Supplementary Document: Inverse Global Illumination using a Neural Radiometric Prior

Saeed Hadadan
University of Maryland, College Park
NVIDIA
MD, USA
saeedhd@umd.edu

Geng Lin
University of Maryland, College Park
MD, USA
geng@umd.edu

Jan Novák
NVIDIA
Czech Republic
jnovak@nvidia.com

Fabrice Rousselle
NVIDIA
Switzerland
frousselle@nvidia.com

Matthias Zwicker
University of Maryland, College Park
MD, USA
zwicker@cs.umd.edu

ACM Reference Format:
Saeed Hadadan, Geng Lin, Jan Novák, Fabrice Rousselle, and Matthias Zwicker. 2023. Supplementary Document: Inverse Global Illumination using a Neural Radiometric Prior . In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Proceedings (SIGGRAPH '23 Conference Proceedings), August 6–10, 2023, Los Angeles, CA, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3588432.3591553

A USING MLPs FOR SCENE PARAMETERS

We show in Figure 1 that representing SVBRDF parameters with MLPs yields reconstructions of higher quality than using dense grids. For this reason, although it requires the use of additional libraries to Mitsuba and is not compatible with mega kernels at this time, we still favor using MLPs.

B ADDITIONAL TIME AND MEMORY ANALYSIS

As described in Section 6 of the paper, we measure the time and memory consumption of different methods as the path lengths change in a scene. To enable the comparison with mega kernels, as shown in the main paper, we used dense grids in Mitsuba as radiance and scene parameter representation. In contrast, here we present the results using MLPs in PyTorch in Figure 2. The VRAM consumption is the sum of the peak allocated memory reported by both Dr.JIT and PyTorch. While our method has a larger VRAM overhead due to additional radiance MLP queries, its time and memory usage remains constant as path length increases, while the costs of other methods grow rapidly.

In all experiments, we initialize the radiance grid values to the albedo of the walls and never update them during the measurement, i.e., the back-propagation and gradients are computed as usual but not applied to the grid values. This ensures that the measurements are from fixed albedos. We obtain the peak VRAM numbers from the Dr.JIT memory allocator. For methods that solve path integrals (AD-PT and PRB), we enable Russian-Roulette with a minimum termination probability of 0.05, and cap the maximum path length to the 99.9 percentile when the scene is rendered with path tracing.

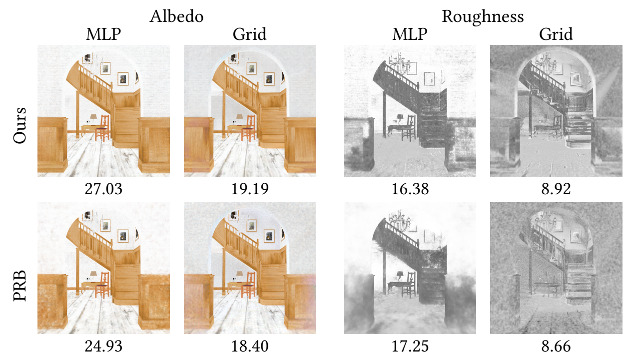

Figure 1: Grid SVBRDF. We compare the results of using a dense grid with resolution 256^3 to store the scene parameters versus using an MLP. MLP results are superior regardless of rendering method. PSNR is reported.

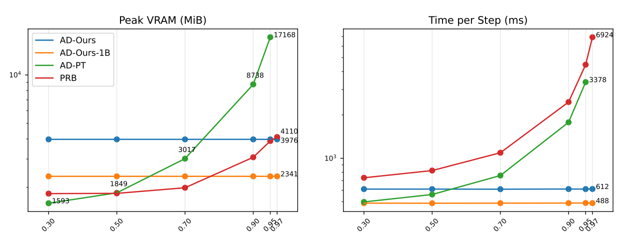

Figure 2: *Cube* scene measurements. We compare the time and memory consumption of each training step, with all parameters represented as MLPs. Note that the Y-axis is in log scale. AD-PT runs out of VRAM (24GB) at albedo 0.97. Our method uses constant amounts of VRAM and time.

C ADDITIONAL RESULTS

We present results of additional scenes in Figure 3.

