Diffusion-Based User-Guided Data Augmentation for Coronary Stenosis Detection

Sumin Seo[1 [0000−0001−8703−0322]⋆], In Kyu Lee[1,2 [0000−0001−5554−808X]], Hyun-Woo Kim[1 [0009−0003−2740−0397]], Jaesik Min[1 [0000−0002−0007−0637]], and Chung-Hwan Jung[1]

[1] Medipixel, Inc., Republic of Korea
{sumin.seo,hyunwoo.kim,jaesik.min,danny.jung}@medipixel.io
[2] University of California San Diego, USA
kl002@ucsd.edu

Abstract. Coronary stenosis is a major risk factor for ischemic heart events leading to increased mortality, and medical treatments for this condition require meticulous, labor-intensive analysis. Coronary angiography provides critical visual cues for assessing stenosis, supporting clinicians in making informed decisions for diagnosis and treatment. Recent advances in deep learning have shown great potential for automated localization and severity measurement of stenosis. In real-world scenarios, however, the success of these competent approaches is often hindered by challenges such as limited labeled data and class imbalance. In this study, we propose a novel data augmentation approach that uses an inpainting method based on a diffusion model to generate realistic lesions, allowing user-guided control of severity. Extensive evaluations show that incorporating synthetic data during training enhances lesion detection and severity classification performance on both a large-scale in-house dataset and a public coronary angiography dataset. Furthermore, our approach maintains high detection and classification performance even when trained with limited data, highlighting its clinical importance in improving the assessment of stenosis severity and optimizing data utilization for more reliable decision support.

Keywords: Stenosis detection · Coronary angiography · Diffusion models.

1 Introduction

Coronary artery disease (CAD), a leading cause of global mortality [1,7], is characterized by narrowed artery often caused by buildup of plaque, reducing blood flow to the heart. Accurate quantification of stenosis severity with coronary angiography (CAG) is essential for guiding clinical decisions and planning interventions. Percentage diameter stenosis (%DS) is the principal metric for quantifying coronary narrowing: values above 50% indicate clinical significance, and

⋆ Corresponding author: sumin.seo@medipixel.io

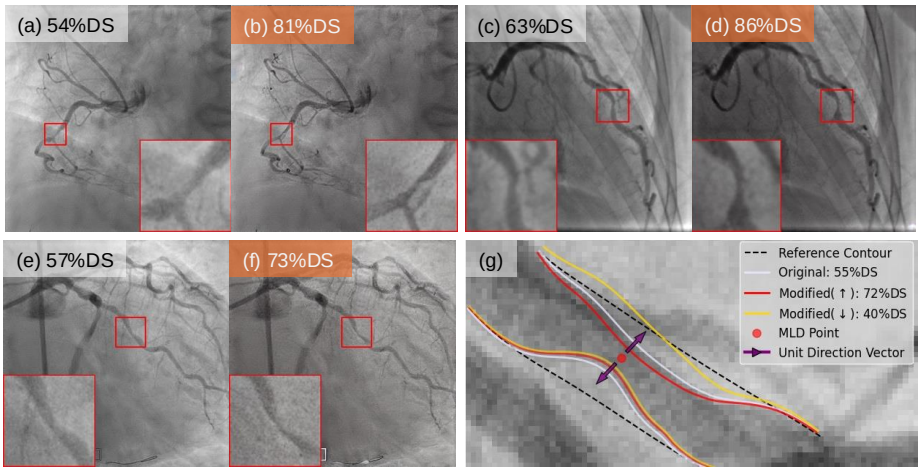

Fig. 1: **Examples of coronary angiograms with various levels of stenosis.** (a,c,e) are the original angiograms, and (b,d,f) are the synthetic angiograms generated with specific %DS values applied to the vessel segmentation masks. (g) describes how the algorithm generates a lesion mask when numeric value of %DS is given. Minimum lumen diameter (MLD) point is marked as a red dot.

those above 70% are generally classified as severe. According to the ACC/AHA guideline for the management of CAD [16], different severity levels may lead to different clinical decisions or therapeutic interventions. Specifically, the %DS severity level is used to guide decisions on revascularization and the need for further investigation.

