

Non-Orthogonal Reduction for Rendering Fluorescent Materials in Non-Spectral Engines





A. Fichet^{1,*} , L. Belcour^{1,*} , and P. Barla² ¹Intel Corporation, France, ²Inria, France

Figure 1: Fluorescent effects in a tristimulus render. (a) We introduce a method to render fluorescent materials in a non-spectral renderer, exhibiting typical increased luminance and hue-shifting effects when the same materials are placed under a different lighting. (b) Our method (left insets) is based on a new reduction technique that accurately matches the spectral reference (center insets), significantly improving on previous work [HHA*10] (right insets) where a naive formulation of the fluorescence reradiation matrix reduction was used.

Abstract

We propose a method to accurately handle fluorescence in a non-spectral (e.g., tristimulus) rendering engine, showcasing color-shifting and increased luminance effects. Core to our method is a principled reduction technique that encodes the reradiation into a low-dimensional matrix working in the space of the renderer's Color Matching Functions (CMFs). Our process is independent of a specific CMF set and allows for the addition of a non-visible ultraviolet band during light transport. Our representation visually matches full spectral light transport for measured fluorescent materials even for challenging illuminants.

1. Introduction

Fluorescent materials absorb light and re-emit it at longer wavelengths. This effect has two significant implications in Computer Graphics. First, fluorescence introduces color shifts: after an interaction with a fluorescent material, the light hue changes. Second, the absorption/re-emission process may transfer light from the ultraviolet (UV) range toward the visual range. This has the effect of increasing the perceived luminance of fluorescent materials, compared to other surrounding materials that only reflect light in the visual range. Fluorescent effects are especially noticeable during dawn or dusk, which is due to Rayleigh scattering: sun light (depleted in UV) is less pronounced or absent, making light scattered from the sky (rich in UV) relatively more prominent.

Fluorescence thus has a significant potential to increase artistic control over a scene appearance [JWH*19]. However, the few rendering engines that accurately support fluorescence (e.g., ART [Tob*18], Ocean [Dig22]) require spectral light transport, and assume this is a prerequisite for reproducing wavelength-shifting

effects. In contrast, traditional tristimulus light transport employs component-wise vector multiplications that do not readily model hue-shifting. Fluorescence is instead merely mimicked in such rendering engines, most often through the addition of an emission term or the use of a higher albedo. Not only may this break energy conservation, but it also neglects the impact of reflections from the surrounding environment on fluorescent appearance, and it entirely neglects light outside the visible spectral range. Even though artists may manually adjust material parameters to visually reproduce a fluorescent appearance, this would be restricted to a specific lighting, forbidding light interaction such as shown in Figure 1 (a).

Our goal in this paper is to add support for the accurate reproduction of fluorescence in non-spectral (e.g., tristimulus) renderers, with a focus on Lambertian materials. The central idea consists in replacing the classic vector-vector multiplication between the incoming radiance and diffuse albedo, with a matrix-vector multiplication where the matrix is obtained by a reduction of the fluorescence reradiation matrix. A reduction approach has been suggested

by Hullin *et al.* [HHA*10], but it is based on an empirical argument and produces inaccurate results (see Figure 1 (b), right insets). This method injects unwanted energy since color matching functions (e.g., RGB or XYZ) are non-orthogonal. Our main contribution lies in a principled reduction method, which relies on spectral upsampling and downsampling operators according to an arbitrary set of non-orthogonal functions and produces significantly more accurate results (see Figure 1 (b), left insets). Since our method is generalizable to any basis, we show that an additional UV band is easily incorporated in our framework to account for reradiation from non-visible to visible light.

2. Related work

Fluorescence rendering. With fluorescence, the reflected radiance of a material is no longer a mere function of the input wavelength λ_i , but also accounts for a reradiated wavelength λ_o [Gla95]:

$$L_o(\omega_o, \lambda_o) = \iint \rho(\omega_i, \omega_o, \lambda_i, \lambda_o) L_i(\omega_i, \lambda_i) \langle \omega_i, \omega_n \rangle d\omega_i d\lambda_i, \quad (1)$$

where $\rho(\omega_i, \omega_o, \lambda_i, \lambda_o)$ is the Bidirectional Reflectance and Reradiation Distribution Functions (BRDF), L_o (resp. L_i) the outgoing (resp. incoming) radiance with ω_o (resp. ω_i) the outgoing (resp. incoming) direction, and ω_n the shading normal. Due to this formulation, fluorescence is restricted to spectral renderers, such as uni-directional path tracers [WTP01; BHD*08; MFW18], and bi-directional path tracers thanks to wavelength mollification [JHD20]. To the best of our knowledge, reproducing fluorescence in non-spectral renderers hasn't received much attention.

