Achieving Perceptually-Acceptable Early Renders in Spectral Progressive Rendering by Introducing Bias

Tan Dao
Supervisor(s): Mark van de Ruit, Elmar Eisemann
EEMCS, Delft University of Technology, The Netherlands

June 19, 2022

A Dissertation Submitted to EEMCS faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering

Achieving Perceptually-Acceptable Early Renders in Spectral Progressive Rendering by Introducing Bias

T. Dao, M. van de Ruit, and E. Eisemann

Delft University of Technology, The Netherlands

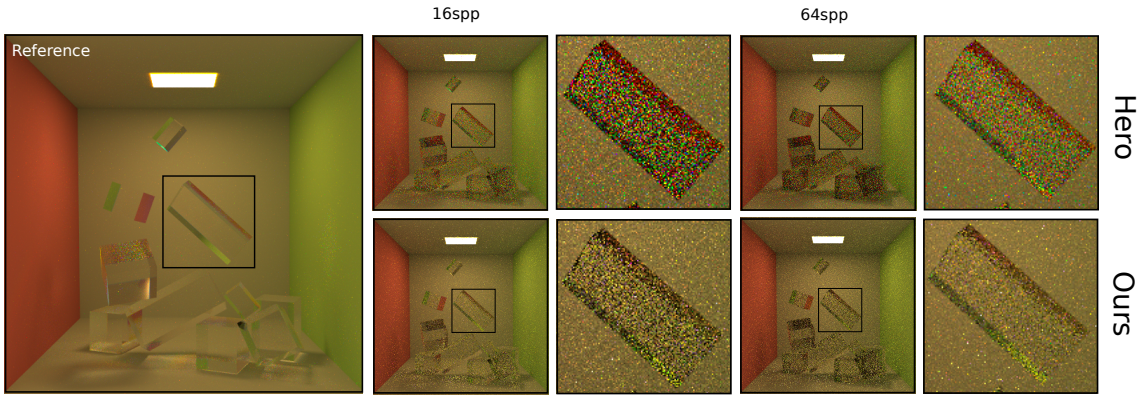

Figure 1: *Our method decreases variance and increases the perception of the early renders. We compare our method with the current state-of-the-art (hero wavelength spectral sampling) [WND*14] at 16 and 64 samples per pixel, respectively. The selected area depicts a cube made of wavelength-dependent material.*

Abstract

Spectral Monte-Carlo rendering can simulate advanced light phenomena (e.g., dispersion, caustics, or iridescence), but require significantly more samples compared to trichromatic rendering to obtain noise-free images. Therefore, its progressive variant typically exhibits an extreme amount of chromatic noise in early renders. To that end, we propose a two-stage progressive approach. We initially restrict the original wavelength distribution, then slowly relax it. In the process of relaxing the range of wavelengths, all wavelengths that are outside of that restricted range will be propagated. Thereby, we lower variance and increase the perception of these early renders with little overhead.

1. Introduction

Ray tracing is a popular light simulation technique for rendering high-quality, photo-realistic images. While traditional forms of ray tracing simulate light in the RGB color space, it is not an accurate representation of light, which is a continuous spectrum of wavelengths. With this limitation, a number of advanced light phenomena, such as dispersion, caustics, or iridescence can only be properly handled with spectral rendering.

Unlike trichromatic rendering, spectral rendering utilizes the full range of visible wavelength to faithfully represent light. Monte-Carlo path tracing are particularly attractive as spectral rendering can be achieved by simply adding the spectral domain as another dimension of the path-space integral equation [Vea98], which is straight-forward and un-

biased. However, this infers a notable increase in required sampling rates to obtain a noise-free result.

While recent approaches reduce chromatic noise by sampling multiple wavelengths for each path [EM99, RBA09, WND*14, WGGH20], the number of sampled paths required still remains drastically higher than that of tristimulus rendering and is, therefore, undesirably slow. Achieving global illumination at interactive rates with spectral rendering remains a challenging problem.

A potential approach to partially solve this problem is to transform normal path tracing into a progressive process. Common approaches involve reduced sample rates or taking a selected subset of image pixels. If the quality of the generated image is not satisfactory, additional samples are created to output a better image. This process will iterate until the image has reached the desired quality. In the case of spectral rendering, however, the first few iterations commonly display a notable increase in variance compared to tristimulus rendering.