Deep learning-based approaches to stenosis detection and severity classification using angiograms have shown promising results [6,9,14,17] in recent years. However, practical adaptation often encounters several critical challenges: data scarcity, class imbalance, and high labeling cost. While large-scale medical datasets are publicly available for other imaging modalities, such as chest X-ray, CT, and MRI [13,19,24], the largest publicly available angiogram dataset contains only around 4,000 images from at most 42 patients [9]. In addition, typical CAG datasets indicate that a majority of the cases are moderate stenosis [20]. This distribution biases neural networks toward the majority class and impedes their ability to detect more severe cases [15]. Another major limitation is the high labeling cost associated with stenosis severity annotation. Manual severity grading is labor-intensive and subject to significant variation across observers [23], which introduces labeling ambiguity. This ambiguity necessitates additional expert review and consensus-building, further increasing the overall cost of obtaining reliable annotations for robust model training [25].

One promising strategy to address these issues is to augment the data by synthesizing rare class images. Previous studies have employed segmentation masks to synthesize anatomically realistic medical images [5,12]. However, unlike clas-

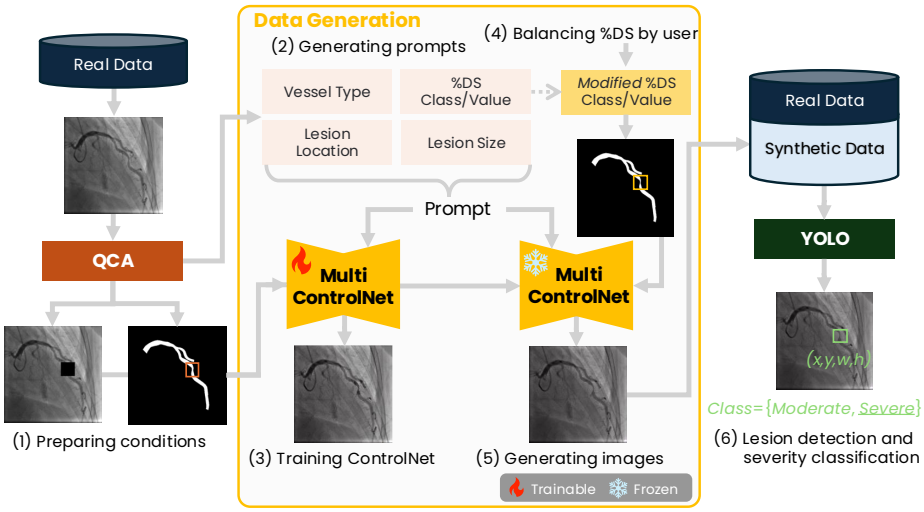

Fig. 2: **Overall pipeline of the proposed method.** (1,2) We prepare conditioning images and generate text prompts from metadata. (3) ControlNet is trained to reconstruct images and (4,5) inferred to generate new images with modified segmentation masks. (6) YOLO model captures lesions and predicts their severity.

sification problems in other medical imaging domains, lesions in angiograms can occur at any location within the vessels, making it impractical to rely solely on segmentation masks for lesion detection. Since segmenting all vessel branches is infeasible, using only main branch segmentation masks for image generation can lead to unintended lesion artifacts in non-target regions. If these unintended lesions are not explicitly labeled, they introduce noise into the dataset, potentially degrading model performance [8]. In contrast, our synthesis method aims at specific regions of interest (ROIs), allowing user-defined modifications to the lesion areas. By limiting changes to the ROIs, our approach synthesizes realistic variations in stenosis severity without introducing unintended artifacts outside ROIs. This targeted synthetic image, which is suitable for lesion detection training data, better reflects the diverse manifestations of severe lesions without the need for additional labeling cost.

To our knowledge, this work is the first to integrate an inpainting strategy with a generative model that synthesizes angiograms for lesion detection. The generated images preserve realistic vascular structure, as illustrated in Fig. 1. Utilizing user-defined masks to make a balanced class distribution enhances model performance compared to training without them. Extensive experimental results consistently demonstrate our pipeline outperforms the model without any synthetic data under data-scarce conditions.