Fluorescent models. Most BRDF models assume a Lambertian reflectance [JHMD18], which confines fluorescence to a diffuse appearance. A common way to represent fluorescence is through the reradiation (or Donaldson) matrix [Don54]. We denote it by $\mathcal{P}(\lambda_i, \lambda_o)$, with its horizontal and vertical axes corresponding respectively to the incoming and outgoing wavelengths. The reflected radiance (Equation 1) is now expressed as:

$$L_o(\omega_o, \lambda_o) = \int \frac{\mathcal{P}(\lambda_i, \lambda_o)}{\pi} \int L_i(\omega_i, \lambda_i) \langle \omega_i, \omega_n \rangle d\omega_i d\lambda_i. \quad (2)$$

For a non-fluorescent material, \mathcal{P} is a diagonal matrix holding the spectral diffuse albedo ρ_d (i.e., $\mathcal{P}(\lambda_i, \lambda_o) = \delta(\lambda_i - \lambda_o) \rho_d(\lambda_i)$).

Hullin *et al.* [HHA*10] modeled the BRDF by an expansion into a series of products between angular and spectral functions, using a Principal Component Analysis (PCA). This goes beyond diffuse fluorescence, and allows for the rendering of data-driven fluorescent glossy materials. Recently, Benamira and Pattanaik [BP23] have extended the Microfacet theory [TS67] with an analytic fluorescent reflectance term. While our formulation should extend to non-diffuse fluorescence, we leave it to future work and focus on Lambertian appearance.

Our approach is based on spectral upsampling and downsampling operators. However, this should not be confused with the spectral upsampling method of Jung *et al.* [JWH*19]: in their approach, upsampling is used to create a new fluorescent material that yields vivid colors, while in ours it is used to accurately account for fluorescent effects of an *existing* material.

Fluorescence acquisition. One way to measure the reradiation matrix is to illuminate a sample with a monochromatic wavelength while a spectral sensor captures the spectral reflectance response. Such devices require costly and well-calibrated equipment, and are practically restricted to a single incoming-outgoing direction pair. Gonzalez and Fairchild [GF00] measured a large range of Lambertian samples using such a device (Labsphere BFC-450); we evaluate our method based on their database. The more recent work of Iser *et al.* [ILW23] (which improves on the method of Blaszki *et al.* [BFW20]) simplify the measurement process by using a Gaussian Mixture Model in place of a dense reradiation model. Our method could be equally applied to their database.

The method of Hullin *et al.* [HHA*10] tackles the challenge of measuring both angular and spectral variations in a BRDF. To this end, they simplify the acquisition process by measuring dense spectral data for a few angular configurations, and dense angular data with a reduced spectral sampling. They then combine these measurements through the aforementioned dedicated BRDF model based on a series expansion.

3. Naive Reradiation Matrix Reduction

In an RGB renderer, we track a triplet of outgoing and incoming radiance \mathbf{c}_o and \mathbf{c}_i in place of the spectral quantity L_o and L_i :

$$\mathbf{c}_o(\omega_o) = \int L_o(\omega_o, \lambda_o) \mathbf{s}(\lambda_o) d\lambda_o, \quad (3)$$

where \mathbf{s} is a vector of sensitivity functions, one per color channel. The definition for \mathbf{c}_i is similar. In the case of a non-fluorescent Lambertian material with an albedo vector $\boldsymbol{\rho}_d$, the rendering equation is approximated using the component-wise product \odot :

$$\mathbf{c}_o(\omega_o) \approx \frac{1}{\pi} \int \boldsymbol{\rho}_d \odot \mathbf{c}_i(\omega_i) \langle \omega_i, \omega_n \rangle d\omega_i. \quad (4)$$

By analogy, one can approximate the fluorescent counterpart (Equation 2) using vector-matrix multiplications:

$$\mathbf{c}_o(\omega_o) \approx \frac{1}{\pi} \int \mathbf{P}^{\text{Naive}} \times \mathbf{c}_i(\omega_i) \langle \omega_i, \omega_n \rangle d\omega_i, \quad (5)$$

where $\mathbf{P}^{\text{Naive}}$ is a reduced 3×3 reradiation matrix that accounts for reradiation from one color channel to the other. To render their BRDF measurements, Hullin *et al.* [HHA*10] suggested to construct such a reduced matrix by integrating the reradiation:


$$\mathbf{P}_{k,l}^{\text{Naive}} = \iint \bar{s}_k(\lambda_i) \bar{s}_l(\lambda_o) \mathcal{P}(\lambda_i, \lambda_o) d\lambda_i d\lambda_o, \quad (6)$$

where the $\bar{s}_k(\lambda) = \frac{s_k(\lambda)}{\|s_k\|}$, $k \in \{r, g, b\}$ are normalized RGB Color Matching Functions (CMFs). In matrix form, this yields:

$$\mathbf{P}^{\text{Naive}} = \bar{\mathbf{S}}^T \times \mathcal{P} \times \bar{\mathbf{S}} \quad (7)$$

where $\bar{\mathbf{S}} = [\bar{s}_r, \bar{s}_g, \bar{s}_b]$ is the $N \times 3$ matrix of normalized RGB sensitivity functions (with $N = 431$ wavelength bins in their case).

