3D Gaussian Splatting for Real-Time Radiance Field Rendering

November 20, 2023

Bernhard Kerbl*, Georgios Kopanas*, Thomas Leimkühler, George Drettakis
In SIGGRAPH 2023

Presenter: Donghwan Kim 20233091

Contents

1. Background

Rasterization

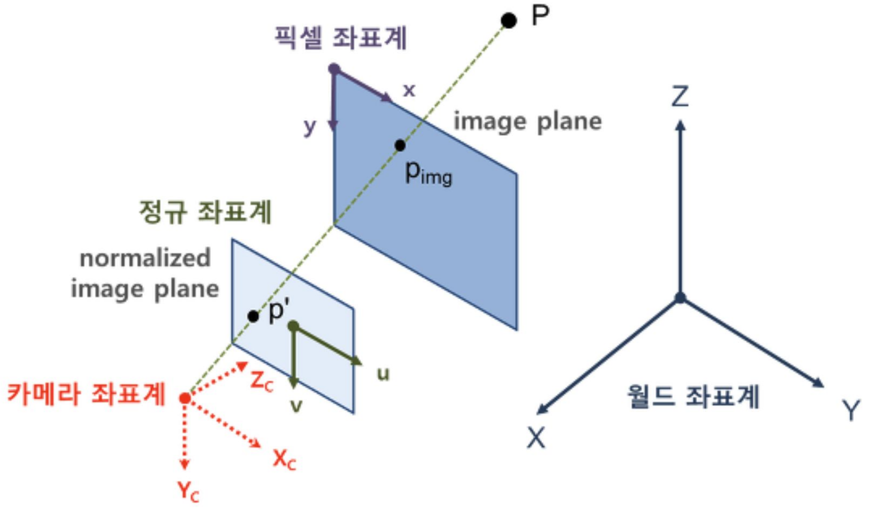

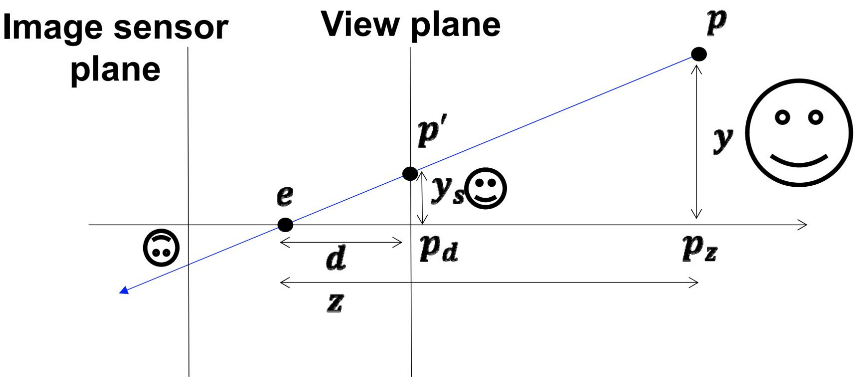

https://darkpgmr.tistory.com/77

Lecture Slide for <3D Transformation> from CS380 (Spring 2020) Sung-eui Yoon

1. Background

Ray Casting

- **For each pixel, find closest object along the ray and shade pixel accordingly**

- **Advantages**
  - **Conceptually simple**
  - **Can be extended to handle global illumination effects**
- **Disadvantages**
  - **Renderer must have access to entire retained model**
  - **Hard to map to special-purpose hardware**
  - **Less efficient than rasterization in terms of utilizing spatial coherence**

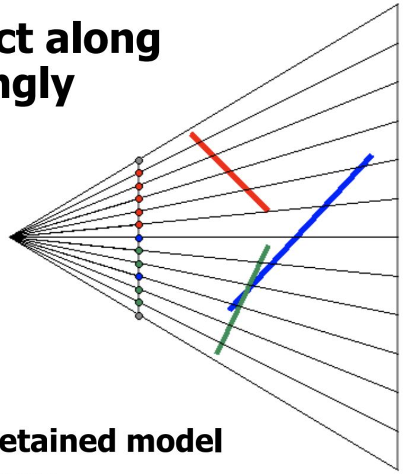

[Lecture Slide for <Ray tracing> from CS380 (Spring 2020) Sung-eui Yoon]

1. Background

NeRF (Neural Radiance Fields): Ray-marching

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*, 2020.

