Learning to Optimize 3D Gaussian Splatting

Haofei Xu

ETH Zurich and University of Tübingen

Helsinki, 09.03.2026

3D Gaussian Splatting (3DGS)

Kerbl et al. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. SIGGRAPH 2023

Existing 3DGS Methods

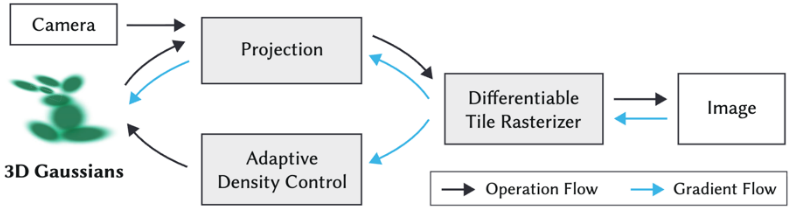

Iterative Optimization

Images

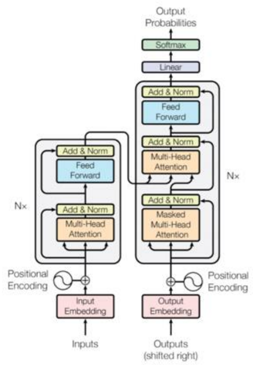

Feed-Forward Inference

Kerbl et al. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. SIGGRAPH 2023

Feed-Forward 3DGS

- Feed-forward prediction of 3D Gaussian parameters, without optimization

Early Works

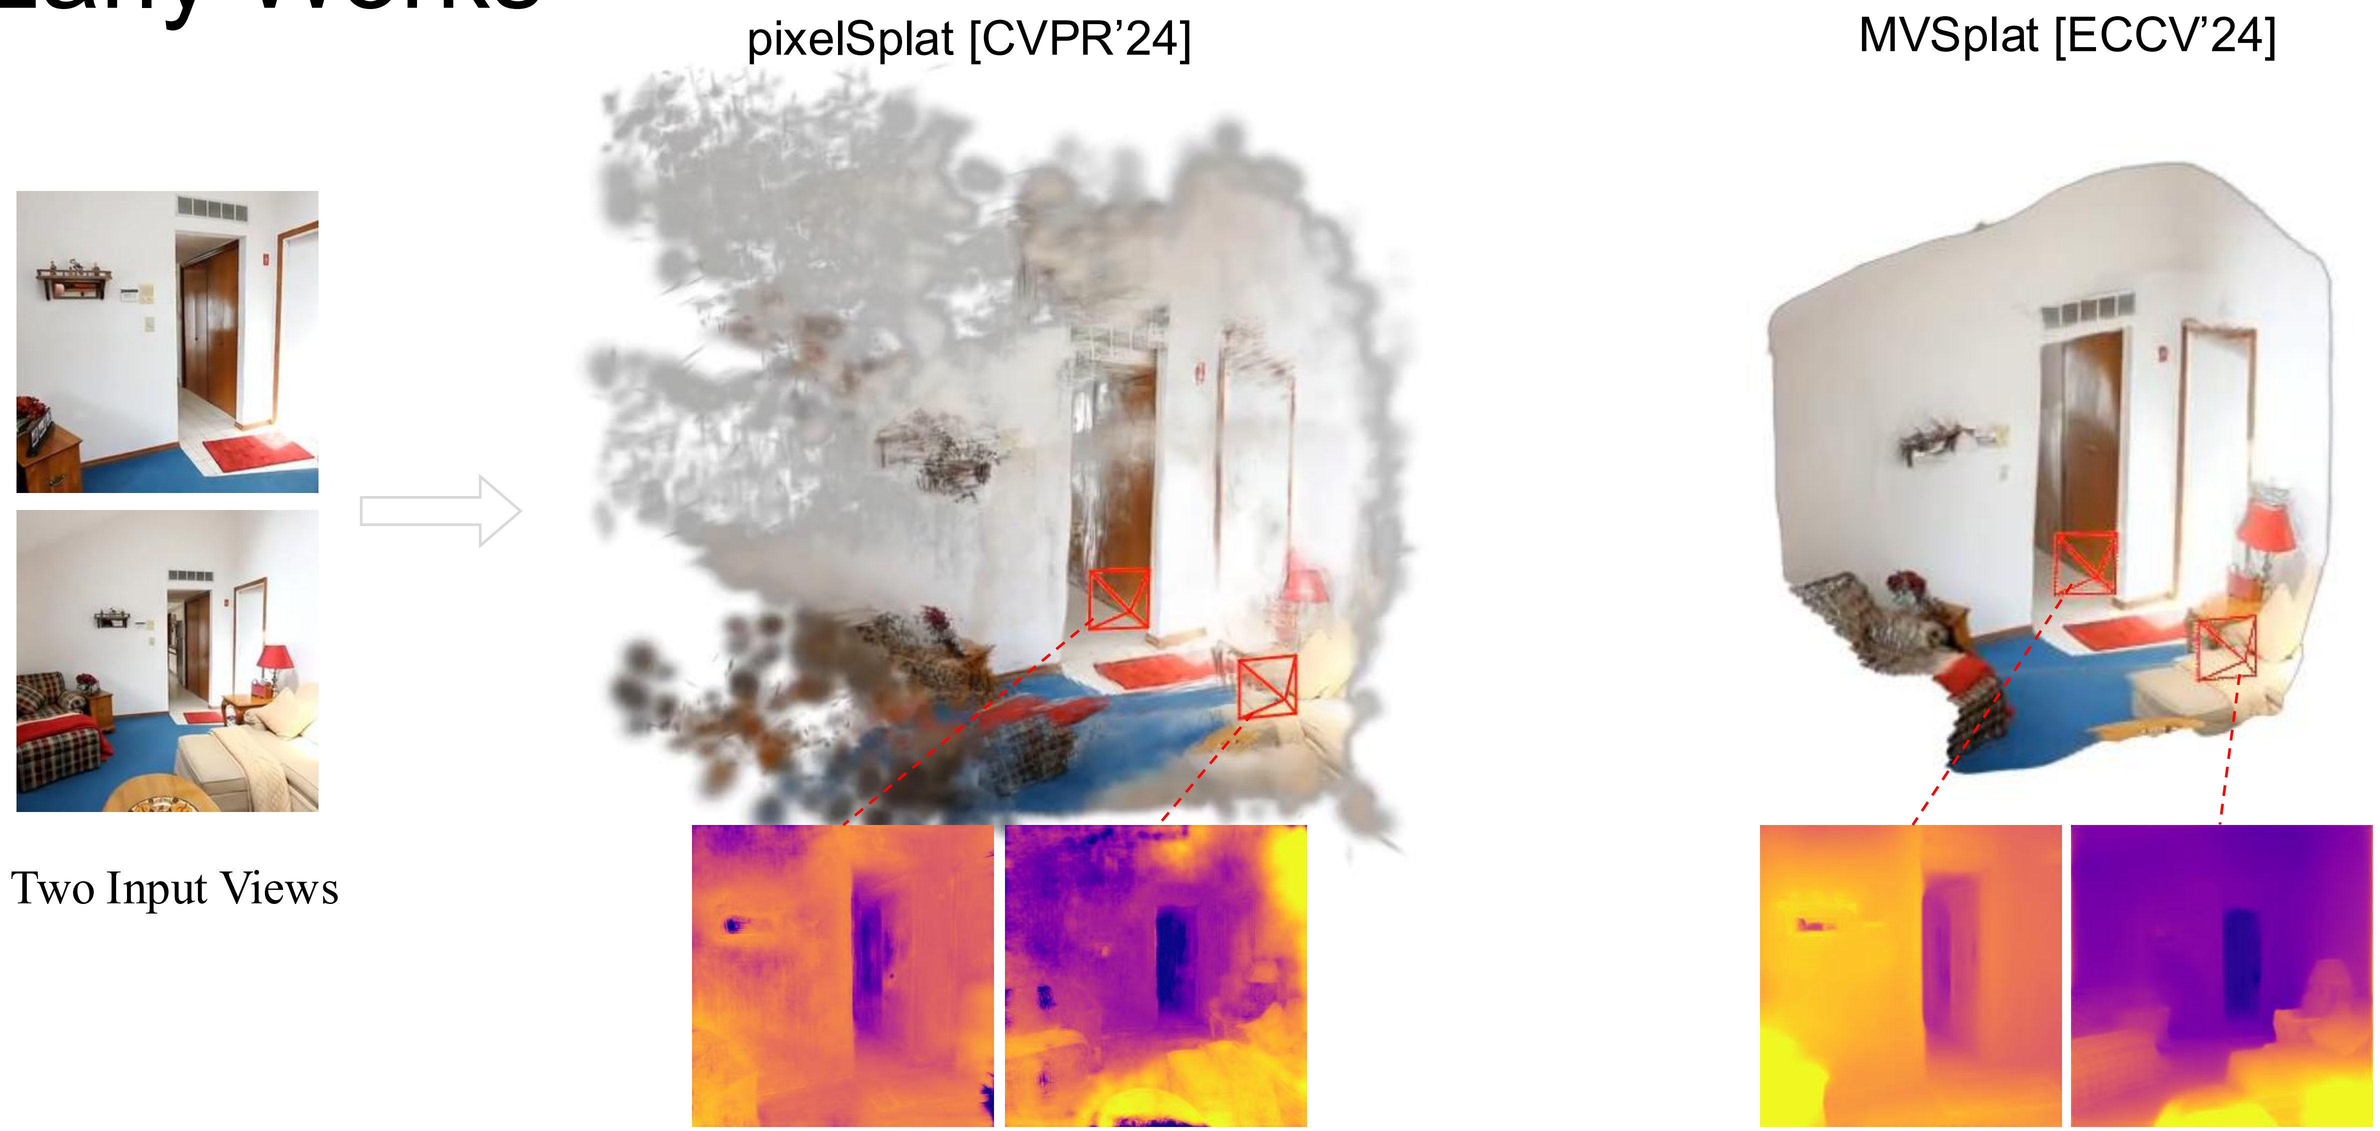

Chen et al. MVSplat: Efficient 3D Gaussian Splatting from Sparse Multi-View Images. ECCV 2024 (Oral)

MVSplat vs. pixelSplat

- 10x fewer parameters, 2x faster speed, +0.5dB PSNR

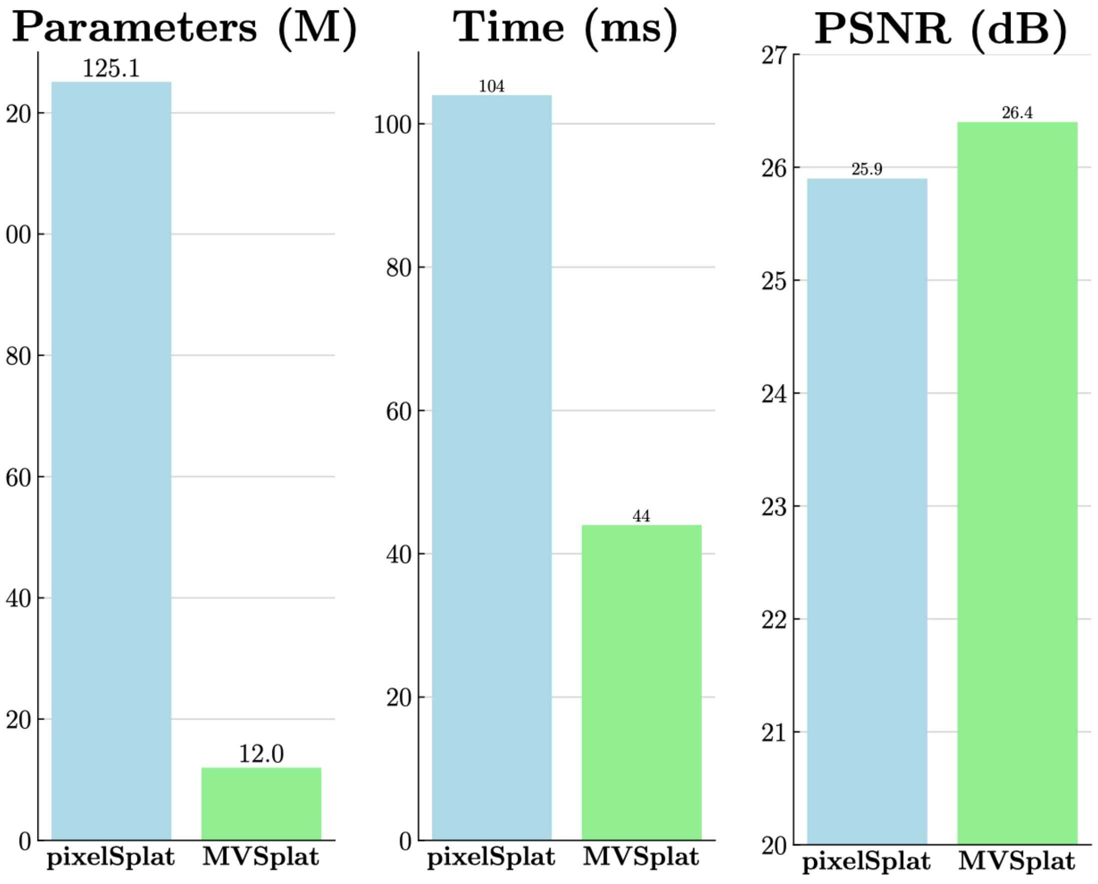

Chen et al. MVSplat: Efficient 3D Gaussian Splatting from Sparse Multi-View Images. ECCV 2024 (Oral)

MVSplat

- Multi-View Stereo (MVS) + Gaussian Splatting

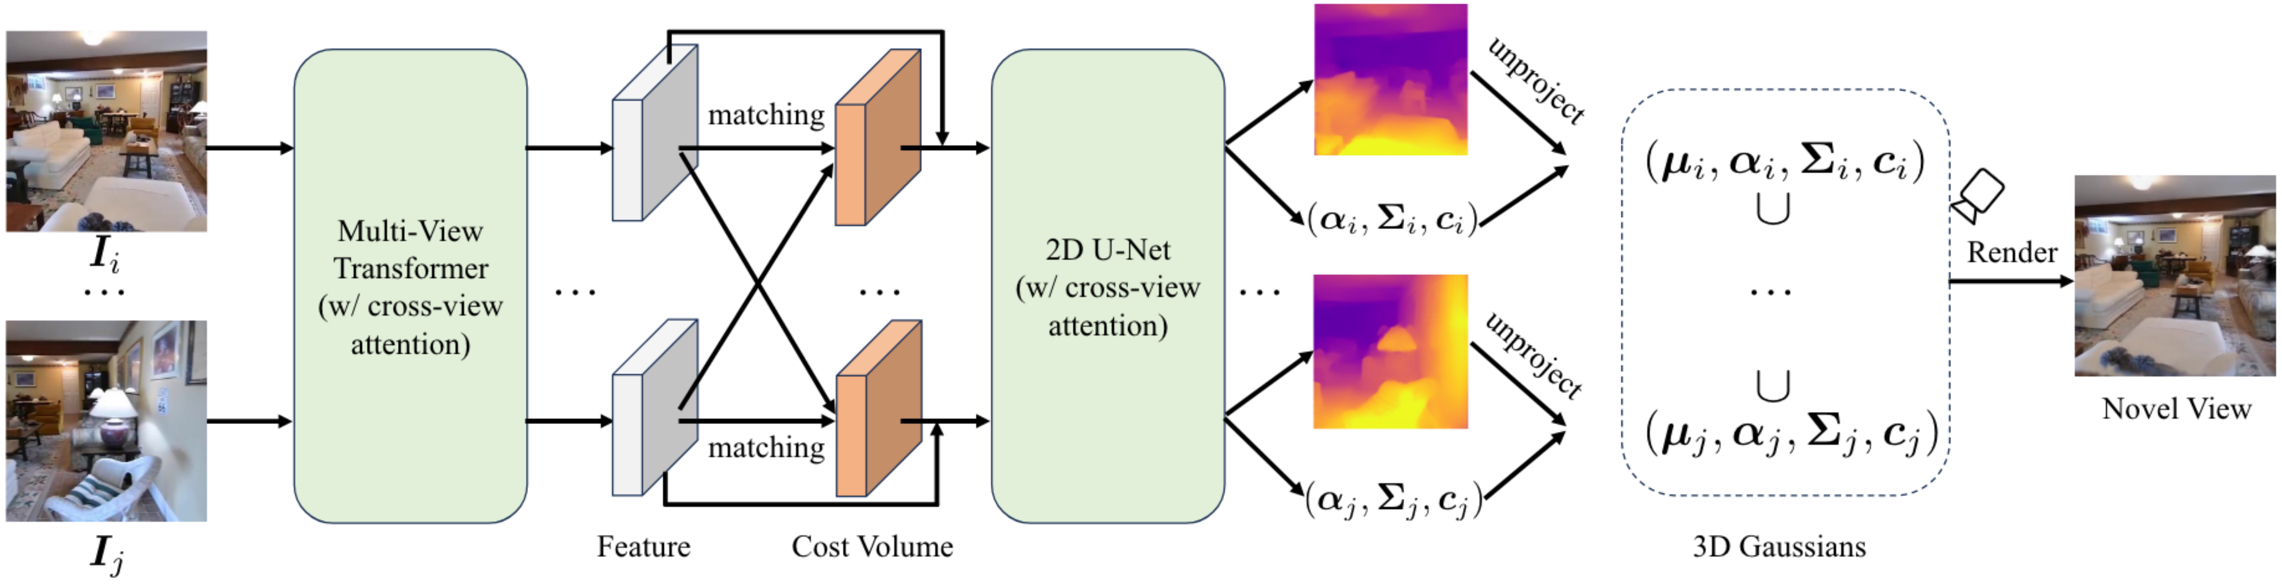

Xu et al. Unifying Flow, Stereo and Depth Estimation. TPAMI 2023

Chen et al. MVSplat: Efficient 3D Gaussian Splatting from Sparse Multi-View Images. ECCV 2024 (Oral)

