Efficient Emitter Sampling for Spectral Path Tracing

Piyush Deshmukh
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

Efficient Emitter Sampling for Spectral Path Tracing

P. Deshmukh, M. van de Ruit E. Eisemann

[1]Delft University of Technology, The Netherlands

Abstract

Spectral Monte-Carlo methods are powerful physically-based techniques for simulating wavelength-dependent phenomena such as dispersion. However, compared to tristimulus rendering, they involve sampling the spectral domain, which adds substantial overhead, requiring significantly more samples for noise-free, realistic-looking renders. Thereby, we propose a simple approach to efficiently sample emitters. We precompute a simple 2-dimensional data structure using spectral power distributions of scene emitters. We use it to model a probability distribution function to sample an emitter that yields high path throughput at every intersection when using Next Event Estimation. Our method handles various geometries and spectral distributions of scene emitters, improves convergence, and reduces noise with negligible overhead.

Path tracing is a renowned physically-based rendering method to produce photorealistic images. Many implementations of path tracing are trichromatic. They model light as a combination of three (RGB) values. However, in reality, light is a continuous spectrum of wavelengths. A typical example that illustrates the contrast between these approaches is the dispersion of light when it passes through a prism. Such an approximation fails to account for the distinct behavior of rays with varied wavelengths when traversing a scene. Trichromatic approaches inadequately reproduce color in wavelength-based phenomena [Bor91]. This leads to color inaccuracy in scenes with the omission of wavelength-dependent phenomena such as caustics, fluorescence, polarization, and iridescence [JF99, GM99].

Spectral path tracing is a variant of path tracing that aims to faithfully model light as a spectrum of wavelengths to enable capturing physically-based spectral phenomena mentioned previously. Since path tracing relies on Monte Carlo methods, spectral rendering requires sampling the spectral domain in addition to sampling paths, which introduces a computational overhead. As the domain for sampling grows larger with the addition of another (spectral) dimension, more samples are required to produce high-quality images. This drawback of spectral rendering makes it an unattractive approach for VFX production.

Over the years, numerous techniques have been developed to reduce noise in tristimulus path tracing implementations. Common strategies include stratification, sampling Poisson-disc distributions (blue noise), importance sampling, and Next Event Estimation (NEE). Evans et al. and van de Ruit et al. and Wilkie et al. [EM99, vdRE21, WFM18] demonstrate improvements in spectral path tracing, to note a few. Our approach integrates importance sampling improvements in NEE, which were previously leveraged in tristimulus rendering, into a spectral pipeline.

Unidirectional path tracing involves shooting rays across the scene and following bounces until they end up at an emitter, which propagates energy across the ray path. NEE is a popular technique applied to path tracing, reducing Monte Carlo estimation variance. It samples the direct illumination by constructing secondary paths to one scene emitter at every surface intersection. For NEE, an emitter itself needs to be chosen before sampling a point on the emitter. In a scene with multiple light sources, a probability distribution function (PDF) determines the emitter sampled at every bounce. However, for a spectral renderer, the amount of energy along a path is additionally influenced by the spectral power distribution of the sampled emitter for a propagated wavelength. For instance, when a green wavelength is propagated in a scene, it is not beneficial to sample a red light (despite its visibility) as it does not emit energy in the green wavelength range.

In this paper, we introduce a novel technique for wavelength-based sampling of emitters. Our technique describes a distribution function over a scene's emitters and their spectral power distributions (SPDs). It enables impor-

tance sampling emitters that maximize path energy for a given propagated wavelength. We demonstrate that it improves the convergence behavior of NEE techniques for extended light transport algorithms in wavelength-dependent scenarios.

After reviewing spectral light transport equations and summarising existing literature (Equation 1), we explore the different components of our method (Section 2) and its implementation using an existing path tracer (Section 3). We then then evaluate our emitter sampling strategy (Section 4) and discuss its results (Section 6), before concluding (Section 7).

1. Background

This section highlights related work and interrelates our contribution to the existing literature.

