

Splat-Based Metal Artifact Reduction in Cone-Beam CT via Compact Attenuation Modeling

Kiseok Choi

Jaemin Cho

Inchul Kim

Min H. Kim

KAIST

{kschoi; jmcho; ickim; mhkim}@vclab.kaist.ac.kr

Abstract

X-ray computed tomography (CT) suffers from severe metal artifacts when high-attenuation objects such as dental fillings or orthopedic implants are present. These artifacts originate from the polychromatic nature of X-rays, where attenuation varies strongly with photon energy and material composition, breaking the monochromatic assumption used by conventional reconstruction algorithms. Recent neural rendering approaches attempt to address this mismatch through differentiable polychromatic projection models, but they still struggle with smoothness bias, loss of fine structures, and prohibitive computation when extended to large-scale cone-beam CT. We introduce a splat-based metal artifact reduction framework that incorporates a physically grounded polychromatic forward model into a continuous Gaussian representation for cone-beam CT. Each Gaussian encodes the energy-dependent attenuation of the underlying material using a compact material parameterization, which enables efficient joint optimization of geometric and material properties without relying on a metal mask. This compact attenuation formulation captures the essential variation across biological tissues and metallic implants, allowing our model to explain metal-induced nonlinearity while preserving high-frequency structure. Experiments on simulated and real cone-beam CT scans show that our method converges significantly faster and suppresses metal artifacts more effectively than existing reconstruction and neural field-based approaches.

1. Introduction

X-ray computed tomography (CT) is a non-destructive imaging technique that recovers the internal structure of an object by transmitting X-rays through it and measuring the attenuated energy. The geometry of CT systems has evolved from parallel-beam to fan-beam, and more recently to cone-beam configurations [13, 14]. Cone-beam CT (CBCT) en-




Figure 1. Qualitative comparison of CBCT reconstructions for a real walnut specimen containing inserted metal pins (highlighted in red). Each method is shown using a horizontal slice (top row) and a vertical slice (bottom row). FDK [8], when applied directly to the metal-inserted scan, produces strong beam hardening artifacts including dark streaks, severe shading, and distorted intensities. LIMAR [15] reduces some streaks but noticeably degrades overall structural fidelity. Polyner [35] removes most streaks, however it introduces blurring and oversmoothing that wipe out fine details near metal boundaries. In contrast, our method delivers the most faithful reconstruction, suppressing metal-induced artifacts while preserving crisp structural detail in both axial and sagittal views, demonstrating strong robustness under severe beam hardening in cone-beam CT.

ables volumetric reconstruction from a single rotation, making it widely used in medical diagnostics and industrial inspection.

Despite its efficiency, CBCT reconstruction is highly sensitive to inaccuracies in the underlying X-ray physics. In practice, the detected intensity follows a *polychromatic* Beer–Lambert integration over the source spectrum and material-dependent attenuation coefficients. The attenuation of photons varies strongly with their energy and the material composition they traverse. Conventional reconstruction algorithms, including Feldkamp–Davis–Kress (FDK) and analytical back-projection methods, assume *monochromatic* attenuation, an assumption that breaks down in the presence of high-attenuation metals. This mismatch produces nonlinear distortions in the projection data, result-

ing in severe streaking and shading artifacts around metallic regions. The problem becomes especially pronounced in medical applications, where dental fillings or orthopedic implants introduce extreme beam hardening [3, 12, 27].

Many approaches have attempted to mitigate metal artifacts in fan-beam CT [15, 19, 22, 35]. Early methods corrected corrupted projection bins or interpolated missing regions, sometimes combined with image-domain priors [15, 19]. Although simple, these heuristics are fragile, geometry dependent, and unable to explain the nonlinear physics induced by metal. More recent work has shifted toward model-based reconstruction using differentiable polychromatic forward projections. Neural field-based methods draw inspiration from NeRFs [20] and optimize a continuous attenuation function to match polychromatic measurements [22, 35]. However, such representations often exhibit noise amplification, and their computational cost becomes prohibitive in full 3D cone-beam geometry. Gaussian splatting offers a more efficient and expressive alternative for CBCT reconstruction, enabling faster optimization and improved structural detail [39]. However, existing splat-based methods are still limited to monochromatic attenuation models. In this work, we introduce a splat-based metal artifact reduction framework that incorporates not only a differentiable polychromatic forward projection model into a continuous Gaussian representation, but also a differentiable reconstruction formulation that enables effective image-domain regularization. A key component of our method is a compact representation of energy-dependent material attenuation that captures essential variations across biological tissues and metallic implants using a physically grounded, low-dimensional model. This formulation enables efficient joint optimization of geometry and material properties, improving optimization stability, accelerating convergence, and effectively reducing metal-induced artifacts. Our main contributions are as follows:

- We introduce a Gaussian-splatting-based CBCT reconstruction framework that incorporates both a physically grounded polychromatic forward-projection model and a differentiable reconstruction pipeline to achieve effective metal artifact reduction.
- We propose a compact representation of energy-dependent attenuation that maintains physical realism while reducing the dimensionality of material optimization.
- We provide a 3D CBCT dataset with realistic metal-induced artifacts and plan to release both the dataset and code upon publication.

