KB-SPLATTING: DIFFERENTIABLE KAISER-BESSEL SPLATTING FOR REAL X-RAY PROJECTION-ENABLED CT RECONSTRUCTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Sparse-view cone-beam computed tomography (CBCT) can reduce harmful X-ray patient exposure, but recovering an accurate three-dimensional attenuation field from limited X-ray projections remains severely ill-posed. To address this, recent differentiable CT reconstruction methods have adopted anisotropic Gaussian primitives. Yet CT volumes contain sharp attenuation transitions, whereas Gaussian profiles are smooth and couple spatial extent with kernel decay. Accurately representing these transitions can therefore require several small primitives, driven by adaptive control heuristics. Thus increasing memory and optimization cost. These methods are also evaluated mainly on synthetic digitally reconstructed radiographs (DRRs), so their performance on physical clinical scanners remains unclear. Consequently, we introduce KB-Splatting, a differentiable CT reconstruction framework based on anisotropic generalized Kaiser-Bessel primitives. Our KB primitives have exact compact support per kernel and independently control support size and profile concentration, decoupling where attenuation is represented from how it varies within that tissue region. Their analytic cone-beam projection also allows us to optimize these kernel parameters directly from real X-ray projections. We also introduce a simple reconstruction-guided population control scheme that uses volumetric residuals and local parent kernel selection to add and relocate primitives in poorly reconstructed regions. We evaluate KB-Splatting on established DRR benchmarks and introduce a paired multi-system CBCT dataset containing real measured projections, documented acquisition geometries with coordinate-convention checks, and reference CT volumes from one commercial CT system and one experimental laboratory system. Experiments compare KB-Splatting with neural-field and Gaussian-splatting baselines, where it surpasses current state-of-the-art (SOTA) methods in attenuation PSNR, with a mean gain of 0.93 dB at 2x training speed. Finally, we extend KB-Splatting to RGB novel-view synthesis, testing the same representation under both additive X-ray formation and occlusion-aware rendering.

1 INTRODUCTION

Cone-beam computed tomography (CBCT) recovers a three-dimensional attenuation field from X-ray projections acquired around a patient. Reducing the number of projections can lower radiation exposure and shorten acquisition, but it also removes measurements from an already ill-posed inverse problem. The resulting reconstructions are more sensitive to noise, angular undersampling, and model error, often producing streaks, blurred boundaries, or missing structures. Classical filtered backprojection methods such as Feldkamp-Davis-Kress (FDK) are fast but degrade sharply in this regime (Feldkamp et al., 1984). Per-scan neural fields instead fit attenuation directly to observed projections (Zha et al., 2022; Cai et al., 2024b), showing that the representation determines structural recovery.

Explicit splatting offers a faster alternative to sampled neural fields. In RGB rendering, 3D Gaussian Splatting combines anisotropic primitives, differentiable rasterization, and adaptive density control (Kerbl et al., 2023). R^2-Gaussian adapts this formulation to tomography with X-ray line integration, FDK initialization, and differentiable voxelization (Zha et al., 2024), demonstrating that

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

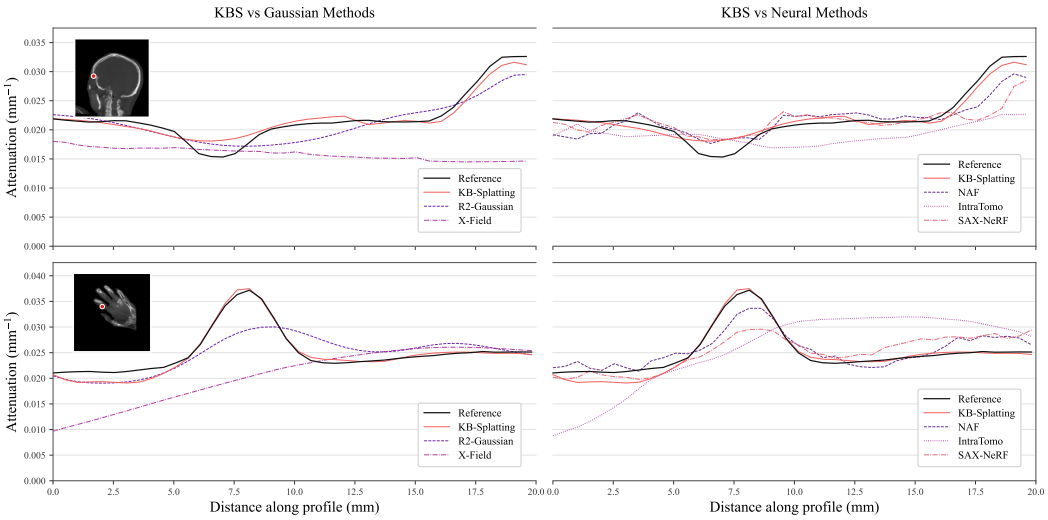

Figure 1: Local attenuation profiles on our cadaver skull (top) and hand (bottom) data. KB-Splatting better follows abrupt tissue transitions than Gaussian and ellipsoidal (left) and neural-field (right) baselines. Red insets mark the sampled locations.

trainable explicit primitives can reconstruct attenuation volumes efficiently. Once the image formation model is corrected, however, a basic question remains: *is the Gaussian profile itself the right primitive for sparse-view CT?*

A Gaussian's scale couples spatial coverage to profile decay. Large Gaussians cover homogeneous tissue efficiently but smooth across sharp attenuation transitions, whereas small Gaussians preserve detail but require more primitives. Moreover, a Gaussian primitive's infinite support also forces practical rasterizers to impose numerical tail truncation. Sharp CT transitions can therefore require many small primitives and aggressive population control, increasing memory and optimization cost.

We instead revisit generalized Kaiser–Bessel (KB) functions, historically called blobs (Hsieh, 2022), which provide compact support and analytic X-ray transforms (Lewitt, 1990). Their support size and concentration independently control where a primitive exists and how attenuation varies within it. Based on this separation, we introduce **KB-Splatting**, whose anisotropic primitives have trainable center, rotation, support, density, and concentration. We derive their cone-beam line integrals and gradients and implement a support-aware differentiable GPU rasterizer. The model is optimized directly from measured X-ray projections and voxelized into a CT volume, retaining the practical structure of 3DGS while replacing its spatial basis with one designed for tomographic imaging.

We initialize from FDK and use volumetric residuals to propose three-dimensional seeds. Nearby high-density primitives serve as split parents, while weak primitives are relocated. Inspired by 3DGS-MCMC (Kheradmand et al., 2024), this adapts population-based optimization to additive X-ray formation.

Furthermore, most volumetric pipelines train on digitally reconstructed radiographs (DRRs) synthesized from known volumes. Conversely, real X-ray studies often rely on FDK pseudo-reference evaluations (Zha et al., 2024). DRRs provide exact references but simplify detector response, scatter, truncation, and calibration error (Staub & Murphy, 2013). We therefore introduce a paired multi-system CBCT dataset with measured projections, documented acquisition geometries, and corresponding reference volumes from one commercial Loop-X system (medPhoton GmbH, Salzburg, Austria) and an experimental laboratory system.

In summary, our contributions are:

- We introduce a differentiable anisotropic generalized KB splatting representation with exact compact support and independent control of support size and profile concentration;
- We derive analytic cone-beam projection formulas and implement a support-aware differentiable GPU rasterizer for CBCT reconstruction;

- We develop reconstruction-guided population control that combines FDK initialization, three-dimensional residual seeding, local parent selection, splitting, and relocation;
- We introduce a paired measured CBCT dataset with measured projections, documented acquisition geometry, and corresponding reference volumes from two scanner systems; and
- We extend KB primitives to RGB novel-view synthesis, testing it under a second imaging regime.

2 Related Work

Neural attenuation fields. IntraTomo represents a scanned object with a coordinate network and alternates sinogram prediction with geometry-aware refinement (Zang et al., 2021). NAF accelerates this formulation with multiresolution hash encoding (Zha et al., 2022), while SAX-NeRF models dependencies along X-ray segments and adds projection-context sampling (Cai et al., 2024b). These methods show that parameterizing the attenuation field can provide a useful prior under sparse views. Their forward models, however, repeatedly sample a neural field along each ray, and their capacity remains distributed through network weights rather than localized in inspectable spatial elements. KB-Splatting follows the same per-scan, projection-supervised setting, but represents the volume with a finite population of explicit primitives whose line integrals are evaluated analytically.

Explicit primitives for X-ray imaging. X-Gaussian first adapts radiative Gaussian splatting to X-ray novel-view synthesis and uses the synthesized projections for downstream reconstruction (Cai et al., 2024a). R^2-Gaussian then makes Gaussian density directly queryable, corrects the scale factor required by X-ray integration, and introduces differentiable rasterization and voxelization for direct CT reconstruction (Zha et al., 2024). More recent work further improves this pipeline. X-Field partitions each ray into material-bearing ellipsoidal intervals (Wang et al., 2025). Furthermore, X^2-Gaussian extends radiative Gaussians to continuous-time 4D CT (Yu et al., 2025). Our question is complementary: keeping an explicit differentiable renderer, what radial basis is appropriate for representing the attenuation field itself? KB-Splatting answers with compact primitives whose support and concentration are separate parameters.

Alternative splatting bases and adaptive populations. Recent RGB work has questioned whether either a Gaussian radial profile or heuristic split-and-clone rules are essential to differentiable splatting. 3DGS-MCMC recasts population changes as Markov-chain transitions and interprets opacity regularization as a sparsity prior (Kheradmand et al., 2024). GES replaces the Gaussian footprint with a generalized exponential family and reports improved quality–efficiency trade-offs (Hamdi et al., 2024). These results motivate separating the primitive family from the population mechanism. KB-Splatting makes that separation explicit for tomography.

Real CBCT datasets. NAF and SAX-NeRF popularize controlled benchmarks in which projections are rendered from known CT volumes (Zha et al., 2022; Cai et al., 2024b). These DRRs provide exact reference volumes and repeatable view selection, but their scanner effects are simulated. Real X-ray datasets expose different limitations. The FIPS releases provide corrected projections and acquisition geometry for individual laboratory scans (Meaney, 2022). The 42-walnut collection includes three source trajectories from one laboratory setup and computes high-quality references from the combined acquisition (Der Sarkissian et al., 2019). A recent photon-counting CBCT collection adds raw multi-energy projections, calibration data, and example reconstructions for objects acquired on a custom system (Zhou et al., 2025).

3 KB Splatting for Cone-Beam CT

KB-Splatting represents attenuation with a finite set of anisotropic, compactly supported primitives. This section develops the representation in the order required by the reconstruction problem. We first state the cone-beam observation model, then define the three-dimensional KB field and derive its analytic line integral. We finally describe how we optimize the resulting renderer and adapt its primitive population based on volumetric reconstruction error.

3.1 Problem Formulation

For a view v, the detector coordinate $\mathbf{u}$ maps to an X-ray on a pixel with world-space position $\mathbf{p}_v(\mathbf{u})$, with $\mathbf{o}_v \in \mathbb{R}^3$ denoting the X-ray source. The source and detector positions define the unit

ray direction $\mathbf{d}_v(\mathbf{u}) = (\mathbf{p}_v(\mathbf{u}) - \mathbf{o}_v)/\|\mathbf{p}_v(\mathbf{u}) - \mathbf{o}_v\|_2$. The flat-field-corrected transmission measured at detector coordinate $\mathbf{u}$ is $T_v(\mathbf{u})$, with corresponding log attenuation $y_v(\mathbf{u}) = -\log T_v(\mathbf{u})$. KB-Splatting predicts the same quantity as

$$\widehat{y}_v(\mathbf{u}; \Theta) = \sum_{i=1}^{N} \int_0^{\infty} \mu_i(\mathbf{o}_v + t\mathbf{d}_v(\mathbf{u}))\, \mathrm{d}t, \tag{1}$$

where $\mu_\Theta(\mathbf{x}) = \sum_{i=1}^{N} \mu_i(\mathbf{x})$ is the attenuation field represented by N learnable primitives and Θ denoting their parameters.