Spectral Light Transport

Physically-based renderers are concerned with evaluating the light transport equation

$$L_o(p,\omega_o) = L_e(p,\omega_o) + \int_\Omega f(p,\omega_o,\omega_i) L_i(p,\omega_i) |cos\theta_i| d\omega_i \tag{1}$$

where f is the Bidirectional Scattering Function (BSDF), L_d is the radiance from direction ω_i and the cosine term applies a foreshortening factor as per Lambert's cosine law. Using Monte Carlo estimation, we get

$$\frac{1}{N}\sum_{i=1}^{N} \frac{f(p,\omega_o,\omega_i) L_d(p,\omega_i) |cos\theta_i|}{p(\omega_i)} \tag{2}$$

We use an extended form for spectral path-integral formulation [Vea98]. The spectral radiance I entering a single-pixel j is given by

$$I_j = \int_\Lambda \int_\Omega f_j(\bar{x},\lambda) d\mu(\bar{x}) d\lambda \tag{3}$$

where Λ denotes the spectral domain of wavelengths, and Ω is the path space. $f_j(\bar{x},\lambda)$ denotes path throughput. Since this integral is hard to solve analytically in non-trivial scenes, Monte Carlo estimation is applied, yielding

$$\widehat{I_j} = \frac{1}{N}\sum_{i=1}^{N} \frac{f_j(\bar{x_i},\lambda_i)}{p(\bar{x_i},\lambda_i)} \tag{4}$$

which converges towards the correct solution as $N \to \infty$. Here $p(\bar{x_i},\lambda_i)$ describes a PDF which can be decomposed as

$$p(\bar{x},\lambda) = p(\lambda) \cdot p(\bar{x}|\lambda) \tag{5}$$

Multiple wavelength sampling

As stated in Equation 4, the spectral light transport equation is based on calculating a unique wavelength per path. Wilkie et al. formulated *hero wavelength spectral sampling* (HWSS) [WND14, RBA09]. This approach clusters and propagates a group of N wavelengths until a wavelength dependency occurs. The dependency could arise from wavelength-based light phenomena, for example - a change of propagation media. If such a wavelength dependency arises, N new paths would need to be spawned at such an intersection which could lead to an exponential path growth. Instead, a wavelength among the cluster, called the *hero wavelength*, is chosen for further propagation. The primary or hero wavelength is sampled randomly in the visible range (400-700nm), and other wavelengths in the cluster are then stratified across this spectral range using a rotation function. The hero wavelength determines the light transport path, and the results are combined using multiple importance sampling (MIS).

Next Event Estimation and Multiple Importance Sampling

At every surface interaction during path tracing, a new path is drawn by sampling the BSDF. To maximize the path throughput without prior knowledge of scene specifics, importance sampling is used to choose directions ω_i. From Equation 2 it can be seen that directions ω_i normal to a surface carry more energy as evidenced by the $cos\theta$ term. Multiple Importance Sampling is often used to draw samples utilizing the BSDF term and lighting term. Next Event Estimation is also used to split the total domain of sampling a direction in the hemisphere strictly into a part that contains direct illumination and the other that accounts for indirect illumination. As a path now always attempts to sample emitters, it improves convergence with reduced variance and noise. These noise reduction techniques are especially effective when the emitters have small areas, as it is improbable that uniformly sampling the hemisphere yields a ray that intersects a light source (the probability is proportional to the solid angle). However, they do not account for a specific sampled wavelength (equation 4) in proportion to the emitter SPDs in a scene, which is the focus of our work.

Visibility PDF and Related Works

Since choosing paths that carry a high amount of energy is beneficial, current approaches involve picking emitters based on an estimate of visibility, thereby reducing sampling error [KNK*16]. Nabata et al. [OHH18] demonstrate an approach to compute probabilities of sampling emitters from a given vertex by computing these for a subset of points and interpolating the results. Pre-computed visibility caching approaches have been examined earlier [War94, TB94, ZK95]. Guo et al. [GEE20] demonstrate collecting visibility information with a low-cost preprocess to improve NEE. PBRT,

the path tracer we extend upon, uses a simple preprocessed visibility distribution (hereon, referred to as visibility distribution or estimate visibility distribution) to formulate a PDF, which is used to sample emitters during NEE.

2. Methodology

We now present our method. We begin by explaining the crux of modeling an SPD-based PDF. Next, we provide an account of how this PDF can be combined with the existing visibility PDF, followed by the modeling of emitters to test the obtained PDF.