Limits of Naive Reduction. As shown by Hullin *et al.*, this naive reduction technique does not faithfully reproduce fluorescent materials (see Figure 2), from which they conclude that dense Donaldson matrices and spectral rendering are required.



(a) Naive Reduction

(b) Spectral Reference

Figure 2: Naive reduction. Using (a) the reduction method of Equation 7 does not correctly capture the color of the fluorescent material compared to (b) the spectral reference. We use the HERPIOYE material lit by a D65 spectrum.

We do not share their conclusion and show that a different formulation can more correctly reproduce fluorescence. Indeed, with this naive approach, an identity Donaldson matrix does not yield an identity reduced matrix, and introduces unwanted energy (see the first row of Figure 7). In the next section, we show that the correct construction of the reduced matrix is devoid of this issue, and hence produces much more faithful renderings.

4. Correct Reradiation Matrix Reduction

Here, we show that the correct way to model the reduced matrix is through spectral upsampling & downsampling with non-orthogonal bases. We start by describing our approach using XYZ triplets, as the corresponding sensitivity functions are non-negative.

4.1. Reduction through upsampling & downsampling

In the following, we approximate the bi-spectral rendering equation by a series of spectral downsampling and upsampling (see Figure 4). While our approach is applicable to the general form (Equation 1), we demonstrate it on the diffuse case of Equation 2 for simplicity.

We start by describing our spectral upsampling and downsampling operators, using the CMFs as spectral bases. The k -th color channel after spectral downsampling is then defined as:

$$\mathbf{down}[f]_k = \int f(\lambda_o) s_k(\lambda_o) d\lambda_o. \quad (8)$$

We can discretize downsampling with $S = [s_x, s_y, s_z]$ the $N \times 3$ matrix of sensitivity functions, and \mathbf{f} the N -D vector of the spectrum:


$$\mathbf{down}[\mathbf{f}] = S^T \times \mathbf{f}, \quad (9)$$

Since the s_k ($k \in \{x, y, z\}$) form a non-orthogonal basis, we must use its dual basis \tilde{s}_l ($l \in \{x, y, z\}$) to define upsampling from a color vector $\mathbf{c} = [\dots, c_i, \dots]$:

$$\mathbf{up}[\mathbf{c}](\lambda_i) = \sum c_l \tilde{s}_l(\lambda_i). \quad (10)$$

The dual basis also has an associated $N \times 3$ matrix $\tilde{S} = (SS^T)^{-1}S$ (i.e., verifying $S^T \tilde{S} = I_3$, the 3×3 identity matrix¹) to obtain the

¹ Since we have $S^T \times \tilde{S} = I_3 \iff (SS^T) \times \tilde{S} = S \iff \tilde{S} = (SS^T)^{-1}S$.



(a) Our Reduction

(b) Spectral Reference

Figure 3: Optimized reduction. Using (a) our formulation, we obtain a new reduced re-radiation matrix that is better conditioned for rendering fluorescent materials compared to (b) the spectral reference, using the same material and illumination as in Figure 2.

vector formulation for upsampling:

$$\mathbf{up}[\mathbf{c}] = \tilde{S} \times \mathbf{c}. \quad (11)$$

In a tristimulus renderer, instead of transporting spectral radiance, we transport tristimulus radiance $\mathbf{c} = \mathbf{down}[L(\omega, \cdot)]$, as illustrated in Figure 4. Using a downsampled incoming radiance $\mathbf{c}_i = \mathbf{down}[L_i(\omega_i, \cdot)]$ as input, we approximate Equation 2 by upsampling \mathbf{c}_i and multiplying the resulting spectrum by the reradiation matrix:

$$L_o(\omega_o, \lambda_o) \approx \int \frac{\mathcal{P}(\lambda_i, \lambda_o)}{\pi} \int \mathbf{up}[\mathbf{c}_i](\lambda_i) \langle \omega_i, \omega_n \rangle d\omega_i d\lambda_i, \quad (12)$$

with the final pixel color given by $\mathbf{c}_o = \mathbf{down}[L_o(\omega_o, \cdot)]$, which is equivalent to Equation 3. Since all operators are linear, we re-write \mathbf{c}_o in matrix form by swapping the order of integrals:

$$\mathbf{c}_o \approx \frac{1}{\pi} \int \left(\mathbf{P}^{\text{Ours}} \times \mathbf{c}_i \right) \langle \omega_i, \omega_n \rangle d\omega_i, \quad (13)$$

with the matrix-vector product given by:

$$\mathbf{P}^{\text{Ours}} \times \mathbf{c}_i = \mathbf{down} \left[\int \mathcal{P}(\lambda_i, \cdot) \mathbf{up}[\mathbf{c}_i](\lambda_i) d\lambda_i \right]. \quad (14)$$