2. Related Work

CBCT reconstruction. Cone-beam CT reconstruction has been a central topic in X-ray imaging since the mid-20th

century. The FDK algorithm [8], a 3D extension of filtered back-projection based on the Fourier slice theorem, remains the industrial standard due to its simplicity and computational efficiency. However, FDK relies on a monochromatic attenuation assumption and uniform sampling, which makes it sensitive to interpolation errors and geometric inconsistencies, resulting in radial artifacts and limited reconstruction quality.

Iterative reconstruction methods [1, 9, 10] cast CBCT as an inverse problem by repeatedly updating voxel intensities to minimize projection discrepancies. These methods allow the integration of regularizers, such as total variation (TV) [6, 7, 37], which improve edge preservation and noise robustness. Yet, their reliance on voxel discretization and repeated forward and back projections leads to high computational cost, making them impractical for high-resolution CBCT volumes.

To move beyond voxel-based representations, recent work incorporates continuous neural fields. SAX-NeRF [4] uses transformer-based attention along projection rays to reconstruct high-quality volumes, but optimization remains slow and memory-intensive. NAF [38] and NeAT [25] improve efficiency using octrees and the instant-NGP [21], achieving faster reconstructions but still exhibiting over-smoothing and blurry structures. R^2 -Gaussian [39] introduces Gaussian splatting with differentiable projection and Beer-Lambert modeling, providing both speed and fine detail. However, the method assumes a normalized volume space and relies on a monochromatic attenuation model, which prevents accurate modeling of real polychromatic X-ray physics and limits its ability to handle metal-induced nonlinearity.

Metal artifact reduction. Early metal artifact reduction techniques [15, 19] identify metal-corrupted regions in the sinogram and replace them with interpolated neighboring values to suppress streaking artifacts. Although simple and fast, these heuristic approaches are highly sensitive to metal size, shape, and placement, which leads to inconsistent behavior across scans and an inability to model the underlying physics.

Learning-based post-processing methods [16, 30–33] attempt to remove artifacts directly from reconstructed slices using neural networks. While they can improve visual quality within the distribution of the training data, their dependence on supervised examples limits generalization under different imaging geometries, object compositions, or acquisition conditions.

More physically grounded approaches incorporate differentiable CBCT models to enable physics-aware MAR [4, 39]. Polyner [35] embeds the source spectrum and metal masks as priors to estimate energy-dependent attenuation, but its reliance on metal segmentation and its NeRF backbone results in amplified noise and slow conver-

gence. Diner [36] simplifies the problem using an energy-independent attenuation parameter, which reduces complexity at the cost of missing key polychromatic behavior. Park et al. [22] assume a linearly energy-dependent model, avoiding metal masks but oversimplifying the attenuation physics, which leaves noticeable residual artifacts. In contrast, our method introduces a splat-based, physics-driven polychromatic projection and differentiable reconstruction framework that does not require metal masks or computationally heavy volumetric rendering networks. Furthermore, our formulation naturally supports incorporating image-domain priors. By leveraging a compact material representation within Gaussian splatting, our approach preserves sharp structural boundaries while accurately modeling metal-induced nonlinear attenuation, resulting in more effective and efficient artifact reduction.

3. Background

Polychromatic forward projection. X-ray photons attenuate according to the Beer–Lambert law:

$$I = I_0 \exp\left(-\int_L \mu(l) dl\right), \quad (1)$$

where I_0 is the source intensity, $\mu(l)$ is the linear attenuation coefficient along the ray path L , and I is the detector measurement. Taking the logarithm yields the standard monochromatic projection model:

$$P = -\log \frac{I}{I_0} = \int_L \mu(l) dl. \quad (2)$$

Real X-ray sources, however, emit polychromatic spectra, and both the source and attenuation depend on photon energy. The monochromatic model therefore fails to describe the true physical process. The correct formulation integrates attenuation over the entire energy spectrum \mathcal{E} :

$$P = -\log \int_{\mathcal{E}} \eta(E) \exp\left(-\int_L \mu(l, E) dl\right) dE, \quad (3)$$

where $\eta(E)$ denotes the spectral distribution and filtration, and $\mu(l, E)$ is the energy-dependent linear attenuation coefficient (LAC). This coupling between energy and material composition produces the characteristic beam-hardening behavior observed in polychromatic CT and creates a strong physical prior for inferring material-specific attenuation.

Material attenuation. The linear attenuation coefficient can be factorized into mass density ρ and the mass attenuation coefficient (MAC) $\mu_\rho(E)$ [26]:

$$\mu(E) = \rho \mu_\rho(E). \quad (4)$$

The MAC is an intrinsic quantity determined solely by elemental composition and photon energy. It does not depend on the material’s geometry or microstructure, and tabulated MAC curves are available for most materials through

NIST [26]. The density ρ , in contrast, varies with physical structure, such as variations in bone density across individuals or the difference between soft tissue and cortical tissue. This decomposition separates a fixed, energy-dependent material signature from a variable, scene-specific density term. In our Gaussian-splatting representation, this structure-dependent density is encoded in the density parameter of each Gaussian primitive, enabling us to model heterogeneous attenuation fields.