3.2 Generalized Kaiser–Bessel Primitives

Compact Radial Profile. Let I_m denote the modified Bessel function of the first kind and let $m \in \{0, 1, 2\}$ be the KB order. For squared normalized radius q, we use the peak-normalized generalized Kaiser–Bessel profile (Lewitt, 1990)

$$b_{m,\beta}(q) = \begin{cases} (1-q)^{m/2} \dfrac{I_m\left(\beta\sqrt{1-q}\right)}{I_m(\beta)}, & 0 \le q < 1, \\ 0, & q \ge 1, \end{cases} \tag{2}$$

where $\beta > 0$ is the concentration. The profile equals one at the center and is exactly zero outside its support (Fig. 2). Its parameters have distinct roles: density sets attenuation magnitude, support sets the region occupied by the primitive, β controls how attenuation is distributed inside that region, and m controls behavior at the boundary. This boundary behavior matters for reconstruction. At $m = 0$, the interior limit at $q = 1$ is nonzero, and the three-dimensional field is discontinuous at its support boundary. For $m > 0$, the profile decays to zero at the boundary, and higher order gives greater smoothness. We therefore treat order as a representation choice rather than a cosmetic kernel setting, and evaluate its effect on boundary error and homogeneous-region artifacts in our experiments as described in Sec. 4.

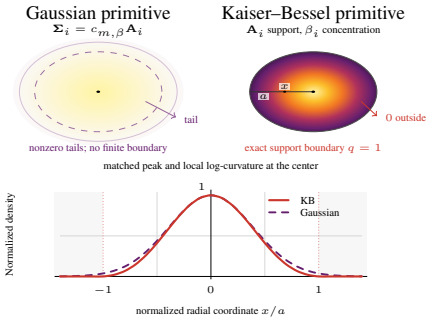

Figure 2: **Gaussian and KB primitives.** Their central peak and curvature are matched. KB has exact support and independently controls support and concentration.

Anisotropic Spatial Parameterization. Each primitive i has center $\mathbf{c}_i$, rotation $\mathbf{R}_i \in \mathrm{SO}(3)$, positive support semiaxes $\mathbf{a}_i = (a_{i1}, a_{i2}, a_{i3})$, density $\rho_i \ge 0$, and concentration β_i. We then define ellipsoidal support $\mathbf{A}_i = \mathbf{R}_i \,\mathrm{diag}(\mathbf{a}_i^2)\mathbf{R}_i^\top$. The squared normalized radius at a local point $\mathbf{x}$ is defined by Mahalanobis distance $q_i(\mathbf{x}) = (\mathbf{x} - \mathbf{c}_i)^\top \mathbf{A}_i^{-1}(\mathbf{x} - \mathbf{c}_i)$. As such, the primitive field is:

$$\mu_i(\mathbf{x}) = \rho_i\, b_{m,\beta_i}(q_i(\mathbf{x}))\,. \tag{3}$$

Its support is the ellipsoid $q_i(\mathbf{x}) < 1$. Unlike 3DGS scales, the semiaxes $\mathbf{a}_i$ are literal spatial boundaries. As such, changing β_i changes the profile within the anisotropic primitive without moving its extent. This supports the decoupling argument in our central representation. A KB primitive's support can cover a broad tissue region while β_i controls how attenuation decays inside it.

Furthermore, learning β introduces a practical ambiguity because support and concentration can both change the profile's effective width. We stabilize this optimization by parameterizing support directly, constraining $\beta_i \in [0.1, 20]$ with a bounded sigmoid, and using a separate learning rate for concentration. Finally, we keep the order m that governs kernel boundary behavior fixed within a reconstruction so that changes in boundary regularity are not confounded with population adaptation. We provide the analytic cone-beam projection and projected KB profile in Appendix B.1.

3.3 Relationship to Gaussian Splatting

Local Gaussian Approximation. KB and Gaussian primitives are locally analogous since both allow a controlled way to separate profile choice from scale. Expanding the logarithm of Eq. 2

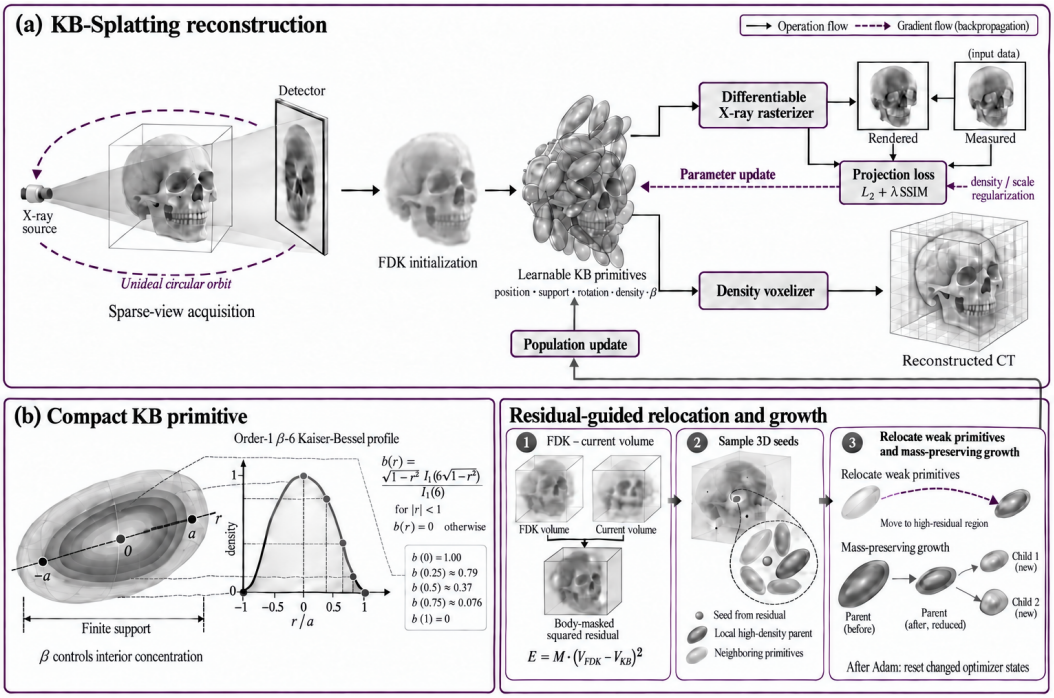

Figure 3: **KB-Splatting reconstruction pipeline.** (a) An FDK volume initializes a learnable population of KB primitives. Differentiable X-ray rasterization optimizes the primitives against measured projections, while voxelized residuals guide relocation and population growth. (b) A Kaiser–Bessel primitive with decoupled spatial support and profile concentration. The support scale controls spatial extent, while β controls how attenuation is distributed within that support.

around $q = 0$ gives

$$\log b_{m,\beta}(q) = -\frac{q}{2c_{m,\beta}} + \mathcal{O}(q^2), \qquad c_{m,\beta} = \left(2m + \beta\frac{I_{m+1}(\beta)}{I_m(\beta)}\right)^{-1}. \tag{4}$$

A Gaussian primitive with covariance $c_{m,\beta}\mathbf{A}_i$ and peak density ρ_i therefore matches the KB primitive's peak and local log-curvature at its center. We use this relation to initialize KB and Gaussian controls at comparable effective widths. After rescaling support to hold this effective covariance fixed, the KB core approaches the matched Gaussian as β increases.

As highlighted in Fig. 3b, the local match does not erase the structural distinction. A Gaussian extends beyond every finite boundary, and its covariance controls both coverage and decay. A KB primitive remains exactly zero at $q = 1$, while $\mathbf{A}_i$ and β_i separately control the support and its interior profile. This yields a direct experimental hypothesis: if that separation is useful for CT, differences should appear most clearly in attenuation transitions (Fig. 1), the number of primitives needed to represent them, and computation that depends on support overlap. Our matched Gaussian and truncated-Gaussian controls test these mechanisms separately.

3.4 Reconstruction-Guided Population Control

In the 3DGS domain, gradient-based adaptive density control is the standard mechanism for allocating scene primitives, but it depends on hand-designed thresholds for when and where Gaussians are cloned, split, or pruned. 3DGS-MCMC (Kheradmand et al., 2024) instead interprets the Gaussian population as samples from an underlying probability distribution over scene representations. Adding noise to conventional gradient updates yields stochastic gradient Langevin dynamics (SGLD), while approximately probability-preserving state transitions relocate low-opacity primitives under a prescribed population budget.

This sampling view successfully replaces most discrete densification and pruning rules with a principled population-update process leading to rendering speed gains in the RGB domain. Yet in our experiments, a naive implementation of MCMC SGLD in the CT domain produces severe radio-density artifacts and ambiguity in 3D attenuation placement.

Volumetric Residual Seeding. To address CT density placement ambiguity, we devise a simple yet effective *Coarse Radiative Densification* (CRD) scheme. At each population update during training, we voxelize the current primitive field on a coarse grid and compare it with the FDK initializer (Fig. 3a). Let M_{body} be the thresholded FDK body mask. We define

$$E(\mathbf{x}) = M_{\mathrm{body}}(\mathbf{x}) \left(\mu_{\mathrm{FDK}}(\mathbf{x}) - \mu_{\Theta}(\mathbf{x})\right)^2 . \tag{5}$$

Therefore, we sample three-dimensional seed locations in proportion to E. It is worth noting that CRD does not force the optimized volume to reproduce FDK because E is used only to propose locations, and the parameters at those locations are subsequently fit to measured projection gradients.

A seed still needs a plausible parent from which to inherit density, support, rotation, and concentration. To do this, we first retain a high-density subset of the current population, find its K nearest primitives to the seed, and sample a parent with probability proportional to parent candidate density. Selecting among nearby candidates makes the inherited parameters more likely to reflect the seed's local anatomy. Weighting those candidates by density favors primitives that already make a meaningful contribution to the attenuation field.

3.5 Differentiable Optimization

Parameterization and Gradients. We optimize each KB center directly and, following R^2-Gaussian (Zha et al., 2024), represent rotation with a normalized quaternion before rendering. We obtain primitive support and density from unconstrained variables through positive activations. Furthermore, the profile concentration uses the bounded parameterization described in Section 3.2. At fixed evaluation intervals, we voxelize the primitive field and compare it with the corresponding reference CT volume for qualitative reporting only. As stated in Fig. 3a, the FDK initializer used for residual seeding is computed from the training-view projections. The reported CT benchmark runs use center-ray rectification: a detector-space footprint and a path-length factor evaluated along the primitive-center ray. The backward pass differentiates this approximation with respect to the primitive parameters, including concentration β. We retain the exact per-ray formulation in Eq. 11 for validation and renderer ablations.

For a sampled training view, we minimize mean squared error in log-attenuation space, augmented by an optional structural term and ℓ_1 penalties inspired by 3DGS-MCMC (Kheradmand et al., 2024):

$$\mathcal{L}(\Theta) = \tfrac{1}{|\Omega_v|} \sum_{\mathbf{u}\in\Omega_v} \left(\widehat{y}_v(\mathbf{u};\Theta) - y_v(\mathbf{u})\right)^2 + \lambda_{\mathrm{s}} \left(1 - \mathrm{SSIM}(\widehat{y}_v, y_v)\right) + \lambda_\rho \sum_i |\rho_i| + \lambda_a \sum_{ij} |\widetilde{a}_{ij}|. \tag{6}$$

Here $\widetilde{a}_{ij} = \sqrt{c_{m,\beta_0}}\, a_{ij}$ is the curvature-matched semiaxis at the reference concentration β_0. Thus, the final two terms adapt the MCMC opacity and scale penalties to CT attenuation density and KB width. Ω_v represents a detector-space ray-validity mask. For a view v, this mask retains detector pixels whose source-to-pixel rays intersect the axis-aligned bounding box of the CT reconstruction volume.

