SpecGen: Neural Spectral BRDF Generation via Spectral-Spatial Tri-plane Aggregation

Supplementary Material

In this supplementary document, we first present additional visual results spanning a broader wavelength range, complementing the teaser figure in the main paper. Next, we present extended robustness analyses and visual comparisons with existing hyperspectral image reconstruction approaches under diverse environmental lighting conditions. Finally, we summarize the current limitations of our framework and outline potential directions for future work, including scene-level spectral BRDF estimation from casual multi-view imagery, spectral-domain neural sampling schemes, adaptive wavelength discretization, and more flexible or lightweight model architectures.

1. Qualitative Experiments

To complement the examples shown in the teaser of the main paper, and to further highlight the advantages of our method over existing hyperspectral image reconstruction approaches [1, 2], which are limited to reconstructing spectral images within a narrow wavelength range, we demonstrate that our method can generate spectral BRDFs capable of reconstructing images across the full wavelength range of 360 nm to 1000 nm. As illustrated in Figure S2, we average 20 sampled wavelengths for multispectral re-rendering and perform a qualitative comparison with the ground truth (GT), verifying the accuracy of our method across different wavelengths.

2. Robustness Analyses

To complement the robustness tests introduced in the main text, as shown in Figure S3, we present experiments under two environmental light conditions that were not used during training. For each case, we showcase rendered results at different wavelengths. The results demonstrate that, even with unseen environmental light inputs, our method maintains strong robustness and generalizability.

3. Limitations and Future Work

Although the proposed SpecGen framework delivers encouraging results, there are a few aspects that could benefit from additional investigation.

Limitations First, since the sampling patterns used during training do not fully match the illumination distributions at test time, renderings under complex environment maps occasionally exhibit mild artifacts, as shown in Fig. S1. Existing neural sampling techniques developed for RGB BRDFs provide valuable insights, yet the higher dimensionality of spectral BRDFs means that these RGB-oriented strategies do not always carry over directly. Second, even with the aid of a total-variation term that promotes smoothness, faint noise can still appear on highly continuous surfaces (e.g., spheres). We attribute this mainly to the present tri-plane formulation, which may be less expressive for extremely smooth functions. Third, the current prototype is tailored to a single RGB photograph; this design choice simplifies network conditioning but also restricts immediate applicability to scenes that feature more complex geometry or spatially varying materials.

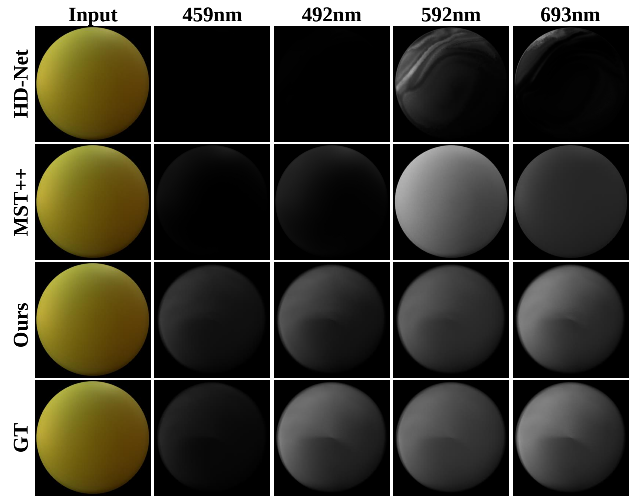

Figure S1. Qualitative comparisons of reconstructed images under environmental illumination are shown alongside results from state-of-the-art hyperspectral image reconstruction methods, including HDNet [2] and MST++ [1]. Our method demonstrates superior spectral accuracy and robustness across different wavelength bands and lighting conditions, yielding reconstructed images that closely match the GT and outperform existing approaches in visual fidelity.

Future Work We plan to extend SpecGen toward scene-level spectral BRDF estimation so that spatially varying materials in real-world environments can be reconstructed from casual multi-view imagery. Another promising direction is the development of neural sampling schemes specifically tuned to the spectral domain, which may better accommodate its higher dimensional characteristics. We are also interested in incorporating adaptive wavelength discretization strategies to cover broader spectral ranges and to model anisotropic reflectance effects. Finally, we will explore addi-

