Learning Dual Prior Presentation for Opportunistic Screening of Visceral Artery Aneurysms on Non-Contrast CT

Jianfeng Zhang[1] *, Tianyi Liu[1,2] *, Long Bai[1,2] **, Zhuowei Xu[1], Yujian Hu[3], Yide Zhou[3], Yangyan He[3], Zirui Wang[1], Min Zhou[3], Hongkun Zhang[3], and Minfeng Xu[1]

[1] DAMO Academy, Alibaba Group, Hangzhou, China
[2] Hupan Lab, Hangzhou, China
[3] The First Affiliated Hospital of ZJU School of Medicine, Hangzhou, China
bailong.bai@alibaba-inc.com

Abstract. Visceral artery aneurysms (VAAs) are relatively rare but potentially life-threatening vascular dilatations that are frequently detected incidentally on abdominal CT, making opportunistic screening on non-contrast CT (NCT) clinically important. However, automated VAA detection on NCT is fundamentally challenging due to weak vessel–tissue contrast, complex vascular branching, and the continuous transition between normal vessel wall and aneurysmal dilation. We reformulate VAA screening as a dual-prior representation learning problem rather than a pure segmentation task. Built upon a multi-label vascular segmentation backbone, the proposed DeepVAA jointly learns two complementary priors defined in continuous target spaces: a geometric confidence prior that models aneurysmal dilation as a smooth structural response along the vessel wall, and an anatomical reliability prior that encodes geometry-induced ambiguity, such as bifurcations and high-curvature segments, as a continuous penalty field. During inference, the learned reliability prior explicitly calibrates the confidence map through structured prior projection and multiplicative suppression, yielding a prior-calibrated lesion probability. Extensive experiments on multi-center cohorts demonstrate consistent performance improvements and strong cross-site generalization, highlighting the robustness and effectiveness of the proposed dual-prior framework for opportunistic VAA screening on NCT.

Keywords: Visceral artery aneurysm · Non-contrast CT · Multi-task learning · Opportunistic screening

1 Introduction

Visceral artery aneurysms (VAAs) are dilatations of the celiac, superior mesenteric, inferior mesenteric, renal arteries or their branches [13, 20]. Although often asymptomatic, they carry a risk of rupture that may result in fatal intra-abdominal hemorrhage [3]. In contrast to abdominal aortic aneurysm (AAA),

* Co-first authors. ** Corresponding author.

which benefits from established screening recommendations and population-based strategies [2], there is no standardized screening protocol for VAAs. With the widespread use of abdominal CT in routine practice, many aneurysms are detected incidentally rather than through dedicated surveillance [20]. Therefore, opportunistic screening from routine CT examinations is clinically meaningful. Non-contrast CT (NCT) is fast, widely available, and frequently performed in general diagnostic workflows [9, 29]. An automated VAA detection framework on NCT could enable early identification without additional imaging burden.

Most automated aneurysm analysis pipelines are designed for AAA under contrast-enhanced CTA [4, 16, 23, 30]. Early systems relied on rule-based or semi-automatic lumen and thrombus segmentation [15]. More recent approaches adopt deep learning to segment the abdominal aorta and perform screening based on diameter measurement [25], while advanced architectures further refine boundary delineation and volumetric assessment on CTA [21, 26, 27]. Radiomic analysis of perivascular tissue has also been explored for progression evaluation [17]. Overall, existing automated frameworks are largely tailored to AAA and leverage explicit vascular enhancement for reliable lumen visualization. In contrast, studies on visceral artery aneurysms (VAAs) are comparatively limited and often rely on 4D-CTA acquisition or hemodynamic simulation [5, 24], which are not feasible for routine opportunistic screening in clinical practice using NCT.

Screening VAAs on non-contrast CT is fundamentally difficult. Without vascular enhancement, aneurysms cannot be separated from vessels by intensity contrast and must be inferred from subtle geometric dilation [18]. The complex branching patterns and high anatomical variability of visceral arteries further degrade vessel extraction and structural continuity [6, 20]. More critically, aneurysmal dilation is not a discrete object but a continuous deformation along the vessel wall. Under conventional voxel-wise classification, vessel and aneurysm labels compete at their interface, leading to unstable boundary learning. In addition, VAAs occupy extremely small spatial support, which weakens gradient accumulation and amplifies false positives at bifurcations and high-curvature segments where normal geometry mimics dilation. These factors make direct segmentation insufficient for reliable VAA screening on NCT.

