Faster-GS: Analyzing and Improving Gaussian Splatting Optimization

Florian Hahlbohm[1] Linus Franke[2] Martin Eisemann[1] Marcus Magnor[1,3]
[1]TU Braunschweig [2]Inria [3]University of New Mexico

CVPR JUNE 3-7, 2026 DENVER COLORADO

Goal: An updated baseline for fast Gaussian Splatting [1]

Why?

Cited by 9675

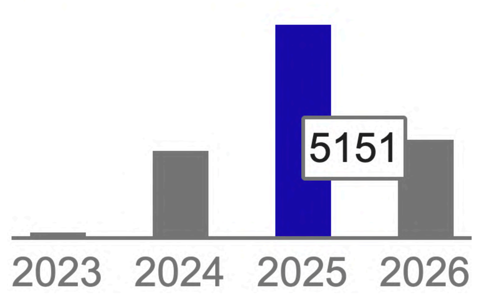

- Follow-up work proposed many improvements
- Difficult to keep track
- Must build on top of them to advance further
- Integration is often not straightforward

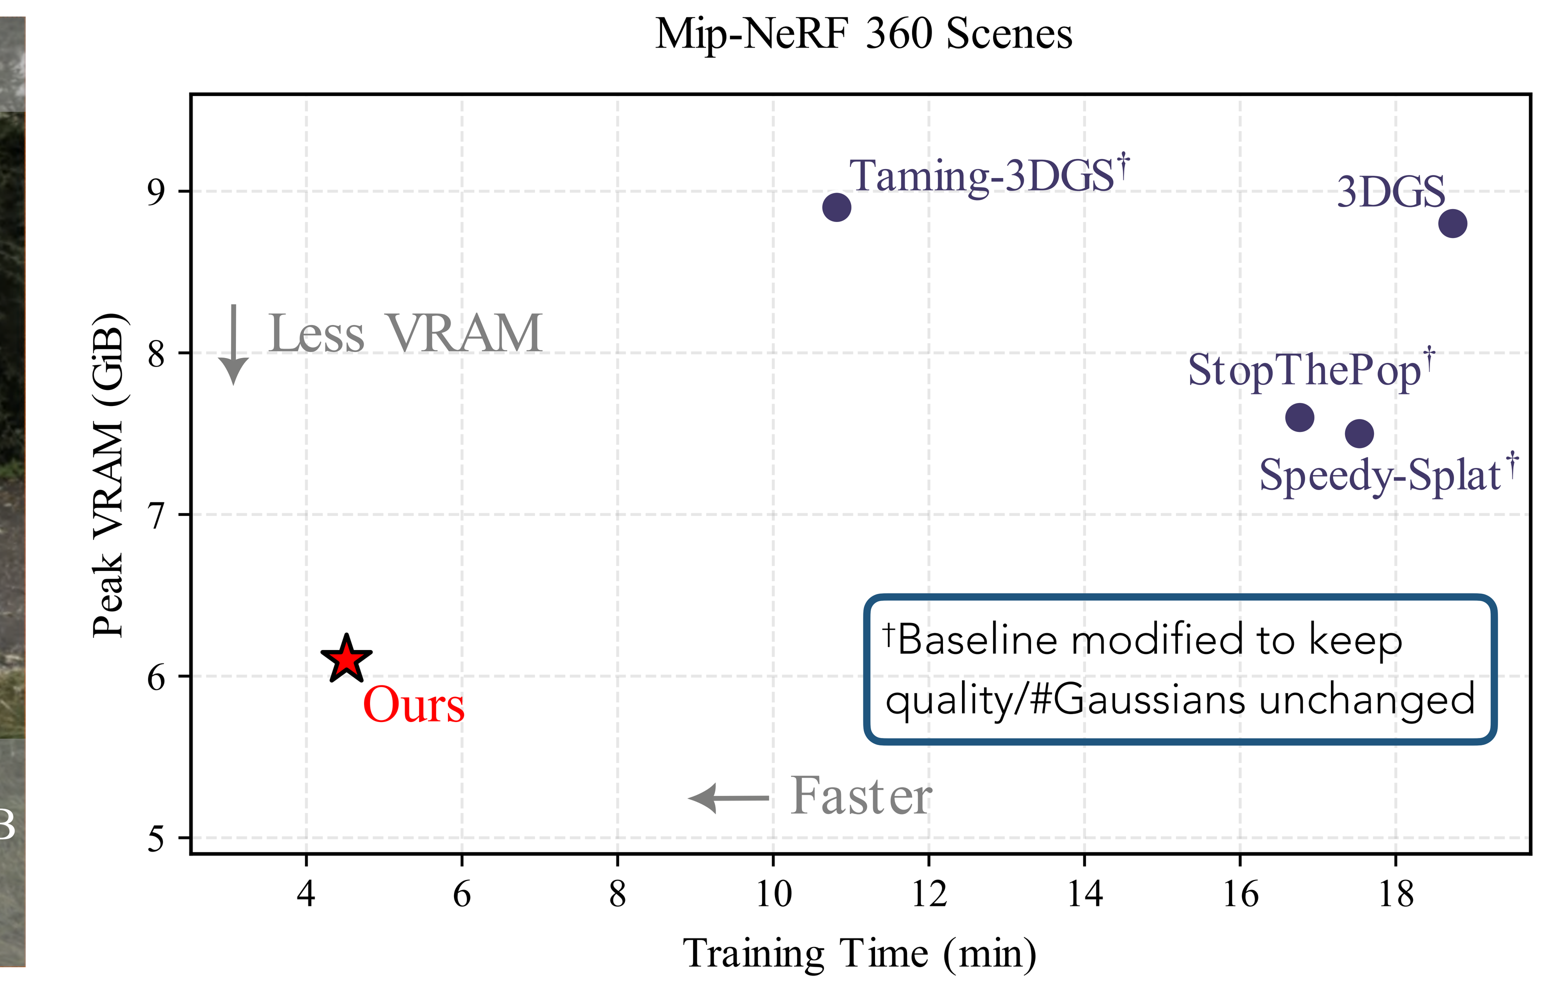

Where Does Training Time Go?

