywords: Multi-Task Learning · Detection · MRI.

* Corresponding author.

† Joint senior authors.

1 Introduction

Cerebral small vessel disease (CSVD) is a primary cause of vascular dementia and a significant factor in stroke occurrence [2, 15]. Among the neuroimaging manifestations of CSVD, enlarged perivascular spaces (EPVS) and lacunes of presumed vascular origin serve as critical biomarkers for the assessment of disease burden and progression [22]. EPVS represent fluid-filled channels that reflect glymphatic dysfunction [8], whereas lacunes indicate focal tissue infarction and permanent brain injury [22]. Although these markers have distinct pathological origins, they frequently co-occur and exhibit significant radiological similarities. On all MRI sequences, both appear isointense to cerebrospinal fluid (CSF), manifesting as focal lesions with comparable signal intensities that complicate differentiation [4]. Distinguishing the linear, tubular topology of EPVS from the ovoid morphology of lacunes is essential for accurate diagnosis, yet remains a challenge for automation [21, 22].

Current automated solutions predominantly quantify CSVD markers as isolated tasks. For instance, Dubost et al. [4] pioneered weakly supervised networks specifically for EPVS, while top-performing methods in the VALDO challenge [21] often employ single-task segmentation architectures for lacunes. By ignoring the physiological correlation between these lesions, isolated models miss crucial contextual guidance. Extreme data imbalance, a recognized challenge for small CSVD markers [20], and the absence of biological priors further exacerbate these architectural flaws. While recent single-task methods have begun incorporating anatomical constraints [11], standard optimization on skewed distributions inherently biases joint models against sparse lacunes. Purely data-driven baselines routinely generate false positives in anatomically implausible regions [13].

To resolve these coupled challenges, we propose a morphology-decoupled multi-task framework. At its core, a Cross-Task Gated Attention module establishes a unidirectional information flow, leveraging the dense perivascular network as a spatial prior to guide sparse lacune detection without feature entanglement. Furthermore, we integrate explicit topological supervision and an adaptive distance-field calibration, effectively preserving structural integrity and suppressing clinically invalid predictions. Our contributions are as follows:

1. **Morphology-Decoupled Architecture:** We propose a shared-encoder, dual-decoder network featuring a Gated Cross-Task Attention mechanism. This design effectively disentangles tubular and spherical features, utilizing EPVS context to guide lacune detection while preventing feature collapse.
2. **Mixed-Supervision & Anatomy-Aware Constraints:** We introduce a hybrid optimization strategy that integrates fully supervised dense masks with weakly supervised regional counts for EPVS. Furthermore, it enforces biological consistency via Mutual Exclusion Loss and vascular topology through Centerline Dice supervision, mitigating the impact of label scarcity.
3. **Anatomical Calibration:** We develop a Distance-Field Calibration module that modulates prediction confidence based on tissue semantics. This mechanism strictly suppresses false positives in exclusion zones.

We then evaluated performance on the VALDO dataset [21] and the generalisability of our proposed solution on a large multicentre external dataset [17] in comparison with multiple state-of-the-art algorithms.

2 Methods




Fig. 1. Overview of the Proposed Morphology-Decoupled Framework.

2.1 Multi-Task Network Architecture

We present a multi-task 3D framework based on a Dynamic U-Net architecture (Fig. 1). To capture the distinct yet correlated morphological features of EPVS and lacunes, the network employs a shared encoder that branches into two task-specific decoders. This design enables the encoder to learn a common representation of the cerebral anatomy, while the decoders specialize in the fine-grained segmentation of tubular EPVS and ovoid lacunes, respectively [12].

We hypothesize that regions exhibiting a high density of EPVS indicate severe local micro-vascular dysfunction, thereby acting as a spatial map that guides the network to locate the rarer, co-occurring lacunes. To model this dependency without the introduction of circular feedback, we introduce a Cross-Task Gated Attention module [14]. This module enforces a unidirectional information flow from the EPVS stream to the Lacune stream at every up-sampling stage. By leveraging the local feature responses from the EPVS branch, the network explicitly guides the identification of sparse lacunar regions using the spatial context of perivascular spaces.