To address these limitations, we build a dual-prior representation learning framework on top of a multi-label vascular segmentation backbone. Rather than modeling aneurysms as an additional discrete class, we augment the segmentation network with two prior-oriented branches that operate in continuous target spaces. The first branch learns a geometric confidence prior that represents aneurysmal dilation as a smooth structural response along the vessel wall, transferring lesion-specific geometric cues from the label domain into the feature space. The second branch learns an anatomical reliability prior that encodes geometry-induced ambiguity, such as bifurcations and tortuous segments, as a continuous penalty field. During inference, the reliability prior explicitly calibrates the confidence map to suppress anatomy-driven false responses. By integrating these dual priors into a unified segmentation framework, the pro-

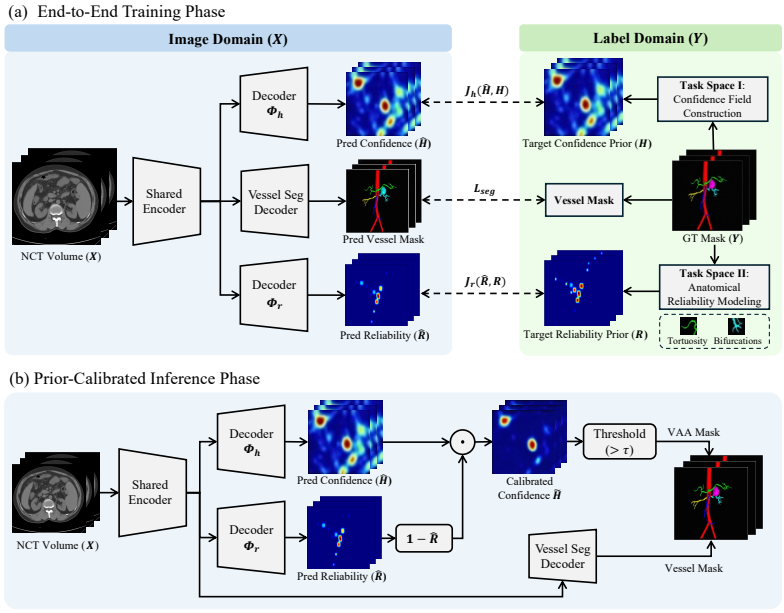

Fig. 1. Overview of DeepVAA. During training, a shared encoder learns vascular anatomy through multi-label segmentation while two auxiliary branches learn a geometric confidence prior and an anatomical reliability prior in continuous target spaces. During inference, the reliability map calibrates the confidence map via multiplicative suppression, and the calibrated response is thresholded to obtain the final VAA mask.

posed method separates dilation-sensitive responses from geometry-induced activations, leading to more stable and precise VAA screening on NCT.

Specifically, our contributions are three-fold: **First**, we introduce a dual-prior representation learning framework for NCT-based VAA screening, which jointly learns a geometric confidence prior for aneurysmal morphology and an anatomical reliability prior for ambiguity suppression. **Second**, we design a structured prior projection and calibration mechanism that amplifies lesion-sensitive responses through confidence modeling while explicitly suppressing geometry-induced false activations via reliability calibration, enabling stable detection under severe sparsity and low contrast. **Third**, we conduct extensive evaluations on multi-center cohorts, demonstrating consistent performance improvements and strong cross-site generalization for opportunistic VAA screening on NCT.

2 Methodology

2.1 Problem Formulation

Given an abdominal NCT volume $X \in \mathbb{R}^{H\times W\times D}$ defined over a 3D spatial domain Ω, and its corresponding discrete anatomical label Y, conventional multi-

label segmentation often struggles with the severe foreground sparsity and low contrast of VAAs. Let $\mathbf{x} \in \Omega$ denote a specific 3D voxel coordinate. Modeling VAAs directly as an independent categorical class introduces label competition at vessel–aneurysm interfaces, disrupting vascular continuity.

