

CTGS: Hybrid Gaussian Splatting for Compact Industrial CT Representation

Shuyang Zhu^{1,*} and Yukie Nagai¹

¹ School of Engineering, the University of Tokyo, 7-3-1 Hongo Bunkyo-ku Tokyo, Japan

*Corresponding Author: derulo-zsy@g.ecc.u-tokyo.ac.jp

Abstract:

Computed tomography (CT) is widely used for inspecting complex manufactured parts in industrial settings. However, dense CT volumes are costly to store, render, and analyze. Existing compact representations often trade off geometric fidelity, internal intensity accuracy, and interactive visualization. A practical CT representation should preserve material boundaries, support compact storage, and allow efficient rendering of internal structures. Therefore, we propose CTGS, a hybrid Gaussian representation for industrial CT volumes. CTGS uses two complementary primitives. Surface Gaussians are thin and anisotropic primitives aligned with material boundaries to represent geometric surfaces, normals, and topology. Bulk Gaussians are compact volumetric primitives placed inside material regions to represent the CT intensity field within the object. Given a reconstructed CT volume, CTGS initializes surface primitives from material boundaries and initializes bulk primitives as a contained lattice inside the material domain and two branches are optimized with separated roles. CTGS achieves significant data reduction compared with dense CT volumes. It also preserves surface geometry for downstream analysis and provides a compact bulk field for fast intensity visualization. Experiments show that the hybrid design gives a practical balance between compression, geometry fidelity, and CT intensity reconstruction.

Keywords: X-ray Computed Tomography, CT Volume, 3D Gaussian Splatting

1. Introduction

Industrial CT is a common tool for inspecting manufactured parts with complex internal structures. It provides a dense three-dimensional voxel volume of the object, which is useful for defect inspection, geometric metrology, and digital archiving. To further broaden these applications, it is important to convert the scan into a representation that is compact, queryable, and still faithful to material-air boundaries.

Voxel volumes are direct and reliable, but they are not always convenient as working representations. Large volumes are expensive to store and stream, especially when repeated inspection, visualization, or remote analysis is required. Downstream geometry operations such as surface extraction, distance queries, and local editing also remain tied to the resolution of the grid. A surface mesh gives a more explicit boundary, but it removes the interior attenuation field. For industrial CT, both properties are often needed: the boundary should remain sharp enough for inspection, and the interior should still provide material-side intensity information.

3D Gaussian splatting [1] provides a useful starting point. It stores a scene as explicit anisotropic primitives, supports differentiable queries, and can be much more compact than a dense grid. However, a direct transfer from 3DGS to CT is not sufficient. CT volumes do not contain view-dependent color, and their main signal is material attenuation. A single undifferentiated Gaussian role therefore tends to

mix two different tasks: representing a thin material boundary and representing a volumetric material interior. This mixing can blur material-air transitions or spread attenuation into void regions, which is undesirable for parts with cavities, holes, and thin walls.

Based on these considerations, we propose CTGS, a role-separated Gaussian framework for industrial CT representation. We first build a geometry bundle from the CT volume, including a material mask and a signed distance field. We then initialize surface Gaussians on the material boundary and bulk Gaussians inside the material region. During training, the surface branch is responsible for boundary geometry and normal consistency, while the bulk branch models material occupancy and attenuation.

This separation makes the representation easier to interpret and control. Surface Gaussians preserve material boundaries for downstream analysis and mesh extraction, while bulk Gaussians reconstruct the interior attenuation distribution for material-side inspection. The loss design further improves robustness to outliers, scatter, beam hardening, and streak artifacts. Since bulk primitives are placed only inside the surface-guided material region and can adapt their size during optimization, CTGS avoids representing large redundant air regions and provides a compact CT scene without discarding the boundary geometry. This design gives a practical balance between compact storage, geometric fidelity, and CT intensity reconstruction.




