

# DIGS: Dynamic CBCT Reconstruction using Deformation-Informed 4D Gaussian Splatting and a Low-Rank Free-Form Deformation Model

Yuliang Huang<sup>1(✉)</sup>, Imraj Singh<sup>2</sup>, Thomas Joyce<sup>3</sup>, Kris Thielemans<sup>1,4</sup>, and Jamie R. McClelland<sup>1</sup>

<sup>1</sup> UCL Hawkes Institute, Department of Medical Physics and Biomedical Engineering, University College London, United Kingdom  
yuliang.huang.21@ucl.ac.uk

<sup>2</sup> Department of Computer Science, University College London, United Kingdom

<sup>3</sup> Elekta, AB, Stockholm, Sweden

<sup>4</sup> Institute of Nuclear Medicine, University College London, United Kingdom

**Abstract.** 3D Cone-Beam CT (CBCT) is widely used in radiotherapy but suffers from motion artifacts due to breathing. A common clinical approach mitigates this by sorting projections into respiratory phases and reconstructing images per phase, but this does not account for breathing variability. Dynamic CBCT instead reconstructs images at each projection, capturing continuous motion without phase sorting. Recent advancements in 4D Gaussian Splatting (4DGS) offer powerful tools for modeling dynamic scenes, yet their application to dynamic CBCT remains underexplored. Existing 4DGS methods, such as HexPlane, use implicit motion representations, which are computationally expensive. While explicit low-rank motion models have been proposed, they lack spatial regularization, leading to inconsistencies in Gaussian motion. To address these limitations, we introduce a free-form deformation (FFD)-based spatial basis function and a deformation-informed framework that enforces consistency by coupling the temporal evolution of Gaussian’s mean position, scale, and rotation under a unified deformation field. We evaluate our approach on six CBCT datasets, demonstrating superior image quality with a  $6\times$  speedup over HexPlane. These results highlight the potential of deformation-informed 4DGS for efficient, motion-compensated CBCT reconstruction. The code is available at <https://github.com/Yuliang-Huang/DIGS>.

**Keywords:** Dynamic CBCT · 4D Gaussian Splatting · Deformation Informed.

## 1 Introduction

Cone-beam CT (CBCT) is widely used in image-guided radiotherapy [5] but is prone to artifacts caused by patients’ involuntary motion, such as respiration [25]. A common clinical approach to address this issue is 4DCBCT, which bins

projections into different respiration phases and reconstructs a separate image for each phase [24]. However, 4DCBCT requires extra acquisition time and radiation dose, and is susceptible to artifacts caused by phase sorting errors [15]. More critically, it assumes that patients’ breathing motion is regular, which is often not the case [10]. To account for irregular motion, dynamic CBCT reconstruction aims to estimate both the motion and the image on a per-projection basis [11,12], where a fast motion-compensated reconstruction algorithm is needed to satisfy the demands of clinical feasibility.

Traditional reconstruction methods typically rely on a discrete voxel grid representation. In recent years, implicit neural representations (INR) have gained attention for their inherent continuity and resolution independence [29]. However, INR suffers from low computation efficiency because each query point needs to be evaluated by a multilayer perceptron (MLP). More recently, 3D Gaussian Splatting (3DGS) has emerged as a promising alternative [13], representing the image volume as a Gaussian point cloud and enabling fast rendering through efficient splatting techniques [30]. While 3DGS has been applied to static CBCT [28], applying 4D Gaussian Splatting (4DGS) for dynamic CBCT reconstruction remains underexplored, which is the primary objective of this paper.

To the best of our knowledge, only one concurrent work [9] applies 4DGS for 4DCBCT reconstruction, which still requires phase sorting. In this work, we introduce several methodological innovations upon existing approaches:


1. Most current 4DGS methods, including those in non-medical applications, model deformation with implicit representations, such as Neural Field [19] or HexPlane [4], which incurs significant computational overhead due to MLP inference. In contrast, we propose a more efficient explicit motion representation based on Free-Form Deformation (FFD).

