GROUP-WISE SIMPLIFICATION OF RADIATIVE GAUSSIAN SPLATTING FOR COMPUTED TOMOGRAPHY GUIDED BY CONVEX DENSITY OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Radiative Gaussian splatting achieves high fidelity computed tomography (CT) reconstruction, but heuristic densification causes severe primitive redundancy. Pruning redundant primitives is challenging because naive removal damages subtle structural boundaries and induces streak artifacts. To resolve this issue, we propose CDGS, a structure preserving group-wise simplification framework guided by convex density optimization. By exploiting the linearity of CT projection physics under fixed geometry, we formulate candidate pricing as a convex quadratic subproblem solved via sketching. Primitives are partitioned into topological structural groups that confine density compensation within contiguous structural boundaries, which prevents non local ray interference. In addition, an autonomous in-group stopping rule dynamically terminates simplification when redundancy is exhausted. Experiments across multiple CT benchmarks demonstrate that CDGS substantially reduces primitive counts while maintaining reconstruction fidelity. Our framework outperforms existing state-of-the-art methods in both numerical accuracy and structural preservation.

1 INTRODUCTION

3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) represents 3D scenes using a collection of anisotropic Gaussian primitives, enabling high-fidelity scene reconstruction. Recent work has extended this paradigm to computed tomography (CT) by modeling the attenuation volume as a sum of *radiative Gaussian primitives* (Zha et al., 2024; Pieta et al., 2026). These primitives contribute additively to the predicted projections through line integration along X-ray paths and are optimized directly against measured projections. However, existing Gaussian-based CT reconstruction frameworks often exhibit substantial primitive redundancy. In sparse-view CT, the limited number of projection views makes the discrete forward projection problem highly underdetermined. Under these conditions, gradient-based densification can generate excessive overlapping primitives along ambiguous ray trajectories, resulting in redundant representations in which many primitives contribute little to distinct structural features. For standard 3DGS, existing approaches to reducing primitive redundancy mainly follow two strategies: budget-controlled simplification (Fan et al., 2024; Hanson et al., 2025; Wang et al., 2025; Xiong et al., 2026), which prunes or merges primitives toward a prescribed budget, and adaptive density control (Yuluo et al., 2026; Ren et al., 2026; Zhou & Ni, 2025), which adjusts their number and distribution during training.

One major drawback of these methods is that the properties of radiative Gaussian primitives in the CT domain have not been fully considered. In CT, projection measurements integrate attenuation along X-ray paths (Kak & Slaney, 2001), accumulating contributions from multiple Gaussian primitives. Under fixed geometry, these projections depend linearly on density, enabling convex density fitting as a principled reference for evaluating simplification. Existing simplification criteria, primarily developed for nonlinear alpha blending, do not explicitly exploit this CT-specific property. Furthermore, because the forward projection operator in sparse-view CT is severely underdetermined, it possesses an extensive null space (Kuo et al., 2022). Spatially unrelated primitives can therefore produce similar projection responses, allowing non-trivial volumetric perturbations to produce near-zero projection residuals ($\mathbf{Ad} \approx \mathbf{0}$). Under unconstrained simplification, distant primitives along intersecting ray paths can falsely absorb the residual of a removed primitive. This

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.

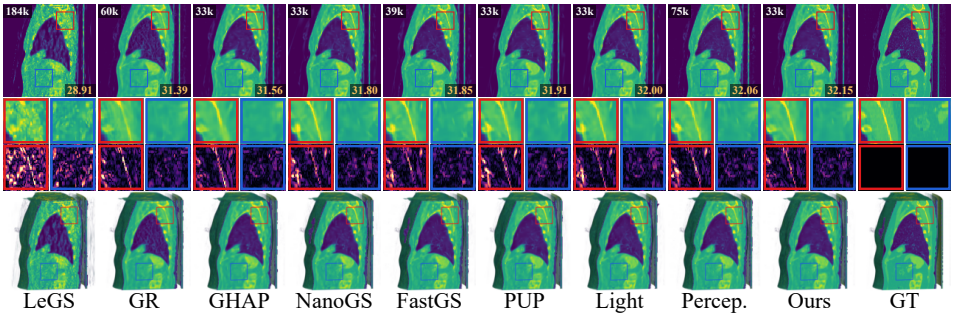

Figure 1: Simplifying Gaussian primitives for CT reconstruction often degrades local structures or introduces artifacts. In comparison with state-of-the-art simplification methods, CDGS achieves higher PSNR at a matched or lower primitive count while preserving fine structural details.

non-local compensation cancels projection discrepancies while introducing streak artifacts (Bubba et al., 2019) and damaging structural boundaries, as illustrated in Figure 1. Without sufficient control over these effects, naive pruning and merging may compromise subtle anatomical details or obscure small material defects, raising safety and reliability concerns in clinical diagnosis and industrial inspection.

To address this limitation, we propose CDGS, a group-wise simplification framework based on convex density optimization and designed for seamless integration into radiative Gaussian CT reconstruction pipelines. By fixing candidate geometry, we decouple geometric actions from density optimization and exploit the linear projection model to formulate density fitting as a convex non-negative problem. We further propose a geometry-based structural grouping that confines simplification within each group, so that unrelated structures cannot compensate for one another in projection space and mask local damage. Our main contributions are summarized as follows:
(1) Physics-Grounded Formulation: We decouple geometric actions from density optimization and formulate primitive simplification under fixed geometry as a verifiable density optimization problem derived from linear X-ray projection physics.
(2) Structural Grouping for Spatial Coherence: We develop a training-free topological grouping mechanism that confines primitive compensation within contiguous structural boundaries and suppresses tomographic null-space streak artifacts.
(3) Autonomous Group-Wise Stopping: We design an in-group relative stopping rule that dynamically detects when representational redundancy is exhausted within each structural group to provide robust model simplification without manual error thresholds or a prescribed target count.
(4) Benchmark Performance and Quantitative Evaluation: Experiments across two CT reconstruction backbones and three sparse-view configurations demonstrate that CDGS removes up to 47.6% of Gaussian primitives while lowering the mean PSNR by only 0.03 dB. Furthermore, it achieves the lowest Gaussian error product in every setting and establishes a new state of the art in rate-distortion performance while preserving fine structures in the reconstructed volumes.

