Diffusion-Based User-Guided Data Augmentation for Coronary Stenosis Detection (DiGDA)

Sumin Seo[1], In Kyu Lee[1,2], Hyun-Woo Kim[1], Jaesik Min[1], Chung-Hwan Jung[1]
1 Medipixel, Inc. 2 University of California San Diego

Our product MPXA

Stay tuned on our GitHub!

Goal

We generate synthetic angiograms with severe lesions to improve the detection and severity classification of coronary stenosis.

Introduction

Coronary stenosis is a critical risk factor for coronary artery disease, one of the leading causes of global mortality, and is characterized by arterial narrowing due to plaque accumulation.

Challenges in Stenosis Detection

- **Data scarcity**: While deep learning has achieved remarkable progress, its application to coronary angiography remains limited by data scarcity and high labeling cost.
- **Class imbalance**: Coronary angiography datasets are inherently imbalanced and generation-based synthesis often introduces artifacts in non-lesion regions.

Contribution: Our method enables balanced lesion synthesis while preserving vascular structure and keeping the severity of unaffected lesions unchanged.

Methods

Percentage Diameter Stenosis (%DS) defines how severe the lesion is. It depends on the *diameter* such as minimum lumen diameter (MLD) and reference diameter D_{ref}.

$$\%DS = 1 - (MLD/D_{ref}) \times 100$$

- (1,2) We employ our proprietary QCA software to extract lesion information from angiograms and generate the textual prompts.
- (3) Our controllable generation model (multi-input ControlNet [1]) is trained with three inputs: a vessel segmentation mask (subject to modification during inference), a lesion-masked image, and a text prompt describing vessel type and lesion size.
- (4) To enforce user-defined %DS values, we refine the generation process by modifying the vessel mask, specifically by shifting MLD points and applying Gaussian weighting along the vessel contours.
- (5) At inference, the trained model synthesizes angiogram images using the modified masks as guidance.
- (6) Detection algorithm YOLO [2] is then applied to automatically detect lesions and classify their corresponding severities.

