Diffusion Posterior Illumination for Ambiguity-aware Inverse Rendering – Supplemental Material

A PERFORMANCE BREAKDOWN

We report runtime performance comparisons in Tab. 1.

Prior	Time (ms) / iteration	Iterations	Total time (min)
`No Prior`	669	400	4.5
`Global Norm.`	669	400	4.5
`Smoothness`	675	400	4.5
`Chromaticity`	675	400	4.5
`DIP`	698	600	7.0
`GAN`	706	950	11.2
`RENI`	690	1000	11.5
`Ours`	920	1000	15.3 (19.8)

Table 1. Runtime performance comparisons. We report the denoising and joint optimization scheme as `Ours`. Our method includes 4.5 minutes for initialization and material refinement, and 15.3 minutes for optimization, resulting in 19.8 minutes in total.

B EXTRACTABLE INFORMATION FROM HIGHLY SPECULAR MATERIALS

In Fig. 1, we analyze the information that can be extracted from highly specular materials, where small changes in lighting lead to significant changes in the image observations. In the figure, replacing a car in the environment map with a pure red box leads to many differences in the rendered images and shows that the renderings contain significant information about the car. Our method can use this information to reconstruct environment maps close to ground truth.

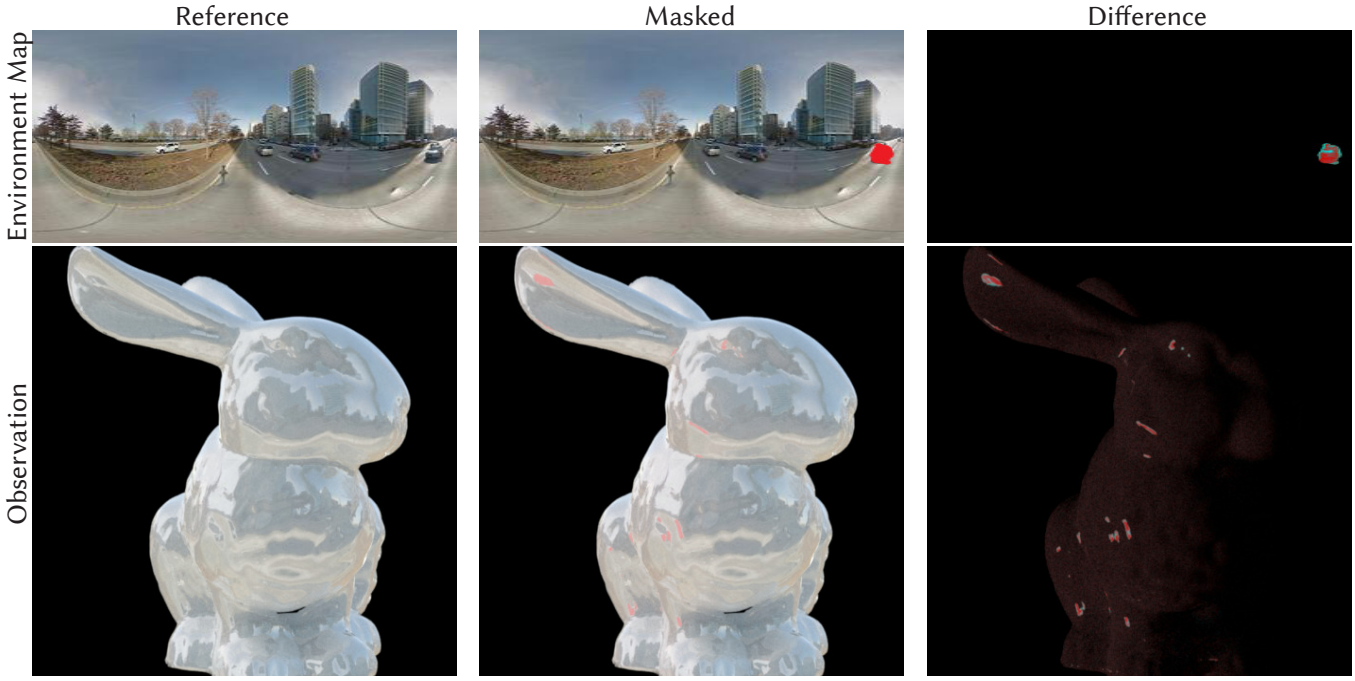

Fig. 1. Masking pixels in an environment map reveals that there is significant information to reconstruct close-to-ground-truth illumination from observations with highly specular materials.

C EVALUATION ON REAL-WORLD DATASETS

In Fig. 2 we provide qualitative comparisons against baseline priors on two real-world datasets.

