# PhyRadGS: Physics-Aware Radiative Gaussian Splatting for Fly-Scanning CBCT Reconstruction

Anwei Li[1] [0000−0002−5386−2689], Zhongyang Zheng[1] [0009−0007−1808−5486], Tianyuan Sun[1] [0000−0002−5051−2552], Yubing Zhang[1], and Rongqiu Wang[2]⋆

[1] CVTE Research, China
{lianwei,zhengzhongyang,suntianyuan,zhangyubing}@cvte.com
[2] Maker-Ray, China
wangrongqiu@maker-ray.com

**Abstract.** Cone-Beam CT (CBCT) systems utilizing continuous gantry rotation and flat-panel detector readout introduce complex acquisition dynamics, such as angular motion blur and rolling shutter distortion. Conventional reconstruction algorithms, however, typically rely on a static geometry assumption, inevitably leading to the blurred structures and geometric inconsistencies. In this paper, we propose **PhyRadGS**, a physics-aware 3DGS framework that integrates a novel spatial temporal discretization strategy. Specifically, We explicitly model rolling shutter by spatially partitioning projections and motion blur by temporally segmenting the exposure duration. Furthermore, we implement a parallelized multi-pose batch rendering pipeline to mitigate computational overhead, maintaining the real-time efficiency inherent to 3DGS. Experiments on synthetic and real-world datasets demonstrate PhyRadGS significantly outperforms state-of-the-art static baselines, improving PSNR by *+2.3 dB* and SSIM by *+0.043*. Qualitative comparisons further confirm that PhyRadGS effectively recovers high-frequency clinical details even under rapid scanning protocols. The source code and dataset is available at https://github.com/liaw05/PhyRadGS.

**Keywords:** CBCT · 3D Gaussian Splatting · Motion Blur · Rolling Shutter · Fly-Scanning.

## 1 Introduction

Motion artifact is one of the main challenges in Cone-Beam Computed Tomography (CBCT) systems. Previous works [13] discussed artifacts caused by involuntary motion for surgical scenarios. Yet another cause of motion artifact remains untouched: modern CBCT systems utilize **fly-scanning**, where the gantry rotates non-stop while the detector acquires projections in a rapid stream.

Conventional reconstruction approaches [1,11,29] explicitly rely on an instantaneous static assumption (stop-and-shoot), presuming that the geometries

⋆ Corresponding Author

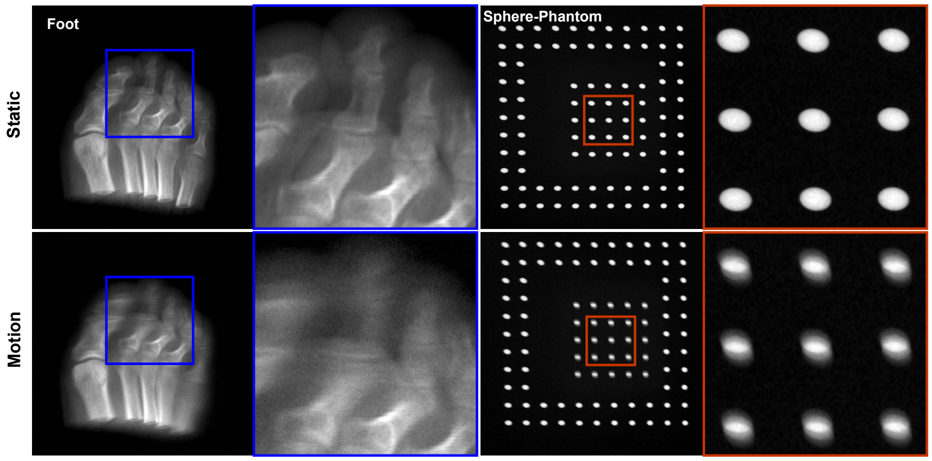


Fig. 1: Comparison of static and motion projections.

of the X-ray source and detector remain constant during the exposure of a single projection. In the context of fly-scanning, this assumption is fundamentally violated, leading to two primary artifacts(see Fig. 1 for sample projection):

- **Motion Blur:** During the detector integration period, the gantry traverses a non-negligible angular arc. The recorded pixel intensity is not a static sample but an integral over a continuous trajectory, causing angular smearing.
- **Rolling Shutter Distortion:** Flat-panel detectors typically read out data row-by-row. Since each row continues to integrate photons until its specific readout time, the effective exposure window shifts temporally across the panel. This results in row-dependent viewing angles within a single projection, leading to geometric warping.

