



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.

3DGR-CAR: Coronary artery reconstruction from ultra-sparse 2D X-ray views with a 3D Gaussians representation

Xueming Fu^{1,2}, Yingtai Li^{1,2}, Fenghe Tang^{1,2}, Jun Li^{1,2}, Mingyue Zhao^{1,2}, Gao-Jun Teng⁵, and S. Kevin Zhou^{1,2,3,4} ✉

¹ School of Biomedical Engineering, Division of Life Sciences and Medicine, University of Science and Technology of China (USTC), Hefei Anhui, 230026, China

² Center for Medical Imaging, Robotics, Analytic Computing & Learning (MIRACLE), Suzhou Institute for Advance Research, USTC, Suzhou Jiangsu, 215123, China

³ Key Laboratory of Precision and Intelligent Chemistry, USTC, Hefei Anhui, 230026, China

⁴ Key Laboratory of Intelligent Information Processing of Chinese Academy of Sciences (CAS), Institute of Computing Technology, CAS

⁵ Center of Interventional Radiology & Vascular Surgery, Department of Radiology, Medical School, Zhongda Hospital, Southeast University, Nanjing 210009, China

Abstract. Reconstructing 3D coronary arteries is important for coronary artery disease diagnosis, treatment planning and operation navigation. Traditional reconstruction techniques often require many projections, while reconstruction from sparse-view X-ray projections is a potential way of reducing radiation dose. However, the extreme sparsity of coronary arteries in a 3D volume and ultra-limited number of projections pose significant challenges for efficient and accurate 3D reconstruction. To this end, we propose 3DGR-CAR, a 3D Gaussian Representation for Coronary Artery Reconstruction from ultra-sparse X-ray projections. We leverage 3D Gaussian representation to avoid the inefficiency caused by the extreme sparsity of coronary artery data and propose a Gaussian center predictor to overcome the noisy Gaussian initialization from ultra-sparse view projections. The proposed scheme enables fast and accurate 3D coronary artery reconstruction with only 2 views. Experimental results on two datasets indicate that the proposed approach significantly outperforms other methods in terms of voxel accuracy and visual quality of coronary arteries. The code will be available in <https://github.com/windrise/3DGR-CAR>.

Keywords: 3D Gaussians Representation · Coronary artery reconstruction · Monocular depth estimation

1 Introduction

Cardiovascular disease, particularly coronary artery disease (CAD), is becoming increasingly prevalent worldwide [13]. Accurate 3D reconstruction of coronary

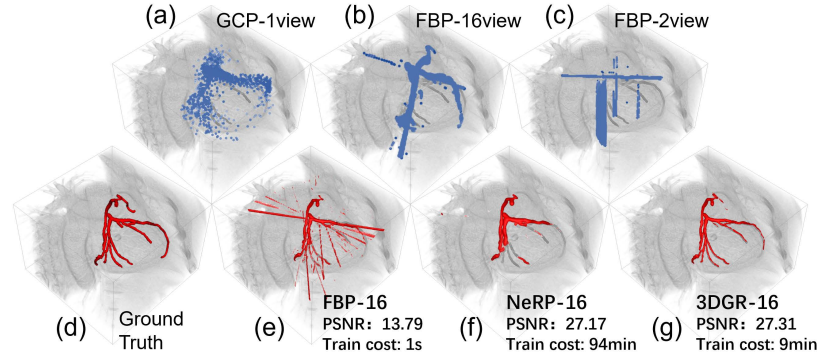


Fig. 1. Coronary arteries occupy only approximately 0.1% of a typical cardiac volume. (a-c): Gaussian centroid initial positions derived from our Gaussian center predictor and FBP results. (d): Ground truth. (e-g) Coronary reconstruction from FBP [20], NeRP [16] and 3DGR.

arteries greatly assists physicians in the diagnosis and treatment planning of CAD [14,15], enabling them to make informed decisions and provide targeted interventions. However, many projections are required during angiography to obtain accurate spatial structures of the vessels.