Our reduced matrix is obtained by using the discretized forms provided in Equations 9 and 11 into Equation 14, yielding:

$$\mathbf{P}^{\text{Ours}} = S^T \times \mathcal{P} \times \tilde{S}, \quad (15)$$

Equation 13 has the same form as Equation 5, but is derived from a principled approach that highlights the importance of using the dual basis. Note that there is no need to normalize sensitivity functions with our approach. More importantly, our solution in Equation 15 yields an identity reduced matrix when given an identity reradiation matrix. Indeed, when $\mathcal{P} = I_N$, $\mathbf{P}^{\text{Ours}} = S^T \tilde{S} = I_3$.

We show in Figure 3 that our solution produces a closer match to the spectral reference. In particular, it does not introduce unwanted energy. As with the naive method, our approach may be applied iteratively to build paths starting from the light source and ending at the camera. However, our solution is justified thanks to our formulation based on downsampling and upsampling operators, as illustrated in Figure 4. As shown in the top row, even though the upsampled spectra (in purple) do not initially match the exact spectra

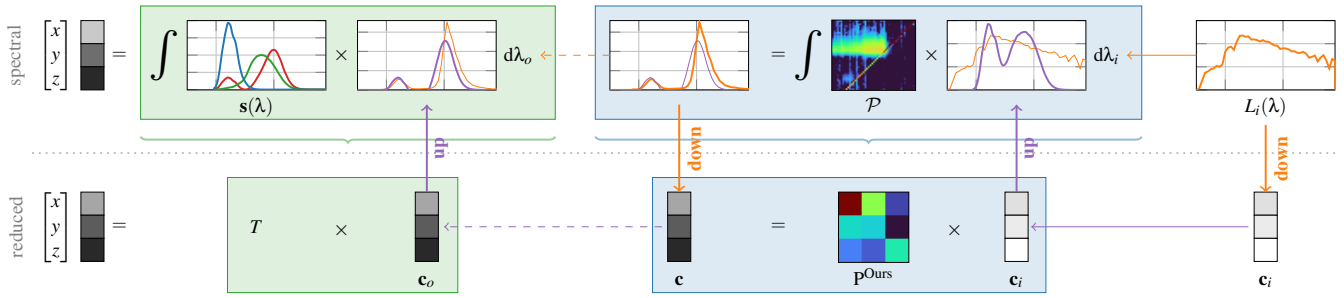


Figure 4: Reduced Light Transport through Upsampling & Downsampling. We illustrate the mapping between reduced transport (bottom row) and its spectral counterpart (top row). At each surface interaction (blue boxes), multiplication by the reduced matrix P^{ours} accounts for upsampling of the incoming reduced radiance, multiplication by \mathcal{P} , and downsampling of the outgoing spectral radiance. We compare the spectra obtained with this process (in purple) to those resulting from exact spectral light transport (in orange). After potentially multiple surface interactions (dashed arrows), spectral radiance must be integrated with the XYZ CMFs $s(\lambda)$ in the spectral pipeline (green boxes). In the reduced pipeline, this is directly done through multiplication with a matrix T , which is the identity for XYZ reduced light transport.

(in orange), they do closely reproduce them after a single multiplication by \mathcal{P} . The main advantage of our method is that this accurate approximation is done implicitly through P^{ours} (blue boxes), as shown in the bottom row. Note that P^{ours} will likely be entirely filled due to the wide basis support, even though \mathcal{P} is close to upper triangular (although not exactly due to measurement errors).

The final step of the rendering process (green boxes in Figure 4) is to integrate spectra over XYZ CMFs. For reduced light transport, this is done through a mere multiplication by a matrix T , which is the identity in the case of XYZ transport. It might differ with a different choice of basis, as discussed in the next section where we consider the extension of the XYZ basis functions to account for the reradiation from the UV-range.

4.2. Extending Reduction to Incorporate Non-Visible Spectra

We can use Equations 8 to 15 with an arbitrary set of non-orthogonal functions. We experimented the use of a fourth sensitivity function to capture the UV-range. To do so, we added the first element of a non-uniform B-spline Partition of Unity of 5 elements and degree 2 as detailed in Appendix A (the pink curve in Figure 5 (b), top row). During rendering, we need to transport an additional float with our XYZ triplet, which we denote by the term XYZU rendering. Figure 5 shows the impact of this additional basis function to handle fluorescence from UV lighting.