Let $F_{Lac}, F_{EPVS} \in \mathbb{R}^{C \times D \times H \times W}$ denote the Lacune (query) and EPVS (context) feature maps, respectively. We project these via 1×1 convolutions to obtain embeddings Q , K , and V . To minimize computational overhead, Q and K are mapped to a bottleneck dimension $C_{int} = C/r$ (with $r = 4$), while V retains the original channel dimension C . The spatial attention gate G and the refined Lacune features $\hat{F}_{Lac}$ are computed as follows:

$$G = \sigma \left(\frac{1}{\sqrt{C_{int}}} \sum_{c=1}^{C_{int}} (Q_c \odot K_c) \right), \quad \hat{F}_{Lac} = F_{Lac} + G \odot V \quad (1)$$

where σ is the sigmoid function and $\odot$ denotes element-wise multiplication. This mechanism aggregates global context across the bottleneck channels to highlight spatial regions where EPVS density structurally correlates with lacune formation.

This selective incorporation of EPVS context suppresses false positives in regions devoid of vascular pathology. Furthermore, we apply *zero-initialization* to the value projection V , ensuring the module initially acts as an identity mapping to facilitate stable task-specific convergence [1]. Finally, we employ deep supervision at three resolution levels to facilitate gradient flow [23].

2.2 Anatomically-Informed Inference Calibration

To mitigate false positives in anatomically implausible regions, we utilize FastSurfer [7] to partition the brain volume into three reliability tiers: Zone 1 (Allowed), comprising white matter, deep grey matter and the brainstem; Zone 2 (Transition), covering the hippocampus and cerebellar white matter; and Zone 3 (Exclusion), encompassing the cortex, ventricles and extra-cerebral tissues.

We compute a one-sided truncated distance map $D(x)$, where $D(x) = 0$ within Zone 1 and increases linearly with the Euclidean distance in exterior regions. To enforce anatomical consistency, we replace the standard static binarization threshold with a spatially adaptive decision boundary $T(x)$. For a given voxel x with the network’s predicted foreground probability $p(x)$, the final binary segmentation mask $M(x)$ is determined by the condition $M(x) = 1$ if $p(x) \geq T(x)$ and 0 otherwise, where $T(x)$ is formulated as:

$$T(x) = \begin{cases} 0.5 & \text{if } x \in \text{Zone 1} \\ 0.5 + \lambda \cdot \tanh(\gamma \cdot D(x)) & \text{otherwise} \end{cases} \quad (2)$$

Here, the base value of 0.5 represents the standard unbiased decision threshold for probabilistic predictions. The parameters $\lambda = 0.5$ and $\gamma = 0.5$ modulate the penalty magnitude and steepness, respectively. This formulation ensures that predictions in exclusion zones are heavily penalized, requiring significantly higher confidence (e.g., $p > 0.9$) to survive binarization, thereby effectively filtering cortical noise. Finally, we transition from voxel-level segmentation to object-level detection via Connected Component Analysis (CCA), extracting all candidate lesions with a minimum size of 1 voxel.

2.3 Objective Function

To address the extreme class imbalance, we employ the Tversky loss ($\mathcal{L}_{Tvk}$) with $\alpha = 0.1$ and $\beta = 0.9$ as the primary segmentation objective [18], penalizing false negatives more heavily. For the EPVS branch, we explicitly incorporate the Soft-Centerline Dice loss ($\mathcal{L}_{clDice}$) [19] to preserve the tubular topology of vascular structures. Additionally, a validity mask is applied to exclude unannotated regions from gradient computation. The task-specific losses, $\mathcal{L}_{Lac}$ and $\mathcal{L}_{EPVS}$, are computed by aggregating these objectives across all deep supervision scales.

To resolve radiological ambiguities between spatially proximate targets, we enforce a Mutual Exclusion loss ($\mathcal{L}_{excl}$). This term penalizes voxel-wise probabilistic overlap between the two classes:

$$\mathcal{L}_{excl} = \frac{1}{|\Omega|} \sum_{i \in \Omega} (p_{EPVS}^{(i)} \cdot p_{Lac}^{(i)}) \quad (3)$$

where Ω denotes the spatial domain and $p^{(i)}$ represents probability at voxel i .

Finally, we balance the multi-task learning using homoscedastic uncertainty weighting [10]. Let s_t denote the learnable log-variance for task $t \in \{EPVS, Lac\}$. The total optimization objective is formulated as:

$$\mathcal{L}_{Total} = \sum_t (e^{-s_t} \mathcal{L}_t + s_t) + \lambda_{excl} \mathcal{L}_{excl} \quad (4)$$

This formulation allows the network to dynamically down-weight tasks with high epistemic uncertainty during the early stages of training.

