



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Multi-Stage NeRF for Efficient 3D Coronary Artery Reconstruction from Two Narrow-Angle Angiographic Projections

Deyu Meng^[0009-0002-0659-827X], Mojtaba Lashgari^[0000-0002-5307-3495], Yiying Wang^[0009-0007-6572-5474], and Abhirup Banerjee^[0000-0001-8198-5128]

Institute of Biomedical Engineering, Department of Engineering Science,
University of Oxford, Oxford, United Kingdom
`deyu.meng@eng.ox.ac.uk`

Abstract. X-ray coronary angiography is the clinical gold standard for coronary artery disease during real-time cardiac interventions, but provides only 2D projections of inherently 3D vessels. Existing learning-based 2D-to-3D reconstruction methods typically require wide angular coverage or multiple views, assumptions that are rarely satisfied in routine practice where only two projections with narrow angular separation are available. To address these challenges, we propose NeCA++, a multi-stage self-supervised neural radiance field (NeRF) framework tailored to clinically realistic acquisition constraints. The framework decomposes reconstruction into two stages that progressively refine spatial support and representation capacity. In the first stage, a coarse 3D representation of the vasculature is reconstructed, restricting the subsequent optimisation to regions with a higher likelihood of vessel presence, termed an active region. Afterward reconstruction is restricted to this region while higher-resolution representations are progressively activated to recover fine vascular details. This multi-stage strategy focuses learning on anatomically plausible regions, mitigates gradient dilution under extreme sparsity, and stabilises global topology before recovering fine vascular branches. Furthermore, two vessel-specific regularisations are introduced: a ray-aligned constraint to reduce projection-induced ambiguity, and a bimodal density penalty to enable early vessel-background separation. Extensive experiments across three datasets (ImageCAS, ASOCA, and Synthetic RCA) and four angular configurations demonstrate consistent superiority over state-of-the-art baselines, particularly under clinically realistic narrow-angle settings, while achieving reconstruction within 58 seconds per case. The code is available at <https://github.com/MultiMeDIA-Oxford/NeCAPP>.

Keywords: 3D reconstruction · X-ray coronary angiography · Neural implicit representation · Self-supervised learning.

1 Introduction

Coronary artery disease (CAD) remains a leading cause of morbidity and mortality worldwide [20,22], highlighting the need for accurate diagnosis and effec-

tive intervention. X-ray coronary angiography is the gold standard for procedural guidance during coronary interventions and is one of the principal imaging modalities for the assessment of CAD, owing to its high spatial and temporal resolution and real-time visualisation of the coronary lumen [13]. However, it provides only two-dimensional (2D) images of three-dimensional (3D) vascular structures, resulting in vessel overlap and foreshortening. Clinicians subjectively infer the 3D structure of arteries from limited views, which can affect stenosis assessment and lesion localisation [9]. Automated 3D coronary artery (CA) reconstruction improves spatial understanding, enabling effective diagnosis and intervention.

Early 3D CA reconstruction methods are predominantly based on computational geometry, recovering vessel centerlines [3,6,12,21] or luminal surfaces [1,4,5,18,25] from multiple projections through explicit inter-view correspondences. Although these approaches can achieve high reconstruction accuracy under controlled settings, most of these methods require manual annotation and exhibit limited scalability and substantial computational time requirements. Learning-based methods, in contrast, mitigate these dependencies by automatically estimating model parameters from data. For example, DeepCA [23] employs a generative adversarial network (GAN) [8] for 3D CA reconstruction. However, such supervised approaches rely on a substantial amount of paired 2D angiographic projections and corresponding 3D ground truth (GT) data. In another approach, neural radiance field (NeRF) [16] introduces a self-supervised implicit representation that enables 3D reconstruction without case-specific annotations. NeRF-CA [15] and its improved version NeRT-CA [14] extend NeRF to dynamic reconstruction with emphasis on temporal consistency, reconstructing CA from four and three projections, respectively. In contrast, NeCA [24], also built on NeRF, addresses static reconstruction from two projections, but assumes a near-orthogonal angle separation between them. Nevertheless, the aforementioned methods differ from routine clinical practice, where only two projections with narrow angular separation are typically available.

