Perception-based Optimization of Wavelength Sampling Distributions for Spectral Rendering

Author: Camil Cristian Dobos
Supervisor(s): Elmar Eisemann
Christoph Peters
Michael Weinmann

1. Introduction and Background

- Spectral Monte Carlo rendering is a state-of-the-art technique for physically based rendering. It estimates the radiance entering a pixel by approximating the rendering equation [2]:

$$I_j = \int_\Lambda \int_\Omega f_j(\bar{x}, \lambda)\, \mathrm{d}\mu(\bar{x})\, \mathrm{d}\lambda,$$

- It can accurately simulate complex optical phenomena by treating light as a spectrum of electromagnetic radiation, instead of relying on trichromatic values such as (R, G, B).
- However, convergence is slow, particularly due to the need to sample the wavelength domain during the Monte Carlo procedure. To address this, multiple techniques were proposed, most of them involving sampling multiple wavelengths per path.
- Most notably, hero wavelength spectral sampling [1] consists of sampling a primary hero wavelength for each path, and additional wavelengths are placed at equal distances from the hero wavelength to ensure they always cover the visible spectrum evenly.

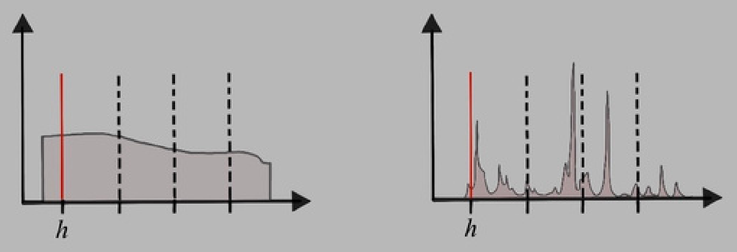

Figure 1: Sampling multiple wavelengths in a smooth versus "spiky" illuminant.

- Although this provides robust spectral coverage, it may prove inefficient in the case of complex illuminant spectra, as pictured in Figure 1.

2. Research Question

The main contribution of this study consists of:

- A preprocessing step, in which we optimize a set of probability distributions for sampling multiple wavelengths from an illuminant spectrum, based on the perceived color difference.
- Assessment of whether this method can improve the perceptual quality of the rendered image in single-emitter scenes when compared to sampling wavelengths uniformly.

3. Methodology

- For a given path, the goal is to sample N wavelengths from N probability distributions $p_0(\lambda), p_1(\lambda), \ldots, p_{N-1}(\lambda)$ using a single random number u. Then a Monte Carlo estimate is computed using multiple importance sampling [2]:

$$c_l \approx \sum_{j=0}^{N-1} \frac{h_l(\lambda_j(u)) R(\lambda_j(u)) I(\lambda_j(u))}{\sum_{k=0}^{N-1} p_k(\lambda_j(u))}$$

- where $l \in \{0, 1, 2\}, h_0 = \bar{x}, h_1 = \bar{y}, h_2 = \bar{z}$ are the color matching functions [3], R is the product of the reflectance spectra encountered along the path, I is the illuminant spectrum, and $\lambda_j(u)$ are the inverse CDFs of the optimized PDFs.
- The sampling strategy should be independent of the reflectance spectra, so we set its value to 1. Therefore the expected value of the estimator becomes the color of the illuminant in the XYZ color space [3].
- Ideally, the estimator dependence on u is small, such that the perceived variance is minimal. This can be visualized in Figure 2.
- Therefore, we optimize the probability distributions based on the objective function:

$$\int_0^1 \Delta_E \left(\left(\sum_{j=0}^{N-1} \frac{h_l(\lambda_j(u)) I(\lambda_j(u))}{\sum_{k=0}^{N-1} p_k(\lambda_j(u))} \right)_{l=0}^{2}, c_l \right) du$$

- where ΔE measures the perceived color difference. [8]
- An initial set of probabilities is handed to a non-linear optimizer, which then iteratively optimizes them to minimize the objective function.

