EUROGRAPHICS 2020 / U. Assarsson and D. Panozzo
(Guest Editors)

Volume 39 (2020), Number 2

Supplemental - Spectral Mollification for Bidirectional Fluorescence

A. Jung[1], J. Hanika[1] and C. Dachsbacher[1]

[1]Karlsruhe Institute of Technology

1. Energy Conservation

In this section we show that r_{mi}

$$r_{mi}(\lambda_i \to \lambda_o) = \frac{1}{2d} r(\lambda_o) \mathbb{1}_{|\lambda_i - \lambda_o| < d} \tag{1}$$

(mollification over incident wavelengths as done in equation 9 in the paper) is not necessarily energy conserving when applied to the diagonal of a physically valid reradiation matrix, or to the full non-fluorescent term

$$\delta_{\lambda_i,\lambda_o}(1 - ca(\lambda_i))r(\lambda_i) \tag{2}$$

of the BBRRDF by [JHD18].

For example, consider

$$c = Q = 1, \lambda_i = 600, d = 50 \tag{3}$$

and the extreme case of a box function for an absorption spectrum

$$a(\lambda_i) = \begin{cases} 1 & \text{if } \lambda_i \in [500, 600] \\ 0 & \text{else} \end{cases} \tag{4}$$

and a reflectance spectrum

$$r(\lambda) = 1 - a(\lambda). \tag{5}$$

Then

$$\begin{aligned}
&\int_\Lambda r_{mi}(\omega_i, \lambda_i, \omega_o, \lambda_o) d\lambda_o \\
&= \int_\Lambda \mathbb{1}_{|\lambda_i - \lambda_o| < d} \cdot \frac{(1 - ca(\lambda_o))r(\lambda_o)}{2d} + ca(\lambda_i) Qe(\lambda_o) d\lambda_o \\
&= \frac{1}{2d} \int_{\lambda_i - d}^{\lambda_i + d} (1 - ca(\lambda_o))r(\lambda_o) d\lambda_o + 1 \\
&= \frac{1}{2d} \int_{\lambda_i - d}^{\lambda_i + d} r(\lambda_o) d\lambda_o - \frac{c}{2d} \int_{\lambda_i - d}^{\lambda_i + d} a(\lambda_o) r(\lambda_o) d\lambda_o + 1 \\
&= \frac{1}{2d} \int_{\lambda_i - d}^{\lambda_i + d} r(\lambda_o) d\lambda_o + 1 \\
&= \frac{1}{2d} \cdot d + 1 \\
&= 1 + \frac{d}{2} \\
&> 1.
\end{aligned} \tag{6}$$

Even though truly box-like spectra do not appear in the real world, this serves to show that in general, whenever the absorption and emission spectrum peak in adjacent areas, energy conservation might be an issue when using f_{mi}. The effect is visualized in Figure 2 with a reradiation matrix constructed from measured reflectance and fluorescence spectra. Figure 3 shows what happens for constant mollification with f_{mo} instead.

Another way to think about mollification in the context of reradiation matrices is to pick a given λ_i and interpret f_{mi} as an additional fluorescent component with $a(\lambda) = \mathbb{1}_{\lambda = \lambda_i}$ and $e(\lambda) = \frac{\mathbb{1}_{|\lambda - \lambda_i| < d}}{2d}$. Mollifying with f_{mi} would then correspond to adding the reradiation matrix of the original fluorescent component to the reradiation matrix of the mollified component. Depending on how extreme the original matrix was, this can now result in columns (corresponding to an incident wavelength) summing up to more than 1, thus breaking energy conservation.

For the same reason, box filtering the full non-fluorescent component (equation 2) may lead to a non-energy conserving reradiation matrix.