4 Experiments

4.1 Experimental settings

Datasets. We evaluate reconstruction from measured projections on 15 scans acquired with two cone-beam CT systems. The commercial Loop-X collection contains 12 scans spanning anonymized human cadaver, phantom, and inanimate object acquisitions: calibrator, forearm, left hand, neck, phantom chest, phantom head, phantom hip, phantom neck, raven, right hand, and skull. The laboratory CT device collection contains a phantom Skull and attenuation-mimicking discs. For every scan, we use 100%, 50%, and 25% of the available projections, preserving identical view subsets across methods. For evaluation, we use corresponding reconstructed reference volumes. The volumes are reconstruction references from proprietary backward projection algorithms (for Loop-X) and an analytical Filtered Back Projection (FBP) method (for the research system).

Table 1: Quantitative evaluation on sparse-view tomography real X-ray CBCT reconstruction against baselines (Zang et al., 2021; Zha et al., 2022; Cai et al., 2024b; Wang et al., 2025; Zha et al., 2024). We colorize the best, second-best, and third-best numbers.

Method	100% views			50% views			25% views		
	PSNR↑	SSIM↑	Time↓	PSNR↑	SSIM↑	Time↓	PSNR↑	SSIM↑	Time↓
Neural Methods									
NAF	15.52	0.58042	34m45s	15.23	0.57168	34m12s	15.05	0.54938	35m00s
IntraTomo	14.13	0.52685	38m30s	14.06	0.52023	38m04s	14.31	0.52394	36m12s
SAX-NeRF	15.34	0.56821	16h19m	14.22	0.55529	14h33m	13.70	0.47186	13h11m
Splatting Methods									
X-Field	42.91	0.98922	**2m52s**	42.87	0.98903	**2m53s**	42.88	0.98915	**2m52s**
R^2-Gaussian	45.66	0.98982	17m46s	45.38	0.98915	16m32s	45.33	0.98912	16m21s
KB-Splatting (ours)	**46.54**	**0.99065**	9m47s	**46.33**	**0.99041**	9m56s	**46.30**	**0.99038**	9m54s

Implementation details. KB-Splatting uses PyTorch with custom CUDA rasterization and voxelization. We optimize each scan with Adam (Kingma & Ba, 2015) for 10,000 iterations, starting from 50,000 primitives sampled from an FDK volume reconstructed using only that run's training views. The final CT recipe uses KB order $m = 1$, center-ray rectification, and per-primitive concentration initialized at $\beta = 6$, bounded to $[0.1, 20]$, with learning rate 10^{-3}. Due to compute restrictions, projections are downsampled by a factor of two. The objective combines projection MSE with $0.10(1 - \text{SSIM})$ and ℓ_1 density (λ_ρ) and scale (λ_a) penalties weighted by 2.5×10^{-5} and 0.5, respectively. We also perform CRD on a 32^3 grid with an FDK body threshold of 0.001. Seed locations are drawn from the squared volumetric residual, and parents are sampled by density from the 55 nearest candidates within the top 20% of primitives by density. As in Kheradmand et al. (2024), population updates occur every 100 iterations from iteration 500 through 10,000, subject to a scan-specific population cap. Updates follow the Adam step and reset optimizer state for structurally changed primitives. We mostly train on the NVIDIA L40S.

Baselines and metrics. We compare with R^2-Gaussian (Zha et al., 2024), X-Field (Wang et al., 2025), NAF (Zha et al., 2022), IntraTomo (Zang et al., 2021), and SAX-NeRF (Cai et al., 2024b). FDK supplies initialization rather than an additional trained baseline in this comparison. We report the original attenuation PSNR and SSIM from the final run records, along with training wall-clock time, in Table 1. PSNR is expressed in dB, and SSIM summarizes slice-wise agreement with the reference CT volume. Reported training time excludes final volume reconstruction and evaluation.

4.2 Results and Evaluation

Measured CT reconstruction. Table 1 summarizes our quantitative results. KB-Splatting attains the highest reported PSNR and SSIM in each setting. Its PSNR is 46.54, 46.33, and 46.30 dB at 100%, 50%, and 25% views, which translates to ∼600, ∼300, and ∼150 X-ray projections, respectively, compared with 45.66, 45.38, and 45.33 dB for the current SOTA method, R^2-Gaussian. These differences correspond to gains of 0.88, 0.95, and 0.97 dB, or 0.93 dB on average across the three settings. The corresponding SSIM scores are 0.99065, 0.99041, and 0.99038, with improvements of 0.00083, 0.00126, and 0.00126 over R^2-Gaussian.

Reducing the projection fraction from 100% to 25% decreases KB-Splatting's reported PSNR by 0.24 dB and SSIM by 0.00027. Nevertheless, its 25%-view result remains above R^2-Gaussian's 100%-view result in these original metrics: 46.30 versus 45.66 dB and 0.99038 versus 0.98982 SSIM. This comparison describes the reported reconstruction setting and does not establish an equivalent reduction in clinical dose or diagnostic performance. These results are consistent with KB primitives retaining useful attenuation structure as the number of training projections decreases. This notion is further confirmed in the qualitative results reported in Fig. 4 where KB splatting displays high rendering fidelity across several anatomies.

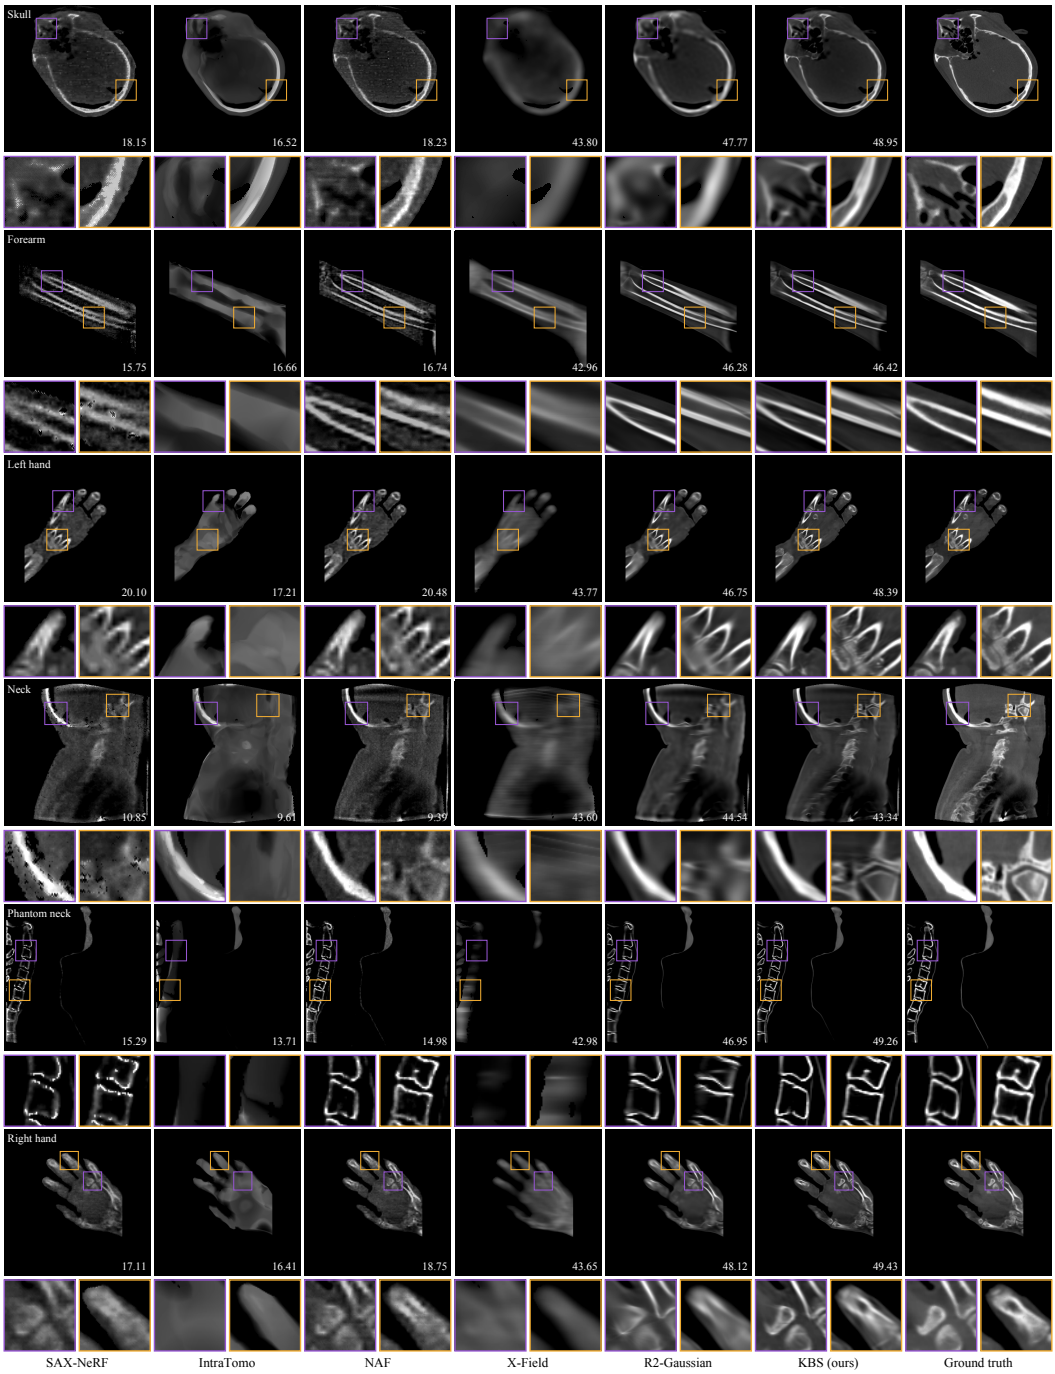

Figure 4: **Loop-X qualitative comparison at 100% views (600 projections).** Matched slices and enlarged regions for six representative anatomies. We report attenuation PSNR in dB.

Table 2: **Matched primitive ablation.** A single controlled configuration isolates the representation.

Primitive	HU PSNR ↑	Atten. PSNR ↑	Time (min) ↓
Matched Gaussian	20.45	51.413	23.72
KB, fixed $\beta = 6$	20.50	51.455	**21.74**
KB, learned β_i	**20.57**	**51.531**	23.47

4.3 Ablation study

For our ablation study, we ask: Are the rendering improvements caused by using Kaiser–Bessel primitives, or did KB-Splatting receive an unfair advantage from using more primitives, a larger computational budget, different initialization, or different optimization settings? We therefore replace the order-one KB primitive with a Gaussian matched at its peak and local log-curvature, while holding initialization, seed, objective, iteration and population budgets, and corrected structural bookkeeping fixed. We then restore a fixed-profile KB ($\beta = 6$) and finally learn β per primitive. Table 2 shows that compact support alone improves attenuation and HU PSNR while reducing observed time: fixed-profile KB gains 0.042 and 0.05 dB, respectively, and is 8.35% faster than the matched Gaussian. Learning concentration raises the gains to 0.117 and 0.12 dB at comparable runtime, supporting the value of separating support from interior profile shape. This is mechanistic evidence from one controlled setting, rather than a second leaderboard result. Additional appendix controls isolate rendering, loss, regularization, and population allocation. In particular, center-ray rendering stays within 0.05 dB of exact ray integration while halving training time. Enforcing spatially dispersed seed locations decreases attenuation PSNR by 0.295 dB with one seed per residual cell and by 0.836 dB with non-maximum suppression, compared with unrestricted squared-residual sampling.

5 Discussion and Conclusion

Generalizes Across Image Formation Regimes. In this work, we investigate the radiative rendering capacity of the Kaiser-Bessel function in the splatting domain. We also explore its volumetric reconstruction capabilities for RGB scene rendering. According to (Hamdi et al., 2024), non-Gaussian splat families have succeeded as rendering surrogates for Gaussian kernels. We report a complete exposition of this viewpoint, along with RGB transfer experiments, including quantitative benchmark tables and qualitative comparisons, in Appendix I.

