RDE-Seg: Role-Disentangled Experts with Residual Routing and Anatomy Constraints for DSA Guidewire Segmentation

Yining Zhang[1], Chengjiang Long[2†], Wenju Xu[3], Tingbao Zhang[4], Yu Feng[4], and Jianhui Zhao[1*]

[1] School of Computer Science, Wuhan University, Wuhan, China
{2020302111300,jianhuizhao}@whu.edu.cn
[2] ByteDance Inc.
[3] Amazon
[4] Zhongnan Hospital of Wuhan University, Wuhan University, Wuhan, China

Abstract. Accurate guidewire segmentation is essential for safe navigation in fluoroscopy-guided interventions but remains challenging in Digital Subtraction Angiography (DSA) images. Beyond pixel-level difficulties, guidewires are tightly constrained by vascular anatomy, a dependency often ignored by existing methods that segment them independently. To address this problem, we propose RDE (Role-Disentangled Experts), a framework combining parameter-efficient domain specialization, role-disentangled decoding, and conservative residual fusion. We instantiate RDE on the Segment Anything Model (SAM), resulting in RDE-SAM. RDE-SAM injects LoRA into the image encoder and role-disentangled experts with residual routing into the decoder. We further incorporate a lightweight anatomy-constrained refinement that uses the predicted vessel mask to define a tolerance band and repair small guidewire discontinuities, improving structural consistency. On a clinically collected DSA dataset of over 3,000 images with dual pixel-level annotations, RDE-SAM consistently outperforms strong baselines, demonstrating the benefit of explicitly modeling anatomical constraints for thin-structure instrument segmentation.

Keywords: DSA Dataset · Guidewire Segmentation · Expert Guidance · LoRA · Anatomy Constraints

1 Introduction

Accurate guidewire segmentation is essential in fluoroscopy-guided interventional procedures, where guidewires serve as structural backbones for device navigation and manipulation [24, 25]. Beyond a conventional thin-structure segmentation problem, guidewire extraction belongs to a broader class of anatomy-constrained

*Corresponding author.
†The work was completed when Dr. Chengjiang Long was at Meta.

instrument segmentation tasks [3]. In clinical practice, feasible guidewire configurations are tightly coupled with vascular anatomy [12], meaning anatomical context imposes structural constraints rather than serving as auxiliary cues. However, most existing approaches treat guidewires as independent segmentation targets [24, 30], overlooking their intrinsic dependency on vessels.

Guidewire segmentation in Digital Subtraction Angiography (DSA) remains challenging due to class imbalance, motion blur, and low signal-to-noise ratios [31, 2, 15]. Early methods relied on geometric modeling and tracking [22, 4, 19], while more recent works adopt UNet-based [17, 1, 20, 26] and Transformer-based architectures [27, 28]. Although these methods enhance pixel-level accuracy, they treat guidewire segmentation as an independent binary task without explicitly modeling anatomical constraints between guidewires and vessels.

Foundation models such as SAM [10] demonstrate strong generalization in natural images, yet domain gaps limit their direct application to medical imaging [6, 16]. Medical adaptations via full fine-tuning or parameter-efficient strategies [14, 23, 29, 5] alleviate domain mismatch but still lack structural disentanglement and constraint-aware modeling. Meanwhile, Mixture-of-Experts (MoE) architectures [8, 18, 11] enable specialized representation learning and have shown effectiveness in large-scale models [13], yet their potential for disentangling anatomical and instrumental representations in medical segmentation remains underexplored [21]. A general and training-stable framework for anatomy-constrained segmentation is still lacking.

To address this gap, we propose RDE, a Role-Disentangled Experts framework for anatomy-constrained instrument segmentation. RDE is architecture-agnostic and integrates three components: parameter-efficient domain specialization via LoRA [7], role-disentangled decoding for anatomical and instrumental structures, and conservative residual routing that preserves baseline stability while enabling expert refinement. RDE-SAM instantiates this framework on SAM, serving as a concrete implementation for guidewire segmentation.