Directly applying static models to such fly-scanning data is fundamentally ill-posed, inevitably resulting in reconstructed volumes with blurred boundaries and geometric distortions, missing high-frequency details such as bone trabeculae.

Addressing these artifacts requires explicitly incorporating the temporal dimension into the reconstruction model. Traditional analytic FDK [11] relies on rigid geometric axioms, making the integration of complex temporal dynamics theoretically intractable. While iterative methods (e.g. SART [1]) and implicit neural representations (e.g. NAF [29]) can theoretically model motion via temporal super-sampling, the computational cost scales linearly with temporal resolution, rendering them impractical for intraoperative use.

To mitigate motion artifacts, early medical approaches explored sinogram pre-processing [6,7] or hardware-based modulation [10,20]. More sophisticated model-based approaches [5,14,23] directly incorporated non-linear forward models of motion blur. Despite limitations in convergence speed and efficiency, these works underscore a critical principle: *enhanced fidelity of the forward model to the physical acquisition process yields superior reconstruction clarity.*

Recently, 3D Gaussian Splatting (3DGS) [15] has emerged as a revolutionary representation with real-time rasterization capabilities. However, pioneering X-ray implementations [3,16,17,28] largely neglect the temporal integration required for fly-scanning. Similarly, works that focus on dynamic surgical scenes [8,12,18] typically assume instantaneous exposure, failing to model the integration blur inherent in X-ray detectors. $X^2$-Gaussian [27] targets 4D CT for self-supervised respiratory motion: motion caused by patient but not the system.

A straightforward adaptation of dynamic 3DGS methods [9,19,24,25] from computer vision (CV) poses fundamental challenges. These approaches typically treat camera motion as an unknown latent variable for blind deblurring, disregarding the precise mechanical kinematics known in CBCT systems. This leads to optimization instability in sparse-view settings. Even methods that explicitly model rolling shutter, such as GS-on-the-Move [21], simplify the physical process by updating only the Gaussian centers while omitting the covariance deformation. This approximation introduces physically inconsistent artifacts in tomographic reconstruction.

In this paper, we propose **PhyRadGS** (**Phy**sics-Aware **Rad**iative **G**aussian **S**platting), the first X-ray 3DGS framework that physically models both motion blur and rolling shutter effects, to the best of our knowledge. Our main contributions are as follows.

1. PhyRadGS integrates **Physics-Aware Acquisition Model** with mechanical kinematics and detector readout characteristics, recovering sharp microstructures from blurred detector projections.
2. By mathematically modeling the **Spatio-Temporal Discretization**, we ensure reconstruction aligns strictly with the physical reality of the acquisition system.
3. Leveraging a custom **Multi-Pose Batch Rendering** kernel, PhyRadGS converts 3DGS rasterization process into parallelized batches of discrete micro-views, and reconstructs from fly-scanning data with negligible computational overhead compared to static methods.

## 2 Methodology

Our goal is to reconstruct the 3D attenuation field from projections acquired under continuous gantry rotation with rolling-shutter readout. We represent the target scene with a group of learnable 3D Gaussians, where the $i$-th Gaussian $\mathcal{G}_i$ is defined by center position $\mu_i \in \mathcal{R}^3$, rotation $R_i \in \mathcal{SO}^3$, scale $\mathbf{s}_i \in \mathcal{R}^3$, and density $\rho_i \in \mathcal{R}$. To physically model both motion blur and rolling shutter, we propose **PhyRadGS**, a framework grounded in a **Physics-Aware Acquisition Model**. As illustrated in Fig. 2, our pipeline realizes this model through two key implementation modules: **Spatio-Temporal Discretization** and **Multi-Pose Batch Rendering**.