2 Related Work

2.1 Gaussian splatting for CT reconstruction

3DGS has been extended to X-ray imaging by modeling each Gaussian primitive as a radiative primitive whose projection accumulates additively rather than through alpha blending (Cai et al., 2024), and subsequently to CT reconstruction with physically consistent line integrals (Zha et al., 2024). Later works focus on accelerating reconstruction throughput via optimized rasterization pipelines (Pieta et al., 2026), or improving sparse-view structural fidelity through graph-regularized gradient refinement (GR-Gaussian; Yuluo et al. 2026). These methods still yield a large number of Gaussian primitives. Because interactive inspection requires repeated rotation and re-slicing, this inflated primitive count incurs per-frame latency and substantial transfer overhead across medical networks, making compact post-reconstruction representation essential for clinical workflow.

2.2 SIMPLIFICATION OF 3DGS

Existing simplification methods of 3DGS fall into two categories. **Budget-controlled simplification** reduces an existing Gaussian representation toward a prescribed size through importance-based pruning or Gaussian mixture reduction. LightGaussian (Fan et al., 2024) combines opacity, volume, and hit count to construct a global significance score for Gaussian simplification. PUP 3D-GS (Hanson et al., 2025) evaluates Gaussian importance using a second-order sensitivity metric derived from a Fisher approximation of the Hessian. GHAP (Wang et al., 2025) performs blockwise Gaussian mixture reduction under an optimal transport divergence. NanoGS (Xiong et al., 2026) merges neighboring Gaussian pairs through moment matching using an I-divergence criterion without requiring training images. **Adaptive density control** adjusts the number and distribution of Gaussian primitives during reconstruction training using reconstruction feedback or learned decision policies. LeGS (Ning et al., 2026) formulates density control as a reinforcement learning problem and trains a policy network to determine simplification actions. FastGS (Ren et al., 2026) further incorporates multi-view rendering errors into the density-control process to identify Gaussian primitives that should be densified or removed. Perceptual-GS (Zhou & Ni, 2025) introduces perceptual sensitivity maps to guide the allocation and pruning of Gaussian primitives. GR-Gaussian combines denoised point-cloud initialization with pixel-graph-aware gradient refinement to improve Gaussian splitting and suppress needle-like artifacts in sparse-view CT reconstruction (Yuluo et al., 2026). Despite their differences, these methods are primarily designed for nonlinear alpha blending, where removing a Gaussian primitive can alter the contributions of those behind it along the same ray. Directly transferring their simplification criteria to radiative Gaussian primitives may compromise CT reconstruction quality, potentially obscuring structural details in safety-critical applications.

3 PRELIMINARIES

3.1 RADIATIVE GAUSSIAN SPLATTING

In radiative Gaussian splatting (Zha et al., 2024; Pieta et al., 2026; Yuluo et al., 2026), the attenuation field is represented by the additive contributions of N radiative Gaussian primitives:

$$\sigma(\boldsymbol{x}) = \sum_{i=1}^{N} G_i^3(\boldsymbol{x} \mid \rho_i, \boldsymbol{p}_i, \boldsymbol{\Sigma}_i), \qquad G_i^3(\boldsymbol{x} \mid \rho_i, \boldsymbol{p}_i, \boldsymbol{\Sigma}_i) = \rho_i \exp\left(-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{p}_i)^\top \boldsymbol{\Sigma}_i^{-1}(\boldsymbol{x}-\boldsymbol{p}_i)\right), \tag{1}$$

where $\boldsymbol{x}, \boldsymbol{p}_i \in \mathbb{R}^3$ denote the coordinate and Gaussian centroid, $\boldsymbol{\Sigma}_i \in \mathbb{R}^{3\times3}$ is a positive-definite covariance matrix, and $\rho_i \geq 0$ is the peak attenuation density. Density is the sole non-geometric attribute, and no appearance parameters are used.

3.2 LINEARITY OF RADIATIVE GAUSSIAN PROJECTION IN DENSITY

In CT, forward projection of the continuous attenuation field $\sigma(\boldsymbol{x})$ is modeled as line integration along X-ray paths. Let $\boldsymbol{r}(t) = \boldsymbol{o} + t\,\boldsymbol{d}$ denote a ray path parameterized by depth t, with source position $\boldsymbol{o} \in \mathbb{R}^3$ and unit direction $\boldsymbol{d} \in \mathbb{R}^3$ toward a detector pixel. Under the Beer-Lambert law (Kak & Slaney, 2001), the idealized projection measurement is given by: $y(\boldsymbol{r}) = \int \sigma(\boldsymbol{r}(t))\,\mathrm{d}t$. Substituting Eq. (1) and factoring out the ray-independent density ρ_i gives

$$\hat{y}(\boldsymbol{r}_k) = \sum_{i=1}^{N} \rho_i \underbrace{\int \exp\left(-\frac{1}{2}(\boldsymbol{r}_k(t)-\boldsymbol{p}_i)^\top \boldsymbol{\Sigma}_i^{-1}(\boldsymbol{r}_k(t)-\boldsymbol{p}_i)\right)\mathrm{d}t}_{a_i(\boldsymbol{r}_k)}. \tag{2}$$

Here, $a_i(\boldsymbol{r}_k)$ is the unit-density projection footprint of Gaussian i along ray $\boldsymbol{r}_k$. Under fixed geometry, these footprints are constant, making the projection linear in the densities. In practice, this footprint is computed with the closed-form Gaussian projection of (Zha et al., 2024). It uses a local affine approximation for cone-beam geometry and includes a covariance-dependent amplitude factor that corrects the integration bias. Each Gaussian thus contributes its own term, and no term depends on any other Gaussian. For ray $\boldsymbol{r}_k$, the sum is an inner product between the footprints and the densities, so each ray contributes one row. Collecting the densities into $\boldsymbol{\rho}$ and all K_{ray} rays,

$k = 1, \ldots, K_{ray}$, into $\boldsymbol{A} \in \mathbb{R}^{K_{ray} \times N}$, the forward model reads:

$$\hat{y}(\boldsymbol{r}_k) = [a_1(\boldsymbol{r}_k), \ldots, a_N(\boldsymbol{r}_k)]\, \boldsymbol{\rho}, \qquad \boldsymbol{\rho} = [\rho_1, \ldots, \rho_N]^\top \in \mathbb{R}^N, \qquad \hat{\boldsymbol{y}} = \boldsymbol{A}\boldsymbol{\rho}, \tag{3}$$

where $\hat{\boldsymbol{y}} \in \mathbb{R}^{K_{ray}}$ stacks the projection values. The i-th column $\boldsymbol{A}_{:,i} = \left[a_i(\boldsymbol{r}_1), \ldots, a_i(\boldsymbol{r}_{K_{ray}})\right]^\top$ is the footprint of the i-th Gaussian over all rays. The projection is therefore a linear function of $\boldsymbol{\rho}$, and removing a Gaussian removes one column of $\boldsymbol{A}$ while leaving the others unchanged.