Fig. 1 Pipeline of CTGS, where a CT volume is converted into a role-separated Gaussian scene in which bulk Gaussians reconstruct the material attenuation field, while surface Gaussians anchor boundary-side attenuation and geometry for regularized optimization and mesh extraction

2. Related Works

2.1 CT Reconstruction

CT reconstruction is driven by a tension between projection coverage and reconstruction quality. Analytical methods such as filtered back-projection and the cone-beam FDK algorithm [2] invert the acquisition in closed form and remain the clinical default, but they presuppose dense angular sampling: as views are reduced to lower dose and scan time, their reconstructions degrade into streak artifacts. Iterative and model-based methods [3,4] respond by reformulating reconstruction as regularized optimization, where total-variation and compressed-sensing priors recover plausible volumes from sparse or noisy projections—at the price of repeated projection passes. Neural attenuation fields [5–7] push this further, replacing the voxel grid with a continuous coordinate network trained against measured projections and reconstructing from only tens of views, yet per-sample network queries keep them slow.

The second limitation is that the reconstructed CT object is still represented as volume rather than as an explicit geometric scene. Recovering material boundaries therefore requires an additional isosurfacing stage, and even dedicated schemes such as dual marching cubes interpolation on octree volume-fraction data [8]. As a result, attenuation reconstruction and boundary geometry extraction remain largely decoupled. This motivates the use of an explicit, differentiable representation that can jointly encode material attenuation and geometric structure.

2.2 3D Gaussian Splatting

3D Gaussian splatting [1] represents scenes with explicit anisotropic primitives and has implemented in CT applications. X-Gaussian models radiative Gaussian splatting for X-ray novel-view synthesis [9], and R²-Gaussian corrects projection integration for tomographic reconstruction [10]. Other recent CT Gaussian systems

aim to accelerate sparse-view reconstruction or to use Gaussians as a differentiable volume prior [11]. These methods demonstrate that Gaussian primitives can be effective for CT imaging, but their main target is usually projection-domain reconstruction.

Surface-oriented Gaussian methods provide another important direction. 2D Gaussian splatting flattens primitives into oriented disks [12], and SuGaR regularizes 3D Gaussians for mesh extraction [13]. Implicit-surface methods such as DeepSDF [14] and NeuS [15] similarly focus on the boundary as a level set. These techniques motivate our surface branch, but a CT volume also needs the material interior. CTGS therefore keeps the surface branch explicit while using a separate bulk branch for interior attenuation.

3. Method

3.1 Overview

We first take a reconstructed CT volume as input. The training process involves five key steps. First, a geometry analysis stage estimates the material support and the signed distance field related to the material boundary. Second, these geometric cues are used to initialize a role-separated Gaussian scene, where surface Gaussians are placed near the material-air interface and bulk Gaussians are placed inside the material domain. Third, the CT volume is queried through a compact bulk attenuation field, while the surface branch remains responsible for geometry and normal consistency. Fourth, CT objective optimizes intensity reconstruction, boundary regularity, and material containment. Finally, adaptive bulk growth and pruning adjust the primitive count, producing a compact representation whose compression level can be controlled by available reconstruction budget.

3.2 Geometry Analysis and Sampling

To obtain a stable geometric prior, we first define the reconstructed CT volume as a scalar field $I: \Omega \rightarrow \mathbb{R}_+$, $\Omega \subset \mathbb{R}^3$, where Ω denote

the three-dimensional CT domain and $I(\mathbf{x})$ denotes the CT intensity at position $\mathbf{x}$. Based on an intensity threshold η , we define a binary material mask $M(\mathbf{x}) = \mathbb{1}[I(\mathbf{x}) \geq \eta]$, where $\mathbb{1}[\cdot]$ denotes the indicator function, which equals 1 if the enclosed condition is satisfied and 0 otherwise. The material region is defined as $\Omega_m = \{\mathbf{x} \in \Omega \mid M(\mathbf{x}) = 1\}$, and air-side region as $\Omega_a = \Omega \setminus \Omega_m$ such as by using Otsu algorithm [16]. After defining M , we can define a signed distance field defined as:

$$D(\mathbf{x}) = \begin{cases} -\min_{\mathbf{y} \in \partial\Omega_m} \|\mathbf{x} - \mathbf{y}\|_2, & \mathbf{x} \in \Omega_m, \\ \min_{\mathbf{y} \in \partial\Omega_m} \|\mathbf{x} - \mathbf{y}\|_2, & \mathbf{x} \in \Omega_a, \end{cases} \quad (1)$$

where $\partial\Omega_m = \{\mathbf{x} : D(\mathbf{x}) = 0\}$ is the material–air boundary, therefore $D(\mathbf{x}) < 0$ indicates material, $D(\mathbf{x}) = 0$ indicates the material–air interface, and $D(\mathbf{x}) > 0$ indicates air or void regions. The mask M and the signed distance field D together provide the cues needed to initialize both surface and bulk Gaussians in the next section.




Fig. 2 Geometry analysis of the CT volume: (a) CT intensity $I(\mathbf{x})$; (b) binary material mask $M(\mathbf{x})$, where white areas denote material and black areas denote air; and (c) signed distance field $D(\mathbf{x})$, whose absolute value gives the distance to the boundary, with negative values inside the material and positive values in air regions.

3.3 Role-Separated Gaussian Representation

In the proposed framework, the CT object is represented as a Gaussian scene $\mathcal{G}$ composed of two disjoint subsets: *surface Gaussians* $\mathcal{G}^s$ and *bulk Gaussians* $\mathcal{G}^b$, which share a generic Gaussian primitive as:

$$g_i = (\boldsymbol{\mu}_i, \mathbf{R}_i, \boldsymbol{\ell}_i, \alpha_i, a_i), \quad (2)$$

where $\boldsymbol{\mu}_i \in \mathbb{R}^3$ is the center, $\mathbf{R}_i \in SO(3)$ is the orientation, $\boldsymbol{\ell}_i \in \mathbb{R}_+^3$ is the per-axis scale, $\alpha_i \in [0, 1]$ is the opacity, and $a_i \in \mathbb{R}$ is a scalar attenuation attribute.

Given the signed distance field D , the two index sets are defined as

$$\mathcal{I}^s = \{i \mid D(\boldsymbol{\mu}_i) = 0\}, \quad \mathcal{I}^b = \{j \mid D(\boldsymbol{\mu}_j) < 0\}, \quad (3)$$

so that surface primitives g_i^s are seeded in the zero level set and bulk primitives g_i^b are seeded inside the material.

Surface Gaussians Surface Gaussians preserve the material–air boundary. Each surface primitive g_i^s is oriented so that its normal axis aligns with the signed distance field (SDF) gradient $\mathbf{n}_i = \nabla D(\boldsymbol{\mu}_i) / \|\nabla D(\boldsymbol{\mu}_i)\|$ and the remaining two axes span the local tangent plane. The normal-direction scale is initialized in a small size, making each surface primitive a thin, boundary-aligned disk. During training, each g_i^s is regularized for geometry and normal consistency.

Bulk Gaussians Bulk Gaussians represent the interior of the object,

where each g_i^b carries a learnable attenuation value a_i that encodes local CT intensity. To predict the attenuation at each position $\mathbf{x}$, we evaluate the collective influence of all g_i^b at that location and aggregate their contributions into a single intensity estimate.

We first determine how strongly each g_i^b influences $\mathbf{x}$ by computing the squared Mahalanobis distance under its local ellipsoidal frame:

$$d_i(\mathbf{x}) = \|\text{diag}(\boldsymbol{\ell}_i)^{-1} \mathbf{R}_i^T (\mathbf{x} - \boldsymbol{\mu}_i)\|_2^2, \quad (4)$$

where $\text{diag}(\boldsymbol{\ell}_i)$ is the diagonal scale matrix, and $\mathbf{R}_i^T$ transforms the offset $\mathbf{x} - \boldsymbol{\mu}_i$ from the global coordinate system into the local principal-axis frame of g_i^b . This quantity measures how far $\mathbf{x}$ lies from the primitive center $\boldsymbol{\mu}_i$ in units of the ellipsoid's own axes, with $d_i(\mathbf{x}) = 1$ corresponding to one scale unit along a principal axis.