To elevate the learning paradigm from simple voxel-wise classification to structurally-aware representation learning, we propose to decompose the VAA screening task into the joint learning of vessel segmentation and two complementary prior representations, as shown in Fig. 1. Specifically, we use nnU-Net [11] as the backbone to extract deep intermediate representations, denoted as F, from the input volume X. We define two label-space transformations: a geometric confidence transformation $\mathcal{H}(\cdot)$ and an anatomical reliability transformation $\mathcal{R}(\cdot)$. Through task-specific branches built upon the Feature F, the framework is optimized to map these intermediate representations into the two transformed prior spaces simultaneously, thereby integrating continuous structural guidance with knowledge-driven anatomical constraints.

2.2 DeepVAA: Dual-Prior Representation Learning

Target Space I: Geometric Confidence Prior ($\mathcal{H}$). Clinically, a VAA is not an isolated lesion but a focal, continuous dilation along the vascular wall. In NCT screening, radiologists typically identify VAAs by tracing vessel continuity and spotting gradual morphological bulges rather than searching for discrete boundaries. To mimic this diagnostic intuition, we define the first transformation $\mathcal{H}(Y)$ to map the discrete VAA mask into a continuous, Gaussian-based spatial confidence field $H(\mathbf{x})$ evaluated at each voxel $\mathbf{x} \in \Omega$:

$$H(\mathbf{x}) = \mathcal{H}(Y)(\mathbf{x}) = \begin{cases} \frac{d(\mathbf{x})}{2 \times d_{\max}}, & \mathbf{x} \in \Omega_{\text{aneurysm}}, \\ 0, & \text{otherwise}, \end{cases}$$

where $d(\mathbf{x})$ denotes the Euclidean distance to the aneurysm boundary, and Ω_{aneurysm} represents the foreground VAA region. Crucially, because small VAAs occupy an extremely sparse sub-region of the spatial domain Ω, standard uniform distance metrics (e.g., standard cross-entropy) fail to accumulate sufficient gradient signals. Therefore, we couple this target space with a prior-modulated objective $\mathcal{J}_h$. Instead of treating the error at each voxel equally, we utilize the continuous target field $H(\mathbf{x})$ itself as a spatial modulator to redistribute the gradient mass toward high-confidence dilation centers:

$$\mathcal{J}_h(\hat{H}, H) = \mathcal{L}_{\text{reg}}(\hat{H}, H) + \alpha \frac{\sum_{\mathbf{x}} (1 - \hat{H}(\mathbf{x}))\, H(\mathbf{x})}{\sum_{\mathbf{x}} H(\mathbf{x})},$$

where $\hat{H} = \Phi_h(F)$ is the predicted geometric confidence, and $\mathcal{L}_{\text{reg}}$ represents the baseline regression loss (a combination of Dice and Focal losses). This intrinsic metric strictly forces the network to learn the aneurysmal dilation as a geometry-aware response along the vessel wall.

Target Space II: Anatomical Reliability Prior ($\mathcal{R}$). On NCT volumes, severe vascular tortuosity (high curvature) or complex bifurcations frequently mimic VAAs, leading to radiological pitfalls. To digitize the radiologist's "rule-out" process, we construct a second target space $\mathcal{R}(Y)$ that explicitly penalizes these anatomically ambiguous regions. Instead of defining hard exclusion rules, we construct a continuous penalty field through three intuitive steps:

1) Curvature Estimation: We first extract the vascular skeleton S from the ground-truth vessel mask Y_{vessel}. After necessary smoothing to ensure differentiability, the local curvature response is computed to highlight tortuous regions:

$$\kappa(\mathbf{x}) = \frac{|\Delta S(\mathbf{x})|}{\left(1 + \|\nabla S(\mathbf{x})\|\right)^2},$$

where ∇S and ΔS denote the spatial gradient and Laplacian, respectively.

2) Hotspot Identification: We identify a sparse set of structurally unstable hotspots, denoted as U. A voxel is assigned to U if it is a vascular bifurcation (detected via neighborhood connectivity) or exhibits abnormally high curvature ($\kappa(\mathbf{x})$ exceeding statistical structural norms).