4 Method

4.1 Overview

As illustrated in Figure 2, CDGS interleaves backbone training with E simplification events. CDGS incorporates a Convex Density Optimizer (CDO), which solves density-fitting subproblems under fixed geometry for both in-group candidate pricing and batch-level density adjustment. The first event follows initial densification, fitting measured projections before reconstruction training. Subsequent events occur at predetermined training iterations and use sketched residuals relative to the current reconstructed projection. Each event constructs structural groups from the current Gaussian geometry. Each round first prices pruning and merging candidates through in-group CDO, then checks each active group once against the stopping criterion. Stopped groups' remaining primitives receive keep status for the rest of the event. Compatible candidates from unfrozen primitives form a batch in increasing price order. Following tentative execution, CDO is applied again to adjust densities within the affected groups under fixed trial geometry, after which batch acceptance is evaluated. Accepted trials commit updates, followed by regrouping and repricing; rejected trials are rolled back and retried according to the prescribed backoff procedure. Keep status persists across regrouping. The event ends when no admissible candidates remain or the trust-region rejection floor is reached. Reconstruction training resumes after each event and continues after the final event until the total training budget is exhausted. Sections 4.2–4.4 detail the method; implementation details appear in Appendix A.3.

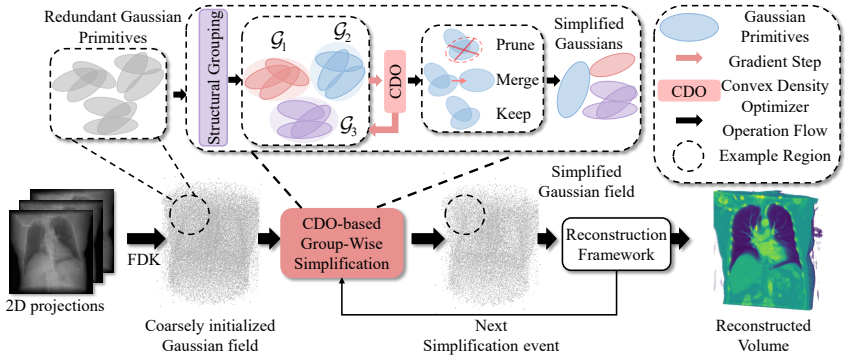

Figure 2: Overview of CDGS. The lower pathway shows the pipeline, initialized by Feldkamp-Davis-Kress (FDK) reconstruction (Feldkamp et al., 1984). The upper inset details our simplification module. Dashed circles mark the same random region, which becomes sparser after simplification.

4.2 Fixed-geometry Convex Density Optimization

Optimizing Gaussian geometry and density simultaneously is challenging because projection footprints depend nonlinearly on centroid positions and covariance matrices via Eq. (2), which produces a severely non-convex landscape. However, as established in Section 3.2, once Gaussian geometry $\{(\mathbf{p}_i, \boldsymbol{\Sigma}_i)\}_{i=1}^N$ is held fixed, the forward projection in Eq. (3) becomes a linear operator $\hat{\mathbf{y}} = \boldsymbol{A}\boldsymbol{\rho}$. Radiometric fitting then reduces to a standard convex non-negative least-squares (NNLS) problem (Lawson & Hanson, 1995) and the minimal density cost C is defined as:

$$C = \min_{\boldsymbol{\rho} \geq \mathbf{0}} f(\boldsymbol{\rho}) = \min_{\boldsymbol{\rho} \geq \mathbf{0}} \frac{1}{2} \|\mathbf{y} - \boldsymbol{A}\boldsymbol{\rho}\|^2 \tag{4}$$

Under fixed geometry, Eq. (4) is a standard convex optimization problem (Vandenberghe & Boyd, 2004). Although the tomographic null space permits non-unique optimal density vectors $\boldsymbol{\rho}^\star$, the fitted projection $\boldsymbol{A}\boldsymbol{\rho}^\star$ and the minimal residual cost C are unique. These properties ensure that C provides an invariant and unambiguous measure of the best projection fidelity achievable by a given geometric configuration.

4.3 Structural Grouping of Radiative Gaussian Primitives

Similar projection-domain costs do not necessarily imply similar effects on 3D structure, as underdetermined tomographic measurements may fail to distinguish spatially unrelated primitives (Kak & Slaney, 2001). We therefore exploit local spatial coherence to organize Gaussian primitives into structural groups. We construct these groups from Gaussian geometry without anatomical labels. To account for both centroid displacement and covariance differences, we adopt the Bhattacharyya distance (Bhattacharyya, 1946):

$$D_B(i,j) = \frac{1}{8}(\mathbf{p}_i - \mathbf{p}_j)^\top \bar{\boldsymbol{\Sigma}}_{ij}^{-1}(\mathbf{p}_i - \mathbf{p}_j) + \frac{1}{2}\ln \frac{\det(\bar{\boldsymbol{\Sigma}}_{ij})}{\sqrt{\det(\boldsymbol{\Sigma}_i)\det(\boldsymbol{\Sigma}_j)}}, \tag{5}$$

where $\bar{\boldsymbol{\Sigma}}_{ij} = (\boldsymbol{\Sigma}_i + \boldsymbol{\Sigma}_j)/2$. This measure accounts for anisotropic separation and differences in kernel scale and orientation. Rather than evaluating all $\mathcal{O}(N^2)$ pairs, we build a sparse graph. For each primitive i, we retrieve K_{cand} candidates by Euclidean centroid distance and retain the K_{nn} nearest under D_B, forming $\mathcal{N}_i$. Mutual neighbors receive affinity weights combining local distance normalization (Zelnik-manor & Perona, 2004) and Jaccard neighborhood overlap (Liben-Nowell & Kleinberg, 2007):

$$w_{ij} = \exp\left(-\frac{D_B(i,j)}{\varsigma_i\varsigma_j + \epsilon}\right)\frac{|\mathcal{N}_i \cap \mathcal{N}_j|}{|\mathcal{N}_i \cup \mathcal{N}_j|}, \quad j \in \mathcal{N}_i,\ i \in \mathcal{N}_j, \tag{6}$$

where K_{cand} is the initial candidate pool size, and K_{nn} is the number of nearest neighbors. $\varsigma_i = \max_{j\in\mathcal{N}_i} D_B(i,j)$, which represents the local neighborhood distance scale. ϵ denotes a positive numerical safeguard against division by zero. This construction yields symmetric weights and removes connections without shared neighbors. All other weights, including self-edges, are zero.