The influence of each g_i^b is then converted into an effective contribution weight w_i defined as:

$$w_i(\mathbf{x}) = \alpha_i e^{-\frac{d_i(\mathbf{x})}{2}} \mathbb{1}[d_i(\mathbf{x}) \leq d_c], \quad i \in \mathcal{I}^b. \quad (5)$$

where $\alpha_i \in [0, 1]$ is the learned opacity, d_c is the compact-support cutoff radius, and $\mathbb{1}[d_i(\mathbf{x}) \leq d_c]$ is the indicator function that truncates the Gaussian tail beyond d_c . Finally, the predicted attenuation $\hat{I}$ at $\mathbf{x}$ is given by the weighted mean over all contributing primitives:

$$\hat{I}(\mathbf{x}) = \frac{\sum_{i \in \mathcal{I}^b} \alpha_i w_i(\mathbf{x})}{\max(\sum_{i \in \mathcal{I}^b} w_i(\mathbf{x}), \epsilon)}, \quad (6)$$

where ϵ is a small constant for numerical stability. This normalized form yields a weighted average of the attenuation values carried by overlapping primitives, ensuring that the predicted intensity reflects local material composition rather than the number of overlapping primitives.




Fig. 3 Role-separated Gaussian representation. (a) Surface Gaussians are seeded on the zero level set $D = 0$ as thin boundary-aligned discs, where each normal axis follows the SDF gradient ∇D . (b) Bulk Gaussians are placed inside the material domain $D < 0$ and carry learnable attenuation values a_i , where the intensity at query point $\mathbf{x}$ is read out from primitives within the compact support d_c .

3.4 Training Objective

After role-separated initialization, we optimize the Gaussian scene while keeping the responsibilities of the two subsets distinct. The total objective combines intensity reconstruction, material occupancy, surface regularization, and bulk containment:

$$\mathcal{L} = \lambda_{\text{ct}} \mathcal{L}_{\text{ct}} + \lambda_{\text{surf}} \mathcal{L}_{\text{surf}} + \lambda_{\text{cont}} \mathcal{L}_{\text{cont}}. \quad (7)$$

Here, $\mathcal{L}_{\text{ct}}$ supervises CT intensity reconstruct, $\mathcal{L}_{\text{surf}}$ regularizes boundary geometry, and $\mathcal{L}_{\text{cont}}$ prevents bulk primitives from leaking into air regions. The coefficients λ_{ct} , λ_{surf} , and λ_{cont} control the relative contributions of the corresponding loss terms.

CT reconstruction loss CT reconstruction loss drives all

parameters of g_i by comparing the predicted attenuation $\hat{I}(\mathbf{x})$ against the ground-truth intensity $I(\mathbf{x})$. To reduce the influence of scatter, beam hardening, and streak artifacts, we use a robust Huber penalty H_β :

$$\mathcal{L}_{ct} = \frac{\sum_{\mathbf{x} \in \mathcal{X}} w(\mathbf{x}) H_\beta(\hat{I}(\mathbf{x}) - I(\mathbf{x}))}{\sum_{\mathbf{x} \in \mathcal{X}} w(\mathbf{x})}, \quad (8)$$

where $\mathcal{X}$ denotes the sampled query set and $w(\mathbf{x})$ gives lower weight to air-side samples to balance material reconstruction and leakage suppression.