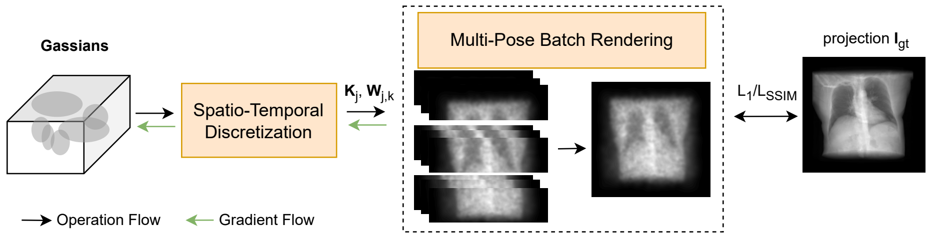


Fig. 2: **Overview of PhyRadGS.** The pipeline includes: **(a) Spatio-Temporal Discretization**, which converts continuous kinematics into discrete micro-states **(b) Multi-Pose Batch Rendering**, which efficiently parallelizes the physics simulation, and then assembled into the final projection.

### 2.1 Physics-Aware Acquisition Model

A typical fly-scanning process is shown in Fig. 3, the gantry angle $\theta(t)$ changes continuously during the exposure duration $T_{exp}$, and the detector rows $y \in [0, H]$ are read out sequentially over a duration $T_{ro}$. This creates a space-time parallelogram in the acquisition domain (yellow shaded area in Fig. 3b), where the vertical axis represents rows in the detector, and the horizontal axis represents the exposure duration of each row. A pixel value $P(x, y)$ in frame $f$ (in the log-attenuation domain) is not a sample at a single timestamp, but a temporal integral of the instantaneous projection operator $\Phi$:

$$P(x,y) = \frac{1}{T_{exp}} \int_{t^f_{start}+\tau_{row}(y)}^{t^f_{start}+\tau_{row}(y)+T_{exp}} \sum_i \Phi(\boldsymbol{\theta}(t), \mathcal{G}_i)\, dt \tag{1}$$

where $\tau_{row}(y) = \frac{y}{H-1} T_{ro}$ denotes the rolling shutter delay for row $y$. Unlike blind deblurring methods that treat motion as latent noise, here $\boldsymbol{\theta}(t)$ is derived from the known mechanical kinematics, ensuring the model respects physical reality.

### 2.2 Spatio-Temporal Discretization

Direct integration of Eq.(1) is computationally intractable. To ensure the reconstruction aligns strictly with the physical reality, we propose a mathematical discretization strategy.

**Spatial Partitioning (Rolling Shutter).** We divide the detector plane along the row dimension into $M$ horizontal strips. For the $j$-th strip, we define a local virtual camera with a shifted principal point to maintain geometric consistency. The intrinsic matrix $\mathbf{K}_j$ is modified such that its optical center $c_{y,j}$ aligns with the physical center of the $j$-th strip: $c_{y,j} = c_y - j \cdot (H/M)$. The exposure start time of strip $j$ is $t^f_j = t^f_{start} + \frac{j+0.5}{M} T_{ro}$, where 0.5 denotes the center of strip.

**Temporal Segmentation (Motion Blur).** We approximate the motion blur integral by subdividing the exposure $T_{exp}$ into $N$ discrete sub-intervals.