3 Experiments and Results

3.1 Datasets

We employed the VALDO 2021 challenge dataset ($N = 40$) for training and voxel-level evaluation [21]. The dataset comprises co-registered T1, T2, and FLAIR sequences, resampled to an isotropic resolution of 1 mm^3 . Annotations include fully segmented Lacune masks and a mixture of dense masks ($N = 12$) and weak regional counts ($N = 28$) for EPVS. To assess clinical generalizability, we further utilized the external EPAD cohort ($N = 1762$) [17]. Given the absence of voxel-wise ground truth, evaluation on EPAD relies on weak labels, patient level lacune presence and EPVS visual rating [16].

3.2 Implementation Details

The framework was implemented in PyTorch utilizing the MONAI library [3], and trained on an NVIDIA A100 (80GB) GPU. Images underwent z-score intensity normalization. To address volumetric imbalance, we employed a dual-stream sampling strategy [9], explicitly balancing Lacune- and EPVS-centered crops to training patches of size 96^3 . Data augmentation included random spatial flips

Table 1. Comparison with SOTA methods on the VALDO and EPAD datasets. **Bold** indicates best numerical results. * denotes statistical significance. Spearman’s rank correlation and Pearson’s correlation coefficient are denoted by ρ and r , respectively.

Task/Model	VALDO						EPAD		
	Segmentation		Detection						
<i>Task 1: EPVS</i>	DSC (%) $\uparrow$	NSD (%) $\uparrow$	Recall (%) $\uparrow$	Precision (%) $\uparrow$	F1 (%) $\uparrow$	FP/Subject $\downarrow$	ρ (BG) $\uparrow$	ρ (CSO) $\uparrow$	ρ (MB) $\uparrow$
DYN Unet	36.9 $\pm$ 9.7	56.9 $\pm$ 12.5	50.2 $\pm$ 1.8	54.9 $\pm$ 16.2	48.6 $\pm$ 9.8	21.0 $\pm$ 11.9	0.13 $\pm$ 0.02	0.17 $\pm$ 0.02	0.10 $\pm$ 0.02
MedNeXt	37.9 $\pm$ 10.2	55.6 $\pm$ 12.0	52.4 $\pm$ 7.5	46.1 $\pm$ 13.2	45.7 $\pm$ 8.7	30.6 $\pm$ 19.5	0.09 $\pm$ 0.02	0.09 $\pm$ 0.02	0.06 $\pm$ 0.02
Swin-UNETR V2	41.9 $\pm$ 9.2	62.2 $\pm$ 11.9	58.4 $\pm$ 5.6	55.0 $\pm$ 19.7	50.5 $\pm$ 10.7	30.4 $\pm$ 23.6	0.02 $\pm$ 0.02	0.05 $\pm$ 0.02	0.03 $\pm$ 0.02
VISTA-3D	36.3 $\pm$ 8.0	57.6 $\pm$ 9.8	51.0 $\pm$ 8.3	49.7 $\pm$ 10.7	47.7 $\pm$ 9.1	23.7 $\pm$ 7.6	0.14 $\pm$ 0.03	0.17 $\pm$ 0.02	0.10 $\pm$ 0.03
VALDO Winner	42.8 $\pm$ 9.9	63.7 $\pm$ 11.5	53.2 $\pm$ 4.4	58.3 $\pm$ 13.5	50.5 $\pm$ 9.2	17.9 $\pm$ 10.0	0.10 $\pm$ 0.03	0.13 $\pm$ 0.03	0.05 $\pm$ 0.02
Our Method	38.1 $\pm$ 6.5	56.7 $\pm$ 8.4	49.8 $\pm$ 10.0	67.4 $\pm$ 17.2	53.7 $\pm$ 9.6	15.6 $\pm$ 10.8	0.22 $\pm$ 0.02	0.29 $\pm$ 0.03	0.11 $\pm$ 0.02
<i>Task 3: Lacune</i>	DSC (%) $\uparrow$	NSD (%) $\uparrow$	Recall (%) $\uparrow$	Precision (%) $\uparrow$	F1 (%) $\uparrow$	FP/Subject $\downarrow$	BAcc (%) $\uparrow$	MAE $\downarrow$	r (Glob) $\uparrow$
DYN Unet	27.7 $\pm$ 7.6	39.1 $\pm$ 9.2	55.8 $\pm$ 14.3	34.6 $\pm$ 18.5	35.7 $\pm$ 11.7	4.3 $\pm$ 2.7	51.6 $\pm$ 0.9	2.6 $\pm$ 0.04	0.11 $\pm$ 0.03
MedNeXt	30.4 $\pm$ 4.9	43.0 $\pm$ 5.1	66.2 $\pm$ 13.3	30.3 $\pm$ 9.7	34.5 $\pm$ 5.5	6.1 $\pm$ 2.5	52.8 $\pm$ 1.0	0.5 $\pm$ 0.01	0.11 $\pm$ 0.07
Swin-UNETR V2	31.3 $\pm$ 7.1	44.2 $\pm$ 8.3	59.1 $\pm$ 9.4	30.5 $\pm$ 12.2	35.3 $\pm$ 9.5	5.1 $\pm$ 1.8	48.7 $\pm$ 2.1	2.0 $\pm$ 0.10	-0.03 $\pm$ 0.01
VISTA-3D	29.1 $\pm$ 6.4	40.2 $\pm$ 8.4	64.5 $\pm$ 20.4	34.5 $\pm$ 10.1	40.1 $\pm$ 9.6	4.0 $\pm$ 1.3	51.7 $\pm$ 2.1	0.8 $\pm$ 0.04	0.06 $\pm$ 0.04
VALDO Winner	38.3 $\pm$ 4.7	45.9 $\pm$ 7.5	59.5 $\pm$ 18.9	37.8 $\pm$ 13.5	42.2 $\pm$ 13.4	2.2 $\pm$ 0.9	52.1 $\pm$ 1.1	2.6 $\pm$ 0.04	0.06 $\pm$ 0.04
Our Method	42.4 $\pm$ 11.3	58.5 $\pm$ 15.9	58.4 $\pm$ 18.1	71.1 $\pm$ 17.3*	62.6 $\pm$ 17.1*	0.7 $\pm$ 0.9	64.9 $\pm$ 2.1	0.2 $\pm$ 0.01	0.24 $\pm$ 0.05

