SPLATTING WITHOUT DISCRETE DENSITY CONTROL: LIFTED GAUSSIAN-MIXTURE KERNEL FOR CT RECONSTRUCTION

Anonymous authors
Paper under double-blind review

ABSTRACT

3D Gaussian Splatting (3DGS) has become one of the key paradigms in 3D reconstruction, offering an efficient explicit representation method with real-time rendering capabilities. It has demonstrated significant potential in X-ray computed tomography (CT) reconstruction, but existing methods rely on a discrete Adaptive Density Control (ADC) mechanism based on hard-threshold cloning, splitting, and pruning decisions. The threshold parameters are coupled with hyperparameters such as the objective function design, leading to potential imbalances in the Gaussian distribution under different reconstruction conditions—particularly in ill-posed reconstruction—which require multiple manual adjustments and thus severely limit the performance ceiling of the reconstruction method. This paper utilizes a dimension-elevated Gaussian mixed-kernel affine transformation to convert the splitting and merging behaviors of Gaussian kernels into continuously differentiable motion trajectories, enabling their integration into the 3D reconstruction optimization chain without reliance on hyperparameter thresholds. Based on this mechanism, we have constructed a 3D CT reconstruction framework that allows to allocate a budget for the number of Gaussian kernels in reconstruction. Experiments show that, with the same number of kernels, our method outperforms the ADC-based 3DGS reconstruction method. Furthermore, when considering larger budgets of Gaussian kernels, ADC begins to show signs of fatigue or even a decline in performance, whereas our method maintains superior performance and continues to improve as the kernel budget increases.

1 INTRODUCTION

Computed tomography (CT) reconstructs a three-dimensional attenuation field from projections measured along different directions. Reducing the number of projections weakens the constraints on this field, making the representation and its optimization particularly important. Three-dimensional Gaussian splatting (3DGS) offers explicit, anisotropic kernels that can be optimized efficiently (Kerbl et al., 2023). For CT, various works such as R^2-Gaussian have been published to fit the integration of these kernels to CT data and provides differentiable projection and volume-query operators (Zha et al., 2024), enabling direct optimization against measurements.

A central question is the efficiency of Gaussian kernel representation. Adaptive density control (ADC) has been a most commonly used strategy to change the representation through cloning, splitting, and pruning when needed. The resulted final kernel count depends on gradient thresholds, initialization, and reconstruction objectives. It is impossible to assign a budget of Gaussian kernel number by specifying a threshold. Under incomplete data situation such as sparse view observations that's of great interest in the field, increasing the number of kernels leaves more freedom to fit structures that the measurements do not uniquely determine. A useful alternative must address both how existing kernels move and how their spatial allocation evolves within the available budget.

We introduce Lifted-GS, which represents a static CT field using Gaussian mixtures in three spatial dimensions and one auxiliary dimension. Evaluating these mixtures at a shared learnable auxiliary coordinate produces spatial Gaussians whose centers and attenuation amplitudes change together. We call this construction a *lifted Gaussian-mixture parameterization*; the auxiliary coordinate has no

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.

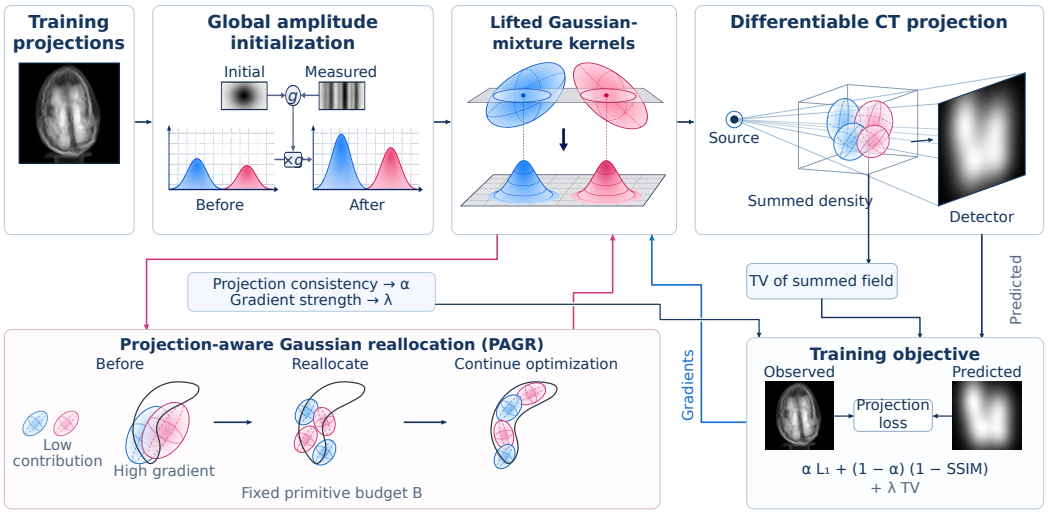

Figure 1: **Overview of Lifted-GS.** Training projections initialize the amplitude scale and determine the data-term weights. Slicing lifted Gaussian-mixture kernels produces a spatial attenuation field for CT projection. Projection fidelity and total-field TV guide optimization; PAGR updates the spatial allocation at fixed budget. The input thumbnail is an acquired Walnut projection. The Gaussian geometry, detector prediction, and refinement contour illustrate the method.

physical time interpretation. To improve spatial allocation, *projection-aware Gaussian reallocation* (PAGR) replaces low-contribution mixtures with refined copies of mixtures selected by position gradients. This operation preserves the kernel count, while using hard selection and parameter resets. A global gain sets the initial attenuation scale, and an adaptive objective combines projection fitting with total variation (TV) of the reconstructed field.

Our contributions are: (i) a Gaussian slicing parameterization for CT reconstruction with an exact spatial-kernel budget; (ii) a reallocation rule that combines integrated attenuation, a projection check, and position gradients; and (iii) a reconstruction procedure linking amplitude initialization, projection-consistency-based data weights, and gradient-based TV weighting.

2 Related Work

Gaussian representations for CT. CT integrates attenuation along rays, whereas appearance rendering accounts for occlusion. R^2-Gaussian adapts Gaussian splatting through a rectified X-ray rasterizer and differentiable voxelizer (Zha et al., 2024). Exact-GS develops a closed-form projection formulation for accurate Gaussian-based X-ray reconstruction (Yang et al., 2026). FaCT-GS improves the efficiency and scalability of rasterization and voxelization, including Gaussian fitting to existing volumes (Pieta et al., 2026). We retain the projection and volume-query interface of R^2-Gaussian and study representation optimization under a specified kernel count.

Gaussian allocation and budgets. Original 3DGS grows and prunes its representation based on Gaussian size, contribution, and position gradients (Kerbl et al., 2023). AbsGS addresses cancellation in positional gradients (Ye et al., 2024). Pixel-GS weights these gradients by pixel coverage to encourage growth in regions with insufficient initialization (Zhang et al., 2024), while TrimGS removes inaccurate or redundant geometry using contribution-based trimming (Fan et al., 2024). Taming 3DGS targets a user-specified final budget (Mallick et al., 2024); 3DGS-MCMC uses stochastic updates and Gaussian relocation (Kheradmand et al., 2024). PAGR evaluates allocation through the CT attenuation field and training projections, with the number of spatial kernels fixed throughout reconstruction.

Higher-dimensional Gaussian representations. Four-dimensional Gaussians have been used to represent dynamic scenes by deriving spatial Gaussians at each timestamp (Yang et al., 2024). 4D-Rotor Gaussian Splatting explicitly describes this operation as temporal slicing (Duan et al., 2024). We use the same Gaussian slicing identities with a learnable auxiliary coordinate. Here the fourth dimension organizes optimization of a static CT field, rather than indexing physical time.

Adaptive reconstruction objectives. Loss weighting can reflect both supervision reliability and optimization dynamics. RIGI uses disagreement between Gaussian models to assign uncertainty-dependent pixel weights in image-to-3D generation (Wang et al., 2024). GradNorm adjusts multi-task weights using gradient magnitudes and relative training progress (Chen et al., 2018). Our data weights instead use consistency diagnostics computed from CT training projections before reconstruction. During optimization, a separate feedback rule adjusts TV according to its gradient magnitude relative to the weighted projection objective.

3 PRELIMINARIES

3.1 CT IMAGING AND GAUSSIAN REPRESENTATION

Let $f : \mathbb{R}^3 \to \mathbb{R}_{\geq 0}$ be the linear attenuation field. Under an effective monoenergetic model, the log-transformed measurement at detector position $\mathbf{u}$ in view v is

$$y_v(\mathbf{u}) = -\log \frac{I_v(\mathbf{u})}{I_{0,v}(\mathbf{u})} = \underbrace{\int_{\ell_{\gamma_v,\mathbf{u}}} f(\mathbf{x})\,\mathrm{d}\ell}_{\mathcal{P}_{\gamma_v}[f](\mathbf{u})} + \epsilon_v(\mathbf{u}) + \delta_v(\mathbf{u}). \tag{1}$$

Here $I_{0,v}$ and I_v are incident and transmitted intensities, γ_v specifies the acquisition geometry, and $\ell_{\gamma_v,\mathbf{u}}$ is its ray. The terms ϵ_v and δ_v represent measurement noise and systematic model discrepancy, respectively. The latter may include geometry error; sparse angular sampling instead reduces the constraints on f even with accurate geometry.

We represent the field by B spatial Gaussians:

$$f_\theta(\mathbf{x}) = \sum_{j=1}^{B} a_j \exp\left[-\tfrac{1}{2}(\mathbf{x}-\mathbf{m}_j)^\top C_j^{-1}(\mathbf{x}-\mathbf{m}_j)\right]. \tag{2}$$

Each kernel has center $\mathbf{m}_j \in \mathbb{R}^3$, covariance $C_j \succ 0$, and peak attenuation $a_j \geq 0$. The eigenvectors and eigenvalues of C_j specify its orientation and squared spatial widths. Thus a_j has the units of f, whereas y_v is dimensionless. The differentiable CT renderer predicts $\widehat{y}_v = \mathcal{R}_{\gamma_v}(\theta) \simeq \mathcal{P}_{\gamma_v}[f_\theta]$ using the supplied geometry; θ denotes the learned representation parameters.

3.2 GAUSSIAN SLICING

Consider a normalized joint Gaussian $p(\mathbf{x}, b)$ with spatial mean $\boldsymbol{\mu}_x$, auxiliary mean μ_b, and covariance blocks Σ_{xx}, $\Sigma_{xb} = \Sigma_{bx}^\top$, and $\sigma_b^2 > 0$. Evaluating the auxiliary variable b at a scalar c gives

$$p(\mathbf{x}, c) = q(c)\,\mathcal{N}(\mathbf{x}; \mathbf{m}(c), C), \qquad q(c) = \mathcal{N}(c; \mu_b, \sigma_b^2), \tag{3}$$

$$\mathbf{m}(c) = \boldsymbol{\mu}_x + \frac{\Sigma_{xb}}{\sigma_b^2}(c - \mu_b), \qquad C = \Sigma_{xx} - \frac{\Sigma_{xb}\Sigma_{bx}}{\sigma_b^2}. \tag{4}$$

The factor $q(c)$ scales the slice, while $\mathbf{m}(c)$ and C are its conditional mean and covariance. Varying c moves the center through the spatial–auxiliary covariance and changes the scale factor; C stays fixed when the joint parameters are fixed. Applying this identity to each component of a mixture yields multiple spatial Gaussians. Section 4.1 converts these probability-density identities into an attenuation parameterization.

4 METHOD

Given projection data of a CT scan, data acquisition geometry, we propose our Lifted-GS reconstruction method to optimize lifted Gaussian-mixture kernels and their slicing coordinate with a

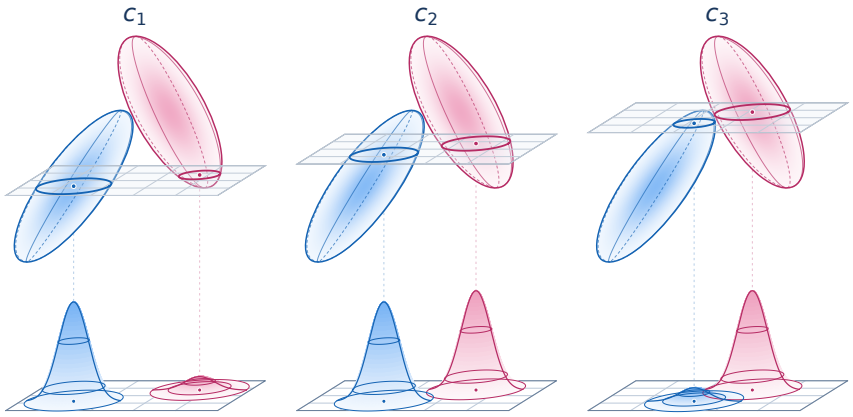

Figure 2: **Effect of the auxiliary slicing coordinate.** This 3D-to-2D schematic illustrates the 4D-to-3D construction. Moving the slicing plane changes the intersection centers and response heights. Conditional widths stay fixed when the Gaussian parameters are fixed. The auxiliary axis is an optimization coordinate for a static field.

spatial-kernel budget B. The global amplitude initialization strategy adjusts the amplitudes of the initialized Gaussian kernels by optimizing the initial volumetric projection values to approximate the overall magnitude of the actual projections. PAGR reallocates the spatial budget, and an adaptive objective controls the contributions of measurement fidelity and density-field regularity.