3) Continuous Penalty Generation: To translate these discrete sparse hotspots into a differentiable target space, we apply a Gaussian-based spatial decay:

$$R(\mathbf{x}) = \mathcal{R}(Y)(\mathbf{x}) = \exp\left(-\frac{\mathcal{D}(\mathbf{x}, U)^2}{2\sigma_{\text{geom}}^2}\right),$$

where $\mathcal{D}(\mathbf{x}, U)$ denotes the shortest Euclidean distance from voxel $\mathbf{x}$ to the nearest hotspot in U. This formulation guarantees that $R(\mathbf{x}) \in (0, 1]$, smoothly assigning a penalty near 1.0 at ambiguous centers, which gradually fades as distance increases. Finally, the optimization objective for this space, $\mathcal{J}_r(\hat{R}, R)$, is defined as a standard regression loss to align the network's predicted reliability map $\hat{R} = \Phi_r(F)$ with this geometry-induced target $R(\mathbf{x})$.

2.3 Training and Inference Strategy

End-to-End Training Phase. The core philosophy of DeepVAA lies in transferring structural knowledge from the label space to the image representation space. During training, the analytical prior transformations $\mathcal{H}(\cdot)$ and $\mathcal{R}(\cdot)$ are applied exclusively to the ground truth Y to construct the continuous target spaces: $H = \mathcal{H}(Y)$ and $R = \mathcal{R}(Y)$. To approximate these targets, the shared nnU-Net encoder Φ_{enc} maps the input volume X into the deep intermediate representation $F = \Phi_{\text{enc}}(X)$. From this shared feature space F, the standard multi-class segmentation head Φ_{seg} outputs the vascular probability map $\hat{P}_{\text{vessel}} = \Phi_{\text{seg}}(F)$. Concurrently, the prior learning branches act as parameterized proxies, predicting the continuous representations via $\hat{H} = \Phi_h(F)$ and $\hat{R} = \Phi_r(F)$.

Let $\mathcal{L}_{\text{seg}}$ be the standard segmentation objective (a composite of Dice and Cross-Entropy losses) to supervise the multi-label vessel prediction against the discrete label Y_{vessel}, the overall loss function shall be defined as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{seg}}(\hat{P}_{\text{vessel}}, Y_{\text{vessel}}) + \mathcal{J}_h(\hat{H}, H) + \mathcal{J}_r(\hat{R}, R).$$

Therefore, the end-to-end optimization can unify the anatomical segmentation with the joint learning of these two prior spaces.

Prior-Calibrated Inference Phase. During inference, the explicit ground truth Y is unavailable. Instead, the optimized network branches independently perform the learned prior transformations directly on the input volume F, yielding the predicted geometric confidence $\hat{H} = \Phi_h(F)$ and the anatomical reliability penalty $\hat{R} = \Phi_r(F)$. To derive the final lesion probability, we perform a prior-calibrated fusion. The learned reliability map spatially modulates the geometric confidence through a voxel-wise inhibitory mapping:

$$\tilde{H}(\mathbf{x}) = \hat{H}(\mathbf{x}) \odot \big(1 - \hat{R}(\mathbf{x})\big).$$

Mathematically, this multiplicative calibration explicitly penalizes activation responses located in geometrically ambiguous or anatomically implausible regions (i.e., where $\hat{R}(\mathbf{x}) \to 1$). The continuous calibrated field is subsequently binarized via a threshold operator $\hat{Y}_\tau(\mathbf{x}) = \mathbb{I}\big(\tilde{H}(\mathbf{x}) > \tau\big)$, followed by connected-component analysis to extract instance-level VAA predictions.

3 Experiments

3.1 Dataset & Implementation Details

Dataset. The dataset was collected from two hospitals: the First Affiliated Hospital, Zhejiang University School of Medicine (Hospital A), and Anji County People's Hospital (Hospital B). **Hospital A** included 1000 cases and was used as the internal cohort, with 800 for training and 200 for internal validation. **Hospital B** with 100 cases was used entirely as the external validation cohort. The internal cohort has two scanner brands GE, UIH, and the external cohort has two more scanner brands Philips and Siemens. Train, valid and external datasets contain 563, 124 and 81 lesions, respectively (celiac trunk 79%, superior mesenteric 11%, right renal 6%, left renal 4%) and follow a negative, positive patient ratio of 1:1 $\pm$ 0.02. For each subject, both NCCT and CTA scans were collected, with NCCT registered to CTA using the deeds algorithm [8, 1, 10]. Visceral arteries and VAAs were manually annotated on CTA by two expert radiologists, and independently reviewed by another radiologist.

