

TG-Field: Geometry-Aware Radiative Gaussian Fields for Tomographic Reconstruction

Yuxiang Zhong^{1,2}, Jun Wei¹, Chaoqi Chen¹, Senyou An², Hui Huang^{1*}

¹College of Computer Science and Software Engineering, Shenzhen University, Shenzhen, China

²State Key Laboratory of Intelligent Construction and Healthy Operation and Maintenance of Deep Underground Engineering, College of Civil and Transportation Engineering, Shenzhen University, Shenzhen, China
{yuxiangzhong0822, hhzhiyan}@gmail.com, {weijun, chencq, senyouan}@szu.edu.cn

Abstract

3D Gaussian Splatting (3DGS) has revolutionized 3D scene representation with superior efficiency and quality. While recent adaptations for computed tomography (CT) show promise, they struggle with severe artifacts under highly sparse-view projections and dynamic motions. To address these challenges, we propose Tomographic Geometry Field (TG-Field), a geometry-aware Gaussian deformation framework tailored for both static and dynamic CT reconstruction. A multi-resolution hash encoder is employed to capture local spatial priors, regularizing primitive parameters under ultra-sparse settings. We further extend the framework to dynamic reconstruction by introducing time-conditioned representations and a spatiotemporal attention block to adaptively aggregate features, thereby resolving spatiotemporal ambiguities and enforcing temporal coherence. In addition, a motion-flow network models fine-grained respiratory motion to track local anatomical deformations. Extensive experiments on synthetic and real-world datasets demonstrate that TG-Field consistently outperforms existing methods, achieving state-of-the-art reconstruction accuracy under highly sparse-view conditions.

Introduction

As a non-invasive imaging modality, Cone Beam Computed Tomography (CBCT) is widely used across medicine, biology, and industry to visualize internal structures (Cormack 1963; De Chiffre et al. 2014; Kherlopian et al. 2008). To achieve high-quality 3D reconstructions, CBCT systems typically require hundreds of X-ray projections (Feldkamp, Davis, and Kress 1984). However, such dense angular sampling raises concerns about radiation exposure, especially in repeated scans or high-dose protocols. To mitigate radiation dose, sparse-view CBCT reconstruction has become a key research focus, aiming to reduce the number of projections while preserving reconstruction fidelity.

Traditional CT reconstruction methods typically rely on a discrete voxel-grid representation. Recently, 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) has emerged as a compelling alternative, modeling the volume with a set of anisotropic 3D Gaussians and supporting fast image formation through efficient splatting. Building on this idea, recent

studies (Cai et al. 2025; Gao et al. 2024; Zha et al. 2024) extend 3DGS to X-ray imaging, showing clear benefits over voxel-based baselines in accuracy and speed.

Nevertheless, existing methods still face two critical challenges in practical applications: (1) **Robustness under extremely sparse views**. While current methods perform well with moderately sparse projections, their quality sharply deteriorates when projections become extremely limited. Without explicit geometric regularization, per-splat optimization fails to maintain geometric consistency and structural coherence, leading to severe artifacts and loss of anatomical detail. This poses challenges for clinical adoption, where minimal-view imaging is crucial for dose reduction. (2) **Struggle with dynamic CT reconstruction**. Most existing methods are designed for static cases and struggle with respiratory motion. Modeling temporal coherence and capturing non-rigid deformations remain challenging, limiting effectiveness in real-world 4D CT.

To address these issues, we propose TG-Field, a geometry-aware Gaussian deformation framework designed for high-fidelity static and dynamic CT reconstruction. Specifically, we employ a multi-resolution hash encoder coupled with a two-stage iterative refinement strategy to capture local geometric context and ensure spatial continuity. To further tackle the ill-posedness of sparse-view projections, we introduce a semantic-consistency regularization module powered by pretrained visual foundation models, enforcing cross-view coherence. For dynamic scenes, TG-Field embeds time-conditioned representations within each Gaussian. Since naïve joint spatiotemporal hashing often induces collisions and temporal drift, we incorporate a spatiotemporal attention block and a motion-flow network.

Our main contributions are summarized as follows:

- A new Gaussian deformation framework that achieves high-quality static and dynamic reconstruction under ultra-sparse views.
- Several technical designs, including a hash-grid encoder, a spatiotemporal attention module, and a motion-flow network to model dynamic local deformations, plus cross-view semantic-consistency regularization.
- Experiments across diverse CT scenarios demonstrate the superiority of TG-Field over existing methods, highlighting robustness and potential for clinical workflows.

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Visual comparison of point clouds obtained by different initialization methods: (a) cubic uniform sampling (Cai et al. 2025), (b) FDK-based initialization (Zha et al. 2024), (c) our iteration-based initialization, and (d) reference. Our method clearly captures more accurate and detailed geometric structures.

