Practical and Accurate Reconstruction of an Illuminant's Spectral Power Distribution for Inverse Rendering Pipelines

Parisha Joshi
Clemson University
Clemson, SC, USA

Daljit Singh J. Dhillon
Clemson University
Clemson, SC, USA

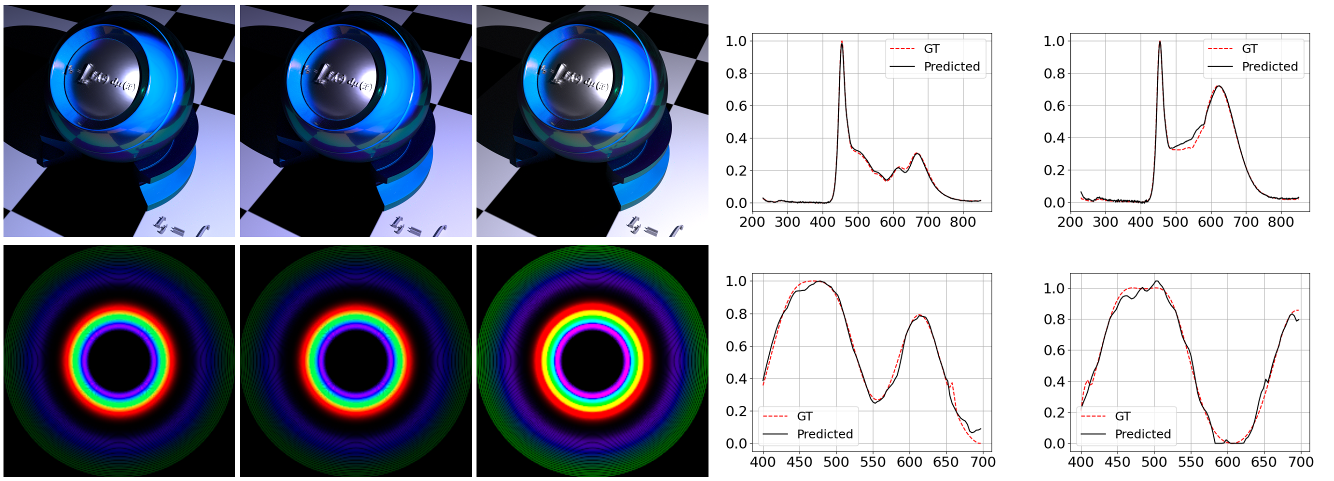

(a) Renderings: Ours Ground truth Reference RGB (b) Few predicted spectra in comparison with their *ground truths*

Figure 1: (a) We use a measured iridescent material BRDF [3] and a simulated CD grating's BRDF to validate our reconstructed spectra against the ground truth under rendering tasks. (b) Shows predicted and ground truth SPDs for four illuminants. X-axes represent the wavelengths in nanometers and Y-axes show relative intensities. The bottom two are synthetic and the top two are real world illuminants. Ground truth for them was acquired using a high-end spectrometer (Hopoocolor OHSP350UV 230–850**nm). The top left spectral profile was used to generate iridescent and CD rendering in figure (a). The reference RGB renderings exhibit visible mismatches in color tones as well as relative intensities. In comparison, our reconstructed spectra produce renderings that are qualitatively similar to corresponding ground truths.**

ABSTRACT

Inverse rendering pipelines are gaining prominence in realizing photo-realistic reconstruction of real-world objects for emulating them in virtual reality scenes. Apart from material reflectances, spectral rendering and in-scene illuminants' spectral power distributions (SPDs) play important roles in producing photo-realistic images. We present a simple, low-cost technique to capture and reconstruct the SPD of uniform illuminants. Instead of requiring a costly spectrometer for such measurements, our method uses a diffractive compact disk (CD-ROM) and a machine learning approach for accurate estimation. We show our method to work well with spotlights under simulations and few real-world examples. Presented results clearly demonstrate the reliability of our approach through quantitative and qualitative evaluations, especially in spectral rendering of iridescent materials.

KEYWORDS

Inverse rendering, spectral illumination, diffraction, SPDs, RNNs

ICVGIP'24, December 2024, Bangalore, India
2024. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

ACM Reference Format:
Parisha Joshi and Daljit Singh J. Dhillon. 2024. Practical and Accurate Reconstruction of an Illuminant's Spectral Power Distribution for Inverse Rendering Pipelines. In *Proceedings of 15th Indian Conference on Computer Vision, Graphics and Image Processing (ICVGIP'24).* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 BACKGROUND

Spectral rendering pipelines that produce high-quality, photo-realistic images require material properties and illumination sources to be represented with accurate spectral characterization. Recent advances in inverse rendering have mainly focused on multi-spectral acquisition of material appearance data [9, 10] for data-driven relighting, indirect environmental light maps [7] and parametric, spectral reflectance (BRDFs) [4] or scattering (BSDFs) [6] functions, in this context. In this paper, instead, we focus on accurate reconstruction of the uniform spectral power distribution (SPD) of common real-world illuminants that a content-creators may wish to replicate in their virtual setups. These SPDs are critical for accurate reconstruction of iridescence and structural coloration effects such as those on insect and reptile bodies [2], scratched or glinty