Surface loss Surface loss is applied only to g_i^s , regularizing them to stay thin and aligning boundaries. It encourages their centers to lie on the material boundary $D = 0$, aligns their normals with the SDF gradient, and constrains their effective normal thickness:

$$\mathcal{L}_{surf} = \frac{1}{|\mathcal{J}_s|} \sum_{i \in \mathcal{J}_s} (|D(\mu_i^s)| + \lambda_n \mathcal{L}_{n,i} + \lambda_{thin} [\max(0, t_i^s - \tau_n)]^2), \quad (9)$$

where t_i^s is the effective thickness of g_i^s along local SDF normal, τ_n is its prescribed upper bound. The normal-alignment term $\mathcal{L}_{n,i}$ is defined as where λ_n controls its strength:

$$\mathcal{L}_{n,i} = 1 - \mathbf{n}_i^s \cdot \frac{\nabla D(\mu_i^s)}{\|\nabla D(\mu_i^s)\|_2}. \quad (10)$$

Containment loss Containment loss prevents g_i^b from leaking into air by regularizing the spatial extent of each bulk primitive. Specifically, it penalizes any g_i^b whose compact ellipsoidal support, determined by μ_i and ℓ_i , approaches or crosses the material-air boundary:

$$\mathcal{L}_{cont} = \frac{1}{|\mathcal{J}_b|} \sum_{i \in \mathcal{J}_b} \text{ReLU}(D(\mu_i^b) + r_i^b + \delta_c)^2, \quad (11)$$

where r_i^b is an isotropic footprint radius computed as the mean of the three scale components of g_i^b , δ_c is a small safety margin. Since $D < 0$ inside the material, this term activates when a bulk primitive's support region extends beyond the interface.

3.5 Compact Bulk Containment and Adaptive Growth

To fill under-covered material regions, we define the **bulk occupancy** as $\text{occ}_b(\mathbf{x}) = 1 - e^{-\sum_{i \in \mathcal{J}_b} w_i(\mathbf{x})}$, which maps the cumulative bulk contribution at $\mathbf{x}$ to $[0,1]$. This quantity does not predict CT intensity; it measures whether the existing bulk primitives collectively reach position $\mathbf{x}$.

Adaptive bulk reseed. Candidates for reseeding are drawn from material-side samples where occupancy is insufficient:

$$\mathcal{C}_{reseed} = \{\mathbf{x} \in \Omega_m \mid \text{occ}_b(\mathbf{x}) < \tau_r\}. \quad (12)$$

A candidate $\mathbf{x}_{new} \in \mathcal{C}_{reseed}$ is accepted only if both its center and its compact ellipsoidal support remain inside the material domain. Each accepted primitive is initialized with center $\mu_i = \mathbf{x}_{new}$, a default isotropic scale ℓ_i , and attenuation attribute $a_i = I(\mathbf{x}_{new})$ taken from the local CT value. This fills narrow structures and cavity shells that gradient-based optimization alone cannot reach, without allocating primitives in exterior air.

Role-separated densification. Existing primitives are split according to their branch role. g_i^s is split tangentially when its tangent-plane scale is large and its center receives a sufficient position gradient; both children are projected back onto the $D = 0$ to maintain boundary alignment. g_i^b is split along its largest scale axis

$\max_k \ell_{i,k}$ only when both children satisfy the material containment condition. Surface splits thus increase boundary resolution, while bulk splits increase interior attenuation coverage.

Controllability. Because bulk primitives are generated only inside the material and bounded by D , we avoid representing large air regions entirely. The compression level is controlled by four parameters: initial bulk spacing, compact support radius, reseeding budget, and pruning threshold. A denser initialization or larger reseeding budget improves material coverage at the cost of more primitives, while tighter settings produce a more compact representation.




Fig. 4 Effect of adaptive growth. Disabling it leads to sparse and incomplete material coverage (b), while the full CTGS model preserves the overall geometry and internal holes (a).

4. Experiments

4.1 Setup

Dataset. We evaluate CTGS on three CT volumes with different geometric characteristics. The *rocker-arm* and *fandisk* cases are simulated industrial CT volumes generated from known CAD meshes. The *rocker-arm* contains holes, thin material bridges, and internal cavities, while the *fandisk* contains sharper corners and stronger outliers and artifacts. In addition, we include a real scanned *figure* case as a qualitative stress case, where the material support is sparser and more fragmented.