Splat-based reconstruction. In \mathbb{R}^2 -Gaussian [39], the attenuation field is modeled as a continuous sum of anisotropic Gaussian primitives:

$$\mu(\mathbf{x}) = \sum_i^M \delta_i \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{p}_i)^T \boldsymbol{\Sigma}_i^{-1}(\mathbf{x} - \mathbf{p}_i)\right), \quad (5)$$

where \mathbf{x} is a spatial query point, M is the number of Gaussians, δ_i is the density of the i th Gaussian, and \mathbf{p}_i and $\boldsymbol{\Sigma}_i$ denote its center and covariance.

The forward projection of a ray is computed by accumulating the contributions of all Gaussians along the ray path:


$$\begin{aligned} P(\hat{\mathbf{x}}) &= \int_L \mu(l) dl \\ &\approx \sum_i^M \sqrt{\frac{2\pi|\tilde{\boldsymbol{\Sigma}}_i|}{|\hat{\boldsymbol{\Sigma}}_i|}} \delta_i \exp\left(-\frac{1}{2}(\hat{\mathbf{x}} - \hat{\mathbf{p}}_i)^T \hat{\boldsymbol{\Sigma}}_i^{-1}(\hat{\mathbf{x}} - \hat{\mathbf{p}}_i)\right), \end{aligned} \quad (6)$$

where $\hat{\mathbf{x}}$ is the pixel location, $\tilde{\boldsymbol{\Sigma}}_i$ is the view-space covariance, $\hat{\boldsymbol{\Sigma}}_i$ is the projected covariance, and $\hat{\mathbf{p}}_i$ is the projected center. The parameters δ_i , \mathbf{p}_i , and $\boldsymbol{\Sigma}_i$ are optimized so that Equation (6) matches the measured projections. After optimization, the final attenuation field is evaluated using Equation (5), yielding a continuous reconstruction.

4. Method

Overview. Our method extends Gaussian-splatting based CBCT reconstruction [39] to perform metal artifact reduction by explicitly modeling polychromatic X-ray physics. Figure 2 summarizes the overall pipeline. Instead of assuming a monochromatic source, we adopt the full polychromatic projection model in Equation (3), which integrates energy-dependent attenuation along each ray. Directly optimizing this formulation, however, is challenging due to its nonlinear coupling between material properties and the X-ray spectrum.

To make this optimization tractable, we assign each Gaussian primitive a compact material parameter that governs its mass attenuation behavior across energy. This parameter controls a low-dimensional approximation of the MAC curve, allowing us to express the polychromatic forward model with significantly fewer variables while maintaining physical realism. As a result, our approach



i : Index of a Gaussian primitive
 E : X-ray photon energy
 M : Number of Gaussian primitives
 \mathbf{p}_i : Center position of the Gaussian primitive
 Σ_i : Covariance of the Gaussian primitive
 δ_i : Density of the Gaussian primitive
 u_i : Compact material parameter of the Gaussian primitive
 $f_R(\cdot)$: Weight function for reconstruction
 $f_P(\cdot)$: Weight function for projection
 $\mu(\mathbf{x}, E)$: Attenuation coefficient at \mathbf{x} and E
 $P(\hat{\mathbf{x}}, E)$: Pixel value of the projection image at $\hat{\mathbf{x}}$ and E
 \mathbf{x} : Voxel position of the reconstruction volume
 $\hat{\mathbf{x}}$: Pixel position of a projection image
 $\eta(E)$: X-ray spectrum response at E

Figure 2. Method overview. Instead of using the conventional monochromatic attenuation assumption, we introduce an energy-dependent formulation and additional compact material parameters for each Gaussian primitive to simulate energy-dependent attenuation. These parameters are jointly optimized along with the baseline Gaussian parameters during reconstruction.

can jointly optimize geometric parameters and material-dependent attenuation in a mask-free manner, enabling efficient and physically consistent metal artifact reduction within the Gaussian-splating framework.

Compact material attenuation. In medical imaging, metallic implants introduce strong energy-dependent attenuation. Among clinically used metals, titanium and iron are most common, with aluminum, chromium, cobalt, and nickel also appearing in specialized applications [3, 12, 27]. Biological tissues, in contrast, are largely composed of water, protein, lipid, and carbohydrate [28, 34], which results in a comparatively narrow range of mass attenuation coefficients.