DepthSplat

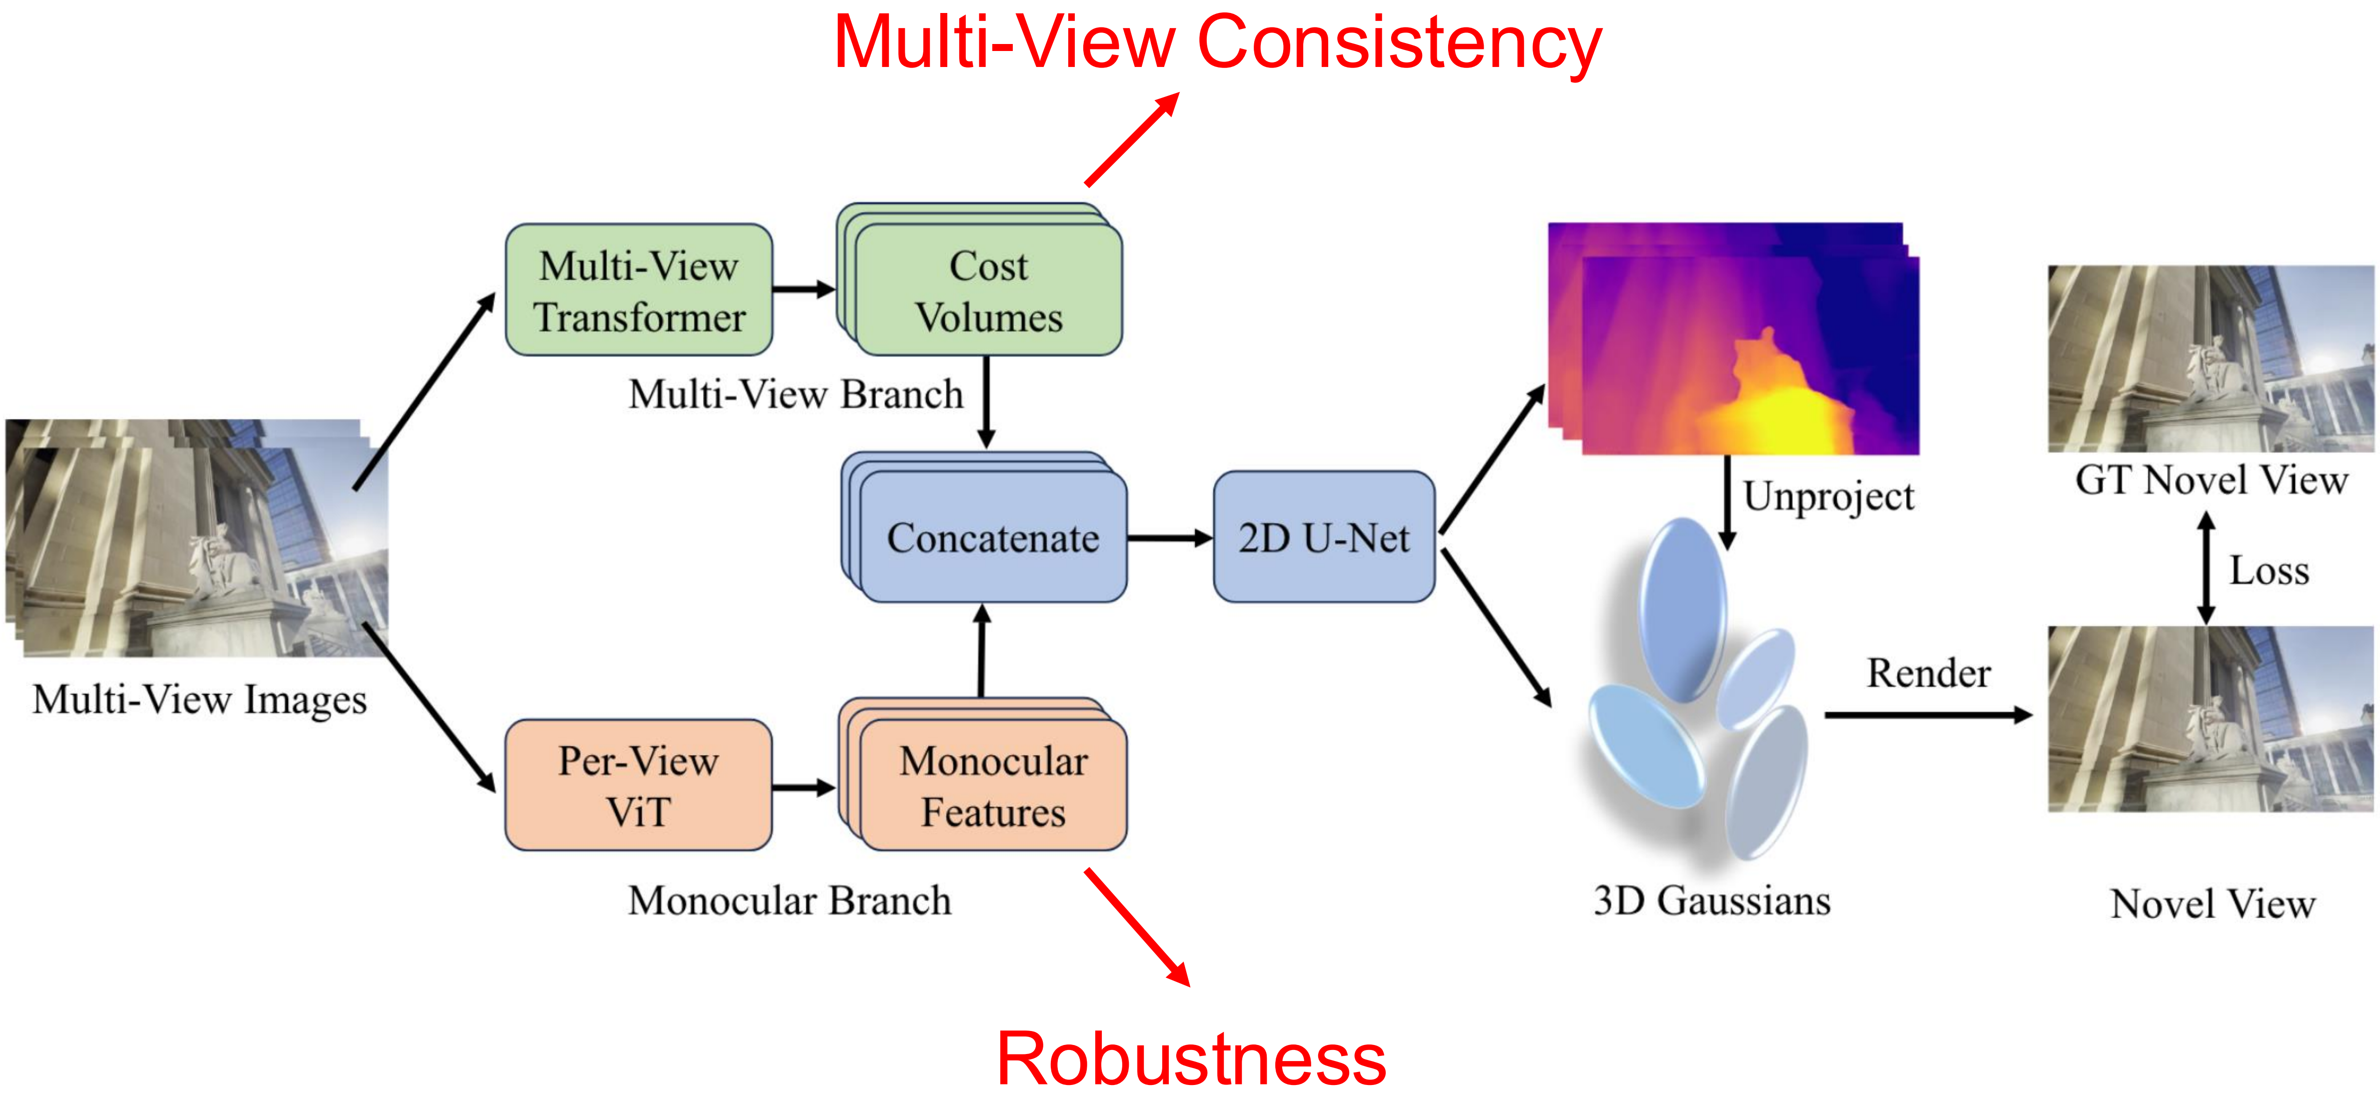

Xu et al. DepthSplat: Connecting Gaussian Splatting and Depth. CVPR 2025

Feed-Forward Scene Synthesis

DepthSplat

6 Input Views

Rendering

Xu et al. DepthSplat: Connecting Gaussian Splatting and Depth. CVPR 2025

Feed-Forward Scene Synthesis

DepthSplat

12 Input Views

Rendering

Xu et al. DepthSplat: Connecting Gaussian Splatting and Depth. CVPR 2025

Feed-Forward vs. Optimization

+ Efficient, feed-forward

+ Data-driven priors, scalable

- Data-dependent, no guarantee for generalization

Feed-Forward

- Slow, many iterative steps

- No priors, not scalable

+ No generalization issue, directly optimize on the test data

Optimization

Feed-Forward vs. Optimization

Feed-Forward	Optimization
+ Efficient, feed-forward	- Slow, many iterative steps
+ Data-driven priors, scalable	- No priors, not scalable
- Data-dependent, no guarantee for generalization	+ No generalization issue, directly optimize on the test data

Learning to Optimize for 3DGS

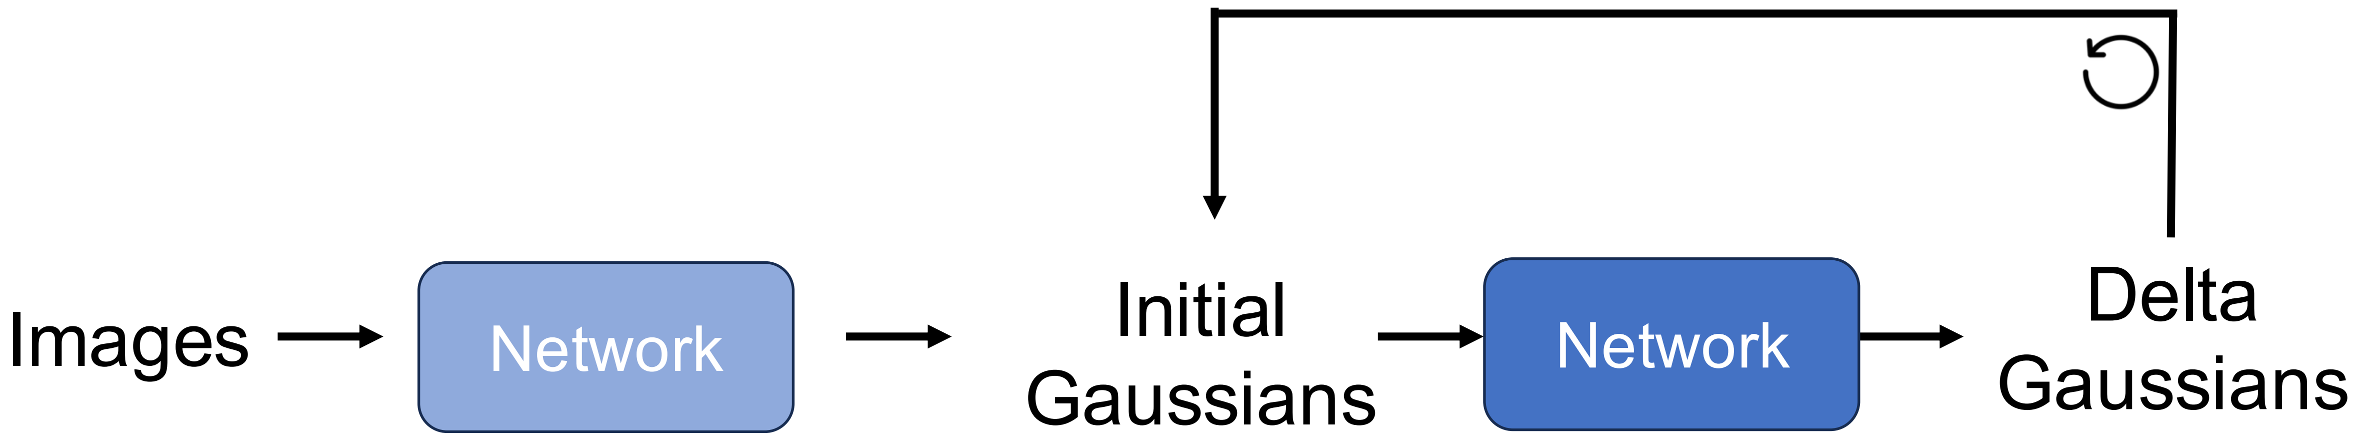

Xu, Barath, Geiger, Pollefeys. ReSplat: Learning Recurrent Gaussian Splatting. arXiv 2026

Recurrent Update from Feedback

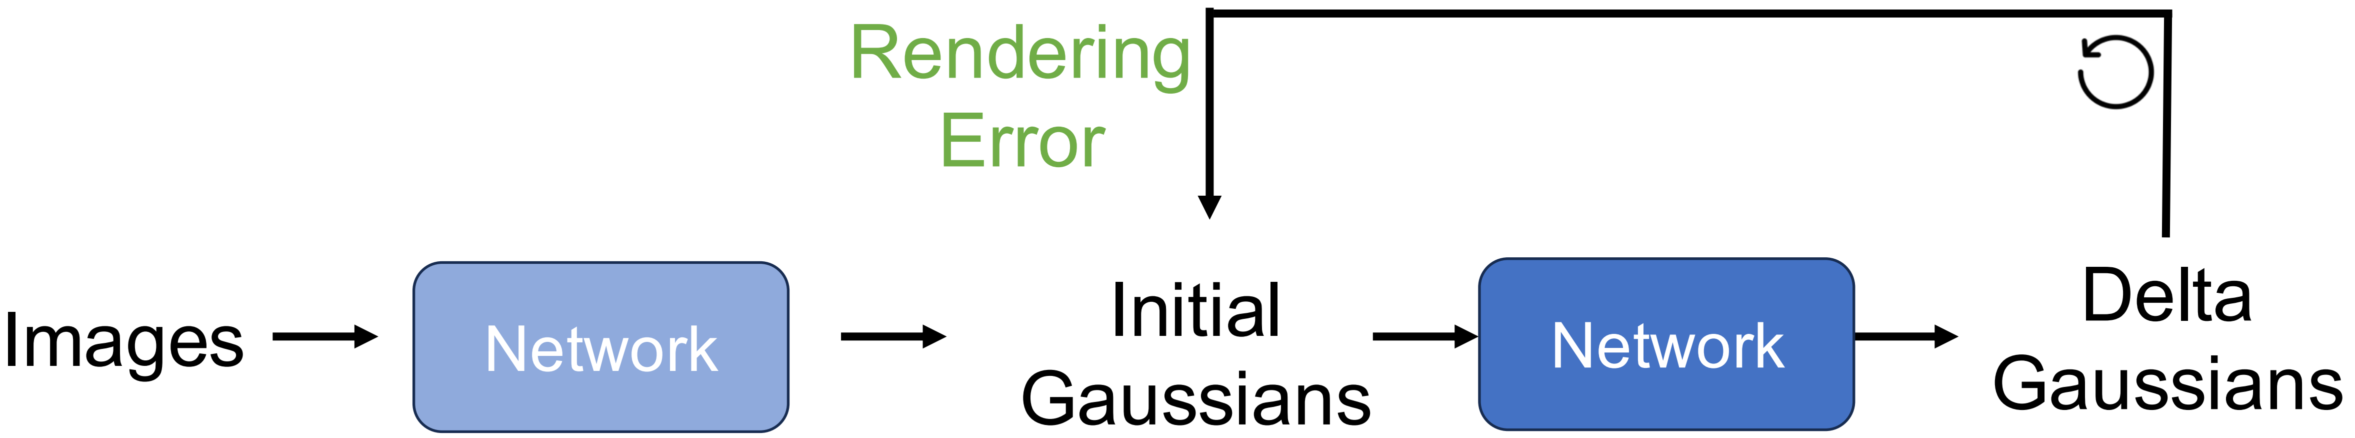

Xu, Barath, Geiger, Pollefeys. ReSplat: Learning Recurrent Gaussian Splatting. arXiv 2026