Data and evaluation. A limitation of our work is the paucity of training and evaluation data. The study contains 15 scans from two systems, which do not span the diversity of clinical scanners, protocols, anatomies, or pathologies. Furthermore, our reconstruction volumes are evaluated against reference volumes that are outputs of analytical or proprietary reconstruction algorithms. As such, quantitative agreement with these reconstructed references does not establish dose equivalence, diagnostic accuracy, or safety. Therefore, released material stemming from this work is intended specifically for reconstruction research.

Conclusion. KB-Splatting shows that differentiable CT splatting need not be restricted to Gaussian primitives. An anisotropic generalized Kaiser–Bessel basis provides exact finite support and decouples support size from interior concentration, while a coarse reconstruction-guided population control assigns capacity using the current volumetric error. On 15 measured scans from two cone-beam CT systems, the KB-Splatting pipeline improves the retained attenuation PSNR over the strongest reported baseline by 0.93 dB on average across three projection fractions. Our KB-Splatting further surpasses the baseline in training speed and fine anatomic detail rendering.

References

Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5470–5479, 2022.

Yuanhao Cai, Yixun Liang, Jiahao Wang, Angtian Wang, Yulun Zhang, Xiaokang Yang, Zongwei Zhou, and Alan Yuille. Radiative gaussian splatting for efficient X-ray novel view synthesis. In *European Conference on Computer Vision*, 2024a.

Yuanhao Cai, Jiahao Wang, Alan Yuille, Zongwei Zhou, and Angtian Wang. Structure-aware sparse-view X-ray 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024b.

Bruno M. Carvalho and Gabor T. Herman. Low-dose, large-angled cone-beam helical CT data reconstruction using algebraic reconstruction techniques. *Image and Vision Computing*, 25(1): 78–94, 2007. doi: 10.1016/j.imavis.2006.03.003.

Henri Der Sarkissian, Felix Lucka, Margo van Eijnatten, Giulia Colacicco, Sophia B. Coban, and Kees Joost Batenburg. A cone-beam X-ray computed tomography data collection designed for machine learning. *Scientific Data*, 6(1):215, 2019. doi: 10.1038/s41597-019-0235-y.

L. A. Feldkamp, L. C. Davis, and J. W. Kress. Practical cone-beam algorithm. *Journal of the Optical Society of America A*, 1(6):612–619, 1984. doi: 10.1364/JOSAA.1.000612.

Abdullah Hamdi, Luke Melas-Kyriazi, Jinjie Mai, Guocheng Qian, Ruoshi Liu, Carl Vondrick, Bernard Ghanem, and Andrea Vedaldi. Ges: Generalized exponential splatting for efficient radiance field rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19812–19822, 2024. doi: 10.1109/CVPR52733.2024.01873.

Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics*, 37(6), 2018.

Jiang Hsieh. *Computed Tomography: Principles, Design, Artifacts, and Recent Advances*. SPIE Press, Bellingham, WA, 4 edition, 2022. ISBN 9781510646872. doi: 10.1117/3.2605933.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023.

Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Weiwei Sun, Jeff Tseng, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. 3d gaussian splatting as markov chain monte carlo. In *Advances in Neural Information Processing Systems*, 2024.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics*, 36(4), 2017.

Robert M. Lewitt. Multidimensional digital image representations using generalized kaiser–bessel window functions. *Journal of the Optical Society of America A*, 7(10):1834–1846, 1990. doi: 10.1364/JOSAA.7.001834.

Yusheng Li. Optimization for blob-based image reconstruction with generalized kaiser–bessel basis functions. *IEEE Transactions on Computational Imaging*, 4(2):257–270, 2018. doi: 10.1109/TCI.2018.2796302.

Chongshan Lu, Fukun Yin, Xin Chen, Wen Liu, Tao Chen, Gang Yu, and Jiayuan Fan. A large-scale outdoor multi-modal dataset and benchmark for novel view synthesis and implicit scene reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023.

Samuel Matej and Robert M. Lewitt. Practical considerations for 3-d image reconstruction using spherically symmetric volume elements. *IEEE Transactions on Medical Imaging*, 15(1):68–78, 1996. doi: 10.1109/42.481442.

Alexander Meaney. Cone-beam computed tomography dataset of a walnut, 2022.

David Staub and Martin J. Murphy. A digitally reconstructed radiograph algorithm calculated from first principles. *Medical Physics*, 40(1):011902, 2013. doi: 10.1118/1.4769413.

Feiran Wang, Jiachen Tao, Junyi Wu, Haoxuan Wang, Bin Duan, Kai Wang, Zongxin Yang, and Yan Yan. X-field: A physically informed representation for 3d X-ray reconstruction. In *Advances in Neural Information Processing Systems*, 2025.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004.

Weihao Yu, Yuanhao Cai, Ruyi Zha, Zhiwen Fan, Chenxin Li, and Yixuan Yuan. X^2-gaussian: 4d radiative gaussian splatting for continuous-time tomographic reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025.

Guangming Zang, Ramzi Idoughi, Rui Li, Peter Wonka, and Wolfgang Heidrich. IntraTomo: Self-supervised learning-based tomography via sinogram synthesis and prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1960–1970, 2021.

Ruyi Zha, Yanhao Zhang, and Hongdong Li. NAF: Neural attenuation fields for sparse-view CBCT reconstruction. In *Medical Image Computing and Computer Assisted Intervention*, 2022.

Ruyi Zha, Tao Jun Lin, Yuanhao Cai, Jiwen Cao, Yanhao Zhang, and Hongdong Li. R^2-gaussian: Rectifying radiative gaussian splatting for tomographic reconstruction. In *Advances in Neural Information Processing Systems*, 2024.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.

Enze Zhou, Wenjian Li, Wenting Xu, Kefei Wan, Yuwei Lu, Shangbin Chen, Gang Zheng, Tianwu Xie, and Qian Liu. A cone-beam photon-counting CT dataset for spectral image reconstruction and deep learning. *Scientific Data*, 12(1):1955, 2025. doi: 10.1038/s41597-025-06246-4.

A Declaration

AI Use Statement

We used generative AI tools for code assistance, experiment bookkeeping, figure assembly, literature organization, and language editing. The authors inspected the resulting code, equations, citations, numerical summaries, and manuscript text, and remain responsible for all claims and errors.

Ethics Statement

The experiments use phantom, inanimate, and cadaveric acquisitions. No released artifact contains direct patient identifiers, and the work does not make clinical decisions. Data redistribution follows the permissions and licenses of the source collections. The intended use is reconstruction research; clinical deployment would require independent validation and appropriate institutional review.

Reproducibility Statement

The supplement records the final renderer, optimization schedule, view subsets, random seed, population rules, evaluation convention, and the provenance of every reported run. We will release the code, configuration files, geometry-conversion and validation scripts, redistributable acquisition metadata, and machine-readable metric manifests. Reference volumes are used only for evaluation, never as optimization targets.

B Kaiser-Bessel Derivations

Kaiser–Bessel and blob basis functions. Generalized Kaiser–Bessel functions are a classical tomographic basis, not a new function introduced by this work. Lewitt derives their Fourier and X-ray transforms and proposes them as finite-support alternatives to voxel basis functions (Lewitt, 1990). Matej and Lewitt show that smoothly tapered spherical volume elements can improve the noise–resolution tradeoff of iterative reconstruction (Matej & Lewitt, 1996). Subsequent work applies KB blobs to low-dose cone-beam reconstruction (Carvalho & Herman, 2007) and studies how global blob parameters and sampling lattices affect approximation error (Li, 2018).

These formulations typically place isotropic translated copies on a prescribed lattice, select the support and shape parameters globally, and optimize only the expansion coefficients. KB-Splatting changes the role of the basis. Each blob becomes a free anisotropic primitive with a learned center, rotation, support, density, and optional concentration. Our contribution is therefore not the KB function; it is the differentiable splatting system built around it: anisotropic parameterization, cone-beam gradients, support-aware GPU rasterization, adaptive population control, and a second renderer for occlusion-aware RGB formation.

B.1 Analytic Cone-Beam Projection

B.1.1 Ray–Ellipsoid Intersection

Consider a unit ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ and primitive i. Let $\mathbf{K}_i = \mathbf{A}_i^{-1}$ and $\boldsymbol{\delta}_i = \mathbf{o} - \mathbf{c}_i$. The quadratic radius along the ray is

$$q_i(t) = (\boldsymbol{\delta}_i + t\mathbf{d})^\top \mathbf{K}_i (\boldsymbol{\delta}_i + t\mathbf{d}). \tag{7}$$

Writing

$$\gamma_i = \mathbf{d}^\top \mathbf{K}_i \mathbf{d}, \qquad t_i^* = -\frac{\mathbf{d}^\top \mathbf{K}_i \boldsymbol{\delta}_i}{\gamma_i}, \tag{8}$$

gives the squared transverse distance at the closest point,

$$q_i^\perp = (\boldsymbol{\delta}_i + t_i^* \mathbf{d})^\top \mathbf{K}_i (\boldsymbol{\delta}_i + t_i^* \mathbf{d}). \tag{9}$$

The ray misses the primitive exactly when $q_i^\perp \geq 1$ or when the intersection lies behind the source. This finite test replaces the numerical tail cutoff required by a Gaussian. It also gives the renderer a direct criterion for rejecting detector pixels and projected tiles that cannot receive attenuation from the primitive.

B.1.2 PROJECTED KAISER–BESSEL PROFILE

For $q < 1$, define $z = \sqrt{1-q}$ and the normalized projected profile

$$\bar{b}_{m,\beta}(q) = z^{m+1/2} \frac{I_{m+1/2}(\beta z)}{I_{m+1/2}(\beta)}, \qquad \bar{b}_{m,\beta}(q) = 0 \text{ for } q \geq 1. \tag{10}$$

The line integral of Eq. equation 3 then has the closed form

$$\int_0^\infty \mu_i(\mathbf{o} + t\mathbf{d})\,\mathrm{d}t = \frac{\rho_i}{\sqrt{\gamma_i}} \underbrace{\sqrt{\frac{2\pi}{\beta_i}} \frac{I_{m+1/2}(\beta_i)}{I_m(\beta_i)}}_{\kappa_{m,\beta_i}} \bar{b}_{m,\beta_i}(q_i^\perp), \tag{11}$$

provided the supported segment lies in front of the source. The factor $\gamma_i^{-1/2}$ accounts for the physical path length through an oriented ellipsoid, while κ_{m,β_i} depends only on order and concentration. Equation equation 11 is differentiable with respect to density, center, rotation, support, and concentration. For the supported orders, the projected half-integer Bessel terms are evaluated with stable closed forms, series, and short recurrences.

Our exact renderer computes $\mathbf{d}$, $q_i^\perp$, and γ_i for each primitive–pixel pair that survives tile binning. We also implement a faster center-ray approximation. It locally projects $\mathbf{A}_i$ to a detector-space ellipse at $\mathbf{c}_i$, evaluates the quadratic distance in that ellipse, and shares the center-ray path-length factor across the footprint. The approximation approaches the exact geometry for small primitives and weak perspective, but can deviate for large or highly elongated primitives near the source. All reported final CT benchmark runs use the center-ray mode; the exact per-ray renderer is used for validation and renderer ablations. The analytic KB profile is retained in both modes, but the center-ray detector geometry is an approximation.

- Generalized order-m profile, normalization, boundary behavior, and volume mass.
- Stable formulas for supported orders and beta derivatives.
- Local curvature match, covariance match, and large-beta Gaussian limit.

C EFFICIENT GPU IMPLEMENTATION

The rasterizer first projects a conservative bound for each support ellipsoid and bins the primitive only into intersecting detector tiles. Within a tile, the exact renderer computes the ray–ellipsoid distance in Eq. equation 9 and rejects the primitive when $q_i^\perp \geq 1$ before evaluating the KB profile. Constants depending only on (m, β_i) are precomputed, and the forward and backward passes share the same support and ray geometry. We implement orders $m \in \{0, 1, 2\}$ with numerically stable Bessel evaluations over the bounded concentration range.

