

Spectral Mollification for Bidirectional Fluorescence

A. Jung¹, J. Hanika¹ and C. Dachsbacher¹

¹Karlsruhe Institute of Technology




Figure 1: An underwater scene which combines fluorescent surfaces and challenging illumination. It features smooth dielectrics (the bottle) and a rough conductor (the can) both with wavelength dependent IOR, and a homogeneous chromatic medium. The light source is a large but focused spot light, similar to caustic illumination under water, which is hard to render with unidirectional methods. The insets show a path tracer with next event estimation at equal time. Right: the scene is illuminated exclusively by UV light, left: visible and UV light. Both pictures are rendered with our spectrally mollified bidirectional path tracer, enabling bidirectional rendering in the presence of fluorescence.

Abstract

Fluorescent materials can shift energy between wavelengths, thereby creating bright and saturated colors both in natural and artificial materials. However, rendering fluorescence for continuous wavelengths or combined with wavelength dependent path configurations so far has only been feasible using spectral unidirectional methods. We present a regularization-based approach for supporting fluorescence in a spectral bidirectional path tracer. Our algorithm samples camera and light sub-paths with independent wavelengths, and when connecting them mollifies the BSDF at one of the connecting vertices such that it reradiates light across multiple wavelengths. We discuss arising issues such as color bias in early iterations, consistency of the method and MIS weights in the presence of spectral mollification. We demonstrate our method in scenes combining fluorescence and transport phenomena that are difficult to render with unidirectional or spectrally discrete methods.

CCS Concepts

• **Computing methodologies** → Ray tracing; Reflectance modeling;

1. Introduction

Light transport simulation is employed extensively to create photo-realistic computer generated imagery. Since this process involves the computation of high dimensional integrals, the method of choice is usually Monte Carlo integration [PJH16]. Due to progress in path sampling methods, appearance modelling, and scalability of the algorithms, the Monte Carlo method has found widespread adoption in the industry in recent years [KFF*15].

For most faithful reproduction of colors, indirect light has to be computed by spectral rendering, not using RGB working spaces [WEV02, FHF*17].

Due to the constraint of energy conservation, there is a limit on how bright a surface of certain color saturation can appear if it only reflects light. If fluorescence is involved, more saturated colors can be achieved [JWH*19]. This is a natural phenomenon that can be observed in many important practical scenarios: garments, printing paper, some flowers, and organic matter dissolved in ocean water all fluoresce.

Including fluorescence into a rendering system requires special care: the inelastic scattering process changes the wavelength of the traced photon. This can heavily affect the geometry of subsequent scattering events: dielectric materials have an index of refraction

that typically varies with wavelength, thus changing the angle of refraction with wavelength. Even more dramatic changes are possible when a participating medium has chromatic extinction, i.e. the mean free path length varies with wavelength. It is thus relatively straight forward to include fluorescence in a unidirectional path tracer [WTP01], but bidirectional methods run into problems.

To support bidirectional path tracing with fluorescence, both sides would sample an individual wavelength, and start to trace their sub-paths. Different wavelengths result in potentially very different path geometry since the refraction angles at object boundaries and the segment lengths in participating media depend on wavelength. A standard spectral path tracer would choose the same wavelength on both sides. This enables connections between vertices on the eye and light sub-paths. With fluorescence, the two wavelengths need to differ, or else the fluorescent part of the transport is neglected. Also, the random walk will change the wavelength even if it was the same at the start. So the requirement to connect sub-paths and the requirement to sample the domain of wavelength change are in fundamental contradiction.

Figure 1 shows that sometimes, however, bidirectional methods provide superior sampling quality. It is thus desirable to maintain the option of using them in combination with all other required features in a renderer. We propose a method to include fluorescence in a bidirectional path tracer by employing *mollification*. Somewhat counter-intuitively, in this setup the regular, non-fluorescent interactions pose a problem. We mollify these, i.e. we replace a Dirac delta in wavelength domain by a finite-width 1D function which integrates to unity. Effectively, this adds a small amount of fluorescence to every scattering event. The rendering algorithm is akin to a collision estimator, or to kernel estimation methods: if the wavelength on the sub-path from the eye is close enough to the wavelength of the sub-path from the light, the connection is performed. To result in a consistent estimator, i.e. to make the resulting bias disappear in the limit, we shrink the mollification width over time.

We discuss the computation of correct weights for multiple importance sampling (MIS), and the effect of mollification bias on non-fluorescent appearance.

In summary, our contributions are

- a bidirectional rendering method for fluorescence based on mollification
- an analysis of mollification of reradiation matrices and continuous fluorescent reflectance
- a discussion of related MIS weights
- an analysis of bias vs. variance tradeoffs.

2. Background and Related Work

For an introduction to Monte Carlo path tracing algorithms, we refer the reader to PBRT [PJH16]. In our notation, we will use PT to signify path tracing from the camera, LT for light tracing from the lights and connecting to the camera, NEE for *next event estimation*, and BDPT for bidirectional path tracing.

Mollification Light transport simulation depends on solving an integration problem. This integrand may contain singularities in the

form of Dirac delta functions (e.g. perfect mirrors). These can pose a problem for instance for NEE: sampling a point on the light source also fixes the outgoing direction from the current shading point, but the Dirac delta function will always evaluate to zero (except for connections carefully crafted to fulfill the angular constraint [HDF15]). To resolve this, the Dirac delta can be replaced by a *mollifier*: a sequence of functions integrating to one which converges to the delta function, but initially smoothes out the singularity. The smoothing radius is iteratively shrunk to ensure the bias goes to zero, i.e. the algorithm stays consistent. The way this shrinking is performed presents a way to trade variance for bias.

Earlier work on mollification [KD13b] is concerned with such specular connections. They regularize the *bidirectional scattering distribution function* (BSDF) with a piecewise constant mollifier. The radius is decreased iteratively to maintain a consistent estimator. They perform *multiple importance sampling* (MIS) in the sense that the least biased strategy gets a weight of one, and they provide more theoretical background on mollification.

In practical scenarios, regularizing materials to aid global illumination has a long history [DH09]. This is often performed by changing the roughness of microfacet surfaces [JG19]. Bidirectional lightcuts [WKB12] regularize the directional spread of eye sub-path tracing in a many-lights rendering context. Photon mapping and derivatives [GKDS12] [HPJ12] can be seen as a spatial type of mollification and also employ radius shrinking techniques [KZ11] to maintain consistency. Kaplanyan and Dachsbacher [KD13a] present a data adaptive strategy to determine radius and shrinking rate. In the context of transient rendering, [JMMn*14] estimate temporal density with a kernel similar to ours.

Fluorescence Fluorescent molecules can absorb light at some wavelength and re-emit it at another, typically longer, wavelength. The emitted wavelength is usually independent of the absorbed wavelength [Lak10]. Due to this independence, the probability of a photon being absorbed by a fluorescent molecule and the distribution of emitted wavelengths can be described by an *absorption* and *emission* spectrum. Whenever the absorption spectrum extends to the ultraviolet range, this means that invisible light may be reradiated as visible light, causing exceptionally bright colors or even a glow effect under ultraviolet light.