ReSplat vs. Standard 3DGS Optimization

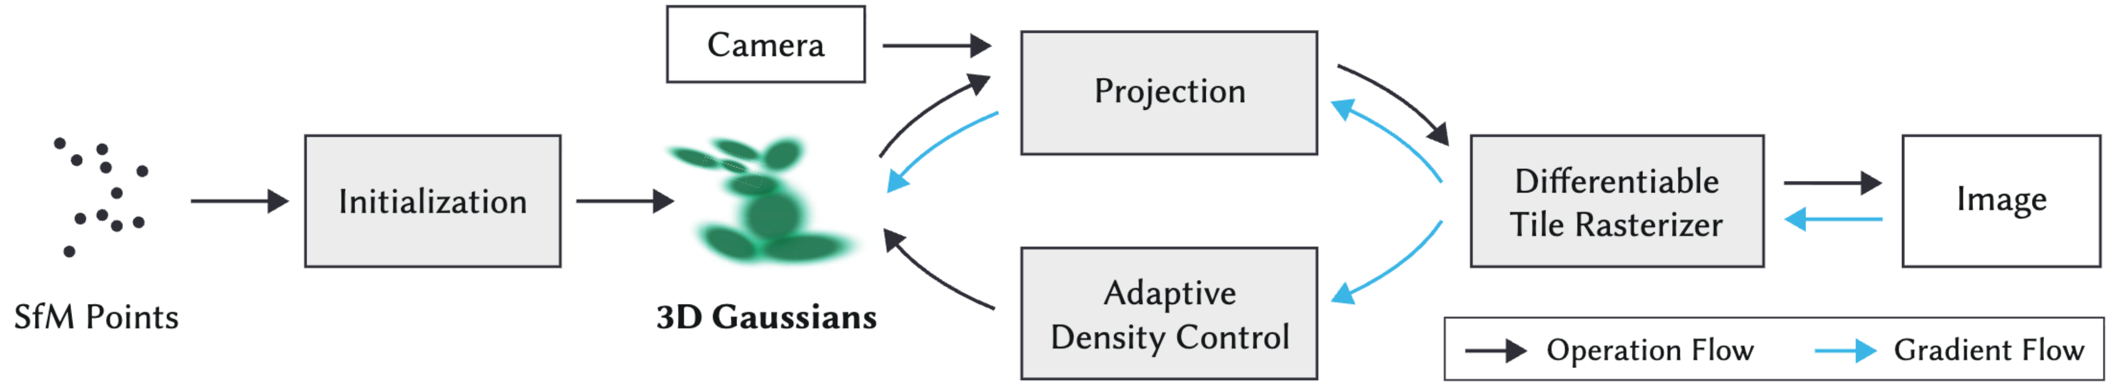

Standard 3DGS Optimization

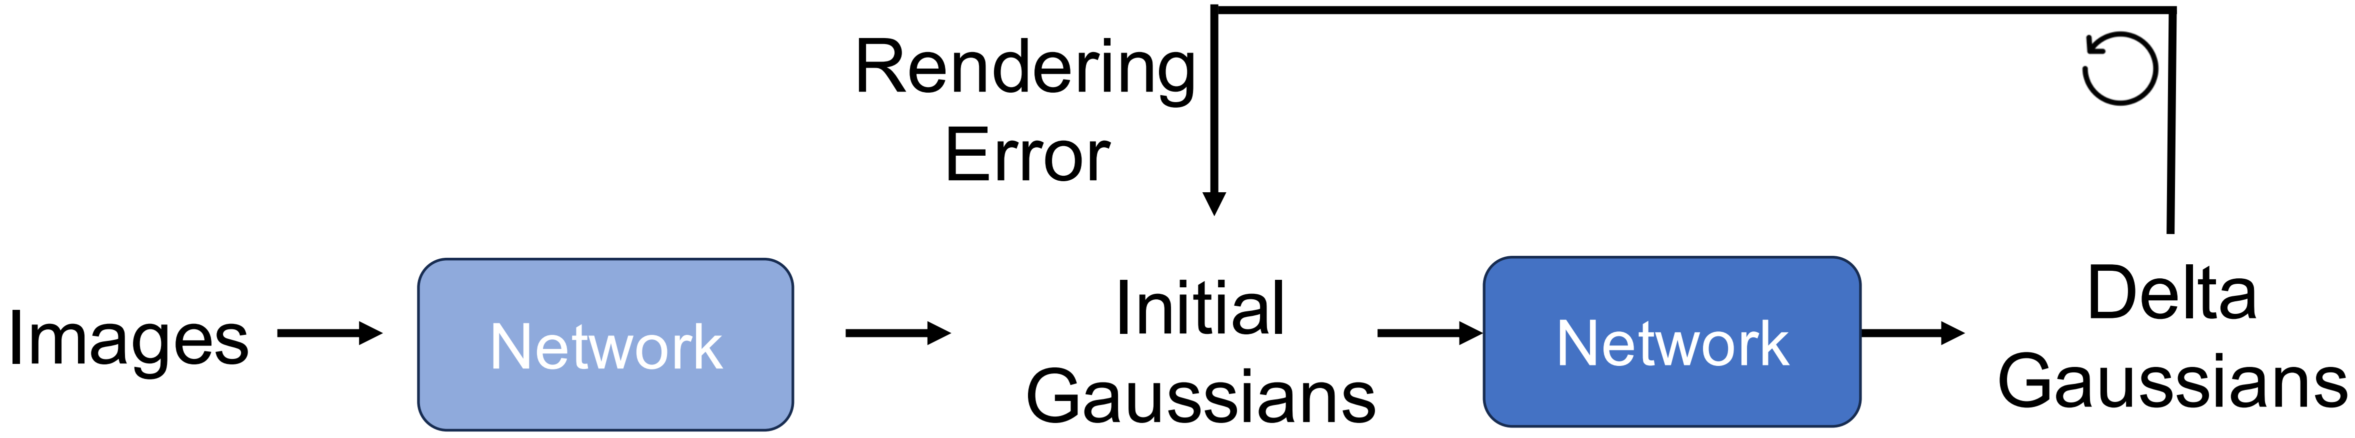

ReSplat

ReSplat vs. Standard 3DGS Optimization

ReSplat

- **Feed-forward inference**
- **Gradient-free update**

$$\theta_{t+1} = \theta_t + \mathrm{RNN}(\theta_t, \mathrm{Error}_t)$$

- **Converges with 4 iterations**

3DGS Optimization

- Per-scene optimization
- Gradient-based update

$$\theta_{t+1} = \theta_t - \alpha \nabla_\theta L(\theta_t)$$

- Thousands of iterations

Results

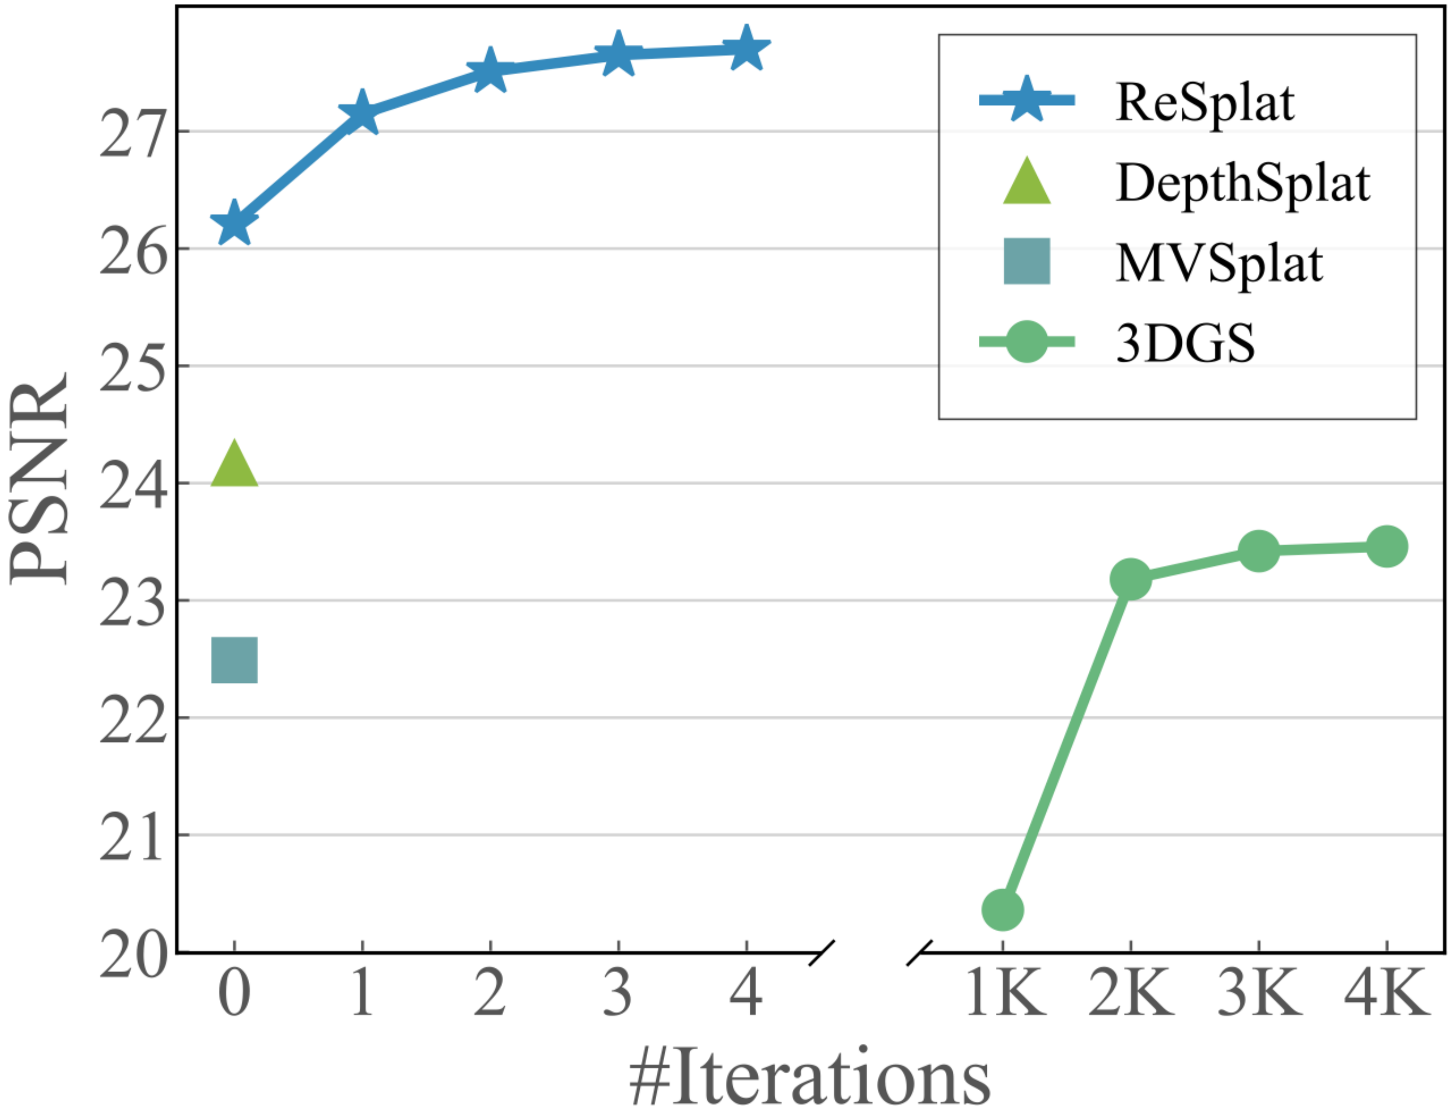

Results for 8 input views (512x960)

