Breaking Annotation Dependency in Coronary Stenosis Segmentation via Physics-Driven Sim2Real Learning

Baochang Zhang[1,2,3], Fatia Kusuma Dewi[1], Shuting Liu[1], Amir Yousefi[1,4], Heribert Schunkert[2,5], Reza Ghotbi[4], and Nassir Navab[1,3]

[1] Computer Aided Medical Procedures, Technical University of Munich, Munich, Germany
baochang.zhang@tum.de & shuting.liu@tum.de
[2] German Heart Center Munich, Munich, Germany
[3] Munich Center for Machine Learning, Munich, Germany
[4] HELIOS Hospital west of Munich, Munich, Germany
[5] German Centre for Cardiovascular Research, Munich Heart Alliance, Munich, Germany

Abstract. Accurate coronary stenosis segmentation in X-ray angiography (XA) remains fundamentally limited by reliance on scarce and biased manual lesion annotations. Clinical datasets, although realistic, provide only incomplete and skewed sampling of pathological variations, constraining generalization. We address this challenge by reframing stenosis segmentation as a controllable synthetic learning problem and propose a physics-driven sim-to-real framework that eliminates dependence on real lesion labels. We generate a scalable synthetic dataset from patient-specific coronary models reconstructed from CTA, where stenosis variations with diverse severities and morphologies are deliberately designed in 3D space. A physics-based projection process produces geometry-preserving synthetic XA images with precise image-annotation alignment, enabling balanced and distribution-aware supervision beyond biased clinical sampling. A segmentation network trained on this synthetic data is subsequently adapted to real XA through uncertainty-guided self-supervised refinement. Predictive uncertainty is estimated online to select reliable pseudo-labels and drive confidence-aware consistency learning, promoting robust adaptation in ambiguous regions. On the ARCADE benchmark, our method surpasses existing self-supervised and few-shot approaches and approaches fully supervised performance without using any real stenosis annotations. These results demonstrate that physics-driven sim-to-real learning offers a principled pathway toward annotation-independent coronary lesion analysis. The code is publicly available at github.com/StenSim2realSeg.

Keywords: Stenosis Segmentation · Physics-driven Simulation · Sim-to-real Learning · Uncertainty-guided Adaptation · X-ray Angiography.

Co-corresponding authors: Baochang Zhang and Shuting Liu.

1 Introduction

Coronary artery disease remains a leading cause of mortality worldwide, and accurate assessment of coronary stenosis from X-ray angiography (XA) is critical for diagnosis and interventional planning. A rapid reduction in vessel diameter is a primary indicator of stenosis severity [2]. However, XA interpretation is challenged by vessel overlap, uneven contrast, and projection ambiguity. Manual assessment introduces inter-observer variability and limited reproducibility [5], motivating the development of automated and objective analysis systems.

Recent advances in deep learning have substantially improved automatic stenosis detection and segmentation [9,7,3]. Nevertheless, these supervised methods fundamentally rely on large-scale, high-quality lesion annotations. In practice, stenosis labeling is costly, subjective, and anatomically inconsistent. Public datasets [11] provide valuable benchmarks but reflect limited pathological diversity and protocol-specific annotation biases, which may constrain model generalization. Efforts to reduce annotation dependence include few-shot and self-supervised approaches [6,12,1,8]. Although promising for vessel segmentation, their effectiveness for subtle and highly localized stenotic patterns remains limited. Sim-to-real learning has emerged as an alternative direction. Domain adaptation techniques [15,10] and adaptive representation learning frameworks for XA [16] demonstrate that synthetic data can facilitate cross-domain generalization. However, existing approaches largely focus on vessel structures rather than controlled pathology modeling and do not explicitly address annotation independence at the lesion level.

In this work, we argue that annotation dependency represents a key bottleneck in coronary stenosis segmentation. We reformulate stenosis segmentation as a controllable synthetic learning problem. By introducing systematically designed stenosis variations into patient-specific coronary models reconstructed from CTA and projecting them through a physics-based digitally reconstructed radiograph (DRR) process, we generate scalable synthetic XA data with geometry-consistent supervision. A segmentation network trained on this data is subsequently adapted to real XA using uncertainty-guided self-supervised refinement, eliminating the need for manual lesion annotations. Our contributions are threefold: (1) We reframe coronary stenosis segmentation as a synthetic, controllable learning paradigm to break annotation dependency. (2) We develop a physics-driven sim-to-real framework that enables balanced modeling of pathological variations with precise image-annotation alignment. (3) We introduce an uncertainty-guided adaptation strategy that achieves performance comparable to fully supervised methods without real stenosis labels.