2. Existing 4DGS methods typically fit motion by independently modifying the Gaussian kernel’s mean positions, scales and rotations. Although these methods can produce plausible dynamic images, they lack an interpretable spatial correspondence between time frames. This is a critical limitation in applications such as tumor tracking in radiotherapy, where the voxel-wise deformation vector field (DVF) itself is as important as the reconstructed images. To address this, we propose a deformation-informed (DI) framework that consistently updates all Gaussian attributes according to a unified DVF. Figure 1 provides a schematic of the proposed method.

## 2 Related Work

### 2.1 Explicit Motion Representation

Several studies have used explicit representation for motion fields in 4DGS. For example, Gaussian-flow [17] parameterizes the motion of each Gaussian as polynomial curves and Fourier series. Other studies [14,3] factorize the motion field into a linear combination of a few learnable basis. Yang et al. [27] also use a low-rank motion model but constrain the temporal basis to be Gaussian functions.



**Fig. 1.** Schematic illustration of the proposed method. A continuous free-form deformation (FFD) is constructed from a set of learnable spatial and temporal basis, which is used to warp the reference Gaussians through a deformation-informed (DI) framework. Projections rendered from the dynamic Gaussians are compared with the measured projections and the optimiser updates the reference Gaussians and the motion basis to minimize their difference.

These methods fit the weights of motion basis for each Gaussian independently, without imposing any smoothness constraint on the motion of adjacent Gaussians. In comparison, our proposed method interpolates the basis weights from a free-form deformation field [18,22] that enforces adjacent Gaussians to move consistently with each other.

## 2.2 Applying Deformation Fields to Gaussian Kernels

Most 4DGS methods learn the temporal changes of the Gaussian’s mean position, scale and rotation independently. In contrast, inspired by continuum mechanics [2], physics-based 3DGS simulates dynamic scenes by explicitly applying continuous DVF to Gaussians, by deriving the transformation of scale and rotation from the Jacobian of mean position deformation [26,16,8]. While these previous works focus on applying synthetic deformation to static objects, we are the first to apply a physics-based DI framework for handling motion in image reconstruction problems.

## 3 Methodology

### 3.1 Preliminary: 3DGS for CBCT reconstruction

Given a set of measured projections  $\{P_t\}_{t=1,\dots,N_t}$  for  $N_t$  time points, the problem of CBCT reconstruction can be formulated as an inverse problem:

$$\min_{\phi_{\mathbf{x}}} \sum_{t=1}^{N_t} L(\mathbb{F}(\phi_{\mathbf{x}}), P_t) \quad (1)$$

where  $\mathbb{F}$  is the X-ray transform,  $\phi_{\mathbf{x}}$  is the attenuation field, and  $L$  is the dissimilarity between re-projection and measured data.

Defining a set of Gaussian kernels  $\{G_n\}_{n=1,\dots,N_g}$  as

$$G_n(\mathbf{x}|\rho_n, \boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n) = \rho_n e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu}_n)^T \boldsymbol{\Sigma}_n^{-1}(\mathbf{x}-\boldsymbol{\mu}_n)} \quad (2)$$

where  $\rho_n$  is the density,  $\boldsymbol{\mu}_n$  is the mean position,  $\boldsymbol{\Sigma}_n = \mathbf{R}_n \mathbf{S}_n \mathbf{S}_n^T \mathbf{R}_n^T$  is the covariance matrix that can be decomposed into scale  $\mathbf{S}_n$  and rotation  $\mathbf{R}_n$ .  $\mathbf{S}_n$  is a diagonal matrix and  $\mathbf{R}_n$  is an orthonormal matrix parameterized by quaternion. The attenuation field can be represented as the superposition of these Gaussian kernels:

$$\phi_{\mathbf{x}} = \sum_{n=1}^{N_g} G_n(\mathbf{x}|\rho_n, \boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n) \quad (3)$$

Cone-beam projection of a 3D Gaussian onto a 2D plane can be approximated by a 2D Gaussian:

$$\tilde{G}_n = \tilde{\rho}_n e^{-\frac{1}{2}(\tilde{\mathbf{x}}-\tilde{\boldsymbol{\mu}}_n)^T \tilde{\boldsymbol{\Sigma}}_n^{-1}(\tilde{\mathbf{x}}-\tilde{\boldsymbol{\mu}}_n)} \quad (4)$$