Initialization: Compact Gaussian Reconstruction

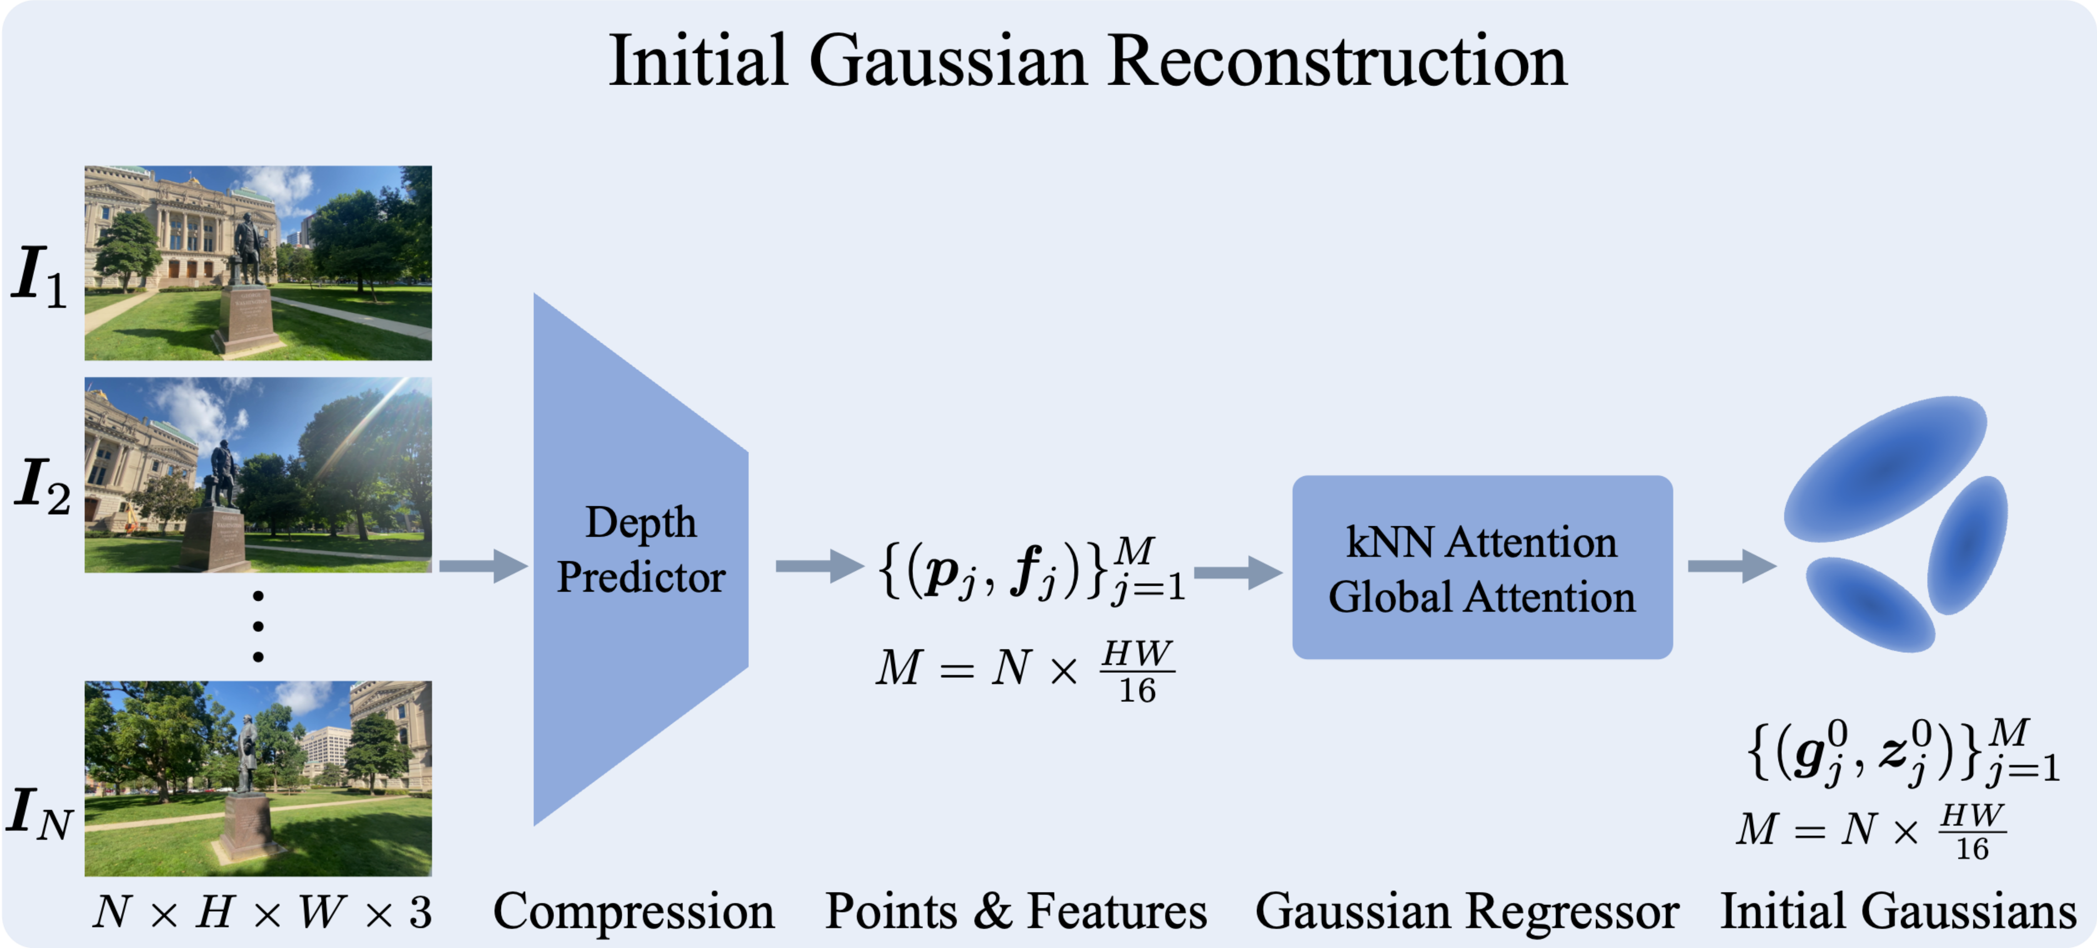

- 16x fewer Gaussians than previous pixel-aligned methods

Ablation of the Initial Model

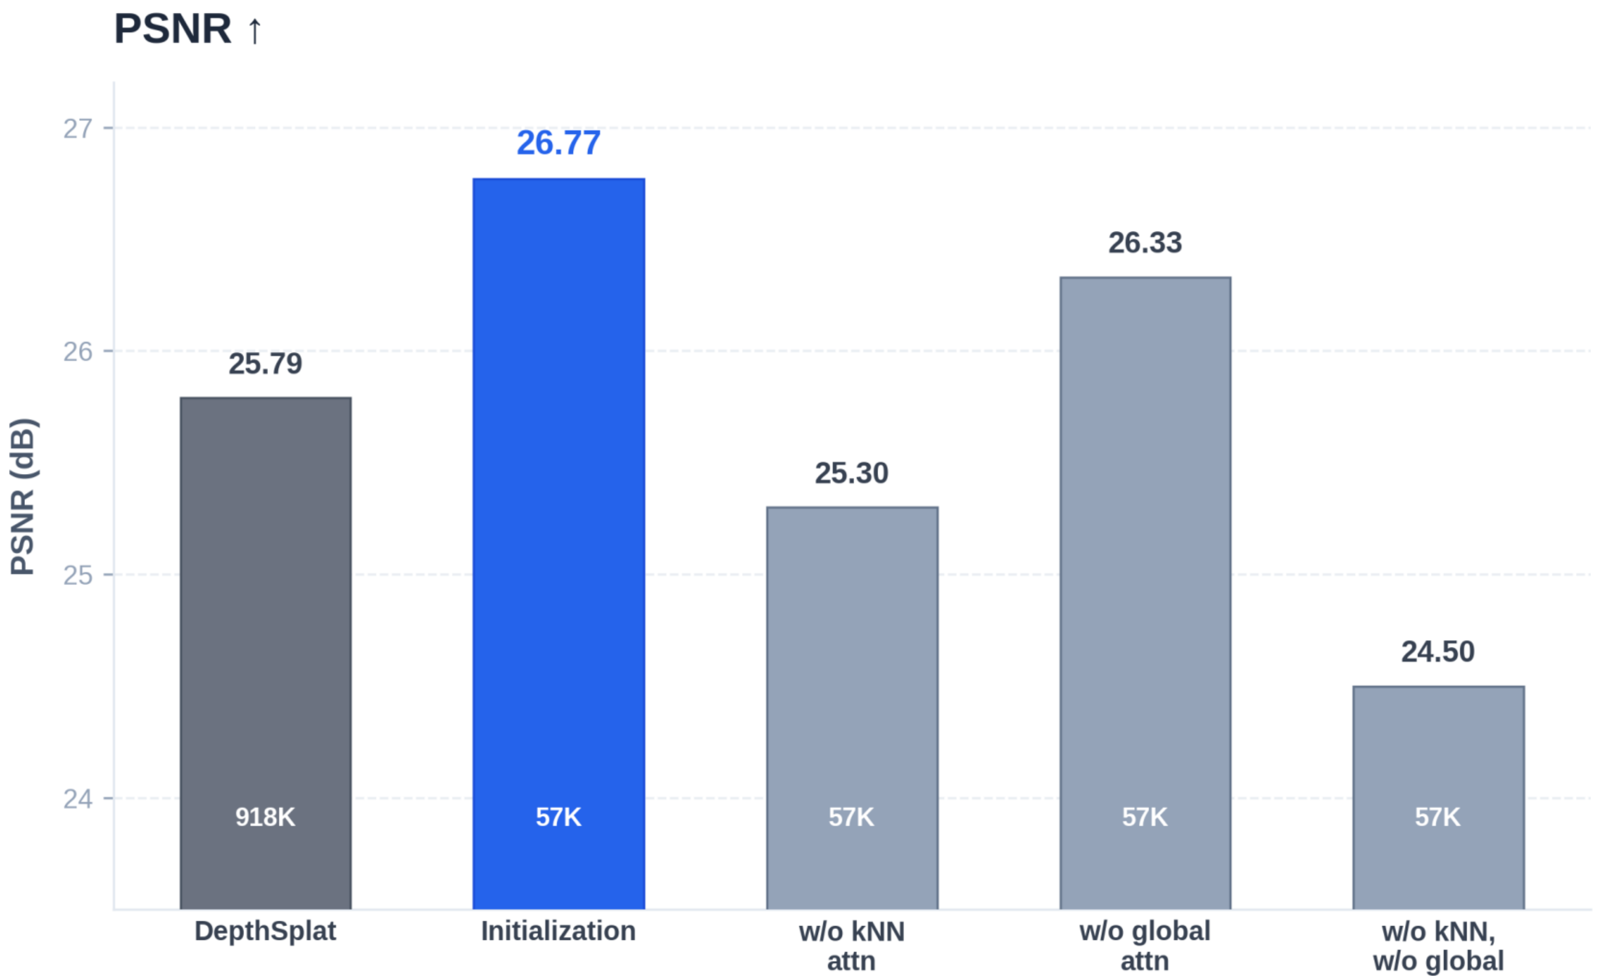

- Key: modeling 3D context with kNN and global attention

Recurrent Gaussian Update

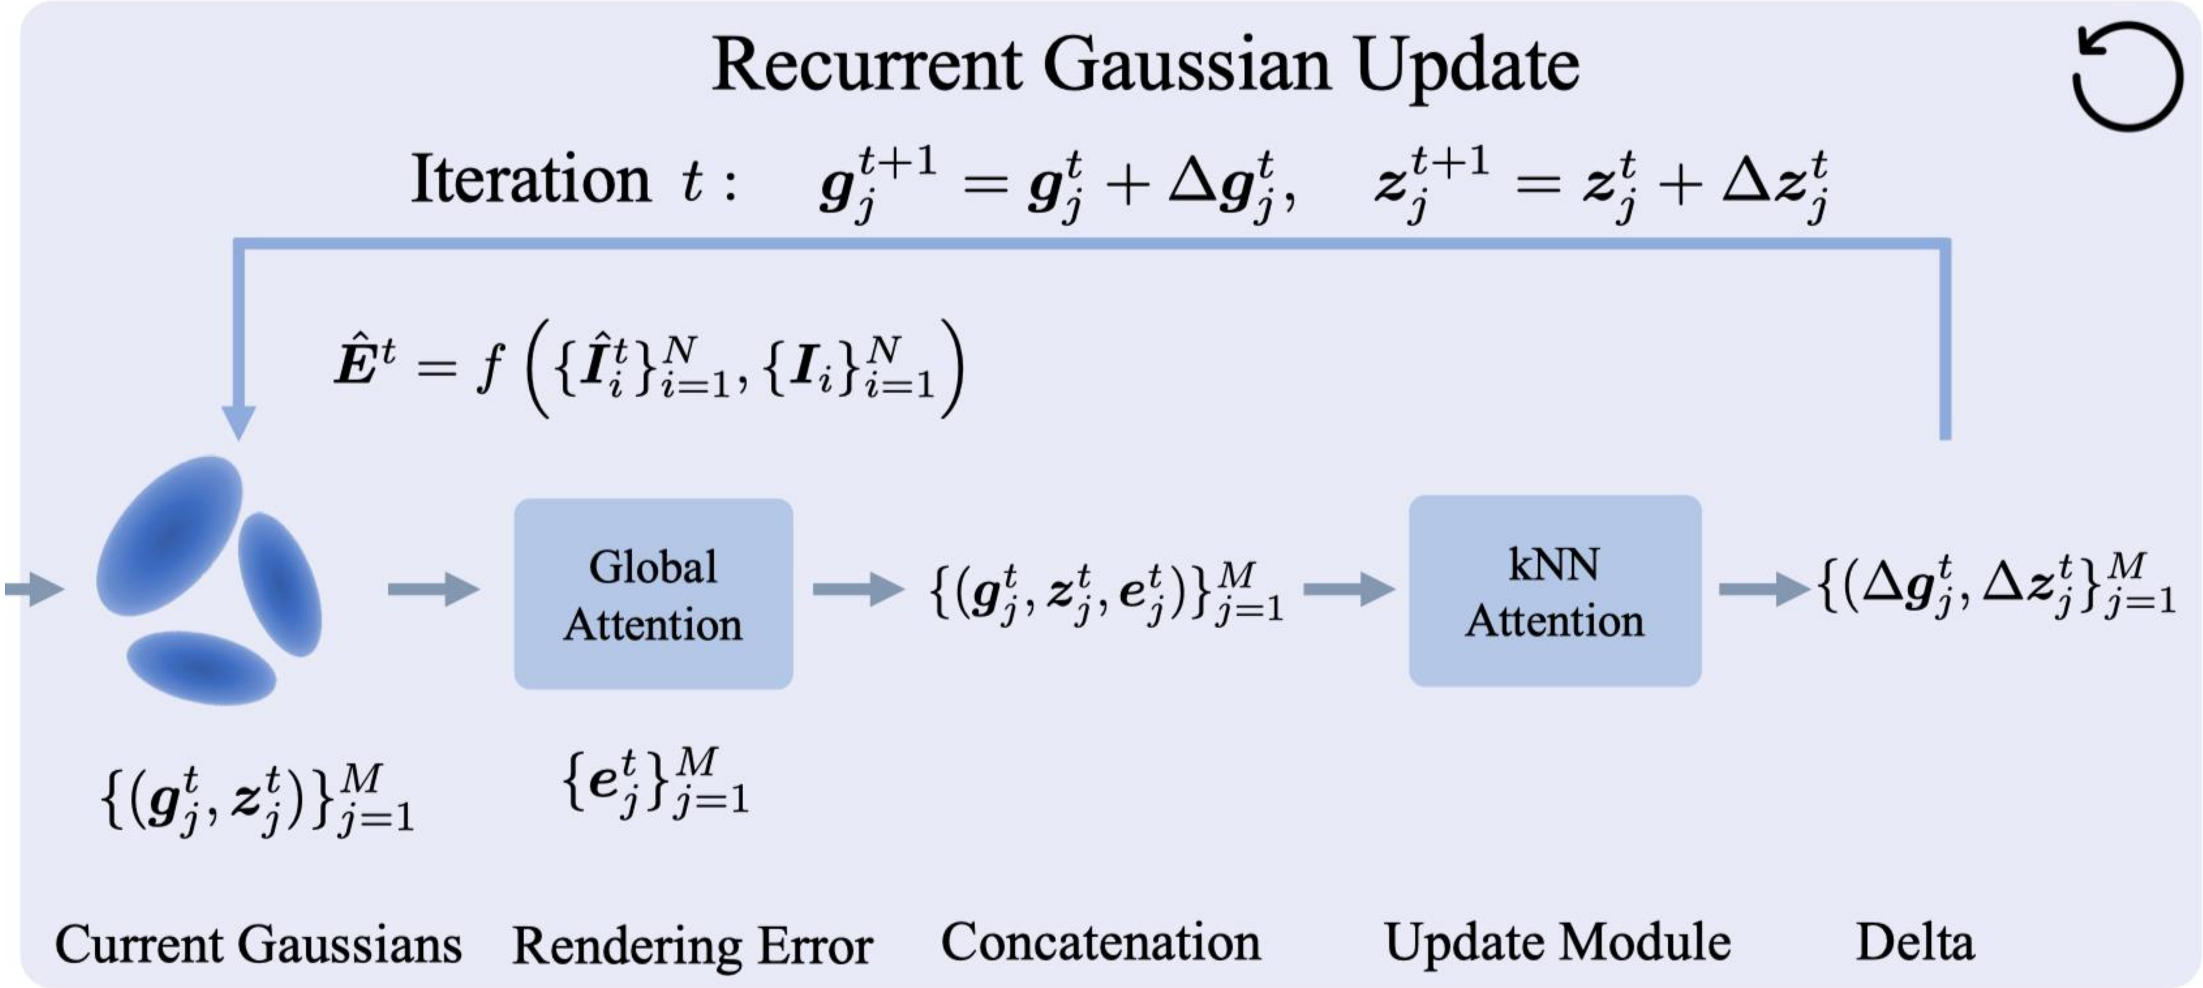

- Converges with 4 iterations

Ablation of the Recurrent Model

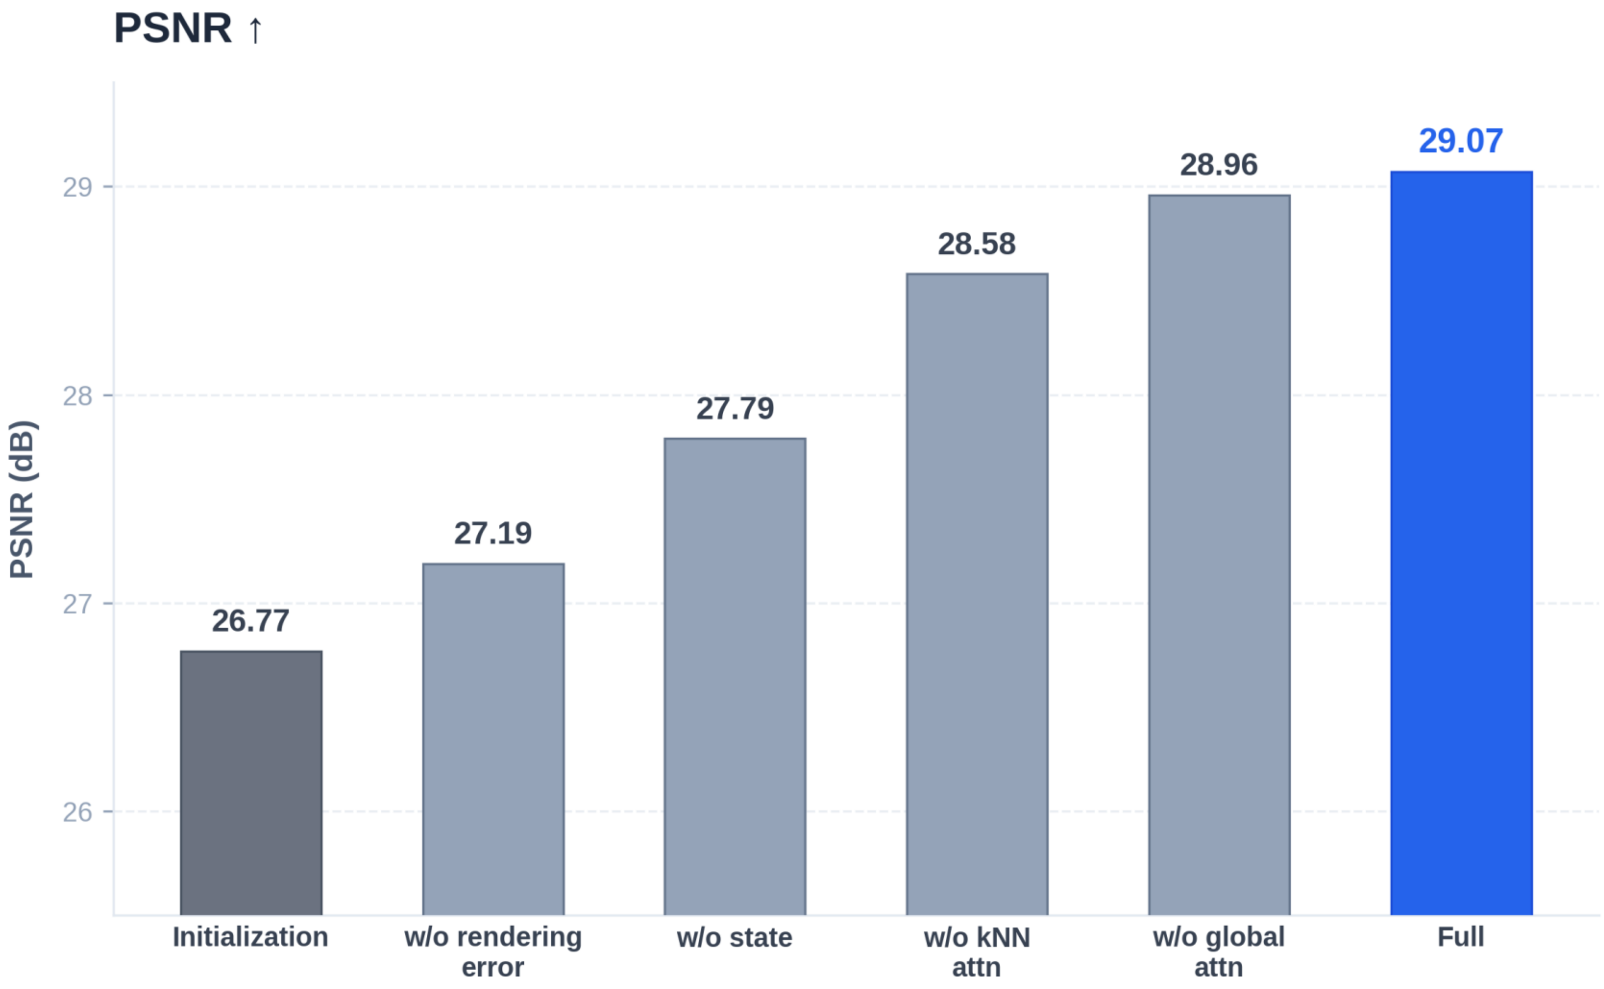