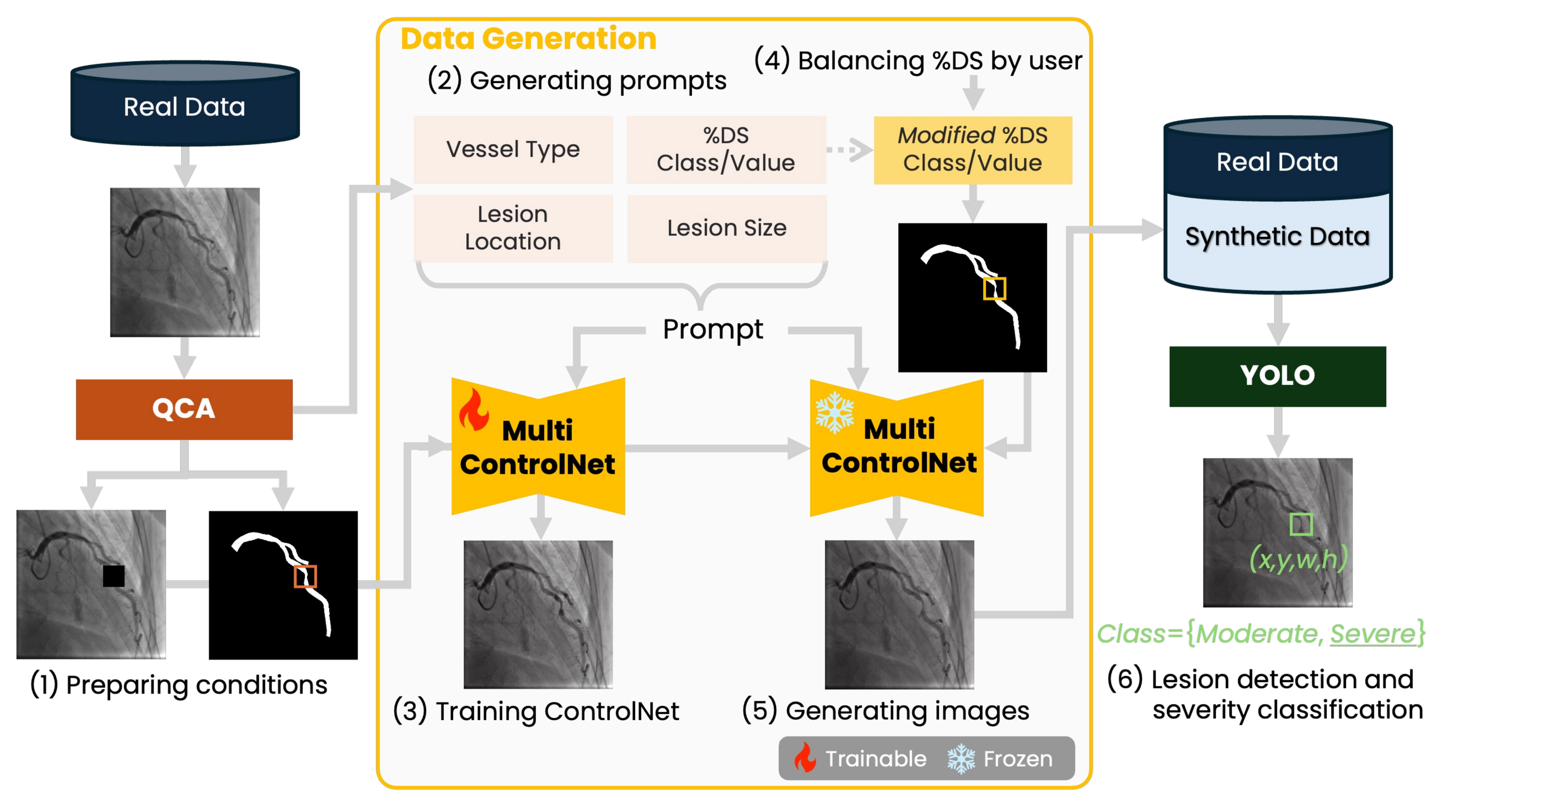

References

[1] Zhang et al. Adding Conditional Control to Text-to-Image Diffusion Models. ICCV, 2023.
[2] Khanam & Hussain. YOLOv11: An Overview of the Key Architectural Enhancements. arXiv, 2024.
[3] Popov et al. ARCADE Dataset for Automatic Region-Based Coronary Artery Disease Diagnostics Using X-ray Angiography Images. Scientific Data, 2024.

Experiments

We compare different synthetic data sizes and class balance settings.

- We observe imbalance in the %DS distribution.
- We balance the ratio of moderate (5-70%) and severe (>70%) by gen-erating images with severe lesions.
- Adding balanced synthetic data gradually improves lesion detection performance up to x4.