2 Method

As illustrated in Fig. 1, our framework mainly consists of two parts: controllable physics-driven stenosis modeling and uncertainty-guided self-supervised refinement in the real domain.

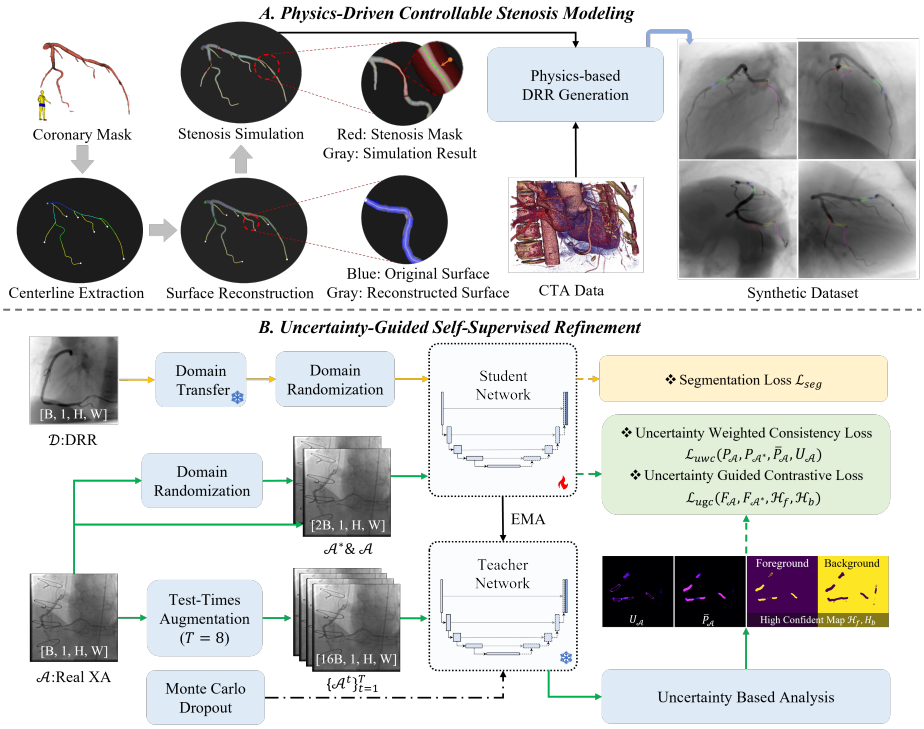

Fig. 1. Overview of the proposed framework. (A) Controllable stenosis modeling and DRR generation with precise masks. (B) Uncertainty-guided self-supervised refinement enables annotation-free stenosis segmentation on real XA.

2.1 Physics-Driven Synthetic Controllable Stenosis Modeling

Given a coronary lumen segmentation from CTA, vessel centerlines are extracted using VMTK [4]. Each branch is represented as a radius-parameterized curve $\mathcal{C} = \{\mathbf{c}(s), R(s)\}$, where $\mathbf{c}(s) \in \mathbb{R}^3$ denotes the centerline coordinate at arc-length s, and $R(s)$ the corresponding lumen radius.

Anatomically Constrained Baseline Geometry. The raw radius profile $R(s)$ may exhibit segmentation-induced oscillations that violate physiological tapering. To enforce anatomical plausibility, we construct a corrected radius field $\tilde{R}(s)$ under a proximal-to-distal tapering constraint, $d\tilde{R}(s)/ds \leq 0$. Specifically, representative control points are selected along each branch, followed by monotonic interpolation and spline-based smoothing to obtain a smooth and continuous radius profile. The refined representation $\{\mathbf{c}(s), \tilde{R}(s)\}$ is subsequently converted into a 3D baseline lumen surface $\mathcal{M} \subset \mathbb{R}^3$ via implicit sphere union modeling and isosurface extraction.