Mollifying r only If we only mollify r in the analytical BBRRDF we have no issues. That is because the unmollified $1 - ca$ regulates how much light is available for reflection in the first place. Consider arbitrary $a, r, Q, c \in [0, 1]$ and $e \geq 0$ with $\int e = 1$:

$$\begin{aligned}
&\int_\Lambda \int_\Omega f(\omega_i, \lambda_i, \omega_o, \lambda_o) d\omega_o d\lambda_o \\
&= \int_\Lambda \mathbb{1}_{|\lambda_i - \lambda_o| < d} \cdot \frac{(1 - ca(\lambda_i))r(\lambda_o)}{2d} + ca(\lambda_i) Qe(\lambda_o) d\lambda_o \\
&= \frac{1 - ca(\lambda_i)}{2d} \int_{\lambda_i - d}^{\lambda_i + d} r(\lambda_o) d\lambda_o + ca(\lambda_i) Q \\
&\leq \frac{1 - ca(\lambda_i)}{2d} \cdot 2d + ca(\lambda_i) \\
&= 1
\end{aligned} \tag{7}$$

2. Photon conservation

2.1. Symmetric kernels

If an elastic BBRRDF is mollified with any symmetric kernel $k(\lambda_o | \lambda_i)$ centered around $r(\lambda_i)$, i.e.

$$f_{moll}(\lambda_i, \lambda_o) = r(\lambda_i) \cdot k(\lambda_o | \lambda_i) \cdot \ldots, \tag{8}$$

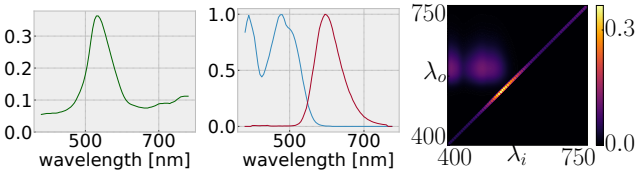

Figure 1: *A reflectance spectrum (left) and fluorescence absorption (blue) and emission (red) spectrum (center), and the corresponding reradiation matrix (right) based on the model by [JHD18].*

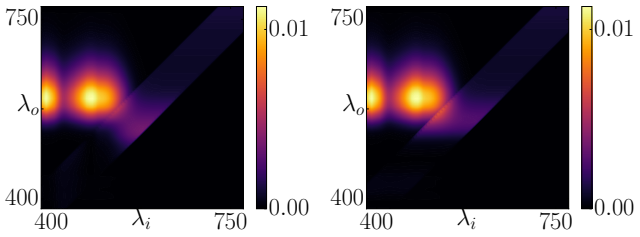

Figure 2: *Based on the reradiation matrix in Figure 1. Left: $r(\lambda)$ mollified with f_{mi}. Right: $(1-ca(\lambda))r(\lambda)$ (i.e., the matrix diagonal) mollified with f_{mi}. The former fulfills energy conservation, but is not directly applicable to the reradiation matrix itself. The latter leaks energy to shorter wavelengths, violating energy conservation.*

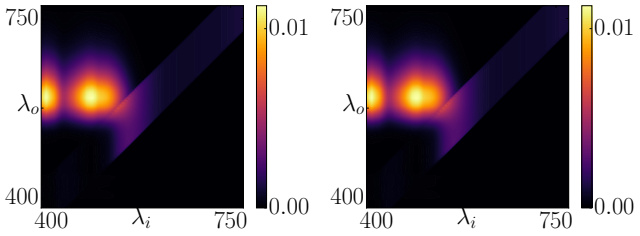

Figure 3: *Based on the reradiation matrix in Figure 1. Left: $r(\lambda)$ mollified with f_{mo}. Right: $(1-ca(\lambda))r(\lambda)$, in other words, the matrix diagonal, mollified with f_{mo}. Both are the same and preserve energy conservation.*

the dots representing an optional geometrical part of the BBRRDF which does not depend on wavelengths, and

$$k(\lambda_i + x|\lambda_i) = k(\lambda_i - x|\lambda_i), \tag{9}$$

f_{moll} is not only energy conserving but also photon conserving.