Classic scanning and reconstruction approaches (such as Feldkamp-Davis-Kress [2] and Filtered Back-Projection (FBP) [20]) that depend on dense view data require patients to be continuously exposed to ionizing radiation [5]. In contrast, sparse-view reconstruction techniques, which utilize a limited number of projections to reconstruct 3D structures, can potentially reduce radiation dose and minimize the risk to patients while still providing valuable diagnostic information [18]. Given the latest technical advances and the fact that coronary arteries occupy only approximately 0.1% of an entire volume for a typical cardiac scan (Fig. 1), a question naturally arises: **is it possible to utilize a really sparse number of 2D X-ray views to reconstruct coronary arteries in 3D?**

Due to its independence from training data, the neural radiation field (NeRF) [10] approach holds a great potential for medical image reconstruction [16,9,22,1]. Shen *et al.* propose an implicit neural representation (INR) with prior embedding (NeRP) to reconstruct images from sparsely sampled views [16]. However, INR-based approaches struggle with coronary artery reconstruction (CAR) due to the extreme sparsity of coronary artery image, resulting in a slow speed and limited performance according to our empirical observation, which hinders their practical application in clinical settings. Recently, the 3D Gaussian Splatting [6] has emerged as a notable strategy for reconstructing 3D scenes from images, with superior quality and faster convergence compared to NeRF methods. We identify 3D Gaussian representation (3DGR) as particularly well-suited for reconstructing extremely sparse objects such as coronary arteries, as they can be initialized from a sparse set of point clouds outlining the object, thereby avoiding

unnecessary computation in empty spaces. A few pioneering attempts have been made to introduce 3DGR to the medical imaging field [7,23,8]. For example, Li *et al.* utilize FBP-reconstructed images for initializing Gaussian parameters, achieving superior performance in sparse-view CT reconstruction compared to neural field methods [7]. However, when directly applying these strategies for coronary artery reconstruction, we empirically observe that the performance falls short of expectations.

We recognize that the performance gap lies in the deterioration of initialized Gaussians centers as a result of increasing sparsity of the object and reduced number of projections. Accurate initialization plays a crucial role in determining the final reconstruction quality. However, the initialization of Gaussian model parameters heavily relies on point cloud data obtained through Structure-from-Motion (SfM) techniques [6] or voxel data from FBP [7]. Thus, the noise level in the initial point cloud generated by SfM or FBP increases dramatically with a decreasing number of projections, leading to highly inaccurate Gaussian parameter initialization, which significantly degrades the quality of the final reconstruction, as shown in Fig. 1.

In this paper, we propose **3DGR-CAR**, a **3D Gaussian Representation** scheme for **Coronary Artery Reconstruction** from ultra-sparse 2D X-ray views. Our main contributions are as follows: (1) We introduce 3D Gaussian representation for coronary artery reconstruction. 3D Gaussians are initialized from a set of sparse point clouds, which avoids computation in empty spaces and improves efficiency, making it well-suited for the sparse nature of coronary artery images. This is **the first attempt** that successfully applies 3D Gaussian to coronary artery reconstruction. (2) We employ a U-Net to provide initialization of Gaussian centers. Point clouds estimated from ultra-sparse view are often extremely noisy and can harm the reconstruction performance of 3D Gaussians. By training the U-Net, we can leverage the knowledge stored in its parameter to place the initialized Gaussian centers in better locations, enhancing the reconstruction accuracy. (3) We conduct comprehensive evaluations on the ImageCAS and ASOCA datasets, demonstrating that our proposed 3DGR-CAR substantially surpasses current INR approaches and vanilla 3DGR method in terms of reconstruction quality with a considerably reduced time cost (Fig. 1).

2 Methods

As in Fig. 2, our proposed 3DGR-CAR consists of two stages: the Gaussian Center Predictor (GCP) training stage and the 3DGR reconstruction stage. In the first stage, a U-Net network [12] is trained to estimate voxel depth from a single simulated artery X-ray projection. In the second stage, a monocular image is processed through the U-Net to obtain the positional parameters $M = (d, \Delta_x, \Delta_y, \Delta_z)$. Subsequently, these predicted positional parameters are utilized for the initialization of 3D Gaussian centers. Afterwards, the parameters of the Gaussian model are optimized based on sparse vascular projection X-ray images, thereby reconstructing 3D coronary arteries from sparse views. In the following