Empirically, the MAC curves of common biological tissues and clinically relevant metals occupy a low-dimensional manifold. Although the absolute magnitude differs substantially between soft tissue and metal, the shape of the MAC curve across energy remains smooth and highly correlated. Leveraging this observation, we approximate the MAC of each Gaussian primitive using a quadratic Bézier curve whose control points represent the minimum, intermediate, and maximum MAC vectors (\mathbf{b}_s , \mathbf{b}_m , \mathbf{b}_f) observed across these materials:

$$\mu_\rho(u_i, E) = (1 - u_i)^2 \mathbf{b}_s(E) + 2(1 - u_i)u_i \mathbf{b}_m(E) + u_i^2 \mathbf{b}_f(E), \quad (7)$$

where $u_i \in [0, 1]$ is a continuous scalar that smoothly interpolates along this material manifold.

This compact representation offers three key benefits. First, it captures the dominant variation of real MAC curves across energy, since the difference between materials lies primarily in scale and curvature rather than in complex high-frequency patterns. Second, it provides a compact and smoothly varying parameterization that stabilizes the optimization of polychromatic attenuation within the Gaussian framework. Third, by encoding MAC variation in one scalar, it avoids spatially inconsistent attenuation estimates that arise when optimizing a high-dimensional material vec-




Figure 3. (a) Representative mass attenuation coefficients (MACs) from the NIST database [26] of biological and metallic materials. (b) A quadratic Bézier curve defined by minimum, intermediate, and maximum MAC vectors closely approximates the MACs of real materials in (a), confirming that the real MACs lie on a smooth, low-dimensional manifold.

tor at every primitive.

To validate this formulation, we compare real MAC curves from NIST with their Bézier-approximated counterparts (Figure 3). The interpolated curves reproduce both the magnitude and the shape of the energy-dependent attenuation with high fidelity, indicating that a scalar parameter u_i is sufficient to approximate the physically relevant variation of MACs. Prior work has also observed similar low-dimensional structure in energy-dependent attenuation for related tasks [17], supporting the feasibility of our compact modeling approach.

Polychromatic forward projection. The baseline approach [39] reformulates the Beer–Lambert law into a Gaussian-based forward projection model (Equation (6)), but it assumes a *monochromatic* X-ray source and therefore cannot capture the energy-dependent attenuation responsible for metal artifacts. We address this limitation by extending the formulation to a fully polychromatic model (Equation (8)), where the X-ray spectrum is explicitly integrated over energy:

$$P(\hat{\mathbf{x}}, E) = -\log \sum_{E=0}^{E_{\max}} \eta(E) \exp\left(-\sum_{i=1}^M f_P(\hat{\mathbf{x}} | \mathbf{p}_i, \Sigma_i) \delta_i \mu_\rho(u_i, E)\right), \quad (8)$$

where $\eta(E)$ is the X-ray spectrum response, and $f_P(\hat{\mathbf{x}} |$

\mathbf{p}_i, Σ_i) is the projection weight of the i -th Gaussian: $f_P(\hat{\mathbf{x}} | \mathbf{p}_i, \Sigma_i) = \sqrt{\frac{2\pi|\hat{\Sigma}_i|}{|\Sigma_i|}} \exp\left(-\frac{1}{2}(\hat{\mathbf{x}} - \hat{\mathbf{p}}_i)^T \hat{\Sigma}_i^{-1}(\hat{\mathbf{x}} - \hat{\mathbf{p}}_i)\right)$.

Here, δ_i denotes the density of the primitive, and $\mu_\rho(u_i, E)$ is the energy-dependent attenuation model controlled by its *compact material parameter* u_i . This formulation embeds a differentiable polychromatic Beer–Lambert law directly into the Gaussian splatting framework.

To model the X-ray spectrum, we use $\eta(E)$ generated by the SPEKTR simulator [24], following prior polychromatic reconstruction studies [35, 36]. By jointly optimizing the compact material parameter u_i along with all Gaussian primitives, our method adaptively captures diverse attenuation behaviors across tissue and metal, which enables effective suppression of metal-induced artifacts.

The attenuation field is reconstructed by aggregating all Gaussian contributions:

$$\mu(\mathbf{x}, E) = \sum_{i=1}^M f_R(\mathbf{x} | \mathbf{p}_i, \Sigma_i) \delta_i \mu_\rho(u_i, E), \quad (9)$$

where the reconstruction weight is

$$f_R(\mathbf{x} | \mathbf{p}_i, \Sigma_i) = \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{p}_i)^T \Sigma_i^{-1}(\mathbf{x} - \mathbf{p}_i)\right).$$

This energy-aware reconstruction enables our method to accurately recover attenuation volumes across both biological tissue and metal regions, providing a physically consistent foundation for metal artifact reduction. Furthermore, we design the reconstruction pipeline to be fully differentiable, enabling seamless integration of image-domain priors. Detailed derivations of the loss gradients with respect to the intermediate variables in the reconstruction pipeline are provided in the supplementary material.