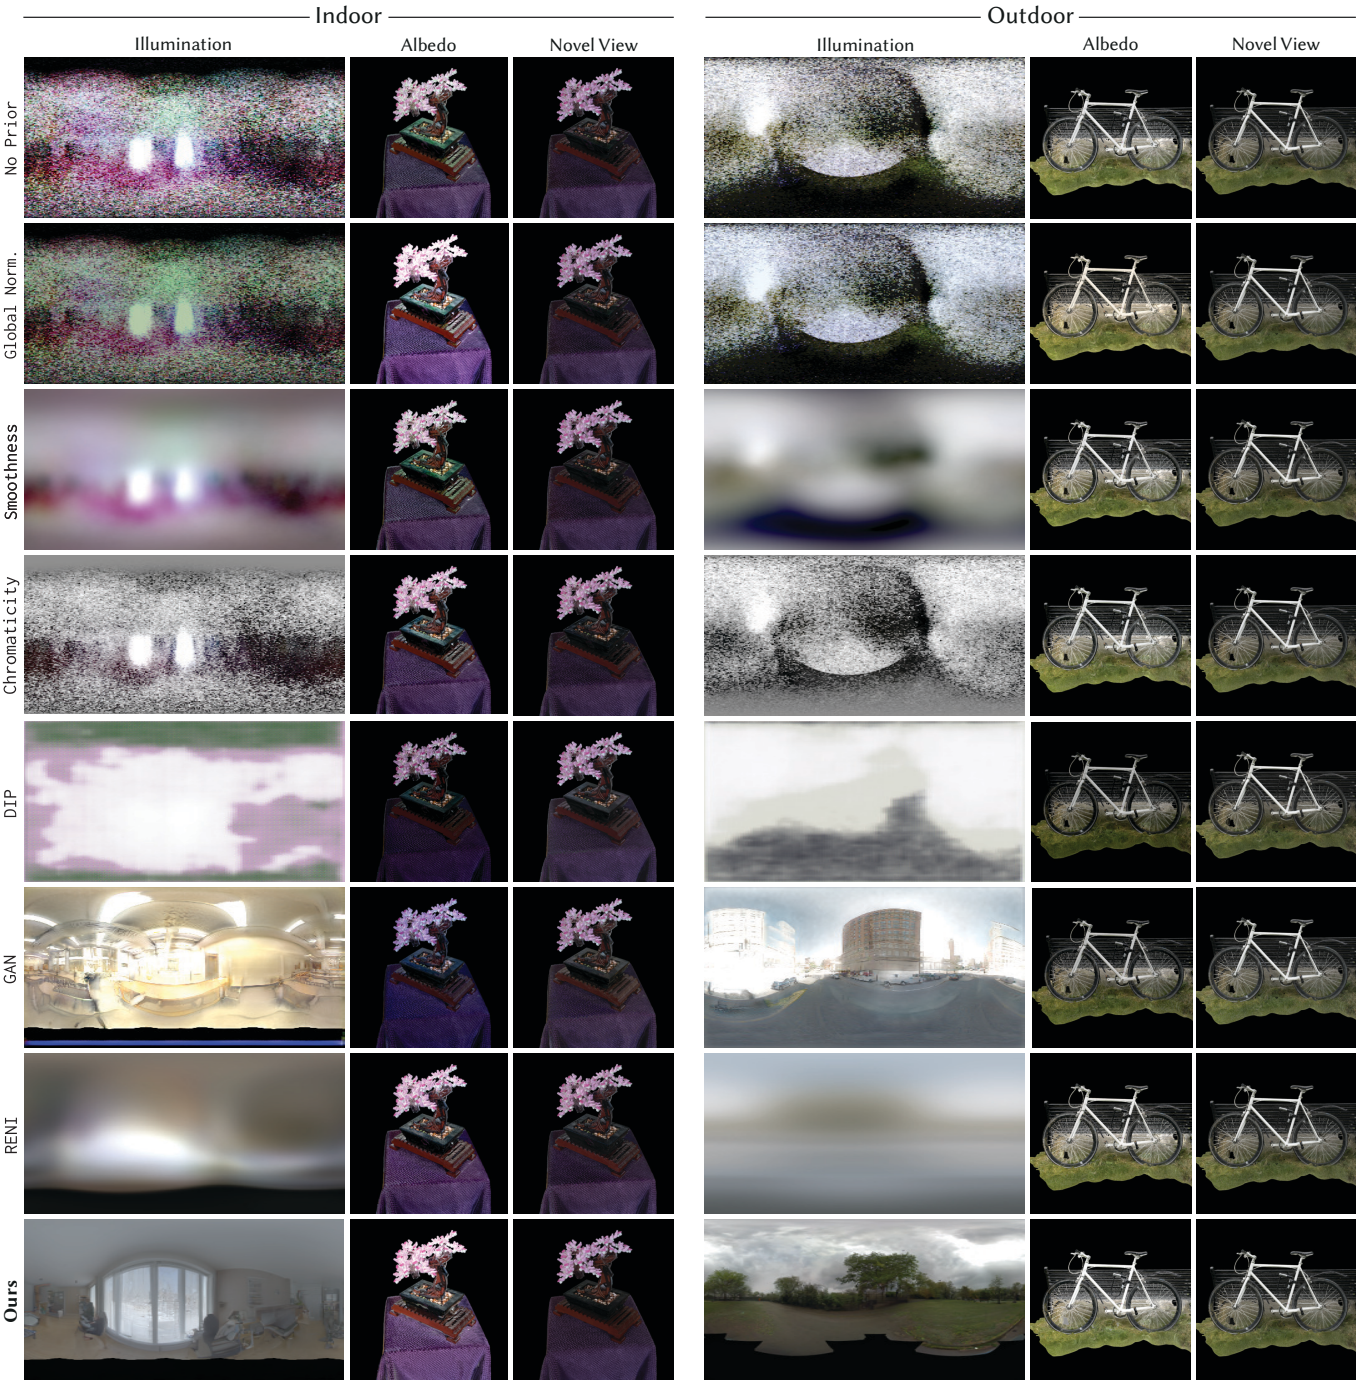

Fig. 2. Scene decompositions using different priors for two real-world scenes with indoor (left) and outdoor (right) illumination.