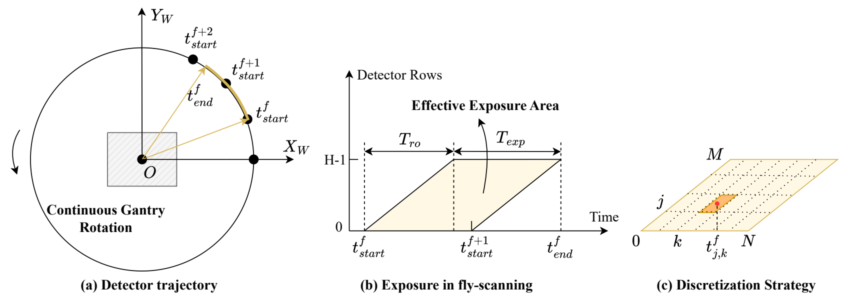


Fig. 3: **Physics of Continuous Acquisition and Discretization Strategy**. (a) The detector executes a continuous circular trajectory, where the exposure duration of frame $f$ is highlighted. (b) The effective exposure forms a parallelogram (yellow shaded area) defined by $T_{ro}$ and $T_{exp}$. (c) The prallelgoram is discretized into $M$ spatial strips and $N$ temporal sub-intervals. The red dot denotes sampling time $t^f_{j,k}$.

Crucially, unlike random sampling, we employ stratified kinematic sampling. As shown in Fig. 3c, for the $k$-th temporal sub-interval of the $j$-th spatial strip in frame $f$, the exposure time is:

$$t^f_{j,k} = t^f_j + \frac{(k+0.5)}{N} T_{exp} \tag{2}$$

We then query the calibrated kinematics function $\boldsymbol{\theta}(t)$ of the scanner to obtain the precise pose $\mathbf{W}_{j,k} = \boldsymbol{\theta}(t^f_{j,k})$. This discretization converts the continuous physical motion into a set of precise, discrete micro-states.

### 2.3 Efficient Multi-Pose Batch Rendering

Simulating physics via $M \times N$ micro-views per projection could imply a massive computational burden. To achieve high fidelity with negligible overhead, we implement a multi-pose batch rendering kernel for X-ray physics based on Gsplat [26]. Instead of looping through states, we construct a "super-batch" of $M \times N$ virtual cameras for each physical projection. The rasterizer processes these micro-cameras concurrently, yielding a tensor of micro-patches $\hat{I}^f_{j,k}$. The final physical projection is synthesized by:

$$\hat{I}^f = \text{Concat}_j \left( \frac{1}{N} \sum_{k=0}^{N-1} \hat{I}^f_{j,k} \right) \tag{3}$$

This design allows PhyRadGS to forward-model complex dynamics purely through parallelized rasterization, avoiding the computational cost of temporal supersampling iteratively.

### 2.4 Objective Function

We optimize the Gaussian parameters by minimizing the discrepancy between the synthesized $\hat{I}$ and the ground truth $I_{gt}$, regularized by structural priors. The total loss $\mathcal{L}$ is formulated as:

$$\mathcal{L} = \|\hat{I} - I_{gt}\|_1 + \lambda_{ssim}\text{SSIM}(\hat{I}, I_{gt}) + \lambda_{tv}\mathcal{L}_{TV} + \lambda_{reg}\mathcal{L}_{reg} \tag{4}$$

where $\mathcal{L}_{TV}$ suppresses noise in the density field, following [28]. $\mathcal{L}_{reg}$ imposes an L1 regularization for the gaussian density to eliminate floating artifacts in empty regions.

## 3 Experiments

### 3.1 Experimental Setup

**Datasets.** We conducted experiments on both synthetic and real-world data to validate performance under controlled and realistic conditions.

(1) Synthetic Flying Dataset: We utilize 15 diverse anatomical volumes (e.g., Chest, Foot, Jaw) from the collection of R2-Gaussian [28]. Unlike previous works that simulate static projections, we employed *TIGRE* toolbox [2] to simulate fly-scanning data with Poisson noise. We simulate two protocols: Rapid-Scan (25 views, 14.4° integration arc) and Standard-Scan (50 views, 7.2° integration arc), with rolling shutter ratio $T_{ro} = 0.86T_{exp}$, resolution is $512 \times 512$.

(2) Real Phantom Data: We acquired real projection data of a double-layer sphere phantom using a commercial CBCT system, and designed a paired study to isolate motion artifacts. The projection resolution is $2048 \times 2048$. **Pseudo-GT:** The reference volume, reconstructed using high-quality R2-Gaussian [28] with 64 "stop-and-shoot" projections. **Fly-Scanning Input:** The 25 fly-scanning projections over 360°. These projections inherently contain realistic mechanical motion blur and rolling shutter distortions.