The tilde over the letters refers to the 2D counterpart of the variable, i.e.  $\tilde{\mathbf{x}}$  and  $\tilde{\boldsymbol{\mu}}_n$  are the projections of  $\mathbf{x}$  and  $\boldsymbol{\mu}_n$ .  $\tilde{\boldsymbol{\Sigma}}_n$  can be obtained by dropping the third row and column of  $\mathbf{J}_n \mathbf{W} \boldsymbol{\Sigma}_n \mathbf{W}^T \mathbf{J}_n^T$  where  $\mathbf{W}$  is the projection-angle dependent viewing transform and  $\mathbf{J}_n$  is the local Jacobian of the perspective transform.  $\tilde{\rho}_n$  is  $\rho_n$  scaled by the integration bias factor proposed in [28]. The X-ray transform of  $\phi_{\mathbf{x}}$  can then be approximated by the superposition of the 2D Gaussians, i.e.  $\mathbb{F}(\phi_{\mathbf{x}}) \simeq \sum_{n=1}^{N_g} \tilde{G}_n(\tilde{\mathbf{x}})$ . A gradient-based method is then used to optimize Eqn(1) to fit the attributes of each Gaussian, i.e.  $\{\rho_n, \boldsymbol{\mu}_n, \mathbf{R}_n, \mathbf{S}_n\}_{n=1,\dots,N_g}$ .

### 3.2 Deformation-Informed 4DGS

4DGS extends 3DGS to dynamic scenes by allowing the attributes of the Gaussians to vary over time. Many 4DGS methods use separate models for estimating the change of each attribute, i.e.

$$\mathbf{v}_{n,t} = \mathbf{v}_{n,0} + f_{\mathbf{v}_n}(t) \quad (5)$$

where  $\mathbf{v}_{n,t}$  refers to the values of any attribute in  $\{\rho_n, \boldsymbol{\mu}_n, \mathbf{R}_n, \mathbf{S}_n\}_{n=1,\dots,N_g}$  at time  $t$ , and  $f_{\mathbf{v}_n}(t)$  is an attribute-specific function that gives the change of  $\mathbf{v}_{n,t}$  against  $\mathbf{v}_{i,0}$  at time  $t$ .

From this, the dynamic images can be generated by

$$\phi_{\mathbf{x}}(t) = \sum_{n=1}^{N_g} G_n(\mathbf{x}|\rho_{n,t}, \boldsymbol{\mu}_{n,t}, \boldsymbol{\Sigma}_{n,t}), \quad t = 1, \dots, N_t. \quad (6)$$

However, if changes of Gaussian parameters are modeled independently, it is impossible to derive a continuous DVF that is consistent with the rendered dynamic images, i.e. one cannot define a transformation that maps a voxel-grid representation of one dynamic image to another. To see this, consider two initially overlapping Gaussian that move apart over time, there is no spatial

transformation that can map the overlapping Gaussians to the two separate Gaussians.

Instead of allowing each attribute to evolve independently, we follow a physics-based 3DGS approach [26] and warp the Gaussian kernels using a unified DVF. More specifically, the mean positions are shifted directly with the DVF  $D$ , i.e.

$$\boldsymbol{\mu}_{n,t} = D(\boldsymbol{\mu}_{n,0}, t) \quad (7)$$

and the covariance matrices are derived by approximating  $D$  by a locally linear deformation, i.e.

$$\boldsymbol{\Sigma}_{n,t} \simeq \mathbf{K} \boldsymbol{\Sigma}_{n,0} \mathbf{K}^T, \quad \mathbf{K} = \nabla_{\boldsymbol{\mu}} D(\boldsymbol{\mu}_{n,0}, t) \quad (8)$$

and  $\rho_n$  remains constant over time. This ensures that all attributes of the Gaussian kernels are updated consistently by the same DVF.

### 3.3 Explicit Low-Rank B-Spline Motion Model

The DVF can be decomposed into a linear combination of a few spatial and temporal basis functions,  $\mathbf{u}_r(\mathbf{x})$  and  $\omega_r(t)$ , i.e.

$$D(\mathbf{x}, t) = \mathbf{x} + \sum_{r=1}^{N_r} \omega_r(t) \mathbf{u}_r(\mathbf{x}) \quad (9)$$

where  $N_r$  is the number of basis functions, much smaller than  $N_g$  and  $N_t$ .