4.1 Lifted Gaussian-Mixture Parameterization

To allow local correlation control, We group spatial kernels into mixtures indexed by i, with K_i components satisfying $\sum_i K_i = B$. The auxiliary coordinate provides a shared way to adjust their centers and relative contributions without changing this count.

For component (i, k), $\mathbf{x}_{ik} \in \mathbb{R}^3$ is its reference center, b_{ik} its auxiliary mean, and $\sigma^b_{ik} > 0$ its auxiliary width. A dimensionless scalar c is shared and learned across mixtures. The slope $\mathbf{d}_{ik} \in \mathbb{R}^3$ specifies spatial displacement per unit change in c:

$$\mathbf{m}_{ik}(c) = \mathbf{x}_{ik} + \mathbf{d}_{ik}(c - b_{ik}), \qquad q_{ik}(c) = \mathcal{N}\big(c; b_{ik}, (\sigma^b_{ik})^2\big)\,. \tag{5}$$

Here $\mathbf{m}_{ik}$ is the spatial center and q_{ik} the auxiliary response at the selected slice. The positive-definite conditional covariance C_{ik} determines spatial extent and orientation. These quantities follow the conditioning construction in Section 3.2; Appendix A gives the corresponding joint covariance.

Let $A_{ik} > 0$ be a base attenuation amplitude and $\pi_{ik} \geq 0$ a learned mixture weight, with $\sum_k \pi_{ik} = 1$. We convert the auxiliary response to a peak attenuation through

$$\bar{q}_{ik}(c) = \frac{q_{ik}(c)}{\max_{1 \leq j \leq K_i} q_{ij}(c)}, \qquad a_{ik}(c) = A_{ik}\pi_{ik}\bar{q}_{ik}(c). \tag{6}$$

Maximum normalization retains relative slice responses while removing their common attenuation as the slice moves away from a mixture. The $\bar{q}_{ik}$ need not sum to one: the mixture weights allocate relative contributions, whereas the base amplitudes set their attenuation scale. Substituting $(\mathbf{m}_{ik}, C_{ik}, a_{ik})$ into equation 2 defines the reconstructed field, whose spatial kernels have unit peak before amplitude weighting. The resulting attenuation field is not a probability density.

Changing c couples center displacement with relative amplitude (Figure 2). The individual centers, slopes, widths, and amplitudes also remain learnable, allowing local adjustments alongside this shared change. The auxiliary dimension thus changes the optimization parameterization of a static three-dimensional field. Appendix A specifies the numerical stabilization and parameter constraints.

4.2 Projection-Based Amplitude Initialization

Spatial initialization provides coverage of the reconstruction domain but does not determine the measured attenuation scale. We separate these roles by fitting a common amplitude gain while holding the initial kernel geometry fixed. Let $p_v^{(0)} = \mathcal{R}_{\gamma_v}(\theta_0)$ denote the initial projection. Linearity of the projection model in attenuation gives

$$g^* = \arg\min_{g \geq 0} \sum_v \|g p_v^{(0)} - y_v\|_2^2 = \max\left(0, \frac{\sum_v \langle p_v^{(0)}, y_v \rangle}{\sum_v \|p_v^{(0)}\|_2^2}\right), \tag{7}$$

where the sum uses training views and each inner product sums detector pixels. For nonzero initial predictions, this gain aligns their overall scale with the measurements. Applying $A_{ik} \leftarrow g^* A_{ik}$ preserves relative kernel contributions, giving subsequent optimization a calibrated starting scale. The spatial initialization and stabilized gain computation are specified in Appendix B.

4.3 Projection-Aware Gaussian Reallocation

Continuous parameter updates can improve existing kernels while leaving some of the fixed budget poorly allocated. PAGR redirects low-contribution mixtures toward regions with sustained sensitivity to position changes. It first identifies reusable mixtures, checks their projection contribution, and then refines selected targets.

Identify and check reusable mixtures. The spatial integral of mixture i is

$$M_i = \sum_k a_{ik}(c)(2\pi)^{3/2}\sqrt{\det C_{ik}}. \tag{8}$$

This integrated attenuation accounts for both amplitude and spatial extent. We rank mixtures by increasing M_i and select a set $\mathcal{D}$ whose cumulative contribution is at most a fraction ζ of the total. Low integrated attenuation alone does not ensure a small effect on measured rays. We therefore temporarily suppress $\mathcal{D}$ and compare its prediction $\widehat{y}_v^{(-\mathcal{D})}$ with the current prediction:

$$E_{\mathcal{D}} = \max_{v \in \mathcal{V}_{\text{check}}} \frac{\|\widehat{y}_v^{(-\mathcal{D})} - \widehat{y}_v\|_1}{\|\widehat{y}_v\|_1} \leq \tau_D. \tag{9}$$

Here $\mathcal{V}_{\text{check}}$ is a set of training views and τ_D the allowed relative projection change. This check connects a representation-level contribution measure to the observed imaging geometry. If suppression changes projections too much, the candidate set is reduced; if no admissible set remains, reallocation is skipped.

Select targets and refine. Among the remaining mixtures, we prioritize those with large accumulated position-gradient norms:

$$S_i = \sum_{t \in \mathcal{W}_r} \sum_k \|\mathbf{g}_{ik}^{(t)}\|_2, \tag{10}$$

where $\mathcal{W}_r$ is a recent optimization window and $\mathbf{g}_{ik}^{(t)}$ is the position gradient associated with reference center $\mathbf{x}_{ik}$ under the full reconstruction objective. Accumulation favors sustained sensitivity over a single update. Each reusable mixture $u \in \mathcal{D}$ is paired with a distinct high-scoring target h having the same component count. We redistribute the target's contribution over two finer spatial configurations:

$$\mathbf{m}_{hk}^+, \mathbf{m}_{uk}^+ \sim \mathcal{N}(\mathbf{m}_{hk}, C_{hk}), \qquad C_{hk}^+ = C_{uk}^+ = C_{hk}/\beta^2, \qquad A_{hk}^+ = A_{uk}^+ = A_{hk}/2, \tag{11}$$

where $\beta > 1$ controls spatial refinement and $+$ denotes the state after reallocation. The two centers are drawn independently from the target's spatial support. Narrower kernels provide finer local coverage; both copies inherit the target's mixture weights and auxiliary response parameters, with its base amplitude divided equally.

Both configurations occupy existing mixtures, so the kernel count remains B. Reallocation is a discrete change between continuous fitting steps and need not preserve the attenuation field. The projection check bounds candidate removal only; optimization subsequently refits the field after the full refinement event. Appendix C specifies the selection and parameter updates.

4.4 Adaptive Reconstruction Objective

We fit the projections while regularizing spatial variation in the reconstructed field:

$$\mathcal{L}_t = \underbrace{\alpha \frac{\|\widehat{y}_{v_t} - y_{v_t}\|_1}{n_{v_t}} + (1-\alpha)\big[1 - \mathrm{SSIM}(\widehat{y}_{v_t}, y_{v_t})\big]}_{\mathcal{L}_{\mathrm{data,v_t}}} + \lambda_t \ \mathrm{TV}(f_\theta|_{\Omega_t}). \tag{12}$$

Here v_t is a training view with n_{v_t} detector pixels and Ω_t a spatial region of the reconstructed volume. L1 measures pointwise projection error, SSIM compares local structure (Wang et al., 2004), and TV penalizes spatial variation of the summed attenuation field. The coefficients α balance data-term for each view, whereas λ_t adapts regularization strength as fitting progresses.

Set data weights from projection consistency. A reconstruction from the training views provides a reference for assessing structured projection residuals. We reproject it under the nominal geometry, align its overall gain, and compare smoothed predictions $p_{v,\sigma}$ with smoothed measurements $y_{v,\sigma}$, where σ is the smoothing scale. A small detector-plane translation $\widehat{\mathbf{\Delta}}$ is fitted on part of each image and evaluated on a disjoint region $\mathcal{E}$:

$$A_{v,\sigma} = \frac{\|p_{v,\sigma} - y_{v,\sigma}\|_{1,\mathcal{E}} - \|\mathcal{T}_{\widehat{\mathbf{\Delta}}} p_{v,\sigma} - y_{v,\sigma}\|_{1,\mathcal{E}}}{\|y_{v,\sigma}\|_{1,\mathcal{E}}}. \tag{13}$$

The operator $\mathcal{T}$ translates an image. Positive improvement indicates a residual component explainable by a small shift, but reconstruction artifacts and noise can also produce such improvement. We therefore compare this diagnostic with controls generated under the same geometry, including matched noise perturbations. A conservative lower estimate of the observed-minus-control improvement defines the excess score e. Stability across smoothing scales and image subsets is required before it affects the loss.

For an accepted score, we set

$$u = \mathrm{clip}\left(\frac{e-\ell}{h-\ell}, 0, 1\right), \qquad \alpha = \alpha_{\min} + (\alpha_{\max} - \alpha_{\min})u, \tag{14}$$

where $\ell < h$ are excess-score anchors calibrated using unshifted and deliberately perturbed controls, and $\alpha_{\min}, \alpha_{\max}$ bound the L1 share. If the stability condition fails, $u = 0$. Larger stable excess increases the L1 share under this empirical calibration rule. The diagnostic chooses a loss weight; it neither corrects the acquisition geometry nor estimates a physical calibration error. Appendix D defines the reference reconstruction, controls, and score estimation.

Adapt TV using gradient magnitudes. The relative influence of data fitting and TV changes as the field evolves, even when their scalar weights stay fixed. We assess this influence on a common block of attenuation-amplitude parameters, denoted $\boldsymbol{\rho}$. Let G_{data} and G_{TV} be the root mean square norms of their respective gradients across training views and spatial regions. Comparing gradients on the same parameter block makes their magnitudes commensurate. The desired regularization weight is

$$\lambda^* = \mathrm{clip}(G_{\mathrm{data}}/G_{\mathrm{TV}}, \lambda_{\min}, \lambda_{\max}), \tag{15}$$

where the bounds limit the range of regularization strengths. A larger data gradient permits stronger TV, while a larger TV gradient reduces its weight. We move λ_t gradually toward this target to avoid abrupt changes in the objective. This feedback balances the two influences without imposing exact gradient equality. Appendix D gives the gradient statistics and update rule.

The reconstruction thus proceeds from amplitude alignment and scan-specific data weighting to joint optimization of the Gaussian parameters and shared coordinate. PAGR redistributes the fixed budget during fitting, after which optimization continues to reconcile the refined field with the projections. Full pseudocode and optimization settings are provided in Appendix B; discrete loss definitions are given in Appendix E.

5 EXPERIMENTS

Training and metrics. Both methods use a fixed 30,000 iterations. The original R^2-Gaussian loss and optimization settings are retained. For Lifted-GS, amplitude initialization, PAGR, and the data-weight and TV-feedback rules follow Section 4. PSNR and average SSIM of two-dimensional slice along all three axes are evaluted.

Data and acquisition. The evaluation protocol covers the acquired FIPS scans Pine, Seashell, and Walnut, and four clinical-source CT volumes: Colon, Dental, Head-and-neck, and Pancreas. Projections of the clinical-source volumes are simulated using cone-beam geometries. We distinguish these measurements from the physically acquired FIPS projections. FIPS conducted sparse-view experiments using 25, 50, and 75 views. For clinical data, experiments were performed using 25, 50, 75, and 180 views.

Table 1: Comparison with matching gaussian kernels based on ADC defaults

Volume	Acquisition	PSNR (dB)		SSIM	
		Ours	R^2	Ours	R^2
colon	25 views	34.41	**34.53**	**0.8938**	0.8925
	50 views	**38.23**	38.11	**0.9375**	0.9365
	75 views	**38.99**	38.86	**0.9437**	0.9436
	180 views	**39.46**	39.38	**0.9461**	0.9460
dental	25 views	34.59	**34.65**	0.9311	**0.9313**
	50 views	**37.30**	37.13	**0.9538**	0.9522
	75 views	**38.04**	37.87	**0.9598**	0.9582
	180 views	**38.95**	38.88	**0.9639**	0.9625
head-and-neck	25 views	**34.23**	34.22	0.9247	**0.9257**
	50 views	**39.56**	39.49	**0.9684**	0.9683
	75 views	**40.42**	40.20	**0.9752**	0.9747
	180 views	**42.10**	41.89	**0.9802**	0.9795
pancreas	25 views	**33.75**	33.55	**0.9257**	0.9221
	50 views	**37.82**	37.46	**0.9634**	0.9612
	75 views	**39.30**	38.84	**0.9717**	0.9699
	180 views	**40.92**	40.31	**0.9769**	0.9754
walnut	25 views	**29.40**	29.01	**0.6707**	0.6660
	50 views	**32.65**	32.19	**0.7143**	0.7072
	75 views	**34.04**	33.89	0.7239	**0.7285**
seashell	25 views	**39.95**	39.72	**0.9411**	0.9406
	50 views	43.40	**43.62**	0.9608	**0.9626**
	75 views	43.74	**43.99**	0.9627	**0.9663**
pine	25 views	**38.90**	38.50	**0.9316**	0.9277
	50 views	**41.54**	41.42	**0.9480**	0.9461
	75 views	41.90	**42.22**	0.9490	**0.9491**

Table 1 presents the 3D PSNR and SSIM for our method and R^2 Gaussian, compared using an equivalent number of Gaussians; the best values are highlighted in bold. The results demonstrate that, when using the same number of Gaussian kernels and default parameters, our method mostly outperforms R^2 Gaussian in terms of both PSNR and SSIM.

