

CardioGS: Deformation-Driven 4D Gaussian Splatting with Self-Supervised Cardiac Clock Learning for Rotational DSA Reconstruction

Yilun Liu^{1,2,3}, Jinzhu Yang^{1,2,3*}, Chaolu Feng^{1,2,3}, Yao Liu^{1,2,3,4}, Yongbo Wang⁵, Zhaoying Bian⁶, and Yan Xu⁷

¹ Key Laboratory of Intelligent Computing in Medical Image, Ministry of Education, Northeastern University, Shenyang, China, 110169

² School of Computer Science and Engineering, Northeastern University, Shenyang, China, 110169

³ National Frontiers Science Center for Industrial Intelligence and Systems Optimization, Shenyang, China, 110819

⁴ School of Computing, Macquarie University, Sydney, Australia, 2109

⁵ Key Laboratory of Biomedical Information Engineering of Ministry of Education, School of Life Science and Technology, Xi'an Jiaotong University, Xi'an 710049, China

⁶ School of Biomedical Engineering, Southern Medical University, Guangzhou 510515, China

⁷ Neusoft Medical Systems Co., Ltd., Shenyang, Liaoning 110094, China
yangjinzhu@cse.neu.edu.cn

Abstract. Rotational digital subtraction angiography (DSA) captures sparse C-arm projections of beating coronary vasculature over a 5–8s sweep [7,11], where each frame conflates two temporally distinct processes: quasi-periodic cardiac motion deforming vessel geometry and monotonic contrast-agent transport modulating opacity. Under sparse views, parameterizing both with a single acquisition timestamp is highly ambiguous. We present **CardioGS**, a deformation-driven 4D Gaussian Splatting framework that reduces this ambiguity through dual temporal coordinates. A self-supervised ClockNet learns a monotonic cardiac clock $u(\tau) \in [0, 1]$ via Gauss–Legendre quadrature of a bounded instantaneous frequency, adapting to heart-rate variability without ECG gating. Since $u(\tau)$ is monotonic, quasi-periodicity is imposed through cycle-consistency regularization on sparse deformation control nodes. Geometry is conditioned on the learned clock while appearance is conditioned on acquisition time in canonical space. On porcine rotational cardiac DSA with 15–60 views, CardioGS outperforms recent baselines in both 2D synthesis and 3D vessel geometry. The code is available at: <https://github.com/131366/cardiogs>

Keywords: Rotational DSA · 4D Gaussian Splatting · Cardiac Clock Estimation · Sparse-View Reconstruction

* Corresponding author.

1 Introduction

Coronary artery disease is the leading cause of mortality worldwide [16], and rotational DSA remains a primary modality for interventional coronary assessment [7,2,11]. Reconstructing a temporally coherent 4D model from a single rotational DSA acquisition would enable quantitative analysis of vessel deformation across the cardiac cycle, yet this task is fundamentally harder than standard sparse-view CT: during a 5–8s C-arm sweep, **every projection is acquired at a different cardiac phase and a different viewing angle**, so the observed intensity variations encode the compound effect of two physically distinct processes.

Two clocks, one timestamp. The first process is **cardiac motion**—quasi-periodic geometric deformation with patient-specific cycle length [13]. The second is **contrast-agent transport**—monotonic opacity evolution over the sweep [7]. Existing methods assign a single timestamp τ to each projection, making the inverse problem highly ambiguous: deformation may absorb photometric changes and vice versa, especially under sparse sampling (as few as 15 views) [4,18]. R^2 -Gaussian [19] ignores temporal dynamics entirely, producing streaky artifacts when the contrast agent is unevenly distributed across views. TOGS [20] models contrast dynamics but assumes static geometry—a design choice that precludes modeling cardiac motion and causes blurry reconstructions when vessels move. 4DRGS [8] also fixes geometry. Crucially, all three methods keep geometry static; under sparse views this single-timeline design allows geometric and photometric changes to freely trade off, producing temporally inconsistent reconstructions.