Assume a surface is illuminated with one unit energy at wavelength λ_i. The distribution of spectral irradiance over outgoing wavelengths λ_o is described by the mollified BBRRDF integrated over the hemisphere:

$$\begin{aligned} E(\lambda_o) &= \int_\Omega f_{moll}(\lambda_i, \lambda_o) d\omega_o^\perp \\ &= r(\lambda_i) \cdot k(\lambda_o|\lambda_i) \end{aligned} \tag{10}$$

The energy of a single photon is proportional to its wavelength:

$$E = \frac{h}{\lambda}, \tag{11}$$

h being Planck's constant. We can therefore express the distribution of outgoing photons for one unit of incident energy at λ_i as

$$P_o(\lambda_o) := r(\lambda_i) k(\lambda_o|\lambda_i) \cdot \frac{\lambda_o}{h}. \tag{12}$$

Integrating this yields the total number of outgoing photons per unit incident energy at λ_i:

$$\begin{aligned} N_{out}(\lambda_i) &= \int_\Lambda r(\lambda_i) k(\lambda_o|\lambda_i) \frac{\lambda_o}{h} d\lambda_o \\ &= \frac{r(\lambda_i)}{h} \int_{-\infty}^{\infty} k(\lambda_i + x|\lambda_i) \cdot (\lambda_i + x) dx \\ &= \frac{r(\lambda_i)}{h} \left(\lambda_i \int_{-\infty}^{\infty} k(\lambda_i + x|\lambda_i) dx + \int_{-\infty}^{\infty} k(\lambda_i + x|\lambda_i) x \;\; dx \right) \\ &= \frac{r(\lambda_i)}{h} (\lambda_i + \int_{-\infty}^{0} k(\lambda_i + x|\lambda_i) x \;\; dx + \int_{0}^{\infty} k(\lambda_i + x|\lambda_i) x \;\; dx \\ &= \frac{r(\lambda_i)}{h} (\lambda_i + \int_{0}^{\infty} -k(\lambda_i - x|\lambda_i) x \;\; dx + \int_{0}^{\infty} k(\lambda_i + x|\lambda_i) x \;\; dx \\ &= r(\lambda_i) \frac{\lambda_i}{h} \end{aligned} \tag{13}$$

On the other hand, the number of photons corresponding to one unit of incident energy at λ_i can be computed according to equation 10 as

$$N_{in} = \frac{\lambda_i}{h}. \tag{14}$$

This means that if $r(\lambda_i) \in [0,1]$, f_{moll} does not produce additional photons, and if $r(\lambda_i) = 1$, the number of outgoing photons is equal to the number of incident photons.

2.2. Asymmetric kernels

Asymmetric kernels do not necessarily result in photon conserving BBRRDFs. For example, the kernel

$$k(\lambda_o|\lambda_i) = \mathbb{1}_{0 \le \lambda_o - \lambda_i < d} \tag{15}$$

for some mollification distance $d > 0$ only mollifies from shorter to longer wavelengths. While this might be more akin to the typical behaviour of fluorescent materials, consider this kernel applied to equation 8. If $r(\lambda_i)$ is large enough, i.e. enough incident energy is reflected, the number of photons required to carry the reflected energy at longer wavelengths is larger than the number of incident photons with wavelength λ_i.

3. Mollified pdfs

In this section we present the full mollified pdfs for fluorescent BBRRDFs. We start with the model by [JHD18] and provide details on the different mollification options discussed in the paper.

3.1. Constant mollification

We start with constant mollification where we always evaluate r for the incident wavelength λ_i. In this case, mollifying r only and mollifying the full non-fluorescent component $(1-ca)r$ results in the same mollified BBRRDF

$$f_{mo}(\lambda_i \to \lambda_o) = \frac{1}{\pi} \cdot \Big(\frac{\mathbb{1}_{|\lambda_i-\lambda_o|<d} r(\lambda_i)}{2d} \cdot (1-ca(\lambda_i)) + ca(\lambda_i) Qe(\lambda_o) \Big). \tag{16}$$

The corresponding mollified pdf in camera direction is

$$p_c(\lambda_i|\lambda_o) = \frac{f_{mo}(\lambda_i \to \lambda_o)}{\int_\Lambda f(\lambda \to \lambda_o) d\lambda} \tag{17}$$

with

$$\int_\Lambda f(\lambda \to \lambda_o) d\lambda = \frac{1}{2d} \int_{\lambda_o-d}^{\lambda_o+d} r(\lambda)(1-ca(\lambda)) d\lambda + cQe(\lambda_o) \int_\Lambda a(\lambda) d\lambda, \tag{18}$$

which can be precomputed.

