

Spectral Gradient Sampling for Path Tracing

Petitjean, Victor; Bauszat, Pablo; Eisemann, Elmar

DOI

[10.1111/cgf.13474](https://doi.org/10.1111/cgf.13474)

Publication date

2018

Document Version

Submitted manuscript

Published in

Computer Graphics Forum: the international journal of the Eurographics Association

Citation (APA)

Petitjean, V., Bauszat, P., & Eisemann, E. (2018). Spectral Gradient Sampling for Path Tracing. *Computer Graphics Forum: the international journal of the Eurographics Association*, 37(4), 45-53.
<https://doi.org/10.1111/cgf.13474>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Spectral Gradient Sampling for Path Tracing

Victor Petitjean, Pablo Bauszat, and Elmar Eisemann

Delft University of Technology, The Netherlands




Figure 1: An equal-time comparison between our proposed spectral gradient sampling and standard sampling within a conventional path tracer. The insets in the bottom right corner show the difference images to the reference, where blue and red indicates small resp. large errors. Our method improves the convergence and reduces chromatic noise in regions affected by wavelength-dependent scattering.

Abstract

Spectral Monte-Carlo methods are currently the most powerful techniques for simulating light transport with wavelength-dependent phenomena (e.g., dispersion, colored particle scattering, or diffraction gratings). Compared to trichromatic rendering, sampling the spectral domain requires significantly more samples for noise-free images. Inspired by gradient-domain rendering, which estimates image gradients, we propose spectral gradient sampling to estimate the gradients of the spectral distribution inside a pixel. These gradients can be sampled with a significantly lower variance by carefully correlating the path samples of a pixel in the spectral domain, and we introduce a mapping function that shifts paths with wavelength-dependent interactions. We compute the result of each pixel by integrating the estimated gradients over the spectral domain using a one-dimensional screened Poisson reconstruction. Our method improves convergence and reduces chromatic noise from spectral sampling, as demonstrated by our implementation within a conventional path tracer.

CCS Concepts

•Computing methodologies → Ray tracing;

1. Introduction

Monte-Carlo light transport algorithms are popular techniques for rendering high-quality, photo-realistic images. While most renderers are RGB-based, several advanced phenomena of light, such as dispersion, diffraction gratings, or thin-film materials, can only be accurately simulated with spectral rendering. Adding spectral sampling to a Monte-Carlo renderer adds another level of complexity to an already costly process and drastically increases the number of samples required for noise-free images. If the sampling rate is insufficient, visual quality is highly degraded as color noise appears. While previous approaches lower chromatic noise by reusing each

path sample for multiple wavelengths [EM99, RBA09, WND*14], the number of required samples still remains drastically higher compared to trichromatic rendering. Further, they do not improve scenes with perfectly-specular materials (e.g., perfect glass).

Recently, gradient-domain rendering has been introduced as a new approach for noise reduction in Monte-Carlo rendering [LKL*13]. It is based on the idea of directly estimating gradients between image pixels using correlated pairs of paths and it has been shown that gradient sampling can significantly reduce variance. In this paper, we extend gradient-domain rendering beyond trichromatic light transport and introduce gradient sampling

for the spectral domain. In contrast to traditional gradient sampling, which estimates differences between pixels, we estimate finite differences between points in the spectral distribution of a single pixel. This is achieved by sampling pairs of paths using a novel shift mapping that correlates path samples in the spectral domain. The final pixel color is then computed from the noisy estimate of the spectral distribution and the gradient of that distribution using a one-dimensional screened Poisson reconstruction. To reduce the computational overhead introduced by pair sampling, we propose a scheme to balance standard and gradient sampling. This scheme favors paths with higher contributions. We connect our method to a conventional path tracer and show that it significantly improves convergence and reduces chromatic noise in image regions which are affected by wavelength-dependent effects. Our approach is orthogonal to previous methods which reuse path samples for multiple wavelengths and we demonstrate how to combine our method with the recently introduced Hero wavelength sampling [WND*14] for improved performance in scenes with glossy and diffuse wavelength-dependent scattering.

2. Related Work

Spectral Light Transport To render an image, physically-based light transport algorithms integrate over the space of all light-carrying paths in a scene. Spectral rendering introduces an additional integration domain over the space of wavelengths. Using the general path-space form of light transport [Vea97], the radiance spectrum of a pixel, denoted by I , can be expressed as a double integral with form:

$$I = \int_{\Lambda} \int_{\Omega_{\lambda}} f(\bar{\mathbf{x}}, \lambda) d\mu(\bar{\mathbf{x}}) d\lambda \quad (1)$$

where Λ denotes the spectral domain, Ω_{λ} is the union of all paths of finite length through which light with the wavelength λ can travel and arrive at the pixel, and $f(\bar{\mathbf{x}}, \lambda)$ denotes the measurement contribution of the path $\bar{\mathbf{x}}$ for the given wavelength λ . The spectral radiance distribution I is often stored in discretized form using a set of uniformly distributed bins, each representing a range of wavelengths. More adaptive representations of the spectral distribution use sets of basis functions, e.g. Gaussian [Mey88], the Fourier-basis [Pee93], or specialized basis functions [DMCP94, RP97, DF03, BDM09]. In the discrete case, there is a straightforward analogy between the classic image gradient and our proposed spectral gradient; the former is the finite difference between two pixels, while the latter is the difference between two bins. Without loss of generality, we assume a discrete representation in the following for an easier intuition.

Spectral Monte-Carlo light transport methods solve Eq. 1 using point sampling and a naive unbiased estimator given by:

$$I \approx \frac{1}{N} \sum_{i=1}^N \frac{f(\bar{\mathbf{x}}_i, \lambda_i)}{p(\bar{\mathbf{x}}_i, \lambda_i)} \quad (2)$$

Here, N denotes the sampling rate and $p(\bar{\mathbf{x}}_i, \lambda_i)$ the probability density of the sample pair $(\bar{\mathbf{x}}_i, \lambda_i)$. The naive estimator uses a path sample $\bar{\mathbf{x}}_i$ for only a single wavelength, which potentially introduces a large amount of color noise and is especially wasteful in cases where a path does not encounter any wavelength-dependent interactions. A wavelength-dependent interaction is here defined as an

intersection with a dispersive material, whose scattering distribution function depends on the wavelength (e.g. glass or water).

