

R²-Gaussian: Rectifying Radiative Gaussian Splatting for Tomographic Reconstruction

Ruyi Zha¹ Tao Jun Lin¹ Yuanhao Cai² Jiwen Cao¹
Yanhao Zhang³ Hongdong Li¹

¹The Australian National University

²Johns Hopkins University ³University of Technology Sydney

NeurIPS 2024



Australian
National
University



JOHNS HOPKINS
UNIVERSITY




UNIVERSITY
OF TECHNOLOGY
SYDNEY




Background: X-ray Imaging and Tomographic Reconstruction


- X-ray: penetrable high-energy beams
- X-ray projection: visualize internal structures
- X-ray imaging function: $I(\mathbf{r}) = \int_{t_n}^{t_f} \sigma(\mathbf{r}(t)) dt$
- Tomographic reconstruction
 - Recover the density volume $\sigma(\mathbf{x})$ from multi-angle projections
 - Challenges: sparse-view, noise, speed, etc



RGB imaging




X-ray imaging



Previous Work

- Traditional methods
 - Fast (< 10 min)
 - Bad quality
- Supervised-learning methods
 - Good quality
 - External dataset required
 - Poor generation
- NeRF-based methods
 - Good quality
 - Good generation (self-supervised)
 - Slow training (> 30 min)




A new method should be...

- Self-supervised: generalizable to arbitrary objects
- Fast training: comparable to traditional methods

Background: 3D Gaussian Splatting (3DGS)


- SOTA method for RGB view synthesis
- Represent objects with trainable 3D Gaussians
- Differentiable rasterization for fast rendering and training



Can we use the idea of 3DGS for tomographic reconstruction?

Challenges


- How to define 3D Gaussians for X-ray imaging?
- How to rasterize X-ray projections?
- How to **directly**¹ get a density volume from Gaussians?



¹There are some 3DGS works for X-ray view synthesis, but they can't obtain density volumes from Gaussians, thus not suitable for tomographic reconstruction.

Representing Density Fields as Radiative Gaussians

- We define radiative Gaussian as a local Gaussian-shaped density field.
- The overall density field is the sum of radiative Gaussians.
- We use radiative Gaussians for both 2D rendering and 3D querying.




X-ray Rasterization


Rasterization pipeline

- Transformation: world space to ray space
- **Projection**: 3D Gaussians to 2D Gaussians
- **Composition**: 2D Gaussians to image

3DGS rasterization



Our X-ray rasterization



Integration Bias in 3DGS

Integration property of normalized Gaussian distribution

Integrating a normalized 3D Gaussian distribution along an axis yields a normalized 2D Gaussian distribution:

$$\int_{\mathbb{R}} \frac{1}{(2\pi)^{\frac{3}{2}} |\Sigma|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{p})^\top \Sigma^{-1}(\mathbf{x} - \mathbf{p})\right) dx_2 = \frac{1}{2\pi |\hat{\Sigma}|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(\hat{\mathbf{x}} - \hat{\mathbf{p}})^\top \hat{\Sigma}^{-1}(\hat{\mathbf{x}} - \hat{\mathbf{p}})\right),$$

where $\mathbf{x} = \begin{bmatrix} x_0 \\ x_1 \\ x_2 \end{bmatrix} \Rightarrow \hat{\mathbf{x}} = \begin{bmatrix} x_0 \\ x_1 \end{bmatrix}$, $\Sigma = \begin{bmatrix} a & b & c \\ b & d & e \\ c & e & f \end{bmatrix} \Rightarrow \hat{\Sigma} = \begin{bmatrix} a & b \\ b & d \end{bmatrix}$, and $\mathbf{p} = \begin{bmatrix} p_0 \\ p_1 \\ p_2 \end{bmatrix} \Rightarrow \hat{\mathbf{p}} = \begin{bmatrix} p_0 \\ p_1 \end{bmatrix}$.

When projecting a 3D Gaussian primitive to its 2D companion, we have:


$$\begin{aligned} G^2(\hat{\mathbf{x}}|\hat{\rho}, \hat{\mathbf{p}}, \hat{\Sigma}) &= \int G^3(\mathbf{x}|\rho, \mathbf{p}, \Sigma) dx_2 = \rho \int \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{p})^\top \Sigma^{-1}(\mathbf{x} - \mathbf{p})\right) dx_2 \\ &= \rho (2\pi)^{\frac{3}{2}} |\Sigma|^{\frac{1}{2}} \int \frac{1}{(2\pi)^{\frac{3}{2}} |\Sigma|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{p})^\top \Sigma^{-1}(\mathbf{x} - \mathbf{p})\right) dx_2 \\ &= \rho (2\pi)^{\frac{3}{2}} |\Sigma|^{\frac{1}{2}} \frac{1}{2\pi |\hat{\Sigma}|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(\hat{\mathbf{x}} - \hat{\mathbf{p}})^\top \hat{\Sigma}^{-1}(\hat{\mathbf{x}} - \hat{\mathbf{p}})\right) \\ &= G^2(\hat{\mathbf{x}}|\underbrace{\sqrt{2\pi|\Sigma|/|\hat{\Sigma}|}}_{\mu} \rho, \hat{\mathbf{p}}, \hat{\Sigma}) \end{aligned}$$