Our approach. We reduce this ambiguity by routing geometry and appearance to different temporal coordinates: a learned monotonic cardiac clock $u(\tau)$ for deformation, and acquisition time τ for opacity. Rather than enforcing periodicity through a discontinuous modular phase $\text{frac}(u(\tau))$, we keep $u(\tau)$ smooth and monotonic, and impose quasi-periodicity explicitly via cycle-consistency regularization on sparse deformation control nodes [14,3]. The learned clock adapts to intra-sweep heart-rate variability and concentrates temporal resolution near periods of rapid cardiac motion. A conditioning curriculum gradually transitions deformation from linear-time to clock-driven conditioning for stable co-optimization. Our key contributions are as follows:

1. We propose **CardioGS**, a dual-time 4DGS framework that separates the modeling of quasi-periodic cardiac motion and monotonic contrast transport, reducing the ambiguity that single-timeline dynamic methods [17,12] face under sparse views.
2. CardioGS comprises: **(i)** ClockNet, a self-supervised monotonic cardiac clock learned via Gauss–Legendre quadrature [1]; **(ii)** a sparse control-node deformation graph with cycle-consistency regularization as the explicit quasi-periodicity mechanism [14,3]; **(iii)** a canonical-space appearance field [10,9]; and **(iv)** a conditioning curriculum for stable co-optimization.




Fig. 1. Overview of CardioGS. ClockNet estimates a monotonic cardiac clock $u(\tau)$ to drive the deformation field, while the appearance field models contrast transport in canonical space conditioned on τ .

- Experiments on porcine rotational cardiac DSA (15/30/60 views) with ablation studies demonstrate state-of-the-art performance against R²-Gaussian, TOGS, and 4DRGS [19,20,8].

2 Method

2.1 Problem Formulation

In rotational cardiac DSA, a C-arm system acquires n log-subtracted projections over $T \approx 5\text{--}8$ s. Let $\mathcal{D} = \{(y_i, \pi_i, \tau_i)\}_{i=1}^n$ denote the dataset, where $y_i \in \mathbb{R}^{H \times W}$ is the DSA frame, π_i the projection geometry, and $\tau_i = i/n \in [0, 1]$ the normalized acquisition timestamp. This protocol induces a fundamental **view-time coupling**: each measurement y_i is a projection of the vasculature at a single time instant τ_i from a single viewpoint π_i . The observed intensity variations encode two distinct temporal processes: (1) quasi-periodic cardiac motion modulating vessel geometry, and (2) monotonic contrast evolution driven by hemodynamics [5]. Under sparse angular sampling, a model parameterized by a single timeline τ cannot uniquely attribute intensity changes to geometry or appearance—geometric deformation can spuriously absorb photometric fluctuations, and opacity changes can compensate for unmodeled motion [4,18].

2.2 Dual-Time Representation with a Learned Cardiac Clock

We separate the temporal coordinates governing geometry and appearance by introducing a learned **monotonic cardiac clock** $u(\tau) \in [0, 1]$. Deformation is

conditioned on $u(\tau)$; opacity on τ . This architectural separation reduces direct cross-contamination between deformation and opacity modeling by design. Importantly, $u(\tau)$ is *not* a modular phase; quasi-periodicity is enforced explicitly via cycle consistency (Sec. 2.4).

ClockNet: Monotonic Cardiac Clock via Differentiable Quadrature. ClockNet is a lightweight MLP f_θ that predicts a bounded instantaneous cardiac frequency:

$$f_\theta(\tau) = f_{\min} + (f_{\max} - f_{\min}) \cdot \sigma(\text{MLP}_\theta(\gamma(\tau))), \quad (1)$$

where σ is the sigmoid, $\gamma(\tau)$ is a Fourier positional encoding with L frequencies [15], and $[f_{\min}, f_{\max}]$ provides physiologically plausible bounds on instantaneous heart rate. Since $f_\theta(\tau) \geq f_{\min} \geq 0$, the **unwrapped cycle count**

$$w(\tau) = \int_0^\tau f_\theta(s) \, ds \quad (2)$$

is monotonic by construction. We evaluate this integral via Q -point **Gauss–Legendre quadrature** with precomputed nodes and weights [1], which is fully differentiable with degree of exactness $2Q - 1$; $Q = 16$ (exactness through degree 31) provides ample accuracy for smooth HRV trajectories while avoiding the overhead and potential instability of adaptive ODE solvers.