Controllable Stenosis Parameterization. For each eligible branch ($\tilde{R}(s) > 1\,\text{mm}$), a stenosis interval $\mathcal{I} = [s_0, s_1]$, its length $|\mathcal{I}|$, and severity $\alpha \in [\alpha_{\min}, \alpha_{\max}]$ are sampled within clinically meaningful ranges. A smooth radius reduction pro-

file $g(s) \in [0, \alpha]$ is generated using Gaussian, Bézier, or spline-based functions to model diverse longitudinal narrowing patterns along the vessel centerline. Let $\mathcal{M}_{\mathcal{I}} = \{\mathbf{v} \in \mathcal{M} \mid s^\star(\mathbf{v}) \in \mathcal{I}\}$ denote the subset of surface vertices whose closest centerline coordinate $s^\star(\mathbf{v})$ lies within $\mathcal{I}$. Only this localized patch $\mathcal{M}_{\mathcal{I}}$ is modified, while the remaining surface remains unchanged. To unify concentric and eccentric lesions, a symmetry parameter $\beta \in [0, 1]$ is introduced. For a vertex $\mathbf{v} \in \mathcal{M}_{\mathcal{I}}$, the deformed position is defined as

$$\mathbf{v}' = \mathbf{c}(s^\star) + \beta\, \mathbf{d}(s^\star) + (1 - g(s^\star)) \left[\mathbf{v} - \left(\mathbf{c}(s^\star) + \beta\, \mathbf{d}(s^\star)\right)\right], \tag{1}$$

where $\mathbf{d}(s)$ is a smooth displacement field constrained to the local normal plane and vanishing at s_0 and s_1. When $\beta = 0$, the deformation reduces to axis-symmetric contraction; for $\beta > 0$, the contraction axis is locally shifted, producing eccentric narrowing. This unified formulation enables explicit control over lesion severity, morphology, symmetry, spatial distribution, and lesion multiplicity while preserving anatomical smoothness and surface topology.

Physics-Consistent Image and Annotation Generation. Following localized deformation, the modified mesh $\mathcal{M}'$ represents the complete coronary anatomy with simulated stenoses, while the pre-deformation patch $\mathcal{M}_{\mathcal{I}}$ defines the anatomical territory of the lesion. Let $\mathcal{V}$ denote the original CTA volume. A physics-based DRR operator $\mathcal{P}_\theta : \mathbb{R}^3 \to \mathbb{R}^2$ is applied to generate the DRR image I_{drr} and its corresponding stenosis mask M_{sten} [14]:

$$I_{\text{drr}} = \mathcal{P}_\theta(\mathcal{V}, \mathcal{M}'), \quad M_{\text{sten}} = \mathcal{P}_\theta(\mathcal{M}_{\mathcal{I}}), \tag{2}$$

where θ denotes the imaging geometry parameters. Thus, the generated 2D mask is geometrically consistent with the simulated image.

2.2 Synthetic-Domain Representation Learning

Domain Transfer and Domain Randomization To reduce the sim-to-real gap, we employ a pretrained CUT model [10,16] trained on synthetic DRRs and unlabeled XA images from the public XCAD dataset [8], independent of ARCADE, to transfer DRRs into XA style while preserving geometry and lesion structure.

To enhance robustness under real-world variability, we apply domain randomization to the synthetic XA data. In addition to general intensity and noise perturbations, we introduce X-ray-specific augmentations to better emulate clinical conditions. Specifically, random tool embedding simulates interventional devices such as catheters and wires using Bézier curve generation, and stochastic collimation simulates X-ray beam restriction by randomly reducing the effective field of view using a central rectangular aperture (80%–95% of the original size).

Supervised Pre-training. Using the resulting synthetic XA data, we train a U-Net [13] segmentation network under full supervision. The segmentation objective $\mathcal{L}_{seg}$ consists of a standard combination of Dice and binary cross-entropy losses. Pre-training on large-scale, geometry-consistent synthetic data enables the model to learn anatomically grounded representations with controlled stenosis morphology and severity variations before exposure to real XA.

2.3 Uncertainty-Guided Self-Supervised Refinement

After supervised pre-training on synthetic XA, we perform real-domain adaptation via a teacher-student paradigm.