However, when using an alternative basis, a final projection by a transfer matrix T is required to project the n-uple in the rendering basis to the triplet of the CMF. In the case of XYZU rendering:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = T \times \mathbf{c}_o, \quad \text{with } T = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

5. Implementation & Results

We validate our method in two rendering scenarios: in a GPU rasterizer for real-time performance; and inside a path tracer to reproduce a spectral reference and test our method with global illumination. Each implementation showcases a distinct methodology to compute light transport.

5.1. Implementing Reduced Light Transport

The reduced Rendering Equation (Equation 13) does not fully detail how we should solve light transport. Depending on whether we solve light transport starting from the light or from the camera, either forward tracing or its adjoint must be implemented.

Light Tracing. This method should be used for photon mapping and direct lighting (see Section 5.2). It is the simplest as it follows Equation 13: we downsample the light spectrum to a reduced radiance vector that is multiplied by a reduced reflectance matrix at each surface interaction, until it reaches the eye.

Adjoint Transport. This method is required for backward path tracing (see Section 5.3). A difficulty here is that the adjoint of transporting reduced radiance vectors is to propagate a reduced matrix of throughputs (see Jarabo and Arellano [JA18] for details). In this case, we start at the sensor with the identity matrix and multiply it at each surface interaction with the corresponding reduced reflectance matrix. When next-event estimation queries the light,

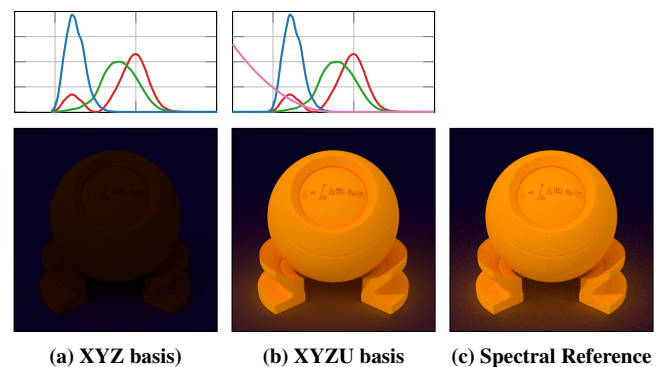


Figure 5: Adding a UV channel. Adding a fourth channel in the ultra-violet range, shown by the (b) pink curve, permits a more faithful reproduction of the (c) spectral reference when lit with UV lights. We used the HERPIOYE material lit by a Gaussian spectrum centred at 350 nm with a standard deviation of 50 nm.

	A	E	D60	D65	FL1	FL2	HP5
XYZ Naive	10.11	9.73	12.77	12.88	14.92	10.88	10.57
XYZ Ours	3.92	5.02	3.88	3.96	1.14	0.62	2.48
XYZU Naive	10.56	11.20	13.43	13.40	15.18	11.51	11.22
XYZU Ours	3.86	3.23	3.36	3.33	1.04	0.52	2.43

Table 1: Average ΔE_{2000}^* . We excluded IXCAXORA from the database [GF00] due to inconsistencies in the measurement.

we simply multiply the downsampled light spectrum with the accumulated throughput matrix.

5.2. Results for Direct Lighting

Unit Tests. We compared our method with a spectral reference first by rendering patches of measured materials illuminated by different lights sources (see Figure 7). On average, our method provides a smaller numerical difference than the naive method (in ΔE_{2000}^* , see Table 1). See our supplemental for a detailed evaluation of every sample of the database of Gonzalez and Fairchild [GF00].

Real-Time Results in Unity. We further integrated our reduced fluorescent method in Unity using the High Definition Render Pipeline (HDRP) by modifying the Lit shader. The working color space in Unity is RGB. We thus choose to convert the reduced reradiation from XYZ to RGB using:

$$\mathbf{P}_{RGB}^{\text{Ours}} = M_{XYZ \rightarrow RGB} \times \mathbf{P}_{XYZ}^{\text{Ours}} \times M_{XYZ \rightarrow RGB}^{-1}, \quad (16)$$

where $M_{XYZ \rightarrow RGB}$ converts a XYZ triplet to RGB. We used the light tracing formulation and stored two additional `vec3` uniforms for the 6 non-diagonal matrix elements. The only other modification is to change the component-wise vector multiplication between the albedo and the illuminant with a matrix-vector multiplication.