We then apply Leiden (Traag et al., 2019) to obtain connected communities by optimizing weighted modularity:

$$Q(\mathbf{c}) = \frac{1}{2W}\sum_{i,j=1}^{N}\left(w_{ij} - \gamma\frac{D_i D_j}{2W}\right)\mathbf{1}[c_i = c_j], \tag{7}$$

where c_i denotes the community assignment, $D_i = \sum_j w_{ij}$ is the weighted degree, $W = \frac{1}{2}\sum_i D_i$ is the total edge weight, and $\gamma > 0$ is the Leiden resolution, which denotes the modularity resolution parameter governing community granularity. The resulting number of groups H is determined by the partition rather than specified beforehand:

$$\mathcal{G}_h = \left\{G_i^3 \in \mathcal{S} \mid c_i = h\right\}, \qquad h = 1, \ldots, H. \tag{8}$$

These disjoint groups define the scope of the in-group pruning, merging, and retention decisions in Section 4.4. Implementation details are provided in Appendix A.3.

4.4 Group-Wise Simplification via Convex Density Cost

4.4.1 Simplification Actions

We consider three action types in the candidate set $\mathcal{A}$: pruning, merging, and keeping, all constructed within each structural group. To avoid $O(|\mathcal{G}_h|^2)$ complexity and preserve local spatial coherence, merge candidates (i,j) are restricted to pairs adjacent in the structural adjacency graph of group $\mathcal{G}_h$. For each candidate action $a \in \mathcal{A}$, the resulting Gaussian configuration $\mathcal{S}_a$ is

$$\mathcal{S}_{\text{prune}(i)} = \mathcal{S} \setminus \{G_i^3\}, \quad \mathcal{S}_{\text{merge}(i,j)} = \left(\mathcal{S} \setminus \{G_i^3, G_j^3\}\right) \cup \{G_c^3\}, \quad \mathcal{S}_{\text{keep}} = \mathcal{S}. \tag{9}$$

where G_c^3 replaces the parents G_i^3 and G_j^3. Its parameters are initialized by matching the zeroth, first, and second spatial moments of the parents, which preserves total attenuation mass, center of mass, and directional dispersion (Runnalls, 2007). Pruning and merging both reduce the Gaussian primitive count by one, so their costs are directly comparable for the same unit reduction. Gaussian primitives not consumed by accepted actions are kept.

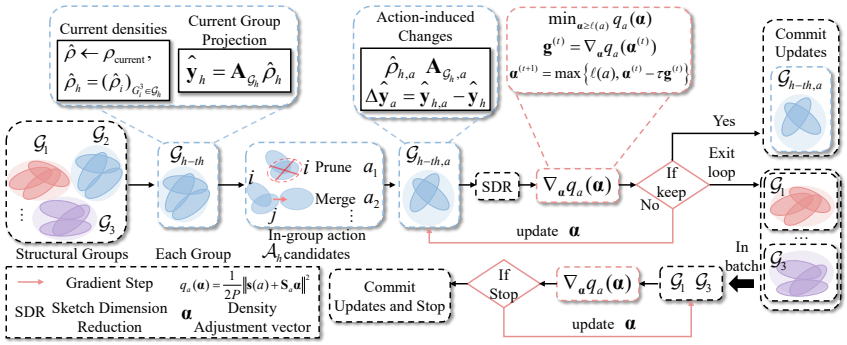

Figure 3: Workflow of CDO based candidate prioritization. Action candidates are constructed within structural groups. Their projection perturbations are evaluated through density adjustment in a shared sketch space.

4.4.2 Candidate Prioritization via In-group Pricing

Figure 3 illustrates CDO-based candidate prioritization. The initial event precedes reconstruction training and fits densities to measured projections through Eq. (4). At subsequent events, the Gaussian field has already been refined by reconstruction training. To preserve this reconstruction during simplification, we use its current density vector $\boldsymbol{\rho}_{\text{current}}$ as the evaluation densities, setting $\hat{\boldsymbol{\rho}} \leftarrow \boldsymbol{\rho}_{\text{current}}$, and take the corresponding projection as the reference. CDO minimizes the sketched residual between the candidate projection and this reference through in-group density adjustment.

For group $\mathcal{G}_h$, each action $a \in \mathcal{A}_h$ is either $\text{prune}(i)$ or $\text{merge}(i, j)$. It produces a trial group $\mathcal{G}_{h,a}$ with densities $\hat{\boldsymbol{\rho}}_{h,a}$ before compensation, retaining surviving densities and initializing the merged primitive. The resulting projection residual is

$$\Delta\hat{\mathbf{y}}_a = \hat{\mathbf{y}}_{h,a} - \hat{\mathbf{y}}_h = \boldsymbol{A}_{\mathcal{G}_{h,a}}\hat{\boldsymbol{\rho}}_{h,a} - \boldsymbol{A}_{\mathcal{G}_h}\hat{\boldsymbol{\rho}}_h. \tag{10}$$

Following the NNLS formulation in Eq. (4), we optimize in-group density increments $\boldsymbol{\alpha}$ to minimize the residual relative to the current group projection. To accelerate this minimization, we apply sketch-based dimensionality reduction (SDR) using a shared unnormalized matrix $\boldsymbol{\Omega} \in \mathbb{R}^{P\times K_{ray}}$, where $P \ll K_{ray}$. With candidate action a and its geometry fixed, CDO solves

$$\min_{\boldsymbol{\alpha}\geq\boldsymbol{\ell}(a)} q_a(\boldsymbol{\alpha}), \qquad q_a(\boldsymbol{\alpha}) = \frac{1}{2P}\left\| \underbrace{\boldsymbol{\Omega}\Delta\hat{\mathbf{y}}_a}_{\mathbf{s}(a)} + \underbrace{\boldsymbol{\Omega}\boldsymbol{A}_{a,\text{adj}}}_{\boldsymbol{S}_a}\boldsymbol{\alpha} \right\|^2 \tag{11}$$

We prove the nonnegative least squares equivalence of the uncompressed pricing problem in Appendix A.1 and validate the effectiveness of sketch based dimensionality reduction empirically in Appendix A.3.3. Here, $\Delta\hat{\mathbf{y}}_a$ is the projection residual induced by action a before density adjustment, and $\boldsymbol{A}_{a,\text{adj}}$ contains the columns of $\boldsymbol{A}_{\mathcal{G}_{h,a}}$ corresponding to adjustable in-group primitives. Thus, $\mathbf{s}(a)$ is the compressed initial residual, while $\boldsymbol{S}_a\boldsymbol{\alpha}$ represents the sketched projection change caused by density adjustment vector $\boldsymbol{\alpha}$. Their sum represents the residual after adjustment in sketch space, and $\boldsymbol{\ell}(a)$ specifies the lower bounds on these increments:

$$\ell_j(a) = -\hat{\rho}_j \text{ (surviving)}, \qquad \ell_j(a) = 0 \text{ (newly merged)}. \tag{12}$$