Bispectral rendering equation In order to accurately simulate fluorescence we rely on spectral rendering, i.e. simulating and evaluating light transport paths for a given wavelength instead of constructing a path independent of its wavelength and evaluating it for a number of discrete color channels, as in traditional RGB rendering or [JA18]. We solve the *bispectral rendering equation*

$$L_o(\omega_o, \lambda_o) = L_e(\omega_o, \lambda_o) + \int_{\Omega} \int_{\Lambda} L_i(\omega_i, \lambda_i) f(\omega_i, \lambda_i, \lambda_o, \omega_o) \cos \theta_i d\lambda_i d\omega_i. \quad (1)$$

described in [HHA*10]. In addition to the traditional form of the rendering equation [Kaj86], this includes an integration over wavelengths Λ , which are absorbed and re-emitted as λ_o as described by a *bispectral bidirectional reflectance and reradiation distribution function* (BBRRDF). Hullin et al. [HHA*10] define the BBRRDF

in terms of incident and exitant energy as

$$f_r(\omega_i, \lambda_i, \lambda_o, \omega_o) = \frac{d^2 L_o(\omega_o, \lambda_o)}{L_i(\omega_i, \lambda_i) \cos \theta_i d\omega_i d\lambda_i}, \quad (2)$$

where $L_o(\omega_o, \lambda_o)$ and $L_i(\omega_i, \lambda_i)$ are the reflected and incident spectral radiance respectively, both measured in $W/(sr \cdot m^2 \cdot nm)$.

Rendering Fluorescence Fluorescence as a visual effect has been explored early on in graphics [Gla95] but most methods so far rely on unidirectional path tracing [WTP01]. Gutierrez et al. [GMAS05] [GSMA08] apply photon mapping to fluorescence, including fluorescent volumes, but rely on discretizing the wavelength domain [GSMA08]. Hullin et al. [HHA*10] use a unidirectional path tracer with Kelemen-style Metropolis sampling to visualize their results. Bendig et al. [BHD*08] use a forward path tracer to simulate fluorescent concentrators for photovoltaic systems. Volumetric forward path tracing poses special challenges when difficult transport such as godrays under water and fluorescence are combined [MW18]. We address this issue by enabling the use of bidirectional methods. Abdellah et al. [ABE*17] also simulate fluorescence in media, and the fluorescence of natural waters has been considered [CS04].

Jarabo et al. [JA18] propose vectorized rendering to enable bidirectional transport with fluorescence. They evaluate path contributions for the full spectrum in discrete steps, whereas we use stochastic wavelength sampling to Monte Carlo integrate the wavelength domain. Their method ignores wavelength dependent refraction angles and segment lengths in participating media.

For typical camera shutter times fluorescence can be assumed to be instantaneous. Nalbach et al. [NSR17] investigate phosphorescence, which is similar to fluorescence except that it can occur over up to several seconds, which requires additional care.

This body of research lead to a few fluorescence-aware BSDF models. Glassner [Gla95] formalized this in the form of reradiation matrices, enabling energy exchange between discrete wavelengths. This has been extended to microfacet and Phong models [WWLP06]. Jung et al. [JHD18] model diffuse re-emission and parameterize the effect with continuous spectra instead of a reradiation matrix.

Acquisition of fluorescence is performed in biology and chemistry by measuring an absorption and emission spectrum [Lak10]. Databases of such measurements are available online, see for instance [May19, The19]. Tominaga et al. [THH18] estimate reradiation matrices from images. Hullin et al. [HHA*10] capture high dimensional bidirectional bispectral reflectance and reradiation distribution functions (BBRRDF) both in the wavelength and angular domains. Suo et al. [SBCD14] reconstruct reradiation matrices from sparse measurements.

Modeling fluorescent surfaces A fluorescent BBRRDF can be based on existing emission and absorption spectra [JHD18], or on a two-dimensional *reradiation matrix* [Gla95] [HHA*10], which describes how much light is converted from incident to exitant wavelengths in discrete steps. In order for a BBRRDF to be energy con-

serving, it has to fulfill the condition

$$\forall \omega_i \forall \lambda_i \int_{\Lambda} \int_{\Omega} f(\omega_i, \lambda_i, x, \lambda_o, \omega_o) d\omega_o^\perp d\lambda_o \leq 1. \quad (3)$$

In case of reradiation matrices with rows corresponding to exitant wavelengths and columns corresponding to incident wavelengths (as in Figure 3), energy conservation means that no column may sum up to more than 1. Not all energy conserving BBRRDFs are photon conserving, especially for reradiation towards longer wavelengths, as a photon's energy is inversely related to its wavelength.

We use the continuous BBRRDF by [JHD18], which covers both elastic and inelastic scattering. It is parameterized by an absorption spectrum $a(\lambda_i)$, an emission spectrum $e(\lambda_o)$, a non-fluorescent reflectance spectrum $r(\lambda)$ and two scalars $c, Q \in [0, 1]$:

$$f(\omega_i, \lambda_i, x, \lambda_o, \omega_o) = \frac{1}{\pi} \begin{cases} \delta_{\lambda_i, \lambda_o} (1 - c \cdot a(\lambda_i)) r(\lambda_i) & \text{if } \lambda_i = \lambda_o \\ c \cdot a(\lambda_i) \cdot Q \cdot e(\lambda_o) & \text{else} \end{cases} \quad (4)$$

This formulation does not allow the fluorescent component to re-emit light at its original wavelength. Since this corresponds to removing a zero-set from the continuous fluorescent part of the BBRRDF, this does not change the result of the integration of f over Λ in the rendering equation (1), but allows for an easier discussion later on. Note that we still need the δ in the non-fluorescent case in order for the integration over Λ to work out as expected.

As we will only mollify the non-fluorescent component of the BBRRDF, the methods we present also apply to mollifying the diagonal of a discrete 2D reradiation matrix instead.

Spectral singularities in the path integral In the context of the bispectral rendering equation, a non-fluorescent BSDF can be written as a BBRRDF using a spectral Dirac delta

$$f(\omega_i, \lambda_i, x, \lambda_o, \omega_o) = \delta_{\lambda_i, \lambda_o} f'(\lambda_i, \omega_i, x, \omega_o). \quad (5)$$

This results in a singularity in the bispectral rendering equation, similar to how a perfectly specular BSDF causes a geometrical singularity in non-spectral rendering. When rendering without fluorescence this singularity is avoided by omitting the integration over Λ and omitting the $\delta_{\lambda_i, \lambda_o}$ component of the BBRRDF. For BDPT without fluorescence, both sub-paths are simply traced for the same wavelength, which allows us to connect non-fluorescent vertices. Conversely, when sampling different wavelengths for eye and light sub-paths, we can still connect fluorescent BBRRDFs, but now regular non-fluorescent BSDF pose a problem.

