

Splat-based Metal Artifact Reduction in Cone-Beam CT via Polychromatic Modeling

Kiseok Choi 

Inchul Kim 

Jaemin Cho 

Hyeongjun Cho 

Min H. Kim 

KAIST

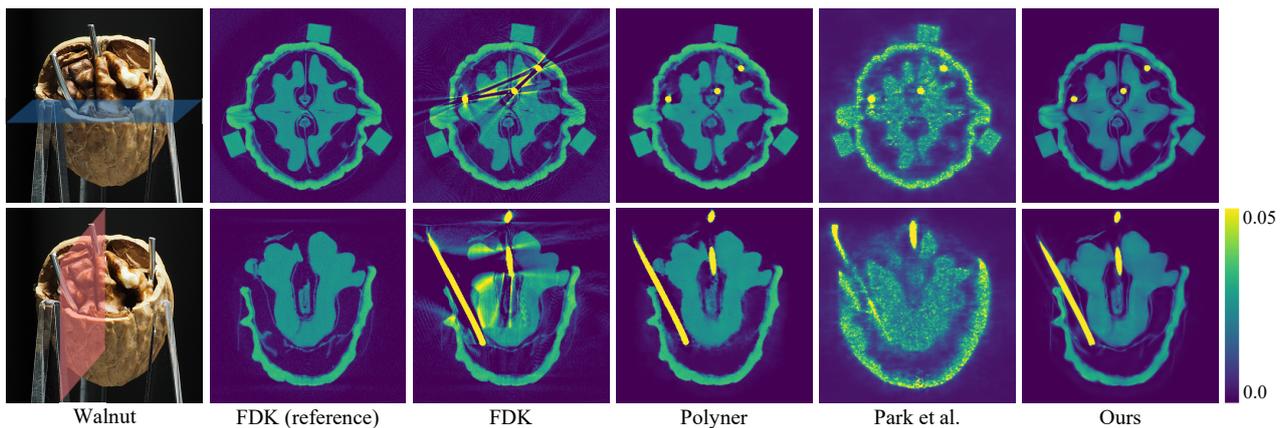


Figure 1: *Qualitative comparison of CBCT reconstruction results for a real walnut object with inserted metal pins. The leftmost column shows the optical photograph of the walnut, where the blue and red planes denote the positions of the horizontal and vertical cuts for visualization. Each reconstruction method is visualized with two slices: a horizontal slice (top row) and a vertical slice (bottom row) from the reconstructed volume. The FDK (reference) result corresponds to a baseline scan of the walnut without metal pins and serves as a ground-truth proxy free from metal artifacts. FDK [FDK84] applied to the metal-inserted scan exhibits severe beam hardening artifacts, including dark streaks and intensity distortions. Polyner [WCW*23] reduces some artifacts but shows noticeable blurring and oversmoothing of fine structures. Park et al. [PSJ] minimizes artifacts but considerably compromises the overall image structures, suffering from noise. Our method delivers the most faithful reconstruction by effectively suppressing artifacts while preserving structural detail in both axial and sagittal views, demonstrating its robustness for cone-beam CT with severe metal-induced beam hardening.*

Abstract

Cone-beam computed tomography (CBCT) enables volumetric reconstruction from X-ray projections, but suffers from severe artifacts—especially beam hardening—when imaging materials with high attenuation such as metals. These artifacts arise from the polychromatic nature of X-rays and are not properly addressed by conventional monochromatic reconstruction algorithms. While recent neural representation-based methods offer improved reconstruction quality, they are computationally expensive and often impractical for deployment. We propose a novel physics-inspired, self-calibrating metal artifact reduction method that efficiently reconstructs 3D CBCT volumes while correcting beam hardening artifacts. Our method integrates a polychromatic X-ray projection model, material-dependent attenuation profiles, and system response modeling into a Gaussian Splatting framework. Unlike prior work, we eliminate the need for manual metal masks or strong prior assumptions, and we optimize both reconstruction parameters and X-ray spectral characteristics jointly during training. We further introduce a high-fidelity synthetic CBCT dataset generation pipeline validated on Monte-Carlo x-ray simulation toolbox and release new datasets with severe metal-induced artifacts to support the community. This is the first splat-based method for reducing beam hardening in CBCT. Extensive experiments on both synthetic and real-world datasets demonstrate that our method outperforms state-of-the-art approaches in artifact suppression and reconstruction accuracy.

CCS Concepts

• **Social and professional topics** → *Medical technologies*; • **Computing methodologies** → *Reconstruction; Volumetric models*;

1. Introduction

X-rays are high-energy electromagnetic waves capable of penetrating various materials. As X-rays propagate through matter, they are attenuated according to the Beer-Lambert law, which models the exponential decay of energy along the path proportional to material density. X-ray computed tomography (CT) reconstructs internal volumetric structures by analyzing the attenuation. Among several CT configurations—parallel-beam, fan-beam, cone-beam, and helical-beam—cone-beam CT (CBCT) has attracted particular interest due to its ability to capture dense projection views in a single scan via conic emission and 2D planar detection, enabling efficient volumetric acquisition. This faster acquisition also contributes to reduced X-ray exposure, making CBCT an attractive modality for low-dose tomography applications.

While beam hardening artifacts are a common challenge in most X-ray CT systems due to the polychromatic nature of X-ray sources, where low-energy photons are absorbed more readily than high-energy ones, CBCT is particularly susceptible to such artifacts. Its wide cone angle and increased scatter amplify these nonlinear attenuation effects, leading to dark streaks, cupping artifacts, and geometric distortions that significantly degrade image quality (Figure 1). These artifacts not only impair diagnostic interpretation but also hinder downstream computational analysis such as segmentation or 3D modeling.

Traditional beam hardening correction (BHC) methods, including sinogram interpolation using predefined metal masks [KHE87, MRL*10], are limited by their reliance on heuristic preprocessing. Data-driven approaches [LLZL20, WLH*22, WLMZ22, WXZ*24] have demonstrated improved artifact removal using supervised learning or dictionary-based priors, but often require metal masks and paired training data, limiting generalization across materials and geometries. More recently, physically-motivated methods have been proposed to address beam hardening during reconstruction. Polyner [WCW*23] reconstructs energy-dependent voxel attenuations by assuming known X-ray spectra, and Park et al. [PSJ] leverage a simplified linear attenuation-energy model under fixed energy assumptions. While these methods show promise, they often depend on prior spectral calibration, operate under restrictive fan-beam geometries, or require computationally expensive neural volumetric representations, making them less suitable for practical CBCT systems.

In this work, we address the long-standing challenge of beam hardening artifact reduction in cone-beam CT by introducing a physics-inspired reconstruction framework that explicitly models both the polychromatic X-ray spectrum and material-specific attenuation behavior within a differentiable Gaussian Splatting (GS) framework. Our method incorporates an energy-aware attenuation model together with a self-calibrating system response formulation, enabling the joint optimization of volumetric attenuation parameters and system characteristics directly from measured projections. This unified design allows us to reconstruct high-quality attenuation volumes while effectively suppressing metal-induced beam hardening artifacts, even in low-dose imaging scenarios. Notably, our approach does not require prior knowledge of the X-ray spectrum or manual metal masking, substantially improving robustness and practicality for real-world CBCT systems. To the best of our knowledge, this is

the first splat-based reconstruction method specifically designed to address beam hardening artifacts in CBCT. Our code and data are available at our repository[†].

2. Related Work

Beam Hardening Artifact Reduction. Beam hardening artifacts are a persistent challenge in CT imaging, particularly when scanning high-attenuation materials such as metals. The underlying cause is the polychromatic nature of X-ray sources: lower-energy photons are more readily absorbed, leading to nonlinear attenuation that violates the assumptions of monochromatic models and results in dark streaks and geometric distortions.

Early correction methods such as LIMAR [KHE87] and NMAR [MRL*10] perform sinogram interpolation using predefined metal masks, but suffer from poor generalizability and reliance on accurate mask segmentation. With the rise of deep learning, data-driven BHC methods have emerged, including ADN [LLZL20], InDuDoNet [WLZ*21], DICDNet [WLH*22], ACDNet [WLMZ22], and OSCNet [WXZ*24]. These approaches attempt to learn artifact-free mappings but often require paired training data and metal masks, limiting their applicability to varied clinical scenarios.