Compact support makes exact culling possible, but it does not by itself guarantee a faster end-to-end system. Runtime also depends on projected footprint size, perspective bounds, special-function evaluation, primitive count, and backward-pass cost. We therefore report both overlap statistics and measured reconstruction time rather than treating compactness as a proxy for speed. We validate the implementation against numerical quadrature, finite-difference gradients, analytic single-ray tests, support-boundary cases, and exact-culling checks; the following section specifies the full tolerances and test configurations.

D RENDERER AND GRADIENT VALIDATION

We validate the analytic ray–ellipsoid integral against dense numerical quadrature over randomized primitive centers, supports, rotations, orders, concentrations, and cone-beam rays. We separately test rays that miss the support, graze the boundary, cross the center, and intersect at shallow angles. The implementation returns exactly zero outside $q^\perp < 1$ and approaches zero continuously at the boundary. Central finite differences for center, log-support, rotation, density, and β agree with automatic differentiation away from the non-differentiable support boundary; one-sided checks are used at the boundary.

The final experiments use center-ray rectification because it preserves the matched development score while substantially reducing cost. For order $m = 0$, exact and center-ray rendering attain 51.20 and 51.15 dB at 32.48 and 15.31 minutes. For $m = 2$, they attain 51.20 and 51.16 dB at 47.98 and 24.36 minutes. These controls indicate errors below 0.05 dB in the tested configuration with a 1.97–2.12$\times$ observed time reduction. They do not imply an analytic equivalence; the exact integral remains the reference for stress tests and for regimes with large footprints or stronger cone divergence.

E Population Control Details

Population updates run every 100 iterations from 500 through 10,000. At an update, we voxelize the current field on a 32^3 grid, apply the FDK body mask at threshold 0.001, and sample seed cells proportionally to the squared residual. For each seed, the parent is sampled by density from the 55 nearest candidates within the top 20% of primitives by density. Empty neighborhoods fall back to the nearest eligible primitive; an update is skipped if no valid body cell or parent remains. Relocation first reuses low-density slots, after which geometric growth adds descendants until the scan-specific cap is reached.

The selected schedule follows targeted controls. Relative to squared-residual sampling, enforcing unique occupied cells lowers the recorded PSNR by 0.295 dB, non-maximum suppression by 0.836 dB, and fourth-power residual weighting by 0.265 dB. We therefore favor concentrated proposals over forced diversity. Population changes are applied after Adam and all moments for relocated destinations and modified parents are reset. Density-gated position noise is active from iterations 1,000–7,000; because its gains were not uniform across development cases, we treat it as an exploration device rather than a primary source of the final accuracy.

E.1 Splitting, Relocation, and Stochastic Motion

We use the selected parents in two structural updates inspired by the population view of 3DGS-MCMC (Kheradmand et al., 2024). First, low-density primitives are treated as dead slots and relocated to residual seeds. Second, new slots are added geometrically until a fixed population cap is reached. If one parent supplies a cluster of n descendants, density and support are adjusted to preserve its integrated attenuation. Since the mass of a KB primitive with fixed order and concentration is proportional to $\rho_i \prod_j a_{ij}$, the update satisfies

$$n\rho_i' \prod_{j=1}^{3} a_{ij}' = \rho_i \prod_{j=1}^{3} a_{ij}. \tag{12}$$

Per-primitive concentration and rotation are inherited. An optional support-shaped Langevin update perturbs low-density centers, allowing weak primitives to explore nearby positions before they are declared dead.

Structural updates occur after the Adam step. Relocated destinations and modified parents no longer represent the variables for which their stored gradients and moments were accumulated, so we reset their optimizer state. Without this reset, stale momentum can immediately move a newly placed primitive away from its residual seed or undo a mass-preserving split.

F Real Multi-Scanner CBCT Dataset

F.1 Scope and records

The measured-data benchmark contains 15 scans from two cone-beam CT systems. The 12 Loop-X records cover abdomen, calibrator, forearm, left-hand short scan, neck, phantom chest, phantom head, phantom hip, phantom neck, raven, right-hand full scan, and skull. The three Research CT System records are Skull, phantom disc 1, and phantom disc 2. Each record pairs projection images, the geometry used by the reconstruction code, deterministic 100%/50%/25% view subsets, and a reference volume on the evaluation grid.

Table 3: Measured CT collections. View counts describe the full and deterministic sparse subsets. The reference volumes are reconstructed references rather than direct attenuation ground truth.

Collection	System type	Scans	Views	Geometry provenance
Loop-X	commercial mobile CBCT	12	acquisition dependent	supplied per view
Research CT System	laboratory cone-beam CT	3	800/400/200	repaired CTorch conversion

F.2 Geometry, preprocessing, and quality control

Loop-X retains the supplied per-view source and detector poses. The Research CT System adapter translates the circular CTorch geometry into the camera convention used by KB-Splatting and the baselines. We explicitly correct detector pixel-center offsets and verify the conversion with off-axis point projections, because an internally consistent orbit can still disagree with CTorch away from the central ray. The same repaired cameras are used for all methods that consume per-view extrinsics.

Projection preprocessing follows each acquisition's available dark- and flat-field records. We retain a provenance warning when object and air scans use different source or detector stage positions, because clipping transmissions above one can introduce shading and attenuation bias. Geometry is checked by reprojecting the supplied reference, inspecting structured edge residuals, and comparing landmark projections. These checks verify the implemented convention; they do not replace an independent physical calibration.

F.3 References, splits, and release boundary

Loop-X references are reconstructed from the associated DICOM volume. Research CT System references are the supplied CTorch FBP volumes reconstructed from the same acquisitions. They are used only for evaluation and for the coarse residual proposal field described in Section 3; they never enter the projection loss. Sparse subsets are deterministic and nested where acquisition ordering permits, and every method receives the same subset at a given fraction.

Released documentation will include the projection-to-volume coordinate convention, detector indexing, axis orientation, units, crop and resampling transforms, view lists, and the checksums of evaluation references. Data whose license prevents redistribution will be accompanied by acquisition identifiers and conversion scripts. Because the benchmark contains cadaveric, phantom, and inanimate material and no direct identifiers, its intended scope is method development rather than clinical decision support.

G Complete CT Experimental Protocol

All final CT runs use 10,000 Adam iterations, an order-one KB profile, center-ray rectification, and learnable per-primitive β initialized at 6 and clamped to $[0.1, 20]$. Position, density, log-support, rotation, and concentration use the recorded final learning-rate schedules. The projection loss combines MSE and $0.10(1 - \mathrm{SSIM})$; density and scale penalties are 2.5×10^{-5} and 0.5, and total variation is disabled. Initial populations contain 50,000 FDK-sampled primitives. Coarse radiative densification follows Appendix E, with seed 42 and scan-specific caps. Training projections are downsampled by two. The final volume is evaluated on the reference grid with the original attenuation PSNR and SSIM emitted by each completed campaign.

H CT Ablations and Additional Results

H.1 Matched representation control

The matched control changes only the primitive family and concentration parameterization under the same corrected bookkeeping. Table 4 and Figure 5 show that fixed KB improves over the Gaussian control, and learned concentration improves further. The fixed profile also reduces observed time; learned concentration recovers most of that speed while providing the best score.

Table 4: Matched primitive control. Attenuation PSNR and HU PSNR use the recorded development evaluation; time is wall-clock training time.

Representation	Atten. PSNR ↑	HU PSNR ↑	Time (min) ↓
Gaussian	51.413	20.45	23.72
KB, fixed $\beta = 6$	51.455	20.50	**21.74**
KB, learned β	**51.531**	**20.57**	23.47

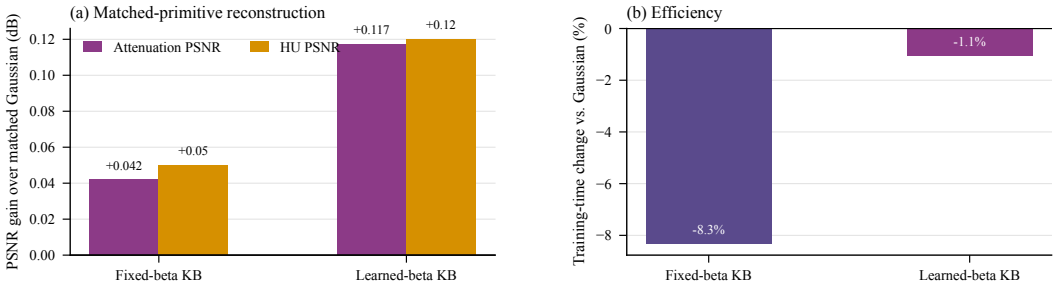

Figure 5: **Controlled CT representation ablation.** Quality and time for Gaussian, fixed-concentration KB, and learned-concentration KB primitives under the matched training recipe.

H.2 Renderer, loss, and allocation controls

The exact and center-ray renderers differ by at most 0.05 dB in the recorded order-zero and order-two controls, while the center approximation is about twice as fast. A DSSIM sweep gives attenuation PSNR 51.05, 51.56, 51.39, and 51.18 dB at weights 0.05, 0.10, 0.15, and 0.25, selecting 0.10. The selected scale and density regularizers, 0.5 and 2.5×10^{-5}, improve the nine-setting aggregate by 0.015 dB and win seven settings. These small changes are reported to delimit the sensitivity of the final recipe rather than to imply universal optima.

Fixed-concentration screening at order one produces HU PSNR 19.90, 20.05, 20.13, and 20.23 dB for $\beta \in \{2, 3, 4, 6\}$, motivating the $\beta = 6$ initialization. The residual proposal study selects squared-residual sampling: unique-cell, non-maximum-suppressed, and fourth-power proposals lose 0.295, 0.836, and 0.265 dB, respectively. Structural changes are therefore applied after Adam with reset moments, and density-gated motion is retained as a bounded exploration mechanism.

H.3 Research CT System quantitative and qualitative results

Figure 6 summarizes the nine repaired-geometry comparisons between KB-Splatting and R^2-Gaussian. KB-Splatting improves attenuation PSNR in every dataset–fraction pair by 1.061 dB on average, while the geometric mean of the observed R^2-Gaussian-to-KB time ratio is 1.74×. Hardware and implementation differences remain a caveat.

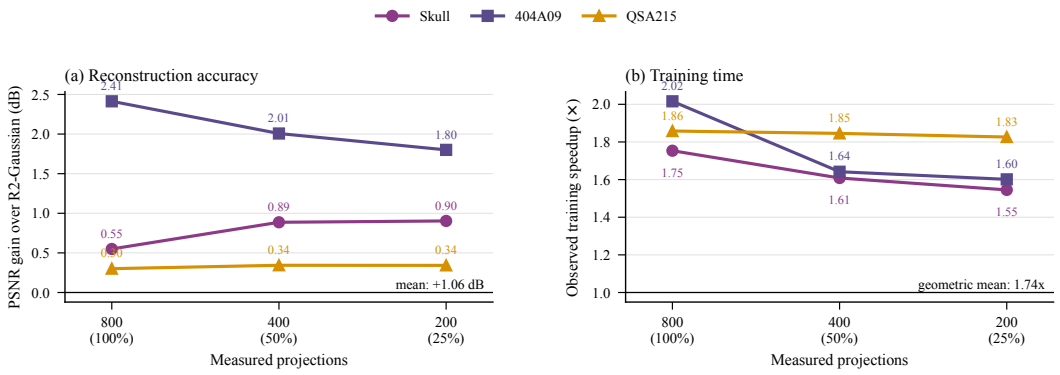

Figure 6: **Research CT System quality and training-time summary.** Results use repaired CTorch camera conventions for all nine dataset–projection-fraction settings.