Primitive-budget comparisons. We conducted experiments on the Gaussian kernel count budget using clinical data for the settings of 25, 50, and 75 projections. The comparison matches the final kernel count of R^2-Gaussian using its default ADC threshold, 5×10^{-5}. Quality–budget curves use 180 projections and ADC thresholds $\{2 \times 10^{-4}, 10^{-4}, 5 \times 10^{-5}, 2 \times 10^{-5}, 10^{-5}\}$. For each

threshold, the final R^2-Gaussian count is used to assign the budget of Lifted-GS. The initial point pool is expanded whenever this budget requires more than 50,000 initial centers.

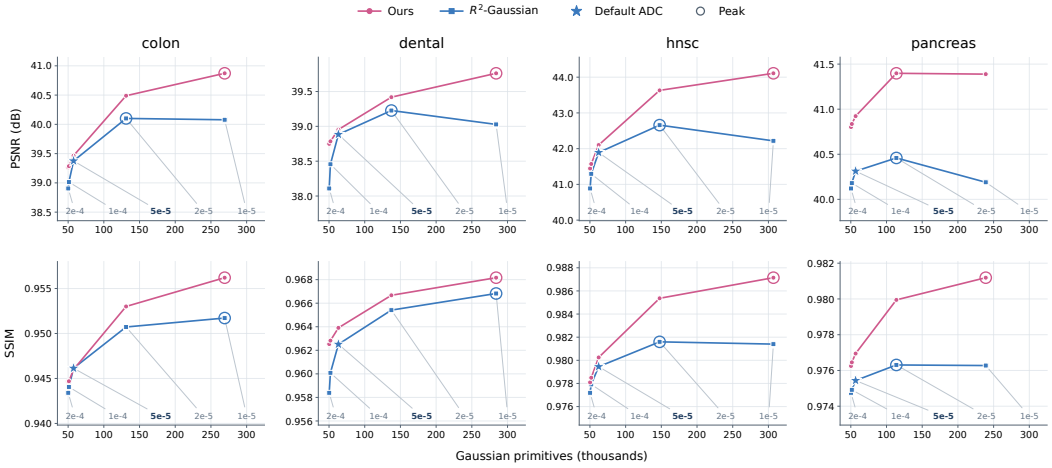

Figure 3: **The relationship between the Gaussian kernel count budget and reconstructed image quality.** The trend of reconstruction quality as the Gaussian kernel budget increases are illustrated across four clinical datasets. We control the Gaussian budget by adjusting the ADC threshold in R^2 Gaussian and use a matching number of Gaussian kernels for our reconstruction experiments. The circles represent the peak values for each method, while the pentagrams indicate the default parameters. The results show that our method consistently improves reconstruction quality as the number of Gaussian kernels increases, whereas the ADC-based R^2 Gaussian method hits a performance bottleneck or even degrades as the Gaussian kernel budget grows.

Figure 3 illustrates the reconstruction performance of four datasets using our method versus the R^2 Gaussian method across various Gaussian kernel budget levels. Threshold labels identify ADC curve points; a star denotes the default threshold and a circle denotes the maximum of each reported metric. While our method holds a slight advantage under default settings, this lead widens as the Gaussian kernel budget increases. The experiments demonstrate that our method effectively supports reconstruction optimization with larger Gaussian kernel budgets, offering a higher performance ceiling. Notably, even within the low-budget range, our method exhibits a more pronounced advantage compared to the R^2 Gaussian method than it does under default configurations; this suggests that our method may achieve superior Gaussian kernel utilization in low-budget scenarios compared to ADC-based approaches. Figure 4 presents a comparison of our method with R^2 Gaussian (showing both default and best reconstruction results) and with the reconstruction results from FDK(Feldkamp et al., 1984), OS-SART(Andersen & Kak, 1984), ASD-POCS(Sidky & Pan, 2008), NAF(Zha et al., 2022), and TensoRF(Chen et al., 2022); additional results are provided in the Appendix F.

Table 2: Comparison of the best image quality within the experimental range

Volume	PSNR3D (dB)		SSIM3D	
	Ours best	R^2 best	Ours best	R^2 best
colon	**40.87**	40.10	**0.9562**	0.9507
dental	**39.76**	39.23	**0.9682**	0.9654
head-and-neck	**44.11**	42.66	**0.9871**	0.9816
pancreas	**41.39**	40.46	**0.9812**	0.9763

Ablation study. We conducted ablation studies using the pancreas synthetic phantom. As shown in the table, the representation method of Lifted-GS inherently possesses certain advantages, and significant performance improvements were achieved through the combination of the amplitude initialization, PAGR and Adaptive reconstruction objective methods.

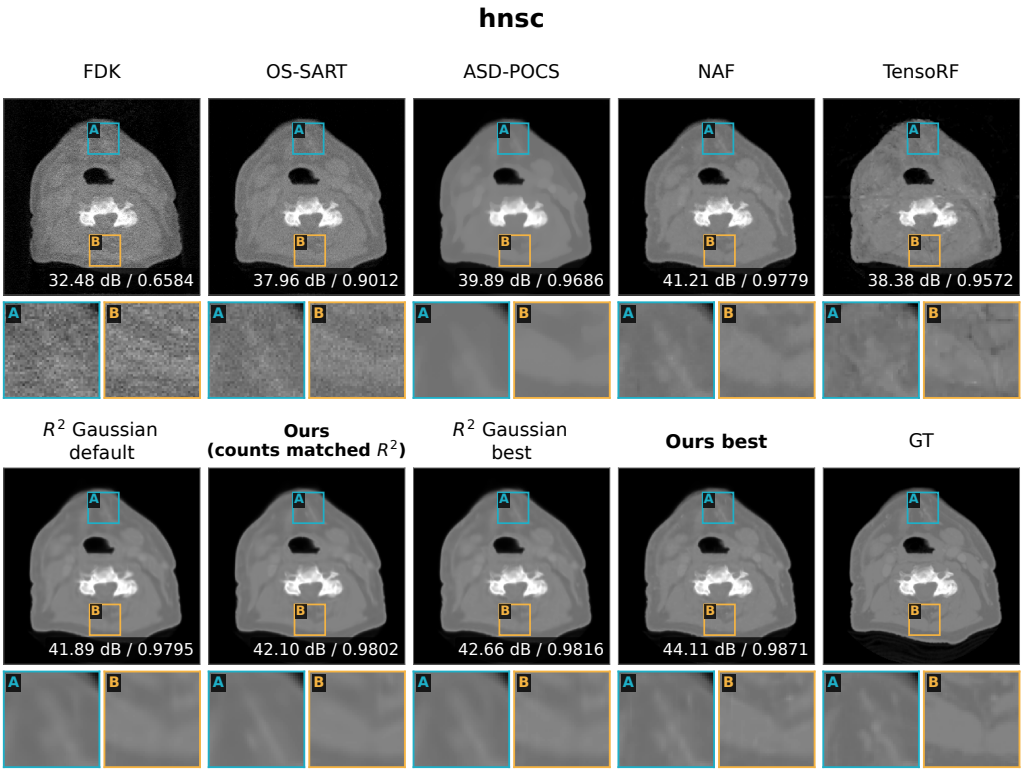

Figure 4: **Example of volumetric data slice reconstruction from 180 angles of the other methods compared with our method.** The metrics PSNR3D, SSIM3D are indicated. Compared to other methods, our approach yields clearer anatomical structures, whether the number of kernels is equal to or greater than the R² Gaussian.

Table 3: Ablation studies of individual modules

Methods	4.2	4.3	4.4	PSNR3D/SSIM3D	Δ
R^2	-	-	-	36.23/0.9548	-
Lifted-GS kernel	-	-	-	37.01/0.9540	+0.78/−0.0008
Lifted-GS kernel	✓	-	-	37.18/0.9549	+0.95/+0.0001
Lifted-GS kernel	✓	✓	-	37.26/0.9578	+1.03/+0.0030
Lifted-GS kernel	✓	✓	✓	**37.39/0.9600**	+1.16/+0.0052

6 Conclusion

Lifted-GS reconstructs a static CT attenuation field with a prescribed number of spatial Gaussians. A lifted mixture parameterization couples center displacement and attenuation amplitude, while projection-aware reallocation redistributes existing kernels. Amplitude initialization and adaptive loss weights connect the representation to projection scale, consistency, and regularization strength. Experiments demonstrate that our method outperforms other comparative reconstruction methods in most cases; furthermore, as the budget for the number of Gaussian kernels increases, our method effectively utilizes them to enhance reconstruction quality. The fixed count enables controlled comparisons across representation budgets.

AI USE STATEMENT

In this work, we used generative AI tools to draft some text. And we use LLMs to assist in implementing the code for our experiments. We have reviewed all AI-assisted work. LLM-generated 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.

REPRODUCIBILITY STATEMENT

We will upload the open-source GitHub projects associated with this paper, enabling direct reproduction of results with performance comparable to that reported in the paper.

REFERENCES

Anders H Andersen and Avinash C Kak. Simultaneous algebraic reconstruction technique (sart): a superior implementation of the art algorithm. *Ultrasonic imaging*, 6(1):81–94, 1984.

Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European conference on computer vision*, pp. 333–350. Springer, 2022.

Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. GradNorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pp. 794–803. PMLR, 2018. URL https://proceedings.mlr.press/v80/chen18a.html.

Yuanxing Duan, Fangyin Wei, Qiyu Dai, Yuhang He, Wenzheng Chen, and Baoquan Chen. 4D-Rotor Gaussian Splatting: Towards efficient novel view synthesis for dynamic scenes. In *ACM SIGGRAPH 2024 Conference Papers*, 2024. URL https://arxiv.org/abs/2402.03307.

Lue Fan, Yuxue Yang, Minxing Li, Hongsheng Li, and Zhaoxiang Zhang. Trim 3D Gaussian Splatting for accurate geometry representation. arXiv:2406.07499, 2024. URL https://trimgs.github.io/.

Lee A Feldkamp, Lloyd C Davis, and James W Kress. Practical cone-beam algorithm. *Journal of the Optical Society of America A*, 1(6):612–619, 1984.

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. URL https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/.

Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Weiwei Sun, Yang-Che 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*, volume 37, 2024. URL https://ubc-vision.github.io/3dgs-mcmc/.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015. URL https://arxiv.org/abs/1412.6980.

Saswat Subhajyoti Mallick, Rahul Goel, Bernhard Kerbl, Markus Steinberger, Francisco Vicente Carrasco, and Fernando De La Torre. Taming 3DGS: High-quality radiance fields with limited resources. In *SIGGRAPH Asia 2024 Conference Papers*. Association for Computing Machinery, 2024. doi: 10.1145/3680528.3687694. URL https://humansensinglab.github.io/taming-3dgs/.

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. arXiv:2604.01844, 2026. URL https://arxiv.org/abs/2604.01844. Accepted at ECCV 2026.

Emil Y Sidky and Xiaochuan Pan. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in Medicine & Biology*, 53(17):4777–4807, 2008.

Jiacheng Wang, Zhedong Zheng, Wei Xu, and Ping Liu. RIGI: Rectifying image-to-3d generation inconsistency via uncertainty-aware learning. arXiv:2411.18866, 2024. URL https://arxiv.org/abs/2411.18866.

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. doi: 10.1109/TIP.2003.819861. URL https://ece.uwaterloo.ca/~z70wang/research/ssim/.

Guangpu Yang, Steffen Kieß, Hanxiang Luo, Xingyu Liu, and Sven Simon. Exact-GS: Mathematically rigorous and accurate 3D Gaussian Splatting for 3D X-ray reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4902–4911, 2026. URL https://openaccess.thecvf.com/content/CVPR2026/html/Yang_Exact-GS_Mathematically_Rigorous_and_Accurate_3D_Gaussian_Splatting_for_3D_CVPR_2026_paper.html.

Zeyu Yang, Hongye Yang, Zijie Pan, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4D Gaussian Splatting. In *International Conference on Learning Representations*, 2024. URL https://fudan-zvg.github.io/4d-gaussian-splatting/.

Zongxin Ye, Wenyu Li, Sidun Liu, Peng Qiao, and Yong Dou. AbsGS: Recovering fine details for 3D Gaussian Splatting. arXiv:2404.10484, 2024. URL https://arxiv.org/abs/2404.10484.

Ruyi Zha, Yanhao Zhang, and Hongdong Li. Naf: neural attenuation fields for sparse-view cbct reconstruction. In *International conference on medical image computing and computer-assisted intervention*, pp. 442–452. Springer, 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*, volume 37, 2024. URL https://arxiv.org/abs/2405.20693.

Zheng Zhang, Wenbo Hu, Yixing Lao, Tong He, and Hengshuang Zhao. Pixel-GS: Density control with pixel-aware gradient for 3D Gaussian Splatting. In *European Conference on Computer Vision*, 2024. URL https://pixelgs.github.io/.