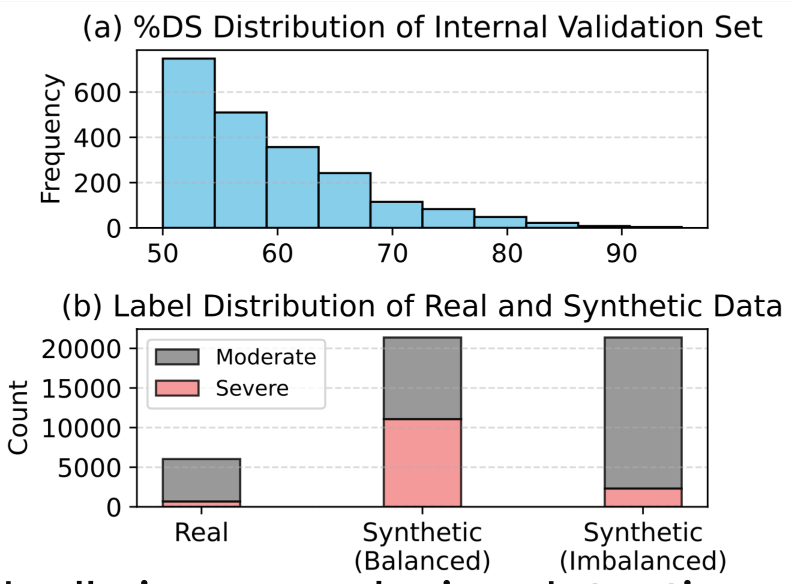

Synthetic data ratio	Internal dataset F1	Internal dataset mAP50	(Moderate	Severe)	ARCADE F1	ARCADE mAP50	(Moderate	Severe)
x0	0.650	0.688	0.731	0.646	0.500	0.464	0.436	0.492
x1	0.656	0.699	0.749	0.649	0.519	0.484	0.440	0.527
x2	0.658	0.710	0.767	0.654	**0.524**	**0.501**	**0.448**	**0.555**
x4	**0.670**	**0.717**	**0.771**	**0.663**	0.513	0.492	0.438	0.546

- The class-balanced synthetic dataset achieved the best detection performance, showing that balancing classes by adding synthetic data substantially improves the results.

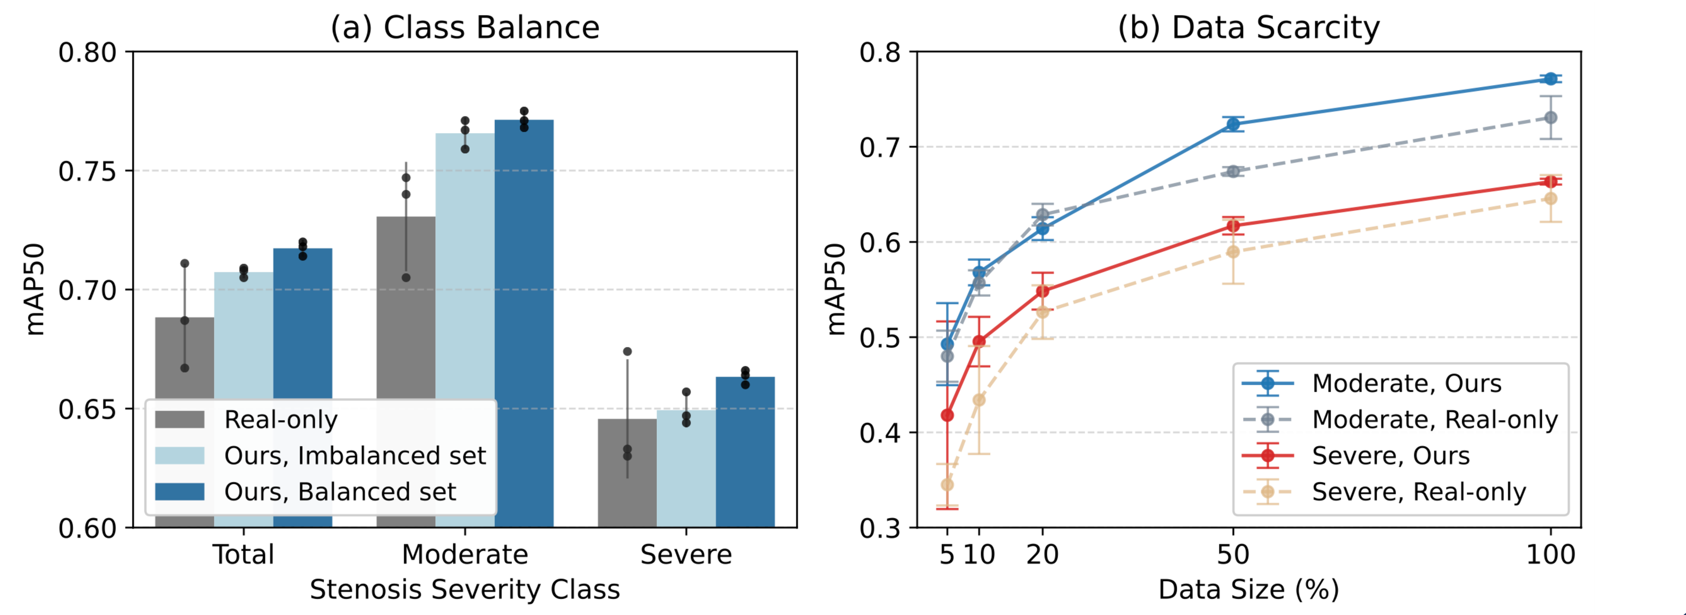

Qualitative Results

Lesion severities in the original images (a,c,e) are modified to (b,d,f) by the mask generation algorithm (g). The synthesized angiograms show lesions *narrowed* according to the specified %DS values, while surrounding vessels are preserved by generative inpainting.