2. Contribution

3D Gaussians

1. preserve **continuous** properties
2. avoid unnecessary computation in **empty space**
3. **rasterization**-based rendering

3. Method

Overview

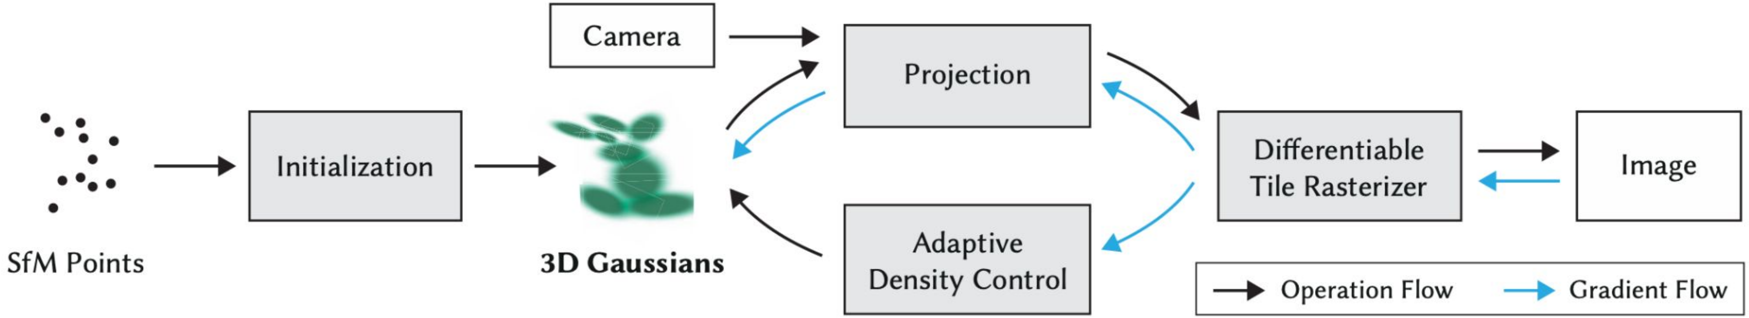

3. Method

Structure-from-Motion

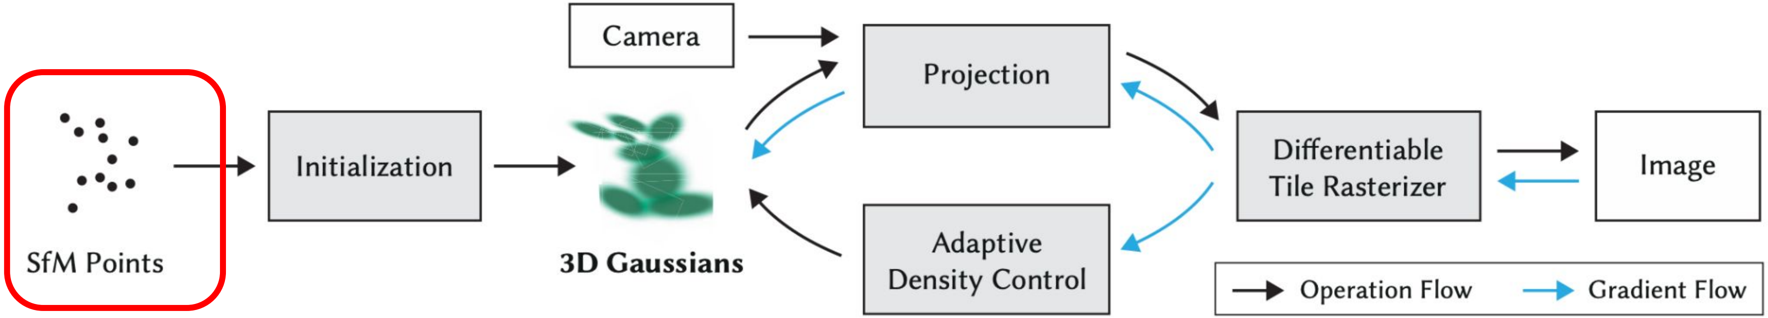

3. Method

Structure-from-Motion

- Camera calibration for in-the-wild NeRF.