Let $C = w(1)$ denote the estimated number of cardiac cycles during the sweep. We define a **normalized cardiac clock coordinate**:

$$u(\tau) = \frac{w(\tau)}{C + \varepsilon} \in [0, 1], \quad (3)$$

where $\varepsilon = 10^{-6}$. To prevent spurious high-frequency oscillations in f_θ , we apply a total-variation regularizer:

$$\mathcal{L}_{\text{tv}} = \mathbb{E}_\tau[|f'_\theta(\tau)|]. \quad (4)$$

Remark (why not a modular phase). A modular phase $\text{frac}(w(\tau))$ introduces a discontinuity at every cycle boundary ($\approx 1 \rightarrow \approx 0$), creating a large numerical gap between physically adjacent cardiac states that prevents smooth MLP interpolation; our monotonic $u(\tau)$ avoids this and delegates quasi-periodicity to cycle-consistency regularization (Sec. 2.4). Critically, because $u(\tau)$ is *not* automatically periodic, the cycle-consistency loss provides a non-trivial gradient signal to ClockNet, enabling the reconstruction objective to calibrate the learned cardiac clock.

2.3 Control-Node Deformation Graph

We represent the canonical vasculature as 3D Gaussians $\mathcal{G}_0 = \{(\mu_k^0, q_k^0, \sigma_k^0)\}_{k=1}^M$ [6], where q_k^0 is a unit quaternion with $R(q_k^0) \in SO(3)$ the corresponding rotation matrix and $\sigma_k^0 \in \mathbb{R}_{>0}^3$ the per-axis scales. To impose spatial smoothness on the deformation field, we parameterize cardiac motion through N sparse learnable control nodes $\{c_j\}_{j=1}^N$ whose transformations are interpolated to individual Gaussians via KNN-RBF weighting [14,3].

Control Node Parameterization. A hash-encoded MLP [10] predicts per-node transformations:

$$[\Delta c_j(t), s_j(t)] = F_\phi(\hat{c}_j, t), \quad (5)$$

where $\Delta c_j \in \mathbb{R}^3$ is a translation offset and $s_j \in (0.5, 2.0)^3$ is bounded anisotropic scale modulation.

Conditioning Curriculum for Stable Co-optimization. Jointly optimizing ClockNet and the deformation field from scratch creates a cold-start problem: a randomly initialized cardiac clock provides uninformative conditioning. We address this with a **conditioning curriculum**:

$$t = (1 - \beta)\tau + \beta u(\tau), \quad \beta \in [0, \beta_{\max}], \quad (6)$$

where β increases linearly during training. At $\beta=0$, the deformation network is conditioned on τ , providing a well-posed initialization since temporally adjacent projections have nearby conditioning values. As $\beta \rightarrow \beta_{\max}$, conditioning gradually shifts to the HRV-adapted cardiac clock, which concentrates temporal resolution near periods of rapid cardiac motion.

KNN-RBF Interpolation. For each Gaussian k , deformation is interpolated from its 4 nearest control nodes $\mathcal{N}_k$:

$$w_{kj} = \exp(-\|\mu_k^0 - c_j\|^2 / (2r_j^2)), \quad \bar{w}_{kj} = w_{kj} / \sum_{j' \in \mathcal{N}_k} w_{kj'} \quad (7)$$

$$\mu_k(t) = \mu_k^0 + \sum_{j \in \mathcal{N}_k} \bar{w}_{kj} \Delta c_j(t), \quad (8)$$

$$\tilde{\sigma}_k(t) = \sigma_k^0 \odot \sum_{j \in \mathcal{N}_k} \bar{w}_{kj} s_j(t), \quad (9)$$

$$\Sigma_k(t) = R(q_k^0) \text{diag}(\tilde{\sigma}_k(t)^2) R(q_k^0)^\top \quad (10)$$