Implementation Details. All data preprocessing was performed within the nnU-Net framework [11]. During training, volumes were cropped to a fixed size of $(160 \times 320 \times 256)$, and the same cropping strategy was applied at inference for consistency. We adopted the standard nnU-Net data augmentation pipeline, including spatial transformations, Gaussian noise, Gaussian blur, brightness and contrast adjustments, and random mirroring. The model was trained using Adam optimizer [14] with an initial learning rate of 0.01 for 1000 epochs. All experiments were implemented in PyTorch and conducted on NVIDIA H200 GPUs.

Table 1. VAA screening performance against the state-of-the-art methods. In addition to reporting the Dice scores for vessel and VAA segmentation, we evaluate VAA precision and recall at the lesion level based on connected components of the predicted VAA segmentation. A predicted component is considered a true positive if it overlaps with a ground-truth VAA by at least one voxel; otherwise, it is counted as a false positive. We also compute the patient-level F1 score to assess case-wise detection performance.

Models	Internal Cohort					External Cohort				
	VAA Prec	VAA Recall	Patient F1	VAA Dice	Vessel Dice	VAA Prec	VAA Recall	Patient F1	VAA Dice	Vessel Dice
nnUNet [11]	72.79	86.29	88.12	61.28	90.16	65.88	77.46	85.98	58.23	**89.92**
SwinUNETR [7]	27.08	87.90	68.73	36.00	87.93	30.15	80.28	69.63	35.09	87.50
SegMamba [28]	45.74	**90.32**	79.48	47.14	89.44	37.43	81.69	78.83	45.79	88.75
MedNEXT [22]	65.85	85.48	85.02	55.98	89.62	59.14	76.06	81.08	51.55	88.99
UXNet [12]	24.87	70.97	65.25	17.30	82.71	23.28	54.93	59.32	17.70	82.40
U-Mamba [19]	69.81	89.52	89.22	62.51	90.06	51.30	**83.10**	84.21	52.28	89.31
Ours ($\tau = 0.05$)	75.00	85.48	89.58	66.14	**90.34**	72.37	78.87	90.20	63.98	89.81
Ours ($\tau = 0.10$)	78.46	85.48	**90.43**	66.08	**90.34**	76.39	77.46	90.20	64.03	89.81
Ours ($\tau = 0.15$)	80.80	82.26	89.84	66.38	**90.34**	80.60	77.46	91.09	64.29	89.81
Ours ($\tau = 0.20$)	81.97	81.45	90.32	66.80	**90.34**	82.81	76.06	91.09	64.57	89.81
Ours ($\tau = 0.25$)	82.64	81.45	91.30	67.43	**90.34**	84.13	76.06	91.09	65.07	89.81
Ours ($\tau = 0.30$)	**86.21**	80.65	91.21	**67.89**	**90.34**	**86.67**	74.65	**92.23**	**65.84**	89.81

Table 2. Ablation study on removing the Geometric Confidence Prior (GCP) and the Anatomical Reliability Prior (ARP). Experiments were conducted with $\tau = 0.10$ on the internal cohort.

GCP	ARP	VAA Prec	VAA Recall	Patient F1	VAA Dice	Vessel Dice
✗	✗	73.05	78.23	82.52	61.41	89.78
✓	✗	70.34	84.68	83.74	62.87	90.12
✗	✓	77.60	77.42	85.13	65.69	89.93
✓	✓	**78.46**	**85.48**	**90.43**	**66.08**	**90.34**

Table 3. Ablation study on the loss ratio α. Experiments were conducted with $\tau = 0.10$ on the internal cohort..

α	VAA Prec	VAA Recall	Patient F1	VAA Dice	Vessel Dice
0.00	77.69	77.42	85.13	65.69	89.93
0.05	78.57	79.84	86.37	65.32	90.07
0.1	**79.36**	81.85	88.02	**66.11**	89.62
0.5	78.46	85.48	**90.43**	66.08	**90.34**
1	76.39	**86.29**	89.58	65.44	89.89