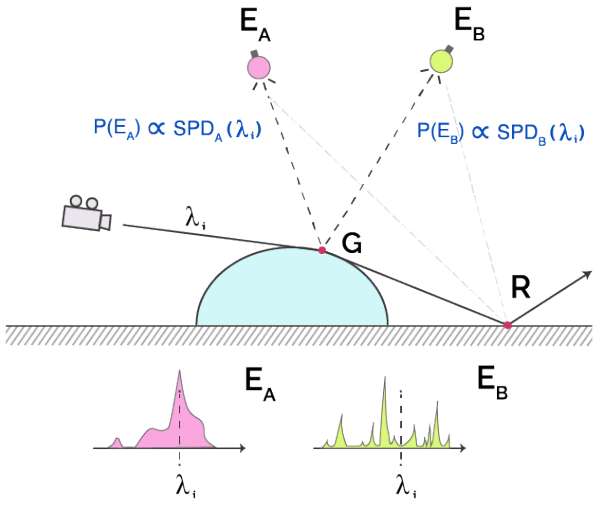

Figure 1: *A light ray intersects a dispersive medium at point G where a secondary ray is shot to either of the emitters (NEE). We propose that the probability of sampling an emitter should be proportional to its emission value λ_i for the wavelength propagated when a wavelength dependency occurs*

Sampling emitters

In contrast to exclusively using a visibility distribution, a more informed and efficient sampling of emitters can produce less color noise when working with a scene involving emitters with varying SPDs. Spectral power distributions (SPDs) model the corresponding power per unit area for each wavelength of light that an emitter emits. When propagating a sampled ray r with wavelength w across the scene, it is more likely that r was emitted from an emitter with a higher emission value for w than another, with a low emission value. This is simply because, in reality, the scene receives more spectral radiance (for w) from the previous light source than the latter. This observation forms the basis of our approach. Without any prior scene geometry information, paths are more likely to carry higher energies if emitters are sampled proportional to their SPDs for all wavelengths propagated as motivated in Figure 1. A PDF that would favorably sample such paths could yield a better convergence.

The scene information and emitter properties are known before calculating the paths. An efficient data structure is precomputed to keep track of SPDs of all emitters. This data structure consists of either emitter power characteristics or artist-specified emission parameters. The former accounts for the shape and size of emitters, whereas the latter includes artist-specified emission parameters that dictate the SPDs per unit area, thus not accounting for geometric attributes. When a ray hits a surface where a wavelength dependency occurs, such as dispersion or other wavelength-specific phenomena, this data structure is then queried for the wavelength under propagation. However, the visibility distribution alone is sampled in cases like reflection, where the intersection does not entail a wavelength dependency. In such cases, there is no difference in interaction with a standard trichromatic renderer. Since HWSS is used, four wavelengths are propagated with a single ray. Using the spectral PDF does not have any benefits in such a case and hence, is not required to do so.

Combining PDFs

Since the product of two PDFs is also a valid PDF, our approach combines naturally with existing sampling distributions (Equation 1). The newly obtained PDF benefits from the spectrum-based emitter preference in cases where multiple light sources are visible. Simultaneously, it prevents occluded sampling emitters, which would negatively impact the convergence and noise. The visibility PDF (in PBRT) already accounts for the solid angle subtended by emitters on the intersection point. Therefore, when the two PDFs are combined, choosing unit emission values over power as a metric for emitters is preferable to avoid accounting for the area twice. This is evaluated further in Section 4.

Emitters

To test spectral techniques under realistic scene lighting, emitter SPD data is required. The technique at hand works best when there is clearly, a better emitter to sample from in a pool of scene emitters, as a wrong choice has a higher difference in path throughput. This occurs when emitters have sharp peaks at wavelengths contrasting to other emitters. For such a case, it would be favorable to sample the emitter with a peak with a high probability as it is more likely that the propagated ray was emitted from this emitter. To ensure that arbitrary scaling of SPDs does not affect the choice of an emitter, the emitter SPDs were normalized.

$$\sum_{w \in W} Emm(w,e) = c \qquad \forall e \in E \tag{6}$$

Where c is a constant scaled based on overall scene brightness, W is the visible spectrum (400-700nm), and E is a set of all emitters.

The method was repeatedly tested for various types of emitters. The scenes used in Section 4 comprise CFLs and LEDs preprocessed in the aforementioned way, as can be seen from figure Figure 2.