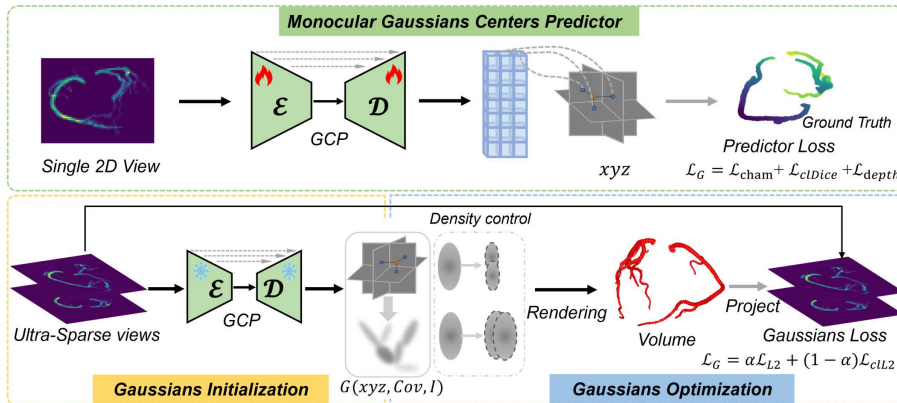


Fig. 2. Overview of 3DGR-CAR. The green box illustrates the pipeline for training GCP; the yellow box denotes the Gaussian position are initialized by GCP; the blue box represents Gaussian parameters optimization with sparse-view projections.

subsections, we first provide a brief overview of the 3DGR for CAR in Sec. 2.1, then detail training of the Gaussian center predictor (in Sec. 2.2).

2.1 3D Gaussians Representation Reconstruction

3D Gaussian is a flexible and expressive scene representation [6]. Gaussians are initialized from a set of sparse point clouds. Each 3D Gaussian is characterized by a set of parameters that define its position, shape, and other task-specific parameters. In this work, we align with [7], parameterize each Gaussian G_i by $\theta_i = \{\mu_i, \Sigma_i, I_i\}$, where μ_i pinpoints the coordinate of the center position for G_i , Σ_i is the covariance matrix that encapsulates information about the spread and orientation of the Gaussian, and I_i is the intensity of G_i . Each 3D Gaussian contributes to the overall representation through a function that gauges its impact on any given point in space, with the influence of each Gaussian described by

$$G_i(X|\theta_i) = I_i \cdot e^{-\frac{1}{2}(X-\mu_i)^T \Sigma_i^{-1}(X-\mu_i)}, \quad (1)$$

where X represents a point in 3D space. Each voxel $V(X)$ is generated from Gaussians within a local region around it:

$$V(X|\theta_{\{i\}}) = \sum_{i: \|X-\mu_i\| \leq d} G_i(X|\theta_i), \quad (2)$$

Parameters of these Gaussians are optimized through successive iterations of comparing the projected image to ground truth acquisitions, interleaved with adaptive density control that creates or destroys geometry based on how well the Gaussian fits the geometry to better represent the scene.

Loss Function. We combine the projection loss $\mathcal{L}_{L2} = \|\hat{P}_i - P_i\|_2^2$ and projected vessel centerline loss $\mathcal{L}_{clL2} = \|(\hat{P}_i - P_i) * M_{cl}\|_2^2$, and get the final loss

$$\mathcal{L}_G = \sum_i^N (\alpha \mathcal{L}_{L2} + (1 - \alpha) \mathcal{L}_{clL2}), \quad (3)$$

where \hat{P} and P represent the cone-beam projections of V and the actual volume, respectively. M_{cl} denotes the binary mask obtained by skeletonizing and binarizing the 2D projection images. $\alpha = 0.5$ set empirically.

2.2 Gaussian Center Predictor Training

To provide Gaussian center initialization with coronary prior and low-noise point cloud data, we propose to train a generalizable network for estimating a rough 3D point cloud from a single view. Specifically, we seek for a function $\mu = C(x)$ to predict Gaussian positions μ for a single projection x using an image-to-image neural network, which we refer to as the GCP Network.

Gaussian Center Predictor Network. More precisely, this network takes the $H \times W \times 1$ grayscale projection image as input and directly outputs a tensor of size $(H/\alpha) \times (W/\alpha) \times K$, where α is the downsampling factor used to reduce the number of initial Gaussians, which can be adjusted based on the sparsity of object to reconstruct. The K -dimensional vector represents the positional parameters $M = (d, \Delta_x, \Delta_y, \Delta_z)$ for each Gaussian, parameterised by depth d and a 3D offset Δ . In practice, the network can learn to automatically predict depth of a given view, providing a rough geometry of coronary artery. The Gaussian center predictor structurally aligns with the U-Net [19]. The final layer is replaced by a 1×1 convolution layer with four output channels, followed by an average pooling layer. The final output is reshaped into $N \times 4$, where N represents the number of predicted Gaussian centers, and 4 corresponds to the positional parameters M , which are transformed into the spatial positions of N Gaussians using the nonlinear activation function.

