Towards Cross-View Consistent Aneurysm Detection in DSA

Yaotian Wang[1] 1767178707@qq.com
Jax Luo[1] luoj2@ccf.org
Scott Raymond[1] raymon3@ccf.org
[1] *Neurological Institute, Cleveland Clinic, OH, USA*

Abstract

Existing aneurysm detection methods in digital subtraction angiography (DSA) typically operate on single-view projections, treating each image independently, which is inconsistent with clinical workflow. In practice, clinicians interpret biplane DSA by jointly examining anteroposterior (AP) and lateral views. As complementary projections of the same anatomy, these cross-views share geometric and anatomical relationships that can be leveraged to improve detection. Our preliminary results support this hypothesis. The code will be made publicly available at the time of the conference.

Keywords: DSA, Aneurysm Detection, Cross-view Consistency

1. Introduction

Accurate detection of intracranial aneurysms is critical for neurointerventional treatment. In high-stakes procedures such as coiling and flow diversion, precise localization directly guides device deployment, reduces procedure time, and improves patient outcomes. This clinical demand has driven increasing interest in automated aneurysm detection methods. Most existing approaches focus on computed tomography angiography (CTA) and magnetic resonance angiography (MRA), where large publicly available annotated datasets facilitate model development(Hu et al., 2024; Wei et al., 2024; You et al., 2025; Ryu et al., 2025; Li et al., 2024). However, these methods are not fully aligned with the intraoperative workflow. In clinical practice, digital subtraction angiography (DSA) remains the gold standard, providing real-time visualization of vascular anatomy during intervention.

Despite its central role, automated aneurysm detection in DSA remains relatively underexplored. Existing studies are typically limited to small, proprietary datasets and, more importantly, operate on single-view projections, treating each DSA image independently(Liu et al., 2025; Hu et al., 2023; Liao et al., 2022; Liu et al., 2021). This assumption is inconsistent with the actual workflow, as clinicians routinely interpret biplane DSA during interventions, simultaneously examining anteroposterior (AP) and Lateral views to localize aneurysms and assess their three-dimensional structure. AP and Lateral views are not independent. Rather, they are complementary projections of the same underlying anatomy acquired from the same patient at the same time. As such, they share intrinsic geometric and anatomical relationships that can provide valuable cues for detection. Leveraging this cross-view relationship has the potential to improve robustness and accuracy, particularly in challenging cases where aneurysms may be obscured or ambiguous in a single view.

In this work, we propose a cross-view consistent framework for aneurysm detection in DSA that jointly models AP and lateral projections. Here, cross-view refers to paired AP and lateral images from the same DSA acquisition, and cross-view consistency is applied to ensure anatomically coherent detections.

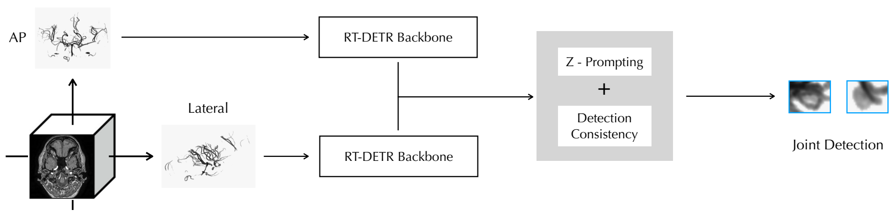

Figure 1: An illustration of the cross-view consistent aneurysm detection framework.

2. Method

2.1. DSA Dataset generation

In DSA, AP and lateral views are orthogonal projections used to visualize cerebrovascular anatomy. The AP view is obtained by projecting along the anterior–posterior axis, yielding a coronal-plane image, whereas the lateral view is generated by projecting along the left–right axis, producing a sagittal-plane image. To approximate DSA from volumetric data, we generate DSA-like projections from the ADAM dataset(Timmins et al., 2020, 2021). Specifically, we first perform skull stripping to remove non-vascular structures, followed by maximum intensity projection (MIP) to obtain two-dimensional angiographic representations. This process collapses the volumetric data while preserving high-intensity vascular signals, thereby approximating the projection mechanism of DSA. AP views are generated by projecting along the anterior–posterior axis, and lateral views along the left–right axis, resulting in paired projections that emulate biplane DSA. To obtain corresponding annotations, the 3D aneurysm labels are subjected to the same projection operations, producing aligned two-dimensional localization targets in both AP and lateral views. The resulting dataset consists of paired AP–lateral projections with spatially consistent annotations.