Similarly, in the direction of a light path the mollified pdf is

$$p_l(\lambda_o|\lambda_i) = \frac{f_{mo}(\lambda_i \to \lambda_o)}{\int_\Lambda f(\lambda_i \to \lambda) d\lambda} \tag{19}$$

with a slightly simpler

$$\begin{aligned} \int_\Lambda f(\lambda_i \to \lambda) d\lambda &= \frac{1}{2d} \int_{\lambda_o-d}^{\lambda_o+d} r(\lambda_i)(1-ca(\lambda_i)) d\lambda \\ &\quad + ca(\lambda_i) Q \int_\Lambda e(\lambda) d\lambda \\ &= r(\lambda_i)(1-ca(\lambda_i)) + ca(\lambda_i) Q. \end{aligned} \tag{20}$$

The main issue with this technique is that we need to evaluate a piecewise integral of $(1-ca)r$ (or the diagonal of a reradiation matrix). If using measured data representing piecewise linear spectra this can be done in constant time by precomputing a piecewise cdf. With reradiation matrices, we can interpret the diagonal as discrete spectral data, which is simple to integrate. For spectral upsampling models this integration depends on the model for reflectance and absorption. For instance, the model by [PMHD19] can be integrated analytically, but the model by [JH19] cannot, in which case an approximate solution is required.

Sampling Note that these pdfs correspond to sampling from the mollified BBRRDF by first sampling whether there is a fluorescent interaction on a camera path

$$\mathbb{P}_c(\lambda_o) = \frac{\int_\Lambda ca(\lambda_i) Qe(\lambda_o) d\lambda_i}{\pi \int_\Lambda f(\lambda_i \to \lambda_o) d\lambda_i} \tag{21}$$

or on a light path

$$\mathbb{P}_l(\lambda_i) = \frac{\int_\Lambda ca(\lambda_i) Qe(\lambda_o) d\lambda_o}{\pi \int_\Lambda f(\lambda_i \to \lambda_o) d\lambda_o}, \tag{22}$$

and then sampling the new wavelength, in case of a fluorescent event, from

$$p_c^{fl}(\lambda_i|\lambda_o) = \frac{a(\lambda_i)}{\int_\Lambda a(\lambda) d\lambda} \tag{23}$$

on a camera path and from

$$p_l^{fl}(\lambda_o|\lambda_i) = e(\lambda_o) \tag{24}$$

on a light path, or, in case of a non-fluorescent event, from the mollified non-fluorescent component using

$$p_c^{moll}(\lambda_i|\lambda_o) = \frac{r(\lambda_i)(1-ca(\lambda_i))}{\int_{\lambda_o-d}^{\lambda_o+d} r(\lambda)(1-ca(\lambda)) d\lambda} \tag{25}$$

on a camera path and using

$$p_l^{moll}(\lambda_o|\lambda_i) = \frac{\mathbb{1}_{|\lambda_i-\lambda_o|<d}}{2d} \tag{26}$$

on a light path. Combining these yields the full pdf

$$p_x(\lambda_a|\lambda_b) = \mathbb{P}_x(\lambda_b) * p_x^{fl}(\lambda_a|\lambda_b) + (1-\mathbb{P}_x(\lambda_b)) p_x^{moll}(\lambda_a|\lambda_b) \tag{27}$$

for $x = c, l$, which equals equations 17 and 19. Also note that for a non-fluorescent BBRRDF this simplifies to $p_x(\lambda_a|\lambda_b) = p_x^{moll}(\lambda_a|\lambda_b)$.

3.2. Box mollification

When mollifying the reflectance component using a 2D box filter, integrating f_{mb} for the pdf evaluation requires evaluating a double integral over the non-fluorescent component $(1-ca)r$. For measured spectra this can still be computed in constant time, based on two precomputed cdfs for $(1-ca)r$ and $\int_{-\infty}^{\lambda} r(\lambda)(1-ca(\lambda) d\lambda$.