and gamma contrast adjustments. Optimization was performed using AdamW (learning rate 1×10^{-4} , weight decay 1×10^{-5}) for 200 epochs. Inference utilized a sliding window approach with an ROI size of 128^3 and overlap 0.6, evaluated via 5-fold cross-validation.

3.3 Evaluation Metrics

Similarly to the VALDO evaluation, for instance-level detection, a Lacune is classified as a True Positive if its centroid falls within 5 mm of the ground truth, whereas EPVS detection requires an Intersection over Union (IoU) $> 10\%$. We report Precision, Recall, F1-score, and False Positives per subject. Segmentation quality over true positive elements was assessed with the Dice Similarity Coefficient (DSC) and Normalized Surface Distance (NSD, tolerance 1 mm). For the EPAD cohort, we evaluate Lacune clinical utility using Balanced Accuracy (presence detection), Mean Absolute Error (MAE, count disparity), and Pearson’s Correlation (r) for global burden. For EPVS, we compute Spearman’s Rank Correlation (ρ) across the basal ganglia (BG), centrum semi-ovale (CSO), and mid-brain (MB) to account for the categorical Potter scale. Furthermore, confidence intervals (CIs) for both tasks were calculated using bootstrapping with 2000 iterations. Statistical significance of performance differences was determined via a paired Wilcoxon signed-rank test ($p < 0.05$).

3.4 Main Results

Table 1 benchmarks our framework against strong MONAI baselines (such as Swin-UNETR [5] and VISTA-3D [6]) and our re-implementations of the task-specific VALDO 2021 challenge winners [21]. For the baseline methods providing only segmentation, probabilistic outputs were thresholded and connected components extracted. For EPVS detection, we achieve an F1-score of $53.7 \pm 9.6\%$,

numerically exceeding the challenge winner ($50.5 \pm 9.2\%$) despite a lower recall ($49.8 \pm 10.0\%$ compared to $53.2 \pm 4.4\%$). In Lacune detection, our approach demonstrates a statistically significant improvement in both Precision ($71.1 \pm 17.3\%$, $p = 0.01$) and F1-score ($62.6 \pm 17.1\%$, $p = 0.03$) versus the challenge winner. External evaluation on the EPAD cohort further validates clinical robustness. For the Lacune task, our model achieves a state-of-the-art Balanced Accuracy of $64.9 \pm 2.1\%$, a global correlation of $r = 0.24 \pm 0.05$, and a minimal MAE of 0.2 ± 0.01 . Regarding EPVS burden estimation, our method achieves the highest regional concordance with visual ratings (e.g., CSO $\rho = 0.29 \pm 0.03$), substantially outperforming best baselines such as VISTA-3D (CSO $\rho = 0.17 \pm 0.02$).