Learning Formulation. For training, we assume a multiview dataset consisting of real or simulated data. At minimum, the dataset comprises quintuplets $(P_{roj}, P_{oint}, V, d)$, where P_{roj} is a single 128×128 X-ray projection, P_{oint} is a point cloud composed of coronary 3D voxel positions in an $N \times 3$ format, V represents the $128 \times 128 \times 128$ voxel data obtained from transforming the point cloud P_{oint} , and d the depth map of P_{roj} . The loss function consists of three parts, each targeting a specific aspect.

1) **Chamfer Distance Loss.** It measures the distance between the set of predicted Gaussian center points and the set of actual label points, enabling the model to accurately localize the coronary artery.

$$\mathcal{L}_{cham} = \frac{1}{|S_1|} \sum_{p \in S_1} \min_{q \in S_2} \|p - q\|^2 + \frac{1}{|S_2|} \sum_{q \in S_2} \min_{p \in S_1} \|p - q\|^2, \quad (4)$$

where S_1 and S_2 are the two sets of predicted and ground truth points. p and q are points belonging to sets S_1 and S_2 , respectively.

Table 1. Quantitative comparison for 2-views coronary reconstruction with FBP, NeRP, 3DGR-CAR on ImageCAS and ASOCA. (mean \pm std, best results in bold).

Dataset	Method	New Projections		Volume	
		DSC(%) \uparrow	PSNR(dB) \uparrow	DSC(%) \uparrow	SSIM(%) \uparrow
ImageCAS	FBP	33.73 \pm 2.98	25.07 \pm 2.13	31.36 \pm 3.33	94.81 \pm 0.76
	NeRP	33.30 \pm 3.53	25.34 \pm 2.11	30.56 \pm 5.87	95.99 \pm 0.52
	<u>3DGR</u>	56.24 \pm 4.21	30.56 \pm 2.41	70.03 \pm 5.95	98.69 \pm 0.41
ASOCA	FBP	31.58 \pm 4.24	26.06 \pm 4.35	40.06 \pm 4.24	95.71 \pm 1.18
	NeRP	31.96 \pm 4.25	26.21 \pm 4.20	29.50 \pm 7.28	97.66 \pm 0.72
	<u>3DGR</u>	59.79 \pm 5.93	30.32 \pm 4.43	73.06 \pm 6.26	97.96 \pm 0.61

2) **Soft-CIDice Loss.** \mathcal{L}_{clDice} is a variant of the Dice Loss, specifically designed for evaluating tubular structure errors [17].

$$\mathcal{L}_{clDice} = \frac{2 \times |Skel_p \cap Skel_{gt}|}{|Skel_p| + |Skel_{gt}|}, \quad (5)$$

where $Skel_p$ and $Skel_{gt}$ represent the skeletons of the predicted voxel and the ground truth voxel, respectively.

3) **Depth Loss.** We employ the Scale-Invariant Logarithmic (SILog) Loss for uniformity and L1 Loss for the accuracy of predicted depth map. To obtain clear edge information, the L1 loss is also applied to the gradient space. The overall depth loss function is formulated as follows:

$$\mathcal{L}_{depth} = \mathcal{L}_{SILog} + \mathcal{L}_{GradL1} + \mathcal{L}_{MaskL1}. \quad (6)$$

In more detail, the SILog loss is defined as $\mathcal{L}_{SILog} = 10\sqrt{var(g) + \beta(mean(g))^2}$, where $g = \log(d_p + \alpha) - \log(d_{gt} + \alpha)$, and var and $mean$ represent the variance and mean, respectively. The gradient L_1 loss is $\mathcal{L}_{GradL1} = \sum |\nabla_x d_p - \nabla_x d_{gt}| + |\nabla_y d_p - \nabla_y d_{gt}|$, and the masked L_1 loss is $\mathcal{L}_{MaskL2} = \sum ((d_p - d_{gt}) \times M_{gt})^2$. The final loss function for training the Gaussian center predictor is as follows:

$$\mathcal{L}(C) = \gamma_1 \mathcal{L}_{cham} + \gamma_2 \mathcal{L}_{clDice} + \gamma_3 \mathcal{L}_{depth}, \quad (7)$$

where $\gamma_1 = 2 \log(\text{Iter}/20000)$, $\gamma_2 = 0.5$, and $\gamma_3 = 0.01$ set empirically.