2.4 Quasi-Periodicity via Cycle Consistency on Control Nodes

Since $u(\tau)$ is monotonic rather than modular, cycle consistency is the **primary mechanism** by which CardioGS captures the quasi-periodic nature of cardiac motion. Let $C = w(1)$ and define a one-cycle shift in clock space:

$$\Delta u = \frac{1}{C + \varepsilon}. \quad (11)$$

For a sampled u , we compare control-node deformations at u and its nearest valid one-cycle-shifted counterpart u_{cmp} within $[0, 1]$ (using $u + \Delta u$ if feasible, otherwise $u - \Delta u$, to avoid boundary extrapolation):

$$\mathcal{L}_{\text{cycle}} = \mathbb{E}_\tau \left[\frac{1}{N} \sum_{j=1}^N (\|\Delta c_j(u) - \Delta c_j(u_{\text{cmp}})\|_1 + \|s_j(u) - s_j(u_{\text{cmp}})\|_1) \right]. \quad (12)$$

Table 1. Quantitative comparison on porcine rotational cardiac DSA. **Bold:** best; underline: second best.

Views	Method	CD↓	HD↓	PSNR↑	SSIM↑	GauN↓
15	R ² -Gaussian [19]	10.03	35.09	23.72	0.510	47 965
	TOGS [20]	24.06	78.63	17.08	0.349	1 098 058
	4DRGS [8]	15.40	70.71	25.17	0.567	218 978
	w/o Clock	<u>10.93</u>	<u>60.77</u>	<u>25.28</u>	<u>0.581</u>	<u>113 187</u>
	CardioGS	11.26	61.96	25.34	0.585	113 522
30	R ² -Gaussian [19]	5.25	12.74	25.07	0.562	48 577
	TOGS [20]	31.54	80.32	17.45	0.359	573 952
	4DRGS [8]	4.56	14.01	28.10	0.673	100 509
	w/o Clock	<u>3.58</u>	<u>9.40</u>	<u>28.94</u>	<u>0.690</u>	89 641
	CardioGS	3.17	9.11	29.30	0.710	<u>87 817</u>
60	R ² -Gaussian [19]	4.27	8.83	26.07	0.609	49 252
	TOGS [20]	16.68	52.12	14.95	0.323	1 241 545
	4DRGS [8]	2.71	6.95	30.59	0.749	80 409
	w/o Clock	2.28	<u>5.80</u>	<u>30.68</u>	<u>0.751</u>	89 417
	CardioGS	<u>2.42</u>	5.74	30.87	0.753	<u>78 719</u>

This regularizer serves a dual role. First, it encourages quasi-periodic deformation by tying states separated by one estimated cardiac cycle. Second, because both $u(\tau)$ and $\Delta u = 1/(C + \varepsilon)$ depend on ClockNet’s parameters through $w(\tau)$ and $C = w(1)$, the loss provides a self-supervised calibration signal for ClockNet’s cardiac clock from the reconstruction objective.

2.5 Canonical-Space Appearance Field

To prevent appearance from explaining away geometric motion, we model opacity in **canonical space**, querying the appearance field at undeformed positions μ_k^0 and conditioning on τ :

$$\alpha_k(\tau) = A_\psi([E_{3D}(x_k), E_{4D}([x_k, \tau])]), \quad (13)$$

where $x_k = (\mu_k^0 - o)/s_{\text{vol}}$ is the normalized canonical coordinate, E_{3D}/E_{4D} are hash-grid encoders [10], and A_ψ is a shallow MLP. Because the appearance field is queried at undeformed positions μ_k^0 rather than deformed positions $\mu_k(t)$, it is not directly conditioned on geometric motion; spatial variation it learns is anchored to canonical anatomy and primarily models contrast-agent dynamics.