Fig. 2 visually corroborates these metrics. While competing methods frequently generate false positives in cortical regions, our anatomically calibrated approach demonstrates superior specificity. It accurately localizes sparse lacunes and dense EPVS with minimal background noise, filtering out radiological mimics where baselines fail. However, this stringent calibration entails a trade-off with sensitivity, as it occasionally misses small or faint lesions.




Fig. 2. Qualitative visualization on the VALDO dataset. Ground truth lesions (cyan), true positive predictions (green), and false positive predictions (red) are shown.

3.5 Ablation Studies

Table 2 evaluates the proposed objective functions. The baseline model, trained solely with standard segmentation losses (L_{MT}), exhibits limited sensitivity for

Table 2. Ablation of objective functions on VALDO. **Bold** indicates numerical best

Task/model			Segmentation		Detection			
L_{MT}	L_{exc}	L_{clDice}	DSC (%) $\uparrow$	NSD (%) $\uparrow$	Recall (%) $\uparrow$	Precision (%) $\uparrow$	F1 (%) $\uparrow$	FP/Subject $\downarrow$
<i>Task 1: EPVS</i>								
$\checkmark$			35.6 $\pm$ 9.4	54.5 $\pm$ 12.0	45.6 $\pm$ 7.6	58.9 $\pm$ 17.3	47.7 $\pm$ 9.6	17.1 $\pm$ 14.2
$\checkmark$	$\checkmark$		35.7 $\pm$ 9.5	54.2 $\pm$ 11.9	46.9 $\pm$ 7.7	59.2 $\pm$ 17.5	48.9 $\pm$ 9.8	15.3 $\pm$ 13.6
$\checkmark$		$\checkmark$	36.9 $\pm$ 9.7	54.7 $\pm$ 12.1	46.6 $\pm$ 7.7	62.6 $\pm$ 19.1	49.2 $\pm$ 9.9	16.4 $\pm$ 12.3
$\checkmark$	$\checkmark$	$\checkmark$	34.3 $\pm$ 8.0	51.1 $\pm$ 12.4	47.1 $\pm$ 7.5	60.4 $\pm$ 10.7	50.0 $\pm$ 6.2	16.5 $\pm$ 11.6
<i>Task 3: Lacune</i>								
$\checkmark$			45.3 $\pm$ 15.5	46.4 $\pm$ 17.3	50.0 $\pm$ 12.8	46.3 $\pm$ 22.0	45.8 $\pm$ 14.9	3.9 $\pm$ 2.7
$\checkmark$	$\checkmark$		36.3 $\pm$ 13.9	51.4 $\pm$ 19.2	68.1 $\pm$ 17.5	50.1 $\pm$ 23.8	52.7 $\pm$ 17.2	2.2 $\pm$ 1.5
$\checkmark$		$\checkmark$	35.0 $\pm$ 13.4	45.2 $\pm$ 16.8	58.4 $\pm$ 15.0	49.1 $\pm$ 23.3	47.0 $\pm$ 15.3	3.3 $\pm$ 2.3
$\checkmark$	$\checkmark$	$\checkmark$	36.5 $\pm$ 14.0	51.4 $\pm$ 19.2	62.7 $\pm$ 16.1	57.8 $\pm$ 27.5	56.1 $\pm$ 18.3	1.0 $\pm$ 0.7

Lacunes (Recall: 50.0 $\pm$ 12.8%). Incorporating the Mutual Exclusion Loss (L_{exc}) explicitly addresses the biological inconsistency of overlapping labels. This constraint yields an improvement in Lacune Recall, rising to 68.1 $\pm$ 17.5%, while simultaneously reducing False Positives (FPs) from 3.9 to 2.2 per subject. Conversely, the Soft-Centerline Dice loss (L_{clDice}) specifically benefits the tubular EPVS structures (Task 1), boosting Precision to its peak of 62.6 $\pm$ 19.1%. The combined objectives provide the most robust trade-off, achieving the lowest FP rate (1.0 $\pm$ 0.7) for Lacunes while maintaining optimal F1-scores across tasks.