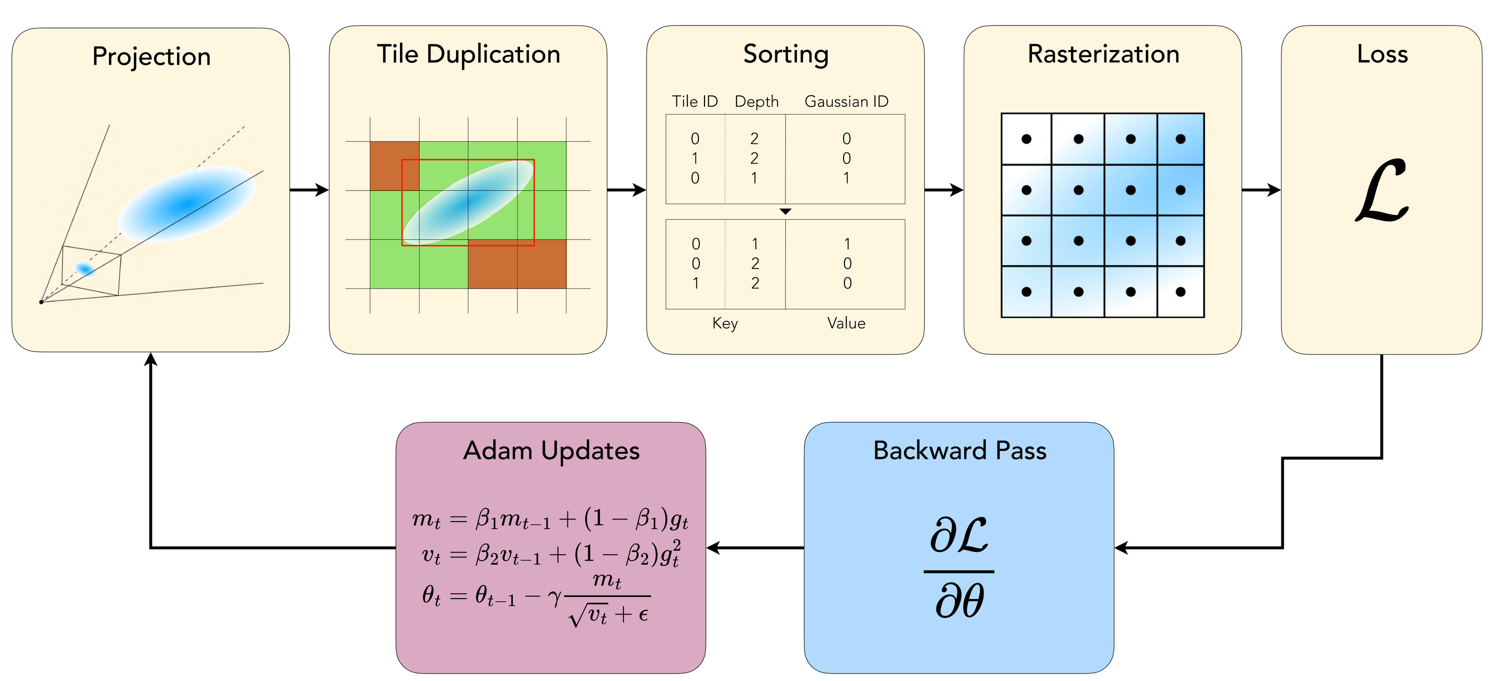

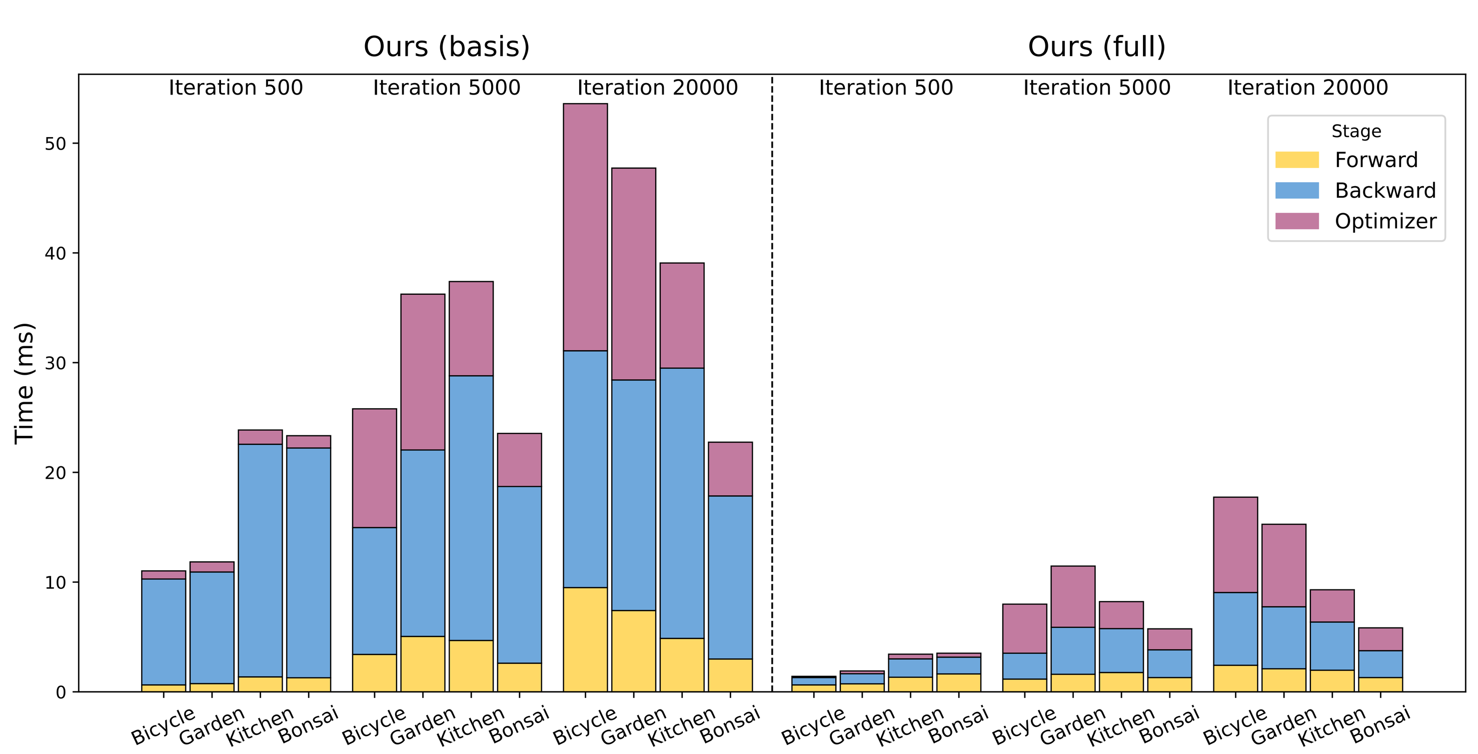

Faster-GS Optimizations

- Fused activations
- Tile-based culling [2]
- Separate depth/tile sorting [4]
- Per-Gaussian backward [3]
- Fused Adam
- Locality-preserving densification

Key Design Choice

No Quality Tradeoff
- ✓ No pruning
- ✓ No compression
- ✓ No mixed precision
- ✓ Same PSNR/#Gaussians