Within RDE, vessel masks are jointly predicted to provide structural constraints for guidewire segmentation. Vessel predictions regularize guidewire outputs to enhance structural continuity and suppress anatomically implausible responses. A lightweight refinement module further enforces tolerance-band consistency and repairs small discontinuities through morphology-aware operations.

To evaluate anatomy-constrained modeling, we construct a clinically collected DSA dataset of over 3,000 images with dual pixel-level annotations for guidewires and vessels. Experimental results demonstrate that RDE-SAM consistently outperforms the strong baselines in Dice and HD95.

Our contributions are three-fold:

- We propose RDE, a general and architecture-agnostic framework for anatomy-constrained instrument segmentation, integrating LoRA-based domain specialization, role-disentangled decoding, and conservative residual routing.
- We explicitly model the anatomical dependency between guidewires and vessels through joint segmentation and constraint-aware refinement, improving structural continuity and spatial plausibility.

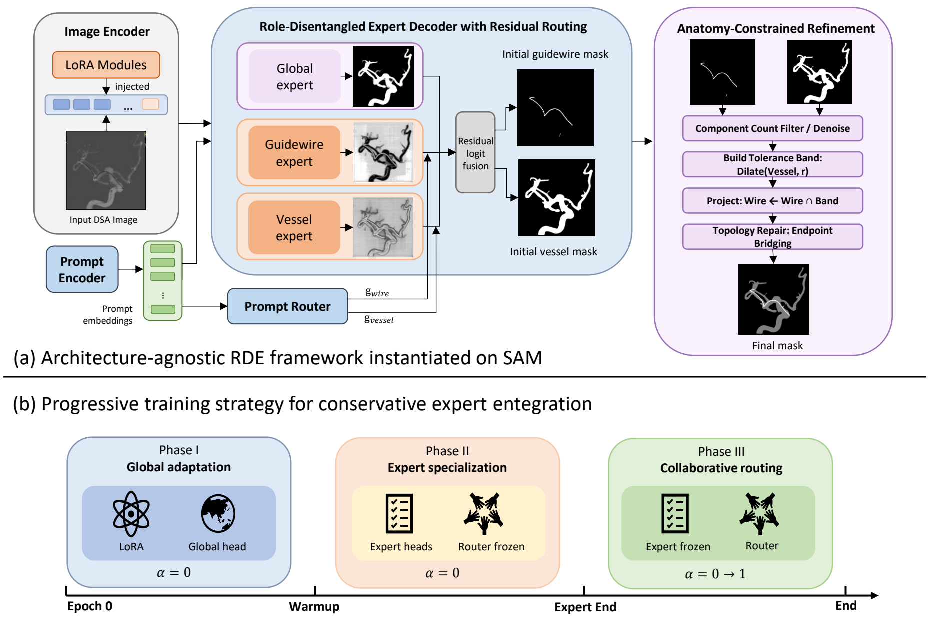

Fig. 1: Overview of the proposed RDE framework instantiated on SAM. (a) Architecture-agnostic design: LoRA-adapted image encoder, role-disentangled expert decoder, and a prompt-routed router that performs conservative residual fusion to generate initial multi-class masks. An anatomy-constrained refinement module further enforces spatial feasibility and topological consistency via tolerance-band projection and endpoint-based repair. (b) Progressive training strategy, including global adaptation, expert specialization, and collaborative routing, to ensure stable optimization and controlled expert integration.

– We construct a dual-annotation DSA dataset to validate anatomy-constrained modeling and demonstrate the effectiveness of RDE-Seg.

2 Method

2.1 Overview

We formulate anatomy-constrained instrument segmentation as a structured prediction problem in which instrumental structures are restricted by anatomical context. To address this class of tasks, we propose RDE, a general and architecture-agnostic framework that integrates parameter-efficient domain specialization, role-disentangled decoding for anatomical and instrumental representations, and conservative residual routing to enable stable expert refinement. As illustrated in Fig. 1, RDE can be instantiated on standard encoder–decoder backbones; in this work, we adopt SAM as a concrete implementation (RDE-SAM) to validate its effectiveness on guidewire segmentation in DSA images.