3. Method

We now present our method for mollifying BBRRDFs to connect camera and light sub-paths for bidirectional path tracing. We start by discussing different strategies for mollifying the spectral δ -component in full reradiation matrices (section 3.1) and in the non-fluorescent part (section 3.2). Depending on the mollifier, this may cause energy loss towards invisible wavelengths, but also violate energy conservation by creating excess energy. Then, even though the mollified BBRRDFs are only evaluated but never sampled, we also need to be able to compute mollified probability density functions (PDFs, section 3.3) corresponding to the mollified




Figure 2: Left: an example reradiation matrix. Center: $r(\lambda)$ mollified with f_{mo} . Right: $(1 - ca(\lambda))r(\lambda)$ mollified with f_{mi} ($d = 50$).

BBRRDFs in order to combine the different sampling strategies inside a BDPT (section 3.4) by computing multiple importance sampling (MIS) weights (section 3.5). The mollifier which replaces the δ -component further needs to shrink over iterations in order to keep the estimator consistent, which we outline in section 3.6.

3.1. Mollifying reradiation Matrices

When working with reradiation matrices, we can choose to mollify the full matrix by using a 2D blur kernel on the 2D matrix, or we mollify the diagonal only. The second option is preferable, since expanding and blurring the full reradiation matrix is expensive and leads to unwanted issues with energy conservation. We discuss this in more detail in the supplemental document. For the purpose of performing bidirectional connections, it is sufficient to mollify the non-fluorescent part of the BBRRDF, i.e. the diagonal of the reradiation matrix. In our implementation, we chose to work with an analytic BBRRDF model [JHD18] which allows us to mollify only the non-fluorescent part.

3.2. Mollifying a non-fluorescent BBRRDF

We consider non-fluorescent BBRRDFs of the form

$$f(\lambda_i \rightarrow \lambda_o, \omega_i \rightarrow \omega_o) = \delta_{\lambda_i, \lambda_o} r(\lambda_i = \lambda_o, \omega_i, \omega_o) \quad (6)$$

consisting of a δ -component and some non-fluorescent reflectance which may depend on incident and exitant direction $\omega_{i,o}$, but only on one wavelength.

We can base our mollification kernels on the indicator function

$$\mathbb{1}_X = \begin{cases} 1 & \text{if } X \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

So for example, a simple mollifier with radius d to replace $\delta_{\lambda_i, \lambda_o}$ could be $\mathbb{1}_{|\lambda_o - \lambda_i| < d} / 2d$, yielding

$$f_{moll}(\lambda_i \rightarrow \lambda_o, \omega_i \rightarrow \omega_o) = \frac{\mathbb{1}_{|\lambda_i - \lambda_o| < d}}{2d} r(\lambda_i, \omega_i, \omega_o). \quad (8)$$

However, this leaves the question how the spectral reflectance $r(\cdot)$ should be evaluated. We consider three options. The first two are to evaluate $r(\cdot)$ for λ_i or λ_o , resulting in a piecewise constant mollifier for either a given λ_o or λ_i , the third one is a box filter based mollification to create a reciprocal BBRRDF. Examples for all three are shown in Figure 3.




Figure 3: From left to right: The reflectance spectrum of a non-mollified non-fluorescent BBRRDF as in equation (6), the BBRRDF mollified over incident (equation (9)) and exitant (equation (10)) wavelengths ($d = 50$), and the box filtered BBRRDF (equation (13)).

Constant mollification Evaluating $r(\cdot)$ for λ_o corresponds to gathering energy reflected to λ_o by $r(\lambda_o)$ over multiple incident wavelengths $[\lambda_o - d, \lambda_o + d]$:

$$f_{mi}(\lambda_i \rightarrow \lambda_o) = \frac{1}{2d} r(\lambda_o) \mathbb{1}_{|\lambda_i - \lambda_o| < d} \quad (9)$$

This strategy has the advantage that, under a constant light source, the apparent reflection of the BBRRDF remains unchanged, as this mollifier can be considered to blur incident light. However, when applying this strategy to fluorescent BBRRDFs the result may not be energy conserving anymore, as is shown in the supplemental document.

The second option, which we decided to use, is

$$f_{mo}(\lambda_i \rightarrow \lambda_o) = \frac{1}{2d} r(\lambda_i) \mathbb{1}_{|\lambda_i - \lambda_o| < d}, \quad (10)$$

which evenly reradiates incident energy at λ_i over a range $[\lambda_i - d, \lambda_i + d]$ according to $r(\lambda_i)$. The main advantage of this strategy is that it is fairly simple to implement as well as energy conserving

$$\int_{\Lambda} f_{mo}(\lambda_i \rightarrow \lambda_o) d\lambda_o = r(\lambda_i) \int_{\lambda_i - d}^{\lambda_i + d} \frac{1}{2d} d\lambda_o = r(\lambda_i). \quad (11)$$

In the supplemental document we show that it is also photon conserving. If working with final reradiation matrices, we recommend mollifying the diagonal using f_{mo} , which corresponds to mollifying $(1 - ca(\lambda))r(\lambda)$ in the BBRRDF model.

2D box filter mollification Both constant mollification strategies result in a non-reciprocal BBRRDF with discontinuities at the mollification borders. Instead, we can produce a continuous and reciprocal mollified BBRRDF by applying a two-dimensional box filter with side length $b = 2d$

$$g_b(\lambda_i, \lambda_o) = \frac{\mathbb{1}_{|\lambda_i - \lambda_o| < b/2}}{b^2} \quad (12)$$

to r , which yields

$$\begin{aligned} f_{mb} &= [f * g_b](\lambda_i \rightarrow \lambda_o, \omega_i \rightarrow \omega_o) \\ &= \frac{1}{b^2} \int_{\max(\lambda_i, \lambda_o) - b/2}^{\min(\lambda_i, \lambda_o) + b/2} r(\lambda, \omega_i, \omega_o) d\lambda. \end{aligned} \quad (13)$$

Note that $\int_x^y \dots$ evaluates to 0 if $x > y$. Also note that while this can be implemented as a 2D box filter operation, we still only mollify in one dimension for either a fixed λ_i or fixed λ_o .

Clamping Note that all three options leak energy out of both ends of the visible range. It is possible to construct a clamped version of f_{mo} , f_{mi} and even f_{mb} , which conserves all energy in the visible range. However we decided against it because a) the mollifier would depend on prior knowledge about the "visible wavelength range" at the camera and b) a clamped BBRRDF can no longer reflect ultraviolet light which might be converted to visible light by some fluorescent surface later on the path. Instead clamping it to visible plus ultraviolet light would in turn require prior knowledge about the minimum wavelength emitted in the scene, and might effectively result in an unclamped BBRRDF for other light sources.

One-sided kernel We also tested the kernel

$$f_{asym}(\lambda_i \rightarrow \lambda_o) = \frac{1}{2d} r(\lambda_i) \mathbb{1}_{\lambda_i < \lambda_o < \lambda_i + 2d} \quad (14)$$

with the intent of creating a BBRRDF more akin to a fluorescent BBRRDF, as real fluorescent materials typically convert shorter to longer wavelengths. However, Figure 6 and the supplemental document reveal that this kernel results in much stronger bias.