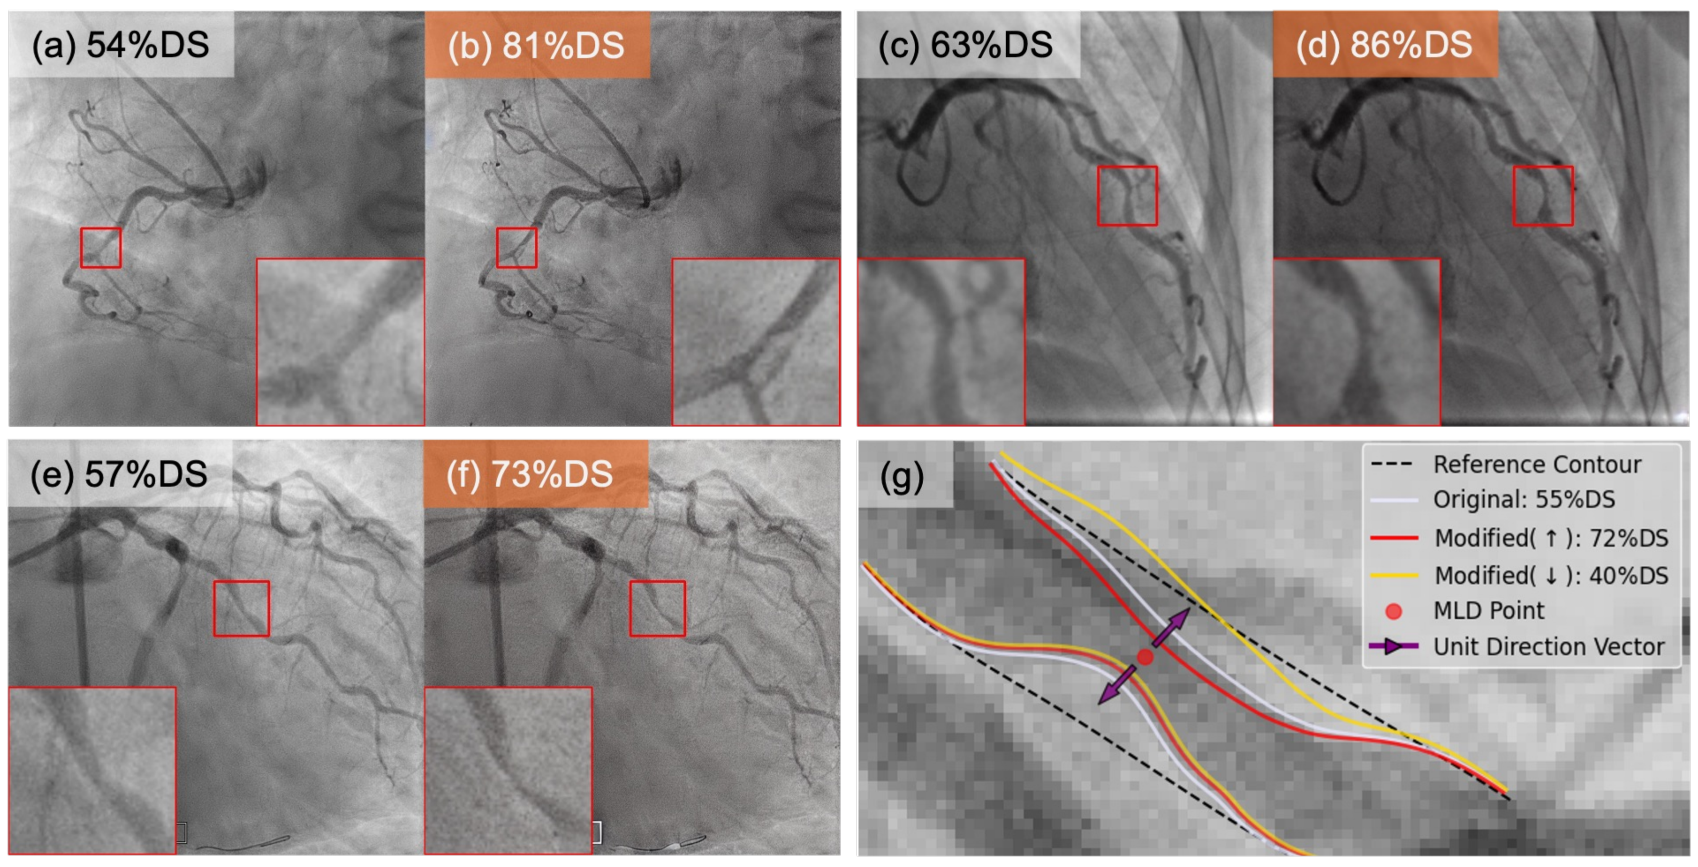

Relabeled Dataset Release

For better model development and robust evaluation, our experienced in-house clinicians re-labeled the public dataset (ARCADE [3]).