Stabilized Refinement with Synthetic Anchors. During real-domain refinement, each mini-batch jointly samples real XA and synthetic XA images. Synthetic samples continue to be optimized using the supervised segmentation objective $\mathcal{L}_{seg}$, serving as an anchor that stabilizes optimization and prevents representation drift caused by noisy pseudo supervision in the real domain.

Uncertainty-aware Soft Pseudo Supervision. For a real XA image $\mathcal{A}$, the teacher network estimates predictive uncertainty using Monte Carlo dropout and test-time augmentation. Given T stochastic predictions $\{P_{\mathcal{A}}^t\}_{t=1}^T$, we compute the predictive mean probability map $\bar{P}_{\mathcal{A}}$ and pixel-wise epistemic uncertainty $U_{\mathcal{A}}$ as

$$\bar{P}_{\mathcal{A}} = \tfrac{1}{T}\sum_{t=1}^{T} P_{\mathcal{A}}^t, \qquad U_{\mathcal{A}} = H(\bar{P}_{\mathcal{A}}) - \tfrac{1}{T}\sum_{t=1}^{T} H(P_{\mathcal{A}}^t), \tag{3}$$

where $H(\cdot)$ denotes pixel-wise entropy. $U_{\mathcal{A}}$ is normalized to $[0,1]$ per image for stable weighting across samples. Rather than constructing hard pseudo masks, we directly use $\bar{P}_{\mathcal{A}}$ as a soft pseudo target and use $U_{\mathcal{A}}$ to modulate student optimization in a reliability-aware manner.

Uncertainty Weighted Consistency. To enforce prediction stability under appearance perturbations, the student processes both the original image $\mathcal{A}$ and a perturbed case $\mathcal{A}^*$, producing predictions $P_{\mathcal{A}}$ and $P_{\mathcal{A}^*}$. The teacher mean prediction $\bar{P}_{\mathcal{A}}$ serves as the supervision target. We define the uncertainty-weighted consistency loss as

$$\mathcal{L}_{uwc} = \sum_{v\in\{\mathcal{A},\mathcal{A}^*\}} \left\| (1-U_{\mathcal{A}}) \circ \left(P_v - \bar{P}_{\mathcal{A}}\right) \right\|_2^2, \tag{4}$$

where $\circ$ denotes element-wise multiplication.

Uncertainty Guided Contrastive Learning. To enhance feature separability between stenosis and background under domain shift, we introduce a pixel-wise contrastive objective. Let $F_{\mathcal{A}}$ and $F_{\mathcal{A}^*}$ denote feature maps of $\mathcal{A}$ and $\mathcal{A}^*$ extracted from the last decoder layer of student network. High-confidence foreground $\mathcal{H}_f$ and background $\mathcal{H}_b$ regions are defined as

$$\mathcal{H}_f = \{i \mid \bar{P}_{\mathcal{A}}(i) > \tau_{fg} \wedge U_{\mathcal{A}}(i) < \tau_u\}, \tag{5}$$

$$\mathcal{H}_b = \{i \mid \bar{P}_{\mathcal{A}}(i) < \tau_{bg} \wedge U_{\mathcal{A}}(i) < \tau_u\}. \tag{6}$$

Queries $\{q_i\}_{i=1}^n$ are sampled from $F_{\mathcal{A}}$ at locations $i \in \mathcal{H}_f$, positives k_i^+ are taken from $F_{\mathcal{A}^*}$ at corresponding spatial locations, and negatives $\{k_j^-\}_{j=1}^w$ are sampled from background locations in $\mathcal{H}_b$ on the feature maps $(F_{\mathcal{A}}, F_{\mathcal{A}^*})$. The contrastive loss is defined as

$$\mathcal{L}_{ugc} = -\frac{1}{n}\sum_{i=1}^{n} \log \frac{\exp(q_i \cdot k_i^+/\tau)}{\exp(q_i \cdot k_i^+/\tau) + \sum_{j=1}^{w} \exp(q_i \cdot k_j^-/\tau)}, \tag{7}$$

where τ is a temperature parameter. Restricting sampling to low-uncertainty regions ensures semantic reliability and reduces noisy feature alignment.