3.3. Mollified PDF

Generally, the PDF p_m corresponding to a mollified BBRRDF f_m is defined as any other PDF:

$$p_m^{cam}(\lambda_i | \lambda_o) = \frac{f_{moll}(\lambda_i \rightarrow \lambda_o)}{\int_{\Lambda} f_{moll}(\lambda \rightarrow \lambda_o) d\lambda} \quad (15)$$

for camera paths, and

$$p_m^{light}(\lambda_o | \lambda_i) = \frac{f_{moll}(\lambda_i \rightarrow \lambda_o)}{\int_{\Lambda} f_{moll}(\lambda_i \rightarrow \lambda) d\lambda} \quad (16)$$

for light paths. Depending on how the BBRRDF was mollified, these equations result in different terms with different simplifications, which we discuss in detail in the supplemental document. In general, we denote spectrally mollified PDFs as $\bar{p}(\lambda_x | \lambda_y)$. In case of constant mollification we get

$$p_{mo}^{cam}(\lambda_i | \lambda_o) = \frac{\mathbb{1}_{|\lambda_i - \lambda_o| < d} r(\lambda_o)}{\int_{\lambda_o - d}^{\lambda_o + d} r(\lambda) d\lambda} \quad (17)$$

$$p_{mo}^{light}(\lambda_o | \lambda_i) = \frac{\mathbb{1}_{|\lambda_i - \lambda_o| < d}}{2d} \quad (18)$$

for f_{mo} and analog PDFs for f_{mi} . This reveals an issue of our mollification technique, which is that for given ω_i, ω_o we need to be able to integrate the mollified BBRRDF over wavelengths in order to compute the correct PDF. However, since we only evaluate the mollified BBRRDF at a connection vertex, but never sample from it, we do not actually need to evaluate the correct PDF. Instead, we will only require a mollified PDF for computing MIS weights, which also work (though less optimally) with approximate values.

In case of the box mollification f_{mb} we get analog PDFs $p_{mb}^{cam}(\lambda_i | \lambda_o)$ and $p_{mb}^{light}(\lambda_o | \lambda_i)$. In this case there are no more simplifications and we end up with a double integration over $r(\cdot)$ in the denominator. The full equation and a solution for piecewise linear spectra $r(\cdot)$ are presented in the supplemental document.

3.4. Sampling strategies

Path sampling works analogous to regular BDPT. We sample an initial wavelength when starting from the camera or the light sources, and trace a path. We associate each edge of a path with an individual wavelength. During the random walk, new wavelengths are sampled at fluorescent interactions. A camera path samples the absorption spectrum, a light path samples the emission spectrum. NEE samples an additional wavelength for the connecting segment from the last vertex on the path before the connection (see Figure 4).

Inner connections (with more than one vertex on both camera and light sub-paths) sample the wavelength of the connecting segment from the camera sub-path. To make the connection work, the last light sub-path vertex is mollified spectrally. This decision is arbitrary, instead we could also mollify at the last camera vertex. The connection is rejected if the wavelengths are too far apart.

Thus the wavelength PDF of a connected path with $k+1$ vertices is $p(\lambda_1)p(\lambda_1 \rightarrow \lambda_2) \cdots p(\lambda_{k-1} \leftarrow \lambda_k)p(\lambda_k)$. The probability of NEE or an implicit path is $p(\lambda_1)p(\lambda_1 \rightarrow \lambda_2) \cdots p(\lambda_{k-1} \rightarrow \lambda_k)$.

3.5. Multiple Importance Sampling weights

Kaplanyan et al. [KD13b] mollify only paths that cannot be sampled by any bidirectional strategy. In such cases, a maximum distance heuristic is used in order to minimize overall bias.

In our test scenes any path can be sampled at least by pure forward path tracing and next event estimation on camera and light paths (i.e. there are no singular emitters). Additionally, especially if there is no fluorescence along the path, Kaplanyan et al.'s strategy of weighting the path with the longest connection distance with 1 does not translate well to mollifying in the wavelength domain as opposed to the geometry domain.

We want to make full use of all bidirectional regularized sampling strategies and therefore use MIS weights similar to the balance heuristic:

$$w_{c,l}(\mathbf{X}) = \frac{\bar{p}_{c,l}(\mathbf{X})}{\sum_{i=0}^{c+l} \bar{p}_{i,c+l-i}(\mathbf{X})}. \quad (19)$$

Note that we use the mollified PDF \bar{p} instead of the actual PDF in all places. This means we pretend that any wavelength except the first along each sub-path was sampled from the mollified BBRRDF at the previous vertex, and evaluate the PDF corresponding to the mollified BBRRDF.

By construction, the weights for a given path sum up to one. We also ensure that each PDF $\bar{p}_{c,l}$ has the same measure in the sense that they all cover the same number of spectral PDFs, even if there is no fluorescent event along the path. Alternatively, path space would have to be extended to consider fluorescent and non-fluorescent events at the same vertex as an additional dimension. This makes the analysis a lot more complex for layered surfaces which can be both at the same time. On the other hand in our approach most of the mollification factors at the vertices will cancel out in the weight computation in equation (19).

Since during path connections we use the wavelength of the camera sub-path for the connecting segment, we evaluate the mollified




Figure 4: Mollified PDFs $\bar{p}_{c,l}$ used in MIS weight computations. Green and red correspond to the camera and light sub-paths. We index the camera vertex as \mathbf{x}_0 and the light vertex as \mathbf{x}_{c+l-1} .

PDF at x_{c-1} for λ_c when computing $p_{c,l}$. An example for 4-vertex-paths is given in Figure 4.

3.6. Consistency of the estimator

Kaplanyan et al. [KD13b] show how the support of the mollifier has to be decreased over iterations depending on dimensionality in order to obtain a consistent estimator. Since we only mollify one dimension once per path, we need to decrease the sequence of mollification distances d_n as $\mathcal{O}(n^{-1}) < d_n < \mathcal{O}(1)$. As suggested in their work we decrease d_n as

$$d_n = d_0 n^{-s}, \quad (20)$$

which satisfies above condition for $s \in (0, 1)$. In section 4.5 we evaluate different choices of d_0 and s , including $s = 0$ (i.e. $d_n \equiv d_0$).

4. Evaluation

Unless mentioned otherwise we use constant mollification as in f_{mo} (equation (10)), an initial mollification distance of $d_0 = 100nm$ and a shrinking coefficient of $s = 1/4$. We further approximate the mollified non-fluorescent PDF $p_{mc}(\lambda_i|\lambda_o) \approx \frac{\mathbb{1}_{|\lambda_i - \lambda_o| < d_n}}{2d_n}$ for MIS weight computations.

4.1. Color bias

Our method produces color bias for two reasons. First, light bleeds to nearby wavelengths, thus blurring the involved spectra (see Figure 5). Yet Figure 7 reveals that even with constant reflectances and illumination, where blurring is not an issue, some bias remains. The reason for that is that camera paths start with an initial wavelength λ_c from some predefined visible range $[\lambda_{min}, \lambda_{max}]$, and light paths start with λ_l from $[\lambda_f, \lambda_{max}]$, where λ_f is the smallest simulated (ultraviolet) wavelength. Assuming λ_l and λ_c are sampled uniformly, we plot the probability of $|\lambda_l - \lambda_c| < d$ in Figure 8.