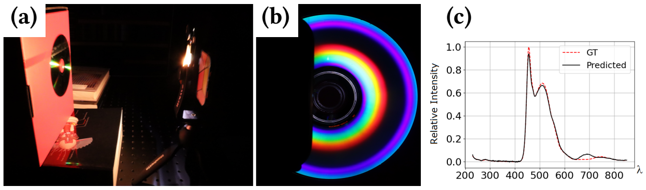

Figure 2: (a) Real World CD Capture setup with Canon Rebel T8i (b)With few adjustments of CD or illumination the rings were captured with the SPD shown in (c) measured with Hoppoocolor OHSP350 .

surfaces [13, 14], Bragg mirrors [5], layered materials with specular sheens [1]. While spectrometers exist to measure illuminant SPDs accurately, they are expensive and not easily available to most artists. We thus devise a simple, effective and affordable method that can be adapted for any camera with one known light source and a set of known transmissive color filters.

2 PROPOSED METHOD

Spectrometers [15] and hyper-spectral imaging methods [8] commonly rely on a diffractive optical element (filter) to profile the spectral distribution of the incident radiance/s. This inspired us to use a simple, diffractive compact disk (CD-ROM) in devising an image-based method for spectral profiling of light sources. CD-ROM are inexpensive, standardized, high-quality grating constructs that are also easily available. Thus, using them as the essential diffractive element has allowed us in devising a simple, reliable and cost-effective method for estimating the SPDs of light sources.

Imaging Setup: We have experimented in a simulated imaging environment (PBR toolkit [11]) that supports spectral rendering. In the next section, we discuss real-world, practical adaption of our method along with a few results. We illuminate an unwritten CD-ROM with a spotlight that is placed fronto-parallelly to it and capture its appearance with a camera that is also placed fronto-parallelly to the CD. The optical axis of the camera passes through the CD's center. The unwritten CDs have fixed circular tracks as a single diffractive layer of known spacing between the gratings. We simulate the reflectance function for the CD gratings by implementing the analytic BRDF given by Toisoul et al. [12]. We set the grating gap to $a = 0.5\mu m$ and its maximum height as $h0 = 0.15\mu m$. Our synthetic camera has known color filter functions and its exposure settings are fixed to avoid color and brightness saturation. Using this setup, we developed a data-driven model for estimating the SPD of any unknown spotlight in exactly the same configuration.

Machine Learning Step: We illuminate this CD with a set of synthetic SPDs to produce a set of images. The SPDs are generated to smooth variations, random noise as well as sharp spikes in different combinations to mimic a large variety of real world SPDs. Each SPD is put on a relative scale to have its peak value as 1. With the data set corresponding of 5000 SPD we train a multilayer perceptron (MLP) network for learning the supervised regression process. Each image is pre-processed to mask out the inner and outer circular regions around the CD that do not contribute to the learning process. With Adam optimizer and leaky RELU activation function, the model generally requires training for up to 100000 epochs. We use a batch size of 64. We train with 4000 random SPD samples and validate against the remaining 1000 SPD samples.

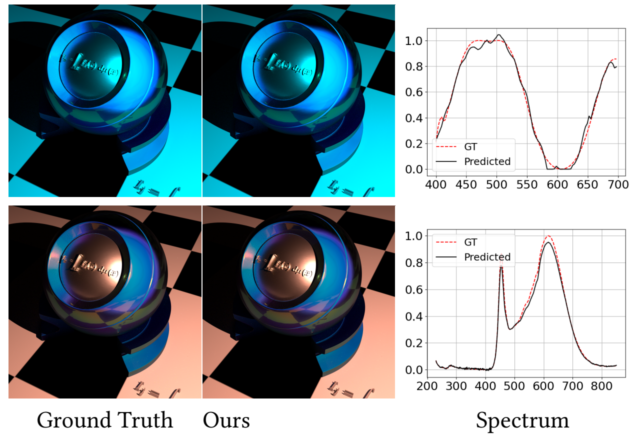

Figure 3: Two more examples in SPD reconstruction and the impact on rendering accuracy. Rendering PSNR for top = 46.596 dB and for the bottom = 55.352 dB

Table 1: Results of the MLP training step

Metric (Average)	Training (4000 sample SPDs)	Validation (1000 unknown SPDs)
MAE	0.0466	0.06771
RMSE	0.007099	0.0105
Correlation	0.936503	0.86411

Results: We performed quantitative evaluations using statistical measures such as MAE, RMSE, and the correlation factor between the ground truth SPD and the respective prediction. Average values across the validation set for all these measures are shown in Table 1. The last column clearly indicates that our MLP network predicts unknown SPDs with very high accuracy. We also, compared our method's performance under rendering. Figure 1 and Figure 3 show that our renderings are visually indistinguishable from the ground truth. Rendering PSNRs are also indicated in Figure 3.