Figure 6 shows the rendering of a fluorescent material using both a reference path tracer with single scattering only and our implementation in Unity. We show that our method produces a similar appearance. We could not fully match the reference rendering since only Disney’s Diffuse model [Bur12] is implemented in Unity HDRP, not a Lambertian model. Fully matching it would require rewriting all custom preintegration code. Figure 1 shows the impact of fluorescence when changing the lighting environment.

5.3. Results in a Path-Tracer

We implemented a backward path tracer with next-event estimation using the adjoint formulation. We compare with a fully stochastic spectral path tracer capable of handling fluorescence by following Mojzik *et al.* [MFW18]. Our path tracer works in XYZ but can take arbitrary basis functions for its transport allowing us to add a 4th component for the UV band.

In Figure 3, we compare our approach with the stochastic reference using a D65 illuminant. The D65 spectrum expands outside the visible range and the chosen fluorescent material partially reradiates UV light in the visible range. We show in Figure 5 that we can add a 4th UV-component when light contains a large amount of energy outside the visible range.

We show our technique works with a wide range of fluorescent materials in Figure 8. It handles well multiple bounces with subtle color shifting as shown in the creases. The slight luminance decrease noted in Figure 3 is also noticeable here in XYZ for D65 and adding a UV band mitigates this issue. We further complement this set of results in our supplemental material.

6. Discussion

6.1. Limitations

Spectral manifold. Our method can only faithfully reproduce light transport within the manifold of the basis functions. Hence, in some cases, our method could not reproduce the same color as spectral rendering. We show in Figure 9 (a) a material where our method fails to correctly capture the appearance.

Number of bases. However, we can improve the quality of the approximation by using more basis functions. This is at the expense of a higher computational cost. For example, we show in Figure 9 (b) using 7 bases by splitting the X and Y channels (as explained in Appendix A) improves the quality of our approximation. Note however that for backward path tracing it requires to track 7×7 matrices, and use Equation 22 for T .

6.2. Future work

Alternate basis. So far, we have experimented with non-orthogonal bases derived from XYZ for ease of reconstruction. However, our method is generic enough to allow any form of basis. It would be interesting to see if less overlapping bases would yield better results. Working in the specific color spaces of cameras’ (or animals’) sensitivity functions might also be a sensible application-dependent choice.

Color transformation. We currently determine T by chaining the upsampling operator of the transport basis \mathbf{b} and the downsampling operator of the sensor’s basis \mathbf{s} with $T = S_s^T \times \tilde{S}_b$. This is a valid




Figure 6: Rendering in a Game Engine. Our formulation is easily integrated into existing RGB rendering pipelines, and is well suited for (b) real-time rendering. We used Unity’s Lit material, which implements the already available Disney Diffuse material [Bur12] and not a strict Lambertian. Hence it does not fully match the (a) path-traced reference of the HERPIOYE material.