More recent strategies tackle beam hardening directly during reconstruction by incorporating physical priors. Polyner [WCW*23] reconstructs energy-dependent attenuation volumes assuming a known X-ray spectrum and metal mask, using a NeRF-style representation. Park et al. [PSJ] adopt a simplified linear attenuation model and fixed energy assumption, avoiding prior calibration but still relying on heavy NeRF-based volumetric rendering. While these methods reduce artifacts to some extent, they are computationally expensive and often constrained to fan-beam geometries, which are incompatible with modern CBCT systems. In contrast, our method targets beam hardening reduction for cone-beam CT by integrating a polychromatic X-ray model and material-specific attenuation profiles into a lightweight Gaussian Splatting framework [ZLC*24]. Our system includes a self-calibrating component that jointly estimates spectral and attenuation parameters without requiring any prior spectral information or metal masks, offering both generality and practical deployability.

CT Reconstruction Paradigms. CT reconstruction algorithms can be broadly categorized into backprojection-based, iterative, and neural volumetric methods. Backprojection-based methods like FBP [Gul79] and FDK [FDK84] use analytical models rooted in the Fourier slice theorem, but they perform poorly when projection inconsistencies arise, such as those induced by metal artifacts. Iterative methods [AK84, Gil72] treat reconstruction as an inverse problem and refine the volume estimate through multiple projection-matching steps. Edge-preserving regularizations [SP08, LMFL12, BT09] such as total variation improve performance, but they still struggle under severe nonlinearity from beam hardening. Neural representation approaches like SAX-NeRF [CWY*24] apply neural rendering to CT reconstruction, achieving high accuracy but incurring substantial computational cost. Variants such as NAF [ZZL22] and multiresolution hash-based methods [MESK22] improve runtime at the

[†] <https://github.com/KAIST-VCLAB/ct-metal-reduction>

cost of fidelity. More recently, Gaussian Splatting (GS) [KKLD23] has been proposed as a neural-free, real-time approach to novel view synthesis. This method has been adapted to CT in works like R²-Gaussian [ZLC*24], which combines explicit Gaussian representations with the Beer-Lambert law to reconstruct attenuation volumes. However, R²-Gaussian assumes a monochromatic model and struggles under real-world polychromatic conditions. Our work builds upon R²-Gaussian by extending it with a physics-inspired polychromatic attenuation model and a self-calibrated system response formulation, enabling artifact-robust cone-beam CT reconstruction under severe beam hardening conditions.

3. Background

3.1. Polychromatic X-ray Projection Model

When an X-ray propagates through an object, its energy is attenuated according to the Beer-Lambert law:

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

where I_0 is the emission intensity of the source, $\mu(l)$ is the attenuation coefficient along the line segment dl of the path L , and I is the intensity measured at the detector. This expression can be rearranged into a log-transformed form:

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

The resulting quantity P is referred to as the projection image or sinogram. However, since real-world X-ray sources emit a spectrum of energies, and the attenuation coefficient μ depends on energy, Equation (2) cannot capture the polychromatic nature of X-ray propagation. To address this, the attenuation process must be expressed as an integral over the energy range \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 system response at energy level E , including the X-ray source distribution, filter characteristics, and detector sensitivity. The attenuation coefficient $\mu(l, E)$ is now energy-dependent and referred to as the linear attenuation coefficient (LAC). The measured polychromatic projection is obtained by integrating across the entire spectrum \mathcal{E} . The LAC is a material-dependent property that varies significantly with energy due to fundamental physical interactions, primarily the photoelectric effect and Compton scattering. This energy dependence of attenuation serves as a strong prior in estimating material-specific attenuation from polychromatic projections, as described in Equation (3).

3.2. R²-Gaussian

In R²-Gaussian [ZLC*24], the attenuation field in 3D space is modeled as a sum of anisotropic Gaussian primitives:

$$\mu(\mathbf{x}) = \sum_i^M \rho_i \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{p}_i)^T \Sigma_i^{-1} (\mathbf{x} - \mathbf{p}_i)\right), \quad (4)$$

where \mathbf{x} is the spatial query point, M is the number of Gaussians, ρ_i is the density of the i -th Gaussian, and \mathbf{p}_i , Σ_i denote its center

and covariance, respectively. The forward projection of a ray is then computed by integrating the contribution of all Gaussians along the ray path:

$$P(\hat{\mathbf{x}}) = \int_L \mu(l) dl \approx \sum_i^M \sqrt{\frac{2\pi|\tilde{\Sigma}_i|}{|\hat{\Sigma}_i|}} \rho_i \exp\left(-\frac{1}{2}(\hat{\mathbf{x}} - \hat{\mathbf{p}}_i)^T \hat{\Sigma}_i^{-1} (\hat{\mathbf{x}} - \hat{\mathbf{p}}_i)\right), \quad (5)$$

where $\hat{\mathbf{x}}$ is the projected pixel position, $\tilde{\Sigma}_i$ is the view-space covariance, $\hat{\Sigma}_i$ is the projected covariance, and $\hat{\mathbf{p}}_i$ is the projected Gaussian center. The parameters ρ_i , \mathbf{p}_i , and Σ_i are optimized such that the computed projections from Equation (5) match the observed projection images. After optimization, the final attenuation volume is recovered by evaluating the sum of the fitted Gaussian components.

4. Method

4.1. Overview

The overall framework is illustrated in Figure 2. Our method consists of two types of integrals: projection and reconstruction. In the projection phase, the Gaussians are aggregated using our differentiable polychromatic forward projector, and the resulting projections are compared against ground-truth transformed measurements acquired from the CT system. Importantly, while the per-Gaussian parameters are optimized, the global X-ray system response $\eta(\bar{E})$ is jointly optimized to achieve more effective reduction of metal artifacts. In the reconstruction phase, all optimized Gaussians are integrated to estimate the linear attenuation coefficient at each voxel location, which constitutes the final CT image output.

4.2. System Response and Polychromatic Attenuation

System Response. To accurately model polychromatic X-ray attenuation, it is essential to incorporate the spectral characteristics of the CT system. The system response is defined as the product of the X-ray source spectrum, filter transmission, and detector sensitivity [CAS*19], which are collectively represented by the function $\eta(E)$. The source spectrum mainly consists of a continuous Bremsstrahlung component with several discrete characteristic peaks [IG24]. Since the contribution of these peaks spans only a narrow energy range, they can be neglected in our case. The filter shifts the Bremsstrahlung spectrum toward higher energies [AS18]; however, for beam hardening correction, the filter thickness is typically small and its effect can be approximated as negligible. The detector sensitivity decreases gradually with photon energy [WMM*11], which, when combined with the Bremsstrahlung spectrum, results in an approximately linear decay. At higher energies, material attenuation becomes very weak, so the exponential term $\exp\left(-\int_L \mu(l, E) dl\right)$ approaches unity. Consequently, spectral components in this range cannot be effectively distinguished by gradient-based optimization. Since the spectrum in this region is also close to linearly decaying, we approximate it with a constant plateau, which improves both computational robustness and physical plausibility. Based on these observations, we approximate the system response using a soft, differentiable piecewise function composed of a linearly decaying component and a constant plateau:

$$\eta(\bar{E}) = \ell_{\bar{E}_{th},r}(\bar{E}) \cdot (1 - \sigma(\bar{E} - \bar{E}_{th})) + \ell_{\bar{E}_{th},r}(\bar{E}_{th}) \cdot \sigma(\bar{E} - \bar{E}_{th}), \quad (6)$$