3 Experiments

3.1 Experimental Setup

Dataset. We evaluate on a porcine rotational cardiac DSA dataset acquired with a Neusoft Medical Systems C-arm: 402 log-subtracted projections span-




Fig. 2. Qualitative comparison across 15, 30, and 60 views. CardioGS produces sharper vessel boundaries and fewer motion-induced artifacts than static-geometry baselines. Boxes highlight regions of notable improvement.

ning 240.72° at 648×474 px. We uniformly subsample 15, 30, and 60 views, leaving the remaining views for 2D synthesis evaluation. Ground-truth 3D vessel volumes are obtained from dense-view FDK reconstruction on all 402 frames. All animal procedures were approved by the Animal Welfare Ethics Committee of Beijing Kuibu Shichuang Biotechnology Co., Ltd. (approval no. KBSC21-2510-04, approved on October 17, 2025) and were conducted in accordance with institutional and national guidelines for the care and use of laboratory animals. The dataset is not publicly released due to proprietary restrictions.

Baselines & metrics. We compare against R^2 -Gaussian [19] (static 3DGS), TOGS [20] (temporal opacity, static geometry), and 4DRGS [8] (neural opacity, static geometry). All baselines are trained with default hyperparameters from their official implementations. To isolate the contribution of the cardiac clock, we additionally evaluate **w/o Clock**, an ablation variant that uses our deformation and appearance modules but conditions deformation on τ directly, omitting ClockNet and cycle-consistency regularization. Metrics: PSNR, SSIM (2D); Chamfer Distance CD (mm), Hausdorff Distance HD (mm) after ICP alignment (3D); Gaussian count GauN as a representation compactness indicator [8].

Implementation. All models train 30k iterations on a single A100 (80 GB). Density control runs iter 500–15k every 200 iter. ClockNet: $L = 3$, $Q = 16$, $f_{\min} = 0$, $f_{\max} = 8$. Deformation: $N = 2048$ control nodes, 4 nearest neighbors. Curriculum: $\beta \rightarrow \beta_{\max}$ over 2k iter warm-up. $\lambda_{\text{cycle}} = 0.2$, $\lambda_{\text{tv}} = 0.01$. All hyperparameters are fixed across view counts.

3.2 Comparison with State-of-the-Art

Table 1 presents quantitative results across three sparsity levels. CardioGS achieves the best PSNR and SSIM at every view count and the best overall 3D geometry at 30 and 60 views. Static-geometry methods (TOGS, 4DRGS) degrade substantially at 15 views because unmodeled cardiac motion produces view-inconsistent

intensity patterns that appearance-only models cannot reconcile. Both CardioGS variants with deformation achieve the highest rendering quality across all sparsity levels, confirming that explicit geometric modeling is essential for temporal consistency. At 15 views, R²-Gaussian obtains the lowest CD and HD, but this reflects phase-averaging rather than accurate geometry—its rendering quality is substantially worse. At $n \geq 30$, CardioGS outperforms R²-Gaussian on all metrics. At 30 views, adding the cardiac clock reduces CD from 3.58 to 3.17 mm (−11.5%) and raises PSNR by 0.36 dB. The advantage is consistent across all view counts, with CardioGS additionally achieving a more compact representation at 60 views (78.7K vs. 89.4K Gaussians).