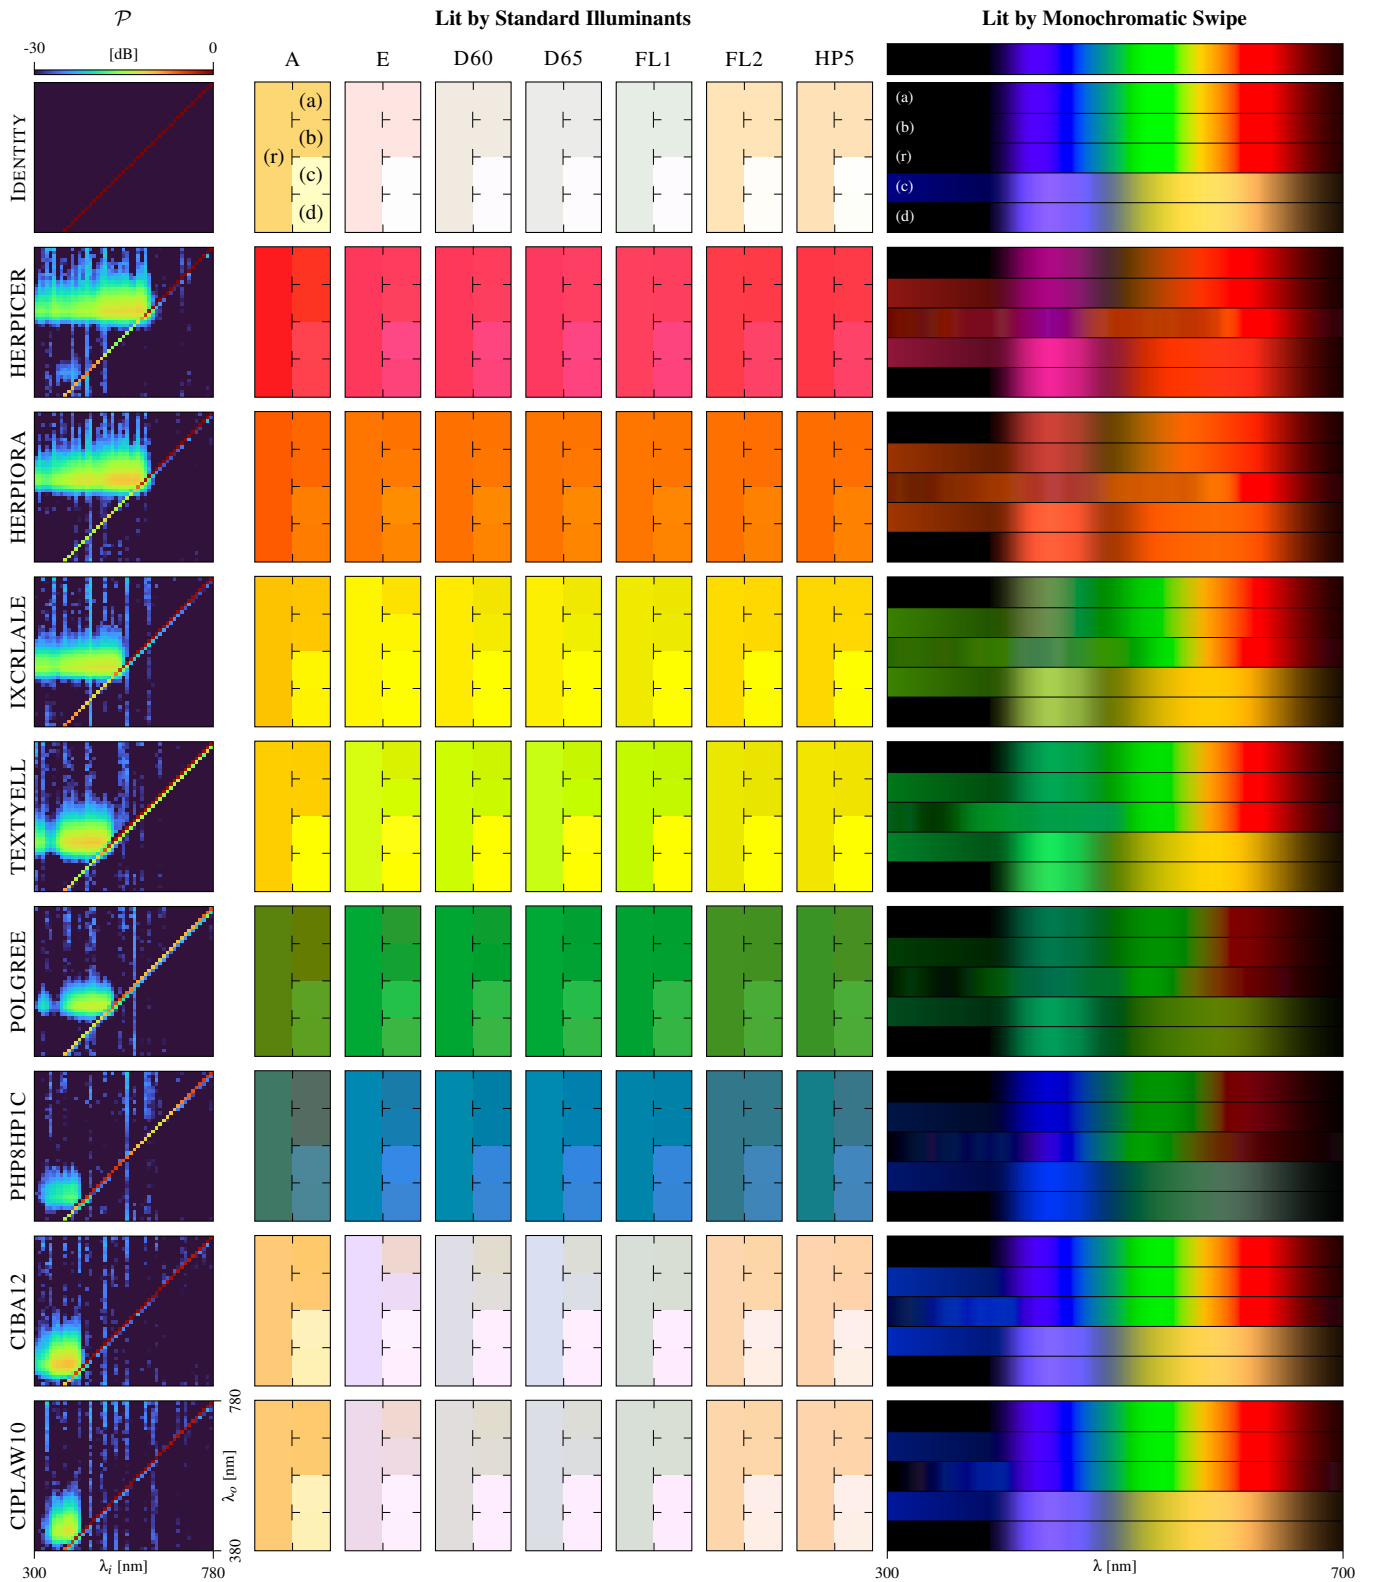


Figure 7: Validation of our method with patches. (r) Reference, (a) Ours XYZ, (b) Ours XYZU, (c) Naive XYZU, (d) Naive XYZ. We display the reradiation matrix (left), color after one bounce for various illuminants (center) and a swipe with a delta illuminant (right).




Figure 8: Validation in a path tracer. We validate our method for different illuminants (left column) and different materials applied to a material probe. We display the method of (a) Naive, (b) our XYZ method, (c) our XYZU method, and (r) the reference.




Figure 9: Limitation. The use of 4 bases in (a) is not sufficient to reproduce the spectral rendering of the HERPIMAG material, lit with a D65 illuminant in (c). When looking at the color swipes (bottom row), we observe a part of the reradiation is missing around 450 nm. This can be addressed by increasing the number of bases. In (b) we split the CIE-XYZ 2006 2° CMFs (as detailed in Appendix A) to better capture color shifting effects while keeping the reconstruction of the XYZ color linear.

solution but may not be optimal for human vision. Alternatively, we could optimize T as done for color space transformations [LF20]. Using perceptual distances such as ΔE_{2000}^* might provide a better result compared to our linear approach.

Glossy fluorescent materials. We have worked with Lambertian fluorescent models for which measured materials are easily available. Our method could apply to the work of Hullin *et al.* [HHA*10] since their decomposition of fluorescent materials is linear. For the model of Benamira and Pattanaik [BP23], more work is needed as the fluorescence is dependent on the incident angle.

Other phenomena. We wonder if our formulation of reduced light transport based on downsampling/upsampling could be applied to other spectral phenomena known to require spectral rendering, such as atmospheric/water rendering [EK10]. Provided that directional and spectral dimensions are independent, our formulation remains applicable.

