

Figure R1: **Qualitative comparison on MIT-Intrinsic (Grosse et al., 2009) with IntrinsicAnything (Chen et al., 2024).** Input image and ground truth have been contrast-adjusted for better visibility.




Figure R2: **Failure Cases.**

## REFERENCES

Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *CVPR*, 2022.




Figure R3: **Results on Mip-NeRF 360 (Barron et al., 2022) (Part 1, outdoor).** We input 4 views for each scene.

Mark Boss, Raphael Braun, Varun Jampani, Jonathan T. Barron, Ce Liu, and Hendrik P. A. Lensch. Nerd: Neural reflectance decomposition from image collections. In *ICCV*, pp. 12664–12674. IEEE, 2021.

Xi Chen, Sida Peng, Dongchen Yang, Yuan Liu, Bowen Pan, Chengfei Lv, and Xiaowei Zhou. Intrinsicananything: Learning diffusion priors for inverse rendering under unknown illumination, 2024. URL <https://arxiv.org/abs/2404.11593>.

Roger Grosse, Micah K Johnson, Edward H Adelson, and William T Freeman. Ground truth dataset and baseline evaluations for intrinsic image algorithms. In *2009 IEEE 12th International Conference on Computer Vision*, pp. 2335–2342. IEEE, 2009.






Figure R4: **Results on Mip-NeRF 360 (Barron et al., 2022) (Part 2, indoor).** We input 4 views for each scene.




Figure R5: **Results on indoor and outdoor scenes.** Input images are collected from the Internet.





Figure R6: **Reconstruction and relighting results on real-world data.**