These bounds preserve non-negative surviving densities and restrict the new primitive to non-negative increments from its initialized density. CDO iterates over $\boldsymbol{\alpha}$ using

$$\boldsymbol{\alpha}^{(t+1)} = \max\left\{\boldsymbol{\ell}(a), \boldsymbol{\alpha}^{(t)} - \tau\nabla_{\boldsymbol{\alpha}} q_a(\boldsymbol{\alpha}^{(t)})\right\} \tag{13}$$

where $\tau = 1/\operatorname{tr}(\mathbf{G}_a)$ is the candidate-specific step size, $\mathbf{G}_a = P^{-1}\boldsymbol{S}_a^\top\boldsymbol{S}_a$ is the Hessian of q_a, and the maximum is taken componen-twise, with $\operatorname{tr}(\cdot)$ is the matrix trace. Upon numerical termination, the returned feasible adjustment $\hat{\boldsymbol{\alpha}}_a$ defines $c(a) = q_a(\hat{\boldsymbol{\alpha}}_a)$. Admissible candidates are ranked by increasing $c(a)$. Sketch construction, adjustable-set selection, and solver settings are provided in Appendix A.3.3.

4.4.3 GROUP- AND BATCH-LEVEL SIMPLIFICATION DECISIONS

Using the in-group prices $c(a)$, each round r checks all active groups once before batch construction. Each primitive stores a fixed reference b_i, initialized to its group's median pruning price when first priced; newly merged primitives follow the same rule. For eligible indices I_h in group $\mathcal{G}_h$,

$$\nu_h = \underset{i \in I_h}{\text{median}}\, c(\text{prune}(i)), \qquad \bar{b}_h = \underset{i \in I_h}{\text{median}}\, b_i, \qquad \text{keep if } \bar{b}_h > 0 \text{ and } \nu_h > \lambda \bar{b}_h. \tag{14}$$

where $\lambda = 1.2$ denotes the relative degradation tolerance multiplier that balances primitive reduction against local reconstruction error. Stopped groups' remaining primitives are ineligible for simplification for the rest of the event. From the remaining eligible candidates, we select mutually compatible actions in increasing price order to form a trial batch $\mathcal{B}_r$, where $b_r = |\mathcal{B}_r|$ is the trial batch size, $1 \le b_r \le b_{\max}$, and $b_{\max}$ is its adaptive upper limit. After tentative execution, a second CDO solve jointly adjusts densities within affected groups under fixed trial geometry. Let Φ_r denote the objective value before trial execution and F_b the value after executing a batch of $b = b_r$ actions and adjusting its densities. Both are evaluated using the same stage-specific target, normalization, and sketch operator when applicable. The batch is accepted if

$$\eta_r = \frac{\max\{0, F_b - \Phi_r\}}{b_r \max\{med_r, \epsilon_c\}} \le M, \tag{15}$$

where med_r is the median pruning price over the remaining eligible primitives before the trial, $M > 0$ is the batch acceptance threshold, and $\epsilon_c = 10^{-30}$ is a positive regularization floor in the denominator that prevents division by zero. Accepted trials commit geometry and densities, followed by regrouping and repricing for the next round while preserving reference prices and keep flags. Rejected trials restore the pre-trial state and retry with the batch size halved. The event ends when no admissible candidates remain or a single-action batch is rejected, after which training resumes. Full hyperparameter settings are given in Appendix A.3.1.

5 EXPERIMENTS

5.1 EXPERIMENTAL SETTINGS

Dataset We used the CT benchmark introduced by R^2-Gaussian (Zha et al., 2024) and also used by FaCT-GS (Pieta et al., 2026). It comprises synthetic projections of 15 CT volumes covering human organs, animals, plants, and artificial objects, together with three real scans from FIPS (The Finnish Inverse Problems Society, 2024): pine, seashell, and walnut (Meaney, 2022a; Kamutta et al., 2022; Meaney, 2022b). Synthetic projections were generated using TIGRE (Biguri et al., 2016) with simulated acquisition noise, and each real scan contains 721 projections. We evaluated sparse-view reconstruction with 75, 50, and 25 input projections.

Implementation details We integrated CDGS into R^2-Gaussian and FaCT-GS, retaining each backbone's initialization, reconstruction objective, optimizer, and other original settings. The reconstruction budgets are 30,000 and 20,000 optimization steps, respectively. After densification, CDGS was applied at steps 16,001, 19,001, and 22,001 for R^2-Gaussian, and at steps 11,001, 14,001, and 17,001 for FaCT-GS. Structural groups were rebuilt at each event. We set the relative-price threshold to $\lambda = 1.2$ and the batch acceptance threshold to $M = 4$ for both backbones, keeping them fixed across objects and view settings. All methods ran on a single NVIDIA RTX PRO 6000 GPU.

Evaluation metrics We evaluated reconstruction quality using PSNR and SSIM, following R^2-Gaussian. PSNR was computed over the three-dimensional volume, and SSIM was averaged over two-dimensional slices in the axial, coronal, and sagittal directions. We also reported the final Gaussian primitive count N and the normalized Gaussian–error product in Eq. (16). Metrics were averaged over objects within each backbone and view setting, with GEP computed per object before averaging. Here is the definition of GEP:

$$\text{GEP} = \frac{N}{N_b}\, 10^{(\text{PSNR}_b - \text{PSNR})/10}, \tag{16}$$

where the subscript b denotes the corresponding backbone baseline. With the same reference volume and intensity range, the exponential term is the relative mean squared error. Lower GEP therefore indicates a smaller count–error product, and the baseline has GEP equal to one.