Optimization. We optimize the per-Gaussian parameters $\{\mathbf{p}_i, \Sigma_i, \delta_i, u_i \mid i = 1, \dots, M\}$, where u_i is the compact material parameter that controls the energy-dependent attenuation of each primitive. The overall objective combines a pixel-wise L_1 loss, a structural similarity (SSIM) loss, and a total variation (TV) loss computed over randomly sampled voxel locations, following the baseline splat-based reconstruction [39]:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_1(P_{\text{GT}}, P) + \lambda_0 \mathcal{L}_{\text{SSIM}}(P_{\text{GT}}, P) + \lambda_1 \mathcal{L}_{\text{TV}}(\mu(V, E_{\text{eff}})), \quad (10)$$

Here, V denotes voxel-grid query points. We set $\lambda_0 = 0.25$ and $\lambda_1 = 3.0$ in all experiments. Full gradient expressions for P and V with respect to the Gaussian parameters are provided in the supplemental material, enabling end-to-end backpropagation through our CUDA-based differentiable projection pipeline.

After optimization, the final voxelized attenuation field is obtained using Equation (9). The reconstructed image is then computed at the effective energy $E_{\text{eff}} =$

$\sum_{i=1}^N \eta(E_i) E_i$, where N is the number of spectral bins, following prior work [36].

5. Evaluation

Implementation. We implement our method in PyTorch [23] with CUDA [11]. The Bézier MAC basis is constructed from the NIST [26] attenuation curves of water, iron, and aluminum, assuming an operating spectrum of 10–90 keV. We use $N = 15$ uniformly spaced energy samples for spectrum integration and adopt the SPEKTR-generated spectral response $\eta(E)$ following Polyner [35]. Gaussian initialization follows R²-Gaussian [39].

5.1. Validation

Dataset. We evaluate our method on three synthetic CBCT phantoms—*Lung*, *Teeth*, and *Broccoli*—constructed from the LIDC [2], X-plant [29], and ZCB100 [18] datasets, each containing inserted metal objects. The datasets are generated using the synthetic projection pipeline introduced in prior work [5].

Evaluation on synthetic dataset. Figure 4 shows qualitative comparisons on three synthetic scenes. FDK [8] produces strong streak artifacts due to its monochromatic assumption, and Polyner [35] suppresses streaks but loses fine structures. Park et al. [22] yields heavily blurred results with incorrect intensity levels. Our method preserves structural details—clearly separating bone and soft tissue in *Lung* and retaining high-frequency textures in *Broccoli*—while effectively removing metal artifacts. As reported in Table 2, our approach achieves the highest 3D PSNR and SSIM across all phantoms. LIMAR [15], NMAR [19], and Polyner improve over FDK but remain limited, whereas supervised baselines [31–33] fail to generalize due to domain gaps. We provide additional results on two synthetic scenes—a pancreas with titanium (Ti) and a pepper with iron (Fe)—in the supplementary material.

5.2. Real Experiment

Dataset. We acquire real CBCT scans using a Bruker SKYSCAN 1273 system at 90 kVp with 1.0 mm aluminum filtration. We construct four real datasets using readily available objects that mimic key anatomical characteristics: garlic (complex structures), enoki mushroom and blueberry (fine textures), and avocado (multiple internal materials). To represent soft-tissue and bone characteristics, we additionally employ the chicken dataset from existing work [5]. Furthermore, to evaluate our method on a human head-like structure, we also include the walnut dataset from the same work. Each specimen is scanned both with and without metal to capture metal-induced beam hardening, with the metal-free scans serving as a reference. Additional system details are provided in the supplemental material.




Figure 4. Axial (yellow), coronal (pink), and sagittal (green) slices from the three synthetic scenes are shown, with metal regions highlighted in red. FDK produces severe streaking and shading artifacts. Polyner and Park et al. reduce some streaks but suffer from detail loss, contrast distortion, and heavy blurring (Park et al. displayed with a 0.12 intensity cap). In contrast, our method preserves fine structural details and effectively eliminates metal-induced artifacts, yielding the most faithful reconstruction across all scenes.

Evaluation on real dataset. Figure 5 shows reconstruction results on the real-world datasets. Similar to the synthetic experiments, FDK [8] produces strong streaking artifacts due to its monochromatic assumption. LIMAR [15] reduces some of the severe artifacts, particularly in the avocado

scene, but leaves noticeable residual streaks. OSCNet [33] struggles to generalize beyond its training distribution, leading to poor overall reconstruction quality. Polyner [35] achieves strong artifact suppression, yet it oversmooths important structural details, failing to capture the garlic’s

Table 1. Efficiency comparison for joint reconstruction and metal artifact reduction methods. Ours is consistently faster across both synthetic and real datasets, achieving significant speedups over existing polychromatic and neural-field-based methods.