Related Work

Cone-Beam CT Reconstruction

CBCT reconstruction methods are commonly categorized as analytical and iterative. Analytical approaches (Feldkamp, Davis, and Kress 1984; Radon 1986) provide fast, closed-form inversion but degrade markedly under sparse-view sampling, whereas iterative approaches (Andersen and Kak 1984; Sidky and Pan 2008) improve robustness by solving data-fidelity-regularization objectives at the cost of higher computation and increased sensitivity to hyperparameter settings and prior selection. Recently, deep learning-based methods have demonstrated substantial improvements across CBCT tasks, including projection completion (Anirudh et al. 2018; Ghani and Karl 2018), volume enhancement (Chung et al. 2023; Lee et al. 2023; Liu et al. 2023, 2020), and direct volume reconstruction from sparse-view or limited-angle data (Adler and Öktem 2018; Jin et al. 2017; Lin et al. 2023, 2024; Ying et al. 2019). Nevertheless, these supervised approaches often struggle to generalize to out-of-distribution scenarios due to their reliance on paired training data.

To mitigate such limitations, recent advances leverage self-supervised, NeRF-inspired implicit representations (Mildenhall et al. 2021), enabling CT reconstruction without paired supervision. Representative methods include IntraTomo (Zang et al. 2021), NAF (Zha, Zhang, and Li 2022), SAX-NeRF (Cai et al. 2024), NEAT (Rückert et al. 2022), and NeRP (Shen, Pauly, and Xing 2022). These methods optimize continuous implicit fields for reconstruction. However, most NeRF-based approaches focus on static reconstruction and require dense ray sampling, incurring significant computational cost. Recently, STNF4D (Zhou, Ye, and Cai 2025) adopts a structured grid representation to model dynamic CT scans, yet it still converges slowly and yields suboptimal reconstruction quality.

3D Gaussian Splatting in Medical Imaging

3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) enables real-time rendering and rapid convergence, and has recently been adapted to medical imaging (Cai et al. 2025; Gao et al. 2024; Zha et al. 2024). For synthesizing X-ray projections, X-Gaussian (Cai et al. 2025) introduces grayscale modeling, while DDGS (Gao et al. 2024) incorporates radiosity

decomposition. In contrast, \mathbb{R}^2 -Gaussian (Zha et al. 2024) directly targets CT volume reconstruction via differentiable voxelization. Building on these static formulations, Spatiotemporal Gaussian (Fu et al. 2025) extends \mathbb{R}^2 -Gaussian with a deformation network and parameterizes the representation with HexPlane to construct a spatiotemporal Gaussian field for dynamic CBCT; \mathbb{X}^2 -Gaussian (Yu et al. 2025) also adopts a HexPlane representation and integrates dynamic radiative GS with self-supervised respiratory motion learning. While effective under moderate sampling, these dynamic models degrade as views become sparser. In contrast, TG-Field explicitly targets robustness under ultra-sparse views and provides a geometry-aware framework for both static and dynamic CT.

Preliminaries

Radiative Gaussian Splatting (Zha et al. 2024) models the 3D CT density field as a finite sum of anisotropic Gaussian kernels $\mathbb{G} = \{G_i\}_{i=1}^N$. Each kernel G_i is parameterized by density ρ_i , center $\boldsymbol{\mu}_i \in \mathbb{R}^3$, and covariance $\Sigma_i \in \mathbb{R}^{3 \times 3}$:

$$G_i(\mathbf{x} \mid \rho_i, \boldsymbol{\mu}_i, \Sigma_i) = \rho_i \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_i)^\top \Sigma_i^{-1}(\mathbf{x} - \boldsymbol{\mu}_i)\right).$$

The density field is the superposition

$$\sigma(\mathbf{x}) = \sum_{i=1}^N G_i(\mathbf{x} \mid \rho_i, \boldsymbol{\mu}_i, \Sigma_i).$$

Let the X-ray be $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ for $t \in [t_{\text{in}}, t_{\text{out}}]$. By the Beer-Lambert law (Swinehart 1962), the detected intensity is

$$I'(\mathbf{r}) = I_0 \exp\left(-\int_{t_{\text{in}}}^{t_{\text{out}}} \sigma(\mathbf{r}(t)) dt\right),$$

and the log-transformed measurement (line integral) is

$$I(\mathbf{r}) = \log I_0 - \log I'(\mathbf{r}) = \int_{t_{\text{in}}}^{t_{\text{out}}} \sigma(\mathbf{r}(t)) dt.$$

Thus each projection pixel $I(\mathbf{r})$ records the integral of σ along its ray.

Method

We begin with an iterative strategy that produces a structurally faithful initial point cloud. We then introduce a geometry-aware deformation framework that regularizes Gaussian primitives via a hash-grid encoder. Next, we extend the framework to dynamic CT by injecting time-conditioned representations together with a spatiotemporal attention module and a motion-flow network to capture respiratory motion. Finally, we impose a semantic-consistency regularization to further enhance reconstruction quality.

Iteration-based Initialization

High-quality point cloud initialization is crucial for ensuring fast convergence and high-fidelity reconstruction in 3DGS-based methods (Wu et al. 2024a). However, under sparse-view CT settings, existing point cloud initialization strategies such as cubic grid sampling (Cai et al. 2025) and FDK-based methods (Zha et al. 2024) struggle to preserve the geometric structures of the target, resulting in degraded reconstructions.