A GAUSSIAN-MIXTURE PARAMETERIZATION

For one four-dimensional Gaussian, suppress the mixture and Gaussian indices. The spatial covariance $C \succ 0$, auxiliary width $\sigma_b > 0$, and slope $\mathbf{d}$ define

$$\Sigma^{4D} = \begin{bmatrix} C + \sigma_b^2 \mathbf{d}\mathbf{d}^\top & \sigma_b^2 \mathbf{d} \\ \sigma_b^2 \mathbf{d}^\top & \sigma_b^2 \end{bmatrix} = \begin{bmatrix} I & \mathbf{d} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} C & 0 \\ 0 & \sigma_b^2 \end{bmatrix} \begin{bmatrix} I & 0 \\ \mathbf{d}^\top & 1 \end{bmatrix}. \tag{A.1}$$

The factorization proves positive definiteness. For joint mean $(\mathbf{x}, b)$, Gaussian conditioning gives

$$\mathbb{E}[\mathbf{X} \mid \widetilde{b} = c] = \mathbf{x} + \mathbf{d}(c - b), \qquad \mathrm{Cov}[\mathbf{X} \mid \widetilde{b} = c] = C. \tag{A.2}$$

The conditional covariance is independent of c when the other parameters are fixed, and remains learnable through the spatial scale and rotation parameters.

Relative slice response. The auxiliary response is $q_{ik}(c) = \mathcal{N}(c; b_{ik}, (\sigma_{ik}^b)^2)$. Within each mixture, we use

$$\bar{q}_{ik}(c) = \frac{q_{ik}(c)}{\max_j q_{ij}(c) + \varepsilon_w}, \qquad a_{ik}(c) = A_{ik}\pi_{ik}\bar{q}_{ik}(c). \tag{A.3}$$

For responses above the numerical floor, maximum normalization compensates for the common decay as the slicing coordinate moves away from the auxiliary means. Overall attenuation can still change through A_{ik}. This is a choice of amplitude parameterization for a CT attenuation field; the rendered kernels are unit-peak spatial Gaussians rather than normalized probability densities. The relative responses do not sum to one.

On an interval with a unique maximizing Gaussian j^*,

$$\frac{\partial \mathbf{m}_k}{\partial c} = \mathbf{d}_k, \qquad \frac{\partial \log q_k}{\partial c} = \frac{b_k - c}{(\sigma_k^b)^2}, \tag{A.4}$$

$$\frac{\partial \log \bar{q}_k}{\partial c} = \frac{b_k - c}{(\sigma_k^b)^2} - \frac{q_{j^*}}{q_{j^*} + \varepsilon_w} \frac{b_{j^*} - c}{(\sigma_{j^*}^b)^2}. \tag{A.5}$$

These identities assume positive auxiliary responses before floating-point underflow. The mapping is continuous at ties, with potentially different one-sided derivatives. The larger relative response need not produce the larger attenuation amplitude, which also depends on $A_k\pi_k$.

Positive parameters and exact counting. We parameterize $A_{ik} = \mathrm{softplus}(\rho_{ik})$ and $C_{ik} = R_{ik}\,\mathrm{diag}(\mathbf{s}_{ik}^2)R_{ik}^\top$. Spatial scales are obtained through a bounded sigmoid, and rotations through normalized four-component quaternions. One sigmoid logit determines the two mixture weights; a second stored logit is redundant. For odd B, the last mixture contains a single spatial Gaussian with mixture weight one. Its unused storage entry is excluded from slice normalization, projection, volume queries, and PAGR. The primitive budget counts all represented spatial Gaussians, including kernels whose amplitudes are small.

B INITIALIZATION AND OPTIMIZATION SETTINGS

For initial point $\mathbf{p}_i$, we set $\mathbf{x}_{i1} = \mathbf{p}_i$ and $\mathbf{x}_{i2} = \mathbf{p}_i + \boldsymbol{\xi}_i$, with $\boldsymbol{\xi}_i \sim \mathcal{N}(0, \sigma_\xi^2 I)$. Both Gaussians start with the same base amplitude and mixture weight, zero slope, and identity rotation. Choosing $c_0 = b_{i1}$ and auxiliary-mean separation Δ_b gives the second Gaussian the smaller initial relative response $\exp[-\Delta_b^2/(2(\sigma_0^b)^2)]$. Both participate in optimization from initialization onward. Spatial scales use the nearest-neighbor statistic of the initial point cloud.

Initial points are sampled from a training-only FDK reconstruction. The point pool contains at least $\lceil B/2 \rceil$ distinct points and is expanded when the requested budget exceeds the capacity of a 50,000-point pool. The amplitude gain is fitted to all training views and stored by inverse softplus before optimizer creation. Because contribution thresholds can affect renderer linearity, all training views are re-rendered after gain fitting; the initialized projection MSE must not exceed its pre-fit value beyond numerical tolerance.

We use Adam (Kingma & Ba, 2015) with moment coefficients $(0.9, 0.999)$, stabilizer 10^{-15}, global gradient clipping at 1, and parameter-group learning rates $\eta_p(t) = \eta_p(0)r^{t/T}$. Table B.1 specifies the common configuration.

Table B.1: Reconstruction configuration. L is the longest side of the reconstruction domain in normalized world coordinates; B is the prescribed spatial-kernel count.

Parameter or operation	Setting
Initial c_0, auxiliary means, auxiliary width	0, $(0, 1)$, 0.5
Initial base amplitude / spatial jitter	0.1 / 0.001 world units
Initial slopes / mixture logit	$\mathbf{0}$ / 0
Scale bounds / auxiliary-width floor	$[0.0005L, 0.5L]$ / 10^{-3}
Normalization and gain stabilizers	10^{-12}
Gain fitting / MSE acceptance ratio	All training views / $\leq 1 + 10^{-6}$
Training iterations / final learning-rate factor	30,000 / 0.033
Initial learning rate: spatial centers / raw scales	10^{-4} / 5×10^{-3}
Initial learning rate: density / mixture logit	10^{-2} / 10^{-2}
Initial learning rate: rotation, b, $\log \sigma^b$, $\mathbf{d}$, c	10^{-3}
L1 coefficient range / SSIM coefficient	$[0.5, 0.8]$ / $1 - \alpha$
Initial TV / TV bounds	0.05 / $[0.0125, 0.05]$
TV patch grid	32^3 voxels
PAGR iteration / gradient accumulation interval	10,000 / 500 steps
Contribution fraction / removal-check threshold	0.01 / 0.05
Removal-check views / maximum fraction halvings	8 / 3
Refinement scale factor / amplitude per copy	$1/1.6$ / $1/2$
Training random seed	0

C PAGR Parameter Updates

The score in equation 10 accumulates the norms of base-center gradients after global clipping and before each optimizer update. It includes all terms in the reconstruction objective. Candidate and target selections operate on complete mixtures; a final singleton is excluded. Each selected target is distinct and is refined once at the reallocation iteration.

To evaluate equation 9, candidate density parameters are temporarily set to a sufficiently negative value to make their softplus amplitudes negligible, then restored. A failed check halves ζ and repeats candidate selection. An empty candidate set, an unsuccessful check after the retry limit, or too few distinct targets leaves the representation unchanged.

For the centers sampled in equation 11, the reference-coordinate updates are

$$\mathbf{x}^+_{hk} = \mathbf{m}^+_{hk} - \mathbf{d}_{hk}(c - b_{hk}), \qquad \mathbf{d}^+_{uk} = \mathbf{0}, \qquad \mathbf{x}^+_{uk} = \mathbf{m}^+_{uk}. \tag{C.1}$$

The target retains its slope and the replacement mixture starts with zero slope. Both inherit the target's auxiliary means, widths, rotations, and mixture weights. Scales are divided by β, clipped inside $[0.0005L, 0.5L]$, and stored through the inverse activation. Base amplitudes are halved through inverse softplus. The shared slicing coordinate remains fixed during this operation and resumes gradient updates afterward.

For each mixture-indexed parameter, both Adam moment rows are copied from the target to the replacement. Target moments, untouched entries, and optimizer step counters are retained. Gradients of the affected entries are cleared. The two center samples are independent draws from the target's pre-update spatial covariance. Projection changes from the complete operation are recorded before optimization continues.

D Adaptive Weight Computation

D.1 Training-Projection Consistency

Let x_{FDK} be the reconstruction from the training views and $p_v = \mathcal{P}_{\gamma_v}[x_{\mathrm{FDK}}]$ its reprojection. A robust global gain aligns their scale with y_v; no per-view gain or additive offset is fitted. At smoothing scales $\sigma \in \{1, 2, 4\}$ detector pixels, a small translation is estimated with four iteratively reweighted least-squares steps on fitting tiles. Each coordinate is restricted to ± 4 pixels. On disjoint

Procedure C.1 Projection-aware Gaussian reallocation (PAGR)

Input: Parameters, optimizer state, scores S_i, contribution fraction ζ, check views, threshold τ_D, retry limit J, shrink factor β

Rank eligible mixtures by increasing M_i
for $j = 0, \dots, J$ **do**
  Select a prefix $\mathcal{D}$ with total contribution at most $\zeta \sum_i M_i$
  if $\mathcal{D} = \varnothing$ **then**
    return unchanged state
  end if
  Suppress $\mathcal{D}$, evaluate $E_{\mathcal{D}}$, and restore its amplitudes
  if $E_{\mathcal{D}} \le \tau_D$ **then**
    break
  end if
  $\zeta \leftarrow \zeta/2$
end for
if the check fails or too few targets are available **then**
  return unchanged state
end if
Select $|\mathcal{D}|$ distinct remaining mixtures with largest S_i
for each replacement–target pair (u, h) **do**
  Copy the target mixture; sample and refine the two spatial configurations
  Set reference centers to recover the sampled centers at the current c
  Transfer optimizer moments to the replacement and clear affected gradients
end for

Output: Updated representation with unchanged spatial-kernel count B

evaluation tiles, its signed improvement is

$$A_{v,\sigma} = \frac{\|p_{v,\sigma} - y_{v,\sigma}\|_{1,\mathcal{E}} - \|T_{\widehat{\boldsymbol{\delta}}} p_{v,\sigma} - y_{v,\sigma}\|_{1,\mathcal{E}}}{\|y_{v,\sigma}\|_{1,\mathcal{E}} + \varepsilon}. \tag{D.1}$$

We use 64×64 checkerboard tiles, two parity choices, and offsets $(0, 0)$ and $(32, 32)$. A guard band of $\lceil 3\sigma \rceil + 4$ pixels separates fitting and evaluation support. Improvements are aggregated by their median over splits and then views.

The self-consistent control starts from $\widetilde{x} = G_1 * \max(x_{\mathrm{FDK}}, 0)$, where G_1 smooths by one voxel. Projecting, reconstructing, and reprojecting $\widetilde{x}$ under the same geometry and view set gives the corresponding control score. Two further controls add Gaussian perturbations whose per-view scale is estimated from the training projection's Laplacian median absolute deviation. These are matched noise surrogates, not calibrated sensor-noise realizations.

At $\sigma = 2$, paired circular angle-block resampling uses blocks of five views and 1,000 repetitions. The score e is the minimum lower 2.5th-percentile endpoint of the observed-minus-control improvement across the self-consistent and two noisy controls. The stability gate requires positive median excess over the self-consistent control at all three scales, at least 75% positive fitting outcomes at scale two, and fewer than 25% clipped or ill-conditioned fits. A fit is ill-conditioned when its normal-matrix condition number exceeds 10^6.

Controlled projection perturbations determine the common anchors: $\ell = 9.56569 \times 10^{-5}$ is the largest no-shift controlled excess and $h = 1.19567 \times 10^{-3}$ is the median excess under a one-pixel perturbation in the calibration controls. The mapping in equation 14 uses $\alpha_{\min} = 0.5$ and $\alpha_{\max} = 0.8$. The score and its resampling interval are conditional on this diagnostic construction; they are not probabilities or confidence intervals for physical calibration error. Training weights use neither evaluation-reference volumes nor test projections.

D.2 TV Gradient Feedback

Let $\boldsymbol{\rho}$ denote the unconstrained density parameter block. For probe views $\mathcal{V}_p$ and voxel patches $\mathcal{Q}_p$, we compute

$$G_{\mathrm{data}} = \left(\frac{1}{|\mathcal{V}_p|} \sum_{v \in \mathcal{V}_p} \|\nabla_{\boldsymbol{\rho}} \mathcal{L}_{\mathrm{data,v}}\|_2^2 \right)^{1/2}, \quad G_{\mathrm{TV}} = \left(\frac{1}{|\mathcal{Q}_p|} \sum_{\Omega \in \mathcal{Q}_p} \|\nabla_{\boldsymbol{\rho}} \mathrm{TV}(f_\theta|_\Omega)\|_2^2 \right)^{1/2}. \tag{D.2}$$

The data term includes the fixed scan-specific L1/SSIM weights. The gradient statistic is the RMS of individual view or patch gradient norms, not the norm of an averaged gradient. The same parameter block is used for both terms.

Starting at iteration 1,000, probes occur every 1,000 iterations. Ten views are distributed over the sorted training angles, with a rotating offset across probes. Sixteen TV patches use one jittered location in each cell of a $4 \times 2 \times 2$ spatial partition. An independent random generator refreshes patch locations without changing training samples. At 50 views, the selected indices are every fifth view with a rotating offset; other view counts use ten approximately uniformly spaced indices.