3.3. Reradiation Matrices

For reradiation matrices M with $m_{i,j} = f(\lambda_i \to \lambda_j)$ we get

$$p_c(\lambda_i|\lambda_o) = \frac{m_{i,o}}{\sum_k m_{k,o}}. \tag{28}$$

This means we can get a mollified version of the pdf by using values from the mollified reradiation matrix.

4. Test scenes

In the following we present the reflectance (Figure 4) and fluorescence spectra (Figures 5 and 6) used in the main paper's test scenes. As we use [JHD18]'s model for fluorescent BBRRDFs we show the spectra directly instead of the corresponding reradiation matrix.

4.1. Scene with spheres

Paper Figures 11, 12 and 13 show scenes with dielectric and conductor spheres. The ceiling and left wall are diffuse and non-fluorescent with the grey reflectance from Figure 4, the remaining planes are fluorescent with spectra described in Figure 5.

4.2. UFO scene

The scene in Figure 14 in the paper is filled with a homogeneous medium with wavelength-dependent mean free path and scattering coefficients depicted in Figure 6 (left). The spectra of the fluorescent coral in the center are shown in Figure 6 (right). The top of the UFO is a smooth dielectric with wavelength-dependent index

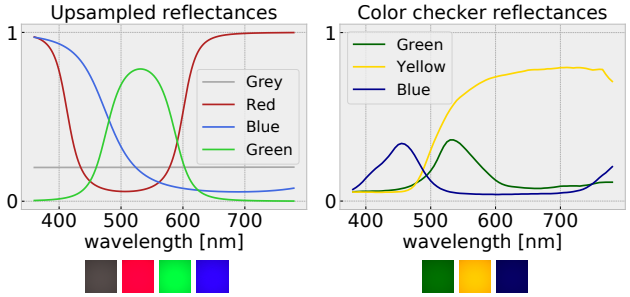

Figure 4: *Reflectance spectra used in color bias tests. The first four spectra were upsampled from RGB=[0.2,0.2,0.2] (grey), [0.8,0.05,0.15] (red), [0.05,0.8,0.15] (green) and [0.05,0.15,0.8] (blue) using the method by [JH19]. The last three are color checker measurements from [MMP*].*

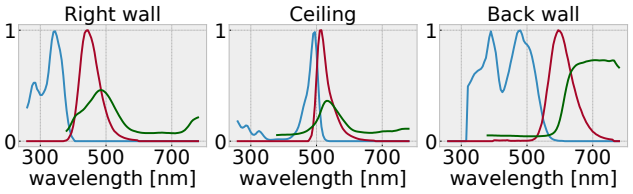

Figure 5: *Reflectance (geen) and fluorescence absorption (blue) and emission (red) spectra in the spheres scene. The fluorescence masurements are from [May19], the spectra are color checker measurements for cyan, green and red from [MMP*].*

of refraction (IOR), with abbe number (dispersive power) 10 and $n_D = 1.3$ [JW57]. The UFO body is a smooth conductor based on the complex IOR of silver. The left and right wall are the same material as the right wall in Figure 5. The floor, back wall and ceiling are constant grey as in Figure 4.

4.3. Lens scene

The lenses are smooth dielectrics with abbe number 23 and $n_D = 1.3$. The floor is orange, with a spectrum upsampled [JH19] from $RGB = [0.9, 0.7, 0.1]$. The light source emits a constant spectrum. The remaining walls are diffuse grey as in Figure 4.

Varying mollification distance and shrinking Figure 7 shows the same setup as Figure 17 in the paper, but for more spp counts. The individual images as well as progression animations are included in the supplemental files. At only 4 samples per pixel, $d_0 = 100nm$ without any shrinking turns out to be a viable strategy for generating quick previews, especially if one of the diffuse surfaces were to be replaced by a fluorescent BBRRDF so that the unmollified BDPT could not render the scene anymore.