3 Experiments

3.1 Setup

Dataset. Due to the unavailability and high cost of accurately calibrated, large datasets with paired X-rays and volumes, we employ digitally reconstructed radiographs (DRR) technology [11] to generate synthetic X-rays from multiple views. We take the 3D cone beam projection to simulate X-ray attenuation adapted from [16] on two coronary computed tomography angiography images (CCTA) datasets. The first is ImageCAS from [21]. The ImageCAS dataset comprises 1000 CCTA images, of which 960 are utilized to train the Gaussian center

predictor, 20 to select GCP, and 20 to test coronary reconstruction. The second is ASOCA from the MICCAI challenge ⁶ [3,4]. The ASOCA dataset includes 40 CCTA images, 20 used for GCP selection, and 20 for one-shot testing. Forty samples from both datasets are utilized to assess the performance of various reconstruction methods. All CCTA voxel data are resampled to $128 \times 128 \times 128$, with the size of the projection grayscale images set at 128×128 .

Evaluation Metrics. The evaluation results are presented in two main aspects: first, the masked Dice similarity coefficient (DSC) and masked Peak Signal-to-Noise Ratio (PSNR) for the new reconstructed coronary projections; second, the masked DSC and Structural Similarity Index Measure (SSIM) for the reconstructed coronary volumes.

3.2 Comparison with Existing Methods

We conduct quantitative comparisons of FBP [20], NeRP [16] and 3DGR-CAR on the ImageCAS and ASOCA datasets. The 2-view coronary reconstruction results are presented in Table 1. Statistical information on the evaluation of new projections and volumes coronary reconstructions from more views (2, 4, 8 and 16 views. The angles interval of views are $\pi/2$, $\pi/4$, $\pi/8$ and $\pi/16$.) on ImageCAS is presented in Fig. 3 and Fig. 4, respectively. The line graph represents the PSNR and SSIM (green font), while the histogram shows the DSC (black font). Statistical significance is indicated by asterisks as follows: a single asterisk (*) denotes $p < 0.05$. A double asterisk (**) indicates $p < 0.01$, while a triple asterisk (***) represents $p < 0.001$.

For both ImageCAS and ASOCA, 3DGR-CAR has significantly outperformed other methods in the evaluation of newly generated projections and reconstructed voxels. It can be observed that as the number of projection views used for reconstruction increases, the discrepancy gradually diminishes from Fig. 3 and Fig. 4. This indicates that 3DGR method possesses stronger and robust representation capabilities in ultra-sparse coronary projection reconstruction.

3.3 Ablation Study

To further assess the impact of different Gaussian center initialization methods, we also compare the performance of the FBP initialization method from [7] and the GCP initialization (abbreviated respectively as 3DGR-FBP and 3DGR-GCP). Table 2 presents the ablation study on both Gaussian initialization and the loss function Gaussian optimization. The 3DGR-GCP initialization method achieves superior performance on ImageCAS. This suggests that transferring prior knowledge from Gaussian center predictor can mitigate the effect of lacking 3D information to some extent, especially in extremely sparse scenarios. Also, both pixel-wise and contour-wise losses significantly impact the results.

Figure 5 demonstrates that the 3DGR-based method achieves the results most similar to the ground truth in sparse coronary projection reconstruction.

⁶ <https://asoca.grand-challenge.org/>

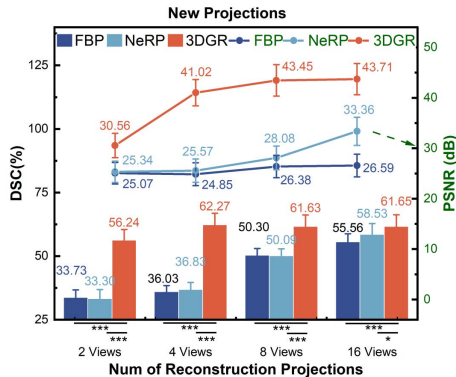


Fig. 3. Comparison of the new projection quality evaluation with increasing numbers of projections on ImageCAS.