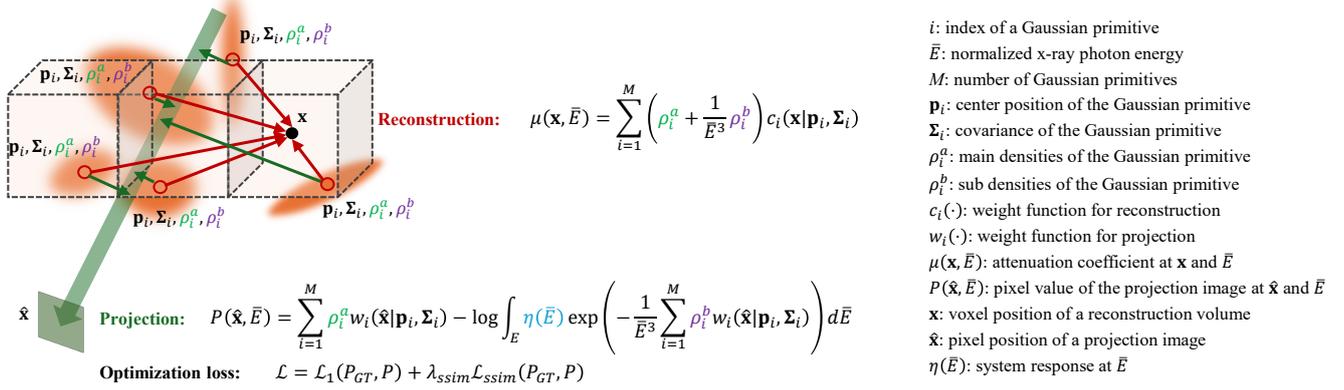


Figure 2: Method overview diagram. Our method consists of two types of integrals: projection and reconstruction. In the projection phase, the Gaussians are aggregated using our differentiable polychromatic forward projector, and the resulting projections are compared against ground-truth measurements acquired from the CT system. Importantly, while the per-Gaussian parameters are optimized, the global X-ray system response $\eta(\bar{E})$ is jointly optimized to achieve more effective reduction of metal artifacts. In the reconstruction phase, all optimized Gaussians are integrated to estimate the linear attenuation coefficient at each voxel location, which constitutes the final CT image output.

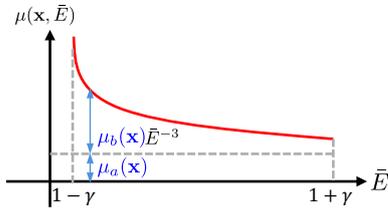


Figure 3: Illustration of our polychromatic attenuation model. The material-dependent linear attenuation coefficient $\mu(\mathbf{x}, \bar{E})$ is decomposed into two components: a constant Compton term $\mu_a(\mathbf{x})$ and an energy-dependent photoelectric term $\mu_b(\mathbf{x})\bar{E}^{-3}$. Each Gaussian primitive is assigned these parameters to model spatially varying material attenuation.

where $\ell_{\bar{E}_{th}, r}$ models the linear decay of the response with increasing energy, and σ is a sigmoid function that smoothly transitions from the decaying region to a constant value. For numerical stability and bounded computation, we normalize the energy E (in keV) to \bar{E} , which ranges over $[1 - \gamma, 1 + \gamma]$. The parameter r controls the relative intensity at the threshold energy \bar{E}_{th} compared to the peak, while γ defines the half-range of the normalized energy domain. Both r and \bar{E}_{th} are optimized during reconstruction. We adopt this simple model for two main reasons: (1) it reduces model complexity and mitigates overfitting, and (2) it decreases the number of estimated parameters, enabling more efficient optimization.

Polychromatic Attenuation. To account for the energy-dependent attenuation characteristics of materials, we incorporate the two dominant interaction mechanisms in CBCT energy ranges: the photoelectric effect and Compton scattering, as discussed in Section 3.1. The photoelectric effect dominates at lower energies and exhibits an inverse cubic dependence on energy, while Compton scattering is more prevalent at higher energies [Spr12]. We model the total attenuation coefficient accordingly:

$$\mu(\mathbf{x}, \bar{E}) = \mu_a(\mathbf{x}) + \frac{1}{\bar{E}^3} \mu_b(\mathbf{x}), \quad (7)$$

where $\mu_a(\mathbf{x})$ approximates the Compton scattering component and $\mu_b(\mathbf{x})$ represents the photoelectric effect. The illustration of our

model is shown in Figure 3. To ensure numerical stability during optimization, we approximate the Compton term with a constant value rather than an energy-dependent formulation. This simplification is justified because the Compton component remains nearly constant within the practical X-ray energy range used in medical CT, where the maximum photon energy is below 120 keV [Spr12, WCW*23]. We incorporate this model into the reconstruction framework by assigning each Gaussian primitive a density ρ_i that varies with energy:

$$\rho_i(\bar{E}) = \rho_i^a + \frac{1}{\bar{E}^3} \rho_i^b, \quad (8)$$

where ρ_i^a and ρ_i^b correspond to the terms μ_a and μ_b in Equation (7) respectively. Using the Gaussian-based attenuation coefficient (Equation (4)) at \bar{E} and Equation (8), we derive Equation (7):

$$\begin{aligned} \mu(\mathbf{x}, \bar{E}) &= \sum_i^M c_i(\mathbf{x}) \rho_i(\bar{E}) = \sum_i^M c_i(\mathbf{x}) \left(\rho_i^a + \frac{1}{\bar{E}^3} \rho_i^b \right) \\ &= \underbrace{\sum_i^M c_i(\mathbf{x}) \rho_i^a}_{\triangleq \mu_a(\mathbf{x})} + \frac{1}{\bar{E}^3} \underbrace{\sum_i^M c_i(\mathbf{x}) \rho_i^b}_{\triangleq \mu_b(\mathbf{x})}, \end{aligned} \quad (9)$$

where $c_i(\mathbf{x}) = \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{p}_i)^T \Sigma_i^{-1} (\mathbf{x} - \mathbf{p}_i)\right)$ that determines the shape of the i^{th} Gaussian. Since the attenuation coefficient in Equation (9) is a multiplication of the geometric Gaussian weight and the density term, it is regularized by the term $c_i(\mathbf{x})$ during optimization.

4.3. Projection

We define the polychromatic forward projection model for beam hardening correction. By substituting the polychromatic attenuation model from Equation (9) into the forward model (Equation (3)), we

obtain:

$$\begin{aligned} P &= -\log \int_{\bar{E}} \eta(\bar{E}) \exp\left(-\int_L \mu(l, \bar{E}) dl\right) d\bar{E} \\ &= -\log \int_{\bar{E}} \eta(\bar{E}) \exp\left(-\int_L \left(\mu_a(l) + \frac{1}{\bar{E}^3} \mu_b(l)\right) dl\right) d\bar{E} \quad (10) \\ &= \int_L \mu_a(l) dl - \log \int_{\bar{E}} \eta(\bar{E}) \exp\left(-\frac{1}{\bar{E}^3} \int_L \mu_b(l) dl\right) d\bar{E}, \end{aligned}$$

where \bar{E} is the integration range in the normalized energy levels (Figure 3). To compute the projection on the projected coordinate $\hat{\mathbf{x}}$ using the Gaussian parameters, we plug in Equation (5) into Equation (10):

$$P(\hat{\mathbf{x}}) \approx \sum_i^M w_i(\hat{\mathbf{x}}) \rho_i^a - \log \int_{\bar{E}} \eta(\bar{E}) \exp\left(-\frac{1}{\bar{E}^3} \sum_i^M w_i(\hat{\mathbf{x}}) \rho_i^b\right) d\bar{E}, \quad (11)$$

where $w_i(\hat{\mathbf{x}}) = \sqrt{2\pi|\tilde{\Sigma}_i|/|\hat{\Sigma}_i|} \exp\left(-\frac{1}{2}(\hat{\mathbf{x}} - \hat{\mathbf{p}}_i)^T \hat{\Sigma}_i^{-1} (\hat{\mathbf{x}} - \hat{\mathbf{p}}_i)\right)$ is the weight at the projected coordinate $\hat{\mathbf{x}}$ of the i^{th} Gaussian primitive. Note that since we approximate the Compton term as a constant variable $\mu_a(\mathbf{x})$, we can take that term out of the exponential function as shown in the third line of the equation above, which is significantly beneficial in numerical stability during optimization. To compute $P(\hat{\mathbf{x}})$, we partition the range of normalized energy levels \bar{E} by N to discretize Equation (11):

$$\begin{aligned} P(\hat{\mathbf{x}}) &= \sum_i^M w_i(\hat{\mathbf{x}}) \rho_i^a \\ &\quad - \log \sum_{k=0}^{N-1} \eta_k \exp\left(-\frac{1}{\left(\left(\frac{2k}{N-1} - 1\right)\gamma + 1\right)^3} \sum_i^M w_i(\hat{\mathbf{x}}) \rho_i^b\right), \quad (12) \end{aligned}$$

where η_k is the normalized system response weight such that $\sum_{k=0}^{N-1} \eta_k = 1$.