**Baselines.** We compare PhyRadGS against three categories of methods, all relying on the static pose assumption: (1) *Analytic/Iterative*: FDK [11] and TV-regularized ASD-POCS [22]. (2) *Implicit Neural Fields*: NAF [29] and SAX-NeRF [4]. (3) *Explicit Splatting*: R2-Gaussian [28], static current SOTA.

**Implementation Details.** Our framework is implemented in PyTorch. The discretization parameters are set to $M = 8$ (spatial strips) and $N = 5$ (temporal sub-intervals). We train the model for 20k iterations using the Adam optimizer for GS methods (R2-Gaussian and our PhyRadGS). The loss weights are set to $\lambda_{ssim} = 0.2$ for SSIM, $\lambda_{tv} = 0.01$, and $\lambda_{reg} = 0.001$. All experiments run on a single NVIDIA RTX 4090 GPU.

### 3.2 Results and Analysis

**Quantitative Performance.** Table 1 reports the quantitative metrics. Our PhyRadGS achieves significant accuracy improvements on both synthetic flying

Table 1: Quantitative comparison. **ST** signifies our Spatio-Temporal module. The top result is **bolded**, and the next two best are <u>underlined</u>. Note that FDK is an analytic method with negligible runtime. Time is measured in **minutes**.

| Method | **Syn 25-view Scan** | | | **Syn 50-view Scan** | | | **Real Sphere-Phantom** | | |
|---|---|---|---|---|---|---|---|---|---|
| | PSNR ↑ | SSIM ↑ | Time ↓ | PSNR ↑ | SSIM ↑ | Time ↓ | PSNR ↑ | SSIM ↑ | Time ↓ |
| FDK [11] | 18.01 | 0.090 | - | 20.72 | 0.152 | - | 30.04 | 0.794 | - |
| ASD-POCS [22] | 27.99 | 0.777 | **0.38** | 30.94 | 0.834 | **0.73** | 30.29 | 0.844 | <u>9.80</u> |
| NAF [29] | 28.18 | 0.776 | 11.56 | 30.96 | 0.844 | 12.30 | 27.52 | 0.779 | 163.60 |
| SAX-NeRF [4] | 28.02 | 0.773 | 61.39 | 30.62 | 0.849 | 115.84 | - | - | - |
| R2-Gaussian [28] | 28.93 | 0.828 | 7.95 | <u>32.12</u> | 0.889 | 7.99 | 35.12 | 0.852 | 18.83 |
| PhyRadGS **w/o ST** | <u>29.00</u> | <u>0.834</u> | <u>3.13</u> | 31.96 | <u>0.891</u> | <u>3.31</u> | <u>35.13</u> | <u>0.855</u> | **4.68** |
| R2-Gaussian **w/ ST** | <u>31.22</u> | <u>0.866</u> | 19.37 | **33.97** | <u>0.909</u> | 20.15 | <u>36.84</u> | <u>0.861</u> | 50.78 |
| **PhyRadGS (Ours)** | **31.23** | **0.871** | <u>5.16</u> | <u>33.54</u> | **0.910** | <u>5.31</u> | **37.43** | **0.862** | <u>9.15</u> |

dataset and real phantom data. On the challenging 25-view synthetic flying dataset, static baselines suffer severe degradation. PhyRadGS outperforms the SOTA static method (R2-Gaussian) by a wide margin, improving PSNR by $+$**2.3 dB** and SSIM by $+$**0.043**. Critically, we perform a controlled ablation by integrating our Spatio-Temporal (ST) module into R2-Gaussian. Although "R2-Gaussian w/ ST" improves quality, its training time increases significantly ($\sim$19 min). In contrast, thanks to our parallelized batch rendering, PhyRadGS incurs negligible overhead ($\sim$2 min) and remains faster than the static baseline ($\sim$5 min vs. $\sim$8 min). This efficiency advantage is further amplified on high-resolution ($2048^2$) real data: PhyRadGS requires only $\sim$9 min, while R2-Gaussian w/ ST balloons to $\sim$51 min and NAF explodes to $\sim$163 min.