SIGGRAPH '23 Conference Proceedings, August 6–10, 2023, Los Angeles, CA, USA

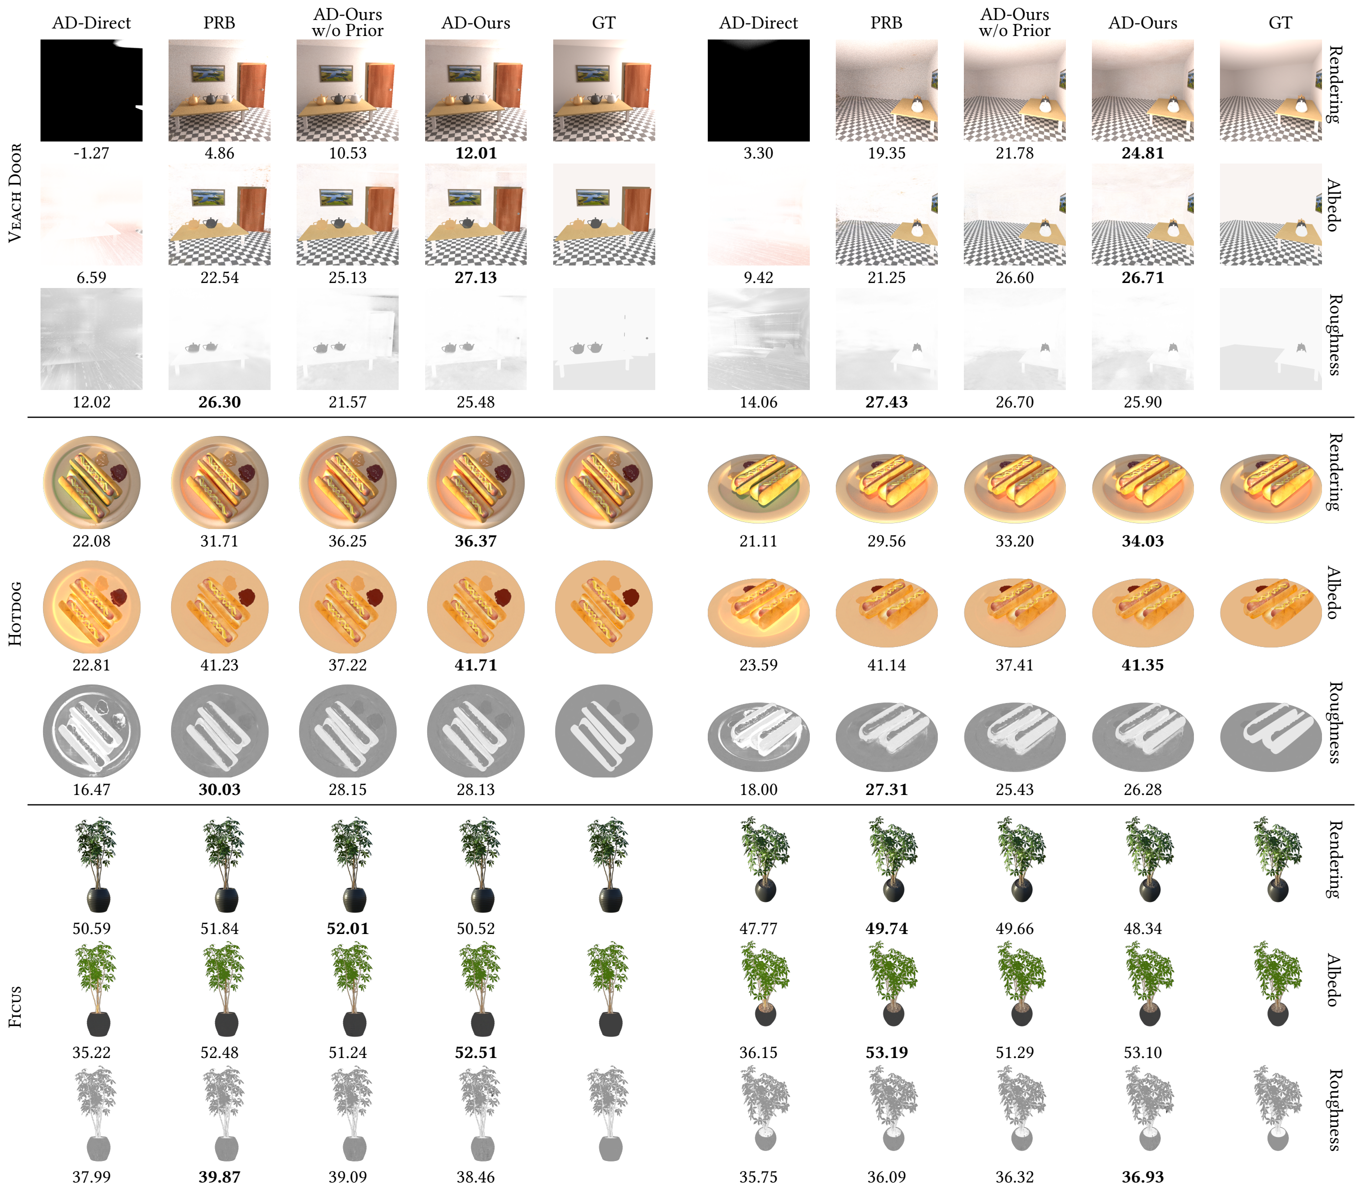

Figure 3: Main results for NeRF scenes. For each scene, we compare the rendering, recovered albedo, and recovered roughness (top to bottom rows) for direct illumination, PRB, and AD-Ours. We also compare to the case where the radiance cache is trained without the prior. We show two different views of each scene, and report PSNR to ground truth (GT).