4.4. Optimization

We jointly optimize the system response parameters r and \bar{E}_{th} , along with the Gaussian parameters $\{\rho_i^a, \rho_i^b, \mathbf{p}_i, \Sigma_i \mid i = 1, \dots, M\}$. The total loss is a weighted combination of pixel-wise L_1 loss and structural similarity (SSIM):

$$\mathcal{L}_{total} = \mathcal{L}_1(P_{GT}, P) + \lambda_{ssim} \mathcal{L}_{ssim}(P_{GT}, P), \quad (13)$$

where λ_{ssim} is fixed to 0.25 in all experiments. We include the gradient derivation of P with respect to the Gaussian parameters in the supplemental material to support backpropagation through the R²-Gaussian CUDA pipeline. After optimization, the final voxelized attenuation field is obtained via Equation (9).

4.5. Implementation

Our method is implemented in PyTorch [PGM*19] with CUDA [Gui13]. We set the normalized energy half-range γ to $\frac{(E_{\max} - 10)}{(E_{\max} + 10)}$, assuming that the minimum photon energy in the spectrum is 10 keV and that E_{\max} (the maximum photon energy) is known in advance. The number of discrete energy components N used for spectrum integration is fixed to 15. The parameters for the system response are initialized as $r = 0.2$ and $\bar{E}_{th} = 1.0$. The

Gaussian parameters in our model are initialized following the same procedure as the base method, R²-Gaussian [ZLC*24]. Specifically, we first reconstruct an initial volume using the FDK from the input projection images. From the reconstructed volume, we randomly sample 50,000 positions in the voxel grid where the voxel intensity is greater than 0.005. One Gaussian is then initialized at each of these sampled positions. For each Gaussian, the value of ρ_a is initialized by multiplying the voxel intensity by 0.075, and ρ_b is initialized with the voxel intensity multiplied by 0.0075. The initialization of the Gaussian covariance follows the same routine used in the original R²-Gaussian algorithm [ZLC*24]. We adopt the original learning rate scheme, setting (ρ_a, ρ_b) to (0.01, 0.1), and (r, \bar{E}_{th}) to 0.001. Learning rates decay exponentially over 20k steps. Our density control scheme follows the baseline method with modification: We have found that a higher activation threshold on the gradient magnitude (5e-3, instead of 5e-5) enables more stable optimization under strong beam-hardening conditions. The other densification control remains unchanged.

5. Validation

5.1. Synthetic Dataset Generation

To validate our method, we implement polychromatic CBCT synthetic dataset generator based on TIGRE [BDHS16], incorporating Poisson and Gaussian noise to mimic realistic acquisition conditions. This simulation pipeline follows established practices in prior work [LLP*19, LLZL20, ZY18], and we further extend it to the CBCT configuration. For validating our simulator, we employ the OpenGATE package [ope25] that is based on Geant4 x-ray simulation library which supports Monte-Carlo (MC) projection. We first generate the X-ray spectrum using XrayPhysics library [KC23]. The spectrum ranges between 10–90keV with the take-off angle of 11 degree. The filter consists of aluminum with 0.5mm thickness. We set the material of the detector using GOS (O2SGd2) scintillator with 0.1mm thickness. We set the X-ray source's shape as isotropic to simulate a conic beam, and the half angle of the cone is set to 13 degree. The total number of particles is set to 1 billion for the MC integration. We use a synthetic scene as Figure 4 (a). In the scene, three aluminum rods are inserted into the large cylinder, and the remaining part of the cylinder is filled with water, and the rest of the space consists of air. The volume's physical size is 60 mm × 60 mm × 60 mm and the voxelized space has 256×256×256 voxels. We run both our projector and the MC-based one to generate raw detector signals and convert them into projection images by taking negative logarithm. Then, we compute the PSNR between them and get 32.94 dB (Figure 4 (b)), which shows strong agreement of our projection with the MC-based result.

Table 1: Validation of our system response model against alternative flexible models (scene: Lung).

System response	PSNR3D	SSIM3D
GMM(w/ 5 Gaussians)	26.51	0.989
Free parameters	27.74	0.991
Our model	29.19	0.993

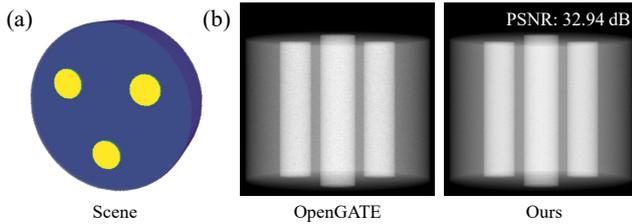


Figure 4: Validation of our synthetic projection image against one generated by Monte Carlo (MC) simulation using the OpenGATE package. A PSNR of 32.94 dB indicates strong agreement between the two results.

5.2. System Response Model Selection

We adopt the simple yet physics-inspired approximation model for the system response described in Section 4.2, which has only two parameters. To validate the superiority of this choice, we evaluate reconstruction performance using more flexible alternatives, including a Gaussian Mixture Model (GMM) and a model with free parameters. As shown in Table 1, our model outperforms these alternatives.

5.3. Generalizability to Multiple CT Systems

We validate the robustness of our method using synthetically generated projections that emulate different detector types and ranges of X-ray energies, as shown in Table 2. Our method shows consistent performance across various hardware configurations.

Table 2: Validation of our method’s generalizability under various acquisition conditions (scene: Lung).

Case	Voltage (kV)	Filter (mm)	Detector	PSNR3D	SSIM3D
1	90	Al 0.5	GOS	29.12	0.993
2	90	Al 1.0	GOS	30.12	0.994
3	90	Al 0.5	CsI	29.38	0.993
4	120	Al 0.5	GOS	28.91	0.990
5	60	Cu 0.5	GOS	31.17	0.995

5.4. Real Attenuation Estimation

In the real scene *Metal Rods* (the third row in Figure 8), water and aluminum rods are used. After running our reconstruction, we segment the reconstructed volume into regions corresponding to water and aluminum via thresholding. We then compare the reconstructed attenuation values for each segment with the ground-truth linear attenuation coefficients of water and aluminum from the NIST database [SHS88] (at the center energy 50keV). As shown in the Figure 5, the estimated values are well aligned with the real-world attenuation coefficients in the homogeneous regions.

6. Results

6.1. Dataset

We validate our method both on synthetic datasets and real datasets. For the synthetic evaluation, we utilize the LIDC dataset [AIMB*11], the X-plant dataset [VDH*22], and the ZCB100 dataset [LSZ*16]. We evaluate our method on three synthetic CBCT phantoms (*Lung*,

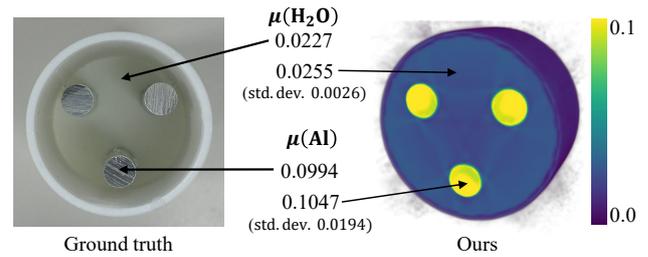


Figure 5: Validation of reconstructed attenuation coefficients against the NIST database [SHS88]. All values are reported in mm^{-1}

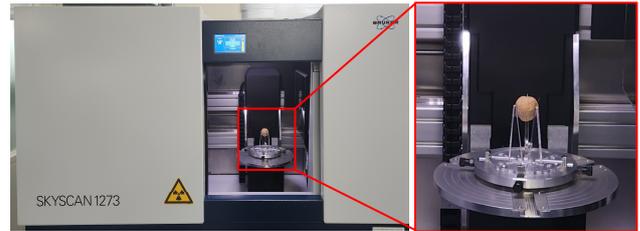


Figure 6: Bruker SKYSCAN 1273 CBCT system.

Teeth, and *Broccoli*) and five real-world scans: a walnut with metal pins (*Walnut*), a metal-embedded cylinder (*Metal Rods*), a chicken wrapped in metal wire (*Chicken*), a bell pepper with metal rivets (*Bell Pepper*), and a broccoli with metal rivets (*Broccoli*). Synthetic projections are generated using our validated simulator introduced in Section 5.1. Real projections are acquired using the Bruker SKYSCAN 1273 CBCT system (Figure 6) at a tube voltage of 90 kVp with a 0.5 mm aluminum filter. Each scan consists of 720 projections at a resolution of 512×512 . We scan each specimen both with and without inserted metal materials to observe metal-induced beam hardening effects. To validate the reconstruction accuracy of our method on real objects, we fabricated a 3D-printed cylindrical phantom containing three aluminum alloy rods to induce severe artifacts (see the lower left of Figure 8). Figure 8 (third column) shows the volume reconstruction exhibiting prominent beam hardening patterns. All projection images are resized to the resolution of 512×512 , and each scene has 720 projections.