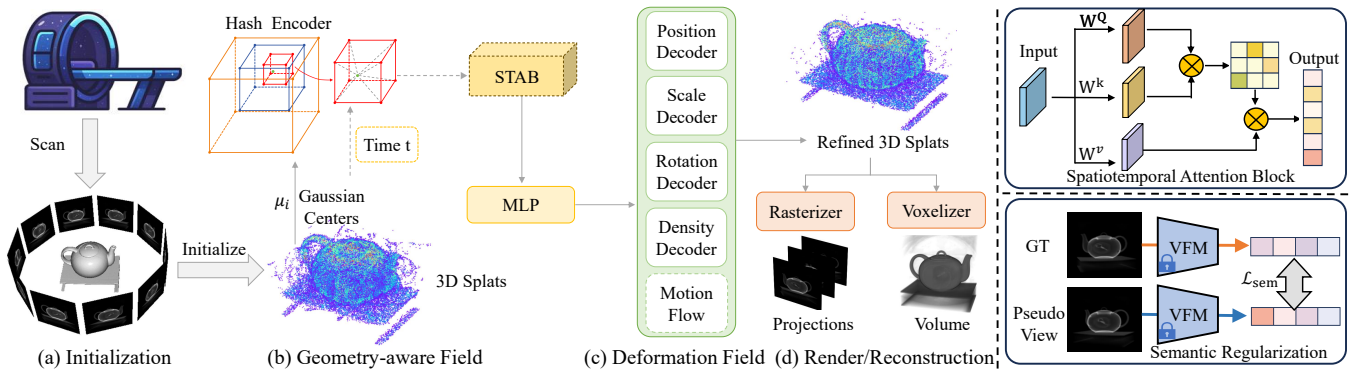


Figure 2: Our TG-Field framework first initializes a point cloud via a two-stage iterative refinement. The refined coordinates are encoded using a hash encoder to capture spatial geometric features, which are then decoded into optimized Gaussian splats by multiple MLP decoders. Additionally, we employ semantic consistency regularization with pretrained visual foundation models to maintain multi-view semantic consistency. Finally, these Gaussian splats are rendered into arbitrary-view X-ray projections and voxelized into a CT volume. For dynamic scenarios, we incorporate a spatiotemporal attention module and a Motion-Flow module to effectively capture temporal variations and fine-grained local movements.

To address this, we propose an iteration-based initialization consisting of two iterative stages. First, we use the algebraic reconstruction method Conjugate Gradient Least Squares (Hestenes, Stiefel et al. 1952) to iteratively obtain a coarse volumetric reconstruction from sparse projections. Subsequently, we refine this coarse reconstruction using the regularization-based Adaptive Steepest Descent-Projection Onto Convex Sets (Sidky and Pan 2008), enforcing total variation (TV) constraints to reduce noise and preserve structural edges. As shown in Figure 1-c, this iterative approach significantly improves initialization quality. The refined volume is then converted into Gaussian primitives characterized by position, covariance, and attenuation derived from voxel intensity, further optimized via our geometry-aware neural field.

Geometry-aware Splat Field

We introduce a geometry-aware deformation field \mathcal{D} that models spatial correlations among Gaussian primitives. By enforcing coherent deformation among the primitives, this framework compensates for missing geometric information under ultra-sparse projections and reduces artifacts.

Specifically, the deformation field predicts deformation parameters for each Gaussian primitive G_i , adjusting their positions and shapes to preserve geometric consistency. We write $G_i \equiv (\mu_i, R_i, S_i, \rho_i)$. The deformed Gaussians G'_i are computed as

$$G'_i = (\mu_i + \Delta\mu_i, R_i + \Delta R_i, S_i + \Delta S_i, \rho_i + \Delta\rho_i), \quad (1)$$

where $\Delta\mu_i$, ΔR_i , ΔS_i , and $\Delta\rho_i$ denote offsets for position, rotation, scale, and density, respectively. Formally, the deformation field is defined as

$$\mathcal{D} = \mathcal{F} \circ \mathcal{E}_{\text{Hash}}, \quad (2)$$

where $\mathcal{E}_{\text{Hash}}$ is a hash-grid encoder capturing spatial correlations, and \mathcal{F} is a lightweight decoder.

Spatially Decomposed Encoder. To explicitly capture spatial correlations under ultra-sparse views, we adopt a compact spatial decomposition with multiple low-dimensional

feature grids inspired by HexPlane (Cao and Johnson 2023) and 4DGS (Wu et al. 2024b). Concretely, the hash encoder is implemented as a set of trainable multi-resolution grids $\{G_s\}_{s \in S}$ that are queried at primitive coordinates. Given a primitive at coordinate $\mu_i \in \mathbb{R}^3$, we obtain per-scale features by differentiable interpolation

$$f_s(\mu_i) = F(G_s, \mu_i), \quad f_s(\mu_i) \in \mathbb{R}^C, \quad (3)$$

and define the hash-encoded embedding as the multi-scale concatenation

$$h_\phi(\mu_i) = \text{concat}_{s \in S}[f_s(\mu_i)] \in \mathbb{R}^{|S| \cdot C}. \quad (4)$$

This encoder provides multi-scale spatial context for each primitive and serves as the input to the deformation decoder.

Multi-head Gaussian Deformation Decoder. Given $h_\phi(\mu_i)$, a lightweight multi-head decoder outputs $\Delta\mu_i = \mathcal{F}_\mu(h_\phi(\mu_i))$, $\Delta R_i = \mathcal{F}_R(h_\phi(\mu_i))$, $\Delta S_i = \mathcal{F}_S(h_\phi(\mu_i))$, $\Delta\rho_i = \mathcal{F}_\rho(h_\phi(\mu_i))$.