2 Proposed Methods

Accurate detection of coronary lesions and measurement of %DS using a quantitative coronary analysis (QCA) tool remain both labor-intensive and susceptible to inter-operator variability. To mitigate these challenges, we develop a one-stage pipeline that performs lesion detection and severity classification (%DS $<$70% and $\geq$70%) using clinically grounded severity thresholds [16], with training enhanced by controllable synthetic data augmentation.

We first apply a conventional QCA tool to each real angiogram to extract lesion location and severity. A user-adjusted severity level is then utilized as a conditional input for a diffusion-based generative model, which synthesizes angiographic images reflecting various degrees of stenosis while preserving the vascular structure. This approach effectively expands the dataset, balances class distributions, and enhances detection model, particularly in light of limited labeled angiographic datasets. Fig. 2 illustrates the overall data augmentation pipeline, including generative model training and inference.

2.1 Vessel Segmentation and Severity Designation

Before generating a realistic angiogram image with a specified stenosis level (%DS), conventional QCA is applied to the original image to extract vessel contours, as shown in Fig. 1 (g). To meet the specific %DS value, we compute %DS using the reference vessel diameter (D_{ref}), which is estimated by interpolating the vessel contour from adjacent non-diseased vessel segments, while the minimum lumen diameter (MLD) is identified at the site of maximum stenosis within the lesion [3,4]. %DS is then computed as $\left(1 - (\text{MLD}/D_{\text{ref}})\right) \times 100$.

Our algorithm is designed to move two control points at the MLD location by shifting them in a direction orthogonal to the vessel direction such that it matches with the desired %DS. After redrawing the vessel contours, these modifications propagate smoothly to nearby points to ensure a natural transition without abrupt changes. The resulting contours are converted into segmentation masks, and they are encoded into a compact latent space and integrated into our conditional generation framework.

2.2 Data Augmentation with Multi-Control Diffusion Model

To generate angiograms with lesions that exhibit varying stenosis severity, we leverage ControlNet [26], an extension of Stable Diffusion [22], which integrates structural guidance to constrain the generation process. Unlike previous diffusion-based data augmentation methods [5,8], our approach incorporates three conditioning inputs across multiple ControlNets: (1) an original image on which a lesion bounding box is overlaid for inpainting, specifying regions that require synthesis, (2) a vessel segmentation mask that guides the model to adhere to the vascular structures, and (3) a lesion-specific text prompt. The text prompt is designed following [18], incorporating structured lesion information such as location, severity class, and vessel type. The guidance image $\mathbf{c}$, derived from

the segmentation mask and the masked image for inpainting, is encoded into latent representations using a lightweight convolution network, yielding $\mathbf{c}_\mathrm{s}$ and $\mathbf{c}_\mathrm{m}$, respectively. A denoising model ϵ_θ, conditioned on vessel structure or lesion location inputs, is fine-tuned to predict the noise added to the latent image $\mathbf{z}_t$, incorporating text prompt $\mathbf{c}_\mathrm{t}$ and timesteps $\mathbf{t}$, by optimizing the following objective:

$$\mathcal{L} = \mathbb{E}_{\mathbf{z}_t,\mathbf{t},\mathbf{c}_\mathrm{t},\mathbf{c}_\mathrm{s},\mathbf{c}_\mathrm{m}\sim\mathcal{N}(0,1)}[\|\epsilon - \epsilon_\theta(\mathbf{z}_t,\mathbf{t},\mathbf{c}_\mathrm{t},\mathbf{c}_\mathrm{s},\mathbf{c}_\mathrm{m})\|^2]. \quad (1)$$

The fine-tuned generative model synthesizes the masked structures through two separate denoising processes, ensuring seamless integration of newly generated regions with the original surroundings.

2.3 Stenosis Detection and Severity Classification

To enhance detection model performance, we add synthesized images to our training dataset, expanding the dataset to N-times the size of real data to balance the ratio of moderate-to-severe stenosis cases. Unlike existing multi-step pipelines [2,11], which require separate modules for vessel segmentation, stenosis detection and severity prediction, our method streamlines the process into a single-step detection framework leveraging the efficiency of YOLO [10]. We trained the model on coronary angiogram images, where the detection module generates bounding boxes of lesions with 50% or higher vessel narrowing, while the classification module assigns a severity class by computing the probability of belonging to each %DS category.