In this paper, we present a reformulation of spectral rendering into a perceptually progressive version. This rendering method focuses on the more important ranges of wavelengths, aiming to achieve more visually acceptable renders after the first few iterations. As time goes on, the method will converge to an output identical to what can be obtained by using normal path tracing.

Section 2 will cover spectral light transport equations and provide an overview of the state-of-the-art. We next expand on how we approached the problem and how we formalized our solution in Section 3. Section 4 goes into detail of our implementation and evaluate it in a variety of scenarios in Section 5. Section 6 reflects on the ethical aspects of our research and discuss the reproducibility of our methods. Lastly, Section 7 and Section 8 will discuss our results, improvements and conclude the research.

2. Background and Related Work

Spectral Light Transport

To obtain a physically-based render, the general path-space light transport equation [Vea98] is evaluated by integrating over all light paths in a scene. Extending upon this equation by having an additional integration over the spectral domain, the spectral radiance of a pixel, denoted by I, can be described as:

$$I = \int_{\Lambda}\int_{\Omega} f(\bar{x},\lambda)d\mu(\bar{x})d\lambda \tag{1}$$

where Λ denotes the *spectral domain* and Ω is the *path space* of all possible light paths $\bar{x} = x_0,...,x_{n-1}$ of finite length n, starting from a light source, travelling through the scene and ending at a camera pixel. $f(\bar{x},\lambda)$ then describes the measurement contribution of path $\bar{x}$ given a wavelength λ.

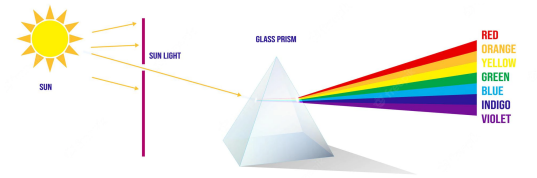

Figure 2: *We show an example of a prism, where different wavelengths are scattered into different directions when entering the prism.*

An estimation of Equation 1 can be given by Monte-Carlo methods to obtain an unbiased yet naive result:

$$\langle I \rangle = \frac{1}{N}\sum_{i=1}^{N}\frac{f(\bar{x}_i,\lambda_i)}{p(\bar{x}_i,\lambda_i)} \tag{2}$$

where N is the *sampling rate* and $p(\bar{x}_i,\lambda_i)$ is the *probability density* of the sample pair $(\bar{x}_i,\lambda_i)$ which can be further decomposed as:

$$p(\bar{x},\lambda) = p(\lambda)\cdot p(\bar{x}|\lambda) \tag{3}$$

The naive approach above only takes into account 1 wavelength per light path. This will inevitably introduce large amount of color noise and is extremely wasteful in cases where a path does not encounter wavelength-dependent interactions. Wavelength dependency here is defined as interactions with a dispersive material, whose scattering distribution function depends on the wavelength, with an example shown in Figure 2.

Multiple Wavelength Sampling

Evans and McCool [EM99] and Radziszewski et al. [RBA09] introduced stratified wavelength sampling, allowing a sampled path to be reused for propagating multiple wavelengths at once (called a cluster). Wilkie et al. [WND*14], more recently, introduces *hero wavelength spectral sampling* (HWSS). A single initial *hero wavelength* is chosen to be responsible for path propagation, and that path is then taken into account for multiple wavelengths. As the same path can now be sampled from multiple valid wavelengths, it is important to weight them appropriately using *multiple importance sampling* [Vea98], leading to the following estimator:

$$\langle I \rangle = \frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{C}\frac{f(\bar{x}_i,\lambda_i^j)}{\sum_{k=1}^{C}p(\bar{x}_i,\lambda_i^k)} \tag{4}$$

where C is the number of wavelengths valid per path.

While reusing sampled paths for multiple wavelengths does not improve the performance for perfectly specular material, it greatly speeds up the convergence rates for diffuse and glossy surfaces.

3. Methodology

The principles of Monte-Carlo estimation applied to Equation 2 states that, the better $p(\bar{x},\lambda)$ matches the integrand, the faster the image converges. However, since some wavelengths contribute more to the image than others, it is more efficient to focus on sampling those areas in the early iterations of progressive rendering, where a smooth and low-variance estimate of the scene is preferable.