5.2 BENCHMARK COMPARISONS

We evaluated CDGS against two categories of Gaussian simplification methods on the R^2-Gaussian backbone: **budget-controlled simplification** (NanoGS, LightGaussian, GHAP, PUP 3D-GS) and **adaptive density control** (GR-Gaussian, LeGS, FastGS, Perceptual-GS). Quantitative results are presented in Table 1. Across all three view settings, CDGS consistently achieves the lowest Gaussian Error Product (GEP). Relative to the unsimplified R^2-Gaussian baseline, CDGS removes up to 47.6% of Gaussian primitives and reduces parameter storage from 2.90 MB to 1.52 MB, while the overall mean PSNR decreases by only 0.03 dB. At comparable primitive counts, CDGS consistently outperforms the four budget-controlled methods in PSNR, surpassing GHAP by up to 1.01 dB. Among adaptive density control methods, CDGS surpasses LeGS by up to 5.11 dB while retaining substantially fewer Gaussian primitives. While Perceptual-GS achieves slightly higher raw PSNR in the dense 75-view configuration (39.09 dB vs. 38.93 dB), it requires 78.93k primitives and 3.31 MB storage (more than double that of CDGS), yielding a substantially inferior rate-distortion trade-off (GEP of 1.138 vs. 0.540). In the challenging 25-view sparse setting, CDGS achieves the highest reconstruction fidelity (35.46 dB PSNR, 0.913 SSIM) and the lowest GEP (0.529) across all evaluated methods. Figure 4 presents qualitative comparisons under different view configurations, where enlarged insets demonstrate CDGS's preservation of fine structures. Per-scene breakdowns and evaluations on FaCT-GS backbone appear in Table A6 of Appendix A.5.1.

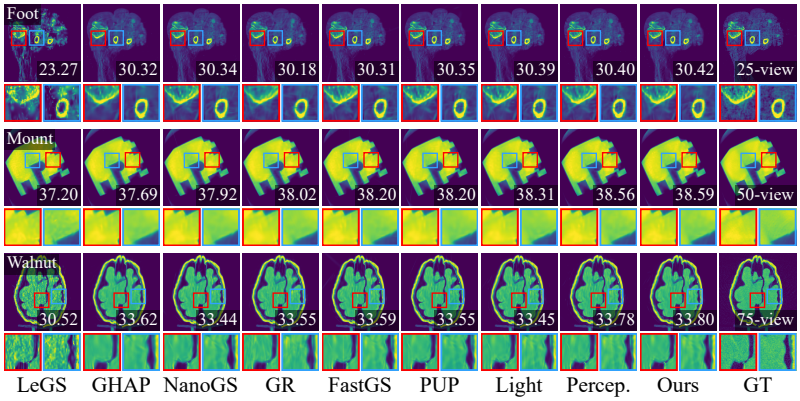

Figure 4: Qualitative comparison with competing methods under different view configurations across three anatomical specimens. Yellow insets highlight regions prone to reconstruction artifacts. Values in each panel denote volumetric PSNR (dB).

Table 1: Quantitative evaluation of CDGS on the R^2-Gaussian backbone in comparison with benchmark methods. PSNR is in dB, SSIM is averaged over axial, coronal, and sagittal planes, #G (N) denotes the mean Gaussian primitive count in thousands, Size is the parameter storage in MB, and GEP follows Eq. (16).

	75-view					50-view					25-view				
Method	PSNR	SSIM	#G (N)	Size	GEP	PSNR	SSIM	#G (N)	Size	GEP	PSNR	SSIM	#G (N)	Size	GEP
Baseline	<u>39.05</u>	<u>0.946</u>	69.00k	2.90 MB	1.000	**38.18**	<u>0.939</u>	67.60k	2.84 MB	1.000	<u>35.40</u>	<u>0.912</u>	66.63k	2.80 MB	1.000
Budget-controlled simplification															
NanoGS	38.39	0.926	<u>36.15k</u>	1.52 MB	0.612	37.68	0.921	**35.74k**	1.50 MB	0.596	35.10	0.896	35.41k	1.49 MB	0.573
LightGaussian	38.73	0.940	36.16k	1.52 MB	<u>0.566</u>	37.98	0.935	<u>35.76k</u>	1.50 MB	<u>0.558</u>	35.37	0.910	35.42k	1.49 MB	<u>0.540</u>
GHAP	37.92	0.939	<u>36.15k</u>	1.52 MB	0.684	37.38	0.934	35.77k	1.50 MB	0.643	35.11	0.909	**35.37k**	1.49 MB	0.573
PUP 3D-GS	38.61	0.932	**36.11k**	1.52 MB	0.580	37.88	0.926	35.77k	1.50 MB	0.570	35.26	0.901	**35.37k**	1.49 MB	0.552
Adaptive density control															
GR-Gaussian	37.15	0.936	57.94k	2.43 MB	1.363	36.57	0.928	57.44k	2.41 MB	1.286	34.48	0.900	57.53k	2.41 MB	1.097
LeGS	35.85	0.890	81.92k	3.44 MB	2.554	34.94	0.879	80.89k	3.39 MB	3.156	30.35	0.808	140.69k	5.90 MB	32.408
FastGS	38.52	0.928	41.25k	1.73 MB	0.686	37.84	0.924	40.93k	1.72 MB	0.661	35.18	0.898	39.94k	1.68 MB	0.635
Perceptual-GS	**39.09**	**0.947**	78.93k	3.31 MB	1.138	<u>38.17</u>	**0.940**	77.41k	3.25 MB	1.157	35.31	0.910	76.09k	3.19 MB	1.175
CDGS (Ours)	38.93	0.944	<u>36.15k</u>	**1.52 MB**	**0.540**	38.15	<u>0.939</u>	**35.74k**	**1.50 MB**	**0.536**	**35.46**	**0.913**	<u>35.41k</u>	**1.49 MB**	**0.529**

Note: **Bold** indicates the best result, and <u>underline</u> indicates the second-best result in each column across all evaluated methods. The lowest GEP in each view configuration is highlighted in red.

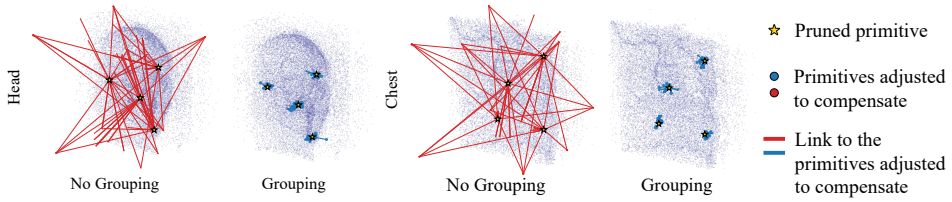

Figure 5: Effect of structural grouping. Without grouping, distant primitives along intersecting rays falsely absorb residuals, whereas CDGS confines compensation within anatomical boundaries.

Table 2: Ablation results of CDGS on the R^2-Gaussian backbone with and w/o structural grouping (Gr) and convex-cost decisions (Cvx). #G (N) denotes the mean Gaussian primitive counts.