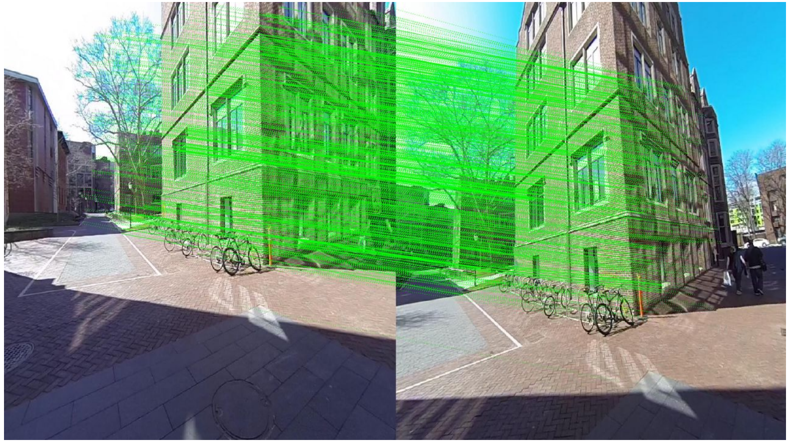

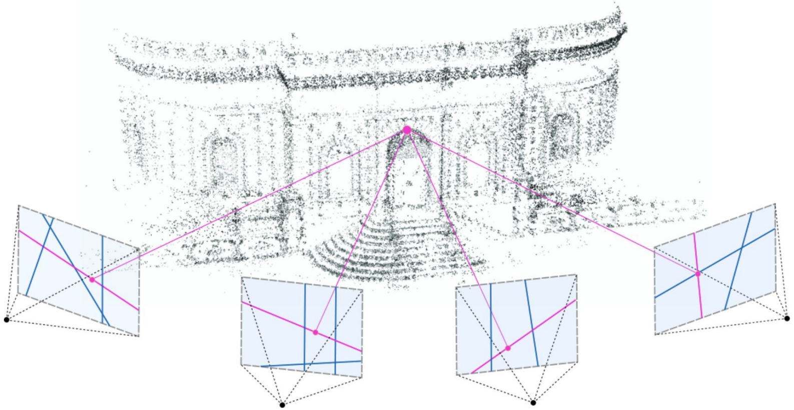

https://mvje.tistory.com/92

https://woochan-autobiography.tistory.com/944

3. Method

Initial spare point cloud

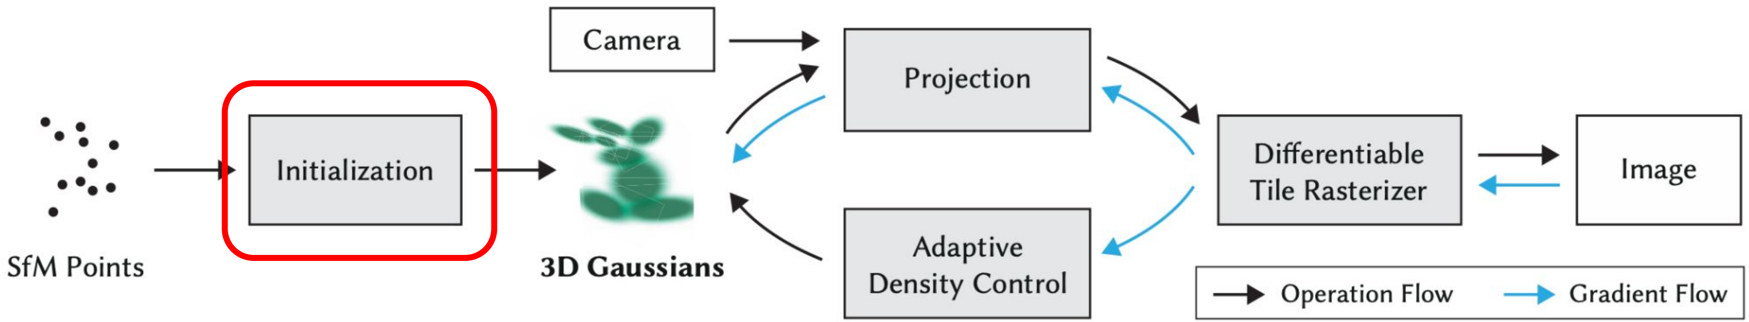

3. Method

3D Gaussians

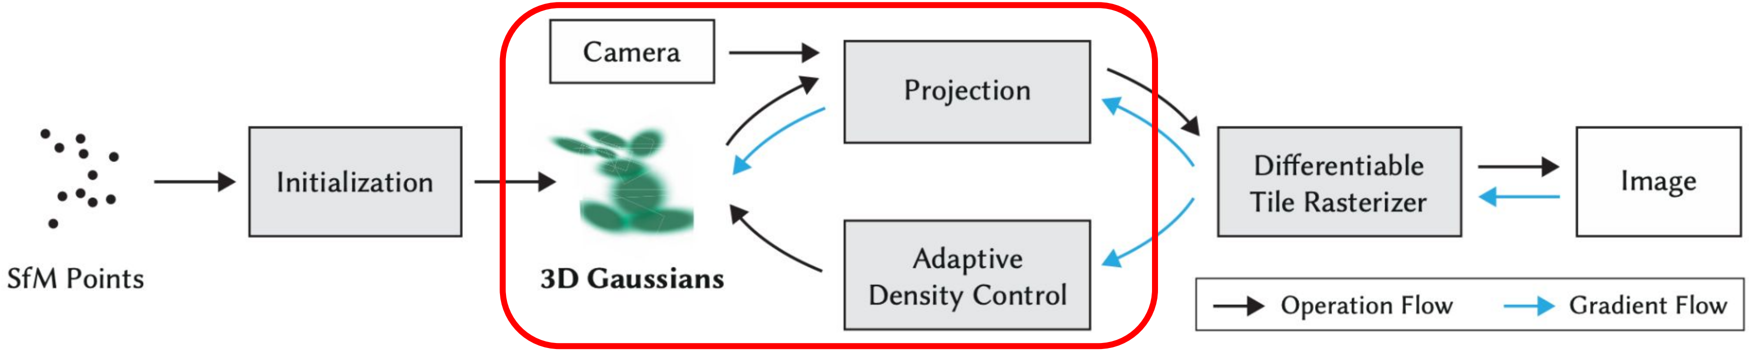

3. Method

3D Gaussian Splatting - Geometry

- 3D Gaussian

$$\mathcal{G}(\mathbf{x}-\mu) = \frac{1}{2\pi\Sigma(\mathbf{x})^{\frac{1}{2}}}\exp(-\frac{1}{2}(\mathbf{x}-\mu)^T\Sigma^{-1}(\mathbf{x}))$$

- 2D projection, where ***W*** is viewing transformation, and ***J*** is the Jacobian of affine projective transformation.

$$\Sigma' = JW\Sigma W^T J^T$$

- Positive semi-definite covariance, where ***S*** is scale and ***R*** is rotation matrix.

$$\Sigma = RSS^TR^T$$

Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, Markus Gross. EWA volume splatting. In *Proceedings of IEEE Visualization*, 2001.