Table 3 validates the necessity of our unified framework design. Standard fully supervised Single-Task Learning (STL-Full) fails completely on the imbalanced EPVS class, resulting in a prohibitive FP rate of 94.8 $\pm$ 91.9 and an F1-score of only 18.3 $\pm$ 14.8%. The introduction of mixed supervision (STL-Mix) effectively mitigates this, recovering the F1-score to 48.6 $\pm$ 9.8%. Furthermore, transitioning from a shared-decoder (MTL-Shared) to an attention-decoupled architecture (MTL-Gated-Attn) resolves feature interference between the divergent targets. This gated-attention split yields consistent gains, improving the F1-score for EPVS (50.0% $\rightarrow$ 51.8%) and Lacunae (56.1% $\rightarrow$ 58.2%) by allowing independent optimization of tubular and round/ovoid representations.

The integration of Anatomically-Informed Inference Calibration refines the output (Table 3, bottom row) by filtering out predictions in implausible regions. For Lacunes, it results in a reduction of the FP rate to 0.7 $\pm$ 0.9 with minimal changes in recall, confirming the validity of the anatomical guidance.

4 Conclusion

In this paper, we presented a morphology-informed framework for the joint detection of lacunes and EPVS, addressing the dual challenges of extreme data imbalance and radiological mimicry. By applying mixed supervision with anatomically-informed inference calibration, our approach effectively enforces biological consistency while preserving the distinct topological features of tubular EPVS and

Table 3. Comparison of supervision strategies, cross-task attention decoupling, and anatomical calibration on the VALDO dataset. **Bold** indicates numerical best results.

Task/model			Segmentation		Detection			
Method	Decoder	Anatomical	DSC (%) $\uparrow$	NSD (%) $\uparrow$	Recall (%) $\uparrow$	Precision (%) $\uparrow$	F1 (%) $\uparrow$	FP/Subject $\downarrow$
<i>Task 1: EPVS</i>								
STL-Full	–	–	12.4 $\pm$ 17.3	20.4 $\pm$ 27.5	25.9 $\pm$ 17.9	21.4 $\pm$ 18.0	18.3 $\pm$ 14.8	94.8 $\pm$ 91.9
STL-Mix	–	–	36.9 $\pm$ 9.7	56.9 $\pm$ 12.5	50.2 $\pm$ 1.8	54.9 $\pm$ 16.2	48.6 $\pm$ 9.8	21.0 $\pm$ 11.9
MTL	Shared		34.3 $\pm$ 8.0	51.1 $\pm$ 12.4	47.1 $\pm$ 8.5	60.4 $\pm$ 10.7	50.0 $\pm$ 6.2	16.5 $\pm$ 11.6
MTL	Shared	✓	35.3 $\pm$ 7.1	52.8 $\pm$ 10.8	47.1 $\pm$ 8.5	60.8 $\pm$ 12.7	49.8 $\pm$ 6.7	16.1 $\pm$ 11.1
MTL	Gated-Attn		37.5 $\pm$ 7.1	56.3 $\pm$ 9.0	50.0 $\pm$ 10.3	63.5 $\pm$ 15.3	51.8 $\pm$ 9.7	15.9 $\pm$ 13.1
MTL	Gated-Attn	✓	38.1 $\pm$ 6.5	56.7 $\pm$ 8.4	49.8 $\pm$ 10.0	67.4 $\pm$ 15.2	53.7 $\pm$ 9.6	15.6 $\pm$ 10.8
<i>Task 3: Lacunae</i>								
STL	–	–	27.7 $\pm$ 7.6	39.1 $\pm$ 9.2	55.8 $\pm$ 14.3	34.6 $\pm$ 18.5	35.7 $\pm$ 11.7	4.3 $\pm$ 2.7
MTL	Shared		36.5 $\pm$ 14.0	51.4 $\pm$ 19.2	62.7 $\pm$ 16.1	57.8 $\pm$ 27.5	56.1 $\pm$ 18.3	1.0 $\pm$ 0.7
MTL	Shared	✓	35.6 $\pm$ 14.5	49.8 $\pm$ 20.5	62.7 $\pm$ 16.1	57.4 $\pm$ 27.4	55.9 $\pm$ 18.3	1.0 $\pm$ 0.6
MTL	Gated-Attn		43.3 $\pm$ 11.6	57.3 $\pm$ 16.0	59.0 $\pm$ 18.3	65.6 $\pm$ 19.2	58.2 $\pm$ 17.9	1.2 $\pm$ 0.9
MTL	Gated-Attn	✓	42.4 $\pm$ 11.3	58.5 $\pm$ 15.9	58.4 $\pm$ 18.1	71.1 $\pm$ 17.3	62.6 $\pm$ 17.1	0.7 $\pm$ 0.9

round/ovoid lacunae. Quantitative benchmarking on the VALDO dataset demonstrates state-of-the-art overall performance, surpassing challenge winners in precision and F1-score. Moreover, successful deployment on a large multi-center EPAD cohort validates the model’s robustness for population-level studies, demonstrating its capacity to extract reliable regional and global metrics across diverse demographics. This work provides an automated, highly reliable tool for quantifying vascular burden, paving the way for broad epidemiological research on cerebrovascular progression.