2.2 LoRA-based Domain Specialization

The first component of RDE performs parameter-efficient domain specialization. Rather than fully fine-tuning the backbone, which risks overfitting to limited medical data, we adopt LoRA to inject domain-specific biases into the image encoder while preserving the general visual priors.

Specifically, we freeze the majority of the SAM encoder and introduce low-rank updates to the query and value projection matrices within each transformer block. For an original projection W, LoRA parameterizes the update as:

$$\hat{W} = W + \Delta W = W + BA, \tag{1}$$

where $A \in \mathbb{R}^{r \times C_{\text{in}}}$ and $B \in \mathbb{R}^{C_{\text{out}} \times r}$ with $r \ll \min(C_{\text{in}}, C_{\text{out}})$. The adapted features are then computed as:

$$\hat{F} = \hat{W} F, \tag{2}$$

where $F \in \mathbb{R}^{C_{\text{in}} \times N}$ denotes the tokenized image feature map with N tokens. This design enables efficient adaptation to angiographic imaging characteristics while maintaining the stability and generalization of the pretrained encoder.

2.3 Role-Disentangled Expert Decoder with Residual Routing

To model the anatomical–instrumental dependency, RDE introduces the role-disentangled expert heads within the decoder. Instead of building independent full decoders, all branches share the SAM Two-Way Transformer and mask up-scaling module. The global branch uses the original SAM hypernetwork to produce $\mathbf{L}_{\text{global}}$, while the guidewire and vessel branches add two lightweight hypernetwork heads applied to the same shared upscaled embedding to produce $\mathbf{L}_{\text{wire}}$ and $\mathbf{L}_{\text{vessel}}$.

A prompt-routed router is implemented as a two-layer MLP that takes the learnable implicit prompt token as input and predicts image-level scalar gating coefficients $g_{\text{wire}}, g_{\text{vessel}} \in [0, 1]$. These coefficients modulate expert contributions globally per image:

$$\mathbf{L}_{\text{final}} = \mathbf{L}_{\text{global}} + \alpha \left(g_{\text{wire}} \mathbf{L}_{\text{wire}} + g_{\text{vessel}} \mathbf{L}_{\text{vessel}} \right), \tag{3}$$

where α is progressively increasing during training.

This residual formulation ensures that the baseline global pathway is always preserved, preventing catastrophic degradation while allowing expert specialization to refine thin-structure continuity and connectivity.

The total loss is defined as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{global}} + \lambda_{\text{wire}} \mathcal{L}_{\text{wire}} + \lambda_{\text{vessel}} \mathcal{L}_{\text{vessel}}. \tag{4}$$

Each loss term consists of a combination of Dice loss and CE loss computed against the corresponding ground-truth mask. The weighting coefficients λ_{wire} and λ_{vessel} balance specialization strength between expert branches.

2.4 Anatomy-Constrained Refinement

We interpret anatomical refinement as a projection of predictions onto a vascular-consistent feasible set. Given predicted masks $\hat{M}_1$ (vessel) and $\hat{M}_2$ (guidewire), we first suppress isolated noise via connected-component filtering. A vascular tolerance band is constructed by dilation:

$$\hat{M}_1^{\mathrm{dil}} = \mathrm{Dilate}(\hat{M}_1; r). \tag{5}$$

Guidewire predictions are then constrained by:

$$\hat{M}_2 \leftarrow \hat{M}_2 \cap \hat{M}_1^{\mathrm{dil}}, \tag{6}$$

which removes anatomically implausible detections outside the vascular region.

To repair small discontinuities, skeleton endpoints are detected and endpoint pairs with $\|p_i - p_j\|_2 < d_{\max}$ are bridged within the tolerance band:

$$\hat{M}_2 \leftarrow \hat{M}_2 \cup \mathrm{Bridge}(\hat{M}_2). \tag{7}$$

This refinement enforces structural continuity while respecting the anatomical constraints during DSA guidewire segmentation.

2.5 Progressive Conservative Training

To ensure stable optimization under capacity expansion, we adopt a progressive training strategy including three phases.

Phase I trains only the global branch with LoRA adaptation, reducing the model to a SAM+LoRA baseline.

Phase II activates expert heads while keeping router weights fixed, allowing role specialization without perturbing the baseline prediction.