Improved Wavelength Sampling More advanced sampling strategies reuse a path sample for multiple wavelengths to reduce color noise. Evans and McCool [EM99] introduced stratified wavelength sampling and separated the number of samples taken for the spectral and path domain. This allows a path sample to contribute to multiple wavelengths until it encounters a wavelength-dependent surface. Given such an interaction, the path is split, which leads to either exponential path growth or selecting a single wavelength for further propagation, which introduces color noise. Furthermore, their approach requires the path sampling to be independent of the wavelength, i.e. $p(\bar{\mathbf{x}}, \lambda) = p(\lambda) \cdot p(\bar{\mathbf{x}})$, which forbids proper importance sampling for wavelength-dependent materials. Radziszewski et al. [RBA09] and more recently Wilkie et al. [WND*14] proposed to use spectral multiple importance sampling [Vea97] to overcome this problem. Each wavelength is treated as a different path sampling strategy and initially a single wavelength, the Hero wavelength, is chosen for path propagation. The sampled path then contributes to a set of wavelengths which are either randomly sampled [RBA09] or deterministically chosen [WND*14]. Since a path-wavelength pair can now be sampled from several wavelengths, multiple importance sampling is required to account for the changed probability density. An unbiased estimator using multiple importance sampling is given as:

$$I \approx \frac{1}{N} \frac{1}{C} \sum_{i=1}^N \sum_{j=1}^C w_{\lambda_h}(\bar{\mathbf{x}}_i, \lambda_j) \frac{f(\bar{\mathbf{x}}_i, \lambda_j)}{p(\bar{\mathbf{x}}_i, \lambda_j)} \quad (3)$$

$$w_{\lambda_h}(\bar{\mathbf{x}}_i, \lambda_j) = \frac{p(\bar{\mathbf{x}}_i, \lambda_h)}{\sum_{k=1}^C p(\bar{\mathbf{x}}_i, \lambda_k)} \quad (4)$$

where C is the number of sampled wavelengths, $w_{\lambda_h}(\bar{\mathbf{x}}_i, \lambda_k)$ denotes the weight (here using the balance heuristic [Vea97]), and λ_h is the sampled hero wavelength of $\bar{\mathbf{x}}_i$. The probability density of the path sampling is now dependent on the wavelength, i.e. $p(\bar{\mathbf{x}}, \lambda) = p(\lambda) \cdot p(\bar{\mathbf{x}} | \lambda)$. Note that although the path is reused for multiple wavelengths, its spatial configuration is not changed which is a major distinction in comparison to our method. Unfortunately, reusing a path for multiple wavelengths has no effect for perfectly-specular materials (e.g., a perfect glass dielectric) which isolate a single wavelength from the spectrum. Nevertheless, these approaches improve convergence for interactions with glossy and diffuse wavelength-dependent surfaces. Since they are orthogonal to our approach, a combination is potentially fruitful and we demonstrate how this is achieved for the Hero wavelength sampling approach from Wilkie et al. [WND*14] in Sec. 3.4.

Gradient-Domain Rendering (GDR) Gradient-domain rendering, originally introduced by Lehtinen et al. [LKL*13] in the context of Metropolis Light Transport [VG97], is a new way of reducing noise for Monte-Carlo rendering. At its core, GDR directly estimates image gradients in addition to pixel colors. By using pairs of paths, which are correlated between pixels, the image gradients are estimated with a significantly lower variance. This occurs due to some of the randomness from the Monte-Carlo process canceling out. A final image is computed by integrating the image gradients together with the estimated pixel colors using a 2D screened

Poisson reconstruction [BCCZ08]. An important role in gradient-domain rendering is played by the shift mapping, which is a bijective function that creates a pair of paths with similar contributions by shifting a base path from one pixel to another. Lehtinen et al. [LKL*13] proposed a shift mapping which employs the powerful Manifold Exploration technique from Jakob and Marschner [JM12]. However, it relies on global information about the path and costly numerical optimizations. Later, Manzi et al. [MRK*14] introduced several improvements for gradient samplings, such as symmetric gradient sampling. This circumvents issues with non-bijective mappings and improves the robustness of the gradient estimator. Manzi et al. [MRK*14] further combine the two ways a gradient can be estimated (standard sampling and gradient sampling) to better handle singularities in the path-space using an additional binary weighting which switches between the two strategies. We make use of this technique for a different purpose and design a binary weighting that allows us to perform spectral gradient sampling only for paths with larger contributions.

The seminal work on gradient-domain path tracing from Kettunen et al. [KMA*15] showed that gradient sampling can also lead to significant improvements for standard Monte-Carlo path tracing. They further proposed a simpler shift mapping which only requires local path information and can be computed sequentially during the path sampling. Lately, the gradient-domain rendering paradigm has received a lot of attention from the computer graphics community and has been successfully applied in the context of bi-directional path tracing [MKA*15], photon mapping [HGNH17], vertex connection and merging [SSC*17], and recently path reusing [BPE17]. Noticeably, temporal gradient-domain path tracing [MKD*16] was the first approach to extend gradient sampling beyond the image domain by computing spatio-temporal differences between frames of animations. Our approach is the first to apply gradient-domain rendering in a non-trichromatic setting and adapt the concept for the spectral domain.

We briefly summarize the analytic definition of the image gradient in gradient-domain rendering. We refer the reader to the original work [LKL*13, MRK*14, KMA*15] for a more detailed description and derivation. The gradient between two pixels i and j can analytically be expressed in a single integration over path-space using a shift mapping $T_{i \rightarrow j}(\bar{\mathbf{x}})$ [LKL*13]:

$$\begin{aligned} \Delta_{ij} &= \int_{\Omega_i} f_i(\bar{\mathbf{x}}) - f_j(T_{i \rightarrow j}(\bar{\mathbf{x}})) |T'_{i \rightarrow j}| d\mu(\bar{\mathbf{x}}) \\ &= \int_{\Omega_i} g_{ij}(\bar{\mathbf{x}}) d\mu(\bar{\mathbf{x}}) \end{aligned} \quad (5)$$

The Jacobian determinant of the shift mapping $|T'_{i \rightarrow j}|$ accounts for the change of integration variable for the pixel j . Since it is non-trivial to find a good path sampler for the pixel i that is also guaranteed to cover the path-space of pixel j after mapping, Manzi et al. and Kettunen et al. performed symmetric sampling of the gradients. In practice, this allows to generate base paths from both pixels using the standard pixel path sampler. The symmetric gradient sampling estimator is given by:

$$\Delta_{ij} = \int_{\Omega_i} w_{ij}(\bar{\mathbf{x}}) g_{ij}(\bar{\mathbf{x}}) d\mu(\bar{\mathbf{x}}) + \int_{\Omega_j} w_{ji}(\bar{\mathbf{x}}) g_{ji}(\bar{\mathbf{x}}) d\mu(\bar{\mathbf{x}}) \quad (6)$$

The weights $w_{ij}(\bar{\mathbf{x}})$ and $w_{ji}(\bar{\mathbf{x}})$ either account for the duplicated




Figure 2: Spectral Gradient Concept. While gradient-domain rendering in the image space (a) estimates differences between pixels, our method (b) computes the gradients of the spectral radiance distribution inside each pixel. Spectral gradient samples are taken using pairs of paths which are correlated by a spectral shift mapping that maps a base path $\bar{\mathbf{x}}$ sampled using the wavelength λ to an offset path $\bar{\mathbf{x}}'$ which contributes to the wavelength λ' .

appearance of path pairs using multiple importance sampling or handle non-invertible shifts. These formulations form the basis for our definition of the spectral gradient.