3.3 Ablation Studies

We ablate three key hyperparameters on the 30-view setting (Fig. 3). **(a) λ_{cycle} .** Removing cycle consistency ($\lambda_{\text{cycle}} = 0$) drops PSNR by 0.85 dB and raises CD by 0.95 mm relative to the optimal $\lambda_{\text{cycle}} = 0.2$, confirming that this regularizer is essential for both calibrating the cardiac clock and inducing quasi-periodic deformation; excessively large values over-constrain the deformation field. **(b) L (ClockNet frequencies).** Without positional encoding ($L = 0$), ClockNet cannot represent intra-sweep heart-rate variability, yielding the worst performance across all ablations (PSNR 27.89, CD 4.65). $L = 3$ provides sufficient spectral bandwidth for the smooth frequency trajectory $f_\theta(\tau)$; larger L introduces unnecessary high-frequency capacity that encourages oscillatory f_θ , which the TV regularizer alone cannot fully suppress. **(c) N (control nodes).** $N = 1\text{K}$ underparameterizes the deformation field. $N = 8\text{K}$ weakens the spatial smoothness prior inherent in the sparse control-node design and risks over-fitting to the limited number of views. $N = 2\text{K}$ achieves the best trade-off between deformation expressiveness and regularization.




Fig. 3. Ablation of key hyperparameters (30 views). Stars mark the chosen configuration. (a) Cycle-consistency weight λ_{cycle} . (b) Fourier frequencies L in ClockNet. (c) Control node count N .

4 Conclusion

CardioGS models quasi-periodic cardiac motion and monotonic contrast transport with separate temporal coordinates via a self-supervised monotonic cardiac clock and cycle-consistency regularization, demonstrating consistent improvements over baselines on sparse-view porcine rotational DSA with the largest gains at clinically relevant sparsity levels. Single-dataset scope and the quasi-periodicity assumption under severe arrhythmias are the main limitations; multi-subject clinical validation and extension to other quasi-periodic organs are natural next steps.

Acknowledgments. This work was supported in part by the National Key Research and Development Program of China (Grant No. 2024YFC2417800), and the Research Startup Fund (Grant No. N25XQD030, No. 01270021922201*111).

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

References

1. Abramowitz, M., Stegun, I.A. (eds.): Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables, National Bureau of Standards Applied Mathematics Series, vol. 55. National Bureau of Standards, Washington, DC (1964), <https://www.nist.gov/mathematics-statistics/handbook-mathematical-functions-abramowitz-and-stegun>
2. Benamer, H., Ramus, L., Haidar, H.A., Perier, M., Saighi Bouaouina, M., Amelot, S., Vaillant, R.: First-in-man study of coronary stent 3D reconstruction in the cathlab using rotational angiography. *European Heart Journal - Imaging Methods and Practice* **3**(2), qyaf065 (2025). <https://doi.org/10.1093/ehjimp/qyaf065>
3. Buhmann, M.D.: Radial Basis Functions: Theory and Implementations, Cambridge Monographs on Applied and Computational Mathematics, vol. 12. Cambridge University Press (2003). <https://doi.org/10.1017/CB09780511543241>
4. Cai, Y., Wang, J., Yuille, A., Zhou, Z., Wang, A.: Structure-aware sparse-view X-ray 3D reconstruction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 11174–11183 (2024), https://openaccess.thecvf.com/content/CVPR2024/html/Cai_Structure-Aware_Sparse-View_X-ray_3D_Reconstruction_CVPR_2024_paper.html
5. Davis, B., Royalty, K., Kowarschik, M., Rohkohl, C., Oberstar, E., Aagaard-Kienitz, B., Niemann, D., Ozkan, O., Strother, C., Mistretta, C.: 4D digital subtraction angiography: Implementation and demonstration of feasibility. *AJNR American Journal of Neuroradiology* **34**(10), 1914–1921 (2013). <https://doi.org/10.3174/ajnr.A3529>
6. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3D gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4), 139:1–139:14 (2023). <https://doi.org/10.1145/3592433>, SIGGRAPH 2023
7. Lauritsch, G., Boese, J., Wigström, L., Kemeth, H., Fahrig, R.: Towards cardiac C-arm computed tomography. *IEEE Transactions on Medical Imaging* **25**(7), 922–934 (2006). <https://doi.org/10.1109/TMI.2006.876166>