D TRAINING PROGRESS CURVES

We present how each method converges during training in Figure 4 and 5.

E ABLATION FOR MORE SCENES

We provide an ablation study for more scenes in Figure 6.

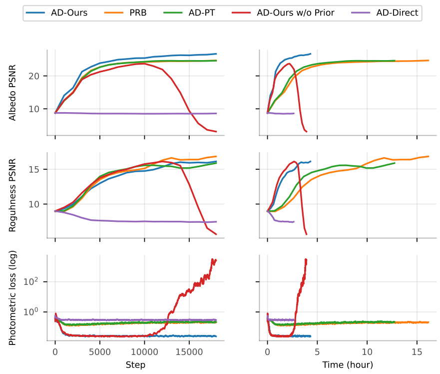

(a) Staircase

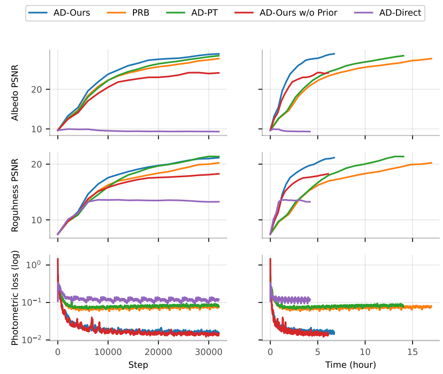

(b) Kitchen

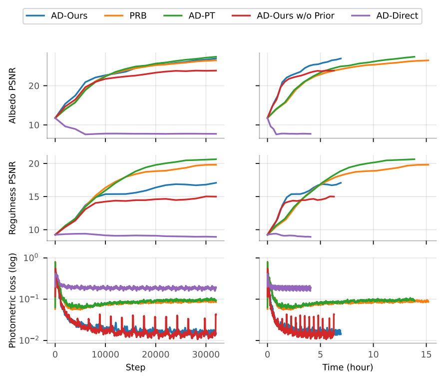

(c) Living Room

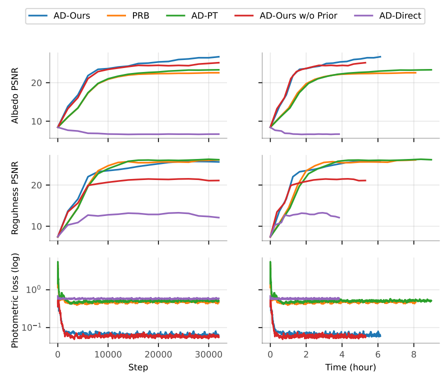

(d) Veach Door

Figure 4: Reconstruction accuracy. Albedo, roughness and photometric error, for all the scenes, as a function of training steps and time. Our method correctly accounts for global illumination thanks to our neural radiometric prior, resulting in comparable accuracy at low computational cost.

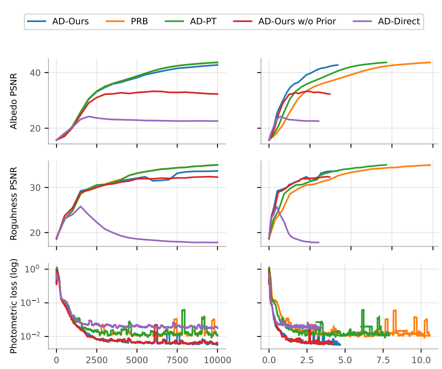

(a) Lego

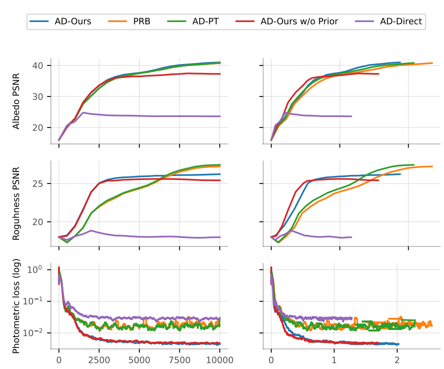

(b) Hotdog

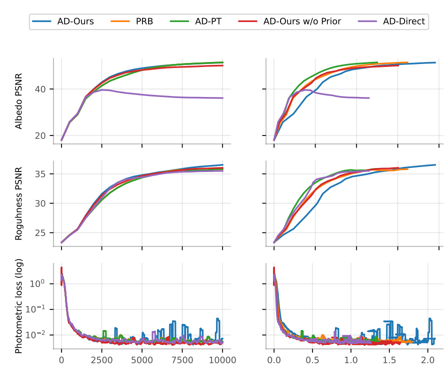

(c) Ficus

Figure 5: Reconstruction accuracy for NeRF scenes. Albedo, roughness and photometric error, for all the NeRF scenes, as a function of training steps and time. Our method correctly accounts for global illumination thanks to our neural radiometric prior, resulting in comparable accuracy at low computational cost.