3. Spectral Gradient Sampling

A typical spectral renderer computes the value of an image pixel by estimating its spectral radiance distribution I . Inspired by gradient-domain rendering in the image domain, we propose to additionally estimate the "gradients" of I for each pixel. Assuming that each point in the spectrum is associated with a certain wavelength λ , we define the spectral gradient at a point in the spectrum to be the finite difference between the spectral (scalar) values I_λ and $I_{\lambda'}$ with $\lambda' = \lambda + \delta$ being another wavelength offset by a small δ . We illustrate our concept in Fig. 2.

Conventionally, we can express the gradient at λ as:

$$\begin{aligned} \Delta_{\lambda\lambda'} &= I_\lambda - I_{\lambda'} \\ &= \int_{\Omega_\lambda} f(\bar{\mathbf{x}}, \lambda) d\mu(\bar{\mathbf{x}}) - \int_{\Omega_{\lambda'}} f(\bar{\mathbf{x}}, \lambda') d\mu(\bar{\mathbf{x}}) \end{aligned} \quad (7)$$

Alternatively, we can follow the derivation of the analytic image gradient using a shift mapping from Lehtinen et al. [LKL*13] and express the spectral gradient in a single integration:

$$\begin{aligned} \Delta_{\lambda\lambda'} &= \int_{\Omega_\lambda} f(\bar{\mathbf{x}}, \lambda) - f(S_{\lambda \rightarrow \lambda'}(\bar{\mathbf{x}}), \lambda') |S'_{\lambda \rightarrow \lambda'}| d\mu(\bar{\mathbf{x}}) \\ &= \int_{\Omega_\lambda} g_{\lambda\lambda'}(\bar{\mathbf{x}}, \lambda) d\mu(\bar{\mathbf{x}}) \end{aligned} \quad (8)$$

Here, $S_{\lambda \rightarrow \lambda'}(\cdot)$ is a spectral shift function which maps a base path associated with a wavelength λ to another wavelength λ' . We denote the spectral shift mapping with the letter S instead of T to emphasize that the mapping is performed between two wavelengths and not pixels. Accordingly, we denote the Jacobian determinant of the spectral shift mapping with $|S'_{\lambda \rightarrow \lambda'}|$. We discuss the requirements for spectral shift mappings and the design of our proposed S in Sec. 3.1.

The benefit of the formulation of Eq. 8 is that it allows our method to directly estimate the spectral gradient with significantly lower variance by sampling it using pairs of correlated paths. In practice, this is achieved by creating a base path sample for

the wavelength λ , mapping it to the wavelength λ' using the correlation-inducing shift S , and measuring the difference in contribution between the paths. For each pixel, we estimate the spectral radiance I and additionally keep track of the spectral gradient distribution G . Both values are estimated using the same set of path samples. After I and G have been estimated, we perform a one-dimensional screened Poisson reconstruction to compute the final pixel color (Sec. 3.2). In contrast to gradient-domain rendering in the image plane, spectral gradient sampling only correlates path samples inside each pixel and only along the spectral dimension. Hence, each pixel can be processed independently.

Since a typical wavelength-dependent pixel path sampler is not guaranteed to fully cover the space of paths $\Omega_{\lambda'}$ through mapping from Ω_{λ} , we also perform symmetric sampling of the gradients using spectral multiple importance sampling:

$$\Delta_{\lambda\lambda'} = \Delta_{\lambda\lambda'}^{\lambda} + \Delta_{\lambda\lambda'}^{\lambda'} = \int_{\Omega_{\lambda}} w_{\lambda\lambda'}(\bar{\mathbf{x}}, \lambda) g_{\lambda\lambda'}(\bar{\mathbf{x}}, \lambda) d\mu(\bar{\mathbf{x}}) + \int_{\Omega_{\lambda'}} w_{\lambda'\lambda}(\bar{\mathbf{x}}, \lambda') g_{\lambda'\lambda}(\bar{\mathbf{x}}, \lambda') d\mu(\bar{\mathbf{x}}) \quad (9)$$

The weights $w_{\lambda\lambda'}$ and $w_{\lambda'\lambda}$ are computed using the wavelength-dependent path probability densities. Using the balance heuristic [Vea97], they are defined as:

$$w_{\lambda\lambda'}(\bar{\mathbf{x}}, \lambda) = \frac{p(\bar{\mathbf{x}}, \lambda)}{p(\bar{\mathbf{x}}, \lambda) + p(S_{\lambda \rightarrow \lambda'}(\bar{\mathbf{x}}), \lambda')} \quad (10)$$

$$w_{\lambda'\lambda}(\bar{\mathbf{x}}, \lambda') = \frac{p(\bar{\mathbf{x}}, \lambda')}{p(\bar{\mathbf{x}}, \lambda') + p(S_{\lambda' \rightarrow \lambda}(\bar{\mathbf{x}}), \lambda)} \quad (11)$$

In practice, a base path sampled by the standard wavelength-dependent pixel sampler using the wavelength λ is mapped twice to the wavelengths $\lambda' = \lambda + \delta$ and $\lambda'' = \lambda - \delta$ to generate a sample for the gradients $g_{\lambda\lambda'}$ and $g_{\lambda\lambda''}$. Although the offset paths are typically cheaper to generate than two completely new paths, gradient sampling introduces a non-negligible computational overhead compared to traditional sampling. We present how we reduce computational overhead by restricting gradient sampling to paths with significant contributions in Sec. 3.3.

3.1. A Spectral Shift Mapping

The shift mapping S receives a path sampled with a wavelength λ and shifts it to a similar path which contributes to the wavelength λ' . As demonstrated by Kettunen et al. [KMA*15], gradient sampling is beneficial when the contributions of the base and offset path are similar: more precisely, the integrand $f(\bar{\mathbf{x}}, \lambda) - f(S_{\lambda \rightarrow \lambda'}(\bar{\mathbf{x}}), \lambda')$ should become as small as possible. The major distinction between existing shift mappings and our new function S is that we do not shift the origin of a path in the image plane but instead its associated wavelength in the spectral domain. This has several implications. First, the two paths in a correlated pair originate at the same image location and will only diverge when they encounter a wavelength-dependent interaction (assuming that the sensor itself does not introduce any dispersion). Second, a spectral gradient path pair can diverge and reconnect multiple times, e. g. by passing through several refractive objects with diffuse interactions in-between. Third, the classification if a base

path vertex is connectable or not also depends on the wavelength-dependency of the vertex's scattering function. For example, a traditional shift mapping would classify a rough dielectric material as connectable, since it is not considered as (near-)specular. However, in the spectral domain such a material potentially causes a wavelength-dependent dispersion.