6.2. Synthetic Experiment

We evaluate the performance of our method on three synthetic CBCT scenes—*Lung*, *Teeth*, and *Broccoli*—each containing embedded metallic objects made of iron (Fe), titanium (Ti), and aluminum (Al), respectively. The volumes are constructed with anatomical and organic structures, and cylindrical or spherical metal implants are synthetically inserted to induce beam hardening artifacts. CBCT projections are simulated under realistic polychromatic conditions using XrayPhysics [CAS*19]. Figure 7 compares the reconstruction quality of various methods on these datasets. FDK [FDK84] exhibits severe streaking artifacts, even in slices not directly intersected by metal due to 3D projection effects inherent in cone-beam geometry. While methods such as LIMAR [KHE87], NMAR [MRL*10], and Polymer [WCW*23] partially mitigate artifacts, they still suffer from residual distortions or blurring. Park et al. [PSJ], originally

developed for fan-beam CT, struggle to generalize to cone-beam settings, resulting in poor reconstruction quality. In contrast, our method produces high-fidelity reconstructions with accurate volume recovery and significantly suppressed metal-induced artifacts across all scenes. Table 3 quantitatively reports 3D PSNR and SSIM metrics. Our method consistently achieves the highest scores in all scenes, surpassing existing classical, learning-based, and physics-guided baselines. These results highlight the effectiveness of jointly modeling polychromatic attenuation and system response within a physics-inspired reconstruction framework.

Table 3: Quantitative results on synthetic CBCT datasets. 3D PSNR and SSIM scores for the Lung, Teeth, and Broccoli scenes corresponding to Figure 7. Our method achieves the highest accuracy across all datasets, demonstrating strong artifact suppression and structural fidelity.

Method	Lung		Teeth		Broccoli	
	PSNR3D	SSIM3D	PSNR3D	SSIM3D	PSNR3D	SSIM3D
FDK [FDK84]	19.04	0.924	32.84	0.917	11.66	0.892
LIMAR [KHE87]	19.83	0.935	32.97	0.921	13.60	0.922
NMAR [MRL*10]	20.49	0.942	33.05	0.923	15.81	0.951
ACDNet [WLMZ22]	13.88	0.644	26.28	0.602	14.37	0.943
DICDNet [WLH*22]	13.48	0.605	25.33	0.423	13.93	0.933
OSCNNet [WXZ*24]	13.28	0.592	25.34	0.422	13.83	0.931
Polyner [WCW*23]	20.43	0.955	30.18	0.959	18.88	0.987
Park et al. [PSJ]	13.59	0.699	23.15	0.524	5.29	0.526
Ours	29.19	0.993	35.97	0.992	23.22	0.995

6.3. Real Experiment

To evaluate the practical applicability of our method, we conduct experiments on real CBCT measurements acquired with the aforementioned X-ray scanner. We test five physical scenes as described in Section 6.1. All scenes except *Metal Rods* are also scanned separately without metal for reference, while the artifact-free version of the cylindrical phantom is synthesized in simulation for ground-truth comparison. Note that the outer shell of the cylindrical phantom is made of polylactide, whose attenuation is very similar to that of water. Figure 8 presents qualitative comparisons across multiple reconstruction methods. FDK and conventional algorithms (e.g., LIMAR [KHE87], NMAR [MRL*10]) exhibit severe streaking and low-frequency distortions that obscure important structures in both axial and sagittal views. The data-driven method [WXZ*24] fails to generalize to volume reconstruction due to out-of-distribution issues. Polyner [WCW*23] achieves reasonable reconstructions but leaves residual artifacts, while Park et al. [PSJ] struggles to adapt to the CBCT setting. In contrast, our method consistently produces volumetrically accurate reconstructions with sharp structural boundaries and substantially fewer artifacts across objects. These results highlight the robustness and real-world transferability of our framework, even under hardware and acquisition conditions not seen during development. Figure 9 further shows results for three organic objects across joint reconstruction and metal artifact reduction methods. Odd columns display representative slices reconstructed by each method, with red boxes marking regions of interest; even columns show the corresponding zoom-in patches. Additional slices are provided in the supplemental material. Our method produces clearer reconstructions with fewer artifacts compared to prior approaches.

6.4. Computation Time

We compare the computation time of our method with state-of-the-art joint reconstruction and metal artifact correction approaches in Table 4. All measurements are performed on an Intel Xeon 4214R CPU and an NVIDIA RTX A6000 GPU. Our method achieves an order-of-magnitude speedup across all scenes.

Table 4: Comparisons of computation time. Our method is faster than the other joint reconstruction and metal artifact reduction methods.

	Scene	Polyner [WCW*23]	Park et al. [PSJ]	Ours
Synthetic	Broccoli	1h 35m 58s	1h 33m 41s	16m 08s
	Lung	1h 18m 19s	1h 15m 35s	22m 16s
	Teeth	1h 19m 50s	1h 18m 34s	23m 38s
Real	Walnut	1h 11m 58s	1h 10m 20s	17m 51s
	Metal Rods	1h 27m 14s	1h 37m 56s	30m 41s
	Broccoli	2h 06m 55s	1h 59m 16s	16m 16s
	Bell Pepper	2h 01m 00s	2h 10m 11s	15m 33s
	Chicken	2h 02m 55s	2h 09m 43s	16m 35s

7. Ablation Study

7.1. Component Contribution

To evaluate the effectiveness of each core component in our framework, we perform an ablation study on the *Lung* scene by systematically adding the proposed modules to the baseline R^2 -Gaussian algorithm. Specifically, we investigate the individual and combined impacts of the system response model and the polychromatic attenuation model. As shown in Table 5, incorporating the system response model alone yields a modest improvement in both PSNR and SSIM, suggesting that response-aware weighting improves projection fidelity. In contrast, the addition of the polychromatic attenuation model leads to a substantial gain in reconstruction quality, highlighting the importance of accurate energy-dependent material modeling in suppressing beam hardening artifacts. When both modules are integrated, we observe the highest reconstruction performance, achieving a PSNR of 29.19 and an SSIM of 0.993, thereby confirming the complementary nature of the two components. Refer to the supplemental material for qualitative results of each component’s contribution.

Table 5: Component contribution analysis. We examine the individual contributions of the system response model and polychromatic attenuation model. Both components contribute positively to reconstruction accuracy, with their combination yielding the best overall performance (scene: Lung).

Module	PSNR3D	SSIM3D
Baseline	20.68	0.954
Baseline with response model	21.21	0.963
Baseline with attenuation model	27.73	0.991
Baseline with both	29.19	0.993

7.2. Number of Spectrum Components

We conduct an experiment to determine the optimal number of spectrum components (N). As shown in Table 6, the best reconstruction quality is achieved when $N = 15$. When N exceeds 15, the performance starts getting saturated and slightly degraded.