3 Experiments

Table 1: **Data statistics.**

	Internal			ARCADE
	Train	Valid	Test	Valid
#Patients	1402	535	232	200
#Images	5156	1854	884	200
#Lesions	6030	2150	1070	262
– Moderate	5350	1932	920	236
– Severe	680	218	150	26

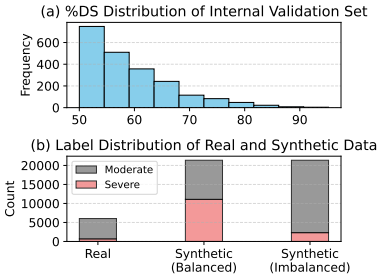

Fig. 3: **Label distributions.**

3.1 Datasets

Internal dataset is composed of 7,894 X-ray coronary angiograms (train 5,156, validation 1,854, test 884 images). We annotated lesions using an automatic QCA tool, with stenosis defined as %DS greater than 50. Unlike many existing datasets that rely solely on segmented vessels without specific lesion details, our dataset includes bounding box annotated lesions and corresponding %DS values, making it uniquely suited for both detection and classification tasks. For clinical relevance and adherence to diagnostic standards, QCA annotations are verified by experienced clinicians. The %DS values are categorized into clinically significant ranges (*e.g.*, %DS $<$70%, $\geq$70%) [16]. Fig. 3 illustrates the highly skewed label distributions of our dataset and detailed data statistics are described in Table 1.

To evaluate the generalizability of our method, we also utilize publicly available dataset as an external validation set for our pipeline. We used 200 validation coronary angiograms from 200 distinct patients with stenosis (defined as %DS $\geq$50%) from the **ARCADE** dataset [20]. One of the limitations of this dataset is that their annotations are segmented into regions, lacking precise regions of interest and reference standards for accurately specifying %DS values. This ambiguity poses a challenge for evaluating %DS classification tasks. To address this, we used the QCA tool to re-annotate the ARCADE validation set and collaborated with our in-house clinicians for refining labels. The labeling process involves identifying lesion locations with corresponding %DS values, ensuring high-quality and clinically relevant annotations. Furthermore, to foster future research in coronary stenosis detection, we will release the re-annotated dataset for public use at https://github.com/medipixel/DiGDA.

3.2 Experimental Setup

All input images are resized to a standard resolution of 512×512 to maintain consistency across experiments. We use segmentation masks along with bounding boxes as conditioning inputs to a ControlNet [26] component, which is encoded into a 64×64 latent. For downstream tasks, we fine-tune a YOLO model on our internal dataset using a learning rate of 2×10^{-4} and a batch size of 32. The model is trained for 50 epochs using the AdamW optimizer to ensure convergence. We use the F1-score metric to evaluate the balance between precision and recall for severity classification task and mean Average Precision at an intersection over union threshold of 50% (mAP50) which evaluate the detection accuracy for lesion detection. We report the average scores over three runs for all downstream experiments. We also conduct paired t-tests with bootstrapping to evaluate the performance gains of our proposed method following [21].

3.3 Quantitative Evaluation

Table 2 summarizes the performance of our method (DiGDA) compared to a baseline detector trained without data augmentation using synthetic images. DiGDA demonstrates superior results on both internal and external dataset,

Table 2: **Performance on different ratios of real to synthetic data.** The real-to-synthetic ratio is expressed as $\times N$, where N denotes the synthetic data scaling factor, while the number of targeted lesions per class is balanced. The mAP50 metrics are also reported for each stenosis severity classes, respectively: Moderate (M, %DS$<$ 70%) and Severe (S, %DS$\geq$ 70%).