It shows that the probability for sampling a λ_l within range d decreases for λ_c close to the borders of the visible range (or towards longer wavelengths only, if UV light is supported). As a result, mollified images are biased towards colors in the center of the visible spectrum and lack red (and violet/blue, if no UV light is supported (Figure 9 right)). At small d , this effect is hardly noticeable, since colors at the borders of what we use as visible spectrum ($[360, 780]nm$) are hardly visible anyway, but for larger d (or if the visible range is smaller to begin with), this color bias becomes increasingly noticeable (see Figure 7 and Figure 9 (left)). Conceptually, this type of bias resembles boundary bias in photon mapping.

The probability in Figure 8 corresponds to the probability that a 2-vertex camera path for λ_c and a 2-vertex light path for λ_l in a scene without fluorescence are successfully connected using mollification as in equation (9), provided a connection is geometrically possible. We visualize this by rendering only such (2,2)-paths in Figures 5 and 7.

Figure 8 also shows that the lower λ_f , the lower the probability for accepting the resulting path. This does not cause bias since the resulting path contribution includes a division by $p(\lambda_f)$, which decreases accordingly.

As explained in section 3.6 we shrink d over time. Therefore any color bias will disappear in the limit. With a centered kernel, even at $d = 50nm$, which is still quite large, the bias is hardly visible. More analysis of the color bias is part of the supplemental material.

Boundary bias could be avoided using a kernel which explicitly depends on λ_o . For example, this can be achieved by clamping the mollified reradiation matrix. We discuss this option and why we decided against it at the end of section 3.2.

One-sided kernel We test the one-sided kernel from equation (14) in Figure 6 and the supplemental document. Instead of blurring light over both shorter and longer wavelengths as the other kernels, which mostly affects saturation and brightness, this kernel only shifts light in one direction, thus deviating strongly from the reference hue. This effect is strongest in the supplemental Figure 13, where light mollified on a colored surface is reflected off a grey surface (instead of off a same-colored surface as in Figure 6, which weakens the hue distortion).

4.2. Constant mollification vs. box filtering

We could not detect any significant difference between box filtered mollification as in equation (13) and constant mollification as in equation (10). Figure 10 shows one such example with fluorescent walls and ultraviolet illumination rendered with both methods.




Figure 5: This scene shows the back wall of a uniformly colored cornell box, rendered with a PT with NEE for 4-vertex-paths (left) and rendered exclusively with mollified (2,2)-paths (right). The mollified paths use a centered (equation (10)) constant mollification kernel (i.e., $s = 0$) with (top to bottom) $d \equiv 10, 20, 50, 100nm$. The supplemental document contains the reflectance spectra and more versions of this scene. With constant spectra (left), only boundary bias is an issue, which is hardly noticeable even at large d . The bleeding bias can become strong for colorful spectra, but considering the visible wavelength range covers roughly 400nm, radii like 50 or 100nm are impractically large anyway.




Figure 6: Same scene and kernel widths as in Figure 5, but with a one-sided kernel (equation (14)). Other than the centered kernel (equation (10)) in Figure 5, this kernel noticeably affects the hue.

4.3. Bidirectional vs. Unidirectional

It depends on the scene which sampling strategy works best, PT, NEE, LT, or BDPT. We present three scenes in which our method is superior to a regular PT with NEE (Figure 1, Figure 11, Figure 14). If in a scene without fluorescence simple PT/NEE is better than BDPT (e.g. if all paths are easily sampled unidirectionally), PT/NEE will also be better than our mollified BDPT in that scene with added fluorescence.

Fluorescence and smooth dielectrics Figure 11 shows a scene with both fluorescent and non-fluorescent grey walls containing smooth dielectric and metal spheres with wavelength dependent indices of refraction. The light source has a constant emission spectrum which extends to the ultraviolet range. Our BDPT does not connect to the perfectly specular spheres, so mollification only occurs for paths with two consecutive bounces off the walls, floor or ceiling. Since this scene is dominated by paths that are easily sampled by next event estimation from camera and light paths, mollification plays only a small role in this setup. Nevertheless our algorithm allows us to render the scene with a full BDPT.

Fluorescence and mollified metals Figure 12 shows a scene containing rough metals which are spectrally mollified by f_{mo} as in equation (10). They are modeled as a microfacet BSDF using a GGX distribution, with a wavelength-dependent complex index of refraction. We do not provide an RMSE plot for this scene since




Figure 7: Boundary bias, illustrated by rendering a grey cornell box with (2,2)-paths (second row), no shrinking (i.e., $s = 0$), compared to a path traced reference (top). The bottom rows show signed difference images multiplied by 64 (third row: mollified - PT, fourth row: PT - mollified). On the right, the constant illuminant includes ultraviolet light ($\lambda_f = 260nm$), so the boundary bias (initially) only affects long wavelengths.




Figure 8: Probability of uniformly sampling a wavelength λ_l for a light path from $[\lambda_f, 780]$ such that the distance to the wavelength from the camera $|\lambda_l - \lambda_c| < d = 100nm$, for two choices of λ_f : with and without UV light.

even at over 100k spp, the PT+NEE reference implementation has not properly converged (orange inset in Figure 12). Instead, equal sample and time comparisons show the superiority of BDPT in this setup. In Figure 13 we compare the same insets for $d_0 = 100nm$ and a shrinking coefficient $s = 1/4$ to $d_0 \equiv 5nm$ and $s = 0$.

Fluorescence in volumes Similar to Figure 1, Figure 14 illustrates a path tracer's issues with non-fluorescent volumes combined with a focused light source. The light tracer works well for volume paths but struggles with the UFO and light source.

4.4. Mollified vs. Unmollified BDPT

In this section we compare our mollified BDPT to the unmollified base version. This is only possible in scenes without fluorescence. The contributions of strategies with at most one camera or light vertex are the same for both renderers. For all other strategies, the mollified BDPT rejects a certain number of connections based on the wavelengths of the sub paths (also see Figure 8), which leads to slower convergence. If we use constant mollification f_{mo} (equation (10)) with current mollification distance d_n , the MIS weights of the mollified BDPT reduce to

$$\frac{\alpha_{c,l} P_{c,l}(\mathbf{X})}{\sum_{i=0}^{c+l} \alpha_{i,c+l-i} P_{i,c+l-i}(\mathbf{X})}, \quad (21)$$




Figure 9: Average image RGB for the grey wall scene in Figure 7 for varying mollification radius d (left) and minimum sampling wavelength λ_f (right).