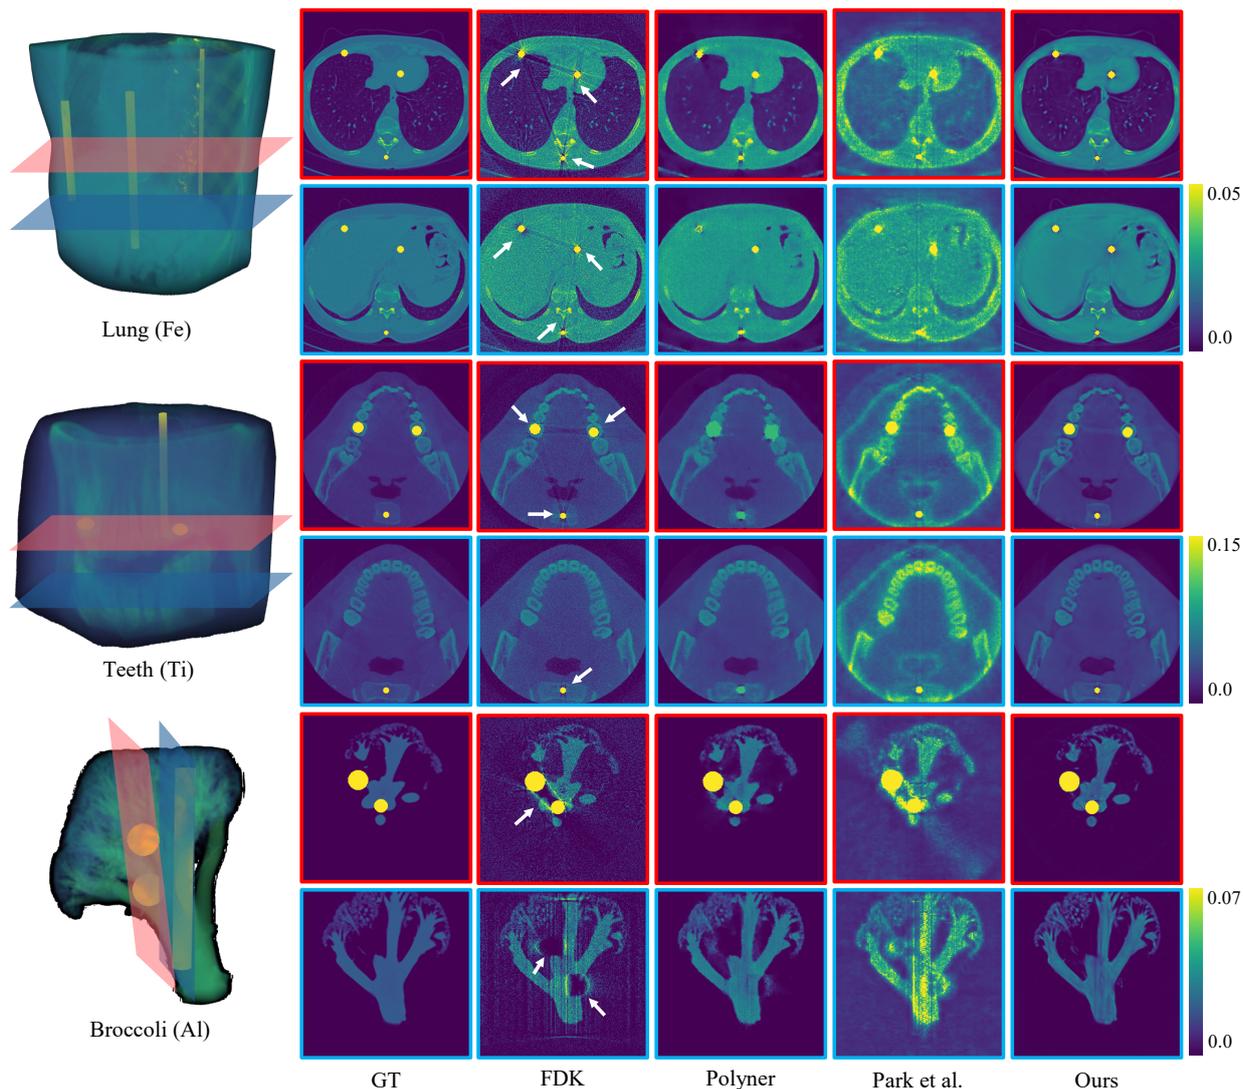


Figure 7: Qualitative comparison on synthetic CBCT datasets. Reconstruction results on Lung, Teeth, and Broccoli scenes with metallic implants of different materials. Our method achieves superior artifact suppression and structural fidelity compared to prior methods.

Table 6: Performance impact of the number of spectral components. The best performance is achieved at $N = 15$ (scene: Lung).

# of Spectrum Components (N)	PSNR3D	SSIM3D
7	27.86	0.988
15	29.19	0.993
31	28.70	0.993
63	28.25	0.992

Table 7: Comparison of the baseline [ZLC*24] and our method on PSNR3D and SSIM3D across different scenes.

Scene (synthetic)	Baseline [ZLC*24]		Ours	
	PSNR3D	SSIM3D	PSNR3D	SSIM3D
Broccoli	17.67	0.984	23.22	0.995
Lung	20.68	0.954	29.19	0.993
Teeth	33.53	0.964	35.97	0.992

7.3. Performance Comparison to Baseline

We compare our method against the baseline approach [ZLC*24] on scenes containing metal artifacts. By incorporating a polychromatic attenuation model and system response, our method consistently achieves superior performance across different scenes, as shown in Table 7. We also compare the computation time between our method and the baseline [ZLC*24] as Table 8. Due to the additional computation of the polychromatic model in our forward projection

compared to the baseline method, each optimization iteration takes slightly longer. To mitigate this overhead, we implement several efficiency-oriented strategies: we analytically simplify the forward projection and its corresponding backward gradient computations, and eliminate redundant operations to reduce computation costs. For faster convergence, we conduct extensive experiments to identify optimal learning rates for ρ_a and ρ_b . As a result, despite the per-iteration overhead, our model converges within 20K iterations,

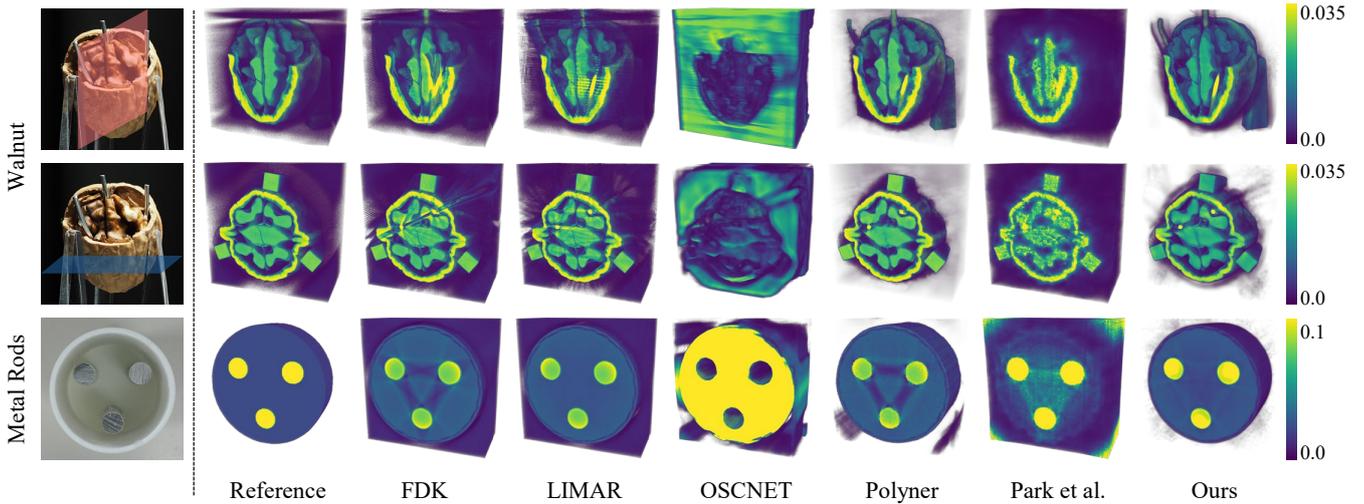


Figure 8: Qualitative comparison on real CBCT datasets. Reconstruction results for walnut and Metal Rods phantoms. Reference volumes are obtained from either separate scans or synthetic generation. Our method achieves clearer structures and stronger artifact suppression compared to prior approaches.

whereas the baseline method typically requires around 30K. This demonstrates that, with proper hyperparameter tuning, our model achieves comparable or even faster convergence.

7.4. Sensitivity on γ

We normalize the energy spectrum so that its center becomes 1 and the minimum and maximum values are set to $1 - \gamma$ and $1 + \gamma$, respectively. If γ is not set correctly and does not match the actual spectrum range, reconstruction quality may deteriorate. We evaluate the sensitivity by running our method using projections (acquired under a 10–90 keV spectrum) with various γ values corresponding to maximum energies of 60, 90, 120, and 150 keV. As shown in Table 9, the model shows the best performance when γ is set near the true spectrum range (i.e., 90 keV), which validates the significance of the choice of γ .

Table 8: Convergence time comparison between the baseline and our method across different scenes.