- Key: the rendering error

ReSplat

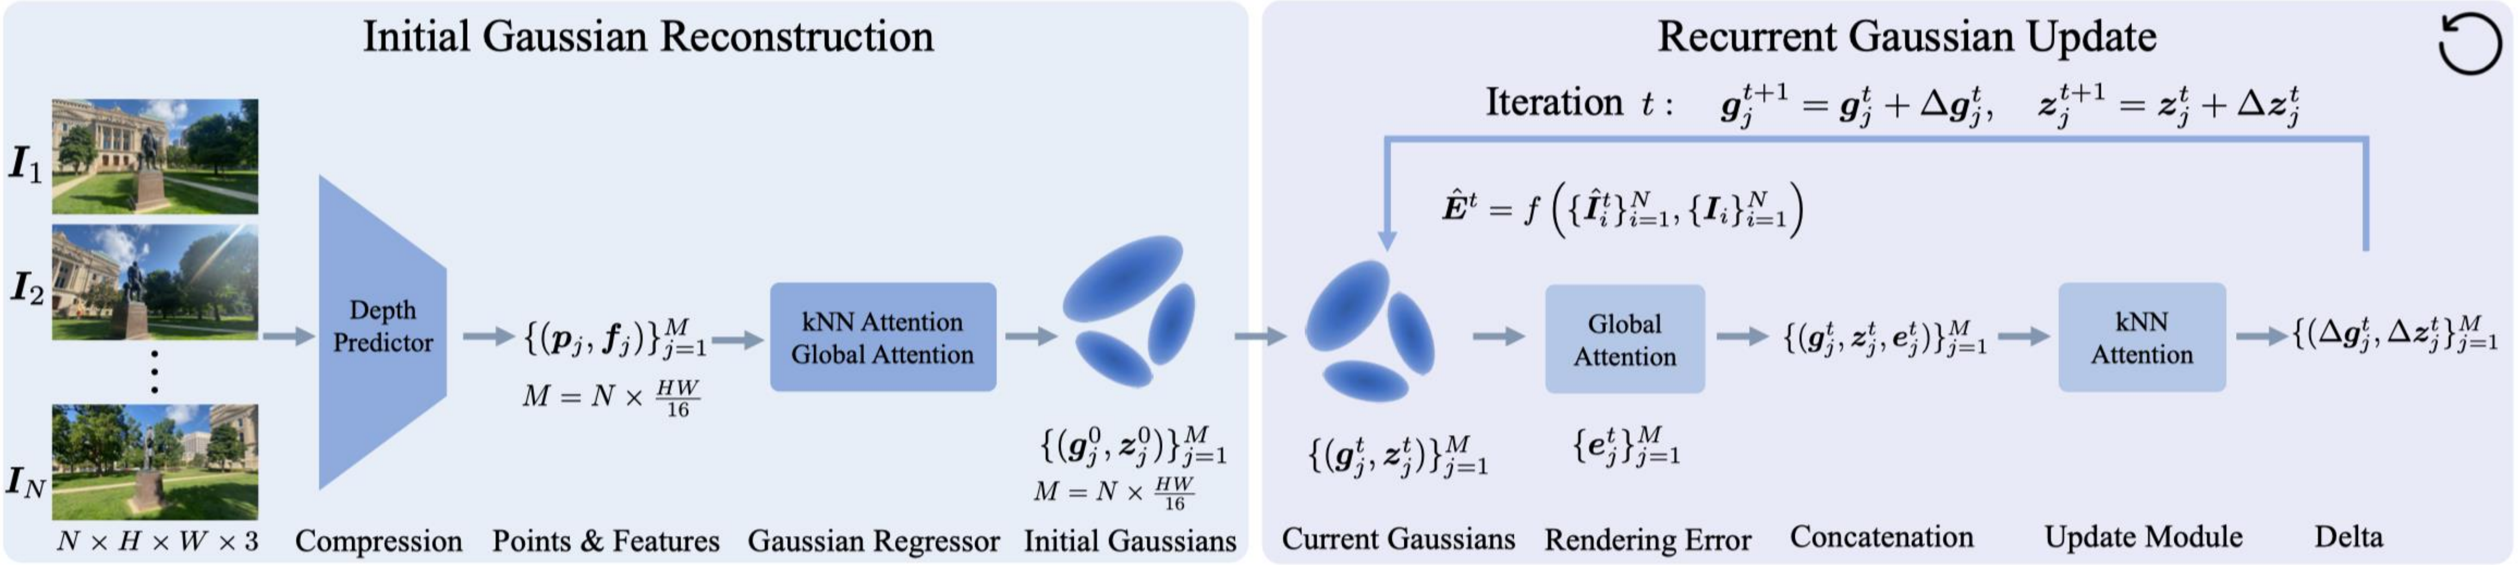

- Two-stage training: initial → recurrent

Comparisons

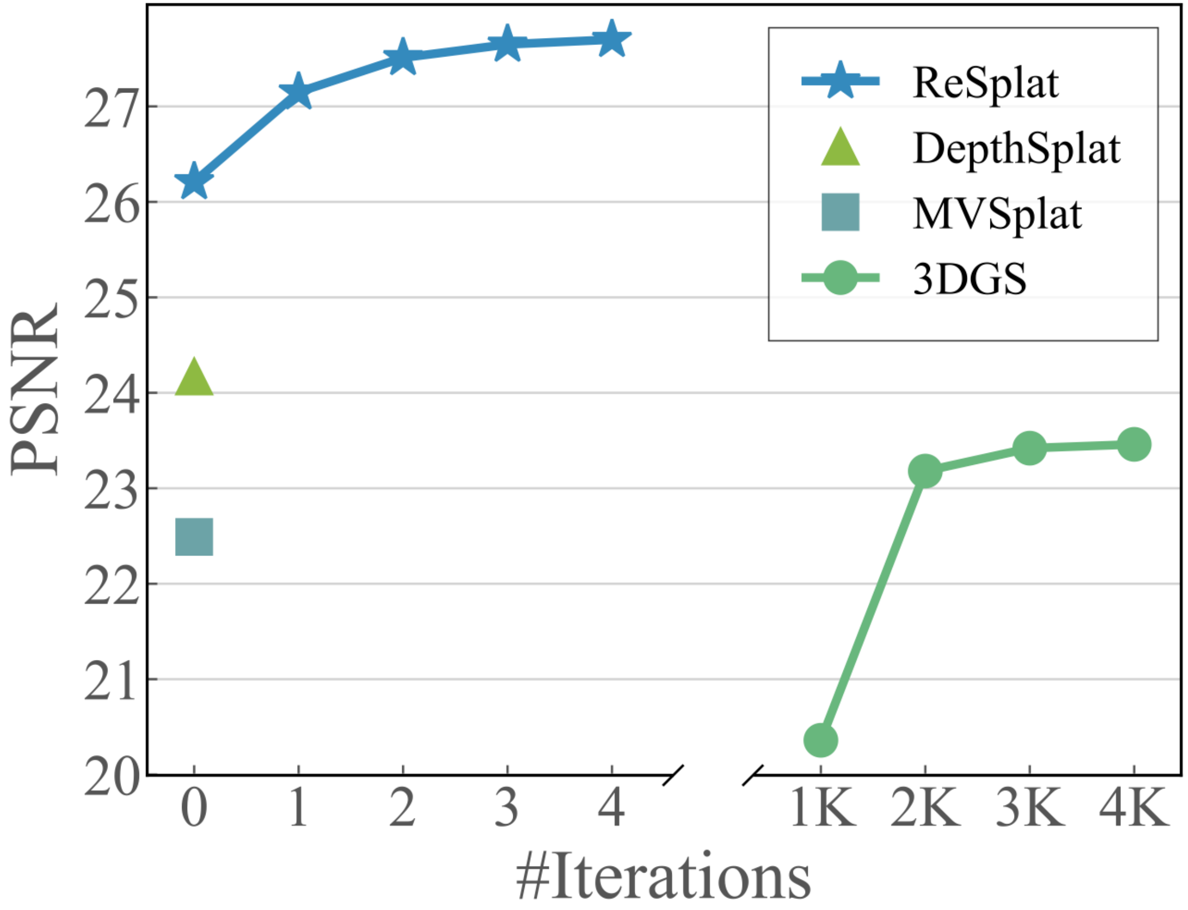

- 16x fewer Gaussians
- 4x faster rendering speed
- Reconstruction in <1 second
- 100x faster than 3DGS

Results for 8 input views (512x960)

Optimization vs. Feed-Forward Refinement

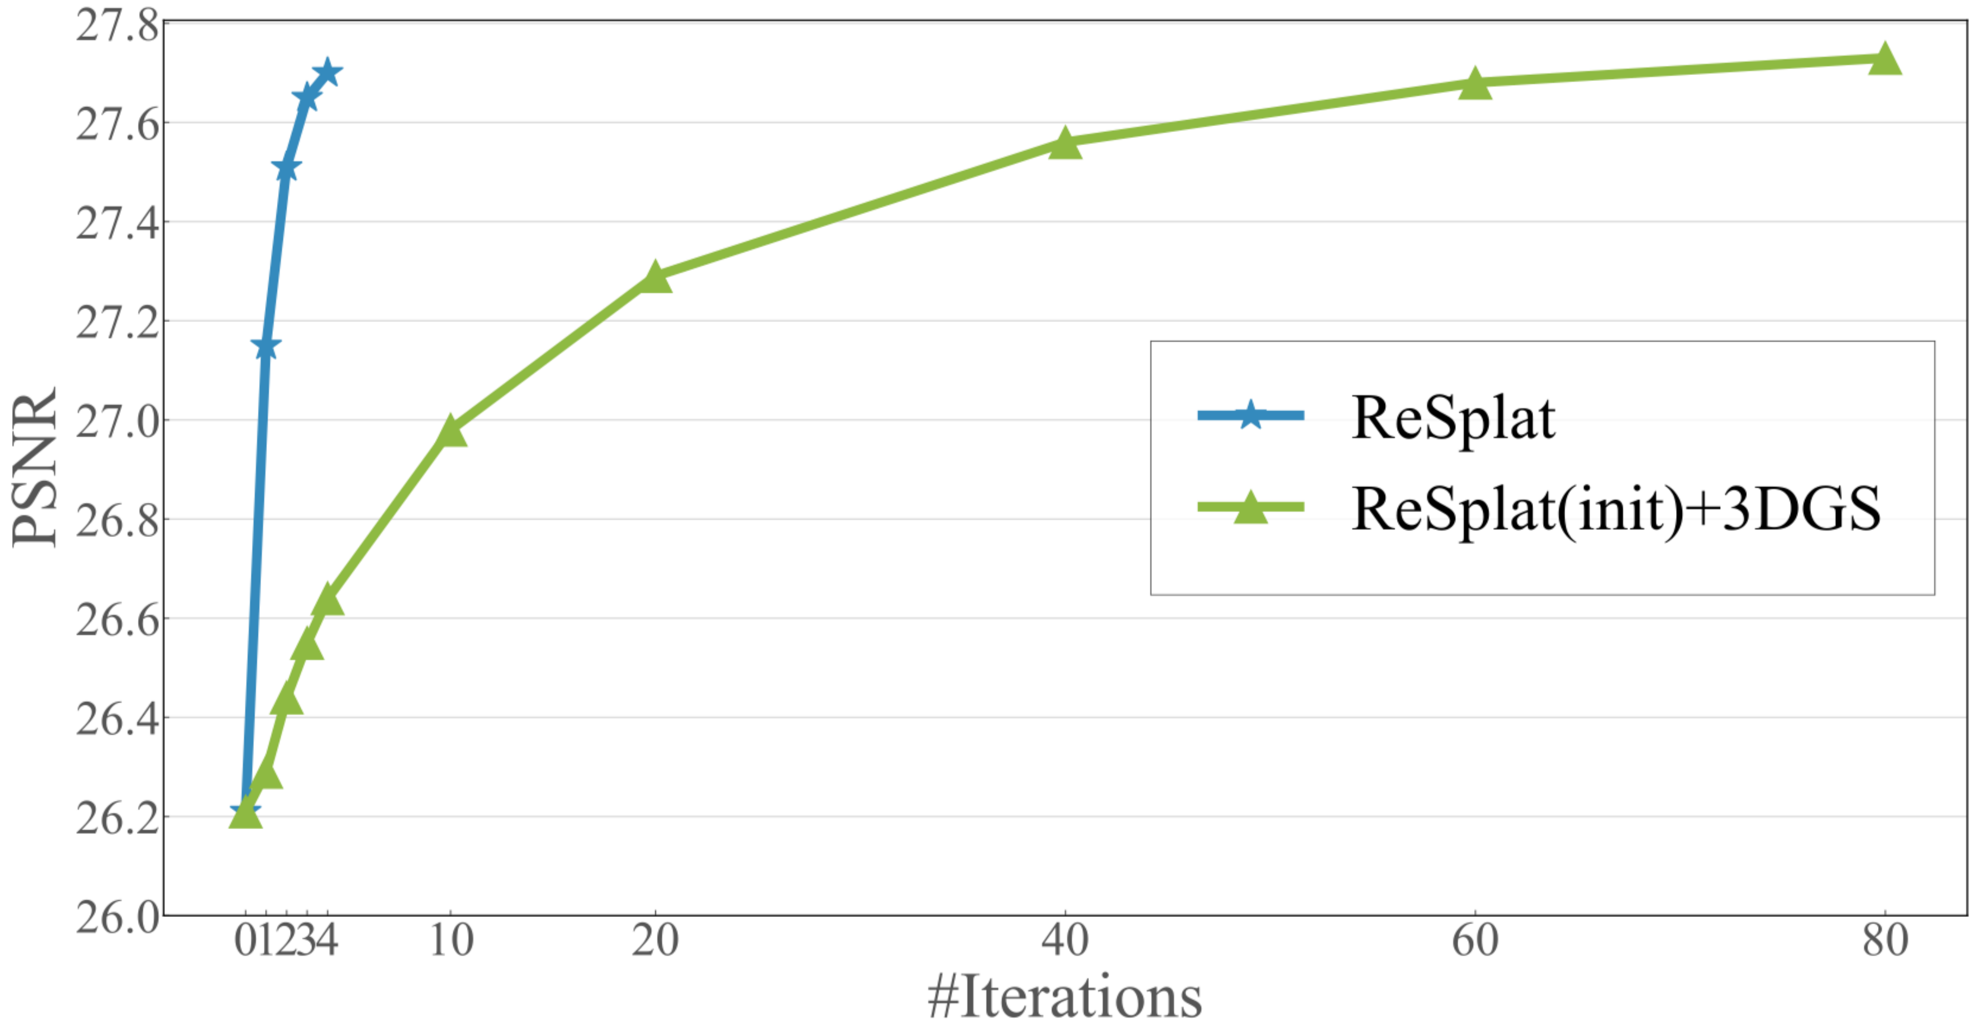

- ReSplat reduces the number of iterations by 20x (4 vs. 80)

