

Inverse Rendering with Spherical Gaussians for Physics-based Material Editing and Relighting Supplemental Document

In this supplemental document, we first provide details of the specular BRDF model used in our pipeline. Then, additional results on the synthetic data are included. Finally, we show additional results with videos.

1. Specular BRDF

Suppose the specular albedo is $\mathbf{s} \in [0, 1]^3$ and roughness is $R \in \mathcal{R}_+$. We use the following simplified Disney BRDF model for specular BRDF $f_s(\omega_o, \omega_i)$ as in prior works [1, 3]:

$$f_s(\omega_o, \omega_i) = \mathcal{M}(\omega_o, \omega_i) \mathcal{D}(\mathbf{h}), \quad (1)$$

$$\mathbf{h} = \frac{\omega_o + \omega_i}{\|\omega_o + \omega_i\|_2}, \quad (2)$$

$$\mathcal{M}(\omega_o, \omega_i) = \frac{\mathcal{F}(\omega_o, \omega_i) \mathcal{G}(\omega_o, \omega_i)}{4(\mathbf{n} \cdot \omega_o)(\mathbf{n} \cdot \omega_i)}, \quad (3)$$

$$\mathcal{F}(\omega_o, \omega_i) = \mathbf{s} + (1 - \mathbf{s}) \cdot 2^{-(5.55473\omega_o \cdot \mathbf{h} + 6.8316)(\omega_o \cdot \mathbf{h})}, \quad (4)$$

$$\mathcal{G}(\omega_o, \omega_i) = \frac{\omega_o \cdot \mathbf{n}}{\omega_o \cdot \mathbf{n}(1 - k) + k} \cdot \frac{\omega_i \cdot \mathbf{n}}{\omega_i \cdot \mathbf{n}(1 - k) + k}, \quad (5)$$

$$k = \frac{(R + 1)^2}{8}, \quad (6)$$

$$\mathcal{D}(\mathbf{h}) = G(\mathbf{h}; \mathbf{n}, \frac{2}{R^4}, \frac{1}{\pi R^4}), \quad (7)$$

where \mathcal{F}, \mathcal{G} are Fresnel and shadowing terms, respectively, and \mathcal{D} is the normalized distribution function.

2. Additional synthetic examples



Figure 1: Results of our pipeline on synthetic data. For a novel test view, we compare our predicted image, estimated diffuse albedo, specular BRDF editing results and relighting results against the ground-truth rendered by Mitsuba [2].

3. Video demos

Please qualitatively inspect our results demonstrated in the supplemental video (included in the zip file) on the real-world SLF fish [6], DeepVoxels coffee [5], and Chips cans [4] datasets.

References

- [1] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, David Kriegman, and Ravi Ramamoorthi. Deep 3d capture: Geometry and reflectance from sparse multi-view images. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5960–5969, 2020. 1
- [2] Wenzel Jakob. Mitsuba renderer, 2010. <http://www.mitsuba-renderer.org>. 1
- [3] Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Learning to reconstruct shape and spatially-varying reflectance from a single image. *ACM Trans. Graph.*, 37(6):1–11, 2018. 1
- [4] Jeong Joon Park, Aleksander Holynski, and Steve Seitz. Seeing the world in a bag of chips. *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 2
- [5] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2437–2446, 2019. 2
- [6] Daniel N Wood, Daniel I Azuma, Ken Aldinger, Brian Curless, Tom Duchamp, David H Salesin, and Werner Stuetzle. Surface light fields for 3d photography. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 287–296, 2000. 2
- [7] Zhiming Zhou, Guojun Chen, Yue Dong, David Wipf, Yong Yu, John Snyder, and Xin Tong. Sparse-as-possible svbrdf acquisition. *ACM Transactions on Graphics (TOG)*, 35(6):1–12, 2016. 1