The overall refinement objective for real-domain adaptation is

$$\mathcal{L}_{real} = \lambda_1 \mathcal{L}_{uwc} + \lambda_2 \mathcal{L}_{ugc}, \tag{8}$$

while synthetic samples continue to be optimized using $\mathcal{L}_{seg}$. This uncertainty-guided refinement selectively enforces reliable supervision and improves domain robustness without relying on hard pseudo labels.

2.4 Implementation Details

Stenosis length is sampled between 5 mm and 75% of the branch length, and radius reduction severity between 30% and 100%. Up to 5 stenoses are randomly simulated per coronary tree. DRR images are rendered under 11 standard angiographic projections with randomized pose perturbations. The final synthetic dataset contains 52,800 image-mask pairs at a resolution of 512×512.

For supervised pre-training on synthetic XA, the network is trained for 200 epochs with batch size 16 using Adam. The initial learning rate is 10^{-3} and decayed by a factor of 0.85 every 10 epochs. For real-domain refinement, the student network is optimized for 20 epochs with batch size 8 using Adam, with an initial learning rate of 10^{-6} decayed by 0.85 every 5 epochs. The teacher network is updated as an exponential moving average (EMA) of the student parameters with momentum 0.99.

Unless otherwise specified, we set $\tau_{fg} = 0.8$, $\tau_{bg} = 0.2$, and $\tau_u = 0.2$. In contrastive learning, the number of sampled queries is $n = 256$, the number of negatives is $w = 2048$, and the temperature parameter is $\tau = 0.1$. The loss weights are set to $\lambda_1 = \lambda_2 = 1.0$.

3 Datasets and Experimental Results

3.1 Dataset and Evaluation Metrics

We evaluate our model on the public ARCADE dataset [11], which contains 512×512 XA images with expert annotations, split into 1000 training and 300 testing samples. In ARCADE, stenosis is defined as a coronary lesion with ≥50% luminal narrowing in vessels of diameter ≥1.5 mm, and mild stenosis below this threshold is not included in the original annotations. Since our simulated stenoses span 30%-100% severity, we re-annotated the testing set using 3D Slicer and reviewed it with experienced surgeons for more comprehensive evaluation. Both the original and revised masks are used for evaluation.

We assess performance using pixel-level and instance-level metrics. Pixel-wise evaluation includes Dice coefficient (Dic.) and Recall (Rec.). Since annotations are provided at the lesion level, instance-wise evaluation is performed via lesion matching: a ground-truth stenosis is counted as a true positive (TP) if at least 50% of its area is covered by a prediction; otherwise, it is treated as a false

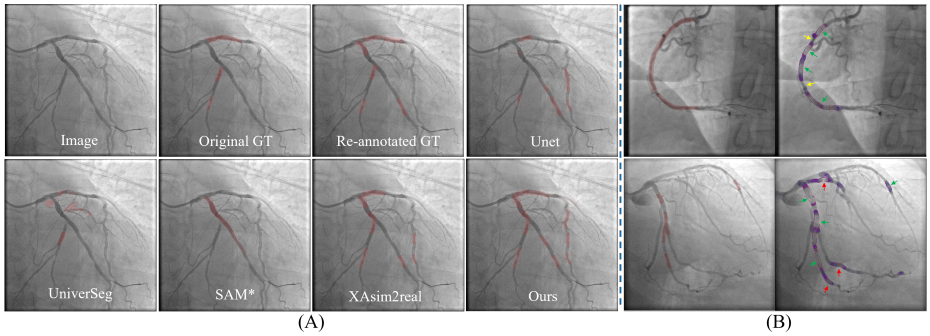

Fig. 2. (A) Qualitative comparison of stenosis segmentation on real XA images. Red regions denote predicted stenosis areas. (B) Our uncertainty map visualization on two real XA test cases. Each case shows (left) annotation and (right) XA image with uncertainty overlay. Green arrows indicate correctly detected lesions, red arrows denote additional detected constrictions, and yellow arrows indicate missed annotated regions.

negative (FN). A predicted region is regarded as a false positive (FP) if it overlaps less than 50% with all ground-truth instances, with one-to-one matching enforced. Based on TP, FP, and FN, we report Instance-wise F1 (Ins-F1) and Instance-wise Recall (Ins-Rec.)