Optimization vs. Feed-Forward Refinement

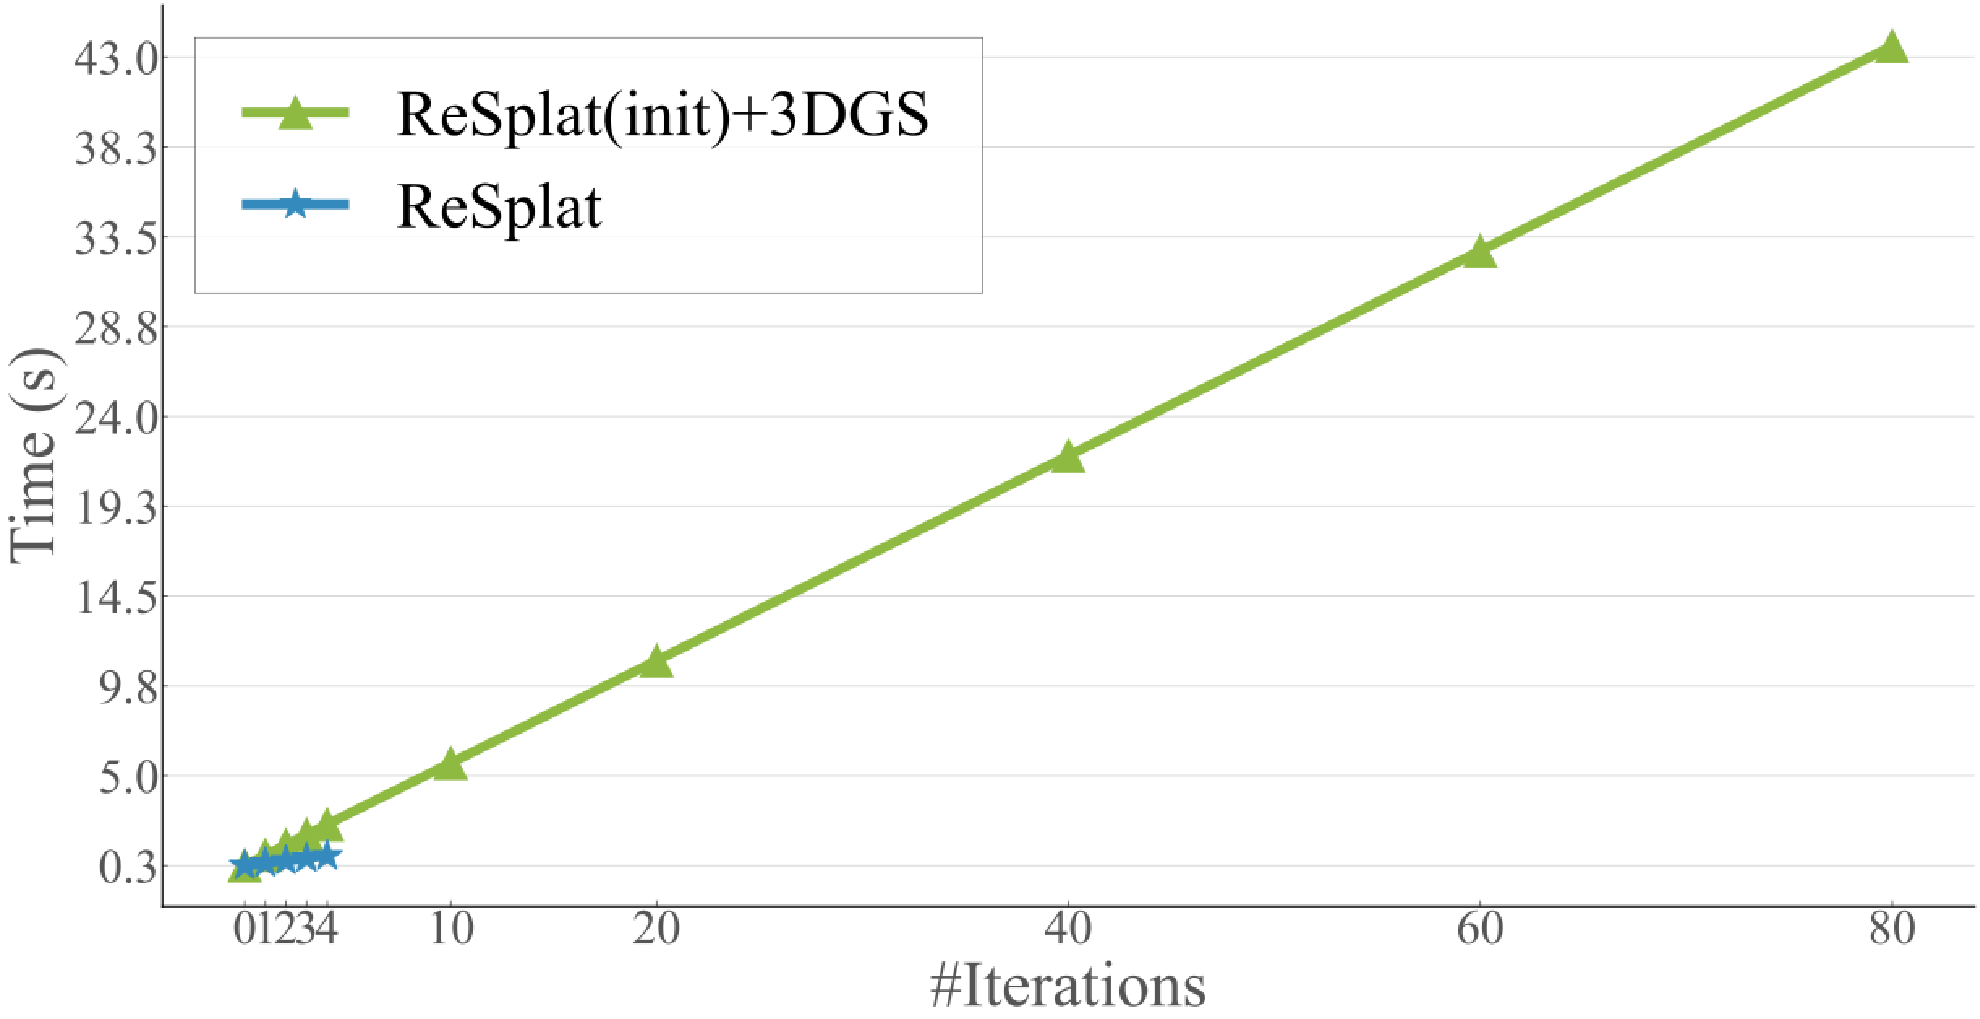

Optimization vs. Feed-Forward Refinement

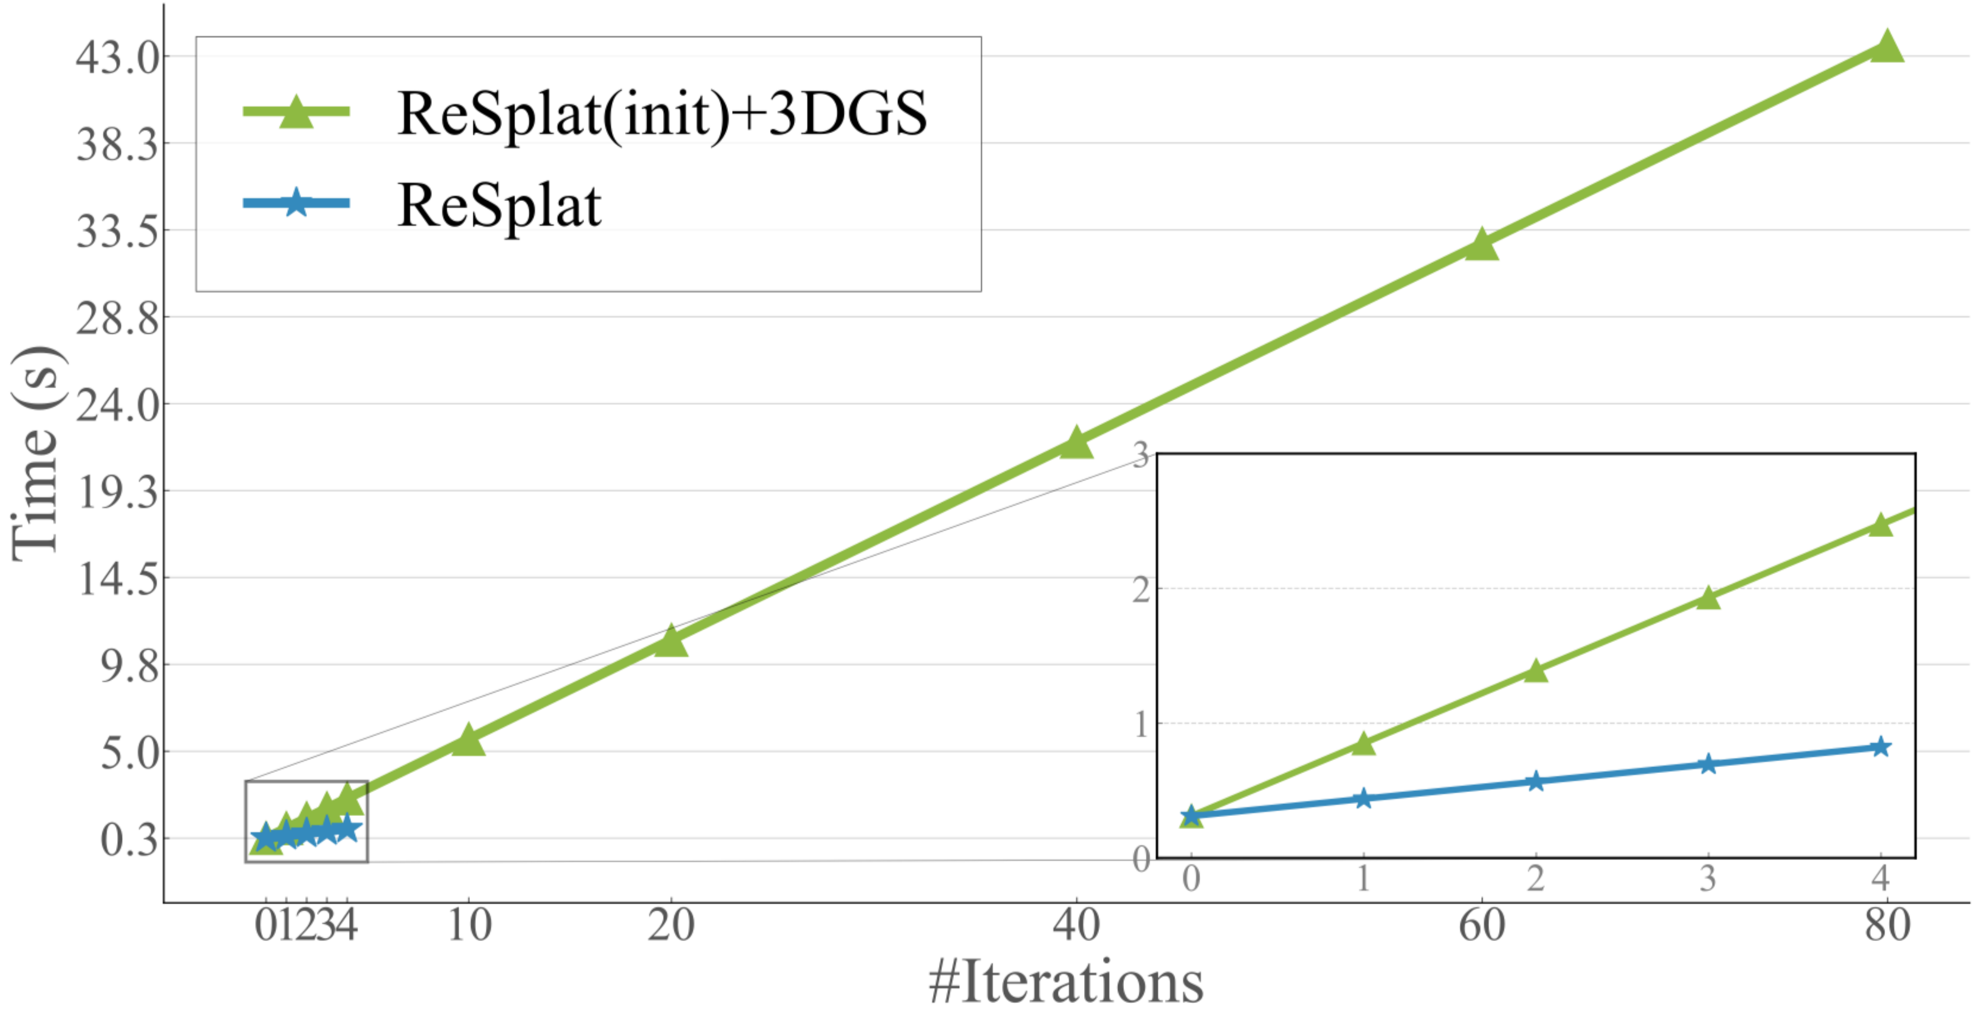

- ReSplat is 53x faster in reconstruction time

Single-Step vs. Recurrent Models

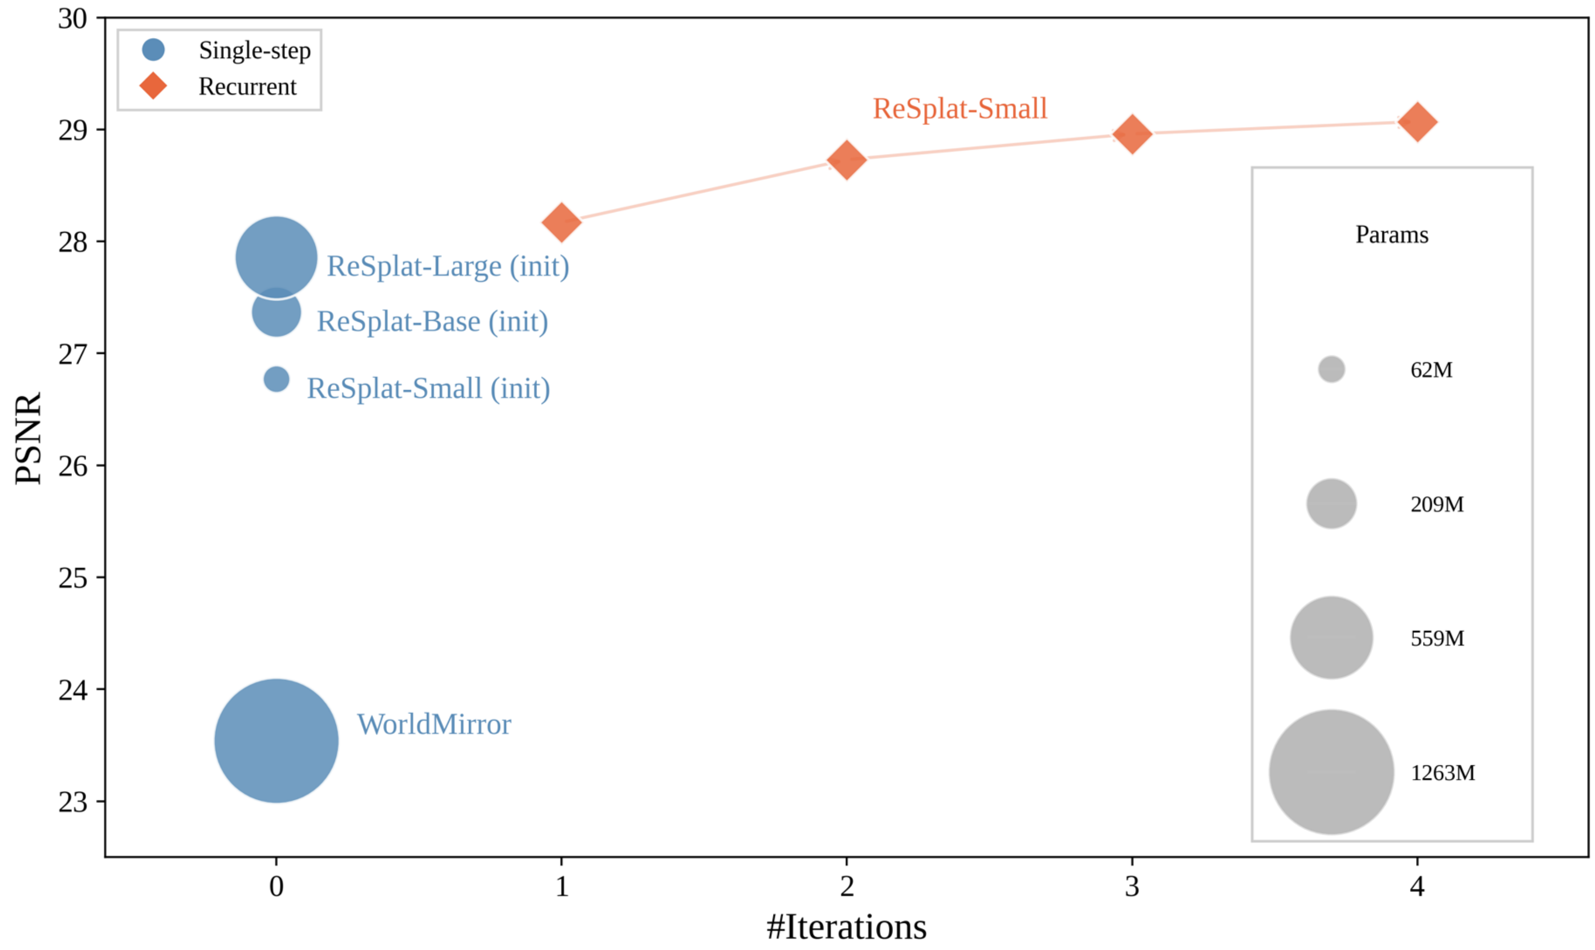

- +1.2dB PSNR
- 7x smaller model

- Recurrent network is much more effective than simply increasing parameters of single-step models

Different Initializations

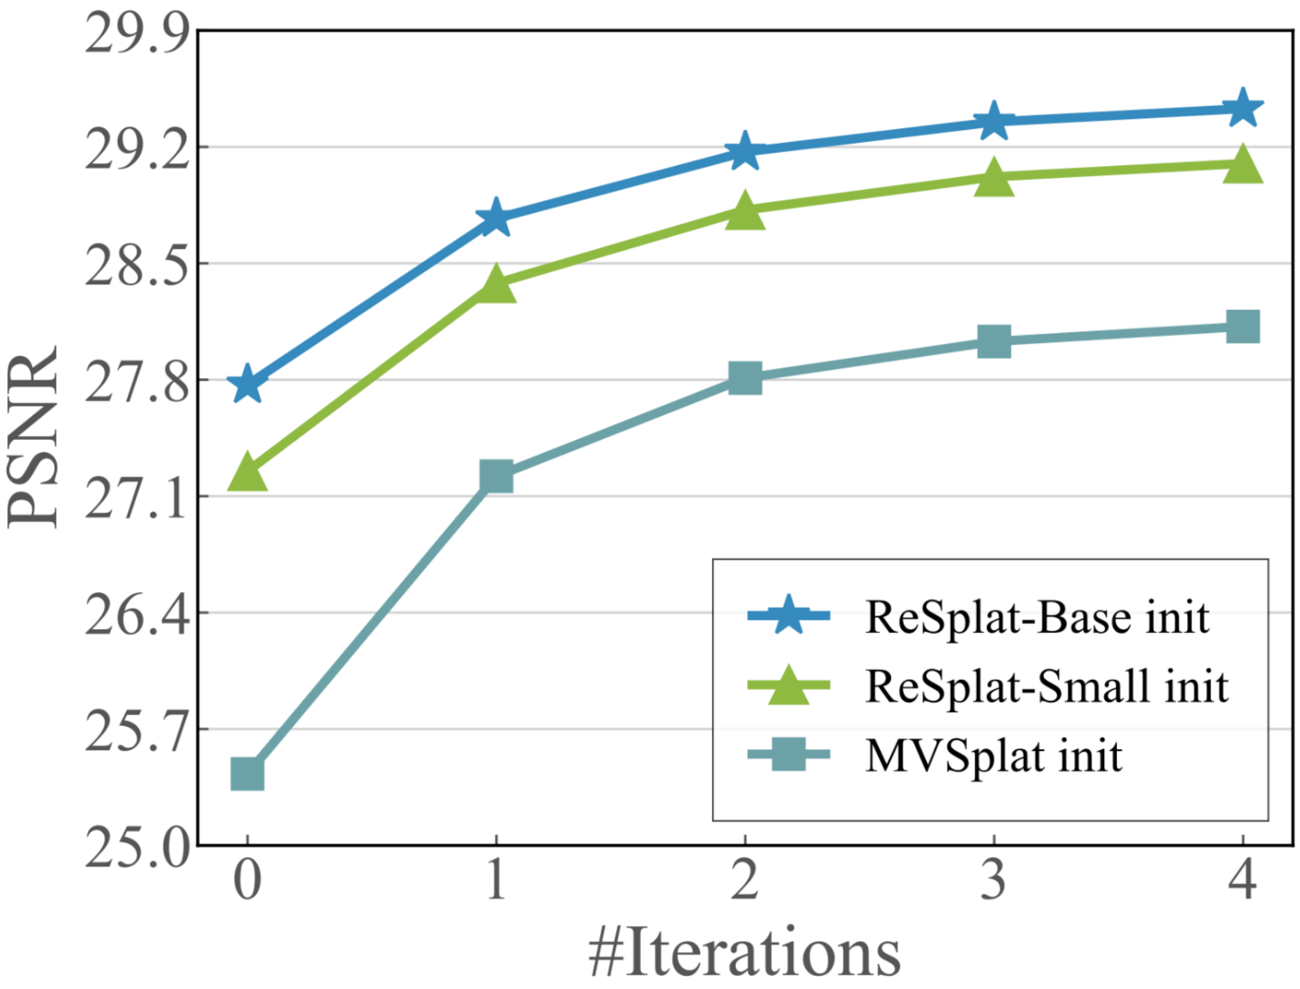

- Works for different initializations
- Better initialization, better final results
- Our initialization enables 13x faster refinement than MVSplat