We design our shift mapping S by building on the main idea behind the shift mapping from Kettunen et al. [KMA*15], which is to replicate the projected half-vectors of the base path for the offset path until two consecutive connectable vertices are found. Hereby, the offset path has a similar contribution as the base path, which, as indicated, is advantageous for the gradient-domain computations. Given a base path $\bar{\mathbf{x}}_{\lambda}$, our shift computes an offset path $\bar{\mathbf{x}}_{\lambda'}$ in the following way (shown in Fig. 3). First, the offset path follows the base path until a wavelength-dependent interaction is encountered. If no such interaction occurs, $\bar{\mathbf{x}}_{\lambda}$ and $\bar{\mathbf{x}}_{\lambda'}$ will be identical. At the first wavelength-dependent interaction, the paths will disperse and we choose a new outgoing direction for the offset path. We denote the incoming directions for the base and offset path at the respective vertices with \mathbf{i}_{λ} and $\mathbf{i}_{\lambda'}$, and the outgoing direction of the base path as \mathbf{o}_{λ} . We choose a new outgoing direction for the offset path $\mathbf{o}_{\lambda'}$ in such a way that the half-vectors for both paths (projected to the local shading space) are identical. The half-vector can be conveniently expressed for reflection and refraction using the *generalized half-vector* formulation [WMLT07]

$$\mathbf{h}_{\lambda} = \hat{\mathbf{h}}_{\lambda} / \|\hat{\mathbf{h}}_{\lambda}\| \quad \hat{\mathbf{h}}_{\lambda} = -\left(n_{\lambda}^i \mathbf{i}_{\lambda} + n_{\lambda}^o \mathbf{o}_{\lambda}\right) \quad (12)$$

where n_{λ}^i and n_{λ}^o are the indices of refraction of the inside and outside media at the base vertex. Note that these indices are different for the half-vector $\mathbf{h}_{\lambda'}$, since they depend on the wavelength λ . With a similar incoming direction but a different wavelength, the same half-vector corresponds to different refracted directions, as seen on the first interaction in Fig. 3. As long as wavelength-dependent materials are encountered by both paths, the shifted path direction is deterministically computed by replicating the base path half-vector. Intuitively, by duplicating the wavelength-dependent half-vector, we ensure that the base and offset path maintain a similar throughput. Furthermore, we apply the half-vector shift for materials which are independent of wavelength but represent specular scattering. Once two consecutive vertices whose scattering functions are neither (near-)specular nor wavelength-dependent are encountered, the offset path can be reconnected to the base path. After the reconnection, the offset path will follow the base path again until the next wavelength-dependent interaction diverges the pair once again, and the above process repeats.

Jacobians To evaluate the integration in Eq. 8, we need to compute the Jacobian determinant of the shift mapping $S_{\lambda \rightarrow \lambda'}$ which reflects the change in density for the path-space of the offset path. In our case, the Jacobian determinant is simply the product of the local changes at each path vertex. A change in path density is only introduced after the base and offset paths diverge and the new outgoing direction at each disconnected vertex of the offset path is chosen by either reflection, refraction or reconnection. The Jacobian determinant for reflection and reconnection events do not depend on the wavelength and we refer to [KMA*15] (Sec 5.2) for their def-

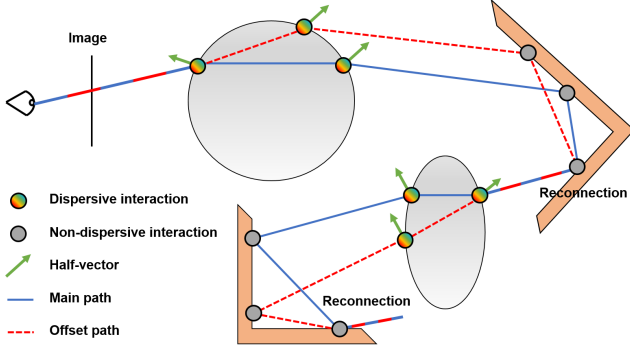


Figure 3: Spectral Shift concept. A base/offset path pair is connected until the first dispersive event, and reconnected as soon as two consecutive connectable and wavelength-independent events are found. The process is repeated as long as dispersive events are encountered.

initiation. However, in case of a refractive event, the local Jacobian determinant needs to reflect the change in path density introduced by the changes in the path geometry and the wavelength, and it is defined as:

$$\left| \frac{\partial \mathbf{o}_{\lambda'}}{\partial \mathbf{o}_{\lambda}} \right| = \left(\frac{n_{\lambda}^o}{n_{\lambda'}^o} \right)^2 \frac{|\mathbf{o}_{\lambda} \cdot \mathbf{h}_{\lambda}|}{|\mathbf{o}_{\lambda'} \cdot \mathbf{h}_{\lambda'}|} \left(\frac{n_{\lambda'}^i (\mathbf{i}_{\lambda'} \cdot \mathbf{h}_{\lambda'}) + n_{\lambda'}^o (\mathbf{o}_{\lambda'} \cdot \mathbf{h}_{\lambda'})}{n_{\lambda}^i (\mathbf{i}_{\lambda} \cdot \mathbf{h}_{\lambda}) + n_{\lambda}^o (\mathbf{o}_{\lambda} \cdot \mathbf{h}_{\lambda})} \right)^2 \quad (13)$$

A detailed derivation of this expression can be found in [Sta01] and [WMLT07].

3.2. Poisson Reconstruction

After computing an estimate of the spectral distribution I of a pixel and its gradients G , we perform a one-dimensional screened Poisson reconstruction to get the final pixel result:

$$\hat{I} = \arg \min_{\hat{I}} \|\nabla \hat{I} - G\|_2^2 + \alpha \|\hat{I} - I\|_2^2 \quad (14)$$

Hereby, the spectral distribution provides the lower frequencies while the higher ones are taken from the gradient estimates. Since the latter are assumed to have a lower variance, chromatic noise in the final pixel value \hat{I} is to be reduced. The gradient operator ∇ is defined as the finite differences between two points in the spectrum with distance δ . The parameter α controls the balancing between the least-squares fitting to I and G . We empirically evaluate the influence of α in the spectral domain in Sec. 4. Using the L2-norm, the final pixel result will be unbiased (see [LKL*13], Sec. 6 for a proof). In contrast to gradient-domain rendering in the image domain, we have no large 2D Poisson problem, but many small and independent one-dimensional ones. We rely on a 1D variant of the Fourier-based method from Bhat et al. [BCCZ08] to solve Eq. 14 efficiently.