With λ^* from equation 15, the update is

$$\log \lambda^+ = \log \lambda + \mathrm{clip}\big(\tfrac{1}{2}(\log \lambda^* - \log \lambda), -\log 1.5, \log 1.5\big). \tag{D.3}$$

Both target and updated coefficient lie in $[0.0125, 0.05]$. When G_{TV} is numerically negligible, the target is 0.05. Iterations 9,500–11,500 retain the existing coefficient to avoid following the immediate PAGR transient. Probe computation does not modify model parameters, existing gradient buffers, or the training random stream. The coefficient floor and smoothing can leave the measured TV-to-data gradient ratio above one.

E Loss and Evaluation Definitions

SSIM uses an 11×11 Gaussian window with standard deviation 1.5, zero padding of width 5, and constants 0.01^2 and 0.03^2 in the supplied normalized projection units. TV sums absolute differences along all neighboring voxel edges in the three axes and divides by the total edge count. It is evaluated on the summed attenuation field in 32^3 patches at the reconstruction voxel spacing, without an additional division by spacing.

Volume PSNR uses a fixed intensity range of one. Volume SSIM is the mean of slice-wise two-dimensional SSIM along the three coordinate axes, excluding reference slices with zero maximum. It is not a three-dimensional sliding-window SSIM.

F Additional Reconstruction Results at Matched Kernel Counts

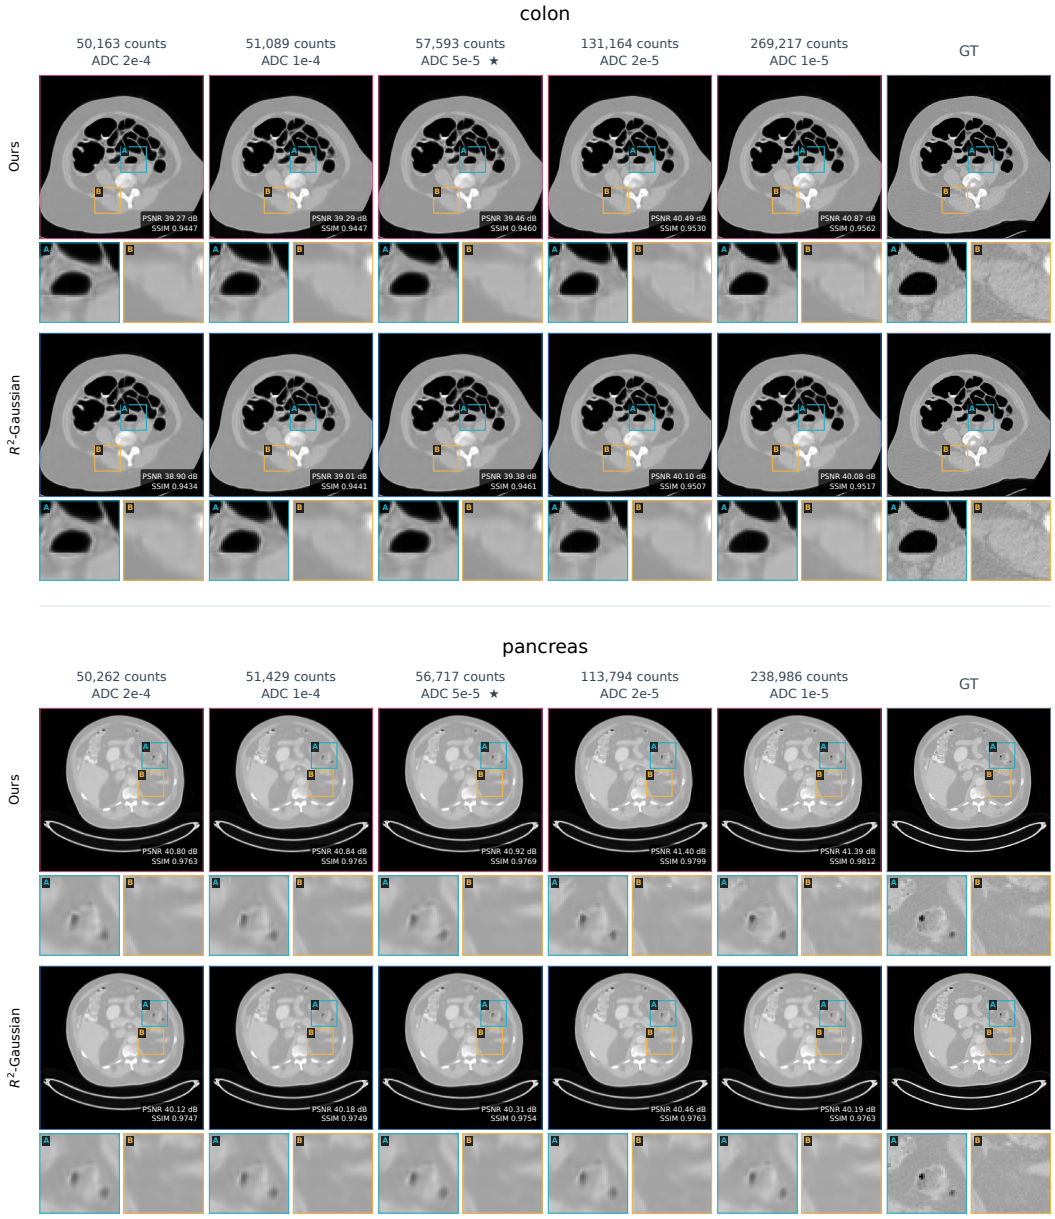

Figure 5: **Comparison of reconstruction results for the colon and pancreas under an equivalent Gaussian count budget.** The PSNR3D and SSIM3D values are displayed in the bottom-right corner of the figure. Magnified views of the ROI reveal that our method yields sharper anatomical structures in the colon reconstruction, although it introduces new artifacts compared to R² Gaussian. In the pancreas reconstruction, our method demonstrates a clear advantage, offering not only sharper anatomical details but also reconstructing actual structures that are absent in the R² Gaussian results.

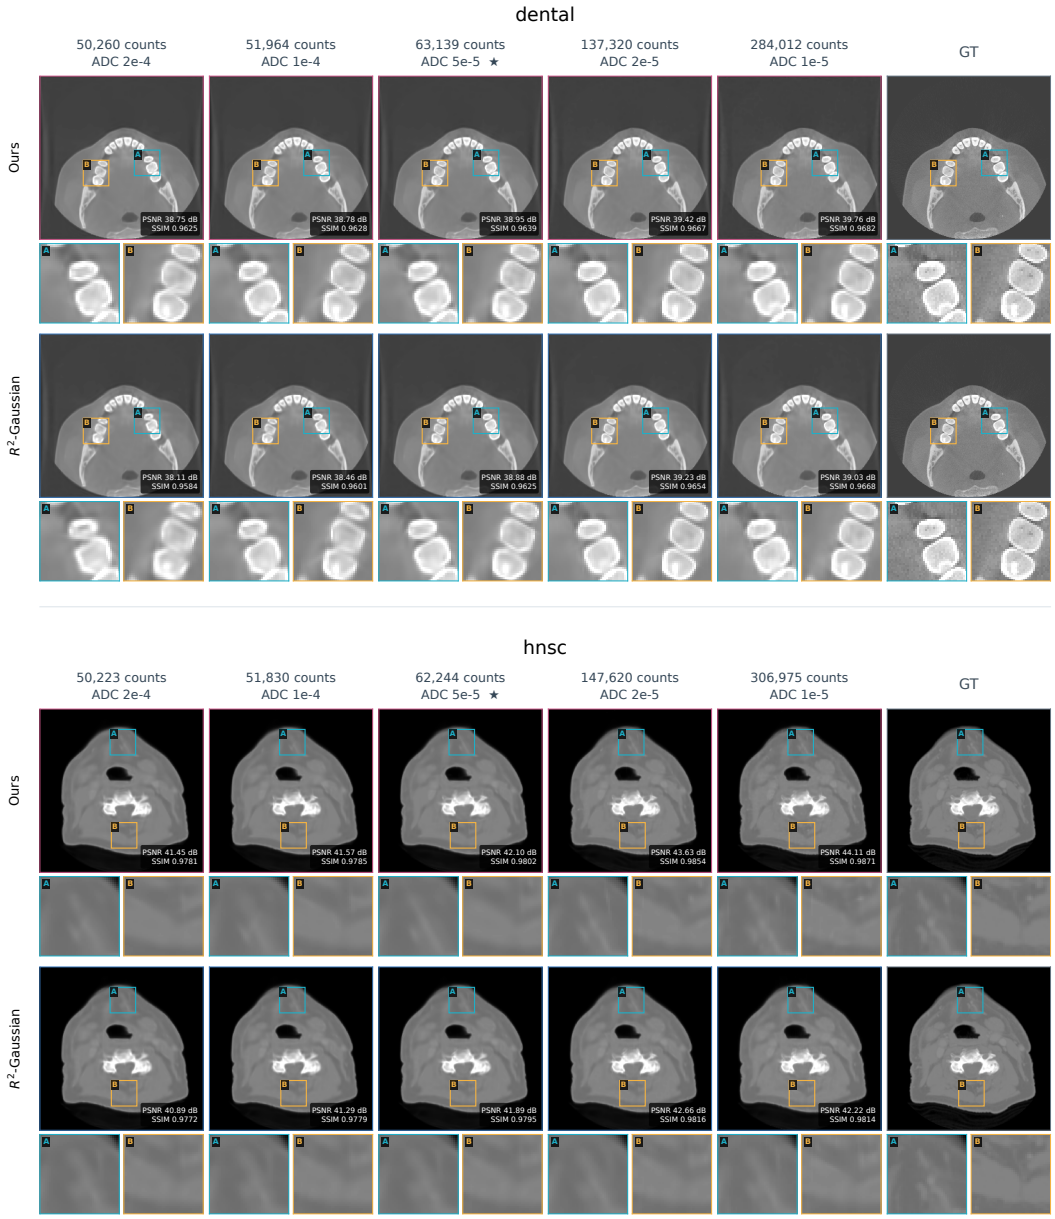

Figure 6: **Comparison of reconstruction results for the dental and head-and-neck under the same Gaussian count budget.** The PSNR3D and SSIM3D values are displayed in the bottom-right corner of the image. Magnified views of the ROI reveal that our method produces fewer artifacts in the dental reconstruction results, with tooth contrast closer to the ground truth. For the head-and-neck reconstruction, our method yields significantly sharper results and recovers more detailed structures.