Dataset	Scene	Polyner [35]	Park et al. [22]	Ours
Synthetic	Broccoli	1h 50m 02s	1h 48m 49s	28m 57s
	Lung	1h 14m 34s	1h 12m 38s	41m 40s
	Teeth	1h 15m 44s	1h 13m 41s	41m 14s
Real	Walnut	1h 11m 58s	1h 10m 20s	30m 10s
	Garlic	1h 30m 19s	1h 26m 34s	19m 54s
	Avocado	1h 56m 49s	1h 50m 24s	27m 02s
	Chicken	2h 02m 55s	2h 09m 43s	23m 36s
	Enoki Mushroom	1h 33m 10s	1h 27m 54s	21m 11s
	Blueberry	1h 33m 12s	1h 28m 12s	18m 52s

Table 2. Quantitative results on the synthetic CBCT datasets. We report 3D PSNR and SSIM for the Lung, Teeth, and Broccoli scenes (Figure 4). Our method achieves the highest reconstruction accuracy across all cases, demonstrating strong metal-artifact suppression and superior preservation of structural detail.

Method	Lung		Teeth		Broccoli	
	PSNR3D	SSIM3D	PSNR3D	SSIM3D	PSNR3D	SSIM3D
FDK [8]	17.21	0.905	17.71	0.885	18.00	0.963
LIMAR [15]	20.02	0.947	18.47	0.899	18.65	0.965
NMAR [19]	20.22	0.950	18.58	0.900	20.40	0.977
ACDNet [32]	13.89	0.655	9.21	0.436	14.36	0.939
DICDNet [31]	13.48	0.622	8.73	0.401	13.99	0.929
OSNet [33]	13.39	0.615	8.67	0.392	13.81	0.922
Polyner [35]	20.63	0.977	20.74	0.970	21.60	0.990
Park et al. [22]	18.80	0.887	14.90	0.768	0.07	0.201
Ours	28.96	0.994	27.40	0.993	27.76	0.997

Table 3. Effect of the number of spectral components on reconstruction performance.

# of Spec. Comp. (N)	7	15	31	63
PSNR3D	28.01	28.04	27.93	27.99
SSIM3D	0.994	0.995	0.994	0.994

fine internal texture and the chicken’s soft-tissue variations. Park et al. [22] also produce overly blurred reconstructions, losing high-frequency content across all scenes. In contrast, our method consistently removes metal-induced artifacts while preserving material-dependent density variations and high-frequency structural details, yielding the most faithful reconstructions across all real-world cases.

5.3. Computation Time

We compare the computation time of our method with state-of-the-art joint reconstruction and metal artifact reduction approaches in Table 1. All measurements are performed on an Intel Xeon 4214R CPU and an NVIDIA RTX A6000 GPU. Across all scenes, our method achieves an order-of-magnitude speedup, demonstrating the computational advantage of the splat-based formulation and the compact material parameterization used in our polychromatic model.

6. Ablation

Component Contribution. We evaluate the impact of our polychromatic model and TV regularization. As shown in

Table 4, the polychromatic model already yields a large improvement over the baseline [39], and adding TV provides a small additional gain.

Table 4. Ablation of the polychromatic model and TV regularization. Both components improve reconstruction quality over the baseline [39].

Module	PSNR3D	SSIM3D
Baseline	23.22	0.984
Baseline + Poly (Ours)	27.97	0.994
Baseline + Poly + TV (Ours)	28.04	0.995

Number of spectrum components. We evaluate the effect of the number of spectral components N used for polychromatic integration. As shown in Table 3, reconstruction quality remains stable across a wide range of N , indicating that our compact material model is robust to the spectral sampling density. The differences between 7, 15, 31, and 63 samples are small, which suggests that the underlying attenuation behavior is well captured even with relatively coarse sampling. We choose $N = 15$ as a practical operating point, since it provides the best balance between numerical accuracy and computational efficiency, while further increasing N yields negligible gains and slightly higher runtime.

7. Limitation

Our compact material model captures the dominant variation in energy-dependent attenuation, yet it cannot represent highly unusual materials or compounds whose MAC curves deviate significantly from the low-dimensional manifold shown in Figure 3. In practice, the X-ray spectrum cannot be measured directly because energy-integrating detectors are used. Following common practice and prior work [35], we therefore rely on physics-based spectrum simulation using SPEKTR [24]. Nevertheless, inaccuracies in spectrum or filtration modeling may introduce residual artifacts or global attenuation bias. In addition, while Gaussian splatting provides substantial efficiency, its performance depends on the density and placement of primitives, which may require tuning for extremely large volumes or highly anisotropic structures. Finally, our evaluation focuses on static CBCT scans; extending the method to dynamic or limited-angle acquisitions remains an open direction for future work.

8. Conclusion

We introduced a CBCT reconstruction framework that mitigates metal-induced beam-hardening artifacts through explicit polychromatic modeling and a Gaussian-splatting formulation. Our method jointly estimates both scene structure and energy-dependent attenuation parameters, all without relying on metal masks or paired supervision. Comprehensive experiments on synthetic and real scans demonstrate that our approach surpasses state-of-the-art baselines




Figure 5. Real-data reconstruction results. FDK [8] produces strong streaking and shading artifacts. LIMAR [15] reduces some of these effects but leaves noticeable residuals. OSCNet [33] fails to generalize to unseen acquisition conditions and produces incorrect intensity distributions. Polyner [35] suppresses streaks but oversmooths fine structural details. Park et al. [22] yields heavily blurred reconstructions and does not recover the correct intensity scale due to its simplified polychromatic model. In contrast, our method effectively removes metal-induced artifacts while preserving material-dependent density variations and high-frequency structural details across all specimens. (For visualization, Park et al. is displayed with an intensity cap of 0.1 for avocado and 0.2 for the remaining scenes.)

in both reconstruction fidelity and artifact reduction. To further facilitate research on artifact-robust reconstruction, we also provide new CBCT datasets featuring realistic metal-induced artifacts.