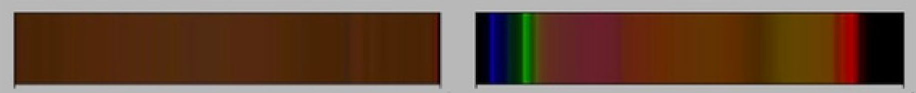

Figure 2: Illustration of the perceptual variance of the estimator using a color bar. The resulting color of the estimator is plotted for a very large number of samples in the interval [0, 1[. Ideally, this color bar should be constant. The left color bar is obtained after optimization, while the right color bar uses the initial distributions

4. Implementation

- The optimization was implemented using the Python libraries SciPy and Ipopt [4], which provide a solid collection of solvers suitable for high-dimensional, non-linear optimization problems.
- For rendering, PBRT-v4 [5] was used, and modified to implement the custom wavelength sampling procedure.

5. Evaluation and Results

- Our method was evaluated against the standard wavelength sampling procedure in PBRT, which first samples a wavelength uniformly. The other wavelengths are selected by placing them at a constant distance from the first one, similar to hero wavelength sampling [1].
- We used a number of pre-made scenes from Benedikt Bitterli's website [6]. The scenes selected have either a single emitter, or a small number of emitters, all of them set to use the same emission spectrum.
- A variety of illuminants with different spectral power distributions were selected from the Lamp Spectral Power Distribution Database (LSPDD) [7].
- The performance was assessed by measuring the perceptual error of the proposed method against a ground truth image rendered at a high sample count using the baseline method. The result was then compared with the error obtained by the baseline method, at the same sample count. Some results of the evaluation are listed in Table 1:

Scene	Illuminant	SPP	Theirs		Ours	
			Avg. ΔE	Time (s)	Avg. ΔE	Time (s)
Cornell Box	Compact Fluorescent	512	2.5311	17.75	**0.9703**	17.14
Glass of Water	Compact Fluorescent	512	5.6761	19.57	**2.2335**	19.17
Staircase	Philips High Bay	512	0.3791	26.58	**0.1908**	25.57
Veach MIS	Philips High Bay	512	0.8455	8.55	**0.3580**	8.56

Table 1: Comparison of perceptual difference (ΔE) and render times across various scenes and illuminants between uniform sampling and our custom method.

- Furthermore, as Table 1 suggests, the only additional computational cost of our approach is the optimization step, which might take longer. However, this is a one-time, offline process performed once for each illuminant spectrum, and does not impact render performance.

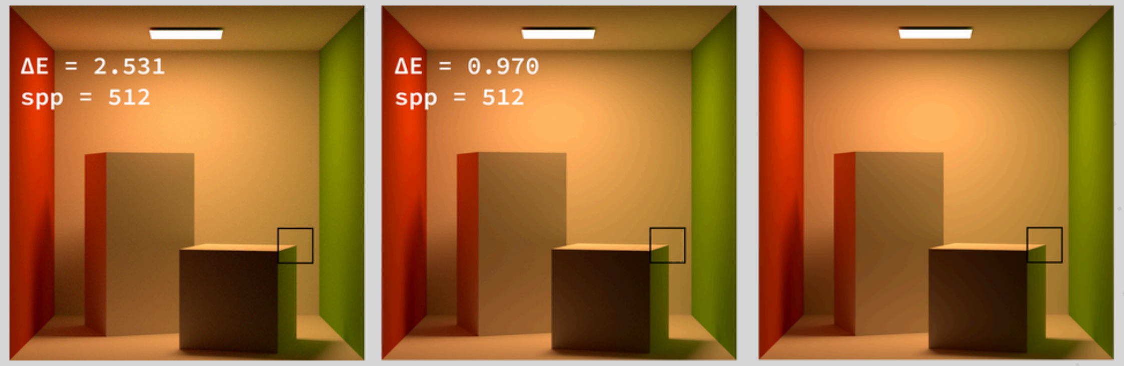

Figure 3: Comparison of baseline versus our method against the ground truth (right), for the Cornell Box scene [6], using a Globe Twister illuminant (LSPDD index 2723) at 512 samples per pixel.