To address these challenges under clinically constrained narrow-angle acquisition, we propose NeCA++, a computationally efficient self-supervised multi-stage neural implicit framework for 3D CA reconstruction from two narrow-angle projections. Rather than directly optimising a high-capacity model over a highly sparse volume in 3D space, NeCA++ decomposes optimisation into two stages that progressively refine spatial support and representation capacity. A coarse 3D vascular structure is reconstructed in the first stage to limit optimisation to a 3D subspace, termed an active region, using low-resolution hash encoding. The subsequent optimisation is restricted to this region, while higher-resolution hash levels are progressively activated to refine structural details. This staged design concentrates learning on anatomically plausible subspaces, mitigates gradient dilution under extreme sparsity, and stabilises topology before recovering fine vascular branches while speeding up the reconstruction. In addition, novel regularisations further constrain the solution under limited angular diversity. Extensive experiments are conducted under clinically realistic narrow-angle ac-

quisition protocols, a setting that remains under-explored in prior reconstruction studies. Results demonstrate robust performance over the ImageCAS, ASOCA, and Synthetic RCA datasets with significant improvements over state-of-the-art methods and an average runtime of 58 seconds per case.

Our main contributions are: **1)** A multi-stage neural implicit framework that localises an initial region of vascular presence and progressively reconstructs 3D coronary arterial tree based on a coarse-to-fine approach, for two projections with narrow-angle separation; **2)** Two vessel-specific regularisation strategies, including a ray-aligned geometry constraint and a bimodal density penalty, explicitly designed to mitigate shape inference ambiguity and facilitate early vessel-background separation; and **3)** Extensive validation under clinical narrow-angle acquisition protocols, demonstrating strong robustness and efficiency.

2 Method

This section presents the proposed NeCA++ for 3D coronary tree reconstruction from two projections with narrow-angle separation. Section 2.1 describes the overall structure of NeCA++. Section 2.2 describes a two-stage optimisation strategy where the first stage estimates a coarse 3D reconstruction of the vascular tree to isolate a highly likely 3D subspace where the vascular tree is present, termed an active region. Then, the second stage progressively refines predictions within the active region. Finally, Sec. 2.3 explains the vessel-specific regularisation terms in NeCA++ in order to reduce projection ambiguity and vessel-background separation under narrow angular separation.

2.1 Implicit Field Representation and Reconstruction Network

The proposed NeCA++ follows the neural field formulation of NeRF and reconstructs the unknown 3D structure of a vascular tree from two X-ray angiography projections Y_1, Y_2 acquired with narrow angular separation [9]. Let $\Omega \subset \mathbb{R}^3$ denote a bounded 3D space discretised into N voxel centers $\{\mathbf{p}_i\}_{i=1}^N$. The corresponding 2D projections of a real vascular tree, acquired using X-ray angiography, satisfy the following image formation: $\mathbf{c}_{1,i} = \mathbf{P}_1 \mathbf{p}_i, \mathbf{c}_{2,i} = \mathbf{P}_2 \mathbf{p}_i$, where $\mathbf{p}_i, \mathbf{c}_{1,i}$, and $\mathbf{c}_{2,i} \in \mathbb{R}^3$ are image coordinates, and $\mathbf{P}_1, \mathbf{P}_2$ are 3×3 projection matrices associated with image planes Y_1 and Y_2 . The objective is to estimate a vascular occupancy value at each spatial location $\mathbf{p}_i \in \Omega$.

A multi-resolution hash encoding inspired by InstantNGP [17] is applied on a general 3D space of Ω with the same $\mathbf{p}_i \in \Omega$, to enhance representation capacity and accelerate convergence. Features are extracted at L resolution levels. For each level ℓ , a hash table stores corresponding feature vectors as learnable parameters Φ . The input coordinate is mapped to the ℓ -th grid to locate the surrounding voxel vertices. Their coordinates are converted into hash indices to access the corresponding feature vectors. The retrieved vertex features are trilinearly interpolated at each level, and the resulting features from all L levels are concatenated to form the multi-resolution encoding $\gamma(x)$. Then, this encoded feature $\gamma(x)$ is




Fig. 1. Overview of the proposed NeCA++, the multi-stage NeRF framework. NeCA++ reconstructs 3D coronary geometry from two narrow-angle X-ray angiographic projections using a self-supervised neural implicit field. In Stage 1, NeCA++ localises an active region for optimisation from a coarse 3D reconstruction of the vasculature. In Stage 2, the coarse 3D reconstruction is optimised progressively by activating higher-resolution hash encoders to refine topology and recover fine vascular branches, while the optimisation is restricted to the active region.

mapped to a vascular occupancy value by a coordinate-based multilayer perceptron (MLP) f_θ . For each voxel centre $\mathbf{p}_i$, this gives $\sigma_\theta(\mathbf{p}_i) = \text{sigmoid}(f_\theta(\gamma(\mathbf{p}_i)))$, and the collection of these values over all voxels forms the reconstructed V .