4.4. Color bias tests

We evaluate the color bias for different mollification radii d and reflectances r (see Figure 4), both for a centered kernel $\mathbb{1}_{|\lambda_i-\lambda_o|<d}/2d$

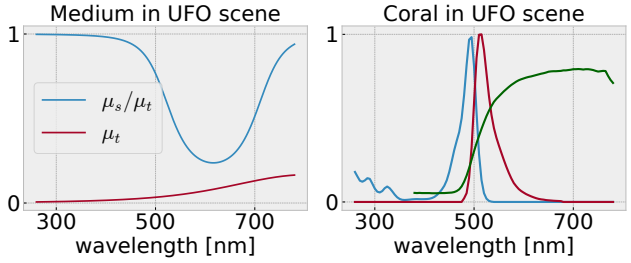

Figure 6: *Left: Volume scattering albedo (red) and μ_t (blue) in the UFO scene. Right: Reflectance (green) [MMP*] and fluorescence absorption (blue) and emission (red) spectrum [May19] of the coral in the UFO scene.*

and an asymmetric kernel $\mathbb{1}_{\lambda_i<\lambda_o<\lambda_i+2d}/2d$. We use an empty cornell box illuminated with a constant light source, though the images only show the center of the back wall (see Figure 8).

We analyze the boundary bias for the asymmetric kernel in Figure 9 for a completely grey box.

In Figures 5, 6, and 7 in the paper all walls, floor and ceiling have the same reflectance. Here, in Figures 10 and 11 only the back wall has a colorful reflectance, the remaining walls are constant grey, such that mollification for $(2,2)$-paths only happens for constant spectra. Figures 7 and 9 in the paper already predict that this has a negligible effect.

In Figures 12 and 13 only the back wall is grey, with the other walls, floor and ceiling having the same colorful reflectance. In this case mollification happens on a colorful surface, and the resulting bias is not further distorted and thereby weakened by a colorful back wall, as was the case in the paper.

5. Guide to supplemental files

5.1. Convergence gifs

The gifs illustrate convergence behaviour by displaying snapshots at 1, 2, 4, ... samples per pixel. The animations have to be taken with a grain of salt due to gif compression artefacts.

5.2. Color bias notebook

color_bias_plots.ipynb is a jupyter python notebook [pyn] to generate various plots on theoretical color bias due to energy leaking out the visible range. It shows expected acceptance rates given a camera or light path wavelength, and overall expected acceptance rate depending on minimal (probably ultraviolet) wavelength λ_f emitted by the light source and mollification distance d.

References

[JH19] JAKOB W., HANIKA J.: A low-dimensional function space for efficient spectral upsampling. *Computer Graphics Forum (Proceedings of Eurographics) 38*, 2 (Mar. 2019). 3, 4