**Visual Analysis.** Fig. 4 presents qualitative comparisons. On synthetic anatomy dataset (Foot), PhyRadGS explicitly disentangles the motion integration, recovering sharp cortical bone edges and fine marrow details. In contrast, baselines exhibit softened boundaries and streaking artifacts. On real phantom data (Sphere), PhyRadGS physically models both motion blur and rolling shutter, producing clear spheres with sharp boundaries. By comparison, FDK and ASD-POCS yield blurred spheres that deviate from circularity; NAF suffers from severe shape deformation; and R2-Gaussian exhibits pronounced motion drag artifacts.

**Ablation Study.** Table 2 analyzes the sensitivity to discretization parameters on the 25-view synthetic flying dataset. **Rolling Shutter ($M$):** With fixed $N = 5$, increasing the spatial strips to $M = 8$ yields a substantial gain of $+$**1.43 dB** (29.79 $\rightarrow$ 31.22 dB), confirming the necessity of geometric correction. We select $M = 8$ as the sweet spot, since $M = 16$ incurs $\sim$33% extra overhead for the marginal gain (+0.31 dB). **Motion Blur ($N$):** As the dominant error source, increasing the temporal samples from 1 to 3 provides a sharp boost of

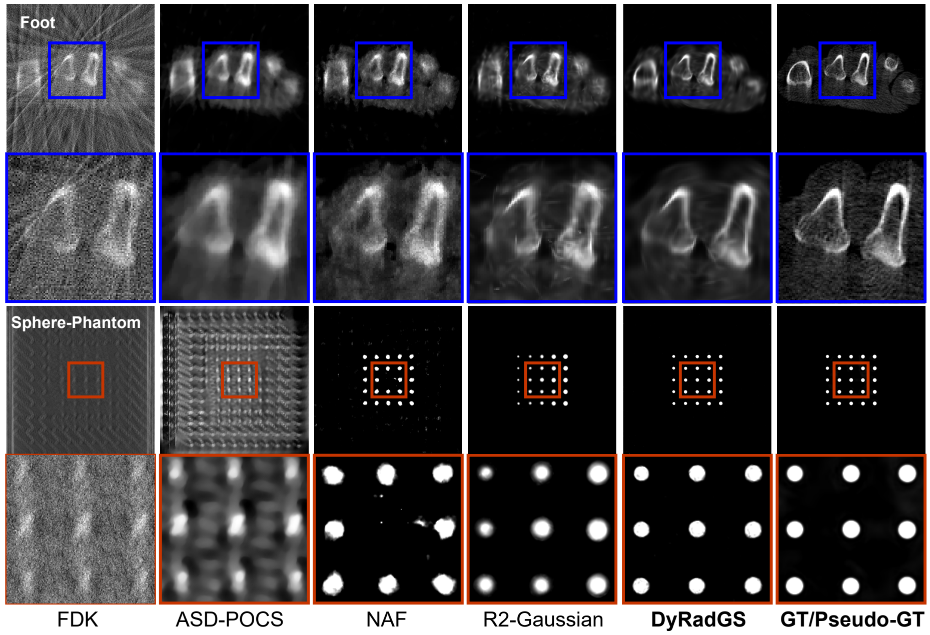


Fig. 4: **Reconstruction Slice Comparison (25 fly-scanning views).** Top Two Rows: Synthetic 'Foot' data. Bottom Two Rows: Real 'Sphere Phantom' data. Static methods fail to resolve mechanical motion, resulting in severe blurring and geometric distortion, while PhyRadGS recovers true geometry.

**+1.89 dB**. Since performance saturates at $N = 5$, we adopt the configuration of $M = 8, N = 5$ to balance high fidelity with computational efficiency.