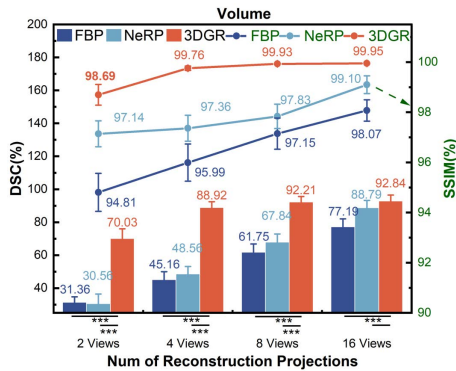


Fig. 4. Comparison of volume quality evaluation with increasing numbers of projections on ImageCAS.

Table 2. Ablation Study on the Gaussian center initialization and loss function of optimization.

Dataset	Method	New Projections		Volume	
		DSC(%) \uparrow	PSNR(dB) \uparrow	DSC(%) \uparrow	SSIM(%) \uparrow
ImageCAS	3DGR-FBP	52.86 \pm 4.95	28.90 \pm 3.11	63.84 \pm 9.73	97.15 \pm 0.58
	3DGR-GCP	56.24 \pm 4.21	30.56 \pm 2.41	70.03 \pm 5.95	98.69 \pm 0.41
ImageCAS	\mathcal{L}_{L2}	52.36 \pm 5.39	29.81 \pm 2.35	63.64 \pm 8.04	97.90 \pm 00.46
	$\mathcal{L}_{L2} + \mathcal{L}_{clL2}$	56.24 \pm 4.21	30.56 \pm 2.41	70.03 \pm 5.95	98.69 \pm 0.41

Furthermore, the 3DGR-GCP method outperforms the 3DGR-FBP approach under extremely sparse conditions (2-views).

4 Conclusion

In this paper, we introduce 3DGR-CAR, a novel 3D Gaussian Representation scheme for accurate coronary artery reconstruction from ultra-sparse 2D X-ray projections. This innovative scheme harnesses the power of 3D Gaussian representation to avoid the inefficiency in computation, and is adeptly tailored to provide accurate initialization from ultra-sparse projections by a U-Net. Extensive experimental results on the ImageCAS and ASOCA datasets demonstrate that the proposed 3DGR-CAR significantly outperforms existing INR methods in reconstruction quality, with much shorter time.

Acknowledgments. Supported by Natural Science Foundation of China under Grant 62271465, Suzhou Basic Research Program under Grant SYG202338, and Open Fund Project of Guangdong Academy of Medical Sciences, China (No. YKY-KF202206).

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

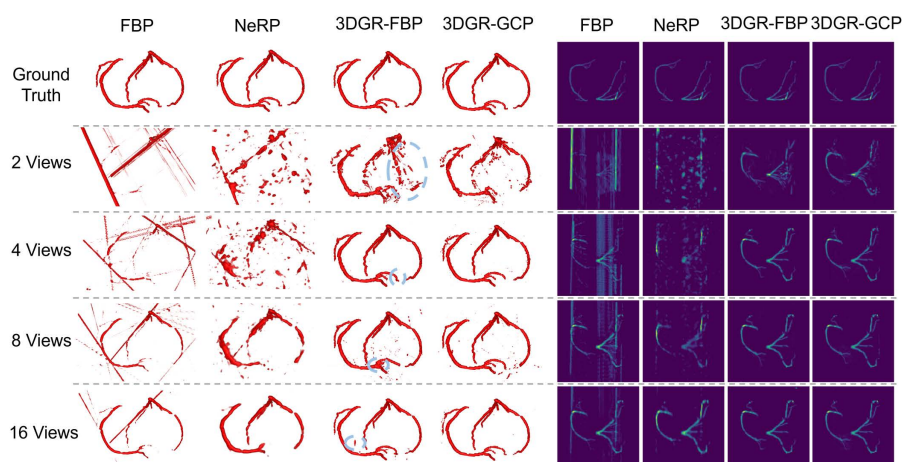


Fig. 5. Comparison of 3D vascular and new projections from different methods. The blue circles marked the most differences between 3DGR-FBP and 3DGR-GCP. (left: 3D coronary reconstruction results; right: new projections of coronary arteries).