Name	Short	LSPDD idx.	Type	Color temp.
Globe Twister	CFL1	2488	CFL	4404
Ledtech PAR20	LED1	2470	LED	5832
Philips Blue-Golite	LED2	2489	LED	7600

Table 1: *Name, LSPDD index, emitter type and color temperature as obtained from the LSPDD [RA19]*

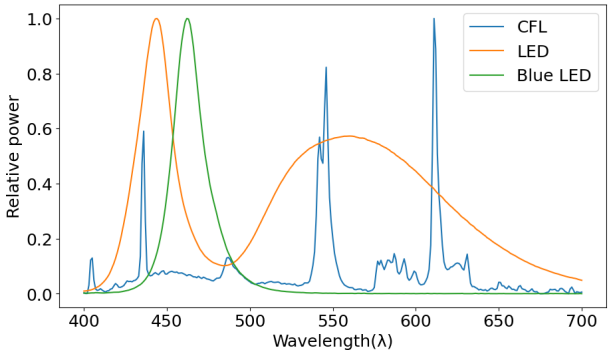

Figure 2: *We showcase the used emission spectra (Equation 2). Note the distinct emission value peaks of the emitters*

3. Implementation

Our method was implemented in Physically Based Rendering framework version 3(PBRT v3) [JPH16]. Our approach analyses emitter sampling when wavelength dependencies occur during intersections. It builds on an implementation that extends specular and non-specular dielectric BSDFs with Cauchy's equation to simulate wavelength dependency. It incorporates 4 wavelength HWSS and CMIS [WGGH20, WND14] as described in Equation 1 to enable sampling the spectral domain efficiently. PBRT has an underlying discrete binned spectral representation that uses 60 equally sized bins over a wavelength range from 400-700nm, with bands of 10nm per bin. This default range is sufficient for spectral-domain sampling and emitter SPDs modeling.

Our method consists of a preprocessing part and a rendering part. The preprocessing part computes a 2D vector array to capture emitter SPDs. This is implemented in the preprocessing part, which is invoked after scene construction from the input *pbrt* file and invoked prior to rendering. The method involved designing a 2D vector array storing the spectral information of emitters in the scene. The x-axis corresponds to the lights in the scene, and the y-axis corresponds to the wavelength bins in the visible range (400nm-700nm in the PBRT implementation). The values in each of the bins correspond to the Cumulative Distribution Function (CDF) computed from the emission intensity of the corresponding emitter at the specific wavelength.

During rendering, if a wavelength dependency occurs at an intersection, the hero wavelength is propagated after the intersection. Hence, the index for the bin containing the hero wavelength is computed, and an array containing the bin-specific emission values over all emitters is obtained. Based on this bin-slice, a distribution is computed by normalizing the power values.

4. Evaluation

We evaluate our method using three scenes with varying combinations of emitter geometry. The first two scenes are a variation of the Cornell box scene [Bit16] while the third showcases Suzzane in a simple scene composed of three emitters. The emission spectra are chosen cover common types of emitters- LED, fluorescent lights and a (hue) colored LED, from the Lamp Spectral Power Distribution Database (LSPDD) [RA19] under the CC-Y-NC-ND 2.5 CA license, displayed in Equation 2 and Figure 2. The reference images for each scene are computed using sufficient samples (N=65K) using an unbiased unidirectional path tracer. We use MSE as an error metric to compare the difference in performance between our method and the current state-of-the-art *Continuous Multiple Importance Sampling* techniques introduced by West et al. [WGGH20]. In all scenes except Suzanne, we use the combined PDF as described in Section 2 to evaluate our method. We use the *hyperfine* benchmarking tool with five warm-up rounds and over five iterations. The preprocessing times are averaged and included in the measure runtimes for our results. We use six render iterations with 16, 32, 64, 128, 256, and 512 SPP to plot the convergence behavior of our method against the HWSS CMIS method [WGGH20]. We evaluate both methods using four wavelengths as originally used [WND14].

Method Evaluation

First, we introduce the two light Cornell scene (Figure 3), which contains two slightly rough dispersive boxes. It uses emitters (CFL1 and LED1 from the table) with distinct peaks in their SPDs. Therefore, for most wavelengths, there is a clear, correct choice to sample one emitter over the other. We include three insets (A,B,C with N=512 spp.), where A and B highlight regions where dispersive phenomena are observable and C indicates a region with a lack of dispersive phenomena. Over the entire scene for N=512 spp, we report a decrease in error of 4.3% in comparison to HWSS. The er-