3.1. Restricting and Relaxing of Wavelength distributions

HWSS gave us the means to improve convergence rates for non-specular wavelength-dependent materials. However, specular materials still project an extreme amount of noise, especially in the first few iterations. Aiming to both utilize HWSS and alleviate the initial noise in specular materials, we propose a two-stage progressive approach: *restriction* and *relaxation* of wavelength sampling distributions.

In the *first stage*, we restrict the range of wavelengths that can be sampled, defined as $p_{restricted}(\alpha,\lambda)$:

$$p_{restricted}(\alpha,\lambda) = p(\lambda)\cdot restrict(\alpha,\lambda) \tag{5}$$

where $restrict(\alpha,\lambda)$ is the restriction function and α is its extent. As $p_{restricted}(\alpha,\lambda)$ will be used as a probability distribution function, there is a need for normalization:

$$\int_{\Lambda} p_{restricted}(\alpha,\lambda) = 1 \tag{6}$$

transforming the PDF of sampling a pair $(\bar{x},\lambda)$ to:

$$p(\bar{x},\lambda) = p_{restricted}(\alpha,\lambda)\cdot p(\bar{x}|\lambda) \tag{7}$$

In the *second stage*, the restricted range gradually relaxes as the algorithm progresses. This expands Equation 2 to the following:

$$\langle I\rangle = \frac{1}{N}\sum_{i=1}^{N}\frac{f(\bar{x}_i,\lambda_i)}{p_{restricted}(\alpha_i,\lambda_i)\cdot p(\bar{x}_i|\lambda_i)} \tag{8}$$

where α_i increases in subsequent iterations. We define a threshold $N_{thold} < N$ that indicates when the relaxation process is completed. As the current iteration approaches N_{thold}, the original wavelength distribution will be regained:

$$\begin{aligned}&\lim_{i\to N_{thold}} p_{restricted}(\alpha_i,\lambda_i)\\ =&\lim_{i\to N_{thold}} p(\lambda_i)\cdot restrict(\alpha_i,\lambda_i)\\ =&\, p(\lambda_i)\end{aligned} \tag{9}$$

With the definition of the restriction and relaxation scheme, shown in Figure 3, a wavelength sample can now be drawn from $p_{restricted}(\alpha_i,\lambda_i)$. The sample will then be evaluated and recorded for its contribution, and an updated render will be written accordingly.

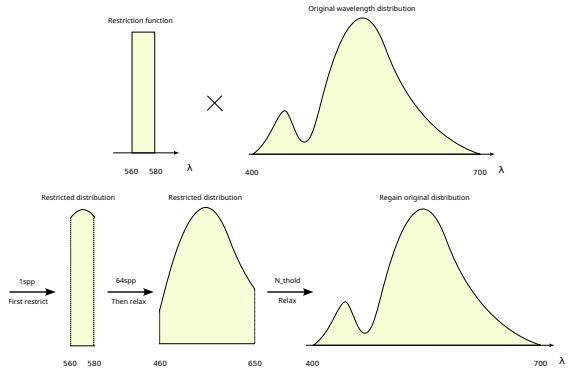

Figure 3: ***Method overview.*** *We initially restrict the original wavelength distribution and then relax it gradually. As the current iteration approaches N_{thold}, we regain the orginal wavelength distribution.*

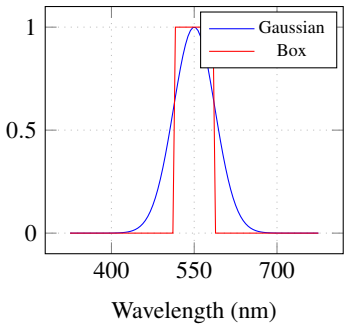

Figure 4: ***Choosing kernels.*** *We show the chosen unimodal functions used to learn their performance.*

3.2. Choosing Restriction Function

Unimodal Function

Out of different unimodal functions, more commonly known as kernels, the Gaussian and box function are chosen to be our restriction function, shown in Figure 4. One represents curves, performing smoothing restrictions while the other represents uniform distributions, performing hard cut-offs.

Each distribution is then multiplied with the original SPD to give $p_{restricted}(\alpha,\lambda)$ in Equation 5. Through experimentation, it can be seen that the box function gave the best performance when chosen as a restriction function, and sampling wavelengths from this restricted distribution gives decreased variance, as shown in Figure 5. On the other hand, the Gaussian function tends to introduce more 'fireflies' and offers no improvement in noise reduction.