We release these new labels at *github.com/medipixel/DiGDA*.

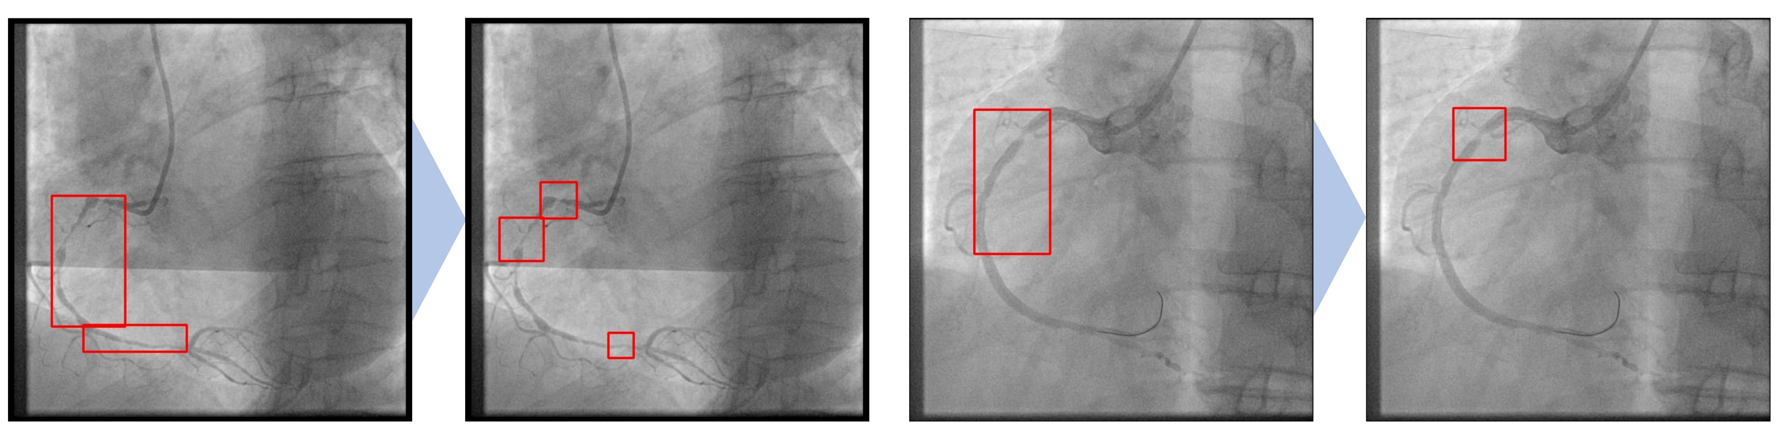

For more information, please contact sumin.seo@medipixel.io or friend me via LinkedIn

Paper #3616