To optimise reconstruction of an accurate 3D vascular structure in NeCA++, we enforce consistency with 2D input projections, Y_v , through L_{proj} along with two other regularisations explained in Sec. 2.3. With L_{proj} , first, 2D projection of the reconstructed volume, $\hat{V}_\theta$, is rendered using a differentiable cone-beam projector implemented with TIGRE [2], as highlighted in Fig. 1 (pink box). Then, the rendered projection $\hat{Y}_v$ is compared with the observed image Y_v (brown box in Fig. 1), to drive optimisation. The two-view projection loss is defined as

$$L_{\text{proj}} = \sum_{v \in \{1,2\}} \|\hat{Y}_v - Y_v\|_2^2. \quad (1)$$

Unlike conventional learning-based approaches that train a network across multiple training cases, the NeCA++ optimises network parameters specifically for each pair of 2D projections at the input.

2.2 Multi-stage Optimisation of Hash Encoding

Vascular structures occupy only a small fraction of Ω . As a result, generating 2D projections of the 3D reconstructed vessels to calculate L_{proj} from the entire Ω is inefficient. Instead, concentrating the optimisation on a subspace with more likelihood of vascular presence, termed the active region Ω^* , avoids unnecessary

computation associated with background voxels, leading to fast convergence. To exclude background voxels from computation, we adopt a two-stage optimisation strategy (Fig. 1). The first stage localises the active region, and the second refines predictions within this region while progressively increasing representational capacity to reconstruct finer details of vascular structure.

In the *first stage*, the network is optimised over the entire Ω and reconstructs a coarse structure of the vascular tree, $\hat{V}_\theta$ (blue box in Fig. 1). $\hat{V}_\theta$ is then binarised to define Ω^* . To reduce the risk of missing thin branches, a $3 \times 3 \times 3$ morphological dilation is applied to $\hat{V}_\theta$ to obtain Ω^* subspace. The first four low hash encoding levels are used in this stage, and the learned weight and low-resolution hash parameters are retained to initialise the second stage.

In the *second stage*, optimisation proceeds in a progressive manner. Here, the number of active hash encoding levels increases monotonically from the lowest level to higher-resolution hash levels. Through this multi-scale procedure, at the low hash levels, the global topology of the vascular structure is estimated. Then, by increasing the hash levels, finer levels are gradually introduced to sharpen boundaries and recover thin vascular branches.

2.3 Vessel-specific Structural Regularisation

The projection loss L_{proj} introduced in Sec. 2.1 cannot enforce accurate 3D shape reconstruction and separation between background and vascular structure due to overlaying the information of voxels along the X-ray beam at a pixel. For this purpose, two regularisation terms of L_{ray} and L_{bin} are introduced to enforce accurate 3D shape reconstruction and separation between background and vascular structure, respectively. To calculate L_{ray} , first we derive the X-ray beam direction from the known source to detector as $\hat{r}$. The gradients of the voxels along the beams are calculated as follows:

$$L_{\text{ray}} = \frac{1}{N} \sum_v \sum_j \sum_k \frac{\partial \sigma_\theta(p_k)}{\partial \hat{r}_{vj}}, \quad (2)$$

where $\sum_v$ is the summation over different views v , $\sum_j$ is the summation over all beams in each view, $\sum_k$ is the summation over all voxel centres p_k along each beam, and N is the total number of voxels.

In addition, to encourage separation between background and vessel in $\hat{V}_\theta$, we increase the influence of σ_θ with values close to 0 and 1, while suppressing influence of ambiguous σ_θ , i.e. $\sigma_\theta \rightarrow 0.5$, using a bimodal density penalty:

$$L_{\text{bin}} = \frac{1}{N} \sum_i \sigma_\theta(\mathbf{p}_i) (1 - \sigma_\theta(\mathbf{p}_i)), \quad (3)$$

where $\mathbf{p}_i \in \Omega^*$. This encourages earlier vessel-background separation during optimisation, which typically reduces the required training iterations. When used together with the ray-aligned constraint, the binarisation effect can improve reconstruction quality by making gradient anomalies more explicit.