Even though variance is indeed reduced, a unimodal function is a poor choice as it tints the output image with the typical color of the initial wavelength range that is hard to remove in later renders. Moreover, it reduces the effective-

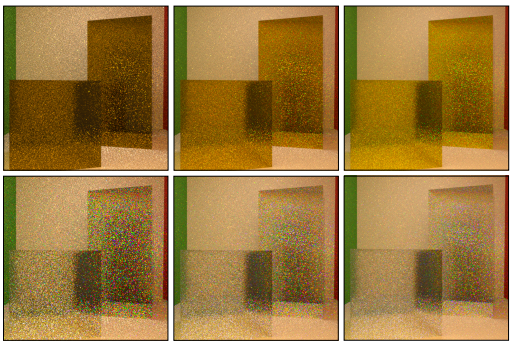

Figure 5: ***Unimodal restriction function.*** *Using a unimodal function as our restriction function, we compare our method with HWSS at 16, 64, and 128 samples per pixel, respectively. The 'peak' is placed at 580nm and starts with an initial width of 20nm.*

ness of HWSS as stratification is limited to a small range of wavelengths. Therefore, there is a need to use multimodal functions.

Multimodal Function

We now consider a distribution with an arbitrary number of box functions placed at different wavelength ranges. The ideal number of 'peaks', through experimentation, is three. It is also important to place these 'peaks' at wavelength ranges that roughly corresponds to red, green, and blue. This can be explained by the trichromatic theory of color vision, stating that all hues can be derived from a mixture of three primary colors [Max57, You02]. Increasing the number of peaks beyond this results in increased variance, as it extends the range of wavelengths that our method can sample from.

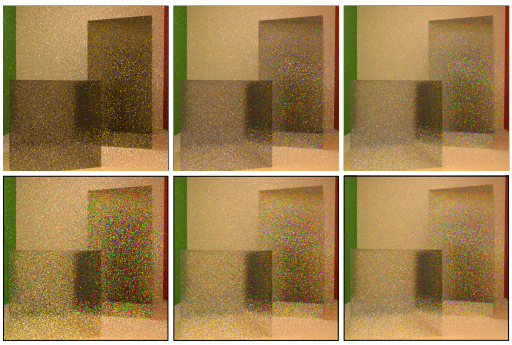

Figure 6: ***Multimodal restriction function.*** *Using a multimodal function as our restriction function, we compare our method with HWSS at 16, 64, and 128 samples per pixel, respectively. The 'peaks' are placed at 465, 570 and 580nm and start with an initial width of 20nm.*

Figure 6 showed that increasing the number of peaks greatly improved the transparency and color of wavelength-dependent materials. However, this multimodal function still suffers from the same problem as the uniform unimodal function. It resulted in the dim appearance of wavelength-dependent materials in initial renders. One way to alleviate this is by scaling the PDFs appropriately as we limit the range of wavelengths. However, this approach's resulting variance is very dependent on the placement of the 'peaks', as shown in Figure 7. Therefore, a different approach is necessary.

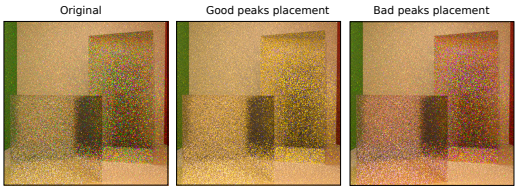

Figure 7: ***Scaling PDFs.*** *We compare examples of when good and bad 'peaks' are chosen, which is dependent on the input scene.*

3.3. Propagating Unused Wavelengths

We now propagate all the wavelengths that are outside of the restricted range and add their contribution to a light path. More formally, if the current number of iterations is lower than N_{thold}, then we extend Equation 8 to:

$$\langle I \rangle = \frac{1}{N_{thold}} \sum_{i=1}^{N_{thold}} \left[\frac{f(\bar{x}_i, \lambda_i)}{p_{restricted}(\alpha_i, \lambda_i) \cdot p(\bar{x}_i | \lambda_i)} + (1 - restrict(\alpha_i, \lambda_i)) \cdot \frac{f(\bar{x}_i)}{p(\bar{x}_i)} \right]. \quad (10)$$

Otherwise, if the iterations exceeded N_{thold} (i.e. $restrict(\alpha, \lambda) = 1$), then $\langle I \rangle$ again relies on the unbiased Equation 8.