Phase III enables the router with a smaller learning rate. Since expert contributions are injected residually, the global pathway is always preserved, guaranteeing that worst-case performance remains comparable to the baseline.

This staged activation ensures a smooth transition from global prediction to expert-enhanced anatomy-constrained segmentation.

3 Experiments

3.1 Dataset

Under the supervision of Medical Ethics Committee with Ethical Approval, we construct a clinically collected DSA dataset comprising 3,069 fluoroscopic images from approximately twenty patients undergoing diverse interventional procedures, including aneurysm repair, internal carotid artery stenosis treatment, and vascular malformation interventions. All images are resized to 688×688 pixels. A strict patient-level split yields 2,584 training images and 485 test images. Pixel-level annotations for guidewires and vessels are performed using LabelMe and cross-verified by two experienced clinicians.

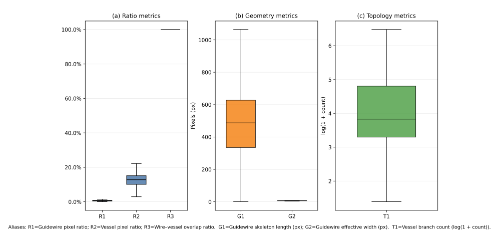

(a) Boxplot visualization of the structural statistics across samples.

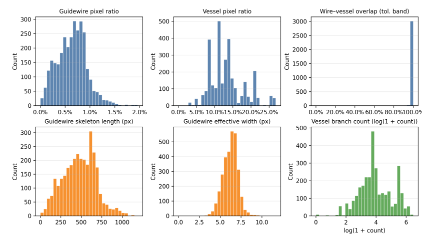

(b) Histogram distribution of the structural metrics across samples.

Fig. 2: Structural statistics of the proposed DSA dataset. The boxplots and histograms illustrate guidewire/vessel pixel ratios, guidewire skeleton length, vessel branch count, and guidewire–vessel overlap ratio, highlighting extreme class imbalance and strong anatomical coupling.

Unlike conventional thin-structure datasets, our dataset explicitly exhibits anatomy-constrained characteristics. We quantitatively analyze its structural properties to characterize task difficulty and anatomical dependency.

Guidewires occupy only $0.64\% \pm 0.30\%$ of image pixels on average, confirming extreme class imbalance. In contrast, vessels occupy $13.27\% \pm 4.66\%$, forming complex tree-like anatomical networks. The average guidewire skeleton length is 480 ± 201 pixels, indicating elongated thin structures spanning large portions of the field of view. The effective annotated guidewire width (estimated by area-to-skeleton ratio) is 6.23 ± 0.90 pixels.

Vessel skeleton analysis reveals an average of 91.7 ± 104.6 branch points per image, reflecting substantial anatomical complexity. Importantly, the ground-truth guidewire–vessel overlap ratio reaches 99.7%, demonstrating that guidewire configurations are strongly constrained by vascular structures.

Structural statistics show a mean guidewire component count of 1.12 and a skeleton continuity ratio of 0.992, indicating nearly contiguous ground-truth structures and a well-defined anatomical feasible set.

Figure 2 summarizes the distribution of pixel ratios, structural lengths, branch counts, and overlap ratios, confirming that guidewire segmentation in this dataset constitutes an anatomy-constrained structured prediction problem.

3.2 Results

Our model is trained with AdamW using a learning rate of 0.005 and a LoRA rank of 4. The anatomy-constrained refinement uses $r = 8$ and $d_{\max} = 20$ by default. Code will be released at https://github.com/singer502/RDE-Seg.