8. Liu, Z., Zha, R., Zhao, H., Li, H., Cui, Z.: 4DRGS: 4D radiative gaussian splatting for efficient 3D vessel reconstruction from sparse-view dynamic DSA images. In: *Information Processing in Medical Imaging – IPMI 2025. Lecture Notes in Computer Science*, vol. 15830, pp. 361–374. Springer (2026). https://doi.org/10.1007/978-3-031-96625-5_24
9. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing scenes as neural radiance fields for view synthesis. In: *Computer Vision – ECCV 2020. Lecture Notes in Computer Science*, vol. 12346, pp. 405–421. Springer (2020). https://doi.org/10.1007/978-3-030-58452-8_24
10. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics* **41**(4), 102:1–102:15 (2022). <https://doi.org/10.1145/3528223.3530127>, SIGGRAPH 2022
11. Neubauer, A.M., Garcia, J.A., Messenger, J.C., Hansis, E., Kim, M.S., Klein, A.J., Schoonenberg, G.A., Grass, M., Carroll, J.D.: Clinical feasibility of a fully automated 3D reconstruction of rotational coronary X-ray angiograms. *Circulation: Cardiovascular Interventions* **3**(1), 71–79 (2010). <https://doi.org/10.1161/CIRCINTERVENTIONS.109.897629>
12. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-NeRF: Neural radiance fields for dynamic scenes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 10318–10327 (2021), https://openaccess.thecvf.com/content/CVPR2021/html/Pumarola_D-NeRF_Neural_Radiance_Fields_for_Dynamic_Scenes_CVPR_2021_paper.html
13. Rohkohl, C., Lauritsch, G., Biller, L., Hornegger, J.: ECG-gated interventional cardiac reconstruction for non-periodic motion. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2010. Lecture Notes in Computer Science*, vol. 6361, pp. 151–158. Springer (2010). https://doi.org/10.1007/978-3-642-15705-9_19
14. Sumner, R.W., Schmid, J., Pauly, M.: Embedded deformation for shape manipulation. *ACM Transactions on Graphics* **26**(3), 80-es (2007). <https://doi.org/10.1145/1276377.1276478>, SIGGRAPH 2007
15. Tancik, M., Srinivasan, P.P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J.T., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. In: *Advances in Neural Information Processing Systems*. vol. 33, pp. 7537–7547 (2020), <https://proceedings.neurips.cc/paper/2020/hash/55053683268957697aa39fba6f231c68-Abstract.html>
16. World Health Organization: The top 10 causes of death. Fact sheet (2024), <https://www.who.int/news-room/fact-sheets/detail/the-top-10-causes-of-death>, accessed 24 February 2026
17. Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Wang, X.: 4D gaussian splatting for real-time dynamic scene rendering. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 20310–20320 (2024), https://openaccess.thecvf.com/content/CVPR2024/html/Wu_4D_Gaussian_Splatting_for_Real-Time_Dynamic_Scene_Rendering_CVPR_2024_paper.html
18. Xu, D., Yang, Y., Liu, H., Lyu, Q., Descovich, M., Ruan, D., Sheng, K.: To-moGRAF: An X-ray physics-driven generative radiance field framework for extremely sparse view CT reconstruction. *PLOS ONE* **20**(8), e0330463 (2025). <https://doi.org/10.1371/journal.pone.0330463>

19. Zha, R., Lin, T.J., Cai, Y., Cao, J., Zhang, Y., Li, H.: R²-Gaussian: Rectifying radiative gaussian splatting for tomographic reconstruction. In: Advances in Neural Information Processing Systems. vol. 37, pp. 44907–44934 (2024). <https://doi.org/10.52202/079017-1427>, https://proceedings.neurips.cc/paper_files/paper/2024/hash/4f550cb7b30b59553e50cd08a9dbf068-Abstract-Conference.html
20. Zhang, S., Zhao, H., Zhou, Z., Wu, G., Zheng, C., Wang, X., Liu, W.: TOGS: Gaussian splatting with temporal opacity offset for real-time 4D DSA rendering. IEEE Journal of Biomedical and Health Informatics **29**(9), 6793–6805 (2025). <https://doi.org/10.1109/JBHI.2025.3575613>