Acknowledgements

Min H. Kim acknowledges the Samsung Research Funding & Incubation Center (SRFC-IT2402-02), the Korea NRF grant (RS-2024-00357548), the

MSIT/IITP of Korea (RS-2022-00155620, RS-2024-00398830, RS-2024-00436680, and 2017-0-00072), the MSIT Advanced GPU Utilization Support Program, and Microsoft Research Asia, with additional support by the KAIST Jang Young Sil Fellow Program.

References

- [1] A. H. Andersen and A. C. Kak. Simultaneous algebraic reconstruction technique (sart): A superior implementation of the art algorithm. *Ultrasonic Imaging*, 6(1):81–94, 1984. 2
- [2] Samuel G Armato III, Geoffrey McLennan, Luc Bidaut, Michael F McNitt-Gray, Charles R Meyer, Anthony P Reeves, Binsheng Zhao, Denise R Aberle, Claudia I Henschke, Eric A Hoffman, et al. The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. *Medical physics*, 38(2):915–931, 2011. 5
- [3] F Edward Boas, Dominik Fleischmann, et al. Ct artifacts: causes and reduction techniques. *Imaging Med*, 4(2):229–240, 2012. 2, 4
- [4] Yuanhao Cai, Jiahao Wang, Alan Yuille, Zongwei Zhou, and Angtian Wang. Structure-aware sparse-view x-ray 3d reconstruction. In *CVPR*, 2024. 2
- [5] Kiseok Choi, Inchl Kim, Jaemin Cho, Hyeongjun Cho, and Min H. Kim. Splat-based metal artifact reduction in cone-beam ct via polychromatic modeling. *Computer Graphics Forum (Proc. EUROGRAPHICS 2026)*, 45(2):1–11, 2026. 5
- [6] Ezgi Demircan-Tureyen and Mustafa E Kamasak. A discretized tomographic image reconstruction based upon total variation regularization. *Biomedical Signal Processing and Control*, 38:44–54, 2017. 2
- [7] Yi Du, Xiangang Wang, Xincheng Xiang, and Zhouping Wei. Evaluation of hybrid sart+ os+ tv iterative reconstruction algorithm for optical-ct gel dosimeter imaging. *Physics in Medicine & Biology*, 61(24):8425, 2016. 2
- [8] Lee A Feldkamp, Lloyd C Davis, and James W Kress. Practical cone-beam algorithm. *Journal of the Optical Society of America A*, 1(6):612–619, 1984. 1, 2, 5, 6, 7, 8
- [9] Peter Gilbert. Iterative methods for the three-dimensional reconstruction of an object from projections. *Journal of Theoretical Biology*, 36(1):105–117, 1972. 2
- [10] Richard Gordon, Robert Bender, and Gabor T. Herman. Algebraic reconstruction techniques (art) for three-dimensional electron microscopy and x-ray photography. *Journal of Theoretical Biology*, 29(3):471–481, 1970. 2
- [11] Design Guide. Cuda c programming guide. *NVIDIA*, July, 29:31, 2013. 5
- [12] Nils Große Hokamp, Brendan Eck, Florian Siedek, Daniel Pinto Dos Santos, Jasmin A Holz, David Maintz, and Stefan Haneder. Quantification of metal artifacts in computed tomography: methodological considerations. *Quantitative imaging in medicine and surgery*, 10(5):1033, 2020. 2, 4
- [13] Avinash C Kak and Malcolm Slaney. *Principles of computerized tomographic imaging*. SIAM, 2001. 1
- [14] Avinash C. Kak and Malcolm Slaney. *Principles of Computerized Tomographic Imaging*. Society for Industrial and Applied Mathematics, 2001. 1
- [15] Willi A Kalender, Robert Hebel, and Johannes Ebersberger. Reduction of ct artifacts caused by metallic implants. *Radiology*, 164(2):576–577, 1987. 1, 2, 5, 6, 7, 8
- [16] Haofu Liao, Wei-An Lin, S. Kevin Zhou, and Jiebo Luo. Adn: Artifact disentanglement network for unsupervised metal artifact reduction. *IEEE Transactions on Medical Imaging*, 39(3):634–643, 2020. 2
- [17] Yong Long and Jeffrey A Fessler. Multi-material decomposition using statistical image reconstruction for spectral ct. *IEEE transactions on medical imaging*, 33(8):1614–1626, 2014. 4
- [18] Ting Luo, Changrong Shi, Xing Zhao, Yunsong Zhao, and Jinqiu Xu. Automatic synthesis of panoramic radiographs from dental cone beam computed tomography data. *PLoS one*, 11(6):e0156976, 2016. 5
- [19] Esther Meyer, Rainer Raupach, Michael Lell, Bernhard Schmidt, and Marc Kachelrieß. Normalized metal artifact reduction (nmar) in computed tomography. *Medical physics*, 37(10):5482–5493, 2010. 2, 5, 7
- [20] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 2
- [21] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, 2022. 2
- [22] Hyoung Suk Park, Jin Keun Seo, and Kiwan Jeon. Implicit neural representation-based method for metal-induced beam hardening artifact reduction in x-ray ct imaging. *Medical Physics*. 2, 3, 5, 7, 8
- [23] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. 2019. 5
- [24] J Punnoose, J Xu, A Sisniega, W Zbijewski, and JH Siewerdsen. Technical note: spektr 3.0—a computational tool for x-ray spectrum. *Medical Physics*, 43(8):4711–4717, 2016. 5, 7
- [25] Darius Rückert, Yuanhao Wang, Rui Li, Ramzi Idoughi, and Wolfgang Heidrich. Neat: Neural adaptive tomography. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022. 2
- [26] EB Saloman, JH Hubbell, and JH Scofield. X-ray attenuation cross sections for energies 100 ev to 100 kev and elements z= 1 to z= 92, 1988. 3, 4, 5
- [27] Mark Selles, Jochen AC van Osch, Mario Maas, Martijn F Boomsma, and Ruud HH Wellenberg. Advances in metal artifact reduction in ct images: A review of traditional and novel metal artifact reduction techniques. *European Journal of Radiology*, 170:111276, 2024. 2, 4
- [28] Jack Valentin. Basic anatomical and physiological data for use in radiological protection: reference values: Icrp publication 89: Approved by the commission in september 2001. *Annals of the ICRP*, 32(3-4):1–277, 2002. 4
- [29] Pieter Verboven, Bart Dequeker, Jiaqi He, Michiel Pieters, Leroi Pols, Astrid Tempelaere, Leen Van Doorselaer, Hans Van Cauteren, Ujjwal Verma, Hui Xiao, et al. www. x-plant.