4D CT Reconstruction. Our deformation-based framework naturally extends to dynamic 4D CT, addressing artifacts caused by respiratory motion and anatomical variations. We jointly encode spatial coordinates μ_i and time t to obtain per-scale features

$$f_s(\mu_i, t) = F(G_s, (\mu_i, t)), \quad f_s(\mu_i, t) \in \mathbb{R}^C, \quad (5)$$

and form a unified spatiotemporal embedding by concatenation

$$h_\phi(\mu_i, t) = \text{concat}_{s \in S}[f_s(\mu_i, t)] \in \mathbb{R}^{|S| \cdot C}. \quad (6)$$

However, jointly hashing spatial and temporal coordinates can induce hash collisions, especially when identical (or near-identical) spatial locations reoccur across time, causing ambiguous embeddings and unreliable deformations.

To mitigate this, we recalibrate features across time with a spatiotemporal attention block (STAB). For each primitive i , we stack a short temporal window of embeddings

$$\mathbf{H}_i = [h_\phi(\mu_i, t_1), \dots, h_\phi(\mu_i, t_T)]^\top \in \mathbb{R}^{T \times |S| \cdot C}. \quad (7)$$

Method	Synthetic (5 Views)		Synthetic (10 Views)		Synthetic (20 Views)		Real (5 Views)		Real (10 Views)		Real (20 Views)	
	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
FDK	11.83	0.112	15.21	0.186	18.48	0.293	13.20	0.109	17.57	0.225	21.03	0.356
SART	22.10	0.683	24.32	0.768	27.24	0.845	26.22	0.771	28.72	0.846	30.94	0.899
IntraTomo	23.79	0.733	26.47	0.781	28.54	0.825	26.59	0.748	29.67	0.802	31.65	0.835
NAF	23.83	0.711	27.41	0.784	30.59	0.849	26.76	0.769	30.94	0.814	33.71	0.853
SAX-NeRF	<u>24.05</u>	<u>0.740</u>	27.55	0.801	31.93	0.875	26.96	0.784	32.26	0.835	35.07	0.863
X-Gaussian	23.62	0.699	27.59	0.804	31.79	0.868	26.71	0.766	31.04	0.819	34.91	0.852
R ² -Gaussian	23.81	0.735	<u>28.15</u>	<u>0.833</u>	<u>32.25</u>	<u>0.923</u>	<u>27.00</u>	<u>0.789</u>	<u>32.73</u>	<u>0.859</u>	<u>36.59</u>	<u>0.901</u>
Ours	24.54	0.779	28.95	0.849	32.92	0.936	27.65	0.794	33.59	0.872	37.20	0.911

Table 1: Quantitative comparisons of sparse-view static CT reconstruction on synthetic and real-world datasets. We **bold** the best results and underline the second-best.

We compute queries, keys, and values as

$$Q = \mathbf{H}_i W_Q, \quad K = \mathbf{H}_i W_K, \quad V = \mathbf{H}_i W_V, \quad (8)$$

where $W_Q, W_K, W_V \in \mathbb{R}^{C \times C}$, and we apply scaled dot-product attention

$$\text{Attn}(\mathbf{H}_i) = \text{softmax}\left(\frac{QK^\top}{\sqrt{C}}\right)V. \quad (9)$$

The recalibrated embeddings are $\tilde{\mathbf{H}}_i = \mathbf{H}_i + \text{Attn}(\mathbf{H}_i)$, from which we take the row corresponding to time t as $\tilde{h}_\phi(\boldsymbol{\mu}_i, t)$. By aggregating temporal context, STAB disambiguates colliding bins at repeated spatial locations and yields more stable, accurate dynamic deformations. We then decode the per-primitive embedding at time t to obtain time-varying offsets:

$$\begin{aligned} \Delta\boldsymbol{\mu}_i(t) &= \mathcal{F}_\mu(\tilde{h}_\phi(\boldsymbol{\mu}_i, t)), & \Delta R_i(t) &= \mathcal{F}_R(\tilde{h}_\phi(\boldsymbol{\mu}_i, t)), \\ \Delta S_i(t) &= \mathcal{F}_S(\tilde{h}_\phi(\boldsymbol{\mu}_i, t)), & \Delta\rho_i(t) &= \mathcal{F}_\rho(\tilde{h}_\phi(\boldsymbol{\mu}_i, t)). \end{aligned} \quad (10)$$

To address non-rigid tissue motion in 4D CT, we add a motion-flow module that predicts fine-grained displacement fields to further refine positions. We adopt a ResFields MLP (Mihajlovic et al. 2024) to capture subtle local deformations that the initial deformation may miss, yielding more accurate Gaussian centers. Formally,

$$\hat{\boldsymbol{\mu}}_i(t) = \boldsymbol{\mu}_i + \Delta\boldsymbol{\mu}_i(t) + \text{Flow}(\boldsymbol{\mu}_i + \Delta\boldsymbol{\mu}_i(t), t), \quad (11)$$

where $\text{Flow}(\cdot, t) \in \mathbb{R}^3$ outputs a per-point displacement at time t .