Type	Scene	Baseline [ZLC*24]	Ours
Synthetic	Broccoli	18m 44s	16m 08s
	Lung	20m 17s	22m 16s
	Teeth	21m 35s	23m 38s
Real	Walnut	20m 52s	17m 51s
	Metal Rods	27m 07s	30m 41s
	Broccoli	17m 52s	16m 16s
	Bell Pepper	19m 53s	15m 33s
	Chicken	21m 59s	16m 35s

Table 9: Effect of varying γ and corresponding maximum energy E_{\max} on reconstruction quality (scene: Broccoli).

γ	Corresponding E_{\max} (keV)	PSNR3D	SSIM3D
0.714	60	21.85	0.9943
0.800	90	23.22	0.9953
0.846	120	22.83	0.9946
0.875	150	22.38	0.9939

8. Discussion

Limitations. Due to limitations in accessing clinical CT datasets, we were unable to conduct comprehensive evaluations across diverse real medical use cases. Instead, we have validated our method on synthetic human-body datasets containing metal implants (e.g., in soft tissue and dental regions). Additionally, we have tested on a real-world dataset where metallic wires were inserted into a chicken to induce metal artifacts in biological tissues. These experiments show that our method generalizes well to various anatomical objects and metals. We are actively seeking access to clinical data to further investigate the generalizability of our method under real-world scenarios for future work.

While our method is built upon a physics-inspired polychromatic projection model, its underlying assumptions may break down under extreme conditions such as photon starvation, and it relies on approximations of both the system response and material attenuation models. As a result, the reconstructed linear attenuation coefficients may deviate from the true physical values for specific materials. In particular, two limitation modes may arise, both of which are inherent to physics-based projection models rather than specific to our approach. First, under extremely large attenuation, the term $\exp(-\int \mu, dl)$ can fall below floating-point precision, causing sinogram values to collapse toward zero and gradients to vanish, a well-known phenomenon in X-ray CT that typically occurs when imaging very large or very dense metals and would require higher-precision numerics or specialized exponential-stabilization schemes to address. Second, our model assumes that the energy-dependent attenuation coefficient $\mu(E)$ is smooth and can be approximated as $\mu(E) = a + bE^{-3}$; while this assumption holds for medically relevant materials such as Ti and Fe, it becomes less accurate for high-Z materials with sharp K-edges, such as Pb or Sn. Since such materials are not used in clinical implants, this limitation does not affect our intended application domain.

Comparison Details. To enable fair comparison, we re-implemented Polyner [WCW*23] and Park et al. [PSJ], as no CBCT-compatible

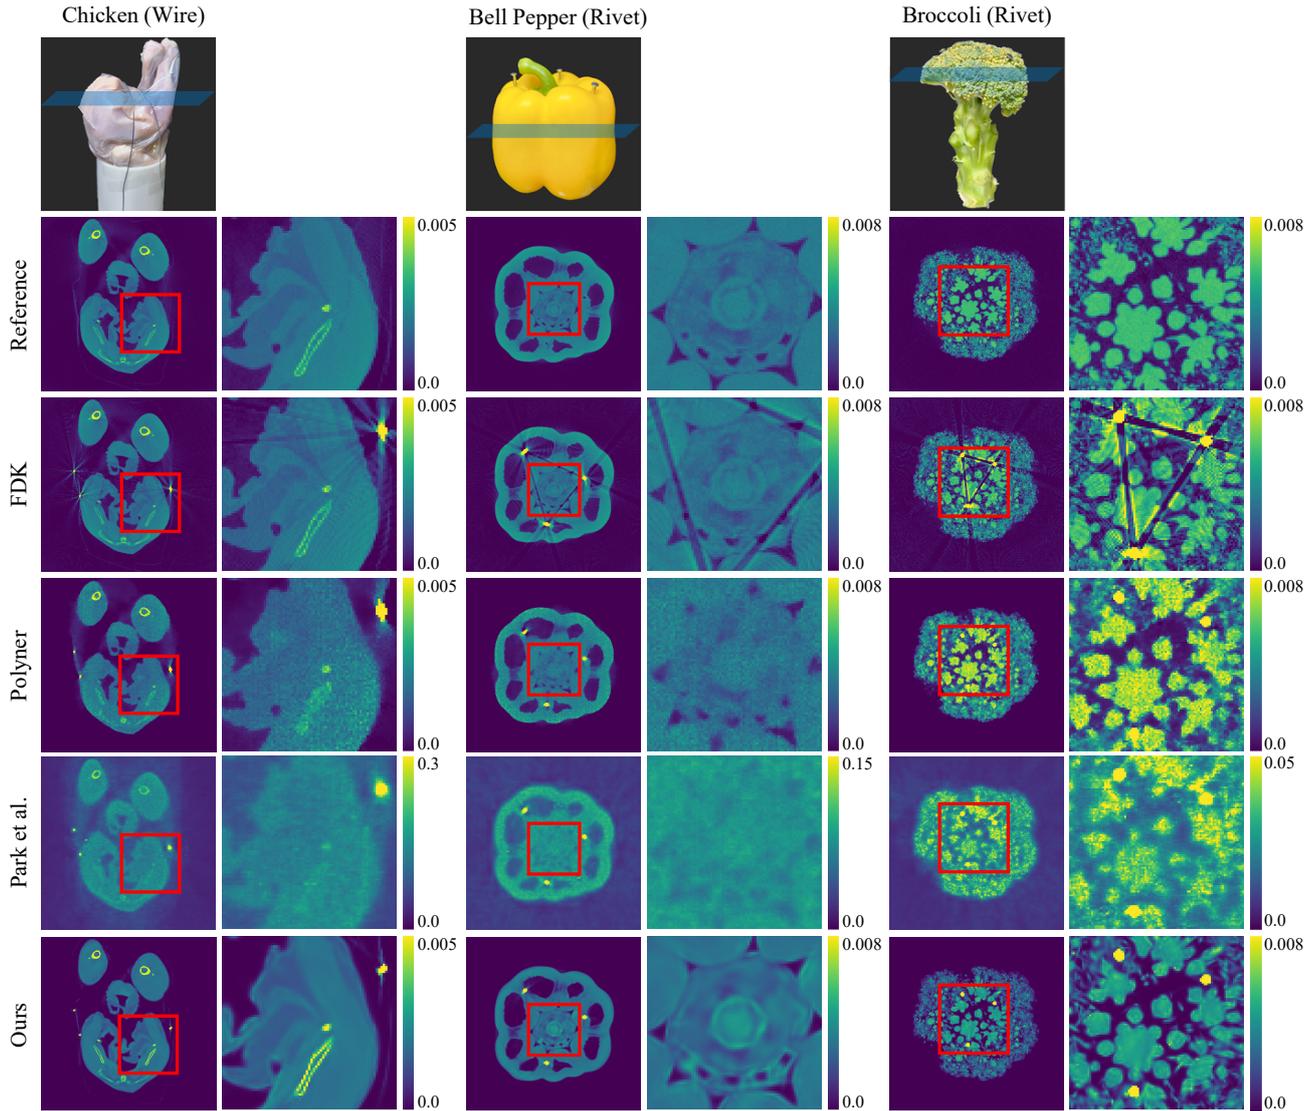


Figure 9: Qualitative comparison on real CBCT datasets. The slicing locations are displayed on each picture of the real objects in the first row. Odd columns show representative slices reconstructed by each method, with red boxes highlighting regions of interest. Even columns display the corresponding zoom-in patches. Additional slices are provided in the supplemental document. Our method produces clearer reconstructions with fewer artifacts compared to prior approaches.

public code was available. After careful tuning, our Polyner reproduction matched or exceeded the original performance in CBCT settings. In contrast, Park et al.’s fan-beam-based method showed limited generalization to CBCT, highlighting the challenge of adapting fan-beam approaches to cone-beam geometries under severe artifacts. Details of the comparison implementation are provided in the supplemental document.

9. Conclusion

We have presented a CBCT reconstruction method that reduces metal-induced beam hardening artifacts via polychromatic modeling and Gaussian splatting. Our framework jointly estimates volume structure, system response, and attenuation parameters—without

requiring metal masks, paired supervision, nor spectral priors. Extensive evaluations on synthetic and real data confirm that our method outperforms state-of-the-art baselines in reconstruction accuracy and artifact suppression. We also release new CBCT datasets with realistic artifacts to support future research in artifact-resilient reconstruction.