Figure 10: Left: mollification of the non-fluorescent reflectance $r(\lambda)$ using a box filter as described in equation (13). Center: constant mollification of $r(\lambda)$ as in equation (10). Right: Difference times 32.

where $p_{c,l}(\mathbf{X})$ is the PDF used during sampling evaluated for the path \mathbf{X} , and

$$\alpha_{c,l} = \begin{cases} p(\lambda_1)p(\lambda_{c+l-1}) & \text{if } c, l > 1 \\ p(\lambda_1)/(2d) & \text{if } l \leq 1 \\ p(\lambda_{c+l-1})/(2d) & \text{if } c \leq 1. \end{cases} \quad (22)$$

Since typically $2d$ is (much) smaller than the range of wavelengths sampled at the camera (e.g. $[360, 780]nm$) or light source (e.g. $[260, 780]nm$), paths with at most one camera or light vertex get weighted overproportionally compared to the unmollified BDPT's MIS weights, and paths with at least 2 vertices on each sub-path receive lower weights.

We tried demonstrating this effect in the scene shown in Figure 16, which relies heavily on such paths. However it is far outweighed by the slower convergence due to rejection of sub-paths with wavelengths that are more than d_n nm apart. The figure also contains an image with a hack replacing $\alpha_{j,k}$ with $\alpha_{1,k}$ wherever $i, j > 1$ to bring the mollified MIS weights closer to the unmollified weights. This shows close to no difference to the original mollified MIS weights in equation (21).

4.5. Convergence rate

Scenes without fluorescence Any scene without fluorescence that can be handled by our BDPT can also be rendered more efficiently by a typical unmollified BDPT. The unmollified version does not reject connections based on wavelengths but instead samples both sub-paths with the same wavelength.

Scenes with fluorescence Variance reduction of our mollified BDPT can be improved by increasing the initial mollification distance d_0 so that fewer paths are rejected, and decreasing the shrinking coefficient s . Vice versa, bias in the long run is reduced by decreasing d_0 and increasing s . Section 4.1 shows bias depending on the current mollification distance d_n .

Too aggressive shrinking, while theoretically consistent, results in too much variance to converge in reasonable time (Figure 17 $s = 3/4$). We illustrate the behavior of d_n in Figure 21.

Root mean squared error (RMSE) Increasing the kernel width increases the number of successfully mollified paths, thus decreasing variance and initially decreasing the RMSE, as shown in Figure 18. In particular, using an n times wider kernel results in roughly n times as many contributing paths. Therefore the mollified method with an n -times wider kernel achieves a similar RMSE with $1/n$ -th as many samples (dotted lines), at least until variance decreases enough for bias to become apparent.

Parameter choice Assuming finite time for rendering an image, it may be preferable to choose a fixed mollification distance for all iterations. In Figure 13 we compare shrinking with $s = 1/4$ starting from $d_0 = 100nm$ to keeping $d = 5nm$ fixed from 1 spp to 1024 spp. In the shrinking progression $d_{1024} = d_0 \cdot 1024^{-1/4} \approx 18nm$ does not quite reach $5nm$, yet there is little noticeable difference, except maybe the upper left sphere. This is explained by the fact that this scene is dominated by paths that are easily sampled by NEE or the LT, and paths requiring mollification play only a minor role here.

In Figure 17 mollified paths are more significant, and we test various combinations of d_0 and s . This clearly illustrates the trade-off between initial noise reduction and long term bias reduction. At 32 spp, $d_0 = 100nm$, $s = 1/4$ outperforms $d_0 = 5nm$, $s = 0$. Then, at 1024spp, both combinations perform similarly well. On the other hand, $d_0 = 100nm$, $s = 3/4$, while initially better than $d_0 = 5nm$, $s = 0$, shows the most variance in the long run due to the drastic shrinking ($100 \cdot 1024^{-3/4} = 0.005nm$). The supplemental document contains the same figure with additional sample counts to qualitatively demonstrate that the optimal d_0 and s depend on the number of available samples.

Figure 19 and Figure 20 show quantitatively that the optimal parameters depend on the scene settings, reflectance spectra and sample count.

5. Discussion

Alternatives During the course of this work, we considered a few alternatives to mollification, in the hope of finding an unbiased solution. Brute force tracing a complete set of wavelengths (quantize the visible spectrum in $5nm$ steps, as the CIE color matching functions, for instance) does not help, since a wavelength dependent scattering event will break the path configuration for all but one wavelength. This results in a lot of wasted computation. There is also the other extreme: first sampling the path *completely without* a wavelength and later re-evaluating it for the required wavelengths. This strategy has the same issue with geometric configurations: a good path cannot be sampled without a wavelength.




Figure 11: A scene with fluorescent and grey walls and smooth dielectric and metal spheres. From left to right: Our BDPT, PT+NEE, LT.




Figure 12: A scene with fluorescent and grey walls and rough conductors, illuminated by constant ultraviolet light. From left to right: PT + NEE (256 spp, 81 seconds), LT (256 spp, 85 seconds), our BDPT ($d_0 = 100\text{nm}$, $s = 1/4$, 256spp, 448 seconds), PT + NEE roughly equal time (2048 spp, 625 seconds), PT + NEE 132k spp. Note that the caustic on the right wall is caused by UV light being refracted through the sphere, which is then made visible by the fluorescent blue material.

Combination with Hero Wavelength Sampling The paths in our implementation transport single wavelengths, but combination with techniques that carry more than one wavelength at a fixed stratified offset [WND*14] are straight forward: if mollification between the two main wavelengths is admissible, this also holds for the copies.

6. Limitations and Future Work

The biggest limitation of our work is that it deteriorates the performance of BDPT in the regular case, where no fluorescence happens. This is presumably a large proportion of the scenes currently encountered in practice. The problem is somewhat alleviated by the fact that light tracing and path tracing with next event estimation are not affected by this, only the inner connections require mollification. This could potentially be addressed by a specialized variant of the hero wavelength sampling scheme [WND*14]. Another potential shortcoming is the use of approximate MIS weights, but we could not prove a big impact here.

Evaluating the mollified PDF requires integrating the mollified BBRRDF over a range of wavelengths. For most BBRRDFs, a separable approximation such as the one we used should work well. This may be interesting for rough dielectric transmission, where the direction and the wavelength show strong correlation. Interest-

ingly, mollifying a non-fluorescent BBRRDF and designing new fluorescent BBRRDF models require similar mathematics.

The convergence of kernel estimation methods deteriorates with increasing dimensionality [KD13b]. As we only mollify one dimension, our approach is in this regard more efficient than e.g. photon mapping which requires at least a 3D lookup. However, photon mapping methods leverage multiple cached sample points.

Another interesting direction is to derive optimal initial mollification distance and shrinking rate parameters. Previous work [KD13a] shows that for photon mapping optimal parameters depend on the scene. Our experiments suggest that the situation is similar for the spectral domain. A straight forward extension is to combine our spectral mollification scheme with directional mollification [KD13b] and spatial mollification [GKDS12].

7. Conclusion

We proposed a mollification-based method to combine bidirectional path sampling methods with wavelength changing, i.e. fluorescent materials. The mollification resolves the fundamental contradiction between sampling new wavelengths and connecting path vertices, but increases variance for inner connections at




Figure 13: Orange and blue insets from Figure 12 rendered with our BDPT at 1, 2, 4, ..., 1024 spp. Top: starting with a lot of mollification and decreasing over time, $d_0 = 100\text{nm}$, $s = 1/4$. Bottom: starting small without decreasing, $d_0 = 5\text{nm}$, $s = 0$. Except towards higher sample counts at the top two rows there is little noticeable difference.