Optimization

Semantic Regularization. Prior work shows that coupling neural scene representations with visual foundation models (VFMs) (e.g., CLIP (Radford et al. 2021) or DINOv2 (Oquab et al. 2023)) improves 3D semantic consistency and robustness (Kerr et al. 2023; Qin et al. 2024). Under sparse views, purely photometric supervision is under-constrained; semantically unrelated textures can explain the projections equally well. We therefore enforce cross-view semantic consistency using frozen VFM features on geometrically corresponding image regions. Specifically, we render novel side views and extract VFM embeddings from their




Figure 3: Ablation of the hash encoder (HE) and semantic regularization (SR) on static CT reconstruction.

reprojected crops in a side view and a training view. Let P'_i and P_i be a pair of corresponding crops for the i -th 3D location; we minimize

$$\mathcal{L}_{\text{sem}} = \frac{1}{N} \sum_{i=1}^N \|f_{\text{vfm}}(P'_i) - f_{\text{vfm}}(P_i)\|_2^2, \quad (12)$$

where f_{vfm} is a pretrained VFM encoder and $\|\cdot\|_2$ denotes the L_2 norm.

Loss Function. We optimize our framework with a compound objective. We use an L1 loss and D-SSIM to supervise the rendered X-ray projections. Following (Zha et al. 2024), we additionally employ a 3D TV regularizer \mathcal{L}_{TV} as a homogeneity prior for tomography. The overall loss is

$$\mathcal{L}_{\text{total}} = \mathcal{L}_1 + \lambda_{\text{SSIM}} \mathcal{L}_{\text{SSIM}} + \lambda_{\text{TV}} \mathcal{L}_{\text{TV}} + \lambda_{\text{sem}} \mathcal{L}_{\text{sem}}, \quad (13)$$

with weighting factors λ_{SSIM} , λ_{TV} , and λ_{sem} controlling the contribution of each component.

Progressive Training Procedure. During training, we first train an X-ray Radiative Gaussian splatting model (Zha et al. 2024) for 5000 iterations. This warm-up phase ensures the model effectively captures underlying anatomical structures from limited projection data. After warm-up, we incorporate our geometry-aware deformation field for refinement.

Experiments

Experiment Settings

Dataset. We conduct experiments on both synthetic and real-world datasets, covering static and dynamic CT scans. For static CT reconstruction, we evaluate our method on both synthetic (Cai et al. 2024) and real data (Zha et al. 2024) under three sparse-view settings (5, 10, and 20 views).



Figure 4: Qualitative comparisons of static CT reconstruction results across different view numbers (5, 10, and 20 views). Compared to SAX-NeRF and R^2 -Gaussian, our method consistently provides sharper structural details and fewer artifacts, especially under extremely sparse-view conditions.

For dynamic CT reconstruction, we use the 4D extended cardiac torso (XCAT) phantom (Segars et al. 2008), real-patient 4D CT images from the 4D-Lung collection of TCIA (Hugo et al. 2017) (for simulated sparse-view projections), and clinically acquired 4D CBCT data from the SPARE challenge (Shieh et al. 2019). For the XCAT and 4D-Lung datasets, each 4D CT sequence consists of 10 respiratory phases, and we simulate a total of 100 projections, evenly distributed across the phases, which is significantly fewer than the hundreds or thousands typically used in clinical protocols. These X-ray projections are synthesized using the TIGRE tomography toolbox (Biguri et al. 2016) with a full 360° circular trajectory and added noise to mimic realistic acquisition.

Implementation Details. TG-Field is implemented in PyTorch (Paszke et al. 2019) with CUDA (Sanders 2010). Training is performed on an NVIDIA L40 GPU using the Adam optimizer (Kinga, Adam et al. 2015) for 30k itera-

tions. The initial learning rates for the Gaussian parameters are set as follows: position at 2×10^{-4} , density at 1×10^{-2} , scale at 5×10^{-3} , and rotation at 1×10^{-3} , all exponentially decaying to 10% of their initial values. The hash encoder and decoder use initial learning rates of 2×10^{-3} and 2×10^{-4} , respectively, also decaying to 10% of their original values. Loss weights are set to $\lambda_{\text{SSIM}} = 0.25$, $\lambda_{\text{TV}} = 0.05$, and $\lambda_{\text{sem}} = 0.1$. We adopt a pretrained DINO-ViT (Zhang et al. 2022) model as the VFM for computing the semantic consistency loss. During evaluation, we report PSNR and SSIM (Wang et al. 2004) to measure volumetric reconstruction quality.

Baselines. For static CT reconstruction, we compare our method with several state-of-the-art baselines. Traditional methods include FDK (Feldkamp, Davis, and Kress 1984) and SART (Andersen and Kak 1984). NeRF-based methods include IntraTomo (Zang et al. 2021), NAF (Zha, Zhang, and Li 2022), and SAX-NeRF (Cai et al. 2024), among

Method	XCAT		TCIA		SPARE		Average	
	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
Hex-plane	21.79	0.866	23.91	0.835	26.43	0.856	24.04	0.852
K-plane	20.57	0.847	24.59	0.855	26.59	0.876	23.92	0.859
STNF4D	25.73	0.928	29.37	0.919	28.75	0.887	27.95	0.911
4DGS	<u>33.95</u>	<u>0.955</u>	<u>34.44</u>	<u>0.948</u>	<u>30.01</u>	<u>0.898</u>	<u>32.80</u>	<u>0.933</u>
Ours	35.51	0.969	35.41	0.955	30.41	0.905	33.78	0.943

Table 2: Quantitative comparisons of our TG-Field with other methods on the 4D CT datasets. We **bold** the best results and underline the second-best.