Cross-Dataset Generalization

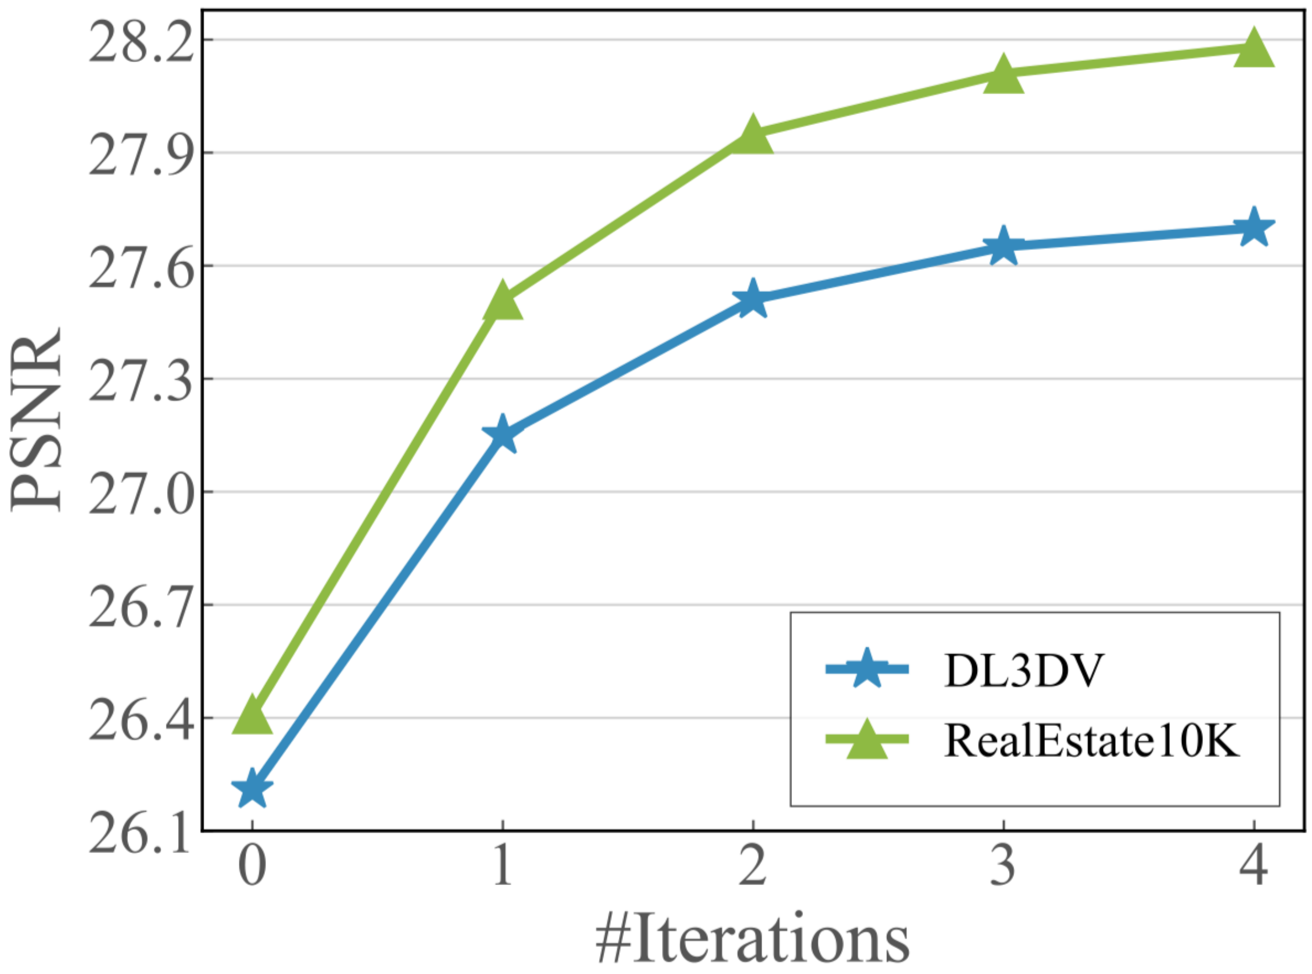

- In domain: +1.5dB PSNR
- Cross domain: +1.7dB PSNR

Cross-View Generalization

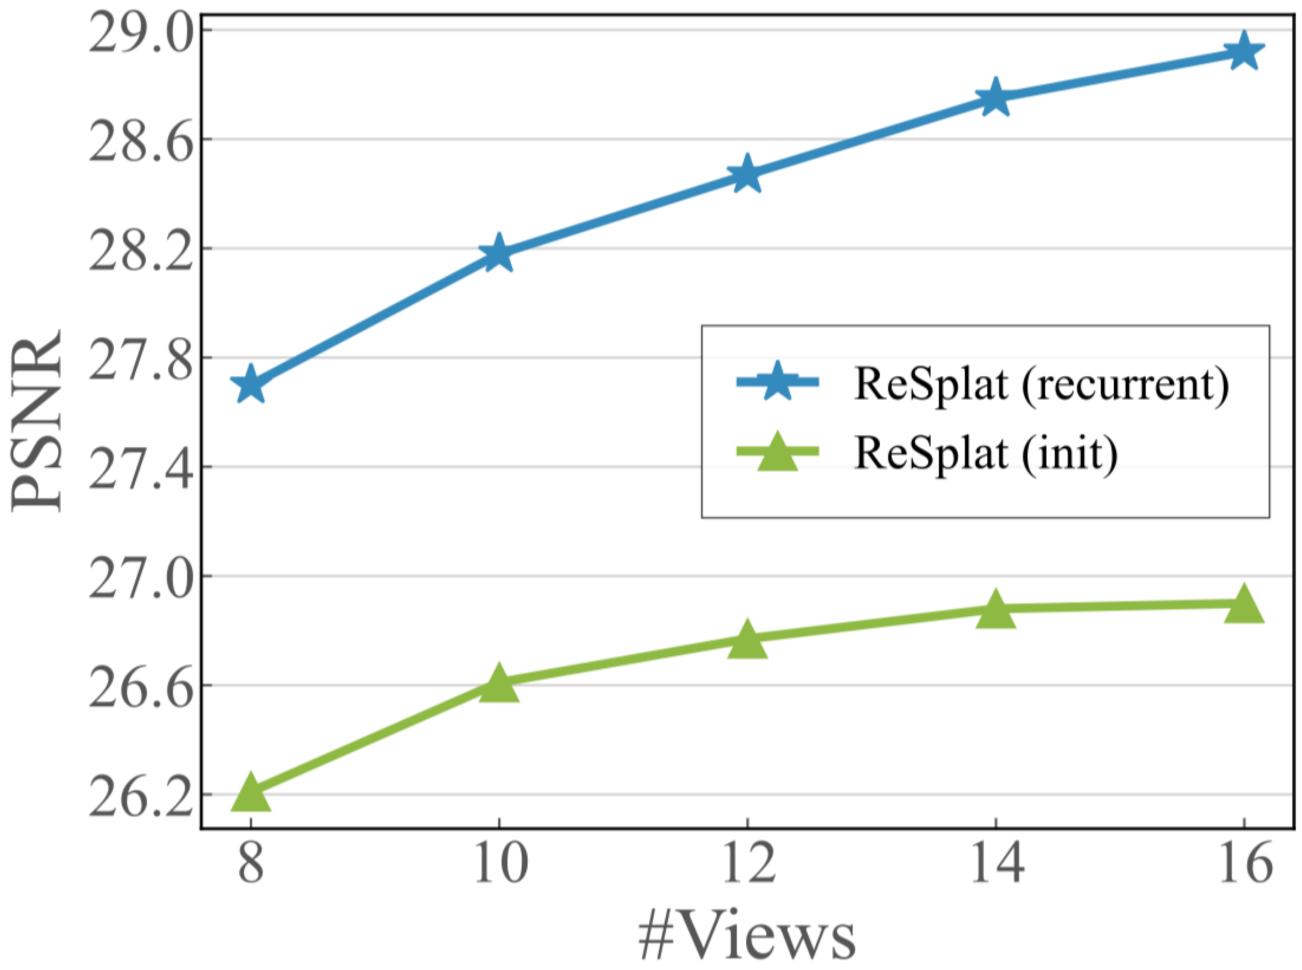

- Single-step: +0.7dB PSNR
- Recurrent: +1.2dB PSNR

Cross-Resolution Generalization

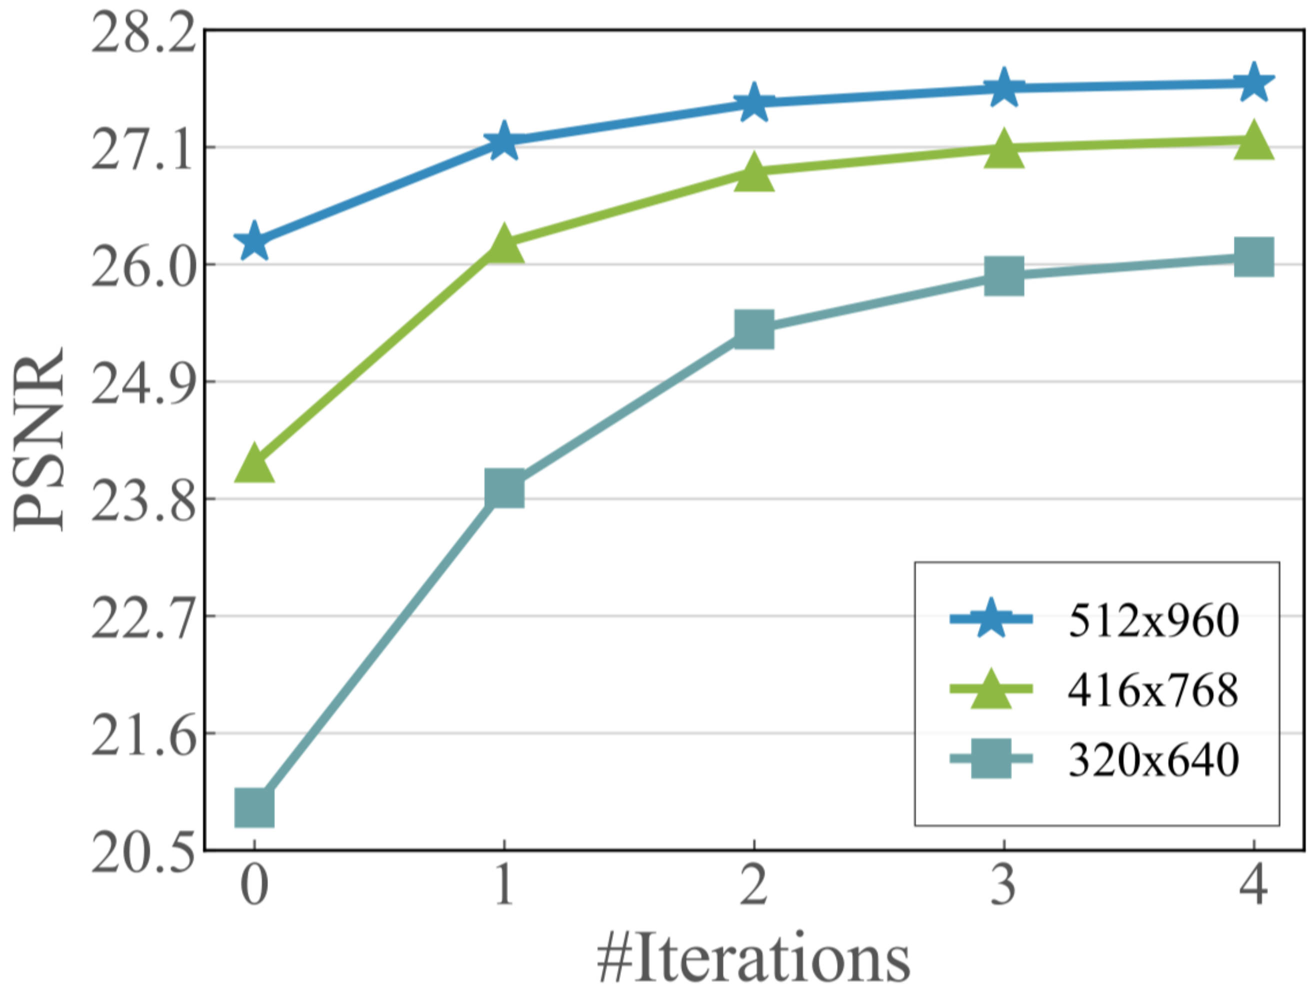

- 512x960: +1.5dB PSNR
- 416x768: +3.0dB PSNR
- 320x640: +5.1dB PSNR

Comparisons

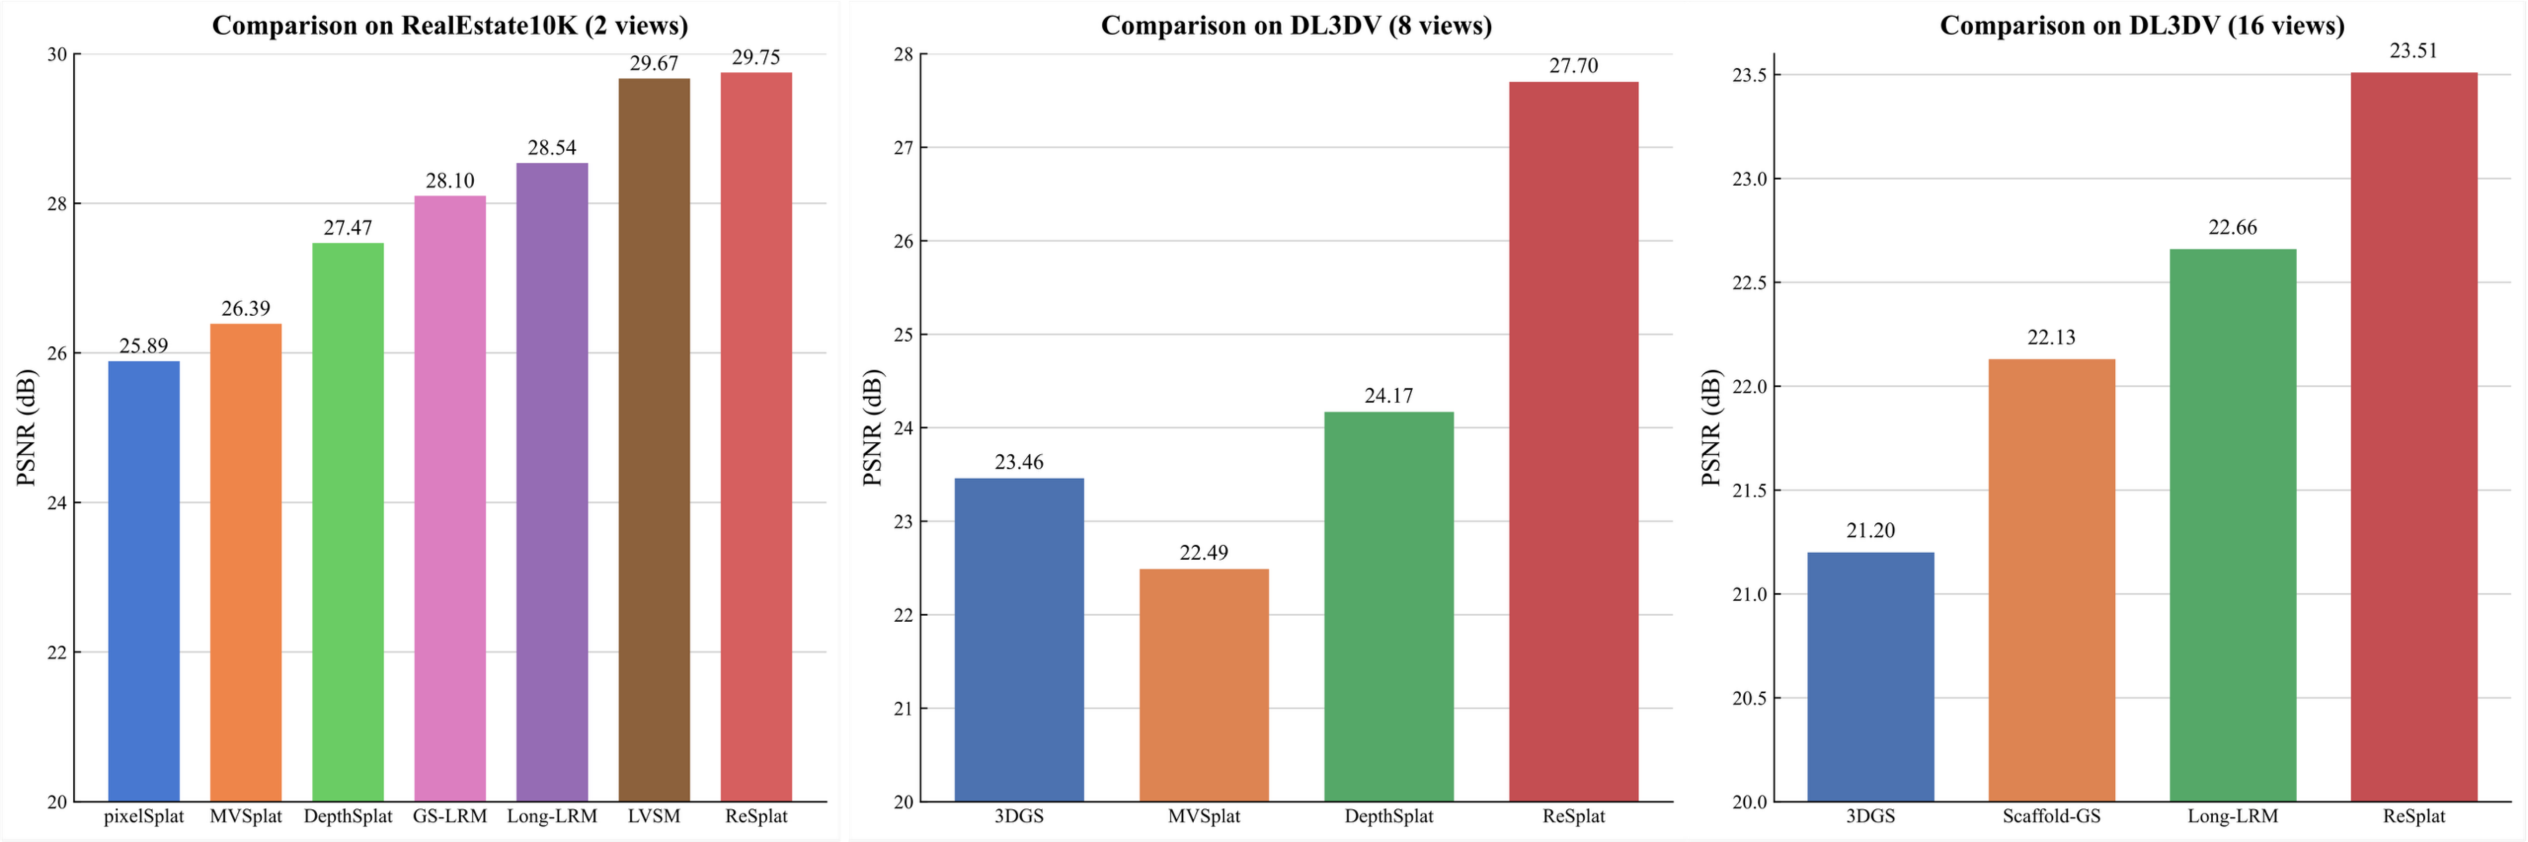

State-of-the-art across 2, 8, 16 views




Summary

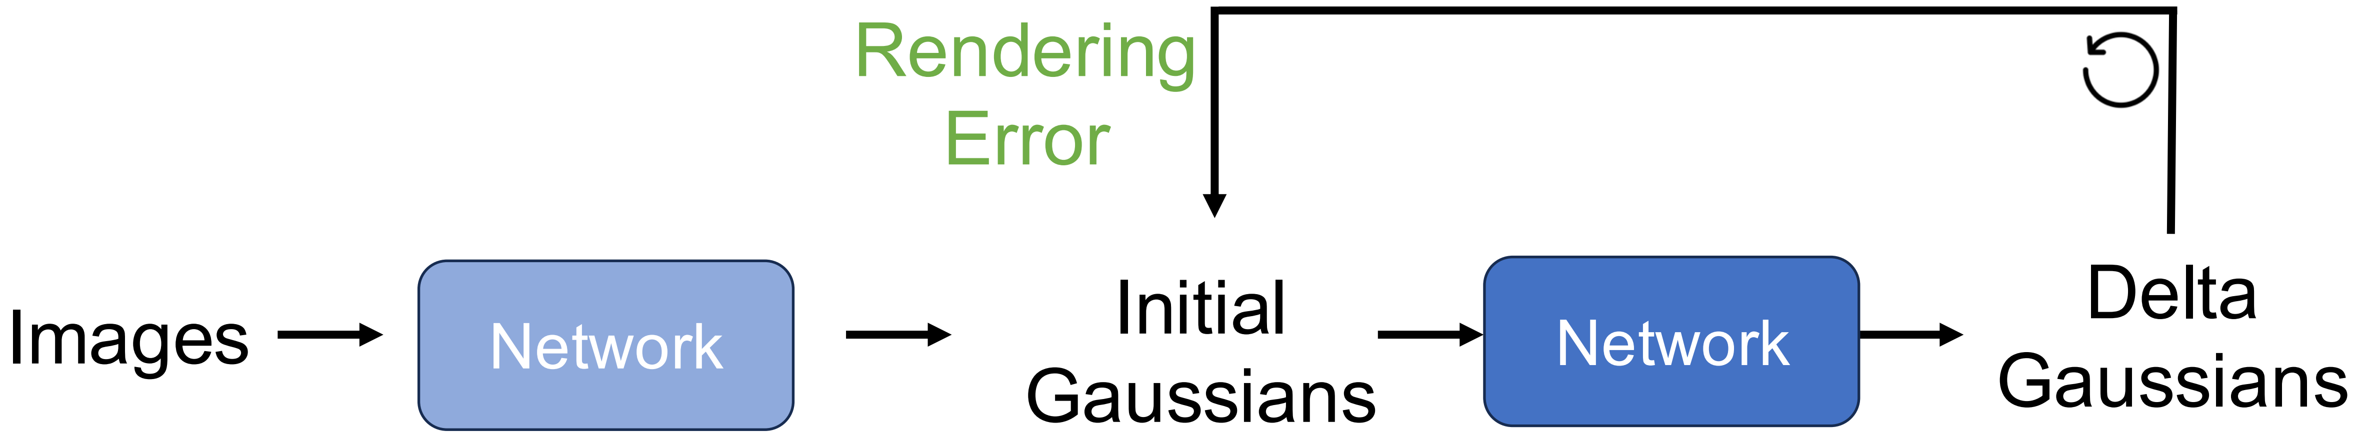

- Feed-forward compact initialization
- Feed-forward recurrent update

Xu, Barath, Geiger, Pollefeys. ReSplat: Learning Recurrent Gaussian Splatting. arXiv 2026

Future Work: Scalability

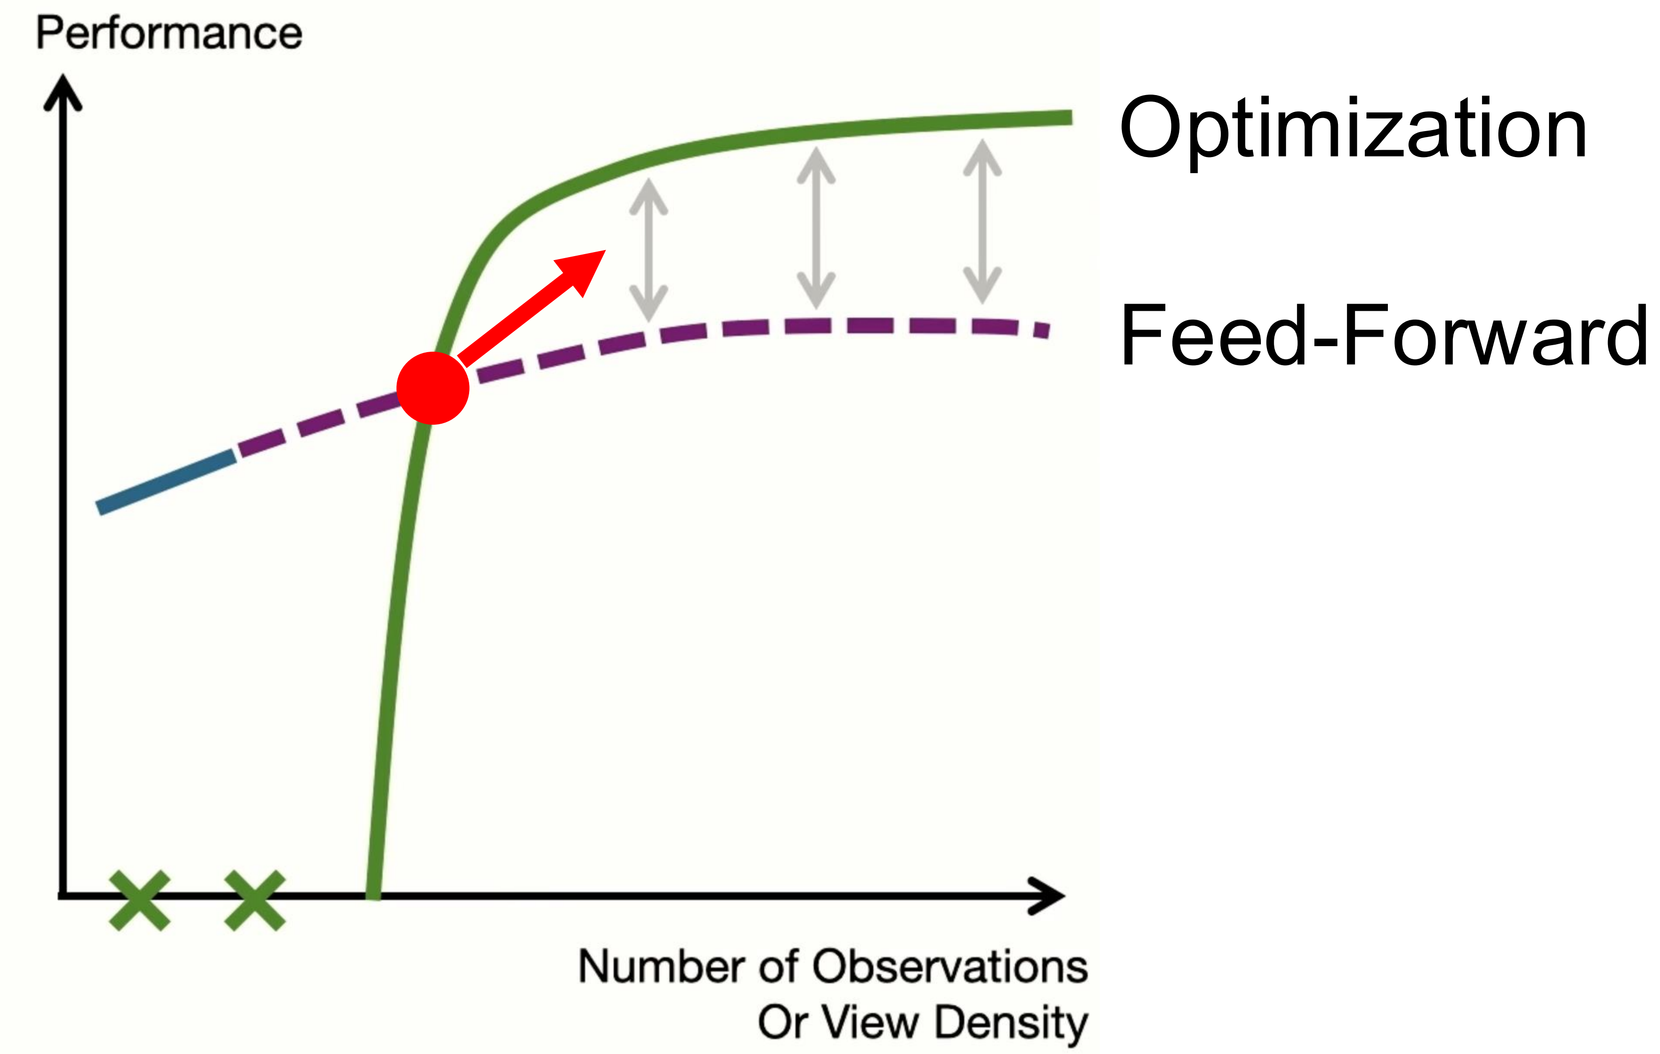

Credit: Angjoo Kanazawa https://youtu.be/Ap5jFj04p1g