We compare RDE-SAM with three representative categories of baselines. First, DARL [9], a vascular segmentation model for elongated structures, serves as a reference for modeling thin and topologically complex patterns. Second, SAM-based medical adaptation models, including Medical-SAM-adapter [23],

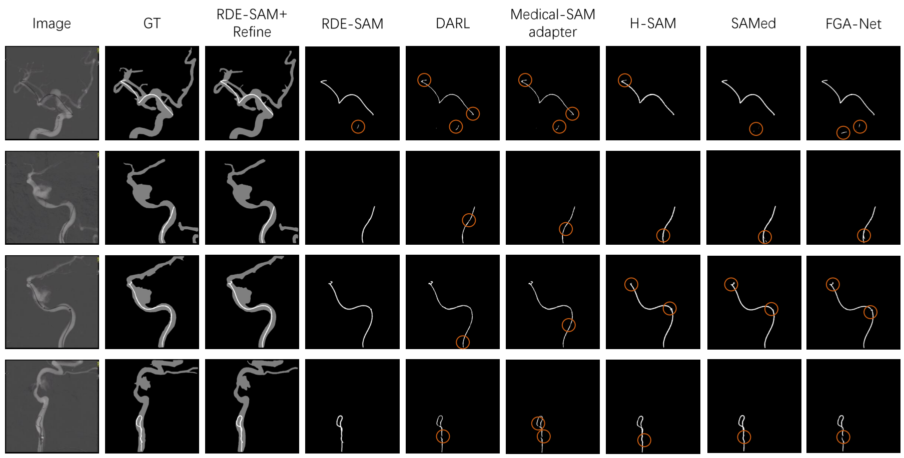

Fig. 3: Qualitative comparison of the segmentation results on representative test cases. From left to right: input image, ground truth, RDE-SAM+Refine, RDE-SAM, DARL, Medical-SAM-adapter, H-SAM, SAMed, and FGA-Net. Orange circles indicate regions of deficiencies from the segmentation methods.

H-SAM [5], and SAMed [29], provide baselines built upon the same foundation-model backbone. Third, FGA-Net [26], a dedicated guidewire segmentation model, provides a task-specific benchmark.

As shown in Table 1, RDE-SAM with refinement achieves the best overall performance, obtaining the highest Dice and the lowest HD95. Compared with the strongest SAM-based baseline (H-SAM), RDE-SAM reduces HD95 from 6.89 to 2.36, indicating improved boundary accuracy and structural continuity. On a single RTX 4090 GPU (batch size 1), our model runs at averaged 176.23 ms per image. The qualitative comparisons are shown in Figure 3.

3.3 Ablation Study

Table 2 presents an ablation of the proposed components. The **SAM+LoRA** baseline performs parameter-efficient domain adaptation of the SAM encoder. Adding **Experts** introduces role-specific heads for guidewire and vessel modeling, improving both Dice and HD95. Enabling the **Router** further injects expert predictions into the global branch through conservative residual fusion, reducing HD95 from 8.32 to 2.98 px and indicating improved thin-structure continuity. Finally, the anatomy-constrained **Refinement** improves Dice from 83.9% to 84.2% and HD95 from 2.98 to 2.36 px, confirming the benefit of tolerance-band projection and endpoint bridging after network prediction.

The refinement is insensitive to moderate changes in its two key hyperparameters, remaining stable for $r \in [6, 10]$ and $d_{\max} \in [15, 25]$. A too small r may over-prune true guidewire pixels, whereas a too large r weakens anatomical con-

Table 1: Quantitative comparison on the test images of DSA dataset. The best results are shown in bold.

Method	Dice (%) ↑	HD95 (px) ↓
DARL (ICLR 2023)	76.8	23.25
Medical-SAM-adapter (MIA 2025)	77.8	5.40
H-SAM (CVPR 2024)	83.1	6.89
SAMed (Preprint 2023)	77.4	18.25
FGA-Net (CMIG 2025)	78.7	17.09
RDE-SAM + Refinement (Ours)	**84.2**	**2.36**

Table 2: Ablation study of the proposed components. All settings share the same backbone, implicit prompt token, and LoRA configuration.

Method	LoRA	Experts	Router	Refine	Dice (%) ↑	HD95 (px) ↓
SAM + LoRA	✓				80.2	9.25
+ Experts	✓	✓			82.8	8.32
+ Experts + Router	✓	✓	✓		83.9	2.98
+ Experts + Router + Refinement	✓	✓	✓	✓	**84.2**	**2.36**

Table 3: Generalization of RDE on different backbones.