Figure 5: Visual ablation of the motion-flow module (MF) for 4D CT reconstruction.

Case	Component				Performance	
	HE	STAB	MF	SR	PSNR \uparrow	SSIM \uparrow
Static	✓				28.71	0.841
	✓			✓	28.95	0.849
Dynamic	✓				34.26	0.939
	✓	✓			34.89	0.945
	✓	✓	✓		35.23	0.952
	✓	✓	✓	✓	35.41	0.955

Table 3: Quantitative ablation of TG-Field components on static and dynamic situation.

which SAX-NeRF employs a Transformer-based architecture. Gaussian-splatting methods include X-Gaussian (Cai et al. 2025), which focuses on novel projection synthesis, and R²-Gaussian (Zha et al. 2024), which enables direct CT reconstruction via differentiable voxelization.

For dynamic CT reconstruction, we compare against recent methods designed for dynamic scenes, including HexPlane (Cao and Johnson 2023), which represents dynamic scenes via hexagonal decomposition, K-Planes (Fridovich-Keil et al. 2023), an spatiotemporal radiance field representation, STNF4D (Zhou, Ye, and Cai 2025), which incorporates spatiotemporal information for dynamic CT, and 4D Gaussian Splatting (4DGS) (Wu et al. 2024b), which focuses on real-time rendering of dynamic scenes.

Main Results

Static CT Reconstruction. Table 1 reports quantitative comparisons for sparse-view CT reconstruction on synthetic and real datasets under 5, 10, and 20 views. TG-Field achieves the best PSNR and SSIM in all settings. On the synthetic dataset, it yields PSNR gains of about 0.5, 0.8, and 0.7 dB over the strongest baseline for the 5-, 10-, and 20-view cases, respectively, together with consistent SSIM improvements. On the real dataset, TG-Field similarly outperforms all baselines, with PSNR gains of about 0.7, 0.9, and 0.6 dB for 5-, 10-, and 20-view settings. Figure 4 compares the reconstruction quality of R²-Gaussian and our TG-Field across different numbers of projection views. Under extremely sparse-view conditions, R²-Gaussian often yields blurred details and noticeable artifacts, whereas TG-Field preserves finer anatomical structures and effectively suppresses artifacts even with very limited projections.

Dynamic CT Reconstruction. We evaluate TG-Field on three 4D CT datasets: XCAT, TCIA 4D-Lung, and SPARE. As shown in Table 2, TG-Field achieves the best PSNR/SSIM on all datasets and the highest average performance. Compared with the strongest baseline 4DGS, it improves the average PSNR from 32.80 to 33.78 dB and SSIM from 0.933 to 0.943, indicating stronger robustness under severe view sparsity. STNF4D performs reasonably but is more prone to noise and blurring, while HexPlane and K-Planes lag behind across all metrics. 4DGS is a strong baseline, yet TG-Field further enhances temporal modeling and local motion handling. Visual results in Figure 6 qualitatively confirm these gains, showing clearer structures and fewer artifacts.

Ablation Study

Initialization Analysis. To evaluate the impact of initialization quality on reconstruction performance, we compare our proposed geometry-informed strategy against two widely used alternatives: cubic uniform sampling and FDK-based point cloud initialization. Quantitative results under varying view numbers on the object CT are visualized in Figure 7-(a). Our method consistently outperforms the baselines across 2, 4, 6, and 8-view settings, achieving notable gains in PSNR. These improvements confirm that accurate, geometry-aware initialization significantly enhances reconstruction fidelity under limited-view conditions.

Projection Numbers. We further examine robustness with respect to the number of projections on the TCIA dataset.




Figure 6: Visual comparisons of dynamic CT reconstruction results on the SPARE dataset. Our TG-Field shows superior performance in modeling anatomical details compared to existing approaches.




Figure 7: (a) Comparison of different initialization strategies under varying numbers of projections. (b) Reconstruction results of TG-Field using different numbers of projections.

As shown in Figure 7-(b), reconstruction quality improves as more projections are available, and TG-Field consistently outperforms both STNF4D and 4DGS across all settings.

Component Analysis. We perform ablations on both static and dynamic CT to quantify the contribution of each module

in TG-Field, including the hash encoder (HE), spatiotemporal attention block (STAB), motion-flow module (MF), and semantic regularization (SR). As summarized in Table 3, on static CT adding SR on top of HE brings a consistent gain in PSNR/SSIM, which is also reflected by sharper bone boundaries and fewer artifacts in Figure 3. On 4D CT, progressively introducing STAB and MF yields steady improvements, and the full model with SR achieves the best overall PSNR and SSIM. Qualitative results in Figure 5 further show that MF mainly refines motion-sensitive regions.