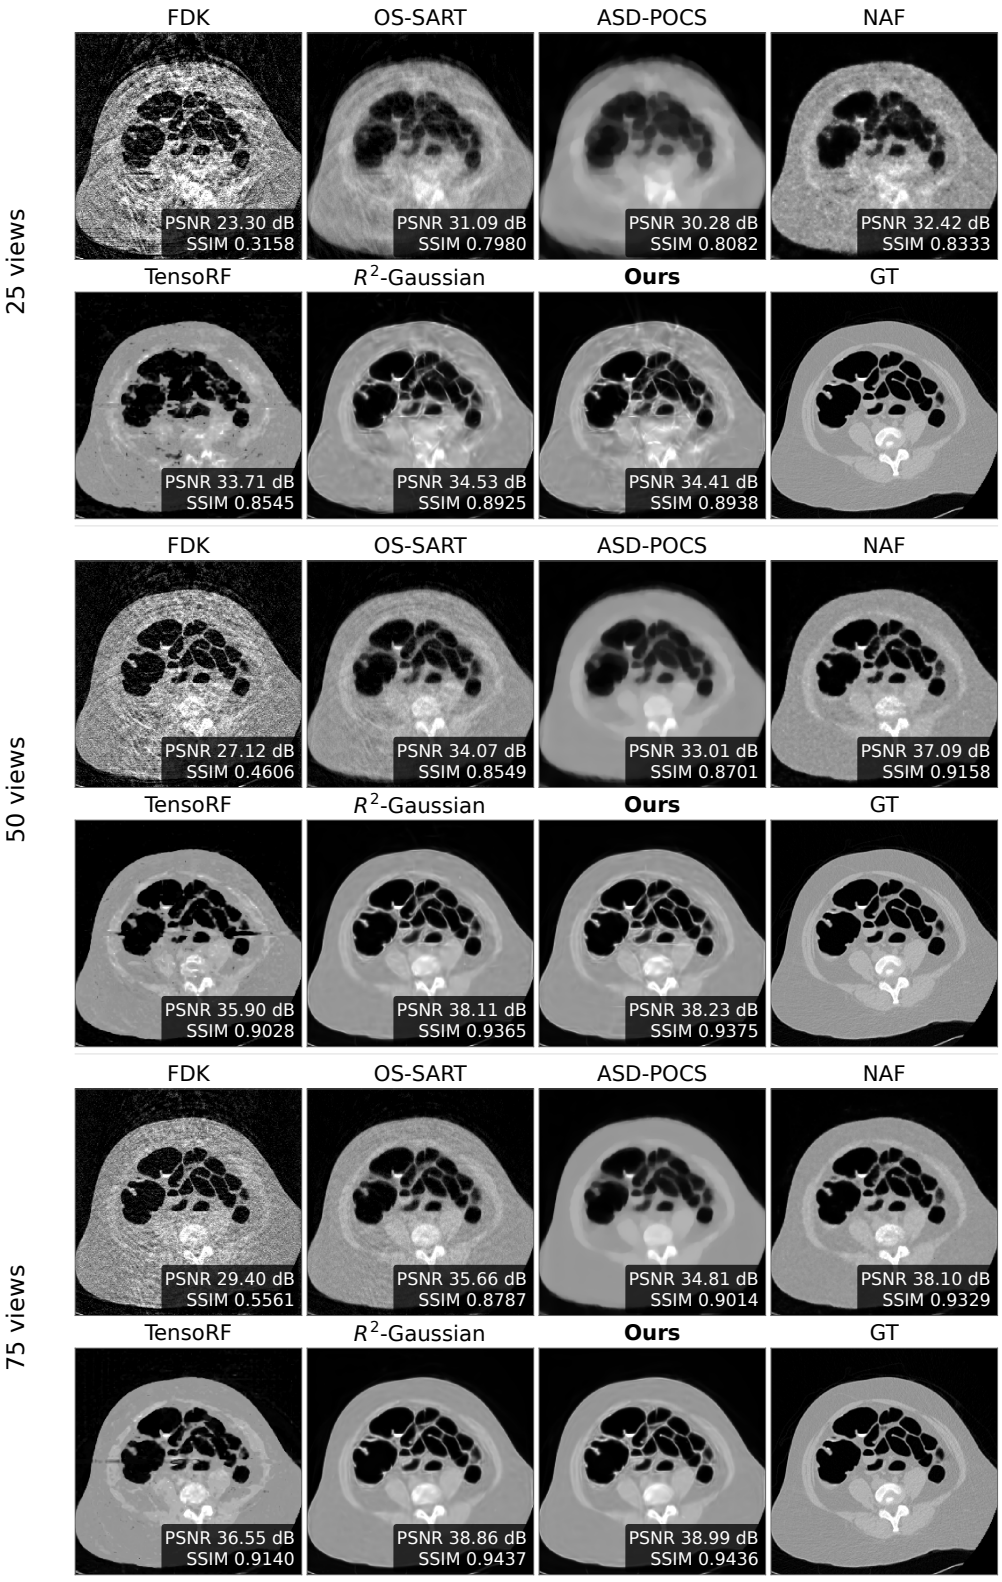

Figure 7: **Comparison of reconstruction results for the colon across various reconstruction methods.** The PSNR3D and SSIM3D values are displayed in the bottom-right corner of the image.

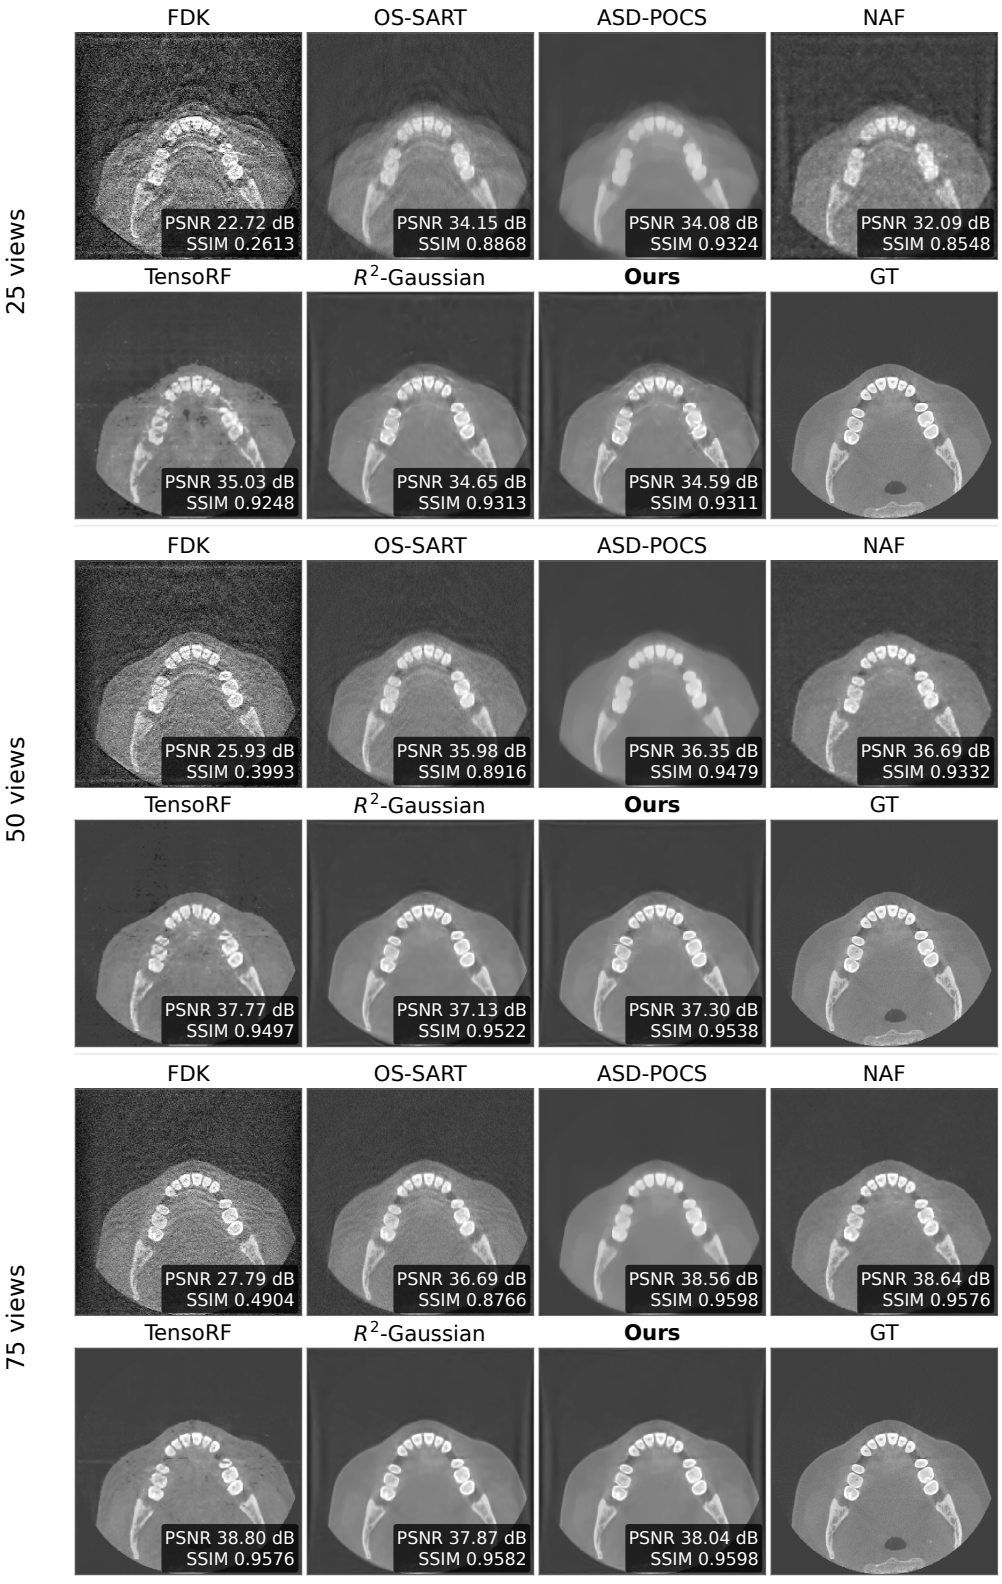

Figure 8: **Comparison of reconstruction results for the dental across various reconstruction methods.** The PSNR3D and SSIM3D values are displayed in the bottom-right corner of the image.

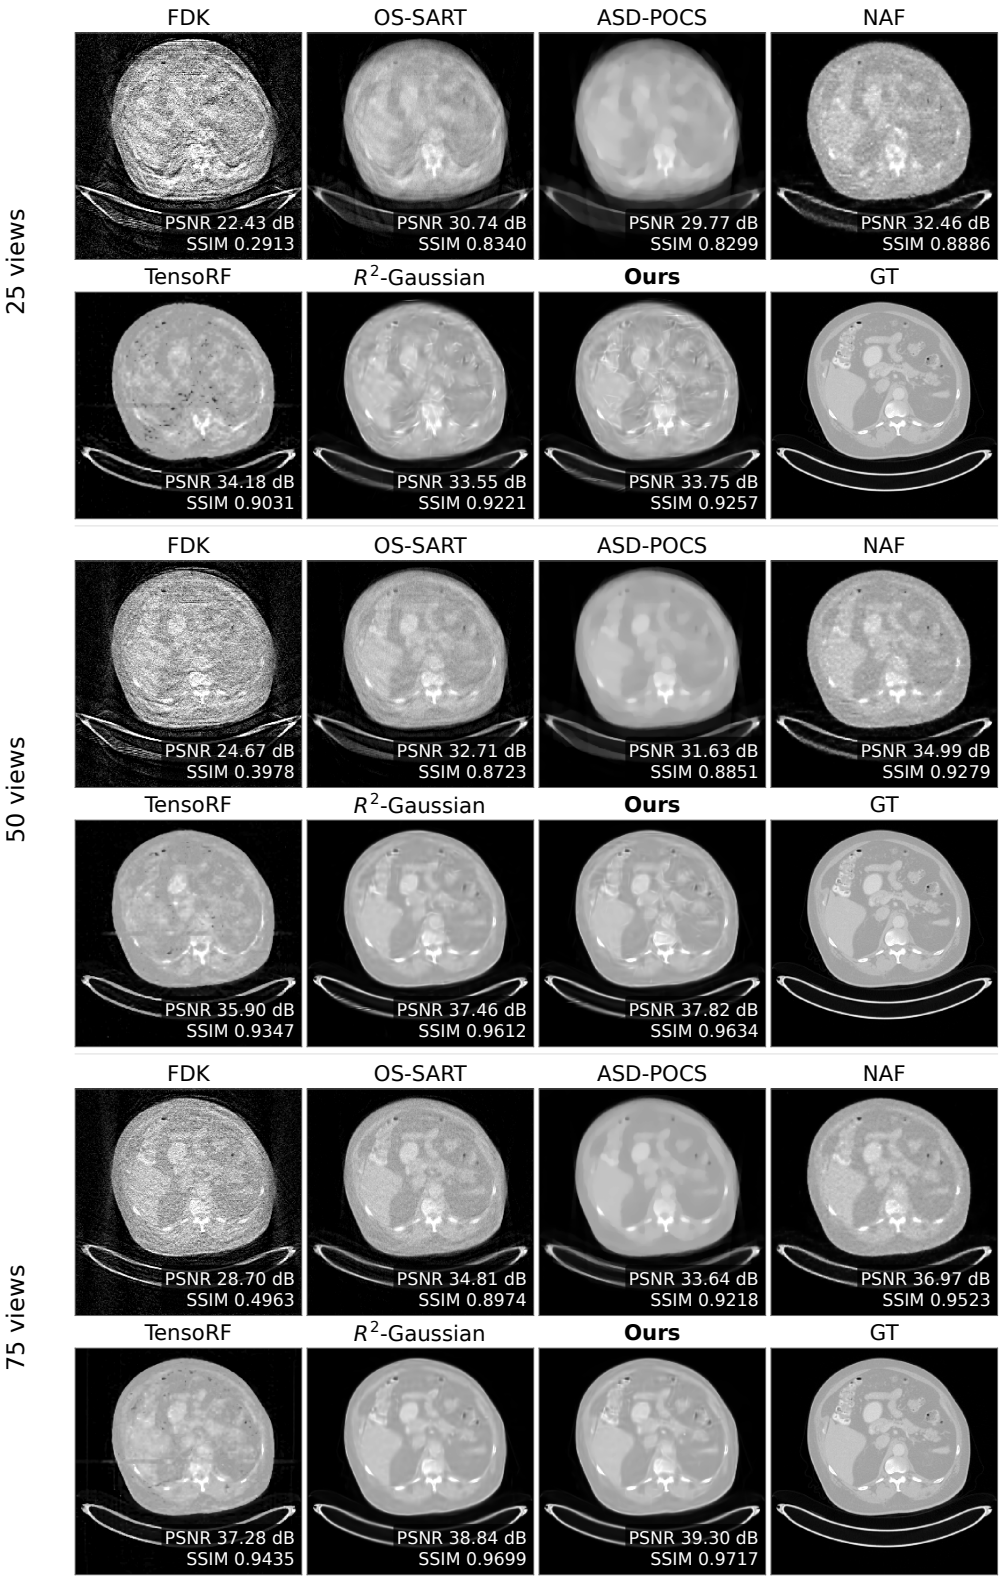

Figure 9: **Comparison of reconstruction results for the pancreas across various reconstruction methods.** The PSNR3D and SSIM3D values are displayed in the bottom-right corner of the image.