Each spatial basis  $\mathbf{u}_r(\mathbf{x})$  is explicitly parameterized by a set of control points arranged on a uniform  $N_x \times N_y \times N_z$  lattice with predefined spacing, and each control point stores a displacement vector  $\mathbf{d}_{ijk}^r$ , where  $i \in \{1, \dots, N_x\}$ ,  $j \in \{1, \dots, N_y\}$ ,  $k \in \{1, \dots, N_z\}$ ,  $r \in \{1, \dots, N_r\}$ . Let  $\mathbf{x}$  be a point in world space and  $\mathbf{p} = (p_1, p_2, p_3)$  its corresponding coordinates in the lattice index space.  $\mathbf{u}_r(\mathbf{x})$  is interpolated from the  $4 \times 4 \times 4$  control points surrounding  $\mathbf{x}$ :

$$\mathbf{u}_r(\mathbf{x}) = \sum_{l=0}^3 \sum_{m=0}^3 \sum_{n=0}^3 B_l(p_1 - \lfloor p_1 \rfloor) B_m(p_2 - \lfloor p_2 \rfloor) B_n(p_3 - \lfloor p_3 \rfloor) \mathbf{d}_{i,j,k}^r \quad (10)$$

where  $i = \lfloor p_1 \rfloor - 1 + l$ ,  $j = \lfloor p_2 \rfloor - 1 + m$ ,  $k = \lfloor p_3 \rfloor - 1 + n$ ,  $\lfloor \cdot \rfloor$  is the floor function and  $B_l$  is the  $l^{th}$  cubic B-spline basis function [22]. The Jacobian is given by

$$\mathbf{K}(\mathbf{x}, t) = \mathbf{I} + \sum_{r=1}^{N_r} \omega_r(t) \nabla \mathbf{u}_r(\mathbf{x}) \quad (11)$$

$\omega_r(t)$  is also parameterized by a 1D B-Spline. A customized CUDA kernel is implemented to efficiently perform differentiable cubic B-spline interpolation.

## 4 Experiment Settings

### 4.1 Datasets

**XCAT Simulation** The 4DXCAT software [23] was used to generate digital phantom images with dimension of  $375 \times 375 \times 345$  and resolution of 1 mm as well as simulated motion controlled by breathing traces. The raw DVFs were post-processed by cidX software [6] to avoid folding and preserve sliding motion. RTK [20] was then used to generate 310 projections in a full circle with dimension of  $512 \times 512$  and resolution of 0.8 mm. The detector was offset by 116 mm. Distances from source to isocenter and source to detector are 1000 mm and 1536 mm respectively. This is the same as the acquisition geometry of a real CBCT scan. Poisson noise ( $\lambda = 10^8$ ) and Gaussian noise ( $\sigma = 4$ ) was added to the projections to simulate quantum and electronic noise. Two sets of breathing traces were measured from cine MR images of real patients and used as input to XCAT to create two datasets, denoted as case 1 and 2 respectively.

**Simulation from 4DCT patient images** Simulated data were generated using clinical 4DCT images from four patients in the 4DLung dataset [1]. All phases were registered to the end-exhalation phase using NiftyReg [18], with default parameters except for the use of sliding registration [7]. Principal Component Analysis (PCA) was then applied to the stack of deformation fields, preserving the mean deformation and the two dominant principal components. To simulate 20 breathing cycles, the fitted principal component weights were replicated while introducing a slow drift. Additionally, the motion magnitude of each breathing cycle was scaled by a random factor between 0.8 and 1.2 to simulate variability in breathing motion. 310 time points were then uniformly sampled from the 20 breathing cycles to generate CBCT projections, following the same procedure as for XCAT data. The four simulations are denoted as case 3-6 respectively.

### 4.2 Implementation Details

The Gaussian point cloud is initialized by sampling 80K points from conjugate gradient least square reconstruction, in line with [9] for fair comparison. The B-Spline control points have a spacing of every 4 time points and 8 voxels.  $N_r$  is set to 2. L2 loss in projection space is used for training. The reference Gaussians and the space-time B-Spline grids are jointly updated with Adam optimizer for 50K iterations. Initial learning rates for the space and time B-Spline grids are set to  $10^{-4}$  and  $10^{-2}$  respectively, and decay linearly to  $10^{-5}$  and  $10^{-3}$  by the final iteration. The gradient threshold for density control is  $5 \times 10^{-8}$ . All the other hyper-parameters are the same as in [28].