[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. 1, 2, 3

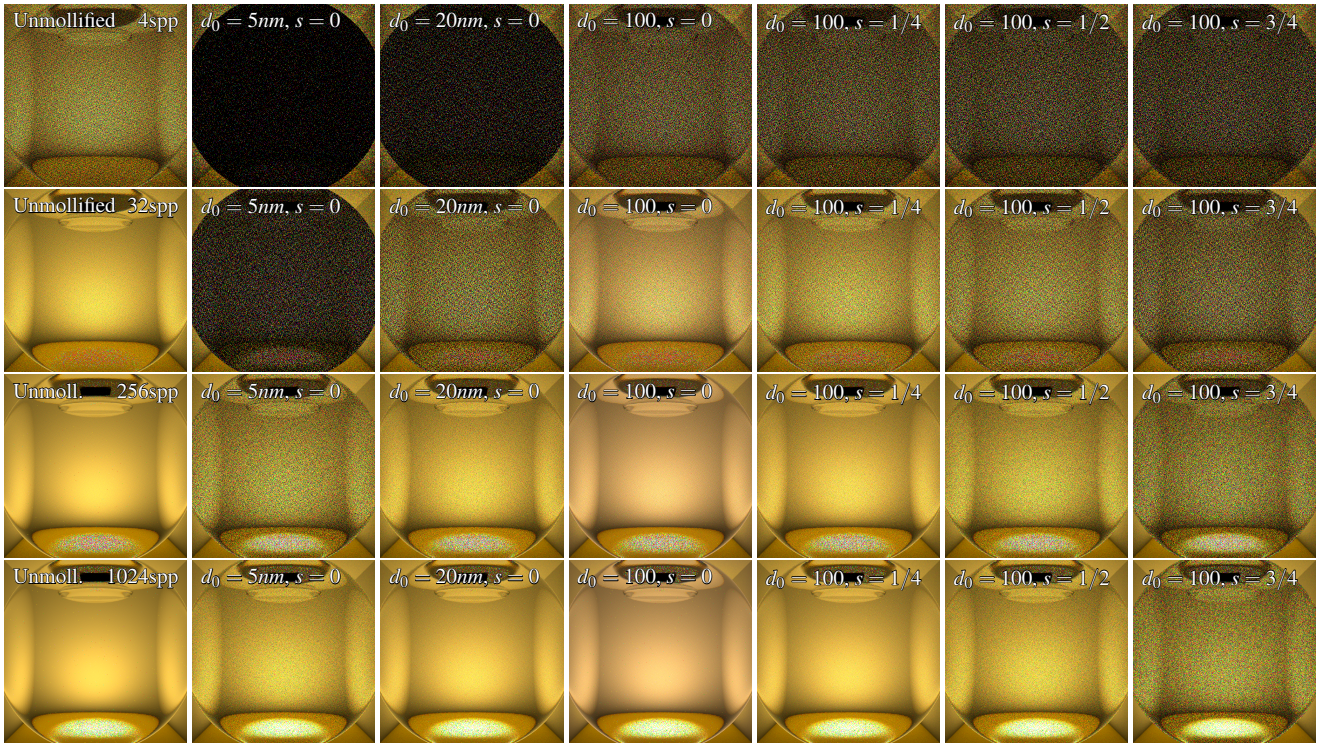

Figure 7: *This scene is rendered with 4, 32, 256 and 1024 (top to bottom) spp and contains no fluorescence, but lenses in front of the camera and focused light source (see Figure 17 in the paper), making LT and PT+NEE mostly useless. Left is the unmollified BDPT, next to it is the mollified BDPT with different values for d_0 and s. This illustrates the tradeoff between initial noise reduction due to higher 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.*

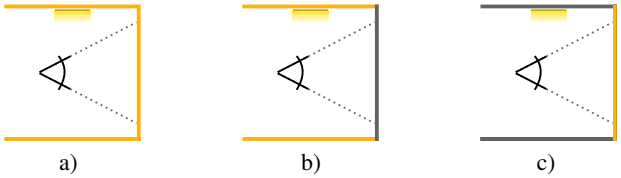

Figure 8: *Scene setups for the color bias tests, with all walls with the same reflectance (a), color reflectances on all walls except a grey back wall (b) and the colorful reflectance only on the back wall, all others being constant grey (c). The reflectance spectra used to color the walls are shown in Figure 4. (a) corresponds to the evaluation in the paper, (b) is used here in Figures 12 and 13, and (c) in Figures 10 and 11.*

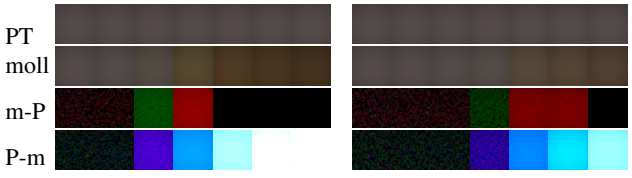

Figure 9: *Boundary bias for the asymmetric kernel ($\mathbb{1}_{\lambda_i<\lambda_o<\lambda_i+2d}/2d$), setup as in Figure 7 in the main paper.*

[JW57] JENKINS F. A., WHITE H. E.: *Fundamentals of Optics*, 3 ed. McGraw-Hill, 1957. URL: http://openlibrary.org/books/OL9259295M. 4