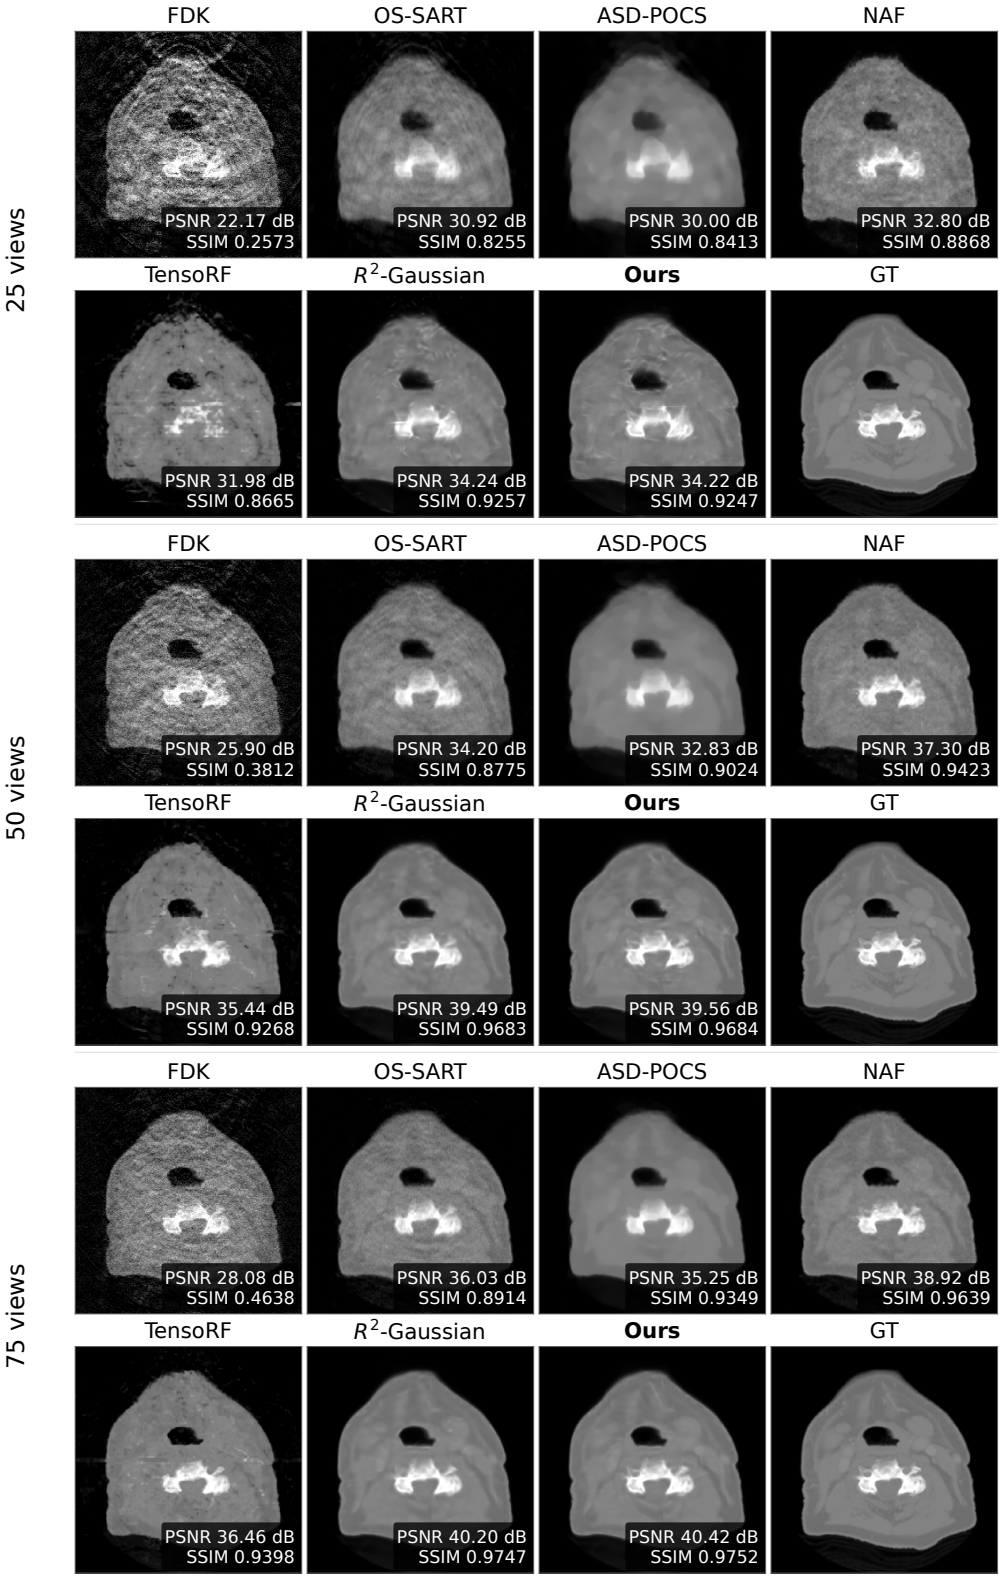

Figure 10: **Comparison of reconstruction results for the head-and-neck across various reconstruction methods.** The PSNR3D and SSIM3D values are displayed in the bottom-right corner of the image.

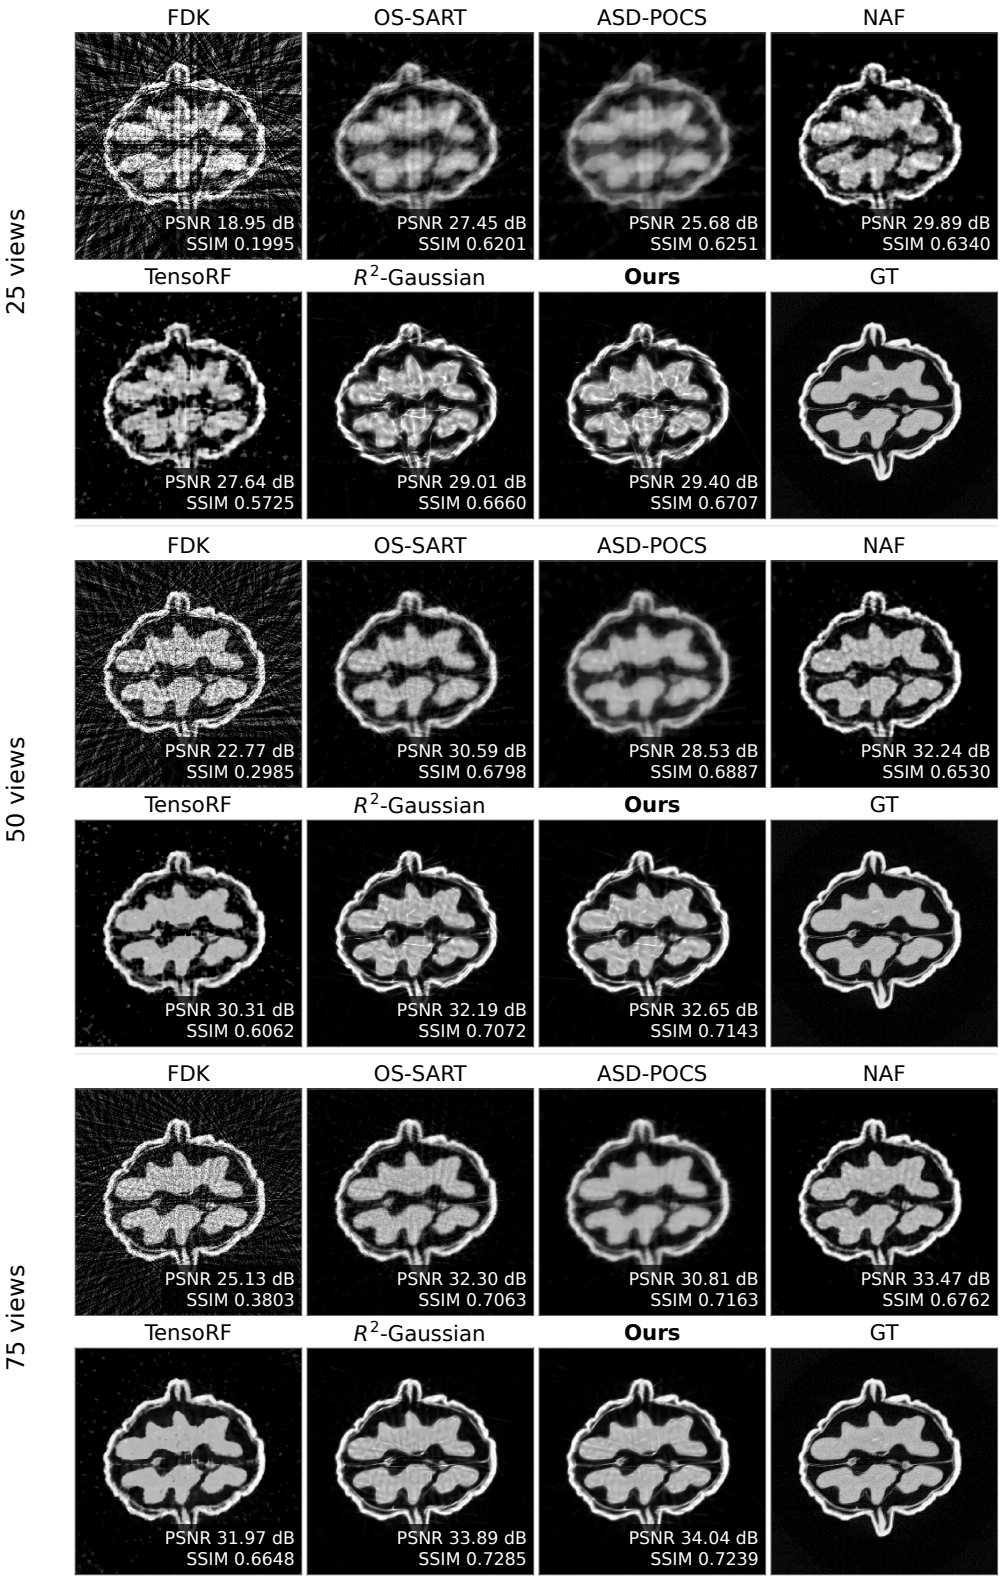

Figure 11: **Comparison of reconstruction results for the walnut across various reconstruction methods.** The PSNR3D and SSIM3D values are displayed in the bottom-right corner of the image.

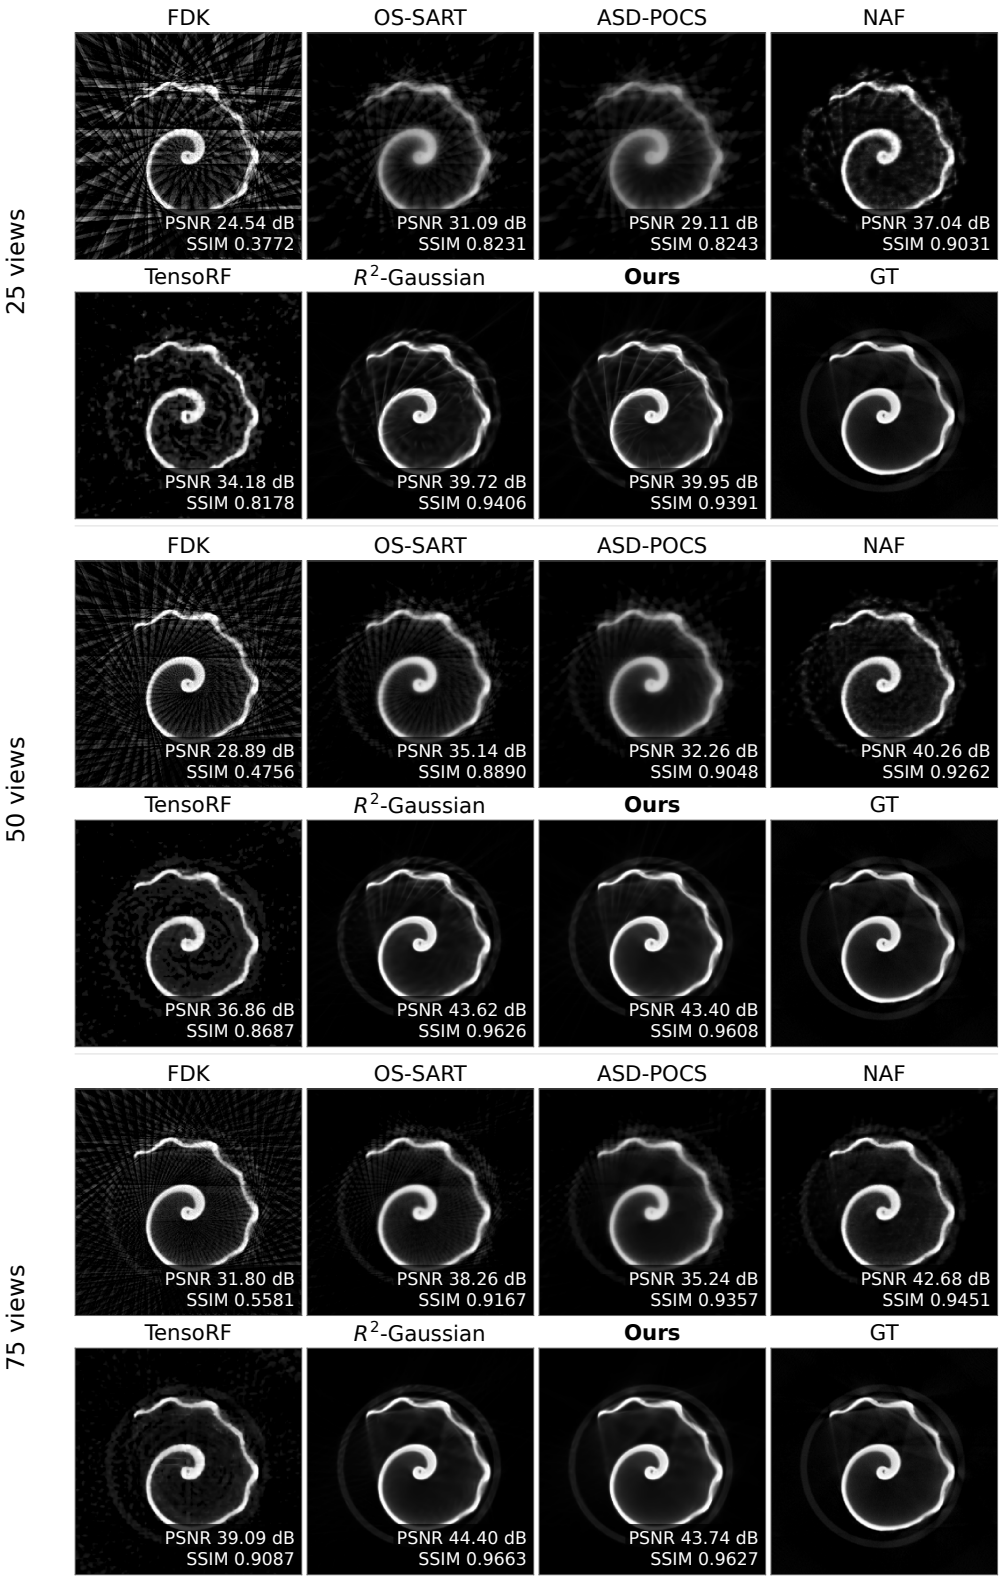

Figure 12: **Comparison of reconstruction results for the seashell across various reconstruction methods.** The PSNR3D and SSIM3D values are displayed in the bottom-right corner of the image.