## 4 Conclusion

PhyRadGS introduces a physics-aware paradigm to 3DGS-based CT reconstruction, explicitly modeling the motion blur and rolling shutter effects inherent in fly-scanning. By enforcing known scanner kinematics via a novel spatio-temporal discretization, our method resolves the ambiguities of static assumptions and recovers high-frequency anatomical details. Furthermore, our multi-pose batch rendering pipeline ensures computational efficiency, enabling rapid reconstruction without compromising fidelity. Experiments on synthetic and real-world data demonstrate that PhyRadGS significantly outperforms state-of-the-art static baselines, restoring sharp tissue boundaries and correcting geometric distortions.

Currently, our method assumes rigid motion. While this covers gantry rotation, it does not account for *non-rigid patient motion* (e.g., breathing or heartbeat). Future work will extend PhyRadGS to model deformable scenes by incorporating time-dependent deformation fields (4D-GS) alongside the acquisition physics, and will further investigate robustness in the presence of mechanical

Table 2: Ablation study for the effect of $M$ (left, with fixed $N = 5$) and $N$ (right, with fixed $M = 8$). Time is measured in minutes. Best results are **bolded**.

| $M$ | PSNR ↑ | SSIM ↑ | Time ↓ | $N$ | PSNR ↑ | SSIM ↑ | Time ↓ |
|---|---|---|---|---|---|---|---|
| 1 | 29.79 | 0.851 | **3.98** | 1 | 29.27 | 0.848 | **3.14** |
| 4 | 30.62 | 0.869 | 4.52 | 3 | 31.16 | 0.870 | 3.98 |
| 8 | 31.22 | 0.871 | 5.16 | 5 | **31.22** | 0.871 | 5.16 |
| 16 | **31.53** | **0.884** | 6.84 | 7 | 31.21 | **0.873** | 6.18 |

jitter or more complex motion patterns. In-vivo clinical data is also planned for future work.

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

## References

1. Andersen, A.H., Kak, A.C.: Simultaneous algebraic reconstruction technique (sart): a superior implementation of the art algorithm. Ultrasonic imaging **6**(1), 81–94 (1984)
2. Biguri, A., Sadakane, T., Lindroos, R., Liu, Y., Landman, M.S., 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**(1), 015011 (2025)
3. Cai, Y., Liang, Y., Wang, J., Wang, A., Zhang, Y., Yang, X., Zhou, Z., Yuille, A.: Radiative gaussian splatting for efficient x-ray novel view synthesis. In: European Conference on Computer Vision. pp. 283–299. Springer (2024)
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. pp. 11174–11183 (2024)
5. Cant, J., Palenstijn, W.J., Behiels, G., Sijbers, J.: Modeling blurring effects due to continuous gantry rotation: Application to region of interest tomography. Medical Physics **42**(5), 2709–2717 (2015)
6. Chang, M., Xiao, Y., Chen, Z., Li, L., Zhang, L.: Preliminary study of rotary motion blurs in a novel industry ct imaging system. In: 2011 IEEE Nuclear Science Symposium Conference Record. pp. 1358–1361 (2011)
7. Chen, D., Li, H., Wang, Q., Zhang, P., Zhu, Y.: Computed tomography for high-speed rotation object. Optics express **23**(10), 13423–13442 (2015)
8. Chen, J., Zhang, X., Hoque, M.I., Vasconcelos, F., Stoyanov, D., Elson, D.S., Huang, B.: Surgicalgs: Dynamic 3d gaussian splatting for accurate robotic-assisted surgical scene reconstruction. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 572–582. Springer (2025)
9. Chen, W., Liu, L.: Deblur-gs: 3d gaussian splatting from camera motion blurred images. Proceedings of the ACM on Computer Graphics and Interactive Techniques **7**(1), 1–15 (2024)
10. Ching, D., Aslan, S., Nikitin, V., Gürsoy, D.: Time-coded aperture for x-ray imaging. Optics Letters **44**(11), 2803–2806 (2019)