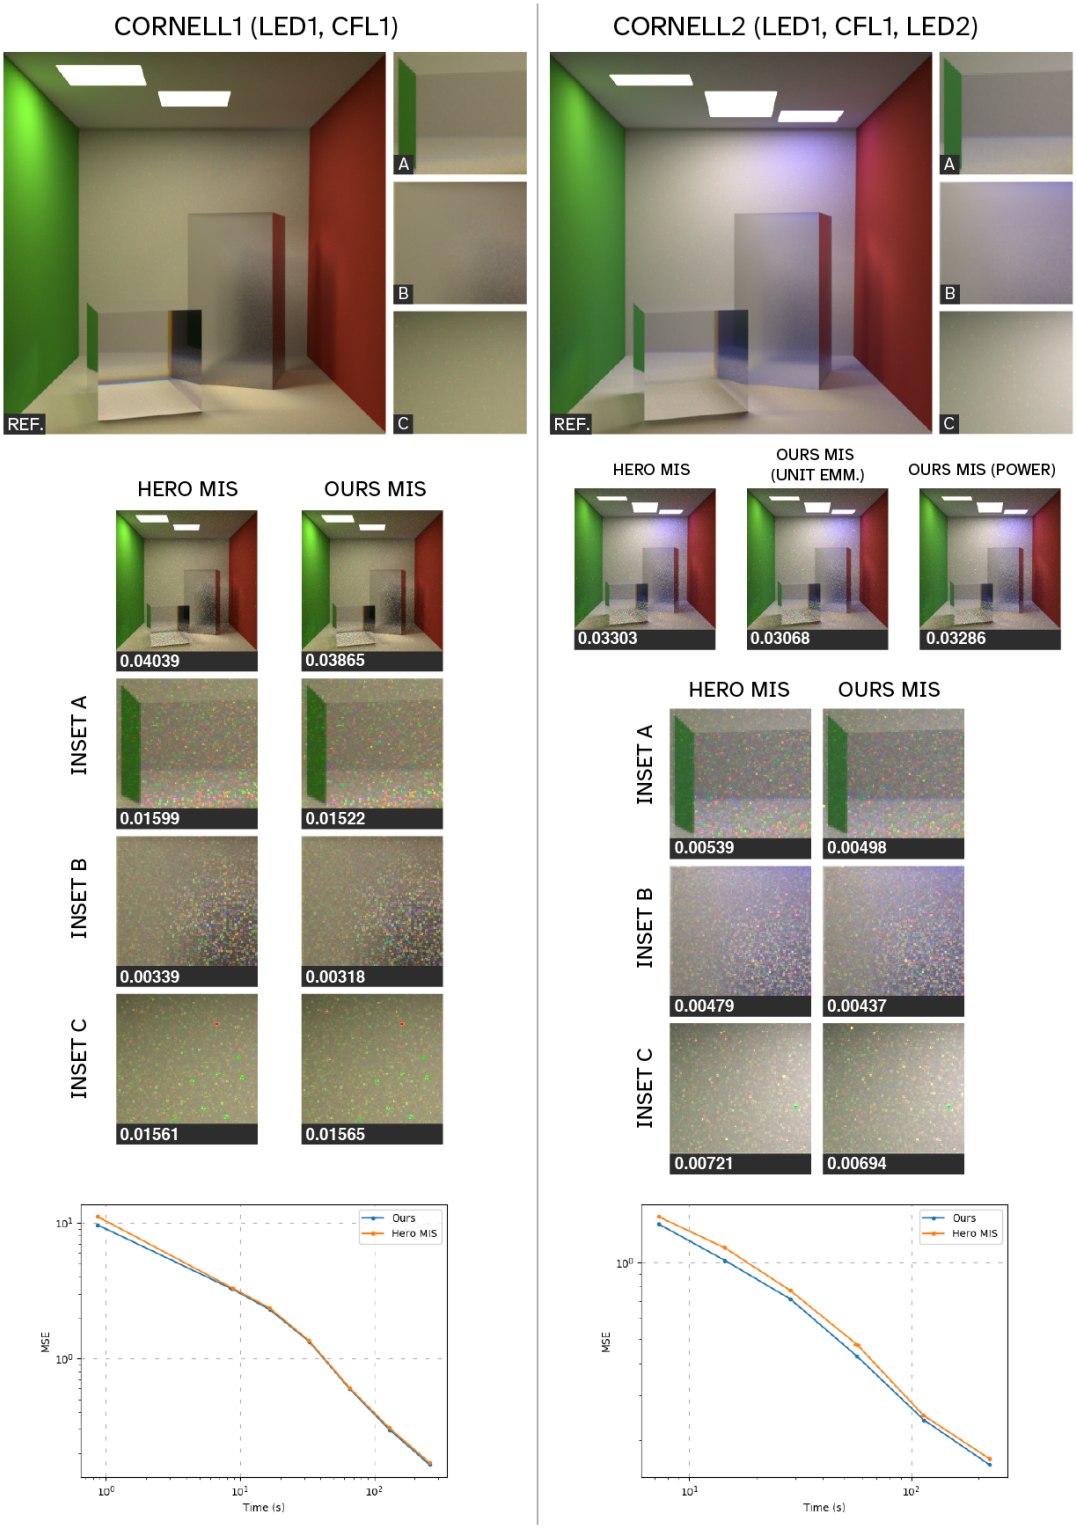

Figure 3: *We compare our method against Hero-Wavelength Sampling across two Cornell scenes using combinations of spectra and sizes of emitters. Insets demonstrate the effectiveness of our method is higher at dispersive surfaces. The convergence graphs show a higher MSE difference with low SPP renders.*

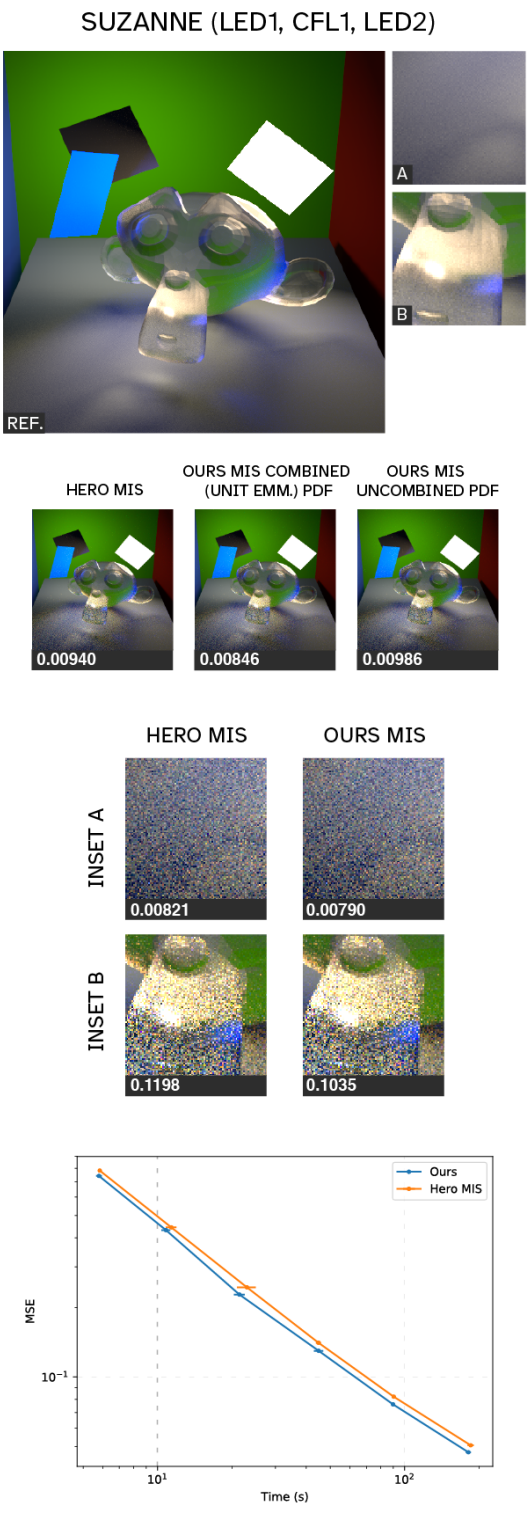

Figure 4: *We evaluate our method using a simple scene with a bright occluded emitter (LED1 from Equation 2) and a dispersive glass Suzanne. Our approach outperforms Hero-wavelength sampling when supplemented with the visibility distribution but performs worse without.*

ror decrease is noticeable in some regions (inset A: 4.8%, inset B:6.2%). However, due to some stray spectral rays, there is a slight increase in noise in inset C (0.3%).