Baselines. We compare CTGS with two Gaussian-based CT reconstruction baselines under a common material-domain protocol. X-Gaussian [9] is included as a 3D Gaussian baseline adapted to sparse-view CT projections, and R²-Gaussian [10] is used as a representative projection-domain radiative Gaussian method. Since these methods consume projections rather than reconstructed CT volumes, we synthesize parallel-beam projections from the CT volume and use the same training and held-out view split for all projection-domain baselines. To avoid a comparison dominated by exterior air or scanner-specific background artifacts, all reported projection metrics are evaluated on the material contribution only.

Implementation details. CTGS is implemented in PyTorch with CUDA acceleration and optimized with Adam [17]. We train CTGS for 500 iterations with learning rates initialized as 1.6×10^{-4} for position, 0.05 for opacity, 0.005 for scale, 0.001 for rotation, and 0.001 for CT attenuation. During training, primitive roles are kept fixed; bulk centers are frozen in the reported setting, while attenuation, opacity, and surface regularization are optimized. We use bulk reseeding for material-coverage recovery. The main loss weights are set to $\lambda_{ct} = 1.0$, $\lambda_{surf} = 0.5$, and $\lambda_{cont} = 1.0$. In $\mathcal{L}_{ct}$, we use a Huber threshold of $\beta = 0.1$, $w(\mathbf{x})$ of 1.0 in material and 0.25 in air, and surface regularization parameters $\lambda_n = 1.0$, $\lambda_{thin} = 0.25$, and $\tau_n = 0.45$. Each iteration samples 4096 query points for Fandisk and Rocker Arm and 8192 query points for the Figure scan. Reconstruction quality is

reported on the original grid using material-domain PSNR, and qualitative results are shown with middle slices.

4.2 Results

Reconstruction and projection. Table 1 reports the reconstruction and projection results under the converted current setting. X-Gaussian and R²-Gaussian are trained from 48 synthesized projection views and evaluated on 12 held-out views, while CTGS is optimized from the reconstructed CT volume. Since CTGS is designed as a material-centered representation, its metrics are computed on the GT material mask for both volume and projection evaluation. In contrast, R²-Gaussian and X-Gaussian are evaluated on the full grid and full detector image, including exterior air and CT background artifacts. Under this protocol, R²-Gaussian fits the held-out projections well on the figure scan, but its 3D structural similarity remains lower than CTGS. On Rocker Arm, CTGS provides higher material-side projection and volume consistency, while X-Gaussian exhibits strong projection artifacts in the held-out views. This result positions CTGS relative to projection-domain Gaussian methods in our reconstructed-volume data regime, with the different input and evaluation domains kept explicit.

Mesh extraction. Following the Poisson-based mesh extraction

strategy used in SuGaR [13], we extract meshes from the surface branch rather than from the optimized attenuation field. To make the two objects comparable, all surface-deviation values are normalized by the reference bounding-box diagonally and reported as percentages. The symmetric Chamfer-L1 distance is 0.230% for *Fandisk* and 0.497% for *Rocker Arm*. The median absolute signed deviation is 0.184% and 0.222%, while the 95th-percentile absolute signed deviation is 0.461% and 1.098%, respectively. The extracted meshes have limited systematic one-sided bias under this normalized evaluation.

Representation compactness. CTGS represents each CT object with a sparse set of Gaussian primitives rather than allocating values over the full voxel grid. This object-centered representation substantially reduces the number of stored elements: across the three test cases, the optimized Gaussian scenes use only 0.055% to 0.214% of primitives compared to the original voxel counts. The saved CTGS point-cloud representations also achieve file-size reduction factors ranging from 9.1 to 35.1 relative to the original DICOM folders. This indicates that CTGS concentrates storage around material regions and boundaries, while avoiding explicit allocation over large exterior-air volumes.