The overall loss function is calculated as:

$$L = L_{\text{proj}} + \lambda_{\text{ray}} L_{\text{ray}} + \lambda_{\text{bin}} L_{\text{bin}}, \quad (4)$$

where $\lambda_{\text{ray}} = 7.5$ and $\lambda_{\text{bin}} = 2$ are weighting factors experimentally selected for the regularisation terms.

3 Experiments and Results

3.1 Experimental Settings

Datasets and Angle Configuration. Experiments are conducted on three clinical and synthetic 3D CA datasets, including the ImageCAS dataset ($N = 40$) [26] and ASOCA dataset ($N = 40$) [7], each comprising segmented left coronary arteries (LCA) and right coronary arteries (RCA). In addition, 100 synthetic RCA cases are generated following [11]. Consistent with prior works [14,15,23,24], 2D X-ray projections are generated from 3D volumes using TIGRE [2].

X-ray acquisition follows a standard C-arm configuration parameterised by (primary angle, secondary angle). Guided by clinical angiography protocols [10], we evaluate four angular separations for each artery type: RCA: orthogonal (0,0)/(0,90); large (40,-20)/(0,30); medium (40,0)/(0,30); small (40,0)/(20,20); LCA: orthogonal (0,0)/(0,90); large (-30,0)/(40,30); medium (-30,0)/(0,40); small (-30,0)/(-40,30).

Baselines, Metrics, and Implementation. We compare NeCA++ with three state-of-the-art 3D CA reconstruction models: NeRT-CA [14], DeepCA [23], and NeCA [24]. As DeepCA is supervised, 75% of the data are used for fine-tuning and the remaining 25% for testing. Since DeepCA was originally designed for RCA reconstruction, its evaluation is restricted to RCA cases for fair comparison. Reconstructed 3D volumes are thresholded into binary masks for evaluation. Dice, intersection over union (IoU), and centerline Dice (clDice) [19] are computed against GT 3D models. Due to GPU memory constraints, all experiments use a resolution of $80 \times 80 \times 80$ and are conducted on an Intel Ultra 9 275HX CPU with a single 16GB NVIDIA RTX 5080 GPU. All hyperparameters were selected based on the performance of the network over the validation set through sensitivity analysis. The hyperparameter settings in which the performance remained stable across reasonable parameter variations ($<2.3\%$ Dice difference) were used for all datasets and angular configurations.

3.2 Results

Reconstruction results are evaluated both qualitatively and quantitatively. Figure 2 shows representative 3D reconstructions. NeCA achieves near-perfect performance under orthogonal views but degrades rapidly as angular separation decreases, producing substantial outliers. DeepCA is less sensitive to angular settings but yields incorrect surface geometry, particularly around thin branches.




Fig. 2. 3D coronary artery reconstruction results on ImageCAS dataset, demonstrating significant outperformance of NeCA++ over state-of-the-art approaches. Designed for RCA reconstruction, DeepCA is evaluated only on RCA data.

NeRT-CA exhibits discontinuities and spurious vessels under two inputs. In contrast, our proposed NeCA++ preserves vascular geometry across both large and small angular separations, maintaining stable performance with limited outliers.

Table 1 reports 3D reconstruction performance under four pairs of angular separation (Orthogonal, Large, Medium, and Small) across three datasets. The data for DeepCA comes from the RCA results, while the other models were tested on both RCA and LCA. Under the Orthogonal setting, where angular separation is sufficient, NeCA achieves robust reconstruction on all datasets, whereas our proposed NeCA++ achieves comparable performance (Dice 0.909 on ImageCAS, 0.992 on ASOCA, and 0.993 on Synthetic RCA). As angular separation decreases (Large $\rightarrow$ Medium $\rightarrow$ Small), state-of-the-art methods degrade substantially, whereas NeCA++ remains consistent across datasets.

On ImageCAS, NeCA++ achieves the best Dice and IoU in the Large, Medium, and Small settings, with clear margins under Medium and Small angles (Dice 0.842 and 0.777, respectively); NeCA is unstable outside Orthogonal and NeRT-CA underperforms, while DeepCA preserves cIDice in some cases but yields lower volumetric accuracy. On ASOCA, our approach again outperforms all baselines, achieving Dice scores of 0.986, 0.941, and 0.932 from Large to Small angular separations; DeepCA attains competitive cIDice but shows lower overlap accuracy and reduced stability as angular separation narrows. On Synthetic RCA, similar trends are observed: NeCA performs strongly only under Orthogonal conditions but drops markedly under realistic angular constraints, and DeepCA achieves high cIDice yet remains inferior in Dice and IoU, whereas our NeCA++ consistently delivers the highest overlap scores across Large, Medium, and Small settings, indicating improved robustness to limited angular diversity.