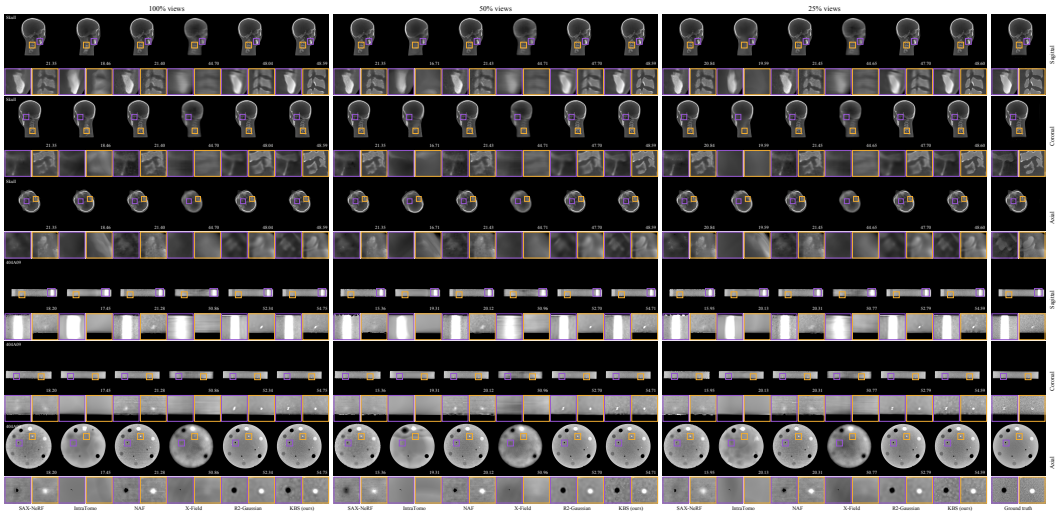

Figure 7: **Research CT System qualitative comparison across projection fractions.** Matched slices are grouped by 100%, 50%, and 25% views; all methods use the same camera correction and display window.

H.4 Additional Loop-X sparse-view comparisons

Figures 9–12 complete the qualitative comparison at 50% and 25% views. Crops and slice locations are fixed across methods and agree with the 100% protocol in Figures 4 and 8.



Figure 8: **Additional Loop-X qualitative results at 100% views.** Matched slices and enlarged regions for the remaining five anatomies; values report original attenuation PSNR in dB.

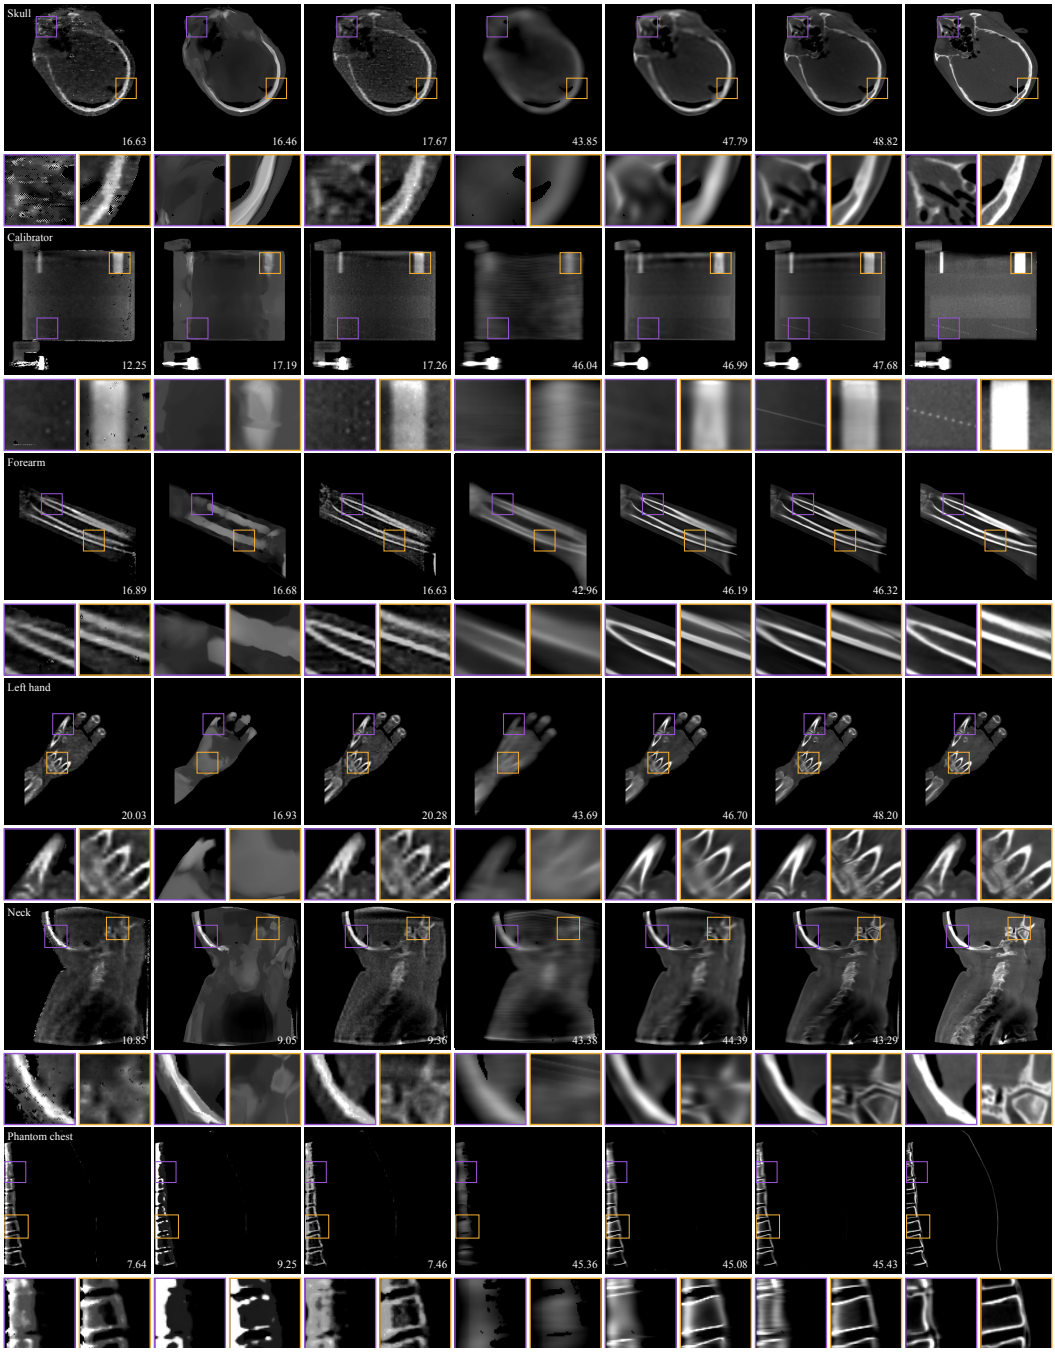

Figure 9: **Loop-X at 50% views, part I.** Matched slices and enlarged regions for the first anatomy group.

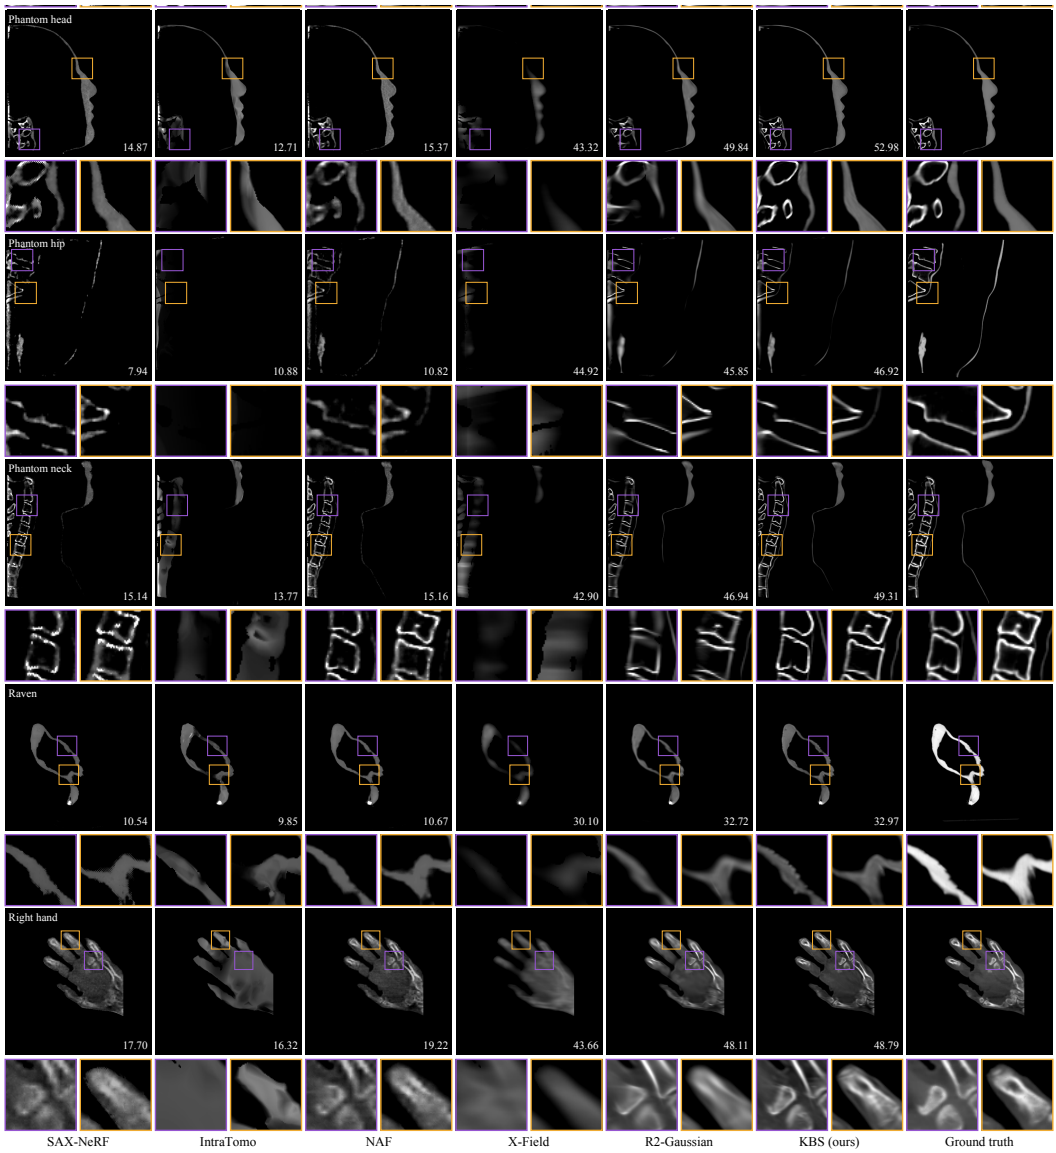

Figure 10: **Loop-X at 50% views, part II.** Remaining anatomies under the same display protocol.