3.2 Quantitative Results and Ablation Study

Table 1 reports results on ARCADE. Without using any real stenosis annotations, our method approaches the fully supervised U-Net [13] trained on all 1000 labeled samples in Dice while substantially improving lesion sensitivity, with Ins-Rec. increasing from 0.658 to 0.886 under the original protocol and from 0.540 to 0.843 under the re-annotated protocol, indicating markedly fewer missed stenoses. Importantly, both our method and XAsim2real [16] are trained on the

Table 1. Performance comparison under the original and re-annotated ground-truth masks. "Anno." denotes the number of annotated XA samples used for training.

Mask	Methods	Anno.	Dic.	Rec.	Ins-F1	Ins-Rec.
Original	Unet [13]	1000	0.431±0.281	0.434±0.314	**0.530**±0.394	0.658±0.451
	Universeg [1]	32	0.128±0.144	0.107±0.133	0.048±0.114	0.153±0.333
	SAM* [12]	32	0.238±0.235	0.251±0.271	0.129±0.203	0.398±0.467
	XAsim2real [16]	0	0.424±0.236	0.432±0.274	0.473±0.357	0.694±0.437
	Ours	0	**0.442**±0.192	**0.592**±0.246	0.434±0.255	**0.886**±0.302
Re-annotated	Unet [13]	1000	0.415±0.249	0.397±0.280	0.533±0.346	0.540±0.376
	Universeg [1]	32	0.118±0.128	0.095±0.117	0.066±0.120	0.133±0.243
	SAM* [12]	32	0.233±0.213	0.237±0.241	0.157±0.200	0.338±0.376
	XAsim2real [16]	0	0.486±0.237	0.457±0.267	**0.578**±0.329	0.645±0.367
	Ours	0	**0.515**±0.189	**0.635**±0.245	0.570±0.248	**0.843**±0.271

Table 2. Ablation study under the re-annotated protocol.

Methods	Dic.	Rec.	Ins-F1	Ins-Rec.
Ours	**0.515**±0.189	**0.635**±0.245	0.570±0.248	**0.843**±0.271
Ours w/o Ref.	0.497±0.233	0.492±0.273	**0.589**±0.321	0.671±0.361
Ours w/o $\mathcal{L}_{uwc}$	0.508±0.192	0.609±0.249	0.562±0.258	0.822±0.295
Ours w/o $\mathcal{L}_{ugc}$	0.497±0.214	0.519±0.265	0.585±0.294	0.723±0.336

same physics-driven synthetic dataset, so the performance gain stems from the proposed uncertainty-guided refinement rather than data scale; compared with XAsim2real, Recall and Ins-Rec. are consistently higher, demonstrating more effective sim-to-real transfer. Although Dice and Ins-F1 remain comparable across methods, the slightly lower Ins-F1 under the original protocol is attributable to annotation definition (only $\geq$ 50% stenoses labeled), whereas our modeling spans 30%–100%, leading to detection of mild stenoses that are unlabeled. Under the re-annotated protocol with more complete lesion coverage, our method achieves the best overall performance, suggesting that physics-driven sim-to-real learning improves lesion coverage without relying on manual stenosis annotations. Fig. 2 shows that our method captures annotated stenoses more completely and highlights additional narrowing-like regions (red arrows) beyond the original masks, consistent with the improved Recall. The uncertainty maps exhibit higher responses around ambiguous or borderline narrowing regions, while remaining confined to the vessel lumen, indicating anatomically meaningful confidence estimation.

Table 2 analyzes the contribution of each component under the re-annotated protocol. Removing real-domain refinement (w/o Ref.) leads to a substantial drop in Rec. (0.635→0.492) and Ins-Rec. (0.843→0.671), confirming that uncertainty guided adaptation is essential for transferring synthetic pathology to real XA. Although Ins-F1 slightly increases in this setting, this is mainly due to reduced lesion sensitivity, which decreases false positives at the cost of more missed detections. Removing $\mathcal{L}_{uwc}$ yields modest degradation, indicating its role in stabilizing pseudo supervision, while removing $\mathcal{L}_{ugc}$ causes broader performance drops across Dice, Recall, and Ins-Rec., highlighting the importance of uncertainty-guided feature separation for both pixel-level accuracy and lesion-level discrimination. Together, these results demonstrate that consistency and contrastive objectives provide complementary benefits for annotation-free sim-to-real learning.