Across datasets, our proposed NeCA++ maintains high cIDice even under narrow-angle settings, indicating preserved vascular connectivity, while achieving balanced improvements in both overlap and structural consistency. Overall, the

Table 1. 3D reconstruction performance under different X-ray angular settings. Dice, IoU, and cIDice are computed for each dataset.

Method	Orthogonal			Large			Medium			Small		
	Dice↑	IoU↑	cIDice↑	Dice↑	IoU↑	cIDice↑	Dice↑	IoU↑	cIDice↑	Dice↑	IoU↑	cIDice↑
ImageCAS Dataset												
NeRT-CA	0.582	0.410	0.486	0.581	0.410	0.657	0.537	0.367	0.526	0.451	0.291	0.424
NeCA	0.966	0.937	0.954	0.569	0.391	0.437	0.688	0.529	0.606	0.568	0.403	0.500
DeepCA	0.582	0.406	0.746	0.593	0.422	0.822	0.593	0.421	0.800	0.628	0.457	0.904
NeCA++	0.909	0.846	0.871	0.917	0.854	0.888	0.842	0.741	0.817	0.777	0.652	0.752
ASOCA Dataset												
NeRT-CA	0.631	0.461	0.510	0.576	0.405	0.664	0.472	0.309	0.525	0.538	0.368	0.572
NeCA	0.995	0.990	0.996	0.618	0.451	0.517	0.773	0.640	0.737	0.698	0.542	0.652
DeepCA	0.736	0.582	0.946	0.734	0.580	0.844	0.762	0.616	0.956	0.756	0.607	0.943
NeCA++	0.992	0.985	0.989	0.986	0.973	0.981	0.941	0.893	0.946	0.932	0.875	0.925
Synthetic RCA Dataset												
NeRT-CA	0.623	0.452	0.638	0.547	0.376	0.565	0.526	0.312	0.505	0.610	0.439	0.656
NeCA	0.990	0.981	0.995	0.582	0.414	0.444	0.735	0.585	0.645	0.642	0.475	0.557
DeepCA	0.560	0.389	0.798	0.660	0.493	0.923	0.755	0.607	0.960	0.718	0.560	0.969
NeCA++	0.993	0.986	0.993	0.930	0.872	0.902	0.914	0.845	0.910	0.880	0.790	0.894

proposed multi-scale framework generalises and maintains stable performance as angular separation decreases, with advantages most pronounced under the clinically realistic Medium and Small angular separation settings. Because all experiments are performed at 80^3 resolution due to memory constraints, the reported metrics may partly benefit from the coarser representation. Hence, we evaluate the upsampled reconstructions against 256^3 ground-truth volumes, achieving Dice scores of 0.71 on ImageCAS and 0.78 on ASOCA and outperforming the strongest baseline (NeCA) by 30% and 27%, respectively.

3.3 Ablation Study

We evaluate the contributions of the proposed multi-stage optimisation and regularisations using four variants: NeCA++ without multi-stage and two regularisations (NeCA), NeCA++ without L_{Ray} (W/o L_{Ray}), NeCA++ without L_{bin} (W/o L_{bin}), and NeCA++ without multi-stage (W/o MS). As shown in Table 2, L_{Ray} is the primary factor driving performance gains under small angular separations. The multi-stage strategy yields limited accuracy improvement but significantly accelerates optimisation, reducing the average training time from 632s to 58s per case. Using L_{Bin} alone degrades performance because it amplifies outliers under sparse supervision. However, when combined with L_{Ray} , it further improves results by encouraging sharper binarisation, which enhances the outlier removal effect of the ray constraint. More importantly, it accelerates network convergence.

Table 2. Ablation results averaged over all datasets for each angular setting.