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), and Microsoft Research Asia.

References

- [AIBM*11] ARMATO III S. G., MCLENNAN G., BIDAUT L., MCNITT-GRAY M. F., MEYER C. R., REEVES A. P., ZHAO B., ABERLE D. R., HENSCHKE C. I., HOFFMAN E. A., 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 (2011), 915–931. 6
- [AK84] ANDERSEN A. H., KAK A. C.: Simultaneous algebraic reconstruction technique (sart): A superior implementation of the art algorithm. *Ultrasonic Imaging* 6, 1 (1984), 81–94. 2
- [AS18] AHMED O., SONG Y.: A review of common beam hardening correction methods for industrial x-ray computed tomography. *Sains Malaysiana* 47, 8 (2018), 1883–1890. 3
- [BDHS16] BIGURI A., DOSANJH M., HANCOCK S., SOLEIMANI M.: Tigre: a matlab-gpu toolbox for cbct image reconstruction. *Biomedical Physics & Engineering Express* 2, 5 (2016), 055010. 5
- [BT09] BECK A., TEOULLE M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.* 2, 1 (2009), 183–202. 2
- [CAS*19] CHAMPLEY K. M., AZEVEDO S. G., SEETHO I. M., GLENN S. M., McMICHAEL L. D., SMITH J. A., KALLMAN J. S., BROWN W. D., MARTZ H. E.: Method to extract system-independent material properties from dual-energy x-ray ct. *IEEE Transactions on Nuclear Science* 66, 3 (2019), 674–686. 3, 6
- [CWY*24] CAI Y., WANG J., YUILLE A., ZHOU Z., WANG A.: Structure-aware sparse-view x-ray 3d reconstruction. In *CVPR* (2024). 2
- [FDK84] FELDKAMP L. A., DAVIS L. C., KRESS J. W.: Practical cone-beam algorithm. *Journal of the Optical Society of America A* 1, 6 (1984), 612–619. 1, 2, 6, 7
- [Gil72] GILBERT P.: Iterative methods for the three-dimensional reconstruction of an object from projections. *Journal of Theoretical Biology* 36, 1 (1972), 105–117. 2
- [Gui13] GUIDE D.: Cuda c programming guide. *NVIDIA, July 29* (2013), 31. 5
- [Gul79] GULLBERG G. T.: The reconstruction of fan-beam data by filtering the back-projection. *Computer Graphics and Image Processing* 10, 1 (1979), 30–47. 2
- [IG24] INIEWSKI K., GADEY H.: *Emerging Radiation Detection: Technology and Applications*. Springer Nature, 2024. 3
- [KC23] KIM H., CHAMPLEY K.: Differentiable forward projector for x-ray computed tomography. *arXiv preprint arXiv:2307.05801* (2023). 5
- [KHE87] KALENDER W. A., HEBEL R., EBERSBERGER J.: Reduction of ct artifacts caused by metallic implants. *Radiology* 164, 2 (1987), 576–577. 2, 6, 7
- [KKLD23] KERBL B., KOPANAS G., LEIMKÜHLER T., DRETTAKIS G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* 42, 4 (2023), 139–1. 3
- [LLP*19] LIN W.-A., LIAO H., PENG C., SUN X., ZHANG J., LUO J., CHELLAPPA R., ZHOU S. K.: Dudonet: Dual domain network for ct metal artifact reduction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 10512–10521. 5
- [LLZL20] LIAO H., LIN W.-A., ZHOU S. K., LUO J.: Adn: Artifact disentanglement network for unsupervised metal artifact reduction. *IEEE Transactions on Medical Imaging* 39, 3 (2020), 634–643. 2, 5
- [LMFL12] LIU Y., MA J., FAN Y., LIANG Z.: Adaptive-weighted total variation minimization for sparse data toward low-dose x-ray computed tomography image reconstruction. *Physics in Medicine & Biology* 57, 23 (nov 2012), 7923. 2
- [LSZ*16] LUO T., SHI C., ZHAO X., ZHAO Y., XU J.: Automatic synthesis of panoramic radiographs from dental cone beam computed tomography data. *PLoS one* 11, 6 (2016), e0156976. 6
- [MESK22] MÜLLER T., EVANS A., SCHIED C., KELLER A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.* 41, 4 (July 2022), 102:1–102:15. 2
- [MRL*10] MEYER E., RAUPACH R., LELL M., SCHMIDT B., KACHELRIBB M.: Normalized metal artifact reduction (nmar) in computed tomography. *Medical physics* 37, 10 (2010), 5482–5493. 2, 6, 7
- [ope25] Opengate documentation, 2025. URL: <https://opengate.readthedocs.io/en/latest/introduction.html>. 5
- [PGM*19] PASZKE A., GROSS S., MASSA F., LERER A., BRADBURY J., CHANAN G., KILLEEN T., LIN Z., GIMELSHEIN N., ANTIGA L., DESMAISON A., KOPF A., YANG E., DEVITO Z., RAISON M., TEJANI A., CHILAMKURTHY S., STEINER B., FANG L., BAI J., CHINTALA S.: Pytorch: An imperative style, high-performance deep learning library. 5
- [PSJ] PARK H. S., SEO J. K., JEON K.: Implicit neural representation-based method for metal-induced beam hardening artifact reduction in x-ray ct imaging. *Medical Physics*. 1, 2, 6, 7, 9
- [SHS88] SALOMAN E., HUBBELL J., SCOFIELD J.: X-ray attenuation cross sections for energies 100 ev to 100 kev and elements z= 1 to z= 92, 1988. 6
- [SP08] SIDKY E. Y., PAN X.: Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in Medicine & Biology* 53, 17 (2008), 4777. 2
- [Spr12] SPRAWLS P.: Interaction of radiation with matter. *Physical Principles of Medical Imaging* (2012), 141–157. 4
- [VDH*22] VERBOVEN P., DEQUEKER B., HE J., PIETERS M., POLS L., TEMPELAERE A., VAN DOORSELAER L., VAN CAUTEREN H., VERMA U., XIAO H., 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). 6
- [WCW*23] WU Q., CHEN L., WANG C., WEI H., ZHOU S. K., YU J., ZHANG Y.: Unsupervised polychromatic neural representation for CT metal artifact reduction. In *Thirty-seventh Conference on Neural Information Processing Systems* (2023). 1, 2, 4, 6, 7, 9
- [WLH*22] WANG H., LI Y., HE N., MA K., MENG D., ZHENG Y.: Dcdnet: Deep interpretable convolutional dictionary network for metal artifact reduction in ct images. *IEEE Transactions on Medical Imaging* 41, 4 (2022), 869–880. 2, 7
- [WLMZ22] WANG H., LI Y., MENG D., ZHENG Y.: Adaptive convolutional dictionary network for ct metal artifact reduction. In *The 31st International Joint Conference on Artificial Intelligence* (2022), IEEE. 2, 7
- [WLZ*21] WANG H., LI Y., ZHANG H., CHEN J., MA K., MENG D., ZHENG Y.: Indudonet: An interpretable dual domain network for ct metal artifact reduction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (2021), Springer, pp. 107–118. 2
- [WMM*11] WANG X., MEIER D., MIKKELSEN S., MAEHLUM G., WAGENAAR D., TSUI B., PATT B., FREY E.: Microct with energy-resolved photon-counting detectors. *Physics in Medicine & Biology* 56, 9 (2011), 2791. 3
- [WXZ*24] WANG H., XIE Q., ZENG D., MA J., MENG D., ZHENG Y.: Oscnet: Orientation-shared convolutional network for ct metal artifact learning. *IEEE Transactions on Medical Imaging* 43, 1 (2024), 489–502. 2, 7
- [ZLC*24] ZHA R., LIN T. J., CAI Y., CAO J., ZHANG Y., LI H.: R²-gaussian: Rectifying radiative gaussian splatting for tomographic reconstruction. In *Advances in Neural Information Processing Systems (NeurIPS)* (2024). 2, 3, 5, 8, 9
- [ZY18] ZHANG Y., YU H.: Convolutional neural network based metal artifact reduction in x-ray computed tomography. *IEEE transactions on medical imaging* 37, 6 (2018), 1370–1381. 5
- [ZZL22] ZHA R., ZHANG Y., LI H.: Naf: neural attenuation fields for sparse-view cbct reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (2022), Springer, pp. 442–452. 2