

360-GS: Layout-guided Panoramic Gaussian Splatting For Indoor Roaming

Supplementary Material

1. Training details

Baseline and datasets. Given that 3D-GS [3] only processes perspective images, we split each training panorama into eight perspective images. As suggested by Tancik et al. [5], we assume a camera field of view of 120 degrees and capture perspective images horizontally at the elevation angles of $[-45^\circ, 0^\circ, 45^\circ]$ with this camera. For testing, 3D-GS generated eight images from test viewpoints and combined them to form panoramas. As MipNeRF-360 [1] currently leads in NeRF rendering quality for perspective images, we trained it with perspective images and evaluated it on panoramas. We adapted INGP [4], a recent real-time rendering NeRF, for panoramic input following Huang et al. [2], to reduce the time for rendering panoramas. We train MipNeRF-360 for 250k iterations using the official code, which takes approximately 12 hours. INGP runs for 30 epochs, taking about 12 minutes per scene. Both 3D-GS and 360-GS are trained for 7k iterations with default parameters in the official code of 3D-GS. All our experiments are conducted on a single GPU Nvidia RTX 3090.

Parameters in the loss function. For 4-view inputs, we set λ_1 , λ_2 , and λ_3 in Eq. 11 as 0.8, 0.2, and 0.1 respectively. For 32-view inputs, λ_3 is set to 0.01 to better fit the sufficient inputs.

2. More results

We provide more results in Fig. 2 and Fig. 3, which supplements Fig. 7 and Fig. 8 in the main paper.

3. Discussion

Robustness to the number of training images. In Fig. 1, we present the variation curve of quantitative results for two scenes under different numbers of training views. With an increasing number of training views, all method exhibits gradual performance improvement, converging to optimal points. Nevertheless, 3D-GS and INGP struggle to handle inadequate training views, leading to diminished performance with 4-view and 8-view inputs. Our method demonstrates robustness against the number of training views. This can be attributed to the effectiveness of room layout priors, which provide valuable information when inputs are sparse. Additionally, our method consistently outperforms others across the majority of configurations.

Limitation. Despite achieving state-of-the-art performance in panoramic rendering, our method has some limitations. We rely on off-the-shelf networks to obtain layouts and depth priors, which may not yield accurate priors for com-

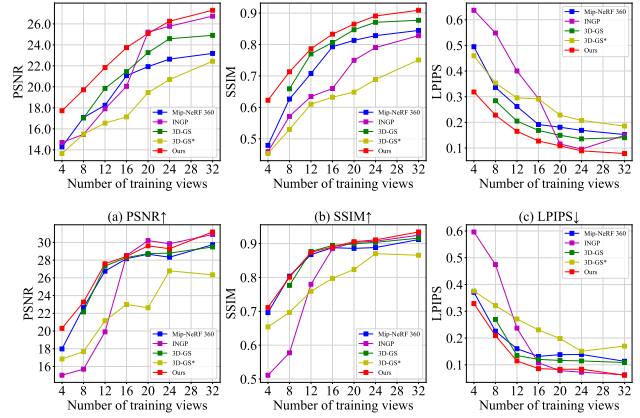


Figure 1. **Impact of varying training views.** We report results for two scenes, with each row corresponding to the results for one scene. Our method is robust to varying training views and achieves superior quantitative results.

plex scenes. This concern could be partially mitigated with a more powerful network or the use of a depth camera. Another limitation is that our initialization point cloud occupies more on-disk space, as it is sampled from the dense planes of room layouts. Balancing storage costs and rendering quality may require a meticulously crafted sampling strategy.

References

- [1] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5460–5469, 2021. 1
- [2] Huajian Huang, Ying-Rui Chen, Tianjian Zhang, and Sai-Kit Yeung. 360roam: Real-time indoor roaming using geometry-aware 360° radiance fields. *ArXiv*, abs/2208.02705, 2022. 1
- [3] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (TOG)*, 42:1–14, 2023. 1
- [4] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)*, 41:1–15, 2022. 1
- [5] Matthew Tancik, Ethan Weber, Evonne Ng, Ruilong Li, Brent Yi, Terrance Wang, Alexander Kristoffersen, Jake Austin, Kamyar Salahi, Abhik Ahuja, David Mcallister, Justin Kerr, and Angjoo Kanazawa. Nerfstudio: A modular framework for neural radiance field development. In *ACM SIGGRAPH*




Figure 2. Qualitative comparison of our methods and some SOTA methods with 32-view inputs.




Figure 3. Qualitative comparison of our methods and some SOTA methods with 4-view inputs.

075 2023 *Conference Proceedings*, New York, NY, USA, 2023.
076 Association for Computing Machinery. 1