References

1. Fang, Y., Mei, L., Li, C., Liu, Y., Wang, W., Cui, Z., Shen, D.: Snaf: Sparse-view CBCT reconstruction with neural attenuation fields. arXiv preprint arXiv:2211.17048 (2022)
2. Feldkamp, L.A., Davis, L.C., Kress, J.W.: Practical cone-beam algorithm. *Josa a* **1**(6), 612–619 (1984)
3. 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)
4. Gharlegghi, R., Adikari, D., Ellenberger, K., Webster, M., Ellis, C., Sowmya, A., Ooi, S.Y., Beier, S.: Computed tomography coronary angiogram images, annotations and associated data of normal and diseased arteries. arXiv preprint arXiv:2211.01859 (2022)
5. Grass, M., Köhler, T., Proksa, R.: 3D Cone-beam CT reconstruction for circular trajectories. *Physics in medicine and biology* **45** **2**, 329–47 (2000)
6. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3D gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4) (2023)
7. Li, Y., Fu, X., Zhao, S., Jin, R., Zhou, S.K.: Sparse-view ct reconstruction with 3d gaussian volumetric representation. arXiv preprint arXiv:2312.15676 (2023)
8. Liu, Y., Li, C., Yang, C., Yuan, Y.: Endogaussian: Gaussian splatting for deformable surgical scene reconstruction. arXiv preprint arXiv:2401.12561 (2024)
9. Maas, K.W., Pezzotti, N., Vermeer, A.J., Ruijters, D., Vilanova, A.: Nerf for 3d reconstruction from x-ray angiography: Possibilities and limitations. In: *VCBM 2023: Eurographics Workshop on Visual Computing for Biology and Medicine*. pp. 29–40. Eurographics Association (2023)

10. 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)
11. Milickovic, N., Baltas, D., Giannouli, S., Lahanas, M., Zamboglou, N.: Ct imaging based digitally reconstructed radiographs and their application in brachytherapy. *Physics in Medicine & Biology* **45**(10), 2787 (2000)
12. 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)
13. Roth, G.A., Mensah, G.A., Johnson, C.O., Addolorato, G., Ammirati, E., Baddour, L.M., Barengo, N.C., Beaton, A.Z., Benjamin, E.J., Benziger, C.P., et al.: Global burden of cardiovascular diseases and risk factors, 1990–2019: update from the gbd 2019 study. *Journal of the American College of Cardiology* **76**(25), 2982–3021 (2020)
14. Ryan, T.J.: The coronary angiogram and its seminal contributions to cardiovascular medicine over five decades. *Circulation* **106**(6), 752–756 (2002)
15. Serruys, P.W., Hara, H., Garg, S., Kawashima, H., Nørgaard, B.L., Dweck, M.R., Bax, J.J., Knuuti, J., Nieman, K., Leipsic, J.A., et al.: Coronary computed tomographic angiography for complete assessment of coronary artery disease: Jacc state-of-the-art review. *Journal of the American College of Cardiology* **78**(7), 713–736 (2021)
16. Shen, L., Pauly, J., Xing, L.: Nerp: implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. *IEEE Transactions on Neural Networks and Learning Systems* (2022)
17. Shit, S., Paetzold, J.C., Sekuboyina, A., Ezhov, I., Unger, A., Zhylka, A., Plum, J.P., Bauer, U., Menze, B.H.: cldice-a novel topology-preserving loss function for tubular structure segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16560–16569 (2021)
18. Sidky, E.Y., Kao, C.M., Pan, X.: Accurate image reconstruction from few-views and limited-angle data in divergent-beam ct. *Journal of X-ray Science and Technology* **14**(2), 119–139 (2006)
19. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020)
20. Wagner, W.: Reconstruction of object layers from their x-ray projections: A simulation study. *Computer Graphics and Image Processing* **5**(4), 470–483 (1976)
21. 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)
22. Zha, R., Zhang, Y., Li, H.: Naf: neural attenuation fields for sparse-view cbct reconstruction. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 442–452. Springer (2022)
23. Zhu, L., Wang, Z., Jin, Z., Lin, G., Yu, L.: Deformable endoscopic tissues reconstruction with gaussian splatting. *arXiv preprint arXiv:2401.11535* (2024)