3.2 Experimental Results

Comparison against the State-of-the-art. Table 1 and Fig. 2 compare our DeepVAA framework against state-of-the-art medical image segmentation models. Our approach consistently achieves the best overall performance across multiple threshold settings. On the internal cohort, it reaches a peak Patient F1 of 91.30% and a VAA Dice of 67.89%, clearly surpassing nnUNet and U-Mamba by optimizing the precision-recall balance rather than merely maximizing recall. On the unseen external cohort, where baseline methods show noticeable performance degradation, our DeepVAA framework demonstrates strong robustness and generalization. While nnUNet and U-Mamba experience significant drops in VAA Dice, our framework maintains stable performance across threshold variants and achieves a peak Patient F1 of 92.23%. By adjusting the threshold parameter τ, the model provides controllable sensitivity–specificity trade-offs while consis-

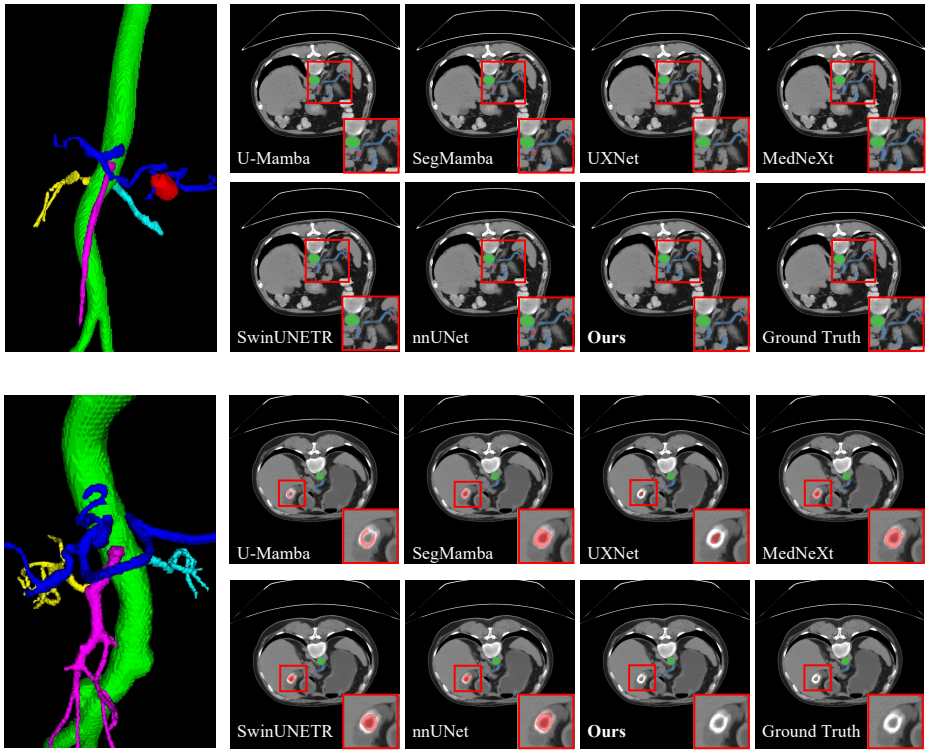

Fig. 2. Qualitative comparison of VAA and vessel segmentation (Red: VAA). Our method effectively detects the small VAA while reducing false positives.

tently delivering superior combined detection quality. τ 0.3 was predefined on validation data, maximizing the Youden Index at the ROC elbow.

Ablation on Prior Components. We performed an ablation study to validate the contributions of the GCP and ARP modules in Table 2. The baseline model achieves a VAA Dice of 61.41% and a Patient F1 of 82.52%. Integrating only the GCP module improves VAA Recall to 84.68%, enhancing sensitivity to visceral artery aneurysms. Conversely, using only the ARP module effectively suppresses false positives, boosting VAA Precision to 77.60%. Adding ARP alone decreased recall because one 7 mm VAA was missed among 124 subjects. However, given the marked reduction in false positives, missing a sub-10 mm VAA may be acceptable, as such lesions are typically not clinically critical. GCP and ARP are explicitly implemented as learnable branches optimized by losses which help the network genuinely learn anatomical representations. Overall, the full architecture yields the best performance across all metrics.

Ablation on Loss Ratio α. Table 3 investigates the impact of the hyper-

parameter α. The results indicate that incorporating this component ($\alpha > 0$) consistently enhances both visceral artery aneurysm detection and segmentation accuracy compared to $\alpha = 0.00$. Specifically, $\alpha = 0.5$ achieves the best overall performance, yielding peak Patient F1 (90.43%) and Vessel Dice (90.34%).