3. Method

3D Gaussian Splatting - Opacity and Color

- Opacity α
- Color - Spherical Harmonics (SH) coefficients

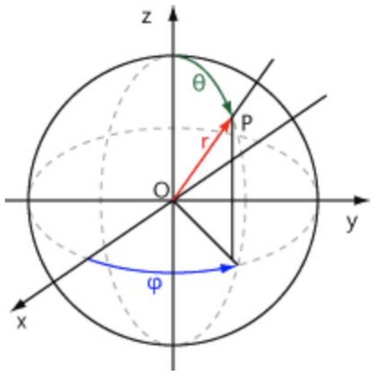

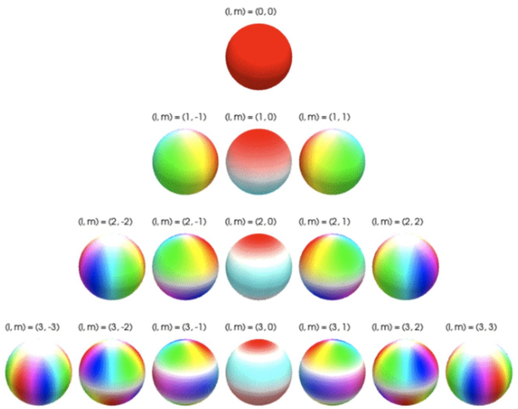

https://xoft.tistory.com/50

Robin Green. Spherical harmonic lighting: The gritty details. In Archives of the game developers conference, 2003.

3. Method

3D Gaussian Optimization

- Position (mean) p
- Covariance Σ
- Opacity α
- Color - Spherical Harmonics (SH) coefficients

→ Gradient Descent

- **There is no neural networks to optimize!**

3. Method

Adaptive Control of Gaussians

- **Clone** 3D gaussians: Under-reconstruction (missing geometric features)
- **Split** 3D gaussians: Over-reconstruction (covering large area)
- **Remove** 3D gaussians: Opacity is lower than threshold