Note the density (opacity) change: $\hat{\rho} = \mu\rho$

Integration Bias in 3DGS

3D density $\Rightarrow$ 2D density: $\hat{\rho} = \underbrace{\sqrt{2\pi|\Sigma|/|\hat{\Sigma}|}}_{\mu} \rho$


- μ is view-dependent, making $\hat{\rho}$ **view-dependent**.
- 3DGS: view-independent $\hat{\rho} = \rho$
 - OK for 2D rendering
 - inconsistency (bias) for 3D querying: $\rho \stackrel{?}{=} \hat{\rho}/\mu$
- Ours: view-dependent $\hat{\rho} = \mu\rho$
 - No inconsistency problem for 3D querying



X-ray Voxelization


We develop a voxelizer to query density volumes from Gaussians.

- CUDA-based, very fast
- Differentiable, support 3D losses




Overall Pipeline

- New initialization strategy by sampling points from FDK volume
- 2D image loss (rasterizer) + 3D regularization (voxelizer)
- Adaptive control for point densification



Experiments: Reconstruction Quality


SOTA reconstruction quality



Experiments: Efficiency

- Outperform baselines in 4 minutes (10k iteration)
 - $\approx$ traditional methods
 - $14\times$ faster than NeRF-based methods
- Optimal results in 15 minutes (30k iteration)
 - $4\times$ faster than NeRF-based methods


Methods	75-view			50-view			25-view		
	PSNR $\uparrow$	SSIM $\uparrow$	Time $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	Time $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	Time $\downarrow$
Synthetic dataset									
FDK [13]	28.63	0.497	-	26.50	0.422	-	22.99	0.317	-
SART [2]	36.06	0.897	4m41s	34.37	0.875	3m36s	31.14	0.825	1m47s
ASD-POCS [55]	36.64	0.940	2m25s	34.34	0.914	1m52s	30.48	0.847	56s
IntraTomo [66]	35.42	0.924	2h7m	35.25	0.923	2h9m	34.68	0.914	2h19m
NAF [67]	37.84	0.945	30m43s	36.65	0.932	32m4s	33.91	0.893	31m1s
SAX-NeRF [6]	38.07	0.950	13h5m	36.86	0.938	13h5m	34.33	0.905	13h3m
Ours (iter=10k)	38.29	0.954	2m38s	37.63	0.949	2m35s	35.08	0.922	2m35s
Ours (iter=30k)	38.88	0.959	8m21s	37.98	0.952	8m14s	35.19	0.923	8m28s
Real-world dataset									
FDK [13]	30.03	0.535	-	27.38	0.449	-	23.30	0.335	-
SART [2]	34.42	0.845	5m11s	33.61	0.827	3m28s	31.52	0.790	1m47s
ASD-POCS [55]	36.33	0.868	2m43s	34.58	0.861	1m49s	31.32	0.810	56s
IntraTomo [66]	36.79	0.858	2h25m	36.99	0.854	2h19m	35.85	0.835	2h18m
NAF [67]	38.58	0.848	51m28s	36.44	0.818	51m31s	32.92	0.772	51m24s
SAX-NeRF [6]	34.93	0.854	13h21m	34.89	0.840	13h23m	33.49	0.793	13h25m
Ours (iter=10k)	38.10	0.872	3m39s	37.52	0.866	3m37s	35.10	0.840	3m23s
Ours (iter=30k)	39.40	0.875	14m16s	38.24	0.864	13m52s	34.83	0.833	12m56s



Ablation: Integration Bias

Correcting integration bias benefits both X-ray rendering and reconstruction.

	75-view		50-view		25-view	
	X-3DGS	Ours	X-3DGS	Ours	X-3DGS	Ours
2D PSNR $\uparrow$	49.97	50.54	47.26	49.70	39.84	46.28
2D SSIM $\uparrow$	0.987	0.986	0.984	0.986	0.967	0.982
3D PSNR $\uparrow$	23.40	38.86	21.24	37.98	14.07	35.17
3D SSIM $\uparrow$	0.660	0.959	0.562	0.952	0.408	0.923




Summary

- First 3DGS-based method for tomographic reconstruction
- SOTA quality and efficiency
- Discovery of integration bias in 3DGS



Paper



Project page



Australian
National
University



JOHNS HOPKINS
UNIVERSITY