- org-the ct database of plant organs. In *6th Symposium on X-ray Computed Tomography: Inauguration of the KU Leuven XCT Core Facility, Location: Leuven, Belgium*, 2022. 5
- [30] Hong Wang, Yuexiang Li, Haimiao Zhang, Jiawei Chen, Kai Ma, Deyu Meng, and Yefeng Zheng. Indudonet: An interpretable dual domain network for ct metal artifact reduction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 107–118. Springer, 2021. 2
- [31] Hong Wang, Yuexiang Li, Nanjun He, Kai Ma, Deyu Meng, and Yefeng Zheng. Dicednet: Deep interpretable convolutional dictionary network for metal artifact reduction in ct images. *IEEE Transactions on Medical Imaging*, 41(4):869–880, 2022. 5, 7
- [32] Hong Wang, Yuexiang Li, Deyu Meng, and Yefeng Zheng. Adaptive convolutional dictionary network for ct metal artifact reduction. In *The 31st International Joint Conference on Artificial Intelligence*. IEEE, 2022. 7
- [33] Hong Wang, Qi Xie, Dong Zeng, Jianhua Ma, Deyu Meng, and Yefeng Zheng. Oscnet: Orientation-shared convolutional network for ct metal artifact learning. *IEEE Transactions on Medical Imaging*, 43(1):489–502, 2024. 2, 5, 6, 7, 8
- [34] DR White, HQ Woodard, and SM Hammond. Average soft-tissue and bone models for use in radiation dosimetry. *The British journal of radiology*, 60(717):907–913, 1987. 4
- [35] Qing Wu, Lixuan Chen, Ce Wang, Hongjiang Wei, S Kevin Zhou, Jingyi Yu, and Yuyao Zhang. Unsupervised polychromatic neural representation for CT metal artifact reduction. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. 1, 2, 5, 6, 7, 8
- [36] Qing Wu, Xu Guo, Lixuan Chen, Yanyan Liu, Dongming He, Xudong Wang, Xueli Chen, Yifeng Zhang, S Kevin Zhou, Jingyi Yu, et al. Solving energy-independent density for ct metal artifact reduction via neural representation. *arXiv preprint arXiv:2405.07047*, 2024. 3, 5
- [37] Ru Xu. Total variation and tight frame iteration regularization based image reconstruction for short-scanning few-view computer tomography. 2019. 2
- [38] Ruyi Zha, Yanhao Zhang, and Hongdong Li. Naf: neural attenuation fields for sparse-view cbct reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 442–452. Springer, 2022. 2
- [39] Ruyi Zha, Tao Jun Lin, Yuanhao Cai, Jiwen Cao, Yanhao Zhang, and Hongdong Li. R²-gaussian: Rectifying radiative gaussian splatting for tomographic reconstruction. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 2, 3, 4, 5, 7