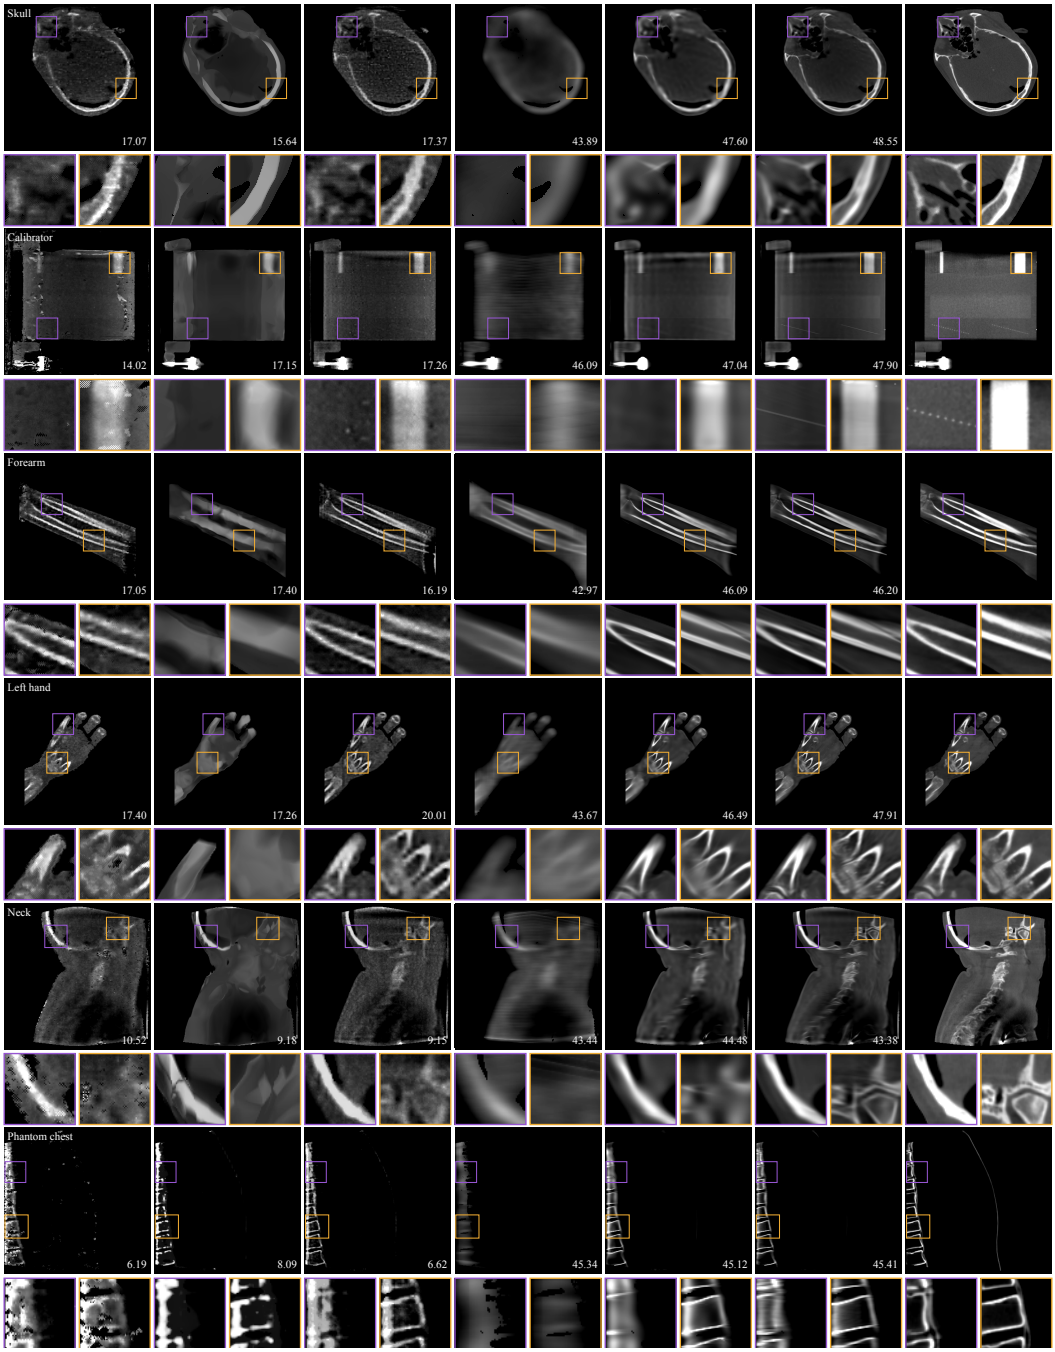

Figure 11: **Loop-X at 25% views, part I.** Matched slices and enlarged regions for the first anatomy group.

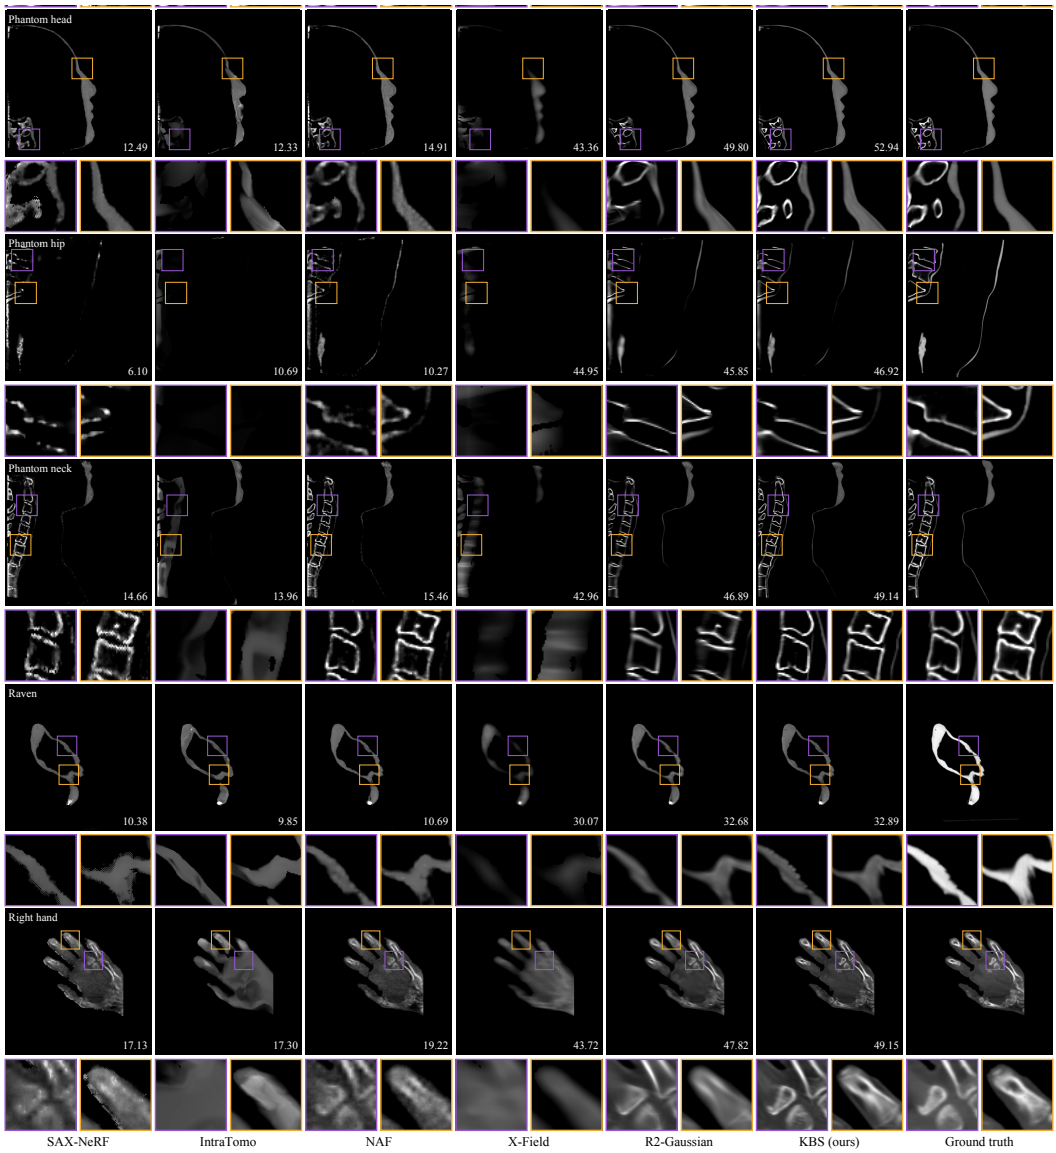

Figure 12: **Loop-X at 25% views, part II.** Remaining anatomies under the same display protocol.

Table 5: **RGB novel-view synthesis results.** Each entry is the mean over all scenes in the corresponding benchmark. Training time is measured in minutes on an NVIDIA L40S GPU. Bold indicates the best value in each benchmark, with higher PSNR/SSIM and lower LPIPS/time preferred.

	Tanks & Temples				Deep Blending				Mip-NeRF 360				OMMO			
Method	PSNR	SSIM	LPIPS	min	PSNR	SSIM	LPIPS	min	PSNR	SSIM	LPIPS	min	PSNR	SSIM	LPIPS	min
3DGS	23.783	.8536	.1692	**16.80**	**29.845**	**.9077**	.2378	**29.27**	29.311	.8801	.1668	**29.18**	29.499	.9047	.1586	**23.55**
3DGS-MCMC	24.573	.8711	**.1466**	26.38	29.423	.9064	.2367	41.53	**29.846**	**.8963**	**.1453**	48.34	**29.625**	**.9120**	**.1455**	32.70
KB-Splatting	**24.633**	**.8735**	.1475	25.30	29.723	.9068	**.2343**	39.44	29.434	.8896	.1536	55.46	29.491	.9094	.1526	34.55

I GENERALIZATION TO RGB NOVEL-VIEW SYNTHESIS

The RGB study tests whether the finite-support representation survives a different image-formation model. It is deliberately separated from the CT claims: no RGB result is used to select the CT recipe, and we do not claim state-of-the-art novel-view synthesis.

I.1 PROJECTED COVERAGE RENDERING

For a depth-sorted primitive, we evaluate its projected KB footprint $G_i(\mathbf{u}) \in [0, 1]$ and define opacity as

$$\alpha_i(\mathbf{u}) = a_i G_i(\mathbf{u}), \qquad \mathbf{C}(\mathbf{u}) = \sum_i \mathbf{c}_i(\mathbf{d})\alpha_i(\mathbf{u}) \prod_{j<i} \big(1 - \alpha_j(\mathbf{u})\big), \tag{13}$$

where a_i is center opacity and $\mathbf{c}_i(\mathbf{d})$ is degree-three spherical-harmonic color. This projected-coverage rule replaces the earlier Beer–Lambert opacity derivation. Population changes preserve center opacity for retained primitives and allocate coverage through the new supports; they do not claim exact preservation of the complete image-space footprint.

Image residuals are lifted into three dimensions by sampling candidate depths along high-error training rays and scoring candidates by local coverage and visibility. This lift is ambiguous under occlusion, so proposals are restricted to the scene bounds, assigned to nearby high-opacity parents, and subjected to the same population cap and optimizer-state reset used by the final RGB implementation. This follows the separation between basis and population dynamics emphasized by 3DGS-MCMC (Kheradmand et al., 2024), while the use of a non-Gaussian footprint is related to the generalized exponential family in GES (Hamdi et al., 2024).

I.2 PROTOCOL

We evaluate 19 scenes: two from Tanks and Temples (Knapitsch et al., 2017), two from Deep Blending (Hedman et al., 2018), seven from Mip-NeRF 360 (Barron et al., 2022), and eight from OMMO (Lu et al., 2023). All methods use their established scene preprocessing and the same held-out views, with every eighth image used for evaluation where the benchmark does not prescribe a split. KB-Splatting uses the projected-coverage renderer above, SfM initialization, 30,000 iterations, and scene-specific population caps. The matched 3DGS-MCMC baseline uses SfM initialization and the same caps; standard 3DGS retains its conventional uncapped population. All final RGB runs use NVIDIA L40S GPUs. We report PSNR, SSIM (Wang et al., 2004), LPIPS (Zhang et al., 2018), and observed training time.

I.3 QUANTITATIVE RESULTS

Table 5 follows the benchmark-grouped presentation used by GES (Hamdi et al., 2024). Across all 19 scenes, KB-Splatting obtains 28.983 dB PSNR, 0.8981 SSIM, and 0.1610 LPIPS in 41.80 minutes. Standard 3DGS obtains 28.865 dB, 0.8906, and 0.1711 in 25.52 minutes; 3DGS-MCMC obtains 29.153 dB, 0.9013, and 0.1551 in 38.73 minutes. KB therefore improves over standard 3DGS by 0.118 dB and 0.010 LPIPS, but trails MCMC by 0.170 dB and 0.006 LPIPS. Benchmark means show the same mixed pattern: KB leads PSNR on Tanks and Temples and Deep Blending, while MCMC leads on Mip-NeRF 360 and OMMO.