### 4.3 Evaluation

Our proposed method is compared with two other methods. SuPreMo [11] fits a linear model between the DVF and low-dimensional surrogate signals extracted

from projections, and uses motion compensated FDK [21] as the reconstruction algorithm. HexPlane-based 4DGS method [9] encodes spatial-temporal features with six 2D grids which are then fed into multi-head MLP to obtain the time-varying attributes per Gaussian kernel. The original method has only been applied on phase-sorted data, where projections of the same phase are treated as the same time-point. To compare with the proposed method, each projection is treated as an individual time-point.

We also perform two ablation studies. To investigate the impact of ensuring that all Gaussian attributes change consistently with the underlying FFD, Ablation Study 1 allows each Gaussian attribute to vary independently over time. Building on this, Ablation Study 2 further removes the FFD representation entirely; that is, the spatial basis is no longer interpolated from the control point grid but instead stored as a separate attribute for each Gaussian.


All the experiments are run on a single RTX3090 GPU. The results are evaluated by the average Root-Mean-Squared-Error (RMSE) and Peak-Signal-Noise-Ratio (PSNR) between the estimated dynamic images at each time-point and the ground-truth. The ground truth is created by warping reference images with the simulated DVFs. All evaluations are within the reconstruction FOV.

## 5 Results

Table 1 lists the average PSNR and RMSE of dynamic reconstruction at all time points obtained by different methods and the corresponding computation time. While SuPreMo [11] achieves the fastest speed, it yields the lowest image quality. Compared with HexPlane [9], our proposed method is about six times faster while achieving better image quality. Ablation Study 1 leads to a slight improve-

**Table 1.** Performance of proposed method on simulation data in comparison with SuPreMo [11] and HexPlane [9] as well as ablation studies excluding deformation-informed (DI) framework and Free-Form Deformation (FFD). Best results are in bold.

Case	PSNR (dB) $\uparrow$			RMSE ( $10^{-3} \text{ mm}^{-1}$ ) $\downarrow$			Time $\downarrow$		
	SuPreMo	HexPlane	Ours	SuPreMo	HexPlane	Ours	SuPreMo	HexPlane	Ours
1	18.49	21.98	<b>25.46</b>	1.63	1.91	<b>1.28</b>	<b>47s</b>	1h31m	16m14s
2	18.26	21.15	<b>21.78</b>	<b>1.88</b>	2.12	1.97	<b>1m</b>	1h27m	15m48s
3	21.46	26.82	<b>29.48</b>	5.08	2.78	<b>2.02</b>	<b>41s</b>	1h43m	14m34s
4	24.79	28.48	<b>30.94</b>	3.46	2.29	<b>1.71</b>	<b>1m2s</b>	1h49m	14m1s
5	24.66	29.35	<b>31.66</b>	3.51	2.10	<b>1.58</b>	<b>45s</b>	1h38m	16m14s
6	20.79	27.36	<b>29.49</b>	5.39	2.68	<b>1.99</b>	<b>48s</b>	1h50m	16m26s
Mean	21.41	25.86	<b>28.14</b>	3.49	2.31	<b>1.76</b>	<b>51s</b>	1h40m	15m32s
Ablation Studies									
	Ablation1	Ablation2	Full	Ablation1	Ablation2	Full	Ablation1	Ablation2	Full
Mean	<b>28.25</b>	27.29	28.14	<b>1.73</b>	1.95	1.76	17m1s	<b>14m50s</b>	15m32s



**Fig. 2.** Ground-truth and reconstructed mean position images for case 2 (top), case 3 (middle) and case 5 (bottom). Dynamic range:  $0.00 - 0.03 \text{ mm}^{-1}$

ment, but remains comparable to results of using a unified DVF. Ablation Study 2 results in a noticeable performance drop.

Figure 2 plots example coronal slices of the reconstructed mean position images for three cases with the lowest, median and highest PSNR values, respectively. Compared with SuPreMo and HexPlane, the proposed method produces sharper diaphragm with less noise and motion artifacts.