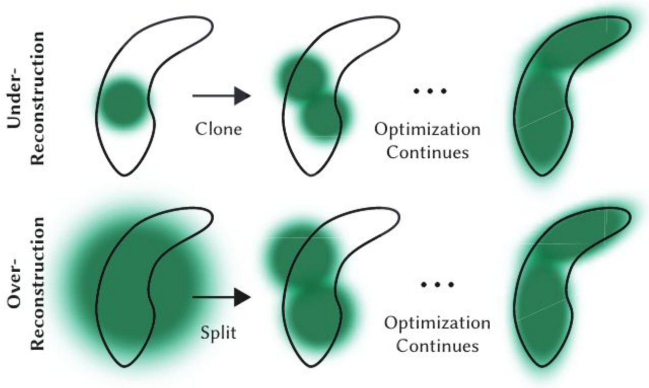

3. Method

Optimization

3. Method

Tile-based Rasterizer

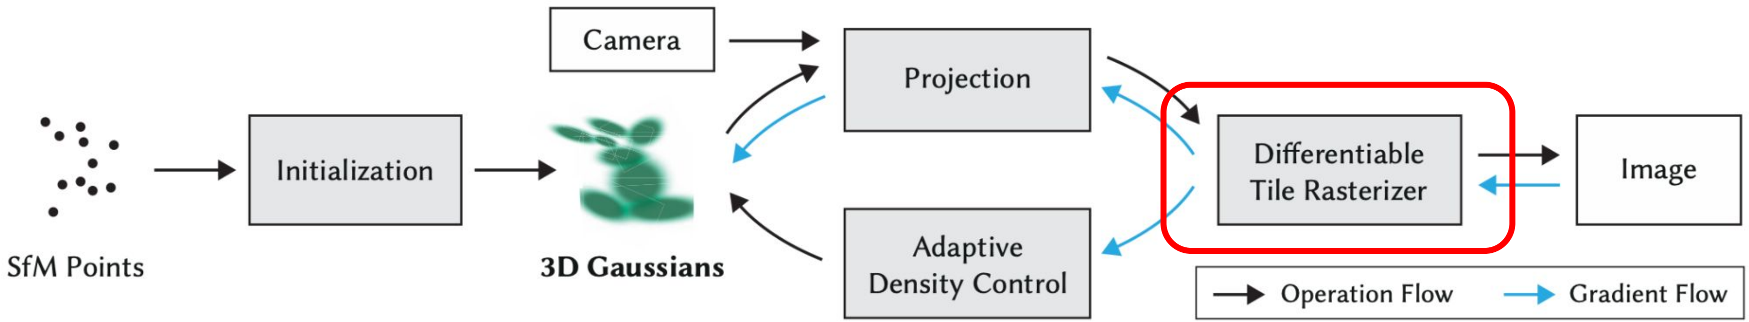

3. Method

Split Image to 16x16 Tiles

1. Parallelize within GPU cores.
2. Avoid the expense of depth sorting per pixel.

3. Method

Rasterization

For NeRF, $C(\mathbf{r}) = \int_{t_n}^{t_f} T(t)\sigma(\mathbf{r}(t))\mathbf{c}(\mathbf{r}(t), \mathbf{d})dt$

, where $T(t) = \exp(-\int_{t_n}^{t} \sigma(\mathbf{r}(s))ds)$

For 3D Gaussian Splatting, the color and opacity α are accumulated per pixel

from the closest to farthest gaussian

until opacity threshold

4. Results

Evaluation Metrics

1. PSNR (Peak Signal-to-Noise Ratio) ↑
2. SSIM (Structural Similarity Index Measure) ↑
3. LPIPS (Learned Perceptual Image Patch Similarity) ↓
4. FPS (Frame Per Seconds for inference rendering with A6000) ↑
5. Train (training time for A6000) ↓

Model Variations

1. Ours7K: 7K iterations
2. Ours30K: 30K iterations

4. Results

1. State of the Art Quality (Equivalent to MipNerf360)
2. Real-Time Rendering (More than 100 FPS)
3. Fast Training (Less than 1h)