By introducing a mixture of wavelength-dependent term and wavelength-independent term in early renders, we unavoidably introduce bias and any spectral effects outside the restricted range of wavelengths are lost. This is more than acceptable, however, as this equation takes advantage of the fact that spectral effects are difficult to observe in initially undersampled renders. This bias will disappear as the algorithm progresses, and at which point, the correct spectral phenomena will be given.

Integration with multiple wavelength sampling

Reusing a light path for multiple wavelengths as in HWSS [WND*14] remains an important technique to reduce chromatic noise for non-specular material, especially in the early renders. Therefore, it is essential that we combine our

method with HWSS. To that end, we will apply our bias term that extends Equation 4 to:

$$\langle I \rangle = \frac{1}{N_{thold}} \sum_{i=1}^{N_{thold}} \left[\left(\sum_{j=1}^{C} \frac{f(\bar{x}_i, \lambda_i^j)}{\sum_{k=1}^{C} p_{restricted}(\alpha_i, \lambda_i^k) \cdot p(\bar{x}_i | \lambda_i^k)} \right) + (1 - restrict(\alpha_i, \lambda_i)) \cdot \frac{f(\bar{x}_i)}{p(\bar{x}_i)} \right]. \tag{11}$$

4. Implementation

Our progressive method was implemented in PBRT-v3 [JPH16], a C++ ray-tracer. It stores spectral power distributions by splitting them over 60 discrete bins from 400 - 700nm. However, PBRT-v3 does not support spectral rendering. We first added a simple simulation of dispersion in wavelength-dependent material interactions by extending specular and non-specular dielectrics BSDFs based on Cauchy's equation [JW01]. HWSS [WND*14] and CMIS [WGGH20] are then incorporated, extending the default path tracer integrator, to implement spectral rendering. HWSS is then transformed into a progressive renderer, where an image is simply written to disk each time a sample contributes to a pixel, and *tev* [Mül22] is used to continuously read the updated image. Lastly, we restrict and relax the original wavelength distribution depending on the current iteration of the progressive method.

5. Evaluation

Our progressive method is evaluated in three test scenes shown in Figure 9 and Figure 10. The first two scenes contain multiple dispersive cubes that would benefit greatly from using our progressive approach. One scene uses the D65 standard illuminant while the other uses an LED light, sourced from the LSPDD [RA]. Their distributions are shown in Figure 8. Finally, the RUBIK scene consists of multiple stacking dispersive glasses that exhibit strong spectral phenomena under D65. Using HWSS for this scene is more beneficial for the viewers. The reference images for these scenes are rendered with sufficient samples per pixel (N = 65K).

5.1. Method evaluation

To compare our approach against HWSS, six different iterations of render were chosen: 16, 32, 64, 128, 256, and 512 samples per pixel (spp). Our comparisons only go up to 512spp as a sufficiently converged render of the scene should already be obtained at this point. Then, our method falls back to HWSS. The default settings for HWSS were used (4 wavelengths per path) as in the original paper [WND*14]. Same-iteration instead of time-based comparison was performed as the process of relaxing and propagating unused

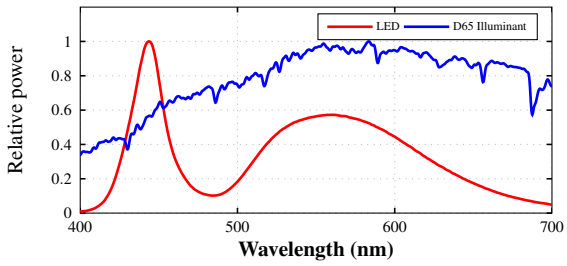

Figure 8: ***Emission spectra.*** *We show used spectra obtained from the LSPDD [RA]. Note the differences between the D65 illuminant and the LED emitter.*

wavelengths produces negligible overhead to the total rendering time. We report errors using the perceptually-based structural similarity metric (SSIM) [WBSS04], but also the commonly-used mean-squared error (MSE) in Figure 9 and Figure 10.