3.3. Performance-Oriented Sampling

Gradient sampling is more expensive than traditional sampling, since an offset path is needed for each base path sample. Although the offset path often shares vertices with its base path, the overhead of tracing the diverging segments is non-negligible. For example, gradient-domain path tracing shifts each base path four times for symmetric gradient sampling in the vertical and horizontal image dimension and Kettunen et al. report an overhead factor of ≈ 2.5 per sample. Since the spectral domain is one-dimensional, symmetric sampling requires only two shifts. However, a spectral shift can be more expensive because an offset path can diverge from the base path multiple times. To reduce this overhead, we choose between standard and spectral gradient sampling depending on the contribution of a path pair via a user-defined threshold τ , as explained below.

The gradient $\Delta_{\lambda\lambda'}$ can be estimated in an unbiased way using a mixture of several sampling strategies, which has been expressed in a general form by Manzi et al. as the Multiple Weighted Gradient Integrals technique (see [MRK*14], Sec. 3.3). While originally used to reduce singularities in the path-space, we employ the technique to improve performance. In our case, we have two sampling strategies: uncorrelated standard sampling -Eq. 7) and sampling using correlated pairs of paths (Eq. 8). We can express the partial gradient from λ in the symmetric formulation of Eq. 9 as a mixture between these two strategies using a binary weight w^s and its complement $\bar{w}^s = 1 - w^s$:

$$\Delta_{\lambda\lambda'}^{\lambda} = \int_{\Omega_{\lambda}} w^s(\bar{\mathbf{x}}, \lambda) w_{\lambda\lambda'}(\bar{\mathbf{x}}, \lambda) g_{\lambda\lambda'}(\bar{\mathbf{x}}, \lambda) + \bar{w}^s(\bar{\mathbf{x}}, \lambda) f(\bar{\mathbf{x}}, \lambda) d\mu(\bar{\mathbf{x}}) \quad (15)$$

The weight w^s is defined as:

$$w^s(\bar{\mathbf{x}}, \lambda) = \begin{cases} 1 & \text{if } \min(f(\bar{\mathbf{x}}, \lambda), f(S_{\lambda \rightarrow \lambda'}(\bar{\mathbf{x}}), \lambda')) > \tau \\ 0 & \text{otherwise} \end{cases} \quad (16)$$

When the gradients are sampled symmetrically, a base path for λ appears either simply as a sample for the integral of I_{λ} or in the gradient formulation, where it is guaranteed that its offset path is also shifted conversely. The value τ balances between standard and gradient sampling and larger values mean that less spectral shifts are performed. Since we designed the offset path to have a similar contribution as its base path, it is likely that both paths either fail or pass the shifting criteria together. We evaluate its influence in Sec. 4.

3.4. Reusing Gradient Samples for Multiple Wavelengths

So far, a gradient sample only contributed to the wavelengths λ and λ' that are associated with the base and offset paths. We combine our approach with Hero wavelength sampling [WND*14]. Hero wavelength sampling chooses a single wavelength λ_0 for path propagation, but lets each path sample contribute to several other wavelengths which are evenly distributed over the spectrum. Similarly, we can select two sets $(\dots, \lambda_{-2}, \lambda_{-1}, \lambda_0, \lambda_1, \lambda_2, \dots)$ and $(\dots, \lambda'_{-2}, \lambda'_{-1}, \lambda'_0, \lambda'_1, \lambda'_2, \dots)$ with C wavelengths which are positioned at equidistant locations around the wavelengths λ and λ' . Each wavelength λ'_j is offset from λ_j by δ . We now let the base and




Figure 4: Hero Wavelength Sampling combination. We show how to combine our Spectral Gradient Sampling technique with Hero Wavelength Sampling. Similar to [WND*14], we evaluate our base and shifted paths for their own wavelength and also the rotated wavelengths, thus covering the whole spectrum. The contributions of each path for each wavelength are combined using an adapted multiple importance sampling weight.

offset path of the gradient sample contribute to all wavelengths in their corresponding sets. Intuitively, this process can be interpreted as computing multiple versions of the gradient which are "rotated" around the spectrum as illustrated in Fig. 4. For each $j \in C$, an estimator for Eq. 8 that performs Hero wavelength sampling is given by:

$$\Delta_{\lambda, \lambda'_j} \approx \frac{1}{N} \sum_{i=1}^N w_{\lambda}(\bar{\mathbf{x}}_i, \lambda_j) g_{\lambda, \lambda'_j}(\bar{\mathbf{x}}_i, \lambda_j) \quad (17)$$

The multiple importance weight $w_{\lambda}(\bar{\mathbf{x}}_i, \lambda_j)$ accounts for every sampling technique that is able to sample the path pair $(\bar{\mathbf{x}}, \lambda_j)$ and $(\bar{\mathbf{y}} = S_{\lambda \rightarrow \lambda'}(\bar{\mathbf{x}}), \lambda'_j)$. This pair can be directly sampled with the wavelength λ_j or it can be obtained by sampling the shifted path $\bar{\mathbf{y}}$ with the wavelength λ'_j , which is then shifted. Additionally, $\bar{\mathbf{x}}$ can also be sampled with any other wavelength λ_k and then rotated, or finally be generated by being shifted after sampling $\bar{\mathbf{y}}$ with any other wavelength λ'_k . Hence, we define the weight as:

$$w_{\lambda}(\bar{\mathbf{x}}, \lambda_j) = \frac{p(\bar{\mathbf{x}}, \lambda_j)}{\sum_{k=1}^C p(\bar{\mathbf{x}}, \lambda_k) + p(S_{\lambda \rightarrow \lambda'}(\bar{\mathbf{x}}), \lambda'_k)} \quad (18)$$

A small overhead is introduced when computing the needed additional probabilities. This sampling technique is most beneficial to glossy and diffuse wavelength-dependent interactions, where a path can contribute to multiple wavelengths.