## 6 Discussion and Conclusion

Our results demonstrate that the proposed method significantly improves both computational efficiency and image quality in dynamic CBCT reconstruction. Compared to HexPlane, our approach achieves a  $6\times$  speedup while producing sharper images with fewer motion artifacts. While SuPreMo is the fastest, it results in the lowest image quality, highlighting the trade-off between speed and accuracy. Qualitative results (Figure 2) further support these findings, showing that our method produces sharper diaphragm boundaries and fewer motion artifacts, even in challenging cases with low PSNR.

Ablation Study 1 assesses the impact of enforcing consistent updates on all the Gaussian parameters, compared to allowing each parameter to evolve independently. It is unsurprising that the decoupled method reconstructs the dynamic images slightly better, as there are more parameters to fit and the Gaussians have more flexibility to match the projection data. The important conclusion is that imposing consistent updates only has a minor impact on reconstruction metrics, but requires fewer parameters, reduces training time, and



more importantly enables consistent continuous DVFs to be generated. Ablation Study 2 eliminates the FFD model and fits motion independently for each Gaussian, which significantly degrades performance. This highlights the critical role of spatial regularization in maintaining motion coherence across neighboring Gaussians.

In conclusion, our DI framework and FFD model effectively balances accuracy, efficiency, and robustness for dynamic CBCT reconstruction. Future work will explore physics-based deformation models and further optimization strategies for clinical applications.

**Acknowledgments.** This study was funded by Elekta Ltd., Crawley and the EPSRC-funded UCL Center for Doctoral Training in Intelligent, Integrated Imaging in Healthcare (i4health) (EP/S021930/1).

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

## References

1. Balik, S., Weiss, E., Jan, N., Roman, N., Sleeman, W.C., Fatyga, M., Christensen, G.E., Zhang, C., Murphy, M.J., Lu, J., et al.: Evaluation of 4-dimensional computed tomography to 4-dimensional cone-beam computed tomography deformable image registration for lung cancer adaptive radiation therapy. *International Journal of Radiation Oncology\* Biology\* Physics* **86**(2), 372–379 (2013)
2. Bonet, J., Wood, R.D.: *Nonlinear continuum mechanics for finite element analysis*. Cambridge university press (1997)
3. Cai, J., Yang, Y., Yuan, W., He, Y., Dong, Z., Bo, L., Cheng, H., Chen, Q.: Gaussian-informed continuum for physical property identification and simulation. *Advances in Neural Information Processing Systems* (2024)
4. Cao, A., Johnson, J.: Hexplane: A fast representation for dynamic scenes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 130–141 (2023)
5. De Los Santos, J., Popple, R., Agazaryan, N., Bayouth, J.E., Bissonnette, J.P., Bucci, M.K., Dieterich, S., Dong, L., Forster, K.M., Indelicato, D., et al.: Image guided radiation therapy (igrt) technologies for radiation therapy localization and delivery. *International journal of radiation oncology, biology, physics* **87**(1), 33–45 (2013)
6. Eiben, B., Bertholet, J., Menten, M.J., Nill, S., Oelfke, U., McClelland, J.R.: Consistent and invertible deformation vector fields for a breathing anthropomorphic phantom: a post-processing framework for the xcat phantom. *Physics in Medicine & Biology* **65**(16), 165005 (2020)
7. Eiben, B., Tran, E.H., Menten, M.J., Oelfke, U., Hawkes, D.J., McClelland, J.R.: Statistical motion mask and sliding registration. In: *Biomedical Image Registration: 8th International Workshop, WBIR 2018, Leiden, The Netherlands, June 28–29, 2018, Proceedings 8*. pp. 13–23. Springer (2018)
8. Feng, Y., Feng, X., Shang, Y., Jiang, Y., Yu, C., Zong, Z., Shao, T., Wu, H., Zhou, K., Jiang, C., et al.: Gaussian splashing: Unified particles for versatile motion synthesis and rendering. *arXiv preprint arXiv:2401.15318* (2024)