Acknowledgments. We acknowledge support from the UCL EPSRC CDT in i4health (EP/S021930/1) and use of the ISambard platform.

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

References

1. Bachlechner, T., Majumder, B.P., Mao, H.H., Cottrell, G.W., McAuley, J.: Rezero is all you need: Fast convergence at large depth (2020), <https://arxiv.org/abs/2003.04887>
2. Cannistraro, R.J., Badi, M., Eidelman, B.H., Dickson, D.W., Middlebrooks, E.H., Meschia, J.F.: Cns small vessel disease: A clinical review. *Neurology* **92**, 1146–1156 (06 2019). <https://doi.org/10.1212/WNL.0000000000007654>, <https://n.neurology.org/content/92/24/1146.abstract>
3. Cardoso, M.J., Li, W., Brown, R., Ma, N., Kerfoot, E., Wang, Y., Myronenko, A., Tang, Y., Yang, I., Zephyr, M., Hashemian, B., Alle, S., Darestani, M.Z., Budd, C., Modat, M., Vercauteren, T., Wang, G., Li, Y., Hu, Y., Bericat, D., Floca, R., Zhou, S.K., Shuaib, H., Farahani, K., Maier-Hein, K.H., Aylward, S., Dogra, P., Ourselin, S., Feng, A.: Monai: An open-source framework for deep learning in healthcare. *arXiv:2211.02701 [cs]* (11 2022), <https://arxiv.org/abs/2211.02701>

4. Dubost, F., Adams, H., Bortsova, G., Ikram, M.A., Niessen, W., Vernooij, M., de Bruijne, M.: 3d regression neural network for the quantification of enlarged perivascular spaces in brain mri. *Medical Image Analysis* **51**, 89–100 (01 2019). <https://doi.org/10.1016/j.media.2018.10.008>
5. Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H., Xu, D.: Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. *arXiv:2201.01266 [cs, eess]* (01 2022), <https://arxiv.org/abs/2201.01266>
6. He, Y., Guo, P., Tang, Y., Myronenko, A., Nath, V., Xu, Z., Yang, D., Belue, M., Harmon, S., Turkbey, B., Xu, D., Li, W.: Vista3d: A unified segmentation foundation model for 3d medical imaging (2024), <https://arxiv.org/abs/2406.05285>
7. Henschel, L., Conjeti, S., Estrada, S., Diers, K., Fischl, B., Reuter, M.: Fastsurfer - a fast and accurate deep learning based neuroimaging pipeline. *NeuroImage* **219**, 117012 (10 2020). <https://doi.org/10.1016/j.neuroimage.2020.117012>
8. Iliff, J.J., Wang, M., Liao, Y., Plogg, B.A., Peng, W., Gundersen, G.A., Benveniste, H., Vates, G.E., Deane, R., Goldman, S.A., Nagelhus, E.A., Nedergaard, M.: A paravascular pathway facilitates csf flow through the brain parenchyma and the clearance of interstitial solutes, including amyloid . *Science Translational Medicine* **4**, 147ra111–147ra111 (08 2012). <https://doi.org/10.1126/scitranslmed.3003748>
9. Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**, 203–211 (12 2020). <https://doi.org/10.1038/s41592-020-01008-z>, <https://www.nature.com/articles/s41592-020-01008-z>
10. Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. *arXiv:1705.07115 [cs]* (04 2018), <https://arxiv.org/abs/1705.07115>
11. Li, B., Bresser, d., Niessen, W., Osch, v., , v., Jan, B.G., Vernooij, M.W., Bron, E.: Prior-knowledge-informed deep learning for lacune detection and quantification using multi-site brain mri (2023), <https://arxiv.org/abs/2306.10622>
12. Moeskops, P., Wolterink, J., Velden, B.H.M.v.d., Gilhuijs, K., Leiner, T., Viergever, M., Išgum, I.: Deep learning for multi-task medical image segmentation in multiple modalities. *MICCAI* (2016). https://doi.org/10.1007/978-3-319-46723-8_55
13. Oktay, O., Ferrante, E., Kamnitsas, K., Heinrich, M., Bai, W., Caballero, J., Cook, S.A., de Marvao, A., Dawes, T., O'Regan, D.P., Kainz, B., Glocker, B., Rueckert, D.: Anatomically constrained neural networks (acnns): Application to cardiac image enhancement and segmentation. *IEEE Transactions on Medical Imaging* **37**, 384–395 (02 2018). <https://doi.org/10.1109/tmi.2017.2743464>
14. Oktay, O., Schlemper, J., Folgoc, L.L., Lee, M., Heinrich, M., McDonagh, S., Hammerla, N.Y., Kainz, B., Glocker, B., Rueckert, D.: Attention u-net: Learning where to look for the pancreas (2018), <https://arxiv.org/abs/1804.03999>
15. Pantoni, L.: Cerebral small vessel disease: from pathogenesis and clinical characteristics to therapeutic challenges. *The Lancet Neurology* **9**, 689–701 (07 2010). [https://doi.org/10.1016/s1474-4422\(10\)70104-6](https://doi.org/10.1016/s1474-4422(10)70104-6)
16. Potter, G.M., Chappell, F.M., Morris, Z., Wardlaw, J.M.: Cerebral perivascular spaces visible on magnetic resonance imaging: Development of a qualitative rating scale and its observer reliability. *Cerebrovascular Diseases* **39**, 224–231 (2015). <https://doi.org/10.1159/000375153>
17. Ritchie, C., Muniz-Terrera, G., Kivipelto, M., Solomon, A., Tom, B., Molinuevo, J.: The european prevention of alzheimer’s dementia (epad) longitudinal cohort study: Baseline data release v500.0. *The Journal of Prevention of Alzheimer’s Disease* pp. 1–7 (2019). <https://doi.org/10.14283/jpad.2019.46>