4 Conclusion

In this work, we propose a dual-prior representation learning framework for opportunistic screening of VAA on NCT. Instead of treating VAA detection as a pure segmentation problem, we model aneurysmal dilation and anatomical ambiguity in two continuous prior spaces. The geometric confidence prior enhances dilation-sensitive responses, while the anatomical reliability prior suppresses geometry-induced false activations through explicit calibration at inference. Experiments on multi-center cohorts demonstrate consistent improvements in lesion-level and patient-level performance, as well as strong cross-site generalization. Our results suggest that structured prior learning provides a reliable and clinically meaningful solution for early VAA detection on routine NCT, with potential to support large-scale opportunistic screening in practice.

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

References

1. Cao, K., Xia, Y., Yao, J., Han, X., Lambert, L., Zhang, T., Tang, W., Jin, G., Jiang, H., Fang, X., et al.: Large-scale pancreatic cancer detection via non-contrast ct and deep learning. Nature medicine **29**(12), 3033–3043 (2023)
2. Carnevale, M.L., Koleilat, I., Lipsitz, E.C., Friedmann, P., Indes, J.E.: Extended screening guidelines for the diagnosis of abdominal aortic aneurysm. Journal of vascular surgery **72**(6), 1917–1926 (2020)
3. Chiesa, R., Astore, D., Guzzo, G., Frigerio, S., Tshomba, Y., Castellano, R., de Moura, M.L., Melissano, G.: Visceral artery aneurysms. Annals of vascular surgery **19**(1), 42–48 (2005)
4. De Bruijne, M., van Ginneken, B., Viergever, M.A., Niessen, W.J.: Interactive segmentation of abdominal aortic aneurysms in cta images. Medical image analysis **8**(2), 127–138 (2004)
5. Gao, F., Chen, B., Zhou, T., Luo, H.: Research on the effect of visceral artery aneurysm's cardiac morphological variation on hemodynamic situation based on time-resolved ct-scan and computational fluid dynamics. Computer Methods and Programs in Biomedicine **221**, 106928 (2022)
6. Golla, A.K., Bauer, D.F., Schmidt, R., Russ, T., Nörenberg, D., Chung, K., Tönnes, C., Schad, L.R., Zöllner, F.G.: Convolutional neural network ensemble segmentation with ratio-based sampling for the arteries and veins in abdominal ct scans. IEEE Transactions on Biomedical Engineering **68**(5), 1518–1526 (2020)
7. Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H.R., Xu, D.: Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. In: International MICCAI brainlesion workshop. pp. 272–284. Springer (2021)