Table 1 Quantitative comparison of 3D reconstruction accuracy on the material domain (3D Mat) and the full volume (3D). PSNR evaluates voxel-wise reconstruction fidelity based on squared error, SSIM measures structural similarity, and MAE reports the mean absolute attenuation error. Darker cells indicate better results.

Dataset	Method	3D PSNR↑	3D SSIM↑	3D MAE↓	3D Mat. PSNR↑	3D Mat. SSIM↑	3D Mat. MAE↓
Fandisk	CTGS	23.47	0.8971	0.0388	17.87	0.9668	0.1251
	R ² Gaussian	18.93	0.3040	0.0815	17.86	0.8787	0.0979
	X-Gaussian	11.51	0.3683	0.1618	13.06	0.8578	0.1692
Figure	CTGS	26.11	0.7125	0.0300	18.84	0.9081	0.0945
	R ² -Gaussian	21.51	0.6814	0.0324	13.74	0.8675	0.1562
	X-Gaussian	22.51	0.5748	0.0464	23.00	0.9748	0.0462
Rocker Arm	CTGS	27.76	0.9229	0.0174	22.18	0.9890	0.0307
	R ² -Gaussian	15.74	0.683	0.1024	11.16	0.8989	0.2317
	X-Gaussian	14.18	0.2224	0.1191	16.89	0.9753	0.1004




Fig. 5 It shows representative slice renderings of the three objects, comparing the ground truth, two Gaussian-based baselines, and CTGS.

Table 2 Mesh extraction accuracy normalized by reference bounding-box diagonal.

Dataset	Chamfer-L1	Signed p50	Abs. p95	Nearest p95
Fandisk	0.497%	-0.042%	1.098%	1.194%
Rocker Arm	0.230%	-0.011%	0.461%	0.476%




Fig. 6 Signed mesh deviation normalized by the reference bounding-box diagonal, using a shared $\pm 1.1\%$ color scale.

Table 3 Summary of the compactness of CTGS.

Dataset	Voxels	Primitives	DICOM Size	CTGS Size	Reduction
Fandisk	71.1M	115.4K	136.0 MB	11.44 MB	11.9×
Figure	496.0M	272.0K	946.7 MB	26.98 MB	35.1×
Rocker Arm	37.4M	79.9K	72.2 MB	7.92 MB	9.1×

4.3 Limitations

CTGS is intended as a compact representation for efficient visualization and approximate reconstruction of industrial CT volumes, rather than as a lossless replacement for the original voxel data.

CTGS still has limitations in regions with high curvature, thin structures, and partial-volume boundaries. The geometry analysis stage converts the CT volume into a binary material support. This support gives a stable geometric prior, but it may also simplify weak or narrow structures before optimization starts. During training, the Huber loss and the material-side filling strategy improve robustness to outliers, scatter, beam hardening, and small false holes. However, the same design can smooth very small protrusions, sharp local corners, or narrow gaps when their CT evidence is weak.

In addition, the current surface branch mainly handles the material-air interface. Multi-material interfaces are not explicitly modeled in the present framework. Future work will extend the surface representation to multi-material boundaries, improve adaptive material-side sampling for small structures, and evaluate CTGS on broader industrial datasets.

Future work will improve small-structure sampling, extend CTGS to multi-material interfaces, and evaluate it on broader and more complex industrial CT datasets.

5. Conclusion

We presented CTGS, a hybrid Gaussian splatting framework for compact industrial CT representation. CTGS converts a reconstructed CT volume into a role-separated Gaussian scene. It first builds a material mask and a signed distance field from the input volume. Surface Gaussians are then initialized on the material-air interface, while bulk Gaussians are initialized inside the material region. With this design, the surface branch preserves boundary geometry and normal consistency, and the bulk branch represents material occupancy and attenuation through compact-support readout.

Our experiments show that this separation is useful for material-side CT reconstruction. CTGS keeps high structural similarity on the material domain and gives stable material projection results. Projection-domain Gaussian baselines can obtain higher full-projection PSNR in some cases because they also fit the air-side background and scanner-like artifacts.