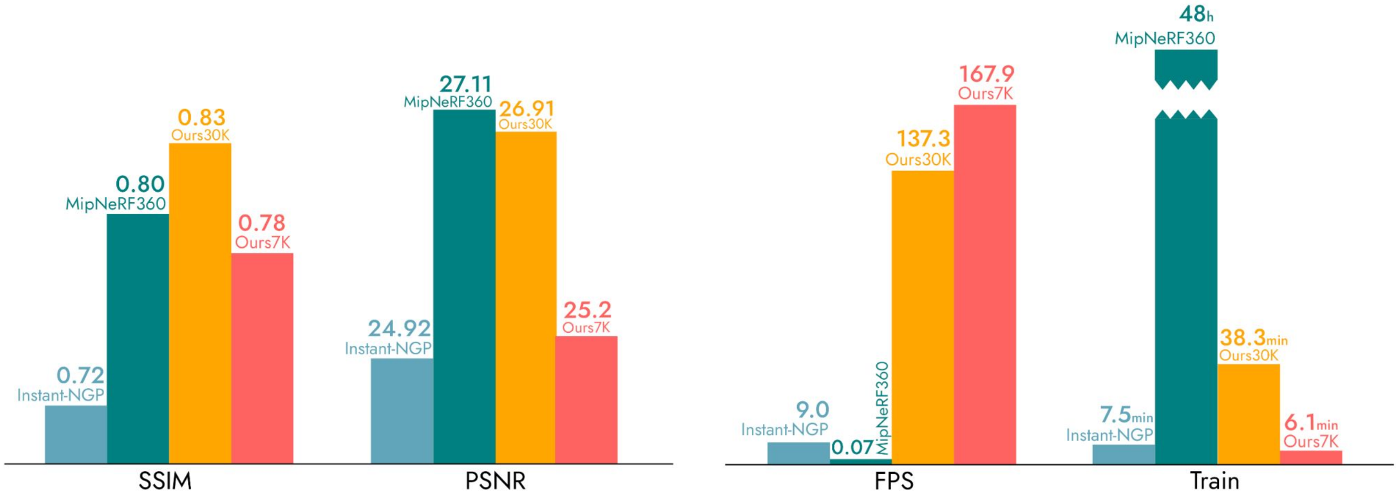

4. Results

Qualitative Evaluation

4. Results

Quantitative Evaluation

Dataset / Method|Metric	Mip-NeRF360 SSIM↑	PSNR↑	LPIPS↓	Train	FPS	Mem	Tanks&Temples SSIM↑	PSNR↑	LPIPS↓	Train	FPS	Mem	Deep Blending SSIM↑	PSNR↑	LPIPS↓	Train	FPS	Mem
Plenoxels	0.626	23.08	0.463	25m49s	6.79	2.1GB	0.719	21.08	0.379	25m5s	13.0	2.3GB	0.795	23.06	0.510	27m49s	11.2	2.7GB
INGP-Base	0.671	25.30	0.371	5m37s	11.7	13MB	0.723	21.72	0.330	5m26s	17.1	13MB	0.797	23.62	0.423	6m31s	3.26	13MB
INGP-Big	0.699	25.59	0.331	7m30s	9.43	48MB	0.745	21.92	0.305	6m59s	14.4	48MB	0.817	24.96	0.390	8m	2.79	48MB
M-NeRF360	0.792†	27.69†	0.237†	48h	0.06	8.6MB	0.759	22.22	0.257	48h	0.14	8.6MB	0.901	29.40	0.245	48h	0.09	8.6MB
Ours-7K	0.770	25.60	0.279	6m25s	160	523MB	0.767	21.20	0.280	6m55s	197	270MB	0.875	27.78	0.317	4m35s	172	386MB
Ours-30K	0.815	27.21	0.214	41m33s	134	734MB	0.841	23.14	0.183	26m54s	154	411MB	0.903	29.41	0.243	36m2s	137	676MB

PSNR scores for Synthetic NeRF

	Mic	Chair	Ship	Materials	Lego	Drums	Ficus	Hotdog	Avg.
Plenoxels	33.26	33.98	29.62	29.14	34.10	25.35	31.83	36.81	31.76
INGP-Base	36.22	35.00	31.10	29.78	36.39	26.02	33.51	37.40	33.18
Mip-NeRF	36.51	35.14	30.41	30.71	35.70	25.48	33.29	37.48	33.09
Point-NeRF	35.95	35.40	30.97	29.61	35.04	26.06	36.13	37.30	33.30
Ours-30K	35.36	35.83	30.80	30.00	35.78	26.15	34.87	37.72	33.32

5. Conclusion

They proposed new radiance field rendering methods using **3D gaussian** as volumetric representation for **real-time rendering** while preserving **state-of-the-art quality**.

References

Author's project page and video

https://darkpgmr.tistory.com/77

https://mvje.tistory.com/92

https://woochan-autobiography.tistory.com/944

https://xoft.tistory.com/50

Lecture Slide of CS380 (Spring 2020) Sung-eui Yoon

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, George Drettakis. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. In *SIGGRAPH*, 2023.

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*, 2020.

Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, Markus Gross. EWA volume splatting. In *Proceedings of IEEE Visualization*, 2001.

Robin Green. Spherical harmonic lighting: The gritty details. In Archives of the game developers conference, 2003.