4. Results

We integrated our method into the Mitsuba renderer [Jak10], building on top of the standard path tracing implementation. While Mitsuba supports spectral rendering, it does not provide any wavelength-dependent materials and we added dispersions based on Cauchy's equation [JW01] to the specular and rough dielectric scattering functions. Mitsuba stores the spectral distribution of each pixel in discretized form and the visible spectrum ranges from 360 nm to 830 nm. The visible spectrum is considered cyclic: if the shifted wavelength λ' falls outside the range, it is put back on the other end of the spectrum. We used a discretization of 15 equally-sized bins, where each bin represents a range of around 31 nm. The




Figure 5: δ and α evaluation. We evaluate the wavelength spacing δ and the Poisson reconstruction parameter α for our test scenes averaged over multiple sampling rates. The relMSE ratio defines the gain of our method over standard sampling. (a) The trend shows that δ values around 0.15 provide the best results. (b) Larger values for α favor the standard estimate during the Poisson reconstruction. A reasonably good value can be found at $\alpha = 0.2$ which coincides with the findings for gradient-domain path tracing.

implementation of Mitsuba and our method is fully CPU-based, implemented in C++ with multi-threading support.

We evaluate our method for three test scenes shown in Fig. 7(a). In the POOL scene, a dispersive water surface is viewed through a wavelength-dependent glass window. This scene should be beneficial for our approach, as it exhibits several paths that can benefit from our solution. The SUZANNE scene is a complex object made of rough glass projecting colored caustics on the ground. This scene is optimal for Hero wavelength sampling. Additionally, the less advantageous configuration makes it a worst-case-like scenario for our solution. The TORUS scene consists of a diffuse object encased in perfect glass, showcasing complex specular-diffuse-specular paths. The perfect glass lets Hero wavelength sampling mostly revert to path tracing, while our approach can handle the material. Still, the geometric configuration only allows a few paths to profit from our approach. We focus our evaluation for each scene on two image regions which are strongly affected by wavelength-dependent phenomena (the regions are shown in Fig. 7). The reference images were rendered on a CPU cluster with 56 cores using 2 to 5 million samples per pixel which required several days. All other results were captured on a Windows 10 PC with an Intel Core i7-4770 CPU with 3.40 GHz and 16GB of system memory. We report errors using the commonly-used relative mean-square error (relMSE) [RKZ12] but also the perception-based structural similarity (SSIM) [WBSS04] in Fig. 7 and Fig. 8, since reducing color noise greatly improves the image's visual quality. The reported timings include the sampling as well as the Poisson reconstruction step. The latter only requires a few hundred milliseconds for a mega-pixel image and is typically negligible compared to the total rendering time.

Parameter Evaluation First, we evaluate the wavelength spacing parameter δ for all three scenes averaged over multiple sampling rates using an α value of 0.2 (proposed in [KMA*15]). We report the ratio between the standard path tracing error and our method for equal sampling rates as an indicator for the achieved gain. We express δ as a percentage over the whole spectrum (360-830 nm) to make the evaluation suitable for other types of spectral repre-

representations besides discretized spectra. The highest gain is reached around $\delta = 0.15$, which corresponds to 15% of the spectral range (Fig. 5(a)). Consequently, we used this value as a default setting for all evaluations.

The parameter of the Poisson reconstruction α was empirically tested (Fig. 5(b)). Again, the gain of our method is compared to path tracing with the same sampling rate. As expected, a large α reduces the impact of gradient sampling and reduces the gain. Following [KMA*15], $\alpha = 0.2$ gives reasonable values and is used for the following evaluations.

Our method can tradeoff gradient sampling and standard path tracing (Sec. 3.3) via a shift threshold parameter τ , as illustrated in Fig. 6. Increasing τ implies that fewer paths are shifted, which decreases the efficiency of our method but also reduces time overhead. However, the gain of our method decreases slower since the paths omitted from shifting have rather small contributions early on. Our method achieves a gain even for larger thresholds, where very few paths are shifted. These paths are typically outliers, which heavily benefit from our spectral gradient sampling. We report statistics on the percentage of performed shifts and redundant shift attempts in Table 1 and it can be seen that the overhead from redundant shifts is small, especially for lower thresholds. The threshold value depends on the spectral complexity of the scene and we choose trade-off values of 2.5 (SUZANNE), 5.0 (TORUS), and 10.0 (POOL) for the following comparisons.

Comparison with previous approaches We compare our method to standard path tracing and Hero wavelength sampling in Fig. 7. We do not include the combination with Hero wavelength sampling (Sec. 3.4), as our tests showed that the difference to standard Hero wavelength sampling is small. For all other techniques, we present visual comparisons with varying sampling rates, as well as convergence plots (in terms of relMSE and SSIM; the axes are given in the log-scale, except for the SSIM axis). It can be seen that our method usually improves convergence over path tracing and gives visually more pleasant results with less chromatic noise, except for the TORUS scene, where both results are similar. An explanation is that our spectral shift is based on the mapping from gradient-domain path tracing which is not as efficient for specular-diffuse-specular paths, as described in [KMA*15]. In the POOL and TORUS scenes, we clearly outperform Hero wavelength sampling, as it cannot handle perfectly specular materials and leads to an overhead. When the first bounce after the camera is diffuse and not dispersive (as in Fig. 7(c) and (g)), our method performs the best since paths mostly bounce on diffuse surfaces and encounter few dispersive materials. In contrast, the SUZANNE scene is optimal for Hero wavelength sampling and a worst-case scenario for our solution, yet our method performs almost on par. In this scene, both shown insets exhibit the same trend since they represent the same conditions: a caustic formed by a rough glass object.

Denosing filtering Given that our approach reduces color noise, it can also prove beneficial when applying filtering to denoise the results – see Fig. 8. The SUZANNE scene’s glass in this example is not rough as in Fig. 7, but instead clear, which is a case that cannot be handled well by Hero wavelength sampling. We employed a bilateral filter whose parameters were chosen manually to obtain




Figure 6: Shift Criteria Threshold. The plot shows the influence of the shift criteria parameter τ on time overhead and achieved gain (relative to standard path tracing for the same sampling rate). While the gain of our method reduces with increasing τ , the time overhead drops faster. By adjusting the parameter for each scene, a good trade-off can be found.

τ	SUZANNE	TORUS	POOL
0.1	48.3% / 8.6%	64.5% / 17.4%	42.52% / 7.9%
1.0	17.8% / 3.9%	32.5% / 10.1%	38.87% / 7.8%
10.0	5.0% / 1.5%	0.04% / 0.02%	28.07% / 6.6%

Table 1: Shifting Statistics. The table shows the average percentage of base paths that fulfill the shifting criteria for our test scenes rendered with 256 samples per pixel. The second reported number is the percentage of shifted paths for which the base path passes the criteria but the shifted path does not.

a good result for Hero wavelength sampling and then kept identical for our approach. The filtered images improve significantly in quality and our improved color correlation manifests itself in terms of an improved relMSE and SSIM.