Conclusion

We presented TG-Field, a geometry-aware Gaussian deformation framework for ultra-sparse-view tomographic reconstruction. It encodes geometric priors into radiative Gaussian optimization without requiring paired training data. Building on the same formulation, we handle dynamic CT via time-conditioned modeling to preserve temporal coherence. Experiments on synthetic and real datasets show consistent gains over recent methods, with higher fidelity and fewer artifacts in both static and dynamic settings.

Acknowledgments

We sincerely thank the anonymous reviewers for their constructive comments and Xueqi Ma for helpful discussions. This work was supported in parts by National Key R&D Program of China (2024YFB3908500), NSFC (U21B2023.52474105), Shenzhen Science and Technology Program (KJZD20240903100022028, KQTD20210811090044003, RCJC20200714114435012), Research Team Cultivation Program (2023QNT004) and Scientific Development Funds from Shenzhen University.

References

- Adler, J.; and Öktem, O. 2018. Learned primal-dual reconstruction. *IEEE transactions on medical imaging*, 37(6): 1322–1332.
- Andersen, A. H.; and Kak, A. C. 1984. Simultaneous algebraic reconstruction technique (SART): a superior implementation of the ART algorithm. *Ultrasonic imaging*, 6(1): 81–94.
- Anirudh, R.; Kim, H.; Thiagarajan, J. J.; Mohan, K. A.; Champley, K.; and Bremer, T. 2018. Lose the views: Limited angle CT reconstruction via implicit sinogram completion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6343–6352.
- Biguri, A.; Dosanjh, M.; Hancock, S.; and Soleimani, M. 2016. TIGRE: a MATLAB-GPU toolbox for CBCT image reconstruction. *Biomedical Physics & Engineering Express*, 2(5): 055010.
- Cai, Y.; Liang, Y.; Wang, J.; Wang, A.; Zhang, Y.; Yang, X.; Zhou, Z.; and Yuille, A. 2025. Radiative gaussian splatting for efficient x-ray novel view synthesis. In *European Conference on Computer Vision*, 283–299. Springer.
- Cai, Y.; Wang, J.; Yuille, A.; Zhou, Z.; and Wang, A. 2024. Structure-aware sparse-view x-ray 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11174–11183.
- Cao, A.; and Johnson, J. 2023. Hexplane: A fast representation for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 130–141.
- Chung, H.; Ryu, D.; McCann, M. T.; Klasky, M. L.; and Ye, J. C. 2023. Solving 3d inverse problems using pre-trained 2d diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22542–22551.
- Cormack, A. M. 1963. Representation of a function by its line integrals, with some radiological applications. *Journal of applied physics*, 34(9): 2722–2727.
- De Chiffre, L.; Carmignato, S.; Kruth, J.-P.; Schmitt, R.; and Weckenmann, A. 2014. Industrial applications of computed tomography. *CIRP annals*, 63(2): 655–677.
- Feldkamp, L. A.; Davis, L. C.; and Kress, J. W. 1984. Practical cone-beam algorithm. *Josa a*, 1(6): 612–619.
- Fridovich-Keil, S.; Meanti, G.; Warburg, F. R.; Recht, B.; and Kanazawa, A. 2023. K-planes: Explicit radiance fields in space, time, and appearance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12479–12488.
- Fu, Y.; Zhang, H.; Cai, W.; Xie, H.; Kuo, L.; Cervino, L.; Moran, J.; Li, X.; and Li, T. 2025. Spatiotemporal gaussian optimization for 4d cone beam ct reconstruction from sparse projections. *arXiv preprint arXiv:2501.04140*.
- Gao, Z.; Planche, B.; Zheng, M.; Chen, X.; Chen, T.; and Wu, Z. 2024. DDGS-CT: Direction-Disentangled Gaussian Splatting for Realistic Volume Rendering. *arXiv preprint arXiv:2406.02518*.
- Ghani, M. U.; and Karl, W. C. 2018. Deep learning-based sinogram completion for low-dose CT. In *2018 IEEE 13th Image, Video, and Multidimensional Signal Processing Workshop (IVMSP)*, 1–5. IEEE.
- Hestenes, M. R.; Stiefel, E.; et al. 1952. Methods of conjugate gradients for solving linear systems. *Journal of research of the National Bureau of Standards*, 49(6): 409–436.
- Hugo, G. D.; Weiss, E.; Sleeman, W. C.; Balik, S.; Keall, P. J.; Lu, J.; and Williamson, J. F. 2017. A longitudinal four-dimensional computed tomography and cone beam computed tomography dataset for image-guided radiation therapy research in lung cancer. *Medical physics*, 44(2): 762–771.
- Jin, K. H.; McCann, M. T.; Froustey, E.; and Unser, M. 2017. Deep convolutional neural network for inverse problems in imaging. *IEEE transactions on image processing*, 26(9): 4509–4522.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4): 139–1.
- Kerr, J.; Kim, C. M.; Goldberg, K.; Kanazawa, A.; and Tancik, M. 2023. LERF: Language embedded radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, 19729–19739.
- Kherlopian, A. R.; Song, T.; Duan, Q.; Neimark, M. A.; Po, M. J.; Gohagan, J. K.; and Laine, A. F. 2008. A review of imaging techniques for systems biology. *BMC systems biology*, 2: 1–18.
- Kinga, D.; Adam, J. B.; et al. 2015. A method for stochastic optimization. In *International conference on learning representations (ICLR)*, volume 5, 6. San Diego, California;.
- Lee, S.; Chung, H.; Park, M.; Park, J.; Ryu, W.-S.; and Ye, J. C. 2023. Improving 3D imaging with pre-trained perpendicular 2D diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10710–10720.
- Lin, Y.; Luo, Z.; Zhao, W.; and Li, X. 2023. Learning deep intensity field for extremely sparse-view CBCT reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 13–23. Springer.
- Lin, Y.; Yang, J.; Wang, H.; Ding, X.; Zhao, W.; and Li, X. 2024. C²2RV: Cross-Regional and Cross-View Learning for Sparse-View CBCT Reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11205–11214.