Measuring Fluorescence. Our solution might be used to design simpler fluorescence measurement setups without dense spectral control on the source and sensor, provided one is only interested in capturing the *reduced* fluorescence matrix.

6.3. Conclusion

We have presented a new formulation to reproduce fluorescence appearance in non-spectral renderers such as RGB rasterizers or path tracers. Our method provides a principled way to reduce the dense reradiation matrix of BRRDF materials to enable correct rendering

of fluorescence. Our method permits to visually reproduce a spectral reference as well as to incorporate non-visible light leaking at a small cost in non-spectral rendering.

References

- [BFW20] BLASINSKI, HENRYK, FARRELL, JOYCE, and WANDELL, BRIAN. “Simultaneous Surface Reflectance and Fluorescence Spectra Estimation”. *IEEE Transactions on Image Processing* 29 (2020), 8791–8804. DOI: [10.1109/TIP.2020.2973810](https://doi.org/10.1109/TIP.2020.2973810) 2.
- [BHD*08] BENDIG, MARION, HANIKA, JOHANNES, DAMMERTZ, HOLLGER, et al. “Simulation of fluorescent concentrators”. *2008 IEEE Symposium on Interactive Ray Tracing*. 2008, 93–98. DOI: [10.1109/RT.2008.4634628](https://doi.org/10.1109/RT.2008.4634628) 2.
- [BP23] BENAMIRA, ALEXIS and PATTANAİK, SUMANT. “A Microfacet Model for Specular Fluorescent Surfaces and Fluorescent Volume Rendering using Quantum Dots”. *Eurographics Symposium on Rendering*. The Eurographics Association, 2023, 11–2313. DOI: [10.2312/ser.20231121](https://doi.org/10.2312/ser.20231121) 2, 8.
- [Bur12] BURLEY, BRENT. “Physically-Based Shading at Disney”. *ACM SIGGRAPH Courses*. 2012 5.
- [Dig22] DIGITAL, ECLAT. *Ocean*. 2022. URL: <https://eclat-digital.com/fluorescence-brighten-your-renderings/1>.
- [Don54] DONALDSON, R. “Spectrophotometry of fluorescent pigments”. *British Journal of Applied Physics* 5.6 (June 1954), 210–214. DOI: [10.1088/0508-3443/5/6/303](https://doi.org/10.1088/0508-3443/5/6/303) 2.
- [EK10] ELEK, OSKAR and KMOCH, PETR. “Real-time spectral scattering in large-scale natural participating media”. *Proceedings of the 26th Spring Conference