Spiky illuminants A limitation that our method shares with previous techniques is efficient handling of luminaires with spiky spectral distributions. When a ray hits a luminaire with a wavelength inside the illuminant spike, the shifted ray can fall off the interesting spectrum region in case of a spiky spectrum. The gradient between the two rays is substantial since the shifted ray conveys very little energy. Nevertheless, covering only a small spectral range, the advantage of gradient sampling for spiky illuminants was too small with respect to its cost, as confirmed by various tests we ran. Additionally, sampling rays that do not contribute to the final image is costly. Hence, with our cost-reduction scheme described in Sec. 3.3, such low-energy rays are actually discarded. Here, our technique automatically falls back to standard spectral path tracing. The behavior of Hero wavelength sampling faces similar difficulties and does not perform well – a rotated wavelength is a quarter of spectrum away from the main wavelength and falls outside of the spectral spike as well.

5. Conclusion

We introduced *spectral gradient sampling* as a new noise reduction technique for spectral Monte-Carlo light transport. Our approach

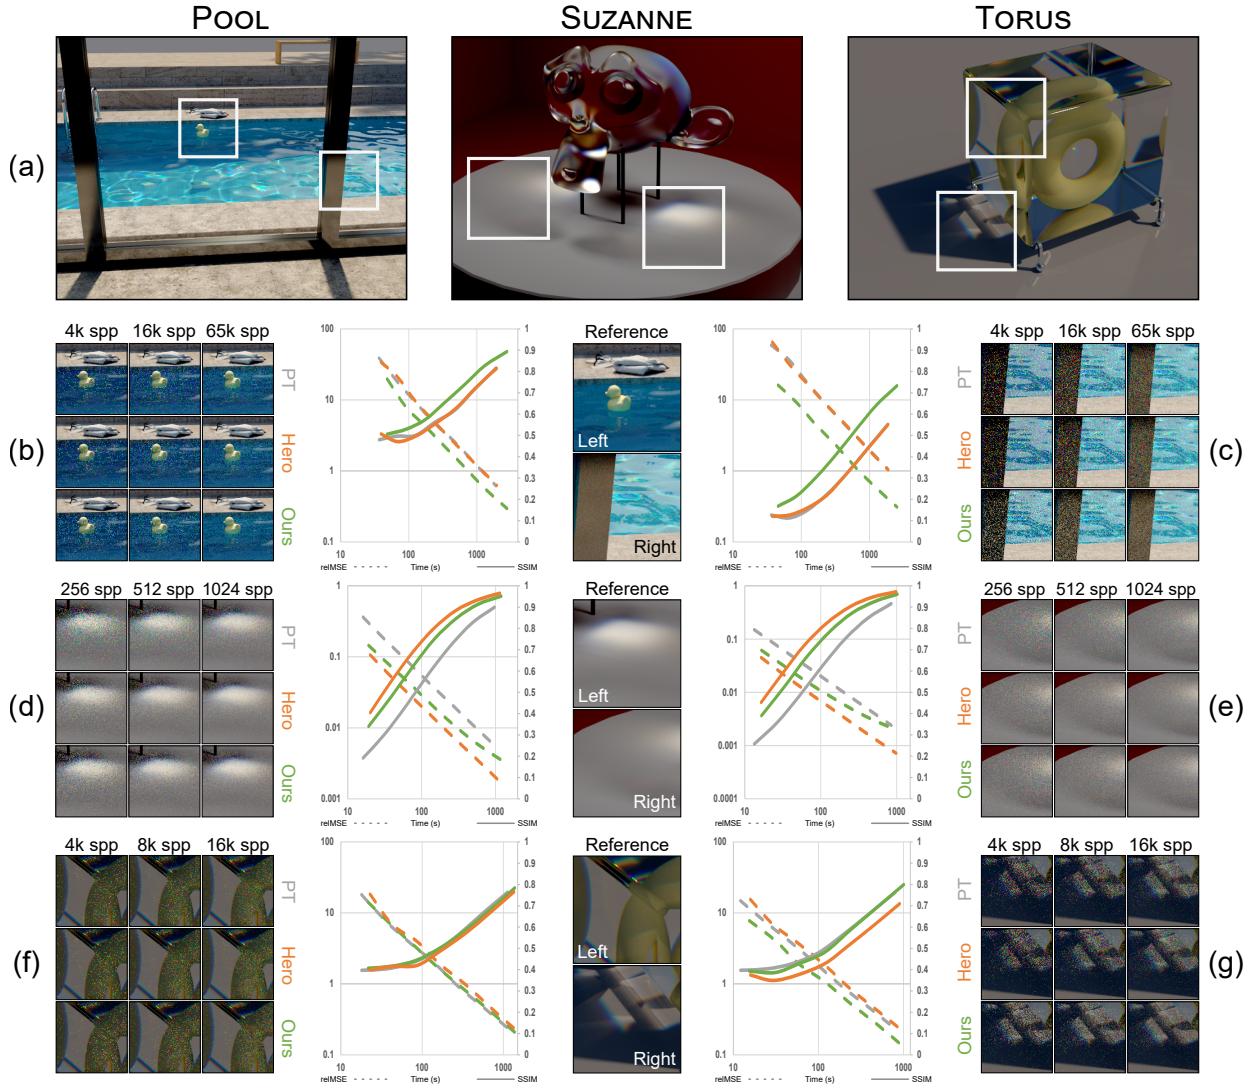


Figure 7: Comparison with previous approaches. The figure shows a visual comparison of our method to standard path tracing (PT) and Hero wavelength sampling (Hero) for multiple sampling rates. Furthermore, convergence plots are provided for a numerical equal-time comparison (note that the "relMSE" and "Time" axis are given in the log-scale). Our method provides significant improvement over path tracing and performs better than Hero wavelength sampling for scenes for which the latter reverts to path tracing (e.g. for the POOL scene).

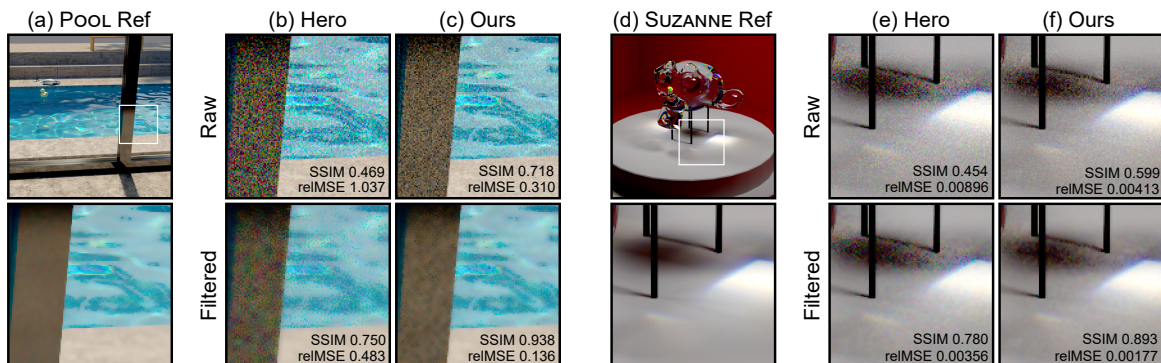


Figure 8: Filtering. With our reduced color noise, denoising (here with a bilateral filter) produces smoother results with our method. Both raw images were rendered with the same amount of samples using Hero wavelength sampling (b and e) or our method (c and f).