- Liu, J.; Anirudh, R.; Thiagarajan, J. J.; He, S.; Mohan, K. A.; Kamilov, U. S.; and Kim, H. 2023. DOLCE: A model-based probabilistic diffusion framework for limited-angle ct reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10498–10508.
- Liu, Z.; Bicer, T.; Kettimuthu, R.; Gursoy, D.; De Carlo, F.; and Foster, I. 2020. TomoGAN: low-dose synchrotron x-ray tomography with generative adversarial networks: discussion. *JOSA A*, 37(3): 422–434.
- Mihajlovic, M.; Prokudin, S.; Pollefeys, M.; and Tang, S. 2024. ResFields: Residual Neural Fields for Spatiotemporal Signals. In *International Conference on Learning Representations (ICLR)*.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Qin, M.; Li, W.; Zhou, J.; Wang, H.; and Pfister, H. 2024. Langsplat: 3d language gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20051–20060.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmlR.
- Radon, J. 1986. On the determination of functions from their integral values along certain manifolds. *IEEE transactions on medical imaging*, 5(4): 170–176.
- Rückert, D.; Wang, Y.; Li, R.; Idoughi, R.; and Heidrich, W. 2022. Neat: Neural adaptive tomography. *ACM Transactions on Graphics (TOG)*, 41(4): 1–13.
- Sanders, J. 2010. *CUDA by Example: An Introduction to General-Purpose GPU Programming*. Addison-Wesley Professional.
- Segars, W. P.; Mahesh, M.; Beck, T. J.; Frey, E. C.; and Tsui, B. M. 2008. Realistic CT simulation using the 4D XCAT phantom. *Medical physics*, 35(8): 3800–3808.
- Shen, L.; Pauly, J.; and Xing, L. 2022. NeRP: implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. *IEEE Transactions on Neural Networks and Learning Systems*, 35(1): 770–782.
- Shieh, C.-C.; Gonzalez, Y.; Li, B.; Jia, X.; Rit, S.; Mory, C.; Riblett, M.; Hugo, G.; Zhang, Y.; Jiang, Z.; et al. 2019. SPARE: Sparse-view reconstruction challenge for 4D cone-beam CT from a 1-min scan. *Medical physics*, 46: 3799–3811.
- Sidky, E. Y.; and Pan, X. 2008. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in Medicine & Biology*, 53(17): 4777.
- Swinehart, D. F. 1962. The beer-lambert law. *Journal of chemical education*, 39(7): 333.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- Wu, G.; Yi, T.; Fang, J.; Xie, L.; Zhang, X.; Wei, W.; Liu, W.; Tian, Q.; and Wang, X. 2024a. 4D Gaussian Splatting for Real-Time Dynamic Scene Rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 20310–20320.
- Wu, G.; Yi, T.; Fang, J.; Xie, L.; Zhang, X.; Wei, W.; Liu, W.; Tian, Q.; and Wang, X. 2024b. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20310–20320.
- Ying, X.; Guo, H.; Ma, K.; Wu, J.; Weng, Z.; and Zheng, Y. 2019. X2CT-GAN: reconstructing CT from biplanar X-rays with generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10619–10628.
- Yu, W.; Cai, Y.; Zha, R.; Fan, Z.; Li, C.; and Yuan, Y. 2025. X²-Gaussian: 4D Radiative Gaussian Splatting for Continuous-time Tomographic Reconstruction. *arXiv preprint arXiv:2503.21779*.
- Zang, G.; Idoughi, R.; Li, R.; Wonka, P.; and Heidrich, W. 2021. Intratomo: self-supervised learning-based tomography via sinogram synthesis and prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1960–1970.
- Zha, R.; Lin, T. J.; Cai, Y.; Cao, J.; Zhang, Y.; and Li, H. 2024. R²-Gaussian: Rectifying Radiative Gaussian Splatting for Tomographic Reconstruction. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Zha, R.; Zhang, Y.; and Li, H. 2022. NAF: neural attenuation fields for sparse-view CBCT reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 442–452. Springer.
- Zhang, H.; Li, F.; Liu, S.; Zhang, L.; Su, H.; Zhu, J.; Ni, L. M.; and Shum, H.-Y. 2022. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*.
- Zhou, Q.; Ye, Y.; and Cai, Z. 2025. Spatiotemporal-Aware Neural Fields for Dynamic CT Reconstruction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 10834–10842.