Synthetic	**Internal dataset**				**ARCADE**			
data ratio	**F1**	**mAP50**	**(M**	**S)**	**F1**	**mAP50**	**(M**	**S)**
×0	0.650	0.688	0.731	0.646	0.500	0.464	0.436	0.492
×1	0.656	0.699	0.749	0.649	0.519	0.484	0.440	0.527
×2	0.658	0.710	0.767	0.654	**0.524**	**0.501**	**0.448**	**0.555**
×4	**0.670**	**0.717**	**0.771**	**0.663**	0.513	0.492	0.438	0.546

achieving higher F1 and mAP50 scores, for the increased size of balanced set with $\times 1$, $\times 2$, and $\times 4$ data ratio. For the synthetic data setting ($\times 4$), mAP50 increased from 0.688 to 0.717 (paired t-test, $p < 0.01$), indicating a significant improvement in lesion detection performance. Furthermore, on the ARCADE dataset ($\times 2$), our method outperforms the no-augmentation baseline with statistical significance ($p < 0.01$). For each setting ($\times N$), synthetic data equivalent to N times the amount of real data is first generated, followed by a removal of a portion of the generated data to ensure that the number of samples for each label remains equal across settings. This highlights the effectiveness of incorporating synthetic data generated using our approach.

Furthermore, we analyze the effect of dataset composition by comparing a class-balanced dataset to an imbalanced dataset for detection task, as shown in Fig. 4 (a). Unlike previous experiments that maintained class balance, this ablation introduces class imbalance, mirroring the label distribution of real data where class frequencies are naturally skewed. DiGDA, when trained on a balanced dataset, consistently achieves superior detection performance for underrepresented severe lesions, outperforming models trained on imbalanced data. DiGDA trained with a balanced dataset augmented by $\times 4$ yields a mAP50 gain with large gap compared to the model trained solely on real data, and a mAP50 increase from 0.707 to 0.717 compared to the imbalanced dataset ($p < 0.01$).

As illustrated in Fig. 4 (b), the proposed method improves performance under different subset size settings. While we collected a large-scale coronary angiography dataset, many researchers only have access to limited-scale datasets. Thus, we show that our approach consistently maintains higher mAP50, with notable improvements in severe stenosis class, which has been extensively enriched through data augmentation. Note that both the generative model and the detection model are trained on a subset.

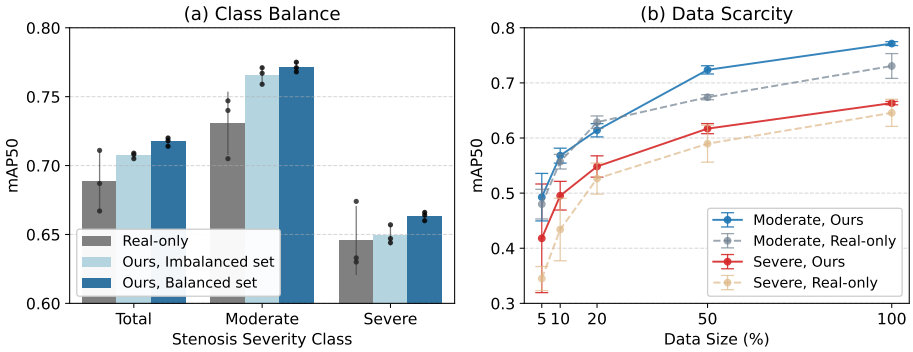

Fig. 4: **Performance comparisons on lesion detection across data compositions and dataset sizes.** (a) examines the effect of class balance: the balanced set has equal target lesions per class, while the imbalanced set maintains the original distribution with the same amount of the total data. (b) compares model performance across different subset sizes. Error bars indicate the standard deviation across three runs.

3.4 Qualitative Results

We conduct a qualitative analysis to visualize the capability of our method as illustrated Fig. 1. Our approach successfully generates realistic vessel narrowing patterns corresponding to various levels of %DS, ensuring anatomical plausibility and alignment with clinical expectations. Compared to the original images with %DS values computed by conventional QCA, our generated images exhibit higher %DS which represents the severe stenosis category, underrepresented in the dataset. The results show that our method accurately captures vessel constriction patterns while maintaining structural consistency across different vessel types. By using narrower vessel masks modified from original ones by user-defined %DS, our approach enables precise severity modulation overcoming the limitations of existing datasets. Furthermore, our inpainting strategy ensures that background regions, except for the areas within the red-outlined bounding boxes, remain unchanged, preserving the original angiographic context. This targeted modification ensures that an augmented dataset does not incur additional labeling cost for refinements.