For geometry extraction, The extracted meshes of CTGS remain close to the reference geometry in most exterior regions, while larger errors mainly appear near high-curvature edges and thin boundary parts. These results suggest that the role-separated representation provides a practical balance between compact storage, boundary preservation, and material-side attenuation reconstruction.

However, CTGS may smooth very thin structures due to the Huber loss and filling strategy and does not explicitly handle multi-material interfaces. Despite these limitations, CTGS provides a compact and interpretable representation that preserves material boundaries while reconstructing material-side attenuation.

REFERENCES

1. Kerbl B, Kopanas G, Leimkühler T, Drettakis G (2023) 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans Graph* 42(4):Article 139, 1–14.
2. Feldkamp LA, Davis LC, Kress JW (1984) Practical Cone-Beam Algorithm. *J Opt Soc Am A* 1(6):612–619.
3. Andersen AH, Kak AC (1984) Simultaneous Algebraic Reconstruction Technique (SART): A Superior Implementation of the ART Algorithm. *Ultrason Imaging* 6(1):81–94.
4. Wang G, Ye JC, De Man B (2020) Deep Learning for Tomographic Image Reconstruction. *Nat Mach Intell* 2(12):737–748.
5. Mildenhall B, Srinivasan PP, Tancik M, Barron JT, Ramamoorthi R, Ng R (2020) NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In: *Proc. ECCV*, LNCS 12346, pp. 405–421. Springer.
6. Müller T, Evans A, Schied C, Keller A (2022) Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM Trans Graph* 41(4):Article 102, 1–15.
7. Zha R, Zhang Y, Li H (2022) NAF: Neural Attenuation Fields for Sparse-View CBCT Reconstruction. In: *Proc. MICCAI*, LNCS 13436, pp. 442–452. Springer.
8. Kim, Seungki, et al. "A novel interpolation scheme for dual marching cubes on octree volume fraction data." *Computers & Graphics* 66 (2017): 169-178.
9. Cai Y, Liang Y, Wang J, Wang A, Zhang Y, Yang X, Zhou Z, Yuille A (2024) Radiative Gaussian Splatting for Efficient X-Ray Novel View Synthesis. In: *Proc. ECCV*, LNCS 15059, pp. 275–292. Springer.
10. Zha R, Lin TJ, Cai Y, Cao J, Zhang Y, Li H (2024) R²-Gaussian: Rectifying Radiative Gaussian Splatting for Tomographic Reconstruction. In: *Proc. NeurIPS*, pp. 44907-44934
11. Pieta PT, Pedersen RJ, Borgi S, Jørgensen JS, Andreassen JW, Dahl VA (2026) FaCT-GS: Fast and Scalable CT Reconstruction with Gaussian Splatting. *arXiv:2604.01844*.
12. Huang B, Yu Z, Chen A, Geiger A, Gao S (2024) 2D Gaussian Splatting for Geometrically Accurate Radiance Fields. In: *Proc. ACM SIGGRAPH Conference Papers*. ACM.
13. Guédon A, Lepetit V (2024) SuGaR: Surface-Aligned Gaussian Splatting for Efficient 3D Mesh Reconstruction and High-Quality Mesh Rendering. In: *Proc. CVPR*, pp. 5354–5363.
14. Park JJ, Florence P, Straub J, Newcombe R, Lovegrove S (2019) DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In: *Proc. CVPR*, pp. 165–174.
15. Wang P, Liu L, Liu Y, Theobalt C, Komura T, Wang W (2021) NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-View Reconstruction. In: *Proc. NeurIPS*. *Int J Autom Techn* 11:5:781-786.
16. Otsu N (1979) A Threshold Selection Method from Gray-Level Histograms. *IEEE Trans Syst Man Cybern* 9(1):62–66.
17. Kingma DP, Ba J (2015) Adam: A Method for Stochastic Optimization. In: *Proc. ICLR*.