Future Work: Unposed Images

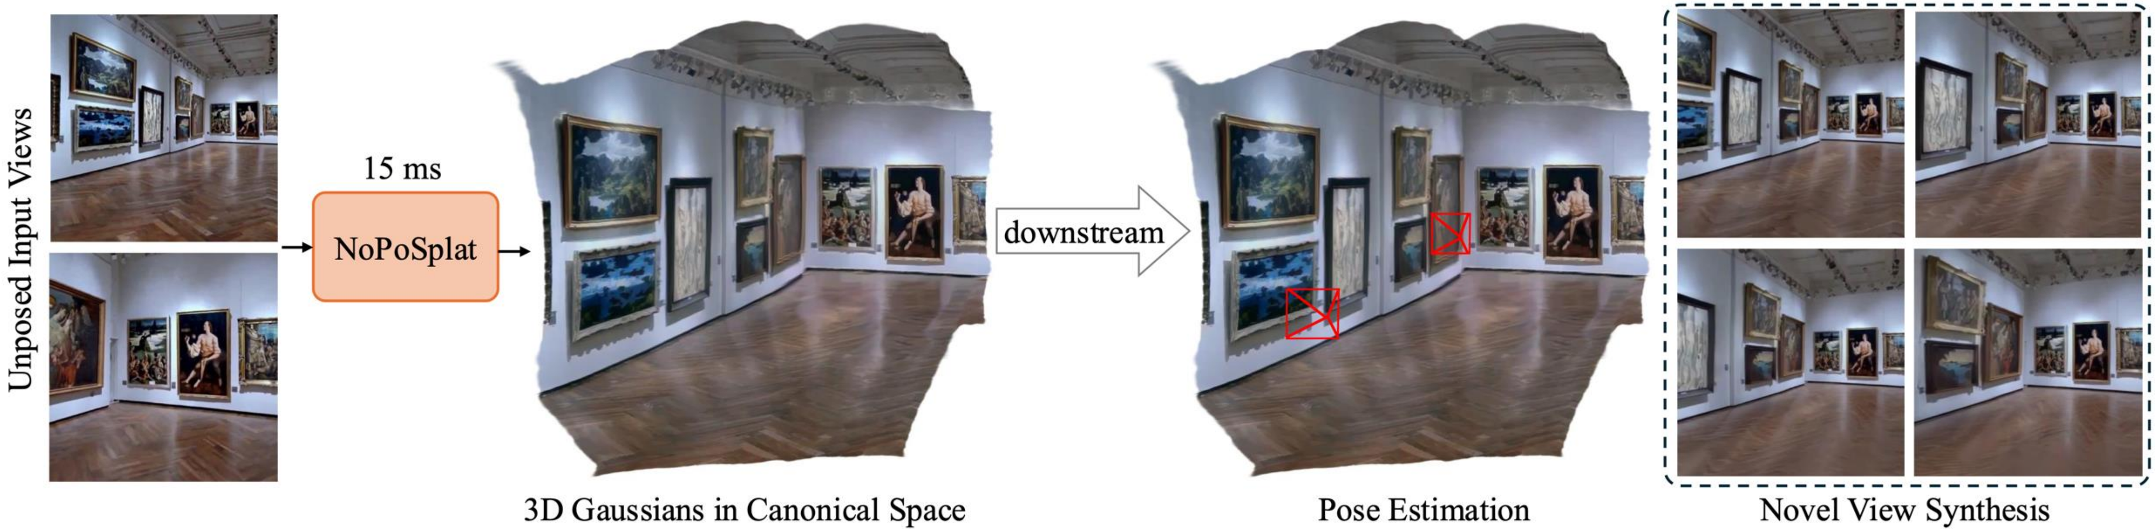

Ye et al. No Pose, No Problem: Surprisingly Simple 3D Gaussian Splats from Sparse Unposed Images. ICLR 2025 (Oral)

Ye et al. YoNoSplat: You Only Need One Model for Feedforward 3D Gaussian Splatting. ICLR 2026

Future Work: Generative Models

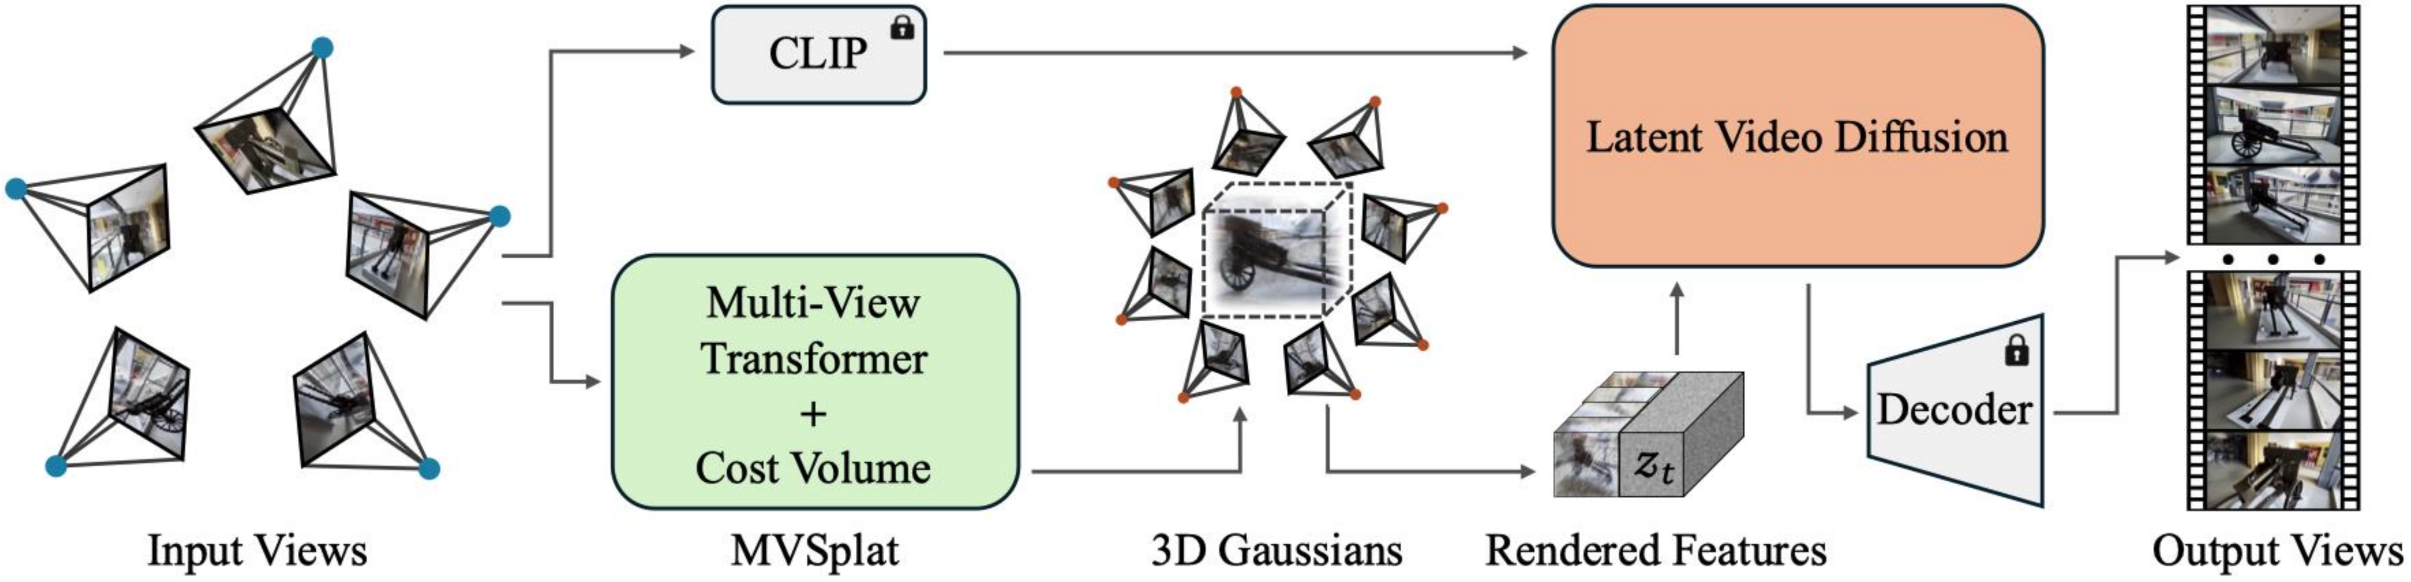

MVSplat360

Chen et al. MVSplat360: Feed-Forward 360 Scene Synthesis from Sparse Views. NeurIPS 2024

Takeaways

- ReSplat: learning to optimize 3DGS
  - Combine data-driven priors with feedback-driven iterative refinement
  - New scaling axis for future feed-forward 3D models

- Towards a unified feed-forward reconstruction system
  - Integrate recent advances into a single framework
  - Robust, instant 3D reconstruction from arbitrary in-the-wild images

Thank you!