11. Feldkamp, L.A., Davis, L.C., Kress, J.W.: Practical cone-beam algorithm. Journal of the Optical Society of America A **1**(6), 612–619 (1984)
12. Huang, Y., Bai, L., Cui, B., Li, Y., Chen, T., Wang, J., Wu, J., Lei, Z., Liu, H., Ren, H.: Endo-4dgx: Robust endoscopic scene reconstruction and illumination correction with gaussian splatting. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 181–191. Springer (2025)
13. Huang, Y., Singh, I., Joyce, T., Thielemans, K., McClelland, J.R.: Digs: Dynamic cbct reconstruction using deformation-informed 4d gaussian splatting and a low-rank free-form deformation model. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 131–141. Springer (2025)
14. Huyge, B., Renders, J., Sanctorum, J.G., De Beenhouwer, J., Sijbers, J.: X-ray image reconstruction for continuous acquisitions with a generalized motion model. Optics Express **32**(22), 39192–39207 (2024)
15. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G., et al.: 3d gaussian splatting for real-time radiance field rendering. ACM Trans. Graph. **42**(4), 139–1 (2023)
16. Li, Y., Fu, X., Li, H., Zhao, S., Jin, R., Zhou, S.K.: 3dgr-ct: Sparse-view ct reconstruction with a 3d gaussian representation. Medical Image Analysis **103**, 103585 (2025)
17. Lin, Y., Wang, H., Chen, J., Li, X.: Learning 3d gaussians for extremely sparse-view cone-beam ct reconstruction. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 425–435. Springer (2024)
18. Liu, H., Liu, Y., Li, C., Li, W., Yuan, Y.: Lgs: A light-weight 4d gaussian splatting for efficient surgical scene reconstruction. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 660–670. Springer (2024)
19. Lu, Y., Zhou, Y., Liu, D., Liang, T., Yin, Y.: Bard-gs: Blur-aware reconstruction of dynamic scenes via gaussian splatting. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 16532–16542 (2025)
20. Majee, S., Aslan, S., Gürsoy, D., Bouman, C.A.: Codex: A modular framework for joint temporal de-blurring and tomographic reconstruction. IEEE Transactions on Computational Imaging **8**, 666–678 (2022)
21. Seiskari, O., Ylilammi, J., Kaatrasalo, V., Rantalankila, P., Turkulainen, M., Kannala, J., Rahtu, E., Solin, A.: Gaussian splatting on the move: Blur and rolling shutter compensation for natural camera motion. In: European conference on computer vision. pp. 160–177. Springer (2024)
22. Sidky, E.Y., Pan, X.: Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. Physics in Medicine & Biology **53**(17), 4777 (2008)
23. Steven Tilley, I., Sisniega, A., Siewerdsen, J.H., Stayman, J.W.: High-fidelity modeling of detector lag and gantry motion in ct reconstruction. In: Conference proceedings. International Conference on Image Formation in X-Ray Computed Tomography. vol. 2018, p. 318 (2018)
24. 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. pp. 20310–20320 (2024)
25. Yang, Z., Gao, X., Zhou, W., Jiao, S., Zhang, Y., Jin, X.: Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 20331–20341 (2024)

26. Ye, V., Li, R., Kerr, J., Turkulainen, M., Yi, B., Pan, Z., Seiskari, O., Ye, J., Hu, J., Tancik, M., et al.: gsplat: An open-source library for gaussian splatting. Journal of Machine Learning Research **26**(34), 1–17 (2025)
27. Yu, W., Cai, Y., Zha, R., Fan, Z., Li, C., Yuan, Y.: X2-gaussian: 4d radiative gaussian splatting for continuous-time tomographic reconstruction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 24728–24738 (2025)
28. Zha, R., Lin, T.J., Cai, Y., Cao, J., Zhang, Y., Li, H.: $R^2$-gaussian: Rectifying radiative gaussian splatting for tomographic reconstruction. In: Proceedings of the 38th International Conference on Neural Information Processing Systems. pp. 44907–44934 (2024)
29. 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)