Method	Large			Medium			Small			Time↓
	Dice↑	IoU↑	clDice↑	Dice↑	IoU↑	clDice↑	Dice↑	IoU↑	clDice↑	
NeCA	0.587	0.418	0.466	0.732	0.585	0.662	0.636	0.473	0.569	633s
W/o L_{Ray}	0.502	0.339	0.394	0.688	0.524	0.608	0.587	0.422	0.523	57s
W/o L_{Bin}	0.925	0.873	0.868	0.860	0.803	0.868	0.839	0.752	0.833	86s
W/o MS	0.931	0.882	0.920	0.887	0.822	0.877	0.859	0.768	0.856	632s
NeCA++	0.944	0.900	0.924	0.899	0.827	0.891	0.863	0.773	0.857	58s

4 Conclusion

This paper presents NeCA++, a multi-stage neural implicit framework for 3D coronary artery reconstruction from two angiographic projections with narrow angular separation. By progressively refining spatial support and representation capacity, and incorporating novel regularisation, the proposed approach improves stability under sparse and narrow-angle separation. Extensive experiments over three datasets demonstrate consistent superiority over existing methods, particularly under clinically realistic non-orthogonal configurations. The framework achieves efficient reconstruction within approximately one minute per case, highlighting its practical potential for real-world interventional settings.

Acknowledgments. The authors acknowledge the use of the services/facilities of the Institute of Biomedical Engineering (IBME), Department of Engineering Science, University of Oxford. The work of AB was supported by the Royal Society University Research Fellowship (grant no. URF\R1\221314) and the British Heart Foundation (BHF) Oxford Centre of Research Excellence (RE/24/130024). The work of ML was supported by the Royal Society Enhanced Research Expenses Grant, awarded to AB.

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

References

1. Banerjee, A., Galassi, F., Zacur, E., De Maria, G.L., Choudhury, R.P., Grau, V.: Point-cloud method for automated 3D coronary tree reconstruction from multiple non-simultaneous angiographic projections. *IEEE Transactions on Medical Imaging* **39**(4), 1278–1290 (2020)
2. Biguri, A., Sadakane, T., Lindroos, R., Liu, Y., Sabaté Landman, M., Du, Y., Lohvithee, M., Kaser, S., Hatamikia, S., Bryll, R., et al.: TIGRE v3: Efficient and easy to use iterative computed tomographic reconstruction toolbox for real datasets. *Engineering Research Express* **7**, 015011 (2025)
3. Blondel, C., Malandain, G., Vaillant, R., Ayache, N.: Reconstruction of coronary arteries from a single rotational x-ray projection sequence. *IEEE Transactions on Medical Imaging* **25**(5), 653–663 (2006)
4. Chen, S.J., Schäfer, D.: Three-dimensional coronary visualization, part 1: modeling. *Cardiology Clinics* **27**(3), 433–452 (2009)