9. Fu, Y., Zhang, H., Cai, W., Xie, H., Kuo, L., Cervino, L., Moran, J., Li, X., Li, T.: Spatiotemporal gaussian optimization for 4d cone beam ct reconstruction from sparse projections. arXiv preprint arXiv:2501.04140 (2025)
10. Huang, Y., Eiben, B., Thielemans, K., McClelland, J.R.: Resolving variable respiratory motion from unsorted 4d computed tomography. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 588–597. Springer (2024)
11. Huang, Y., Thielemans, K., Price, G., McClelland, J.R.: Surrogate-driven respiratory motion model for projection-resolved motion estimation and motion compensated cone-beam ct reconstruction from unsorted projection data. *Physics in Medicine & Biology* **69**(2), 025020 (2024)
12. Jailin, C., Roux, S., Sarrut, D., Rit, S.: Projection-based dynamic tomography. *Physics in Medicine & Biology* **66**(21), 215018 (2021)
13. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* **42**(4), 139–1 (2023)
14. Kratimenos, A., Lei, J., Daniilidis, K.: Dynmf: Neural motion factorization for real-time dynamic view synthesis with 3d gaussian splatting. In: European Conference on Computer Vision. pp. 252–269. Springer (2024)
15. Leng, S., Zambelli, J., Tolakanahalli, R., Nett, B., Munro, P., Star-Lack, J., Palival, B., Chen, G.H.: Streaking artifacts reduction in four-dimensional cone-beam computed tomography. *Medical physics* **35**(10), 4649–4659 (2008)
16. Lin, J., Wang, Z., Jiang, S., Hou, Y., Jiang, M.: Phys4dgen: A physics-driven framework for controllable and efficient 4d content generation from a single image. arXiv preprint arXiv:2411.16800 (2024)
17. Lin, Y., Dai, Z., Zhu, S., Yao, Y.: Gaussian-flow: 4d reconstruction with dynamic 3d gaussian particle. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21136–21145 (2024)
18. Modat, M., Ridgway, G.R., Taylor, Z.A., Lehmann, M., Barnes, J., Hawkes, D.J., Fox, N.C., Ourselin, S.: Fast free-form deformation using graphics processing units. *Computer methods and programs in biomedicine* **98**(3), 278–284 (2010)
19. Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 5865–5874 (2021)
20. Rit, S., Oliva, M.V., Brousmiche, S., Labarbe, R., Sarrut, D., Sharp, G.C.: The reconstruction toolkit (rtk), an open-source cone-beam ct reconstruction toolkit based on the insight toolkit (itk). In: Journal of Physics: Conference Series. vol. 489, p. 012079. IOP Publishing (2014)
21. Rit, S., Wolthaus, J.W., van Herk, M., Sonke, J.J.: On-the-fly motion-compensated cone-beam ct using an a priori model of the respiratory motion. *Medical physics* **36**(6Part1), 2283–2296 (2009)
22. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L., Leach, M.O., Hawkes, D.J.: Non-rigid registration using free-form deformations: application to breast mr images. *IEEE transactions on medical imaging* **18**(8), 712–721 (1999)
23. Segars, W.P., Sturgeon, G., Mendonca, S., Grimes, J., Tsui, B.M.: 4d xcat phantom for multimodality imaging research. *Medical physics* **37**(9), 4902–4915 (2010)
24. Sonke, J.J., Zijp, L., Remeijer, P., Van Herk, M.: Respiratory correlated cone beam ct. *Medical physics* **32**(4), 1176–1186 (2005)
25. Sweeney, R.A., Seubert, B., Stark, S., Homann, V., Müller, G., Flentje, M., Guckenberger, M.: Accuracy and inter-observer variability of 3d versus 4d cone-beam ct based image-guidance in sbirt for lung tumors. *Radiation Oncology* **7**(1), 1–8 (2012)

26. Xie, T., Zong, Z., Qiu, Y., Li, X., Feng, Y., Yang, Y., Jiang, C.: Physgaussian: Physics-integrated 3d gaussians for generative dynamics. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4389–4398 (2024)
27. Yang, S., Li, Q., Shen, D., Gong, B., Dou, Q., Jin, Y.: Deform3dgs: Flexible deformation for fast surgical scene reconstruction with gaussian splatting. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 132–142. Springer (2024)
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: Advances in Neural Information Processing Systems (NeurIPS) (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)
30. Zwicker, M., Pfister, H., Van Baar, J., Gross, M.: Ewa splatting. IEEE Transactions on Visualization and Computer Graphics **8**(3), 223–238 (2002)