18. Salehi, S.S.M., Erdogmus, D., Gholipour, A.: Tversky loss function for image segmentation using 3d fully convolutional deep networks. arXiv:1706.05721 [cs] (06 2017), <https://arxiv.org/abs/1706.05721>
19. Shit, S., Paetzold, J.C., Sekuboyina, A., Ezhov, I., Unger, A., Zhylka, A., Pluim, W., Bauer, U., Menze, B.H.: cldice - a novel topology-preserving loss function for tubular structure segmentation. arXiv (Cornell University) (06 2021). <https://doi.org/10.1109/cvpr46437.2021.01629>
20. Sudre, C.H., Anson, B.G., Ingala, S., Lane, C.D., Jimenez, D., Haider, L., Smith, L., Ourselin, S., Jäger, R.H., Cardoso, M.J.: 3d multirater rcnn for multimodal multiclass detection and characterisation of extremely small objects. PMLR pp. 447–456 (05 2019), <https://proceedings.mlr.press/v102/sudre19a.html>
21. Sudre, C.H., Wijnen, K.V., Dubost, F., Adams, H., Atkinson, D., Barkhof, F., Birhanu, M.A., Bron, E.E., Camarasa, R., Anson, B.G., Ingala, S., Menze, B., Shi, P., Smith, L., Tillin, T., Tochon, G., Urien, H., Wiestler, B., Wolters, F.J., Yilmaz, P., de Groot, M., Vernooij, M.W., de Bruijne, M.: Where is valdo? vascular lesions detection and segmentation challenge at miccai 2021. Medical image analysis **91**, 103029–103029 (01 2024). <https://doi.org/10.1016/j.media.2023.103029>
22. Wardlaw, J.M., Smith, E.E., Biessels, G.J., Cordonnier, C., Fazekas, F., Lindley, R.I., Barkhof, F., Fox, N.C., Greenberg, S., Hachinski, V., Oostenbrugge, R.v., Pantoni, L., Speck, O., Viswanathan, A., Werring, D., Chen, C., Smith, C., van Buchem, M., Norrving, B., Gorelick, P.B., Dichgans, M.: Neuroimaging standards for research into small vessel disease and its contribution to ageing and neurodegeneration. The Lancet Neurology **12**, 822–838 (08 2013). [https://doi.org/10.1016/s1474-4422\(13\)70124-8](https://doi.org/10.1016/s1474-4422(13)70124-8)
23. Zhou, Z., Rahman, M., Tajbakhsh, N., Liang, J.: Unet++: A nested u-net architecture for medical image segmentation (2018), <https://arxiv.org/abs/1807.10165>