5. Çimen, S., Gooya, A., Grass, M., Frangi, A.F.: Reconstruction of coronary arteries from x-ray angiography: A review. *Medical Image Analysis* **32**, 46–68 (2016)
6. Galassi, F., Alkhalil, M., Lee, R., Martindale, P., Kharbanda, R.K., Channon, K.M., Grau, V., Choudhury, R.P.: 3D reconstruction of coronary arteries from 2D angiographic projections using non-uniform rational basis splines (nurbs) for accurate modelling of coronary stenoses. *PloS one* **13**(1), e0190650 (2018)
7. Gharlegghi, R., Adikari, D., Ellenberger, K., Ooi, S.Y., Ellis, C., Chen, C.M., Gao, R., He, Y., Hussain, R., Lee, C.Y., et al.: Automated segmentation of normal and diseased coronary arteries—the asoca challenge. *Computerized Medical Imaging and Graphics* **97**, 102049 (2022)
8. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. *Communications of the ACM* **63**(11), 139–144 (2020)
9. Green, N.E., Chen, S.Y.J., Hansgen, A.R., Messenger, J.C., Groves, B.M., Carroll, J.D.: Angiographic views used for percutaneous coronary interventions: A three-dimensional analysis of physician-determined vs. computer-generated views. *Catheterization and Cardiovascular Interventions* **64**(4), 451–459 (2005)
10. Green, P., Frobisher, P., Ramcharitar, S.: Optimal angiographic views for invasive coronary angiography: A guide for trainees. *British Journal of Cardiology* **23**, 110–113 (2016)
11. Iyer, K., Nallamotheu, B.K., Figueroa, C.A., Nadakuditi, R.R.: A multi-stage neural network approach for coronary 3D reconstruction from uncalibrated x-ray angiography images. *Scientific Reports* **13**, 17603 (2023)
12. Jandt, U., Schäfer, D., Grass, M., Rasche, V.: Automatic generation of 3D coronary artery centerlines using rotational x-ray angiography. *Medical Image Analysis* **13**(6), 846–858 (2009)
13. Lashgari, M., Choudhury, R.P., Banerjee, A.: Patient-specific in silico 3D coronary model in cardiac catheterisation laboratories. *Frontiers in Cardiovascular Medicine* **11**, 1398290 (2024)
14. Maas, K.W.H., Ruijters, D., Pezzotti, N., Vilanova, A.: NerT-CA: Efficient dynamic reconstruction from sparse-view x-ray coronary angiography. In: *Reconstruction and Imaging Motion Estimation, and Graphs in Biomedical Image Analysis*. pp. 13–22. Springer Nature Switzerland, Cham (2026)
15. Maas, K.W., Ruijters, D., Vilanova, A., Pezzotti, N.: NeRF-CA: Dynamic reconstruction of x-ray coronary angiography with extremely sparse-views. *IEEE Transactions on Visualization and Computer Graphics* **31**(10), 8782–8795 (2025)
16. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
17. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)* **41**(4), 1–15 (2022)
18. Schoonenberg, G., Neubauer, A., Grass, M.: Three-dimensional coronary visualization, part 2: 3D reconstruction. *Cardiology Clinics* **27**(3), 453–465 (2009)
19. Shit, S., Paetzold, J.C., Sekuboyina, A., Ezhov, I., Unger, A., Zhylka, A., Pluim, J.P.W., Bauer, U., Menze, B.H.: cldice - a novel topology-preserving loss function for tubular structure segmentation. In: *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 16555–16564 (2021)
20. Townsend, N., Kazakiewicz, D., Lucy Wright, F., Timmis, A., Huculeci, R., Torbica, A., Gale, C.P., Achenbach, S., Weidinger, F., Vardas, P.: Epidemiology of cardiovascular disease in europe. *Nature Reviews Cardiology* **19**(2), 133–143 (2022)

21. Unberath, M., Taubmann, O., Hell, M., Achenbach, S., Maier, A.: Symmetry, outliers, and geodesics in coronary artery centerline reconstruction from rotational angiography. *Medical Physics* **44**(11), 5672–5685 (2017)
22. Virani, S.S., Alonso, A., Aparicio, H.J., Benjamin, E.J., Bittencourt, M.S., Callaway, C.W., Carson, A.P., Chamberlain, A.M., Cheng, S., Delling, F.N., Elkind, M.S.V., Evenson, K.R., Ferguson, J.F., Gupta, D.K., Khan, S.S., Kissela, B.M., Knutson, K.L., Lee, C.D., Lewis, T.T., Liu, J., Loop, M.S., Lutsey, P.L., Ma, J., Mackey, J., Matchar, D.B., Mussolino, M.E., Navaneethan, S.D., Perak, A.M., Roth, G.A., Samad, Z., Satou, G.M., Schroeder, E.B., Shah, S.H., Shay, C.M., Stokes, A., VanWagner, L.B., Wang, N., Tsao, C.W.: Heart disease and stroke statistics—2021 update: A report from the american heart association. *Circulation* **143**(8), e254–e743 (2021)
23. Wang, Y., Banerjee, A., Choudhury, R.P., Grau, V.: Deepca: Deep learning-based 3D coronary artery tree reconstruction from two 2D non-simultaneous x-ray angiography projections. In: 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 337–346. IEEE (2025)
24. Wang, Y., Banerjee, A., Grau, V.: NeCA: 3D coronary artery tree reconstruction from two 2D projections via neural implicit representation. *Bioengineering* **11**(12), 1227 (2024)
25. Wiesent, K., Barth, K., Navab, N., Durlak, P., Brunner, T., Schuetz, O., Seissler, W.: Enhanced 3-d-reconstruction algorithm for c-arm systems suitable for interventional procedures. *IEEE Transactions on Medical Imaging* **19**(5), 391–403 (2000)
26. Zeng, A., Wu, C., Lin, G., Xie, W., Hong, J., Huang, M., Zhuang, J., Bi, S., Pan, D., Ullah, N., et al.: ImageCAS: A large-scale dataset and benchmark for coronary artery segmentation based on computed tomography angiography images. *Computerized Medical Imaging and Graphics* **109**, 102287 (2023)