Figure 14: A specular UFO casting light on a sponge (see Figure 15) in a medium. Left: LT, 1024 spp, 407 seconds. Center: Our BDPT, 64 spp, 253 sec. Right: PT + NEE, 1024 spp, 391 sec.




Figure 15: Setup for the UFO (Figure 14) and lens (Figures 16 and 17) scene. The lenses are modeled as smooth dielectric flat ellipsoids, in order to create a scene where PT and LT struggle.

non-fluorescing surfaces as compared to regular BDPT. Thus, our method should only be used if fluorescent surfaces are present and if there is a clear gain from such sampling strategies (see Figure 1). Our method relies on MIS, so it will transparently prefer PT or LT in simple cases. As with regular BDPT, additional inner path connections come at an overhead, but with the separable constant mollification we proposed there is very little additional raw performance overhead of our method.

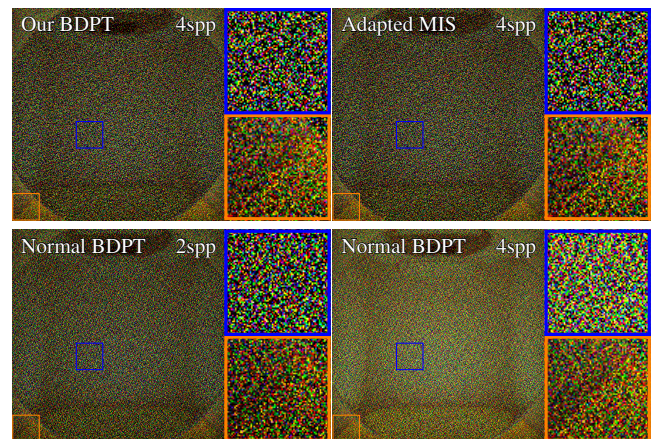


Figure 16: A scene without fluorescence, grey walls and ceiling, and an orange floor. There are perfectly specular lenses in front of the camera and the focused light source (see Figure 15). In this setup NEE and LT are useless. BDPT relies on connecting sub-paths with at least 2 vertices each. From left to right: Our BDPT (4spp, $d_0 = 100\text{nm}$), Our BDPT with adapted MIS weights (4spp, $d_0 = 100\text{nm}$), unmollified BDPT (2spp), unmollified BDPT (4spp).

8. Acknowledgements

This work was funded by DFG grant DA 1200/6-1, project 323377784. The crab, seastar and shell models in the teaser are downloaded from <http://threedscans.com>. The sponge and coral models are part of the digital atlas of ancient life by the paleontological research institution, Ithaca, New York <https://sketchfab.com/DigitalAtlasOfAncientLife>. The can model was created by jackho at <http://www.blendswap.com/blends/view/74793>.




Figure 17: This scene is rendered with 32 (top) and 1024 (bottom) spp and contains no fluorescence, but lenses in front of the camera and focused light source (see Figure 15), making LT and PT+NEE mostly useless. Left is the unmollified BDPT, next to it is the mollified BDPT with different values for d_0 and s . This illustrates the tradeoff between initial noise reduction due to larger d_0 and small s (center, center right) and long-term bias reduction with larger s and smaller d_0 (center left). Note that $s = 0$ is not actually consistent, but with a "good" choice of d_0 may yield better results in limited time. Progression animations of all renders are part of the supplemental material. The grid-like artifacts are noise caused by the Halton point random numbers used to place pixel samples.




Figure 18: RMSE for an empty box scene (see supplemental Figure 8b for setup; color checker green reflectance from supplemental Figure 4) with mollified (2,2)-paths only and no shrinking ($s = 0$). With a 10-times smaller mollification radius d (blue), the variance of mollified techniques increases by a factor of roughly 10. Therefore a similar RMSE is achieved with 10 times as many samples (dotted lines), at least at lower sample counts before bias becomes significant and the lower radius outperforms the larger one (red).




Figure 19: RMSE with same setup as Figure 18, for different d_0 and s . For $d_0 = 100$ we chose $s = 1/4$ which performs best in Figure 20 (right). Here, each combination outperforms the others in a certain wavelength range. At even higher sample counts, the $d \equiv 5$ technique will stop converging as it is biased, and the $d_0 = 100, s = 1/4$ combination will have the lowest RMSE once again.




Figure 20: RMSE with same setup as Figure 18, $d_0 = 100\text{nm}$ and different shrinking parameters s . Initially, no shrinking is best since with larger d more paths are successfully mollified which decreases variance. In the long run, even in this simple setup, the best s depends on the samples per pixel and reflectance spectra involved.




Figure 21: Mollification radii d_n for different shrinking parameters s over samples per pixel, starting from $d_0 = 100\text{nm}$.

References

[ABE*17] ABDELLAH M., BILGILI A., EILEMANN S., MARKRAM H., SCHÜRMANN F.: A physically plausible model for rendering highly scattering fluorescent participating media. *CoRR abs/1706.03024* (2017). URL: <http://arxiv.org/abs/1706.03024>, arXiv: 1706.03024. 3

[BHD*08] BENDIG M., HANIKA J., DAMMERTZ H., GOLDSCHMIDT J. C., PETERS I. M., WEBER M.: Simulation of fluorescent concentrators. pp. 93–98. doi:10.1109/RT.2008.4634628. 3