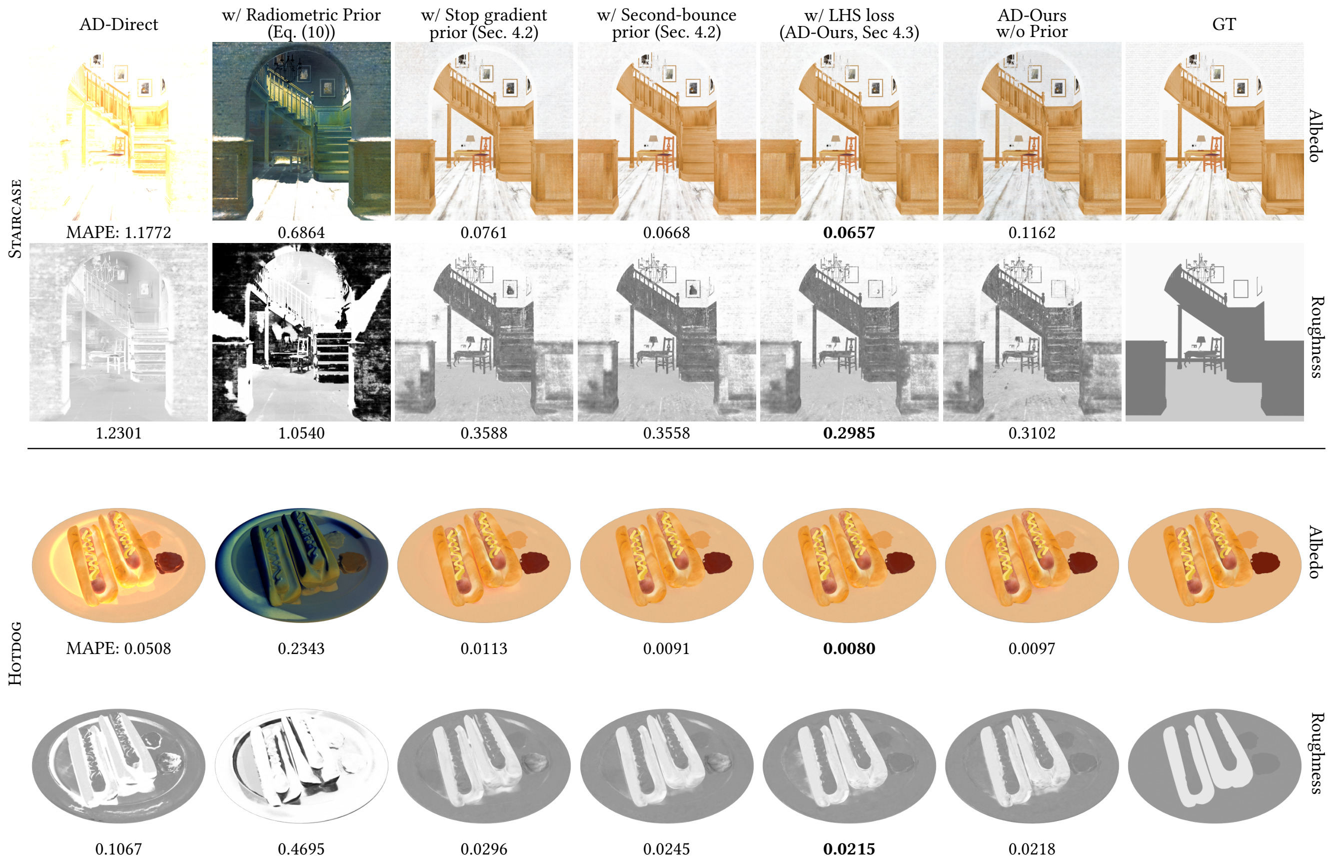

Figure 6: Ablation for more scenes. We start with the direct illumination integrator (left), and add the radiometric prior to it. The results significantly improve when we ignore the gradients of the prior w.r.t scene parameters. Adding the prior to the second bounce better accounts for additional global illumination effects for areas unseen by the input cameras. Using ground truth data to improve the radiance field further improves the quality. Finally, the second column from the right shows our full method, except that we omit the prior.