2.2. Cross-view Consistent Aneurysm Detection

Given a 3D point (x, y, z), the AP view maps it to (x, z), while the lateral view maps it to (y, z). As a result, both views share the same superior–inferior coordinate z. Building upon a transformer-based detection model, RT-DETR (Zhao et al., 2024; Carion et al., 2020), we introduce a z-prompting mechanism that exploits this shared geometry. Specifically, given a detected aneurysm in the AP view, the model attends to the corresponding z-coordinate in the lateral view to facilitate cross-view localization, and vise versa. In addition, we impose a detection consistency loss to encourage coherent detections across AP and lateral views. An illustration of the overall framework is shown in Fig. 1.

3. Experiments

3.1. Setup and Implementation

We evaluate the proposed cross-view approach on paired AP and lateral views and compare it with single-view baselines, where each branch is trained and evaluated independently without access to the counterpart view. This comparison assesses the effectiveness of cross-view consistency in improving detection performance. For single-view baselines, we consider YOLO26n(Jocher and Qiu, 2026) and RT-DETR(Zhao et al., 2024; Carion et al., 2020).

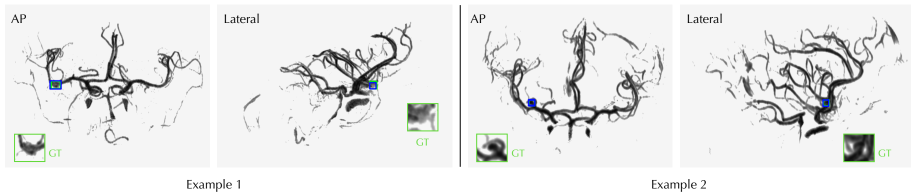

Figure 2: Two examples of aneurysm detection jointly in AP and lateral views. Blue indicates the predicted bounding box.

Table 1: Quantitative evaluation. Best results for each view are highlighted in bold.

	AP				Lateral			
Method	Precision	Recall	mAP50	mAP50-95	Precision	Recall	mAP50	mAP50-95
YOLO26n	0.685	0.353	0.336	0.126	0.556	0.176	0.199	0.071
RT-DETR	0.428	0.294	0.270	0.110	0.516	0.118	0.152	0.074
Cross-View (A→L)	0.750	0.704	0.535	0.213	**0.979**	**0.471**	**0.480**	0.167
Cross-View (L→A)	**0.901**	0.647	0.619	0.220	0.715	**0.471**	0.456	0.159
Cross-View (A↔L)	0.764	**0.765**	**0.643**	**0.229**	0.872	0.412	0.441	**0.177**

For the proposed cross-view approach, we evaluate three prompting variants: AP-to-Lateral (A→L), Lateral-to-AP (L→A), and joint prompting (A↔L). Metrics include Precision, Recall, mAP50, and mAP50-95. We use a total of 113 paired cases, split into training, validation, and test sets with a ratio of 8:1:1. All images are resized to 640×640. Models are trained for 300 epochs with a batch size of 8 using the AdamW optimizer, with an initial learning rate of 1×10^{-4} and weight decay of 1×10^{-4}. A detection-consistency loss with weight 0.05 is applied during training.

3.2. Results

At inference, only a single AP or lateral DSA image is used as input. As shown in Table 1, all models trained with cross-view consistency outperform single-view baselines. Notably, the performance gains are direction-dependent. AP-to-lateral prompting yields the largest improvement on lateral views, suggesting that the AP branch provides higher-quality guidance. In contrast, while joint prompting achieves the best performance on AP views, it does not yield the highest mAP@50 on lateral views, indicating that cross-view assistance is asymmetric and depends on the quality of the source-view prompts. Further inspection suggests that the synthesized lateral DSA images are of lower quality than the AP views, which may explain the observed asymmetry. Qualitative results are presented in Fig.2.

4. Discussion

Motivated by clinical needs and insights, we develop a cross-view consistent framework for aneurysm detection in DSA that leverages the shared anatomical structure between AP and lateral projections. Further studies are warranted to more systematically evaluate this approach.