4 Conclusion

In this work, we reframed coronary stenosis segmentation as a controllable synthetic learning paradigm to break dependency on manual lesion annotations. By coupling anatomically constrained stenosis modeling with physics-driven DRR projection, we enabled geometry-consistent supervision beyond biased clinical

sampling. Through uncertainty-guided refinement, the learned representations are effectively transferred to real XA without requiring any lesion labels. Our results show that physics-driven sim-to-real learning can approach fully supervised performance while improving lesion sensitivity. More broadly, this study demonstrates that controllable pathology simulation offers a principled and scalable alternative to manual annotation for coronary stenosis analysis.

Acknowledgments. The project was supported by the Bavarian State Ministry of Science and Arts within the framework of the "Digitaler Herz-OP" project under the grant number 1530/891 02 and the China Scholarship Council (File No.202004910390).

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

References

1. Butoi, V.I., Ortiz, J.J.G., Ma, T., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Universeg: Universal medical image segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 21438–21451 (2023)
2. Fleming, R.M., Kirkeeide, R.L., Smalling, R.W., Gould, K.L., Stuart, Y.: Patterns in visual interpretation of coronary arteriograms as detected by quantitative coronary arteriography. Journal of the American College of Cardiology **18**(4), 945–951 (1991)
3. 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)
4. Izzo, R., Steinman, D., Manini, S., Antiga, L.: The vascular modeling toolkit: a python library for the analysis of tubular structures in medical images. Journal of Open Source Software **3**(25), 745 (2018)
5. Janssen, J.P., Rares, A., Tuinenburg, J.C., Koning, G., Lansky, A.J., Reiber, J.H.: New approaches for the assessment of vessel sizes in quantitative (cardio-) vascular x-ray analysis. The international journal of cardiovascular imaging **26**, 259–271 (2010)
6. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 4015–4026 (2023)
7. Lin, H., Liu, T., Katsaggelos, A., Kline, A.: Stenunet: automatic stenosis detection from x-ray coronary angiography. arXiv preprint arXiv:2310.14961 (2023)
8. Ma, Y., Hua, Y., Deng, H., Song, T., Wang, H., Xue, Z., Cao, H., Ma, R., Guan, H.: Self-supervised vessel segmentation via adversarial learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7536–7545 (2021)
9. Pang, K., Ai, D., Fang, H., Fan, J., Song, H., Yang, J.: Stenosis-detnet: Sequence consistency-based stenosis detection for x-ray coronary angiography. Computerized Medical Imaging and Graphics **89**, 101900 (2021)
10. Park, T., Efros, A.A., Zhang, R., Zhu, J.Y.: Contrastive learning for unpaired image-to-image translation. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX 16. pp. 319–345. Springer (2020)

11. Popov, M., Amanturdieva, A., Zhaksylyk, N., Alkanov, A., Saniyazbekov, A., Aimyshev, T., Ismailov, E., Bulegenov, A., Kuzhukeyev, A., Kulanbayeva, A., et al.: Dataset for automatic region-based coronary artery disease diagnostics using x-ray angiography images. Scientific data **11**(1), 20 (2024)
12. Qiu, Z., Hu, Y., Li, H., Liu, J.: Learnable ophthalmology sam. arXiv preprint arXiv:2304.13425 (2023)
13. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18. pp. 234–241. Springer (2015)
14. Unberath, M., Zaech, J.N., Lee, S.C., Bier, B., Fotouhi, J., Armand, M., Navab, N.: Deepdrr–a catalyst for machine learning in fluoroscopy-guided procedures. In: International conference on medical image computing and computer-assisted intervention. pp. 98–106. Springer (2018)
15. Yang, Y., Soatto, S.: Fda: Fourier domain adaptation for semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4085–4095 (2020)
16. Zhang, B., Zhang, Z., Liu, S., Faghihroohi, S., Schunkert, H., Navab, N.: Xasim2real: Adaptive representation learning for vessel segmentation in x-ray angiography. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 747–756 (2024)