Next, we evaluate performance in the three-light Cornell scene(Figure 3). In addition to emitters from the Cornell1 scene, this scene consists of a third emitter (LED2 from Equation 2) that partially overlaps with LED1 in the blue wavelength range. This scene is used to demonstrate the difference between power and unit emission as metrics, as the size of the emitters vary. The LED emitter is scaled up, while the CFL emitter is scaled down. We notice that using power as a metric leads to 6.8% more noise than the unit emission metric. We reason that this is caused by oversampling LED1 when using the power metric. The power metric scales the unit emission values with the emitter area. However, since we use a combined PDF that was obtained using the visibility approximation PDF, we account for the emitter area twice. This effect is especially important for the blue wavelengths (400-500nm), where LED1 and LED2 have high emission values. Using the unit emission values as a metric for our combined PDF, we notice a decrease of 7.6%, and 8.8% in insets A and B, respectively, highlighted similarly to the Cornell1 scene.

Finally, we introduce the Suzanne scene (Figure 4) to compare performance between the combined PDF and our PDF, exclusively compiled using emitter SPD values. It consists of LED1, CFL and LED2 from Equation 2. In this scene, LED1 is scaled to be twice as bright (c=2 from Equation 6) as the other emitters (c=1). LED1 is also occluded, which makes this scene a challenging scenario for our method if only the SPD-based PDF is used without supplementing it with the visibility estimate PDF. The power metric is used since we don't account for the area from the visibility estimate distribution. The non-combined PDF tends to sample LED1 with a higher probability in comparison to previous scenes at wavelength-dependent bounces (at the Suzanne surface). However, since the light faces the other way, such a path finds that the emitter is occluded and returns no energy. We analyze the performance of the same scene using the combined PDF and unit emission metric. The non-supplemented PDF performs worse than the baseline, while the supplemented PDF offers an MSE improvement of 9.9%.

Across all tests, we notice that our method performs at par or better than current approaches. We notice that our method performs better in areas where dispersive phenomena are prevalent (insets A, B from the Cornell scenes, and inset B from the Suzanne scene). While some scenes with occluded emitters pose a problem in the efficiency of our approach, supplementing our modeled PDF with the existing visibility estimate distribution yields better results. Additionally, the unit emission metric to construct the data structure is evaluated to perform better than the power metric when using the combined PDF.

5. Responsible Research

Although the ethical concerns in path-tracing are limited, we try to ensure transparency in results and process. We discuss implementation details of our method in section 3. Additionally, we build upon the widely known and public available PBRT v3 renderer [JPH16]. All results presented in Section 4 are run using the *hyperfine* benchmarking tool to reduce the manual errors when running and comparing multiple configurations.

6. Discussion

The implementation of the new sampling distribution is efficient with an *O(N)* construction time in the number of scene emitters, which is generally constant. In practice, it only adds a few milliseconds of overhead. It is important to emphasize that our method is more effective in scenes where emitters have varied SPDs. Fluorescent emitters have sharp spikes in their SPDs, and our method samples them with high probabilities when the corresponding wavelength is propagated in a scene, thus being performant with spiky illuminants. However, the approach provides no improvements in single emitter scenes or a scene with multiple emitters of the same type. Portals and occlusions can be a challenge for this strategy as they subdivide an emitter into multiple emitters to sample with the same SPD. However, for all of the aforementioned cases, the emitter SPD distribution, when combined with the visibility estimation distribution, generally yields similar or better results as they account for both visibility and emitter SPDs. Through the experiments, we conclude that it is always beneficial to use unit emission as a metric when the combined PDF is used.

Our approach can be extended when modeling fluorescent BRDFs. Such materials absorb energies from wavelengths outside the visible range and re-emit them in the visible range. Since the SPD data structure used to compute our PDF can be extended to store additional bins, our approach would efficiently sample an emitter that yields high path throughput when sampling wavelengths outside the visible range.

7. Conclusion

To alleviate the slow convergence behavior of Spectral Monte Carlo rendering, we introduced a simple, efficient emitter sampling strategy that improves NEE in spectral path-tracers by accounting for the correspondence between propagated wavelengths and scene SPDs. Our process involves precomputing a 2D vector array to store emitter SPDs. We combine the PDF generated from our preprocessed distribution with the existing emitter visibility estimation distribution to obtain our final emitter sampling distribution. This distribution is then sampled when choosing an emitter at a (wavelength-dependent) ray intersection during propagation, where NEE is used to estimate direct illumination. We showed that our method handles various geometries and SPDs of emitters. We demonstrated that this yields better convergence in a variety of scenes by reducing spectral noise, especially for low sampling rates.