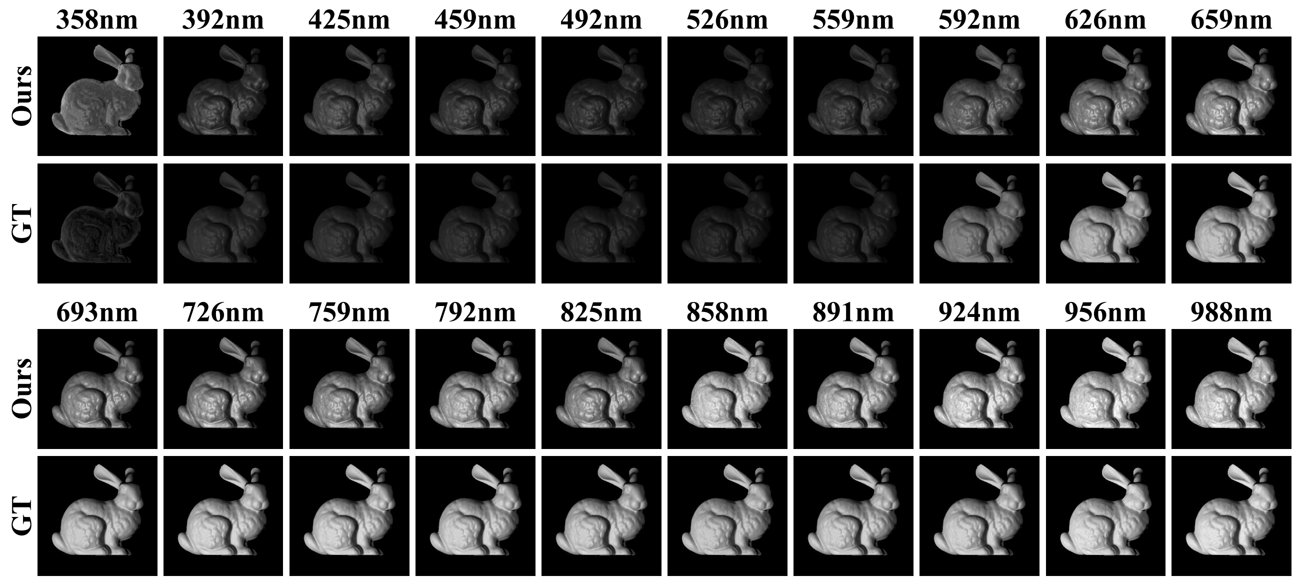

Figure S2. Rendered images of the same material across additional wavelengths (400–1000 nm), generated using the same input as in the teaser figure of the main paper. Comparison with the ground truth (GT) further verifies the accuracy of our method.

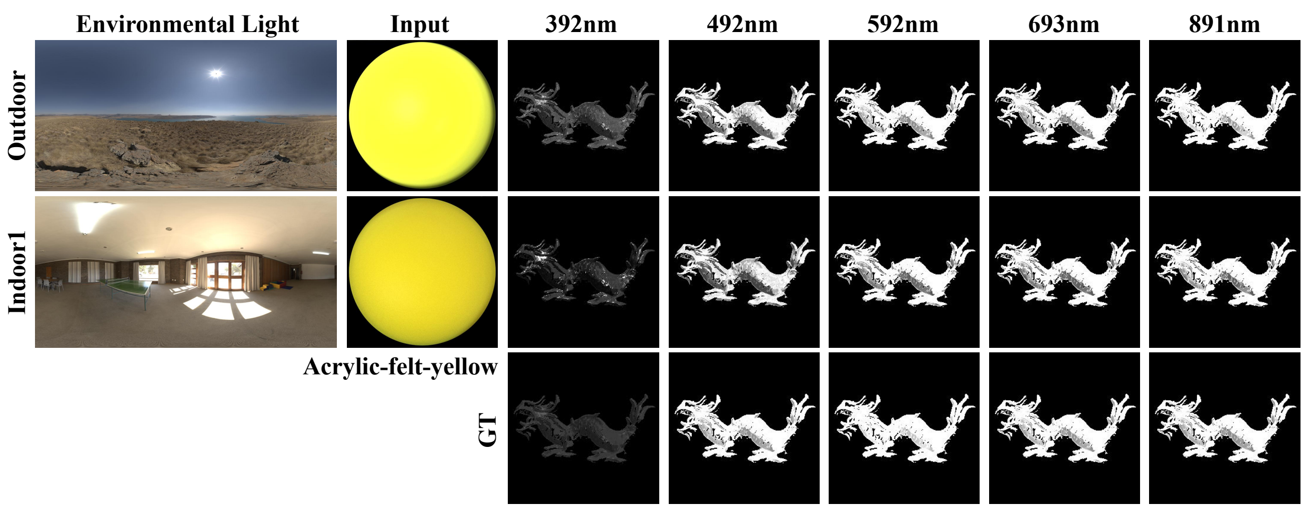

Figure S3. Spectral BRDF results for "acrylic-felt-yellow" generated from inputs rendered under diverse unseen environmental lights (Outdoor, Indoor1), with outputs rendered under distant light. The left side displays the environmental lights used for rendering the input images. Our method demonstrates robust performance across varying environmental lighting conditions, closely matching the GT across wavelengths.

tional regularization techniques, more flexible multi-plane decompositions, and lightweight model variants to reduce memory footprint and inference time while further suppressing residual noise and improving output smoothness.

References

[1] Yuanhao Cai, Jing Lin, Xiaowan Hu, Haoqian Wang, Xin Yuan, Yulun Zhang, Radu Timofte, and Luc Van Gool. Mask-guided spectral-wise transformer for efficient hyperspectral image reconstruction. In *CVPR*, pages 17502–17511, 2022. 1

[2] Xiaowan Hu, Yuanhao Cai, Jing Lin, Haoqian Wang, Xin Yuan, Yulun Zhang, Radu Timofte, and Luc Van Gool. Hdnet: High-resolution dual-domain learning for spectral compressive imaging. In *CVPR*, pages 17542–17551, 2022. 1