- [CS04] CEREZO E., SERÓN F. J.: Rendering natural waters taking fluorescence into account. vol. 15. doi:10.1002/cav.10.3
- [DH09] DIMIAN D., HERBST K.: Practical uses of a ray tracer for *Cloudy with a Chance of Meatballs*. In *International Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 2009, New Orleans, Louisiana, USA, August 3-7, 2009, Talks Proceedings* (2009). doi:10.1145/1597990.1598034.2
- [FHF*17] FASCIONE L., HANIKA J., FAJARDO M., CHRISTENSEN P., BURLEY B., GREEN B.: Path tracing in production – part 1: Writing production renderers. In *SIGGRAPH Courses* (2017). doi:10.1145/3084873.3084904.1
- [GKDS12] GEORGIEV I., KRÍVÁNEK J., DAVIDOVIČ T., SLUSALLEK P.: Light transport simulation with vertex connection and merging. *ACM Trans. Graph.* 31, 6 (Nov. 2012), 192:1–192:10. doi:10.1145/2366145.2366211.2,9
- [Gla95] GLASSNER A. S.: A model for fluorescence and phosphorescence. In *Photorealistic Rendering Techniques* (Berlin, Heidelberg, 1995), Sakas G., Müller S., Shirley P., (Eds.), Springer Berlin Heidelberg, pp. 60–70. doi:10.1007/978-3-642-87825-1_5.3
- [GMAS05] GUTIERREZ D., MUNOZ A., ANSON O., SERON F. J.: Non-linear volume photon mapping. In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques* (Aire-la-Ville, Switzerland, Switzerland, 2005), EGSR '05, Eurographics Association, pp. 291–300. doi:10.2312/EGWR/EGSR05/291-300.3
- [GSMA08] GUTIERREZ D., SERON F. J., MUNOZ A., ANSON O.: Visualizing underwater ocean optics. *Computer Graphics Forum* 27, 2 (2008), 547–556. doi:10.1111/j.1467-8659.2008.01152.x.3
- [HDF15] HANIKA J., DROSKE M., FASCIONE L.: Manifold next event estimation. *Computer Graphics Forum (Proceedings of Eurographics Symposium on Rendering)* 34, 4 (June 2015), 87–97. doi:10.1111/cgf.12681.2
- [HHA*10] HULLIN M. B., HANIKA J., AJDIN B., SEIDEL H.-P., KAUTZ J., LENSCH H. P. A.: Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions. *ACM Trans. Graph.* 29, 4 (July 2010), 97:1–97:7. doi:10.1145/1778765.1778834.2,3
- [HPJ12] HACHISUKA T., PANTALEONI J., JENSEN H. W.: A path space extension for robust light transport simulation. *ACM Trans. Graph.* 31, 6 (Nov. 2012), 191:1–191:10. doi:10.1145/2366145.2366210.2
- [JA18] JARABO A., ARELLANO V.: Bidirectional rendering of vector light transport. *Computer Graphics Forum* 37, 6 (2018), 96–105. doi:10.1111/cgf.13314.2,3
- [JG19] JENDERSIE J., GROSCH T.: Microfacet Model Regularization for Robust Light Transport. *Computer Graphics Forum* (2019). doi:10.1111/cgf.13768.2
- [JHD18] JUNG A., HANIKA J., DACHSBACHER C.: A simple diffuse fluorescent BBRRDF model. In *MAM2018: Eurographics Workshop on Material Appearance Modeling* (2018). doi:10.2312/mam.20181193.3,4
- [JMMn*14] JARABO A., MARCO J., MUÑOZ A., BUISAN R., JAROSZ W., GUTIERREZ D.: A framework for transient rendering. *ACM Transactions on Graphics (SIGGRAPH Asia 2014)* 33, 6 (2014). doi:10/gfznb8.2
- [JWH*19] JUNG A., WILKIE A., HANIKA J., JAKOB W., DACHSBACHER C.: Wide Gamut Spectral Upsampling with Fluorescence. *Computer Graphics Forum* (2019). doi:10.1111/cgf.13773.1
- [Kaj86] KAJIYA J. T.: The rendering equation. In *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1986), SIGGRAPH '86, ACM, pp. 143–150. doi:10.1145/15922.15902.2
- [KD13a] KAPLANYAN A. S., DACHSBACHER C.: Adaptive progressive photon mapping. *ACM Transactions on Graphics* 32, 2 (Apr. 2013), 16:1–16:13. doi:10.1145/2451236.2451242.2,9
- [KD13b] KAPLANYAN A. S., DACHSBACHER C.: Path space regularization for holistic and robust light transport. *Computer Graphics Forum (Proc. of Eurographics 2013)* 32, 2 (2013), 63–72. doi:10.1111/cgf.12026.2,5,6,9
- [KFF*15] KELLER A., FASCIONE L., FAJARDO M., CHRISTENSEN P., HANIKA J., EISENACHER C., NICHOLS G.: The path-tracing revolution in the movie industry. In *SIGGRAPH Courses* (2015). doi:10.1145/2776880.2792699.1
- [KZ11] KNAUS C., ZWICKER M.: Progressive photon mapping: A probabilistic approach. *ACM Trans. Graph.* 30, 3 (May 2011), 25:1–25:13. doi:10.1145/1966394.1966404.2
- [Lak10] LAKOWICZ J.: Principles of fluorescence spectroscopy. *Plenum Press 441* (01 2010). doi:10.1007/978-0-387-46312-4.2,3
- [May19] MAYR T.: Fluorophores.org database of fluorescent dyes, properties and applications, 2019. Accessed: 2019/08/02. URL: fluorophores.tugraz.at.3
- [MW18] MOJZIK M., WILKIE A. F. A.: Handling fluorescence in a unidirectional spectral path tracer. *Computer Graphics Forum 37(4) (Proceedings of Eurographics Symposium on Rendering 201)* 37, 4 (2018). doi:10.1111/cgf.13477.3
- [NSR17] NALBACH O., SEIDEL H.-P., RITSCHEL T.: Practical capture and reproduction of phosphorescent appearance. *Computer Graphics Forum (Proc. Eurographics 2017)* 36, 2 (2017). doi:10.1111/cgf.13136.3
- [PJH16] PHARR M., JAKOB W., HUMPHREYS G.: *Physically Based Rendering: From Theory to Implementation (3rd ed.)*, 3rd ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, Oct. 2016. doi:10.1016/C2013-0-15557-2.1,2
- [SBCD14] SUO J., BIAN L., CHEN F., DAI Q.: Bispectral coding: compressive and high-quality acquisition of fluorescence and reflectance. *Opt. Express* 22, 2 (Jan 2014), 1697–1712. doi:10.1364/OE.22.001697.3
- [The19] THERMOFISHER: Thermofisher spectra viewer, 2019. Accessed: 2019/08/02. URL: <https://www.thermofisher.com/order/spectra-viewer>.3
- [THH18] TOMINAGA S., HIRAI K., HORIUCHI T.: Estimation of fluorescent donaldson matrices using a spectral imaging system. *Opt. Express* 26, 2 (Jan 2018), 2132–2148. doi:10.1364/OE.26.002132.3
- [WEV02] WARD G., EYDELBERG-VILESHIN E.: Picture Perfect RGB Rendering Using Spectral Prefiltering and Sharp Color Primaries. In *Eurographics Workshop on Rendering* (2002), The Eurographics Association, pp. 117–124. doi:10.2312/EGWR/EGWR02/117-124.1
- [WKB12] WALTER B., KHUNGURN P., BALA K.: Bidirectional lightcuts. *ACM TOG* 31, 4 (July 2012), 59:1–59:11. doi:10.1145/2185520.2185555.2
- [WND*14] WILKIE A., NAWAZ S., DROSKE M., WEIDLICH A., HANIKA J.: Hero wavelength spectral sampling. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)* 33, 4 (July 2014), 123–131. doi:10.1111/cgf.12419.9
- [WTP01] WILKIE A., TOBLER R. F., PURGATHOFER W.: Combined rendering of polarization and fluorescence effects. In *Rendering Techniques 2001* (Vienna, 2001), Gortler S. J., Myszkowski K., (Eds.), Springer Vienna, pp. 197–204. doi:10.2312/EGWR/EGWR01/197-204.2,3
- [WWLP06] WILKIE A., WEIDLICH A., LARBOULETTE C., PURGATHOFER W.: A reflectance model for diffuse fluorescent surfaces. In *Proceedings of Graphite 2006* (Nov. 2006), p. 8. doi:10.1145/1174429.1174484.3