	UNet	RDE-UNet	SAM	RDE-SAM
Dice (%) ↑	80.7	**82.0**	80.2	**83.9**
HD95 (px) ↓	41.2	**40.0**	9.25	**2.98**

straints. Similarly, a too small $d_{\max}$ fails to repair local gaps, while a too large $d_{\max}$ may introduce overly long connections.

3.4 Generalization Study

To evaluate the architecture-agnostic property of RDE, we instantiate it on a lightweight UNet backbone (RDE-UNet) using the same expert-and-router design. As shown in Table 3, RDE improves both UNet and SAM backbones. On UNet, it improves Dice by 1.3 and reduces HD95 by 1.2, indicating enhanced structural consistency for convolutional architectures. On SAM, the gains are more pronounced, suggesting that role-disentangled specialization and conservative residual fusion are especially effective with high-capacity pretrained encoders. These results show that RDE is not tied to a specific backbone and can serve as a transferable framework for anatomy-constrained segmentation.

4 Conclusion

In this work, we formulate DSA guidewire segmentation as an anatomy-constrained structured prediction problem and propose RDE, a Role-Disentangled Experts framework integrating parameter-efficient domain specialization, role-disentangled decoding, conservative residual routing, and constraint-aware refinement. By jointly segmenting guidewires and vessels, RDE models their anatomical dependency to regularize and correct guidewire predictions. RDE is architecture-agnostic and can be instantiated on different backbones. Experiments on a clinically collected DSA dataset with dual annotations show that RDE-Seg achieves the best performance. A current limitation is the single in-house dataset; future work will expand the cohort, include multi-center clinical scenarios, and evaluate external generalization under broader acquisition protocols.

Acknowledgments. This work was supported by the National Natural Science Foundation of China (No. 62372338) and the Key Research and Development Projects in Hubei Province (No. 2023BCB133).

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Ambrosini, P., Ruijters, D., Niessen, W.J., Moelker, A., van Walsum, T.: Fully automatic and real-time catheter segmentation in x-ray fluoroscopy. In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 577–585 (2017)
2. Barbu, A.: Mslnet and perceptual grouping for guidewire segmentation and localization. Sensors **25**(20), 6426 (2025)
3. Bian, Z.: Multitask segmentation of vessels, catheters, and guidewires in medical imaging using a yolop-unet hybrid architecture. In: International Conference on Computer Vision, Robotics and Automation Engineering (CRAE). pp. 1–7 (2025)
4. Cheng, E., Pang, Y., Zhu, Y., Yu, J., Ling, H.: Curvilinear structure tracking by low rank tensor approximation with model propagation. In: International Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3057–3064 (2014)
5. Cheng, Z., Wei, Q., Zhu, H., Wang, Y., Qu, L., Shao, W., Zhou, Y.: Unleashing the potential of sam for medical adaptation via hierarchical decoding. In: International Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3511–3522 (2024)
6. Hu, C., Xia, T., Ju, S., Li, X.: When sam meets medical images: An investigation of segment anything model (sam) on multi-phase liver tumor segmentation. arXiv preprint arXiv:2304.08506 (2023)
7. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. In: International Conference on Learning Representations (ICLR). pp. 1–26 (2022)
8. Jacobs, R.A., Jordan, M.I., Nowlan, S.J., Hinton, G.E.: Adaptive mixtures of local experts. Neural Computation **3**(1), 79–87 (1991)
9. Kim, B., Oh, Y., Ye, J.C.: Diffusion adversarial representation learning for self-supervised vessel segmentation. In: International Conference on Learning Representations (ICLR). pp. 1–22 (2023)

10. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. In: International Conference on Computer Vision (ICCV). pp. 3992–4003 (2023)
11. Lepikhin, D., Lee, H., Xu, Y., Chen, D., Firat, O., Huang, Y., Krikun, M., Shazeer, N., Chen, Z.: Gshard: Scaling giant models with conditional computation and automatic sharding. In: International Conference on Learning Representations (ICLR). pp. 1–23 (2021)
12. Li, W., Que, D., Omisore, O.M.: A gaussian-based guidewire segmentation and tracking method for teleoperated robotic intravascular interventions. In: International Conference on Robotics and Automation Sciences (ICRAS). pp. 38–43 (2021)
13. Luo, T., Lei, J., Lei, F., Liu, W., He, S., Zhao, J., Liu, K.: Moelora: Contrastive learning guided mixture of experts on parameter-efficient fine-tuning for large language models. arXiv preprint arXiv:2402.12851 (2024)
14. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. Nature Communications **15**, 654 (2024)
15. Pan, S., Liu, Y., Zhao, L., Chen, E.Z., Chen, X., Chen, T., Sun, S.: Label-efficient data augmentation with video diffusion models for guidewire segmentation in cardiac fluoroscopy. In: AAAI Conference on Artificial Intelligence (AAAI). pp. 6308–6316 (2025)
16. Paranjape, J.N., Nair, N.G., Sikder, S., Vedula, S.S., Patel, V.M.: Adaptivesam: Towards efficient tuning of sam for surgical scene segmentation. In: Annual Conference on Medical Image Understanding and Analysis (MIUA). pp. 187–201 (2024)
17. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 234–241 (2015)
18. Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., Dean, J.: Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In: International Conference on Learning Representations (ICLR). pp. 1–19 (2017)
19. Vandini, A., Glocker, B., Hamady, M., Yang, G.Z.: Robust guidewire tracking under large deformations combining segment-like features (seglets). Medical Image Analysis **38**, 150–164 (2017)
20. Vlontzos, A., Mikolajczyk, K.: Deep segmentation and registration in x-ray angiography video. In: British Machine Vision Conference (BMVC). pp. 1–11 (2018)
21. Wang, G., Ye, J., Cheng, J., Li, T., Chen, Z., Cai, J., He, J., Zhuang, B.: Sam-med3d-moe: Towards a non-forgetting segment anything model via mixture of experts for 3d medical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 552–561 (2024)
22. Wang, P., Chen, T., Zhu, Y., Zhang, W., Zhou, S.K., Comaniciu, D.: Robust guidewire tracking in fluoroscopy. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 691–698 (2009)
23. Wu, J., Ji, W., Liu, Y., Fu, H., Xu, M., Xu, Y., Jin, Y.: Medical sam adapter: Adapting segment anything model for medical image segmentation. Medical Image Analysis **102**, 103547 (2025)
24. Yao, T., Pei, Z., Li, Y., Yuan, Y., Qi, P.: Real-time guidewire tip tracking using a siamese network for image-guided endovascular procedures. Advanced Intelligent Systems **7**(10), 1–12 (2025)
25. Zhang, B., Bui, M., Wang, C., Bourier, F., Schunkert, H., Navab, N.: Real-time guidewire tracking and segmentation in intraoperative x-ray. In: Medical Imaging 2022: Image-Guided Procedures, Robotic Interventions, and Modeling. pp. 1–8 (2022)

26. Zhang, G., Liu, D., Ji, Z., Xie, B., Chen, Q.: Lightweight wavelet convolutional network for guidewire segmentation. Computerized Medical Imaging and Graphics **124**, 102618 (2025)
27. Zhang, G., Wong, H.C., Wang, C., Zhu, J., Lu, L., Teng, G.: A temporary transformer network for guide-wire segmentation. In: International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI). pp. 1–5 (2021)
28. Zhang, G., Wong, H.C., Zhu, J., An, T., Wang, C.: Jigsaw training-based background reverse attention transformer network for guidewire segmentation. International Journal of Computer Assisted Radiology and Surgery **18**, 653–661 (2023)
29. Zhang, K., Liu, D.: Customized segment anything model for medical image segmentation. arXiv preprint arXiv:2304.13785 (2023)
30. Zhao, D., Liu, J., Geng, P., Yang, J., Zhang, Z., Zhang, Y.: Mid-net: Rethinking efficient network architectures for small-sample vascular segmentation. Information Fusion **115**, 102777 (2025)
31. Zhu, J., Wang, C., Zhang, G., Teng, S.A., Lu, L., Teng, G.: A multi-stage guidewire tip tracking framework for cardiovascular robotic interventions. In: International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI). pp. 1–6 (2021)