References

Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I*, pages 213–229. Springer, 2020. doi: 10.1007/978-3-030-58452-8_13.

Bin Hu, Zhao Shi, et al. A deep-learning model for intracranial aneurysm detection on ct angiography images in china: a stepwise, multicentre, early-stage clinical validation study. *The Lancet Digital Health*, 6(4):e261–e271, 2024. doi: 10.1016/S2589-7500(23)00268-6.

Tao Hu, Heng Yang, and Wei Ni. A framework for intracranial aneurysm detection and rupture analysis on dsa. *Journal of Clinical Neuroscience*, 115:101–107, 2023. doi: 10.1016/j.jocn.2023.07.025.

Glenn Jocher and Jing Qiu. Ultralytics yolo26, 2026. URL https://github.com/ultralytics/ultralytics.

Yuanyuan Li, Huiling Zhang, Yun Sun, Qianrui Fan, Long Wang, Congshan Ji, et al. Deep learning-based platform performs high detection sensitivity of intracranial aneurysms in 3d brain tof-mra: An external clinical validation study. *International Journal of Medical Informatics*, 188:105487, 2024. doi: 10.1016/j.ijmedinf.2024.105487.

Jun-Heng Liao et al. Using a convolutional neural network and convolutional long short-term memory to automatically detect aneurysms on 2d digital subtraction angiography images: Framework development and validation. *JMIR Medical Informatics*, 10(3):e28880, 2022. doi: 10.2196/28880.

Ruibo Liu, Ruixuan Zhang, Wei Qian, Guobiao Liang, Guangxin Chu, Hai Jin, Ligang Chen, Jing Li, and He Ma. Intracranial aneurysm segmentation on digital subtraction angiography: a retrospective and multi-center study. *Frontiers in Neurology*, 16:1646517, 2025. doi: 10.3389/fneur.2025.1646517.

Xinke Liu, Junqiang Feng, Zhenzhou Wu, Zhonghao Neo, Chengcheng Zhu, Peifang Zhang, Yan Wang, Yuhua Jiang, Dimitrios Mitsouras, and Youxiang Li. Deep neural network-based detection and segmentation of intracranial aneurysms on 3d rotational dsa. *Interventional Neuroradiology*, 27(5):648–657, 2021. doi: 10.1177/15910199211000956.

Wi-Sun Ryu, Sungmoon Jeong, Jaechan Park, Dougho Park, Heeyoung Kim, Myungjae Lee, et al. Diagnostic accuracy of a deep learning algorithm for detecting unruptured intracranial aneurysms in magnetic resonance angiography: A multicenter pivotal trial. *World Neurosurgery*, 197:123882, 2025. doi: 10.1016/j.wneu.2025.123882.

Kimberley Timmins, Edwin Bennink, Irene van der Schaaf, Birgitta Velthuis, Ynte Ruigrok, and Hugo Kuijf. Intracranial aneurysm detection and segmentation challenge, 2020.

Kimberley M. Timmins, Irene C. van der Schaaf, Edwin Bennink, Ynte M. Ruigrok, Xingle An, others, and Hugo J. Kuijf. Comparing methods of detecting and segmenting unruptured intracranial aneurysms on tof-mras: The adam challenge. *NeuroImage*, 238:118216, 2021. doi: 10.1016/j.neuroimage.2021.118216.

Jianyong Wei, Xinyu Song, Xiaoer Wei, et al. Knowledge-augmented deep learning for segmenting and detecting cerebral aneurysms with ct angiography: A multicenter study. *Radiology*, 312(2):e233197, 2024. doi: 10.1148/radiol.233197.

Wei You, Junqiang Feng, Jing Lu, Ting Chen, Xinke Liu, Zhenzhou Wu, et al. Diagnosis of intracranial aneurysms by computed tomography angiography using deep learning-based detection and segmentation. *Journal of NeuroInterventional Surgery*, 17:e132–e138, 2025. doi: 10.1136/jnis-2023-021022.

Yian Zhao, Wenyu Lv, Shangliang Xu, Jinman Wei, Guanzhong Wang, Qingqing Dang, Yi Liu, and Jie Chen. Detrs beat yolos on real-time object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16965–16974, June 2024.