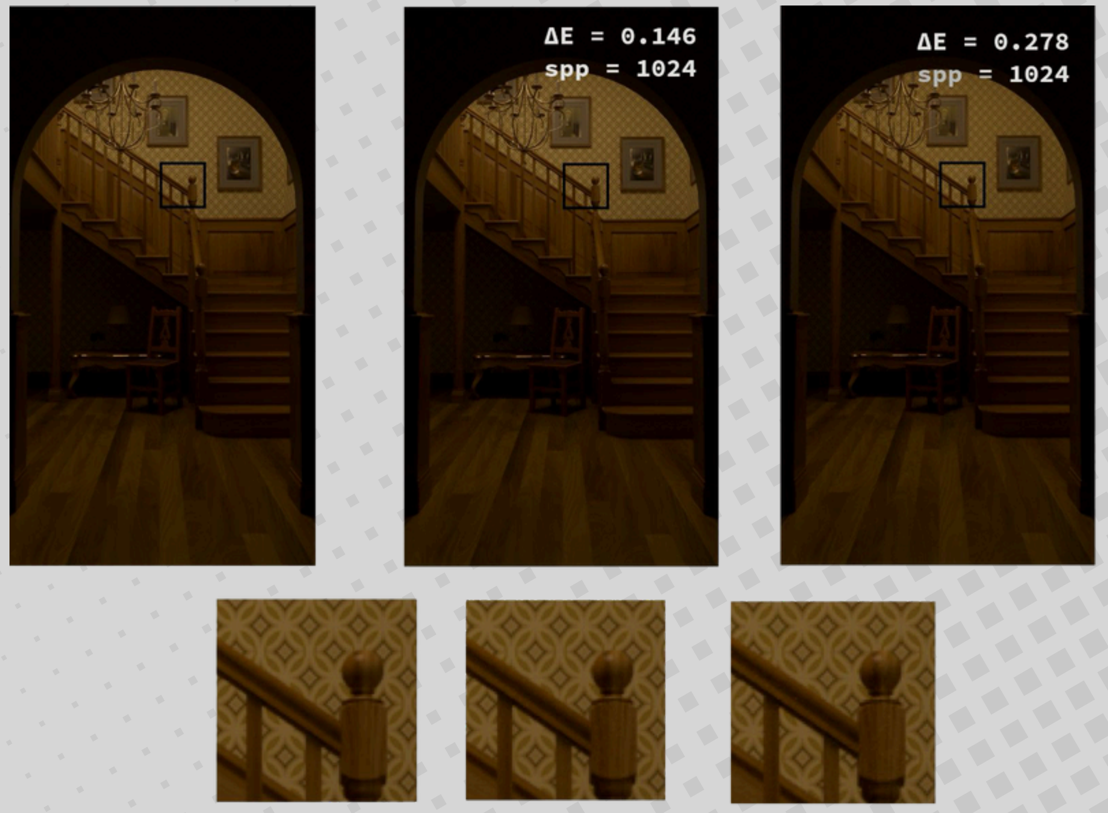

Figure 4: Comparison of our method (middle) versus the baseline (right) against the ground truth for the wooden staircase scene [6] using the Globe Twister emitter at 1024 samples per pixel.

6. Conclusions

- The evaluation demonstrates that this method can significantly reduce perceptual error compared to uniform sampling in single-emitter scenes, particularly scenes featuring illuminants with complex spectral power distributions, while incurring no computational overhead during rendering.

9. Refrences

- [1] Wilkie, A., Nawaz, S., Droske, M., Weidlich, A., & Hanika, J. Hero wavelength spectral sampling. In: In Proceedings of the 25th eurographics symposium on rendering. EGSR '14. Lyon, France: Eurographics Association, 2014, 123–131.https://doi.org/10.1111/cgf. 12419
- [2] Veach, E. (1998). Robust monte carlo methods for light transport simulation [Doctoral dissertation] [AAI9837162]. Stanford University.
- [3] CIE. (1932). Proceedings of the commission internationale de l'eclairage, 8th session [Standard adopted in 1931]. Cambridge University Press.
- [4] Wachter, A., & Biegler, L. T. (2005). On the implementa- ¨ tion of an interior-point filter line-search algorithm for large-scale nonlinear programming. Mathematical Programming, 106(1), 25–57. https://doi.org/10. 1007/s10107-004-0559-y
- [5] Pharr, M., Jakob, W., & Humphreys, G. (2023). Physically based rendering: From theory to implementation (4th). MIT Press. https://pbrt.org
- [6] Bitterli, B. (2016). Rendering resources [https://benediktbitterli.me/resources/].
- [7] Roby, J., & Aube, M. (2019). ´ Lamp spectral power distribution database (lspdd) [Accessed 2025-06-21]. https: //lspdd.org/
- [8] Sharma, G., Wu, W., & Dalal, E. N. (2005). The ciede2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. Color Research & Application, 30(1), 21–30. https://doi.org/10.1002/col.20070