8. Heinrich, M.P., Jenkinson, M., Brady, M., Schnabel, J.A.: Mrf-based deformable registration and ventilation estimation of lung ct. IEEE transactions on medical imaging **32**(7), 1239–1248 (2013)
9. Hu, Y., Xiang, Y., Zhou, Y.J., He, Y., Lang, D., Yang, S., Du, X., Den, C., Xu, Y., Wang, G., et al.: Ai-based diagnosis of acute aortic syndrome from noncontrast ct. Nature Medicine **31**(11), 3832–3844 (2025)
10. Hu, Y., Xiang, Y., Zhou, Y.J., He, Y., Lang, D., Yang, S., Du, X., Den, C., Xu, Y., Wang, G., et al.: Ai-based diagnosis of acute aortic syndrome from noncontrast ct. Nature Medicine **31**(11), 3832–3844 (2025)
11. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. Nature methods **18**(2), 203–211 (2021)
12. Ji, Y., Zhang, R., Li, Z., Ren, J., Zhang, S., Luo, P.: Uxnet: Searching multi-level feature aggregation for 3d medical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 346–356. Springer (2020)
13. Khairallah, M., Morgan, R., Das, R.: Technical considerations of endovascular management of true visceral artery aneurysms. CVIR endovascular **6**(1), 31 (2023)
14. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
15. Lareyre, F., Adam, C., Carrier, M., Dommerc, C., Mialhe, C., Raffort, J.: A fully automated pipeline for mining abdominal aortic aneurysm using image segmentation. Scientific Reports **9**(1), 13750 (2019)
16. Lareyre, F., Caradu, C., Chaudhuri, A., Lê, C.D., Di Lorenzo, G., Adam, C., Carrier, M., Raffort, J., Coscas, R., Jayet, J., et al.: Automatic detection of visceral arterial aneurysms on computed tomography angiography using artificial intelligence based segmentation of the vascular system. In: EJVES Vascular Forum. vol. 59, pp. 15–19. Elsevier (2023)
17. Lv, R., Hu, G., Zhang, S., Zhang, Z., Chen, J., Wang, K., Wang, Z., Jin, Z.: Assessing abdominal aortic aneurysm growth using radiomic features of perivascular adipose tissue after endovascular repair. Insights into Imaging **15**, 232 (2024)
18. Lyu, Z., Mu, N., Rezaeitaleshmahalleh, M., Zhang, X., McBane, R., Jiang, J.: Automatic segmentation of intraluminal thrombosis of abdominal aortic aneurysms from ct angiography using a mixed-scale-driven multiview perception network (m2net) model. Computers in Biology and Medicine **179**, 108838 (2024)
19. Ma, J., Li, F., Wang, B.: U-mamba: Enhancing long-range dependency for biomedical image segmentation. arXiv preprint arXiv:2401.04722 (2024)
20. Obara, H., Kentaro, M., Inoue, M., Kitagawa, Y.: Current management strategies for visceral artery aneurysms: an overview. Surgery today **50**(1), 38–49 (2020)
21. Roby, M., Sakib, A.N.M., Zhang, Z., Muluk, S.C., Eskandari, M.K., Finol, E.A.: Automatic explainable segmentation of abdominal aortic aneurysm from computed tomography angiography. IEEE Access **13**, 178021–178033 (2025)
22. Roy, S., Koehler, G., Ulrich, C., Baumgartner, M., Petersen, J., Isensee, F., Jaeger, P.F., Maier-Hein, K.H.: Mednext: transformer-driven scaling of convnets for medical image segmentation. In: International conference on medical image computing and computer-assisted intervention. pp. 405–415. Springer (2023)
23. Ru, X., Wang, X., Liu, Z., Du, P., Zhao, H., Wu, Z., Ju, X., Liu, S., Zhu, Y.C., Frangi, A.F.: Pose-independent efficient gauge equivariant network for 3d mesh aneurysm segmentation. Neurocomputing **639**, 130188 (2025)
24. Sen, I., Franco-Mesa, C., Erben, Y., DeMartino, R.R.: Abdominal aortic and visceral artery aneurysms. Cardiology clinics **39**(4), 517–525 (2021)

25. Spinella, G., Fantazzini, A., Finotello, A., Vincenzi, E., Boschetti, G.A., Brutti, F., Magliocco, M., Pane, B., Basso, C., Conti, M.: Artificial intelligence application to screen abdominal aortic aneurysm using computed tomography angiography. Journal of Digital Imaging **36**(5), 2125–2137 (2023)
26. Wan, M., Zhu, J., Che, Y., Cao, X., Han, X., Si, X., Wang, W., Shu, C., Luo, M., Zhang, X.: Bif-net: Boundary information fusion network for abdominal aortic aneurysm segmentation. Computers in Biology and Medicine **183**, 109191 (2024)
27. Weiss, D., Hager, T., Lin, M., Sritharan, D., Bousabarah, K., Renninghoff, D., Holler, W., Simmons, K., Haubold, J., Loh, S., Fischer, U., Chapiro, J., Deuschl, C., Aboian, M., Aboian, E., Aneja, S.: Deep learning based volumetric analysis of infrarenal abdominal aortic aneurysms characterized on cta. npj Digital Medicine **9**, 87 (2026)
28. Xing, Z., Ye, T., Yang, Y., Liu, G., Zhu, L.: Segmamba: Long-range sequential modeling mamba for 3d medical image segmentation. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. pp. 578–588. Springer Nature Switzerland (2024)
29. Yao, J., Ye, X., Xia, Y., Zhou, J., Shi, Y., Yan, K., Wang, F., Lin, L., Yu, H., Hua, X.S., et al.: Effective opportunistic esophageal cancer screening using non-contrast ct imaging. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 344–354. Springer (2022)
30. Zhang, B., Lai, Z., Liu, S., Xie, X., Zhou, X., Hou, Z., Ma, X., Liu, B., Li, K., Song, M.: Sdlu-net: A similarity-based dynamic linking network for the automated segmentation of abdominal aorta aneurysms and branching vessels. Biomedical Signal Processing and Control **100**, 106991 (2025)