			75-view			50-view			25-view		
Variant	Gr	Cvx	PSNR	#G (N)	GEP	PSNR	#G (N)	GEP	PSNR	#G (N)	GEP
w/o Grouping	×	✓	<u>38.78</u>	36.15k	<u>0.560</u>	<u>38.05</u>	35.74k	<u>0.549</u>	<u>35.39</u>	35.41k	<u>0.538</u>
w/o Convex Cost	✓	×	38.33	36.15k	0.622	37.68	35.74k	0.598	35.20	35.41k	0.561
w/o Both	×	×	38.32	36.15k	0.623	37.67	35.74k	0.598	35.21	35.41k	0.560
Full (CDGS)	✓	✓	**38.93**	36.15k	**0.540**	**38.15**	35.74k	**0.536**	**35.46**	35.41k	**0.529**

Note: **Bold** indicates the best result, and <u>underline</u> indicates the second-best result in each setting. The lowest GEP is highlighted in red.

5.3 Ablation Study: Component Analysis

We evaluated the effects of structural grouping and convex-cost decisions using the three variants in Table 2. The corresponding FaCT-GS ablations are reported in Table A7 in Appendix A.5.1. *w/o Grouping* removes structural grouping and performs cost-based decisions globally. It lowers PSNR and yields higher GEP across all settings on both backbones, while Figure 5 shows that grouping keeps deletion and compensation structurally localized. *w/o Convex Cost* retains grouping but replaces cost-based prioritization with random in-group selection, with primitive counts matched to CDGS at each event. CDGS improves PSNR by up to 0.60 dB on R^2-Gaussian at the same representation size. *w/o Both* removes grouping and uses random global selection with the same event-wise count targets. CDGS achieves higher PSNR than both random variants across all settings.

5.4 Parameter Sensitivity Analysis

We examined the sketch dimension P, the group stopping threshold λ, the batch acceptance threshold M, and the number of simplification events E, together with comparisons at fixed Gaussian counts. Figure A1 in Appendix A.3.2 shows the joint threshold analysis: colors indicate mean GEP, and cell values report the mean PSNR difference relative to the default CDGS configuration on the three most affected objects. More permissive settings generally lower GEP but reduce PSNR on these objects. We therefore adopted the shared default $\lambda = 1.2$ and $M = 4$. Complete sketch results, including training time and GPU memory, are reported in Appendix A.3.3. Detailed parameter settings, quantitative results, and fixed-count comparisons are provided in Appendix A.3.

6 Conclusion and Limitations

We presented CDGS, a group-wise simplification framework designed for seamless integration into radiative Gaussian CT reconstruction pipelines. Extensive experiments across two backbones and three sparse-view configurations showed that CDGS eliminates up to 47.6% of Gaussian primitives with a minimal fidelity loss of 0.03 dB, establishing a new state of the art in CT rate-distortion performance while preserving delicate structures.

Limitations and Future Work: While CDGS successfully resolves primitive redundancy under fixed trial geometry, candidate action discovery relies on local spatial graphs and moment-matched initialization. Extending convex density optimization to joint geometric refinement, and adapting topological grouping to non-attenuation volumetric imaging modalities (e.g., cryo-electron microscopy and seismic travel-time tomography), represents a promising direction for future research.

AI USE STATEMENT

In this work, we used generative AI tools for assisting to develop theoretical models or conceptual frameworks, to formulate mathematical claims, to provide critical ingredients for proving mathematical claims, to brush up the writing of proofs, to refine hypotheses, to provide feedback on research methodology and experiments, to implement methods, and to assist with translation. We have not used generative AI tools for generating synthetic data sets, cleaning and reformating dataset, supporting qualitative and thematic data analysis, and interpreting results. Additionally, we used generative AI tools for supporting to create and modify scientific figures or images, to create and edit software code, to improve to draft parts of a research paper, to summarize and analyse existing literature, to deeply identify research gaps, to conduct brainstorming, sourcing/searching for information, to edit a research paper to improve readability, to identify relevant literature, and to propose a title or keywords for a research paper. We have reviewed all AI-assisted work; we checked LLM-generated research ideas for potential plagiarism through a manual literature survey, and LLM-generated contents including code was verified and tested for correctness by 2 authors. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REFERENCES

A. Bhattacharyya. On a measure of divergence between two multinomial populations. *Sankhyā: The Indian Journal of Statistics (1933-1960)*, 7(4):401–406, 1946. ISSN 00364452. URL http://www.jstor.org/stable/25047882.

Ander Biguri, Manjit Dosanjh, Steven Hancock, and Manuchehr Soleimani. Tigre: a matlab-gpu toolbox for cbct image reconstruction. *Biomedical Physics & Engineering Express*, 2(5): 055010, sep 2016. doi: 10.1088/2057-1976/2/5/055010. URL https://doi.org/10.1088/2057-1976/2/5/055010.

Tatiana A Bubba, Gitta Kutyniok, Matti Lassas, Maximilian März, Wojciech Samek, Samuli Siltanen, and Vignesh Srinivasan. Learning the invisible: a hybrid deep learning-shearlet framework for limited angle computed tomography. *Inverse Problems*, 35(6):064002, 2019. ISSN 0266-5611. doi: 10.1088/1361-6420/ab10ca.

Yuanhao Cai, Yixun Liang, Jiahao Wang, Angtian Wang, Yulun Zhang, Xiaokang Yang, Zongwei Zhou, and Alan L. Yuille. Radiative gaussian splatting for efficient x-ray novel view synthesis. In Ales Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol (eds.), *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part I*, volume 15059 of *Lecture Notes in Computer Science*, pp. 283–299. Springer, 2024. doi: 10.1007/978-3-031-73232-4_16. URL https://doi.org/10.1007/978-3-031-73232-4_16.

Zhiwen Fan, Kevin Wang, Kairun Wen, Zehao Zhu, Dejia Xu, and Zhangyang Wang. Lightgaussian: Unbounded 3d gaussian compression with 15x reduction and 200+ FPS. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Advances in Neural Information Processing Systems 37: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/fd881d3b625437354d4421818f81058f-Abstract-Conference.html.

L. A. Feldkamp, L. C. Davis, and J. W. Kress. Practical cone-beam algorithm. *J. Opt. Soc. Am. A*, 1 (6):612–619, Jun 1984. doi: 10.1364/JOSAA.1.000612. URL https://opg.optica.org/josaa/abstract.cfm?URI=josaa-1-6-612.