3 DISCUSSION AND FUTURE WORK

The training and validation results conclude that the spotlight illumination spectra can be reconstructed with minimal error using diffraction CD, under simulations. We also experimented to adapt our method for real-world illuminants. Figure 2 shows our imaging setup, one example case and the resulting SPD. For training with the real-world SPDs, we used the spectrometer to obtain the ground truth. We find out method to produce promising initial results that match up to our findings under PBRT simulations. In the future, we want to thoroughly validate our method against a large set of real-worlds images. Also, we would want to extend our method to work with any camera that is color-calibrated using a single MLP. Lastly, we conjecture that our method can be adapted to work with general placements and non-uniform illuminants. Such improvements have the potential to support spectral environment maps for photo-realistic rendering of nuanced structural colorations.

ACKNOWLEDGEMENTS. This work is fully supported by NSF Grant #2007974.

REFERENCES

[1] Laurent Belcour and Pascal Barla. 2017. A practical extension to microfacet theory for the modeling of varying iridescence. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–14.

[2] D. S. Dhillon, J. Teyssier, M. Single, I. Gaponenko, M. C. Milinkovitch, and M. Zwicker. 2014. Interactive Diffraction from Biological Nanostructures. *Computer Graphics Forum* 33, 8 (2014), 177–188. https://doi.org/10.1111/cgf.12425

[3] Jonathan Dupuy and Wenzel Jakob. 2018. An adaptive parameterization for efficient material acquisition and rendering. *ACM Trans. Graph.* 37, 6, Article 274 (dec 2018), 14 pages. https://doi.org/10.1145/3272127.3275059

[4] Jonathan Dupuy and Wenzel Jakob. 2018. An adaptive parameterization for efficient material acquisition and rendering. *ACM Transactions on graphics (TOG)* 37, 6 (2018), 1–14.

[5] Gary Fourneau, Romain Pacanowski, and Pascal Barla. 2024. Interactive Exploration of Vivid Material Iridescence based on Bragg Mirrors. 43, 2 (2024).

[6] Yuliya Gitlina, Giuseppe Claudio Guarnera, Daljit Singh Dhillon, Jan Hansen, Alexander Lattas, Dinesh Pai, and Abhijeet Ghosh. 2020. Practical measurement and reconstruction of spectral skin reflectance. *Computer Graphics Forum* 39, 4 (2020), 75–89. https://doi.org/10.1111/cgf.14055

[7] Giuseppe Claudio Guarnera, Yuliya Gitlina, Valentin Deschaintre, and Abhijeet Ghosh. 2022. Spectral Upsampling Approaches for RGB Illumination. In *Eurographics Symposium on Rendering*, Abhijeet Ghosh and Li-Yi Wei (Eds.). The Eurographics Association. https://doi.org/10.2312/sr.20221150

[8] Daniel S. Jeon, Seung-Hwan Baek, Shinyoung Yi, Qiang Fu, Xiong Dun, Wolfgang Heidrich, and Min H. Kim. 2019. Compact snapshot hyperspectral imaging with diffracted rotation. *ACM Trans. Graph.* 38, 4, Article 117 (jul 2019), 13 pages. https://doi.org/10.1145/3306346.3322946

[9] Chloe LeGendre, Kalle Bladin, Bipin Kishore, Xinglei Ren, Xueming Yu, and Paul Debevec. 2018. Efficient multispectral facial capture with monochrome cameras. In *ACM SIGGRAPH 2018 Posters*. Association for Computing Machinery, New York, NY, USA.

[10] Chloe LeGendre, Xueming Yu, and Paul Debevec. 2016. Efficient Multispectral Reflectance Function Capture for Image-Based Relighting. In *Color and Imaging Conference*, Vol. 24. Society for Imaging Science and Technology, 47–58.

[11] Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2023. *Physically based rendering: From theory to implementation.* MIT Press.

[12] Antoine Toisoul, Daljit Singh Dhillon, and Abhijeet Ghosh. 2018. Acquiring spatially varying appearance of printed holographic surfaces. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–16.

[13] Sebastian Werner, Zdravko Velinov, Wenzel Jakob, and Matthias B Hullin. 2017. Scratch iridescence: Wave-optical rendering of diffractive surface structure. *ACM Transactions on Graphics (TOG)* 36, 6 (2017), 1–14.

[14] Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering specular microgeometry with wave optics. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–10.

[15] ZEISS. 2024. ZEISS Spectrometer Modules: Compendium of products, electronic components and software solutions. https://asset-downloads.zeiss.com/catalogs/download/spc/adb036b4-b2e5-48d2-ae76-3903ddfe9c06/BR_Spectrometer_Modules_EN.pdf.