4 Conclusion

In this study, we propose a novel approach to improve lesion detection and severity classification, integrating synthetic data generated via a diffusion model to vessel narrowing simulations into the training process. User-guided modifications to vessel masks enable the simulation of various severity levels that are underrepresented in the real-world datasets. By leveraging both internal and external datasets, we demonstrate the efficacy of synthetic data in improving

model performance, particularly in scenarios where real-world data is scarce or class-imbalanced. Our method achieves significant improvements in detection and classification, as confirmed in both quantitative and qualitative evaluations. The inclusion of synthetic data allows the model to learn various stenosis patterns while faithfully preserving the vascular structural integrity. Overall, this work demonstrates the potential of integrating synthetic data into lesion detection pipelines, paving the way for enhanced diagnostic tools in medical imaging.

Acknowledgment. The internal dataset was approved by the Institutional Review Boards of the participating hospitals, with informed consent waived, and collected in accordance with the Declaration of Helsinki. A retrospective data collection was performed on consecutive patients who underwent CAG at Sapporo Cardio Vascular Clinic (Sapporo, Japan; 2018–2023), Veterans Health Service Medical Center (Seoul, Republic of Korea; 2013–2023), and Soonchunhyang University Cheonan Hospital (Cheonan, Republic of Korea; before 2024). This work was supported by the Korea Medical Device Development Fund grant funded by the Korean government (the Ministry of Science and ICT, the Ministry of Trade, Industry and Energy, the Ministry of Health and Welfare, and the Ministry of Food and Drug Safety) (Project Number: RS-2023-00223121).

Disclosure of Interests. This study was supported by Medipixel, Inc. (Republic of Korea). The research was conducted using the company's proprietary software. The authors declare no additional competing interests.

References

1. Ahmad, F.B.: Mortality in the United States—provisional data, 2023. MMWR. Morbidity and Mortality Weekly Report **73** (2024) 1
2. Avram, R., Olgin, J.E., Ahmed, Z., Verreault-Julien, L., Wan, A., Barrios, J., Abreau, S., Wan, D., Gonzalez, J.E., Tardif, J.C., et al.: CathAI: fully automated coronary angiography interpretation and stenosis estimation. npj Digital Medicine **6**(1), 142 (2023) 5
3. Beatt, K.J., Luijten, H.E., de Feyter, P.J., van den Brand, M., Reiber, J.H., Serruys, P.W.: Change in diameter of coronary artery segments adjacent to stenosis after percutaneous transluminal coronary angioplasty: failure of percent diameter stenosis measurement to reflect morphologic changes induced by balloon dilation. Journal of the American College of Cardiology **12**(2), 315–323 (1988) 4
4. Brown, B.G., Bolson, E., Frimer, M., Dodge, H.T.: Quantitative coronary arteriography: estimation of dimensions, hemodynamic resistance, and atheroma mass of coronary artery lesions using the arteriogram and digital computation. Circulation **55**(2), 329–337 (1977) 4
5. Cobb, R., Cook, G.J.R., Reader, A.J.: Improved Classification Learning from Highly Imbalanced Multi-Label Datasets of Inflamed Joints in [99mTc]Maraciclatide Imaging of Arthritic Patients by Natural Image and Diffusion Model Augmentation . In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. vol. LNCS 15005, pp. 339 – 348. Springer Nature Switzerland (October 2024) 2, 4
6. Danilov, V.V., Klyshnikov, K.Y., Gerget, O.M., Kutikhin, A.G., Ganyukov, V.I., Frangi, A.F., Ovcharenko, E.A.: Real-time coronary artery stenosis detection based on modern neural networks. Scientific reports **11**(1), 7582 (2021) 2