Comparison with previous approaches

We first make the comparisons in 2 scenes with varying emitter spectral distribution (LED and D65), as shown in Figure 9. Even though emitter sampling is near-optimal in these scenes, its initial renders of wavelength-dependent materials remain extremely noisy, and we expect our method to perform better. As demonstrated, our method results in higher SSIM and lower or similar error rates in all cases. Specifically, an improvement in the uniform emitter (D65) scene can be seen with a 10% overall increase in SSIM for 16spp while MSE stays relatively the same. Our method also offered the same improvement for the LED scene, with a 9% SSIM increase. The benefits of our method gradually disappear as we relax the range of wavelengths to 512spp, where our method returns similar result to that of HWSS.

Early dispersive effects

We next go through a scene where many dispersive hollow cubes are stacked on top of one another, shown in Figure 10. This resulted in strong dispersive effects, which can be observed clearly in the early renders. While our method still offers an improvement for the scene, it drives away any of the desired spectral effects, making the initial renders' inaccuracies visible. Moreover, even though the dispersive effect will be clearer as the algorithm progresses, it is not noticeable. Therefore, in such a scene where spectral phenomena are clearly displayed, it would be more appropriate to use default HWSS.

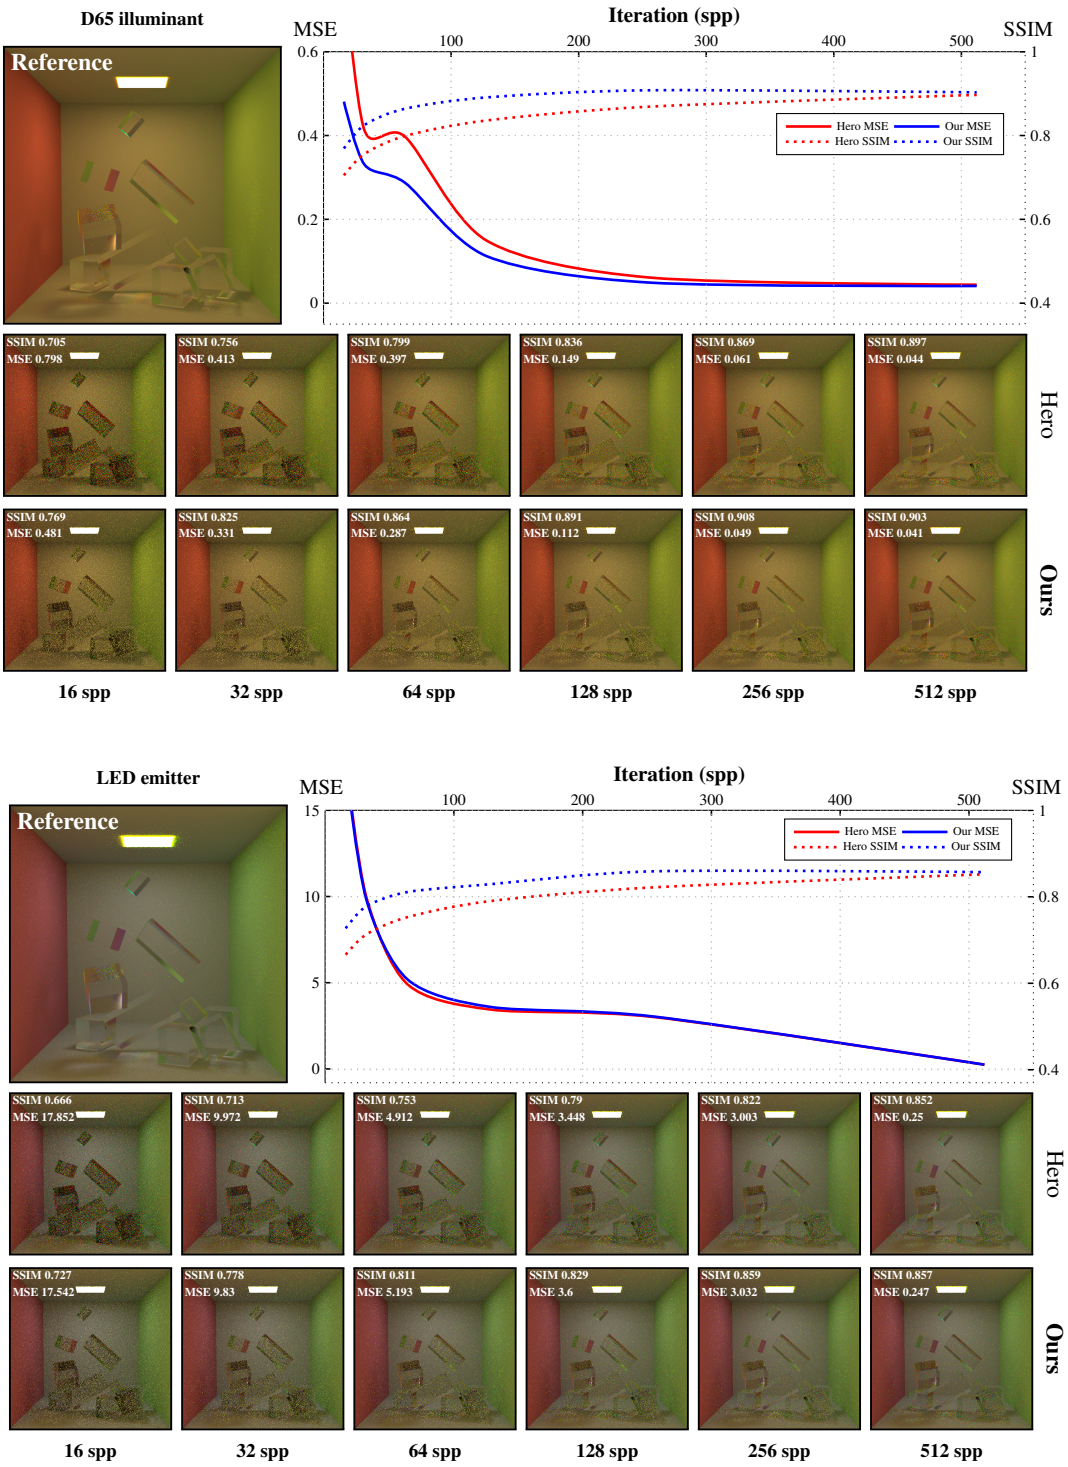

Figure 9: ***Evaluation results.*** *We show the FALLING_CUBE scene and compare our method with HWSS at different iterations: 16, 32, 64, 128, 256, and 512 samples per pixel. The scenes contain a single emitter whose SPD varied. Refer to Figure 8 for the corresponding distribution. Our method greatly outperforms HWSS in SSIM, a perceptual model.*

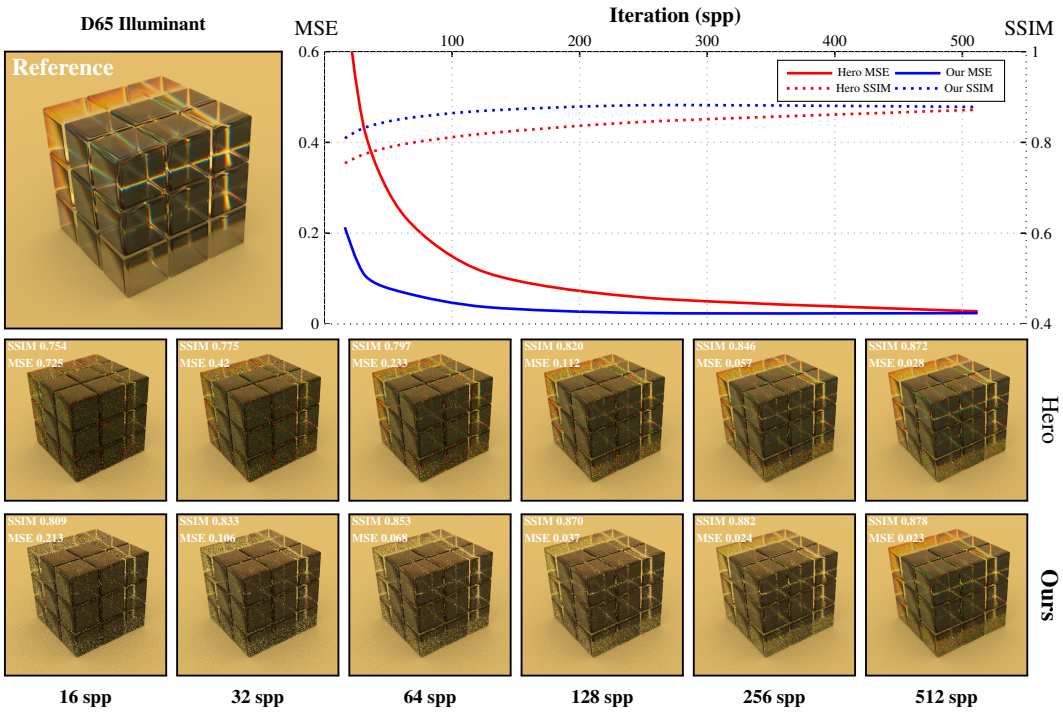

Figure 10: ***Evaluation results.*** *We show the RUBIK scene and compare our method with HWSS at different iterations: 16, 32, 64, 128, 256, and 512 samples per pixel. The scene is under D65 illuminant. Refer to Figure 8 for the corresponding distribution. While our method still outperforms HWSS, dispersion effects' color shown in our method is inaccurate.*