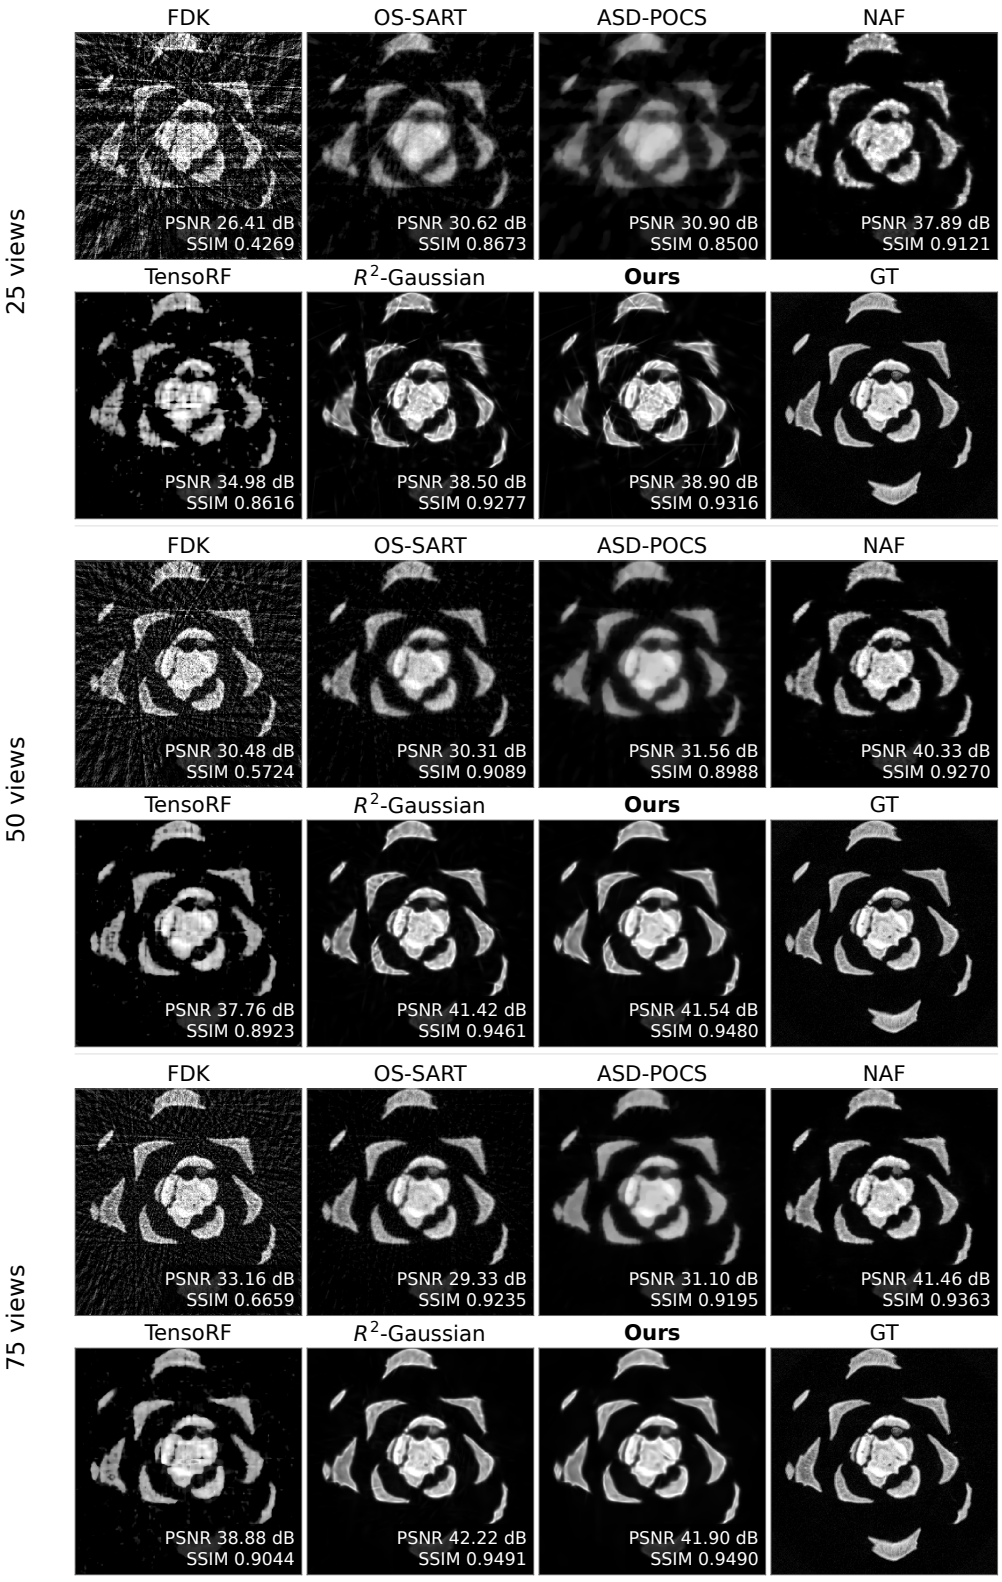

Figure 13: **Comparison of reconstruction results for the pine across various reconstruction methods.** The PSNR3D and SSIM3D values are displayed in the bottom-right corner of the image.

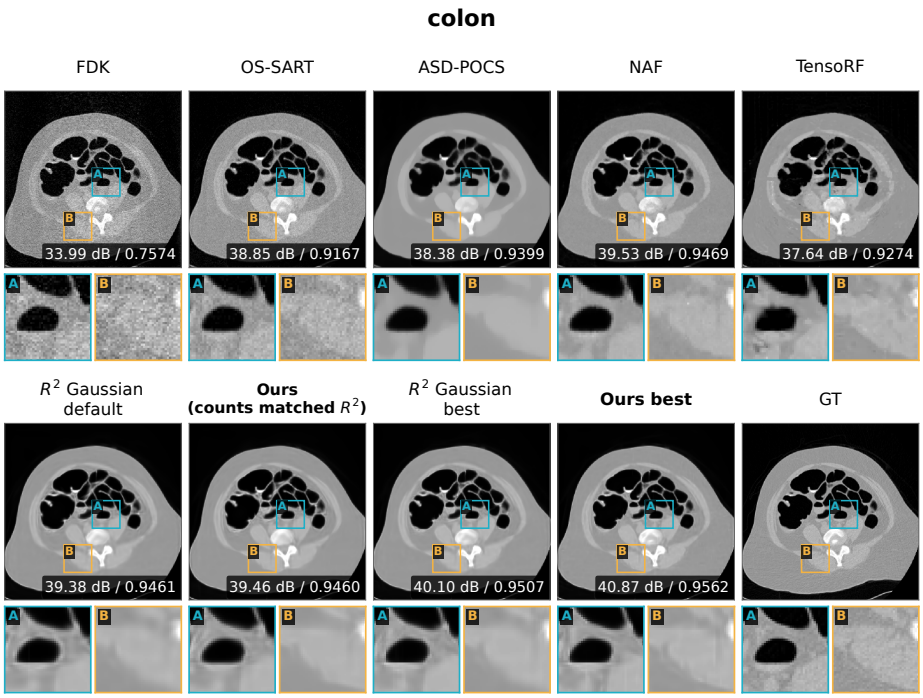

Figure 14: **The colon reconstruction from 180 angles of the other methods compared with our method.** The metrics PSNR3D, SSIM3D are indicated.

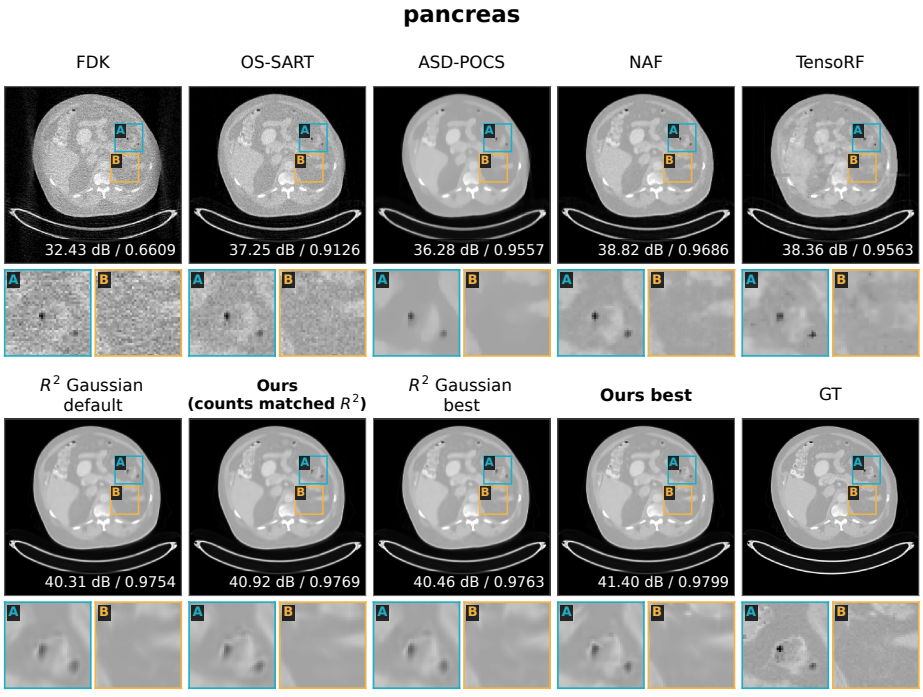

Figure 15: **The pancreas reconstruction from 180 angles of the other methods compared with our method.** The metrics PSNR3D, SSIM3D are indicated.

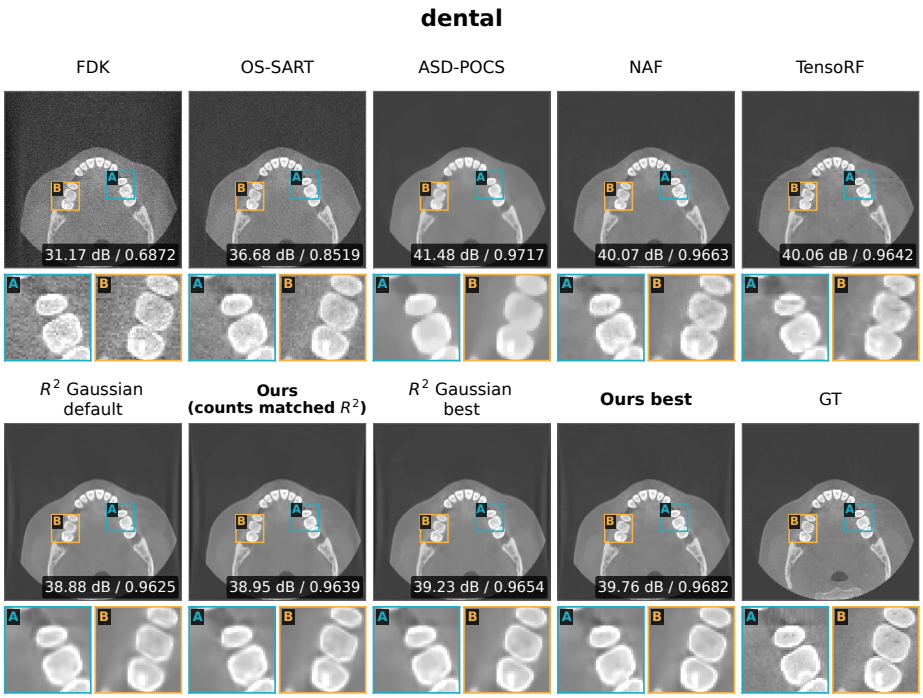

Figure 16: **The dental reconstruction from 180 angles of the other methods compared with our method.** The metrics PSNR3D, SSIM3D are indicated.