[May19] MAYR T.: Fluorophores.org database of fluorescent dyes, properties and applications, 2019. Accessed: 2019/08/02. URL: fluorophores.tugraz.at. 4

[MMP*] MANSENCAL T., MAUDERER M., PARSONS M., SHAW N., WHEATLEY K., COOPER S., VANDENBERG J. D., CANAVAN L., CROWSON K., LEV O., LEINWEBER K., SHARMA S., SOBOTKA T. J., MORITZ D., PPPP M., RANE C., ESWARAMOORTHY P., MERTIC J., PEARLSTINE B., LEONHARDT M.: Colour (python package). URL: https://www.colour-science.org, doi:10.5281/zenodo.3519603. 4

[PMHD19] PETERS C., MERZBACH S., HANIKA J., DACHSBACHER C.: Using moments to represent bounded signals for spectral rendering. *ACM Transactions on Graphics 38*, 4 (2019). URL: https://cg.ivd.kit.edu/compact_spectra.php. 3

[pyn] Jupyter notebook install guide. URL: https://jupyter.org/install.html. 4

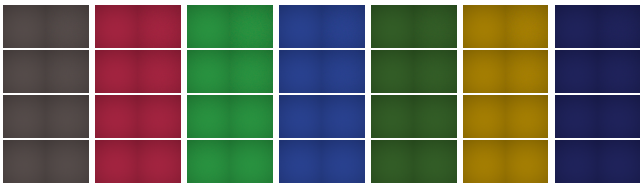

Figure 10: *Color bias evaluation for $(2,2)$-paths for constant d (i.e., $s=0$). Left is the path traced reference for paths of length 4, right shows the mollified $(2,2)$-paths with (from top to bottom) $d=10,20,50,100$. The image shows the colored (see Figure 4 for reflectance spectra) back wall of an otherwise constant grey cornell box (Figure 8c). This means that mollification always happens on constant grey surfaces, where boundary bias is the only issue.*

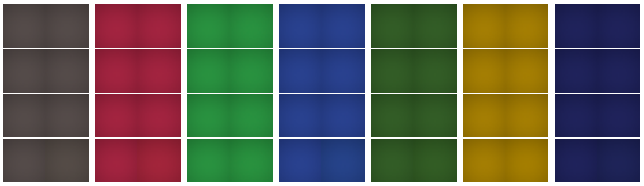

Figure 11: *Same setup and kernel widths ($d=10,20,50,100$) as Figure 10, but with the asymmetric kernel $(1/2d)\cdot\mathbb{1}_{\lambda_i<\lambda_o<\lambda_i+2d}$.*

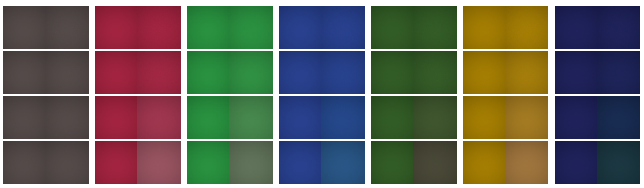

Figure 12: *Color bias evaluation for $(2,2)$-paths for constant d (i.e., $s=0$). Left is the path traced reference for paths of length 4, right shows the mollified $(2,2)$-paths with (from top to bottom) $d=10,20,50,100$. The image shows the grey backwall of an otherwise uniformly colored (see Figure 4 for reflectance spectra) cornell box (see Figure 8b). Thus mollification always happens on colorful reflectance spectra, resulting in blurring and boundary bias.*

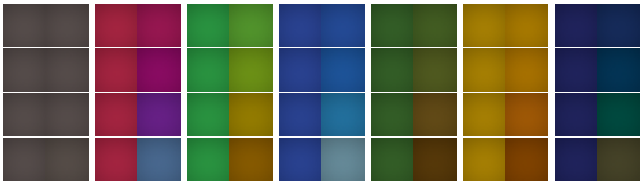

Figure 13: *Same setup and kernel widths ($d=10,20,50,100$) as Figure 12, but with the asymmetric kernel $(1/2d)\cdot\mathbb{1}_{\lambda_i<\lambda_o<\lambda_i+2d}$.*