directly estimates the gradients of a pixel's spectral distribution using pairs of paths, which are correlated in the spectral domain. Final pixel colors are computed from the noisy spectral distributions and their estimated gradients by solving a 1D screened Poisson reconstruction. To generate pairs of correlated paths, we introduced a novel shift function which performs mappings in the spectral domain. Further, we proposed a weighting scheme to focus on high-energy paths, reducing the computational overhead of our method. Our approach can significantly reduce color noise and offers increased convergence when integrated in a conventional path tracer. While currently not beneficial in practice, our method is orthogonal to previous approaches, as demonstrated by combining it with Hero wavelength sampling.

Integration in bidirectional path sampling and application in a Metropolis Light Transport context seem fruitful directions for future work. Finally, we would like to extend our spectral shift mapping beyond the spectral domain (e.g., explore a combination with image and temporal gradients) and investigate more advanced shift mappings to support scenes with complex specular-diffuse-specular transport.

Acknowledgments

This work is part of the research program "LED it be 50%" (project number P13-20), which is funded by the Netherlands Organisation for Scientific Research (NWO) and supported by LTO Glaskracht, Philips, Nunhems, WUR Greenhouse Horticulture. We thank Wenzel Jakob for making the Mitsuba renderer [Jak10] publicly available. We also thank Ondřej Karlík for the POOL scene, and Cline et al. [CTE05] for the TORUS scene.

References

- [BCCZ08] BHAT P., CURLESS B., COHEN M., ZITNICK C. L.: Fourier analysis of the 2D screened poisson equation for gradient domain problems. In *Proceedings of the 10th European Conference on Computer Vision: Part II* (2008). 3, 5
- [BDM09] BERGNER S., DREW M. S., MÖLLER T.: A tool to create illuminant and reflectance spectra for light-driven graphics and visualization. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* (2009). 2
- [BPE17] BAUSZAT P., PETITJEAN V., EISEMANN E.: Gradient-domain path reusing. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH Asia)* (2017). 3
- [CTE05] CLINE D., TALBOT J., EGBERT P.: Energy redistribution path tracing. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* (2005). 9
- [DF03] DREW M. S., FINLAYSON G. D.: Multispectral processing without spectra. *J. Opt. Soc. Am. A* (2003). 2
- [DMCP94] DEVILLE P. M., MERZOUK S., CAZIER D., PAUL J. C.: Spectral data modeling for a lighting application. In *Computer Graphics Forum (Proceedings of Eurographics)* (1994). 2
- [EM99] EVANS G. F., MCCOOL M. D.: Stratified wavelength clusters for efficient spectral Monte Carlo rendering. In *Proceedings of the 1999 Conference on Graphics Interface* (1999). 1, 2
- [HGNH17] HUA B.-S., GRUSON A., NOWROUZSAHRAI D., HACHISUKA T.: Gradient-domain photon density estimation. *Computer Graphics Forum (Proceedings of Eurographics)* (2017). 3
- [Jak10] JAKOB W.: Mitsuba renderer, 2010. <http://www.mitsuba-renderer.org>. 6, 9
- [JM12] JAKOB W., MARSCHNER S.: Manifold exploration: A Markov Chain Monte Carlo technique for rendering scenes with difficult specular transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* (2012). 3
- [JW01] JENKINS F., WHITE H.: *Fundamentals of Optics*. McGraw-Hill Education, 2001. 6
- [KMA*15] KETTUNEN M., MANZI M., AITTALA M., LEHTINEN J., DURAND F., ZWICKER M.: Gradient-domain path tracing. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* (2015). 3, 4, 6, 7
- [LKL*13] LEHTINEN J., KARRAS T., LAINE S., AITTALA M., DURAND F., AILA T.: Gradient-domain metropolis light transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* (2013). 1, 2, 3, 5
- [Mey88] MEYER G. W.: Wavelength selection for synthetic image generation. *Comput. Vision Graph. Image Process.* (1988). 2
- [MKA*15] MANZI M., KETTUNEN M., AITTALA M., LEHTINEN J., DURAND F., ZWICKER M.: Gradient-domain bidirectional path tracing. In *Eurographics Symposium on Rendering - Experimental Ideas & Implementations* (2015). 3
- [MKD*16] MANZI M., KETTUNEN M., DURAND F., ZWICKER M., LEHTINEN J.: Temporal gradient-domain path tracing. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* (2016). 3
- [MRK*14] MANZI M., ROUSSELLE F., KETTUNEN M., LEHTINEN J., ZWICKER M.: Improved sampling for gradient-domain metropolis light transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* (2014). 3, 5
- [Pee93] PEERCY M. S.: Linear color representations for full speed spectral rendering. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (1993). 2
- [RBA09] RADZISZEWSKI M., BORYCZKO K., ALDA W.: An improved technique for full spectral rendering. *Journal of WSCG* (2009). 1, 2
- [RKZ12] ROUSSELLE F., KNAUS C., ZWICKER M.: Adaptive rendering with non-local means filtering. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH Asia)* (2012). 6
- [RP97] ROUGERON G., PÉROCHE B.: An adaptive representation of spectral data for reflectance computations. In *Eurographics Workshop on Rendering* (1997). 2
- [SSC*17] SUN W., SUN X., CARR N. A., NOWROUZSAHRAI D., RAMAMOORTHI R.: Gradient-domain vertex connection and merging. In *Eurographics Symposium on Rendering - Experimental Ideas & Implementations* (2017). 3
- [Sta01] STAM J.: An illumination model for a skin layer bounded by rough surfaces. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques* (2001). 5
- [Vea97] VEACH E.: *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, 1997. 2, 4
- [VG97] VEACH E., GUIBAS L. J.: Metropolis light transport. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (1997). 2
- [WBSS04] WANG Z., BOVIK A. C., SHEIKH H. R., SIMONCELLI E. P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* (2004). 6
- [WMLT07] WALTER B., MARSCHNER S. R., LI H., TORRANCE K. E.: Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques (EGSR)* (2007). 4, 5
- [WND*14] WILKIE A., NAWAZ S., DROSKE M., WEIDLICH A., HANIKA J.: Hero wavelength spectral sampling. In *Proceedings of the 25th Eurographics Symposium on Rendering (EGSR)* (2014). 1, 2, 5, 6