6. Responsible Research

Reproducibility in computer graphics is important as it allows for verifying and improving existing techniques. We ensure our method can be reproduced by including our implementation details in Section 4, which is implemented using the open-source PBRT-v3 renderer [JPH16].

7. Discussion

We have shown that our method increases the perception of early renders, while producing lower or similar error rates when compared to HWSS. In scenes that contain specular wavelength-dependent materials, our method greatly outperforms HWSS while producing little overhead in the process of restricting and relaxing wavelength distributions.

In the future, investigations can be spent on how we relax the range of wavelengths. Given the human visual system, it could be more beneficial to relax in a quadratic fashion rather than linear. In addition, an area that is worth looking into is adding support for participating media to measure the effectiveness of our method with different spectral phenomena.

8. Conclusions

We introduced a *progressive approach* for spectral light transport. Our method focuses on alleviating variance in early renders by initially restricting and gradually relaxing the wavelength distribution. In the process of relaxing, we showed that introducing bias can reduce variance and significantly improve the perception of the output images. We also demonstrated that the bias will go away as we obtain an unbiased final render with accurate spectral phenomena. Moreover, our method introduces little overhead compared to the previous methods.

9. Acknowledgments

We would like to thank Matt Phar and his collaborators for making PBRT-v3 [JPH16] publicly available. We also wish to thank Johanne Roby, Martin Aubé et al. for their work on the Lamp Spectral Distribution Database [RA].

References

[EM99] EVANS G. F., MCCOOL M. D.: Stratified wavelength clusters for efficient spectral Monte Carlo rendering. Morgan Kaufmann Publishers Inc., pp. 42–49. 2

[JPH16] JAKOB J., PHARR M., HUMPHREYS G.: *Physically*

Based Rendering:From Theory To Implementation. Morgan Kaufmann, 2016. 5, 7

[JW01] JENKINS F., WHITE H.: *LSC Fundamentals of Optics*. McGraw-Hill Education, Dec. 2001. 5

[Max57] MAXWELL J. C.: XVIII.—Experiments on Colour, as perceived by the Eye, with Remarks on Colour-Blindness. *Earth and Environmental Science Transactions of The Royal Society of Edinburgh 21*, 2 (1857), 275–298. doi:10.1017/S0080456800032117. 4

[Mül22] MÜLLER T.: Tev — The EXR Viewer, June 2022. 5

[RA] ROBY J., AUBE M.: Lamp spectral power distribution database (lspdd). Accessed 2022-06-10 at https://lspdd.org/app/en/home2019. 5, 7

[RBA09] RADZISZEWSKI M., BORYCZKO K., ALDA W.: An improved technique for full spectral rendering. 2

[Vea98] VEACH E.: *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, 1998. 1, 2

[WBSS04] WANG Z., BOVIK A., SHEIKH H., SIMONCELLI E.: Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing 13*, 4 (Apr. 2004), 600–612. doi:10.1109/TIP.2003.819861. 5

[WGGH20] WESTREX, GEORGIEVILIYAN, GRUSONADRIEN, HACHISUKATOSHIYA: Continuous multiple importance sampling. *ACM Transactions on Graphics (TOG)* (July 2020). doi:10.1145/3386569.3392436. 2, 5

[WND*14] WILKIE A., NAWAZ S., DROSKE M., WEIDLICH A., HANIKA J.: Hero Wavelength Spectral Sampling. *Computer Graphics Forum 33*, 4 (2014), 123–131. doi:10.1111/cgf.12419. 1, 2, 4, 5

[You02] YOUNG T.: II. The Bakerian Lecture. On the theory of light and colours. *Philosophical Transactions of the Royal Society of London 92* (Jan. 1802), 12–48. doi:10.1098/rstl.1802.0004. 4