7. Di Cesare, M., Perel, P., Taylor, S., Kabudula, C., Bixby, H., Gaziano, T.A., McGhie, D.V., Mwangi, J., Pervan, B., Narula, J., et al.: The heart of the world. Global heart **19**(1), 11 (2024) 1
8. Fang, H., Han, B., Zhang, S., Zhou, S., Hu, C., Ye, W.M.: Data augmentation for object detection via controllable diffusion models. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 1257–1266 (2024) 3, 4
9. Jiménez-Partinen, A., Molina-Cabello, M.A., Thurnhofer-Hemsi, K., Palomo, E.J., Rodríguez-Capitán, J., Molina-Ramos, A.I., Jiménez-Navarro, M.: CADICA: A new dataset for coronary artery disease detection by using invasive coronary angiography. Expert Systems **41**(12), e13708 (2024) 2
10. Khanam, R., Hussain, M.: Yolov11: An overview of the key architectural enhancements. arXiv preprint arXiv:2410.17725 (2024) 5
11. Kim, Y.I., Roh, J.H., Kweon, J., Kwon, H., Chae, J., Park, K., Lee, J.H., Jeong, J.O., Kang, D.Y., Lee, P.H., et al.: Artificial intelligence-based quantitative coronary angiography of major vessels using deep-learning. International Journal of Cardiology **405**, 131945 (2024) 5
12. Konz, N., Chen, Y., Dong, H., Mazurowski, M.A.: Anatomically-controllable medical image generation with segmentation-guided diffusion models. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 88–98. Springer (2024) 2
13. Li, P., Wang, S., Li, T., Lu, J., HuangFu, Y., Wang, D.: A Large-Scale CT and PET/CT Dataset for Lung Cancer Diagnosis (Version 5) [Data set] (2020). `https://doi.org/10.7937/TCIA.2020.NNC2-0461`, `https://www.cancerimagingarchive.net/collection/lung-pet-ct-dx/` 2
14. Lin, A., Manral, N., McElhinney, P., Killekar, A., Matsumoto, H., Kwiecinski, J., Pieszko, K., Razipour, A., Grodecki, K., Park, C., et al.: Deep learning-enabled coronary CT angiography for plaque and stenosis quantification and cardiac risk prediction: an international multicentre study. The Lancet Digital Health **4**(4), e256–e265 (2022) 2
15. Lin, H., Liu, T., Katsaggelos, A., Kline, A.: Stenunet: automatic stenosis detection from x-ray coronary angiography. arXiv preprint arXiv:2310.14961 (2023) 2
16. Members, W.C., Lawton, J.S., Tamis-Holland, J.E., Bangalore, S., Bates, E.R., Beckie, T.M., Bischoff, J.M., Bittl, J.A., Cohen, M.G., DiMaio, J.M., et al.: 2021 ACC/AHA/SCAI guideline for coronary artery revascularization: a report of the American College of Cardiology/American Heart Association Joint Committee on Clinical Practice Guidelines. Journal of the American College of Cardiology **79**(2), e21–e129 (2022) 2, 4, 6
17. Moon, J.H., Cha, W.C., Chung, M.J., Lee, K.S., Cho, B.H., Choi, J.H., et al.: Automatic stenosis recognition from coronary angiography using convolutional neural networks. Computer methods and programs in biomedicine **198**, 105819 (2021) 2
18. Oh, H.J., Jeong, W.K.: Controllable and Efficient Multi-class Pathology Nuclei Data Augmentation Using Text-Conditioned Diffusion Models. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 36–46. Springer (2024) 4
19. Pham, H.H., Nguyen, N.H., Tran, T.T., Nguyen, T.N., Nguyen, H.Q.: Pedicxr: an open, large-scale chest radiograph dataset for interpretation of common thoracic diseases in children. Scientific Data **10**(1), 240 (2023) 2
20. 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) 2, 6
21. Ribli, D., Horváth, A., Unger, Z., Pollner, P., Csabai, I.: Detecting and classifying lesions in mammograms with deep learning. Scientific reports **8**(1), 4165 (2018) 6
22. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022) 4
23. Shivaie, S., Tohidi, H., Loganathan, P., Kar, M., Hashemy, H., Shafiee, M.A.: Inter-observer Variability of Coronary Stenosis Characterized by Coronary Angiography: A Single-Center (Toronto General Hospital) Retrospective Chart Review by Staff Cardiologists. Vascular Health and Risk Management pp. 359–368 (2024) 2
24. Sunoqrot, M.R., Saha, A., Hosseinzadeh, M., Elschot, M., Huisman, H.: Artificial intelligence for prostate MRI: open datasets, available applications, and grand challenges. European radiology experimental **6**(1), 35 (2022) 2
25. Sylolypavan, A., Sleeman, D., Wu, H., Sim, M.: The impact of inconsistent human annotations on AI driven clinical decision making. NPJ Digital Medicine **6**(1), 26 (2023) 2
26. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3836–3847 (2023) 4, 6