Results & Impact

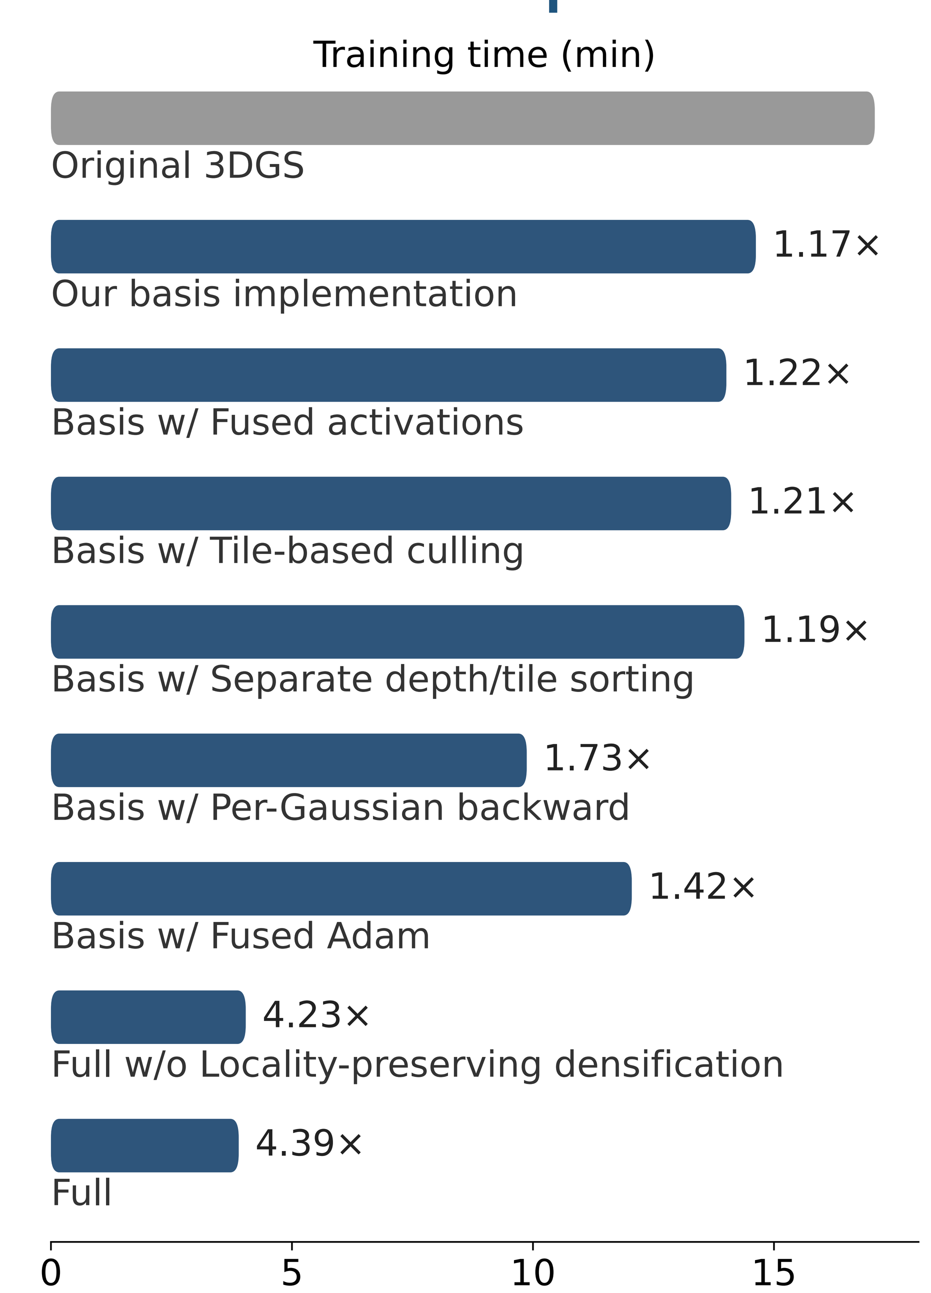

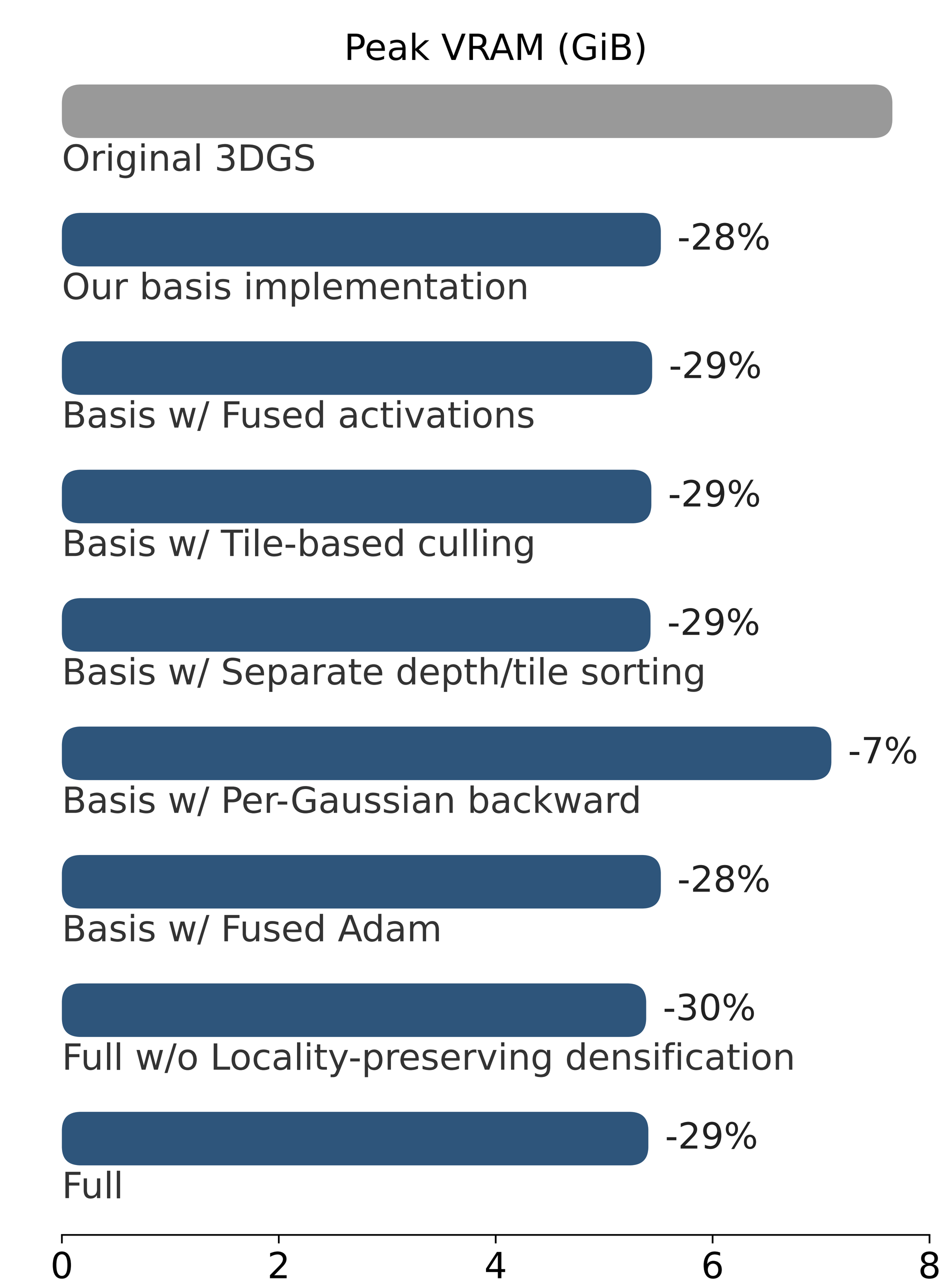

Fusing Backward & Optimizer

	PSNR↑	Train↓	VRAM↓	#Gs↓
Full	24.72	5m31s	6.0GiB	3.87M
+ fused updates	24.73	5m04s	5.6GiB	3.89M
+ fused updates (skip invisible)	24.59	3m03s	5.1GiB	3.34M
+ fused updates (SH degree=0)	24.38	2m24s	3.2GiB	3.85M

Application to 4DGS [5]

12min → 5:30min

104min → 19min

Speedup Over 3DGS

RTX 3090	RTX 4090	RTX 5090
3.9×	4.3×	4.8×

TL;DR

- ✓ Up to 5× faster training
- ✓ 30% lower VRAM
- ✓ Same reconstruction quality

Paper & Code

Contact: flo.hahlbohm@gmail.com

Funding
Deutsche Forschungsgemeinschaft
German Research Foundation
erc European Research Council
Established by the European Commission

References
[1] Kerbl et al. 3D Gaussian Splatting, ACM TOG 2023.
[2] Radl et al. StopThePop, ACM TOG 2024.
[3] Mallick et al. Taming 3DGS, SIGGRAPH Asia 2024.
[4] Schütz et al. Splatshop, CGF 2025.
[5] Yang et al. 4D Gaussian Splatting, ICLR 2024.