We hope that integration with fluorescent BRDFs becomes possible in the future, and our method contributes to making spectral rendering slightly more feasible, given its overhead.

8. Acknowledgments

We would like to thank Matt Phar, Wenzel Jakob and Greg Humphreys for their work on the PBRT v3 renderer [JPH16]. We also extend our gratitude to Johanne Roby et al. for making the Lamp Spectral Distribution Database publicly available. I would also like to thank my supervisors, Mark van de Ruit and Prof. Dr. Elmar Eisemann for their guidance throughout the research project, and my teammates for their feedback on many aspects of the project.

References

[Bit16] BITTERLI B.: Rendering resources, 2016. https://benedikt-bitterli.me/resources/. 4

[Bor91] BORGES C.: Trichromatic approximation for computer graphics illumination models. *ACM SIGGRAPH Computer Graphics 25*, 4 (1991), 101–104. doi:10.1145/127719.122729. 1

[EM99] EVANS G., MCCOOLL M.: Stratified wavelength clusters for efficient spectral monte carlo rendering. *Graphics Interface* (1999). 1

[GEE20] GUO. J., EISEMANN M., EISEMANN E.: "next event estimation++: Visibility mapping for efficient light transport simulation". *Pacific Graphics 39*, 4 (2020). 2

[GM99] GLENN E., MICHAEL D.: Stratified wavelength clusters for efficient spectral monte carlo rendering. *Graphics Interface* (1999), 42–49. 1

[JF99] JOHNSON G., FAIRCHILD M.: Full-spectral color calculations in realistic image synthesis. *IEEE Computer Graphics and Applicaitons 19*, 4 (1999), 47–53. 1

[JPH16] JAKOB J., PHARR M., HUMPHREYS G.: *Physically Based Rendering:From Theory To Implementation*. Morgan Kaufmann, 2016. 4, 7

[KNK*16] KOERNER D., NOVÁK J., KUTZ P., HABEL R., JAROSZ W.: Subdivision next-event estimation for path-traced subsurface scattering. *Eurographics Symposium on Rendering - Experimental Ideas & Implementations (2016)* (2016), 1–6. 2

[OHH18] OTSU H., HANIKA J., HACHISUKA T. D. C.: "geometry-aware metropolis light transport.". *ACM Trans. Graph. 37*, 6 (2018). doi:10.1145/3272127.3275106. 2

[RA19] ROBY J., AUBE M.: Lamp spectral power distribution database (lspdd), 2019. Accessed 2022-06-10 at https://lspdd.org/app/en/home. 4

[RBA09] RADZISZEWSKI M., BORYCZKO K., ALDA W.: An improved technique for full spectral rendering. *WSCG 17* (2009), 1–16. 2

[TB94] TELLIER P., BOUATOUCH K.: Physics-based lighting models: Implementation issues. 110–121. doi:10.1007/978-3-642-57963-9_12. 2

[vdRE21] VAN DE RUIT M., EISEMANN E.: A multi-pass method for accelerated spectral sampling. *Pacific Graphics 40*, 7 (2021). 1

[Vea98] VEACH E.: Robust monte carlo methods for light transport simulation. *PhD thesis, Standford Universit* (1998). 2

[War94] WARD G.: Adaptive shadow testing for ray tracing. *Springer* (1994), 11–20. doi:10.1007/978-3-642-57963-9_2. 2

[WFM18] WILKIE A., FICHET A., MOJZIK: Handling fluorescence in a uni-directional spectral path tracer. *Computer Graphics Forum 37.4* (2018). doi:10.1111/cgf.13477. 1

[WGGH20] WEST R., GEORGIEV I., GRUSON A., HACHISUKA T.: "continuous multiple importance sampling". *ACM Trans. Graph. 39*, 4 (2020). doi:10.1145/3386569.3392436. 4

[WND14] WILKIE A., NAWAZ S., DROSKE M.: Hero wavelength spectral sampling. *Computer Graphics Forum 33*, 4 (2014), 123–131. doi:10.1111/cgf.12419. 2, 4

[ZK95] ZIMMERMAN K. S. P.: A two-pass solution to the rendering equation with a source visibility preprocess. *Springer* (1995), 282–300. doi:10.1007/978-3-7091-9430-0_27. 2