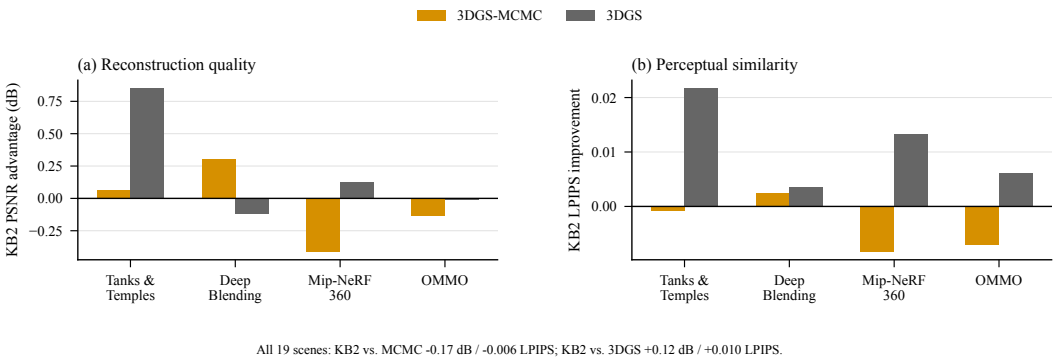

Figure 13: **RGB benchmark trade-offs.** Quality and observed training time for projected-coverage KB-Splatting, standard 3DGS, and 3DGS-MCMC across the four benchmark groups.

I.4 Qualitative comparison

Following the matched-view, region-of-interest layout used by GES, Figure 14–15 places standard 3DGS, 3DGS-MCMC, KB-Splatting, and ground truth in fixed columns. Each scene uses the same held-out frame and two shared crops, chosen where the final KB rendering preserves high-frequency structure. The figure is diagnostic rather than a randomized visual sample and must be read with Table 5.

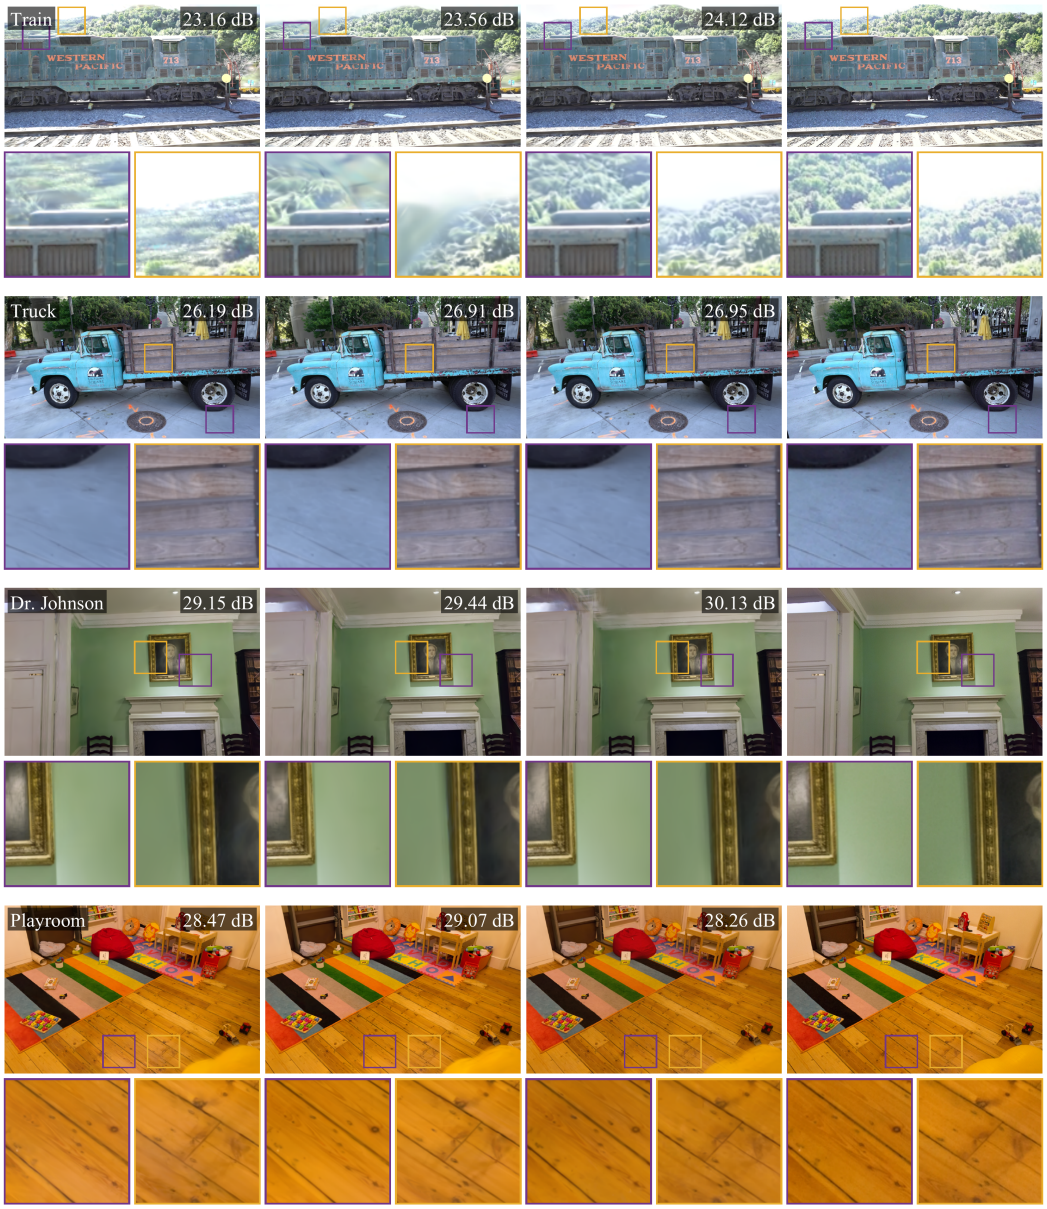

Figure 14: **RGB qualitative comparison, part I.** Matched held-out views, shared regions of interest, and per-view PSNR for the first four scenes.

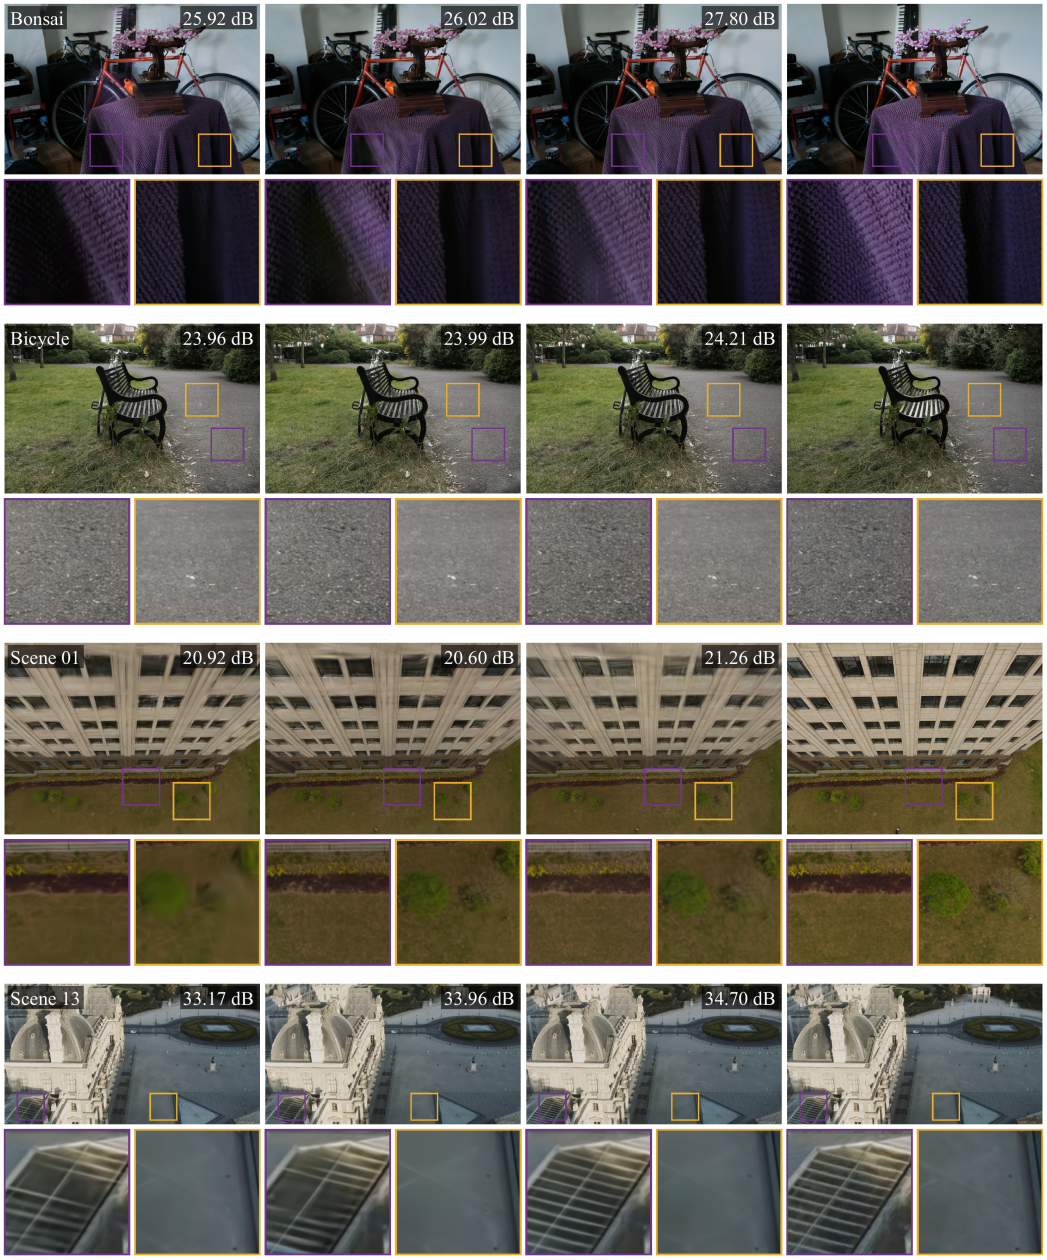

Figure 15: **RGB qualitative comparison, part II.** The remaining four scenes under the same display protocol.

Table 6: Historical RGB allocation control under the earlier exponential-opacity renderer. These results guide population allocation but should not be compared numerically with Table 5.

Policy	PSNR	SSIM	LPIPS	min
Global opacity	22.295	.8291	.1979	19.61
Max residual, direct, 1.0	21.068	.8033	.2256	18.26
Expected residual, direct, 1.0	10.639	.4157	.6462	16.01
Max residual, $K = 32$, pool 40%, 1.0	22.472	.8265	.1973	18.79
Max residual, $K = 32$, pool 20%, 0.5	**22.582**	**.8308**	**.1900**	**17.93**

I.5 RGB ABLATIONS

Table 6 records the allocation study that preceded the final projected-coverage renderer. Because these runs used the earlier exponential-opacity formulation, they establish the direction of the population search but are not direct ablations of the final KB2 renderer. Capping candidates at $K = 32$, pooling the top 20%, and allocating half of the eligible pool gives the best recorded balance among the tested policies. A later K sweep selects $K = 50$ for the final implementation: it reaches 22.816 dB mean PSNR and 0.1729 LPIPS, with a 30% pool retained for stability. The photometric factorial selects L2 color loss with DSSIM weight 0.10, yielding 28.263 dB, 0.9099 SSIM, and 0.1852 LPIPS on its development scenes.

Finally, replacing exponential opacity with projected coverage reduces training time in all three paired implementation checks: 17.60 to 15.94 minutes on OMMO scene 06, 19.70 to 17.13 on Train, and 38.08 to 33.74 on Playroom. PSNR changes are -0.158, -0.039, and -0.075 dB, respectively. This is the intended engineering trade-off: the final renderer removes an unnecessarily expensive CT-derived opacity computation while preserving the RGB operating point.