Alex Hanson, Allen Tu, Vasu Singla, Mayuka Jayawardhana, Matthias Zwicker, and Tom Goldstein. PUP 3d-gs: Principled uncertainty pruning for 3d gaussian splatting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025*, pp. 5949–5958. Computer Vision Foundation / IEEE, 2025. doi: 10.1109/CVPR52734.2025.00558. URL https://openaccess.thecvf.com/content/CVPR2025/html/Hanson_PUP_3D-GS_Principled_Uncertainty_Pruning_for_3D_Gaussian_Splatting_CVPR_2025_paper.html.

Avinash C. Kak and Malcolm Slaney. *Principles of computerized tomographic imaging*, volume 33 of *Classics in applied mathematics*. SIAM, 2001. ISBN 978-0-89871-494-4.

Emma Kamutta, Sofia Mäkinen, and Alexander Meaney. Cone-beam computed tomography dataset of a seashell, aug 2022. URL https://doi.org/10.5281/zenodo.6983008.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139:1–139:14, 2023. doi: 10.1145/3592433. URL https://doi.org/10.1145/3592433.

Joseph Kuo, Jason Granstedt, Umberto Villa, and Mark A. Anastasio. Computing a projection operator onto the null space of a linear imaging operator: Tutorial. *Journal of the Optical Society of America. A, Optics, image science, and vision*, 39(3):470–481, 2022. ISSN 1084-7529. doi: 10.1364/josaa.443443.

Charles L Lawson and Richard J Hanson. *Solving least squares problems*. SIAM, 1995.

David Liben-Nowell and Jon M. Kleinberg. The link-prediction problem for social networks. *J. Assoc. Inf. Sci. Technol.*, 58(7):1019–1031, 2007. doi: 10.1002/ASI.20591. URL https://doi.org/10.1002/asi.20591.

Alexander Meaney. Cone-beam computed tomography dataset of a pine cone, aug 2022a. URL https://doi.org/10.5281/zenodo.6985407.

Alexander Meaney. Cone-beam computed tomography dataset of a walnut, aug 2022b. URL https://doi.org/10.5281/zenodo.6986012.

Zhenhua Ning, Xin Li, Jun Yu, Guangming Lu, Yaowei Wang, and Wenjie Pei. Beyond heuristics: Learnable density control for 3d gaussian splatting. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=J3uM4v9vJv.

Pawel Tomasz Pieta, Rasmus Juul Pedersen, Sina Borgi, Jakob Sauer Jørgensen, Jens Wenzel Andreasen, and Vedrana Andersen Dahl. Fact-gs: Fast and scalable CT reconstruction with gaussian splatting. In Paolo Favaro, Zuzana Kukelova, Atsuto Maki, Anna Rohrbach, Konrad Schindler, and Federico Tombari (eds.), *Computer Vision - ECCV 2026 - 19th European Conference, Malmö, Sweden, September 8-12, 2026, Proceedings, Part XXXV*, volume 17035 of *Lecture Notes in Computer Science*, pp. 225–245. Springer, 2026. doi: 10.1007/978-3-032-37464-6_13. URL https://doi.org/10.1007/978-3-032-37464-6_13.

Shiwei Ren, Tianci Wen, Yongchun Fang, and Biao Lu. Fastgs: Training 3d gaussian splatting in 100 seconds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26094–26103, 2026.

A.R. Runnalls. Kullback-leibler approach to gaussian mixture reduction. *IEEE Transactions on Aerospace and Electronic Systems*, 43(3):989–999, 2007. doi: 10.1109/TAES.2007.4383588.

The Finnish Inverse Problems Society. X-ray tomographic datasets. URL https://fips.fi/category/open-datasets/x-ray-tomographic-datasets/, 2024.

V. Traag, L. Waltman, and Nees Jan van Eck. From louvain to leiden: guaranteeing well-connected communities. *Scientific Reports*, 9:5233, 03 2019. doi: 10.1038/s41598-019-41695-z.

Lieven Vandenberghe and Stephen Boyd. *Convex optimization*, volume 1. Cambridge university press Cambridge, 2004.

Tao Wang, Mengyu Li, Geduo Zeng, Cheng Meng, and Qiong Zhang. Gaussian herding across pens: An optimal transport perspective on global gaussian reduction for 3dgs. In Danielle Belgrave, Cheng Zhang, Laura N. Montoya, Hsuan-Tien Lin, Razvan Pascanu, Piotr Koniusz, Marzyeh Ghassemi, Nancy Chen, Iván Vladimir Meza Ruíz, and Arturo Loaiza-Bonilla (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2025, NeurIPS 2025, San Diego, CA, USA, December 2-7, 2025 / Mexico City, Mexico, November 30 - December 5, 2025*, 2025. URL http://papers.nips.cc/paper_files/paper/2025/hash/e79574cc3355e831cc276c845605ed72-Abstract-Conference.html.

Butian Xiong, Rong Liu, Tiantian Zhou, Meida Chen, Zhiwen Fan, and Andrew Feng. Nanogs: Training-free gaussian splat simplification. In Paolo Favaro, Zuzana Kukelova, Atsuto Maki, Anna Rohrbach, Konrad Schindler, and Federico Tombari (eds.), *Computer Vision - ECCV 2026 - 19th European Conference, Malmö, Sweden, September 8-12, 2026, Proceedings, Part LXI*, volume 17061 of *Lecture Notes in Computer Science*, pp. 530–548. Springer, 2026. doi: 10.1007/978-3-032-37095-2_29. URL https://doi.org/10.1007/978-3-032-37095-2_29.

Yikuang Yuluo, Kuan Shen, Yue Ma, Ao Wang, Tongtong Jin, Yixing Huang, and Fuquan Wang. Gr-gaussian: Graph-based radiative gaussian splatting for sparse-view tomographic reconstruction. *IEEE Trans. Computational Imaging*, 12:1241–1253, 2026. doi: 10.1109/TCI.2026.3701585. URL https://doi.org/10.1109/TCI.2026.3701585.

Lihi Zelnik-manor and Pietro Perona. Self-tuning spectral clustering. In L. Saul, Y. Weiss, and L. Bottou (eds.), *Advances in Neural Information Processing Systems*, volume 17. MIT Press, 2004. URL https://proceedings.neurips.cc/paper_files/paper/2004/file/40173ea48d9567f1f393b20c855bb40b-Paper.pdf.

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 Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Advances in Neural Information Processing Systems 37: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/4f550cb7b30b59553e50cd08a9dbf068-Abstract-Conference.html.

Hongbi Zhou and Zhangkai Ni. Perceptual-gs: Scene-adaptive perceptual densification for gaussian splatting. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025*, volume 267 of *Proceedings of Machine Learning Research*. PMLR / OpenReview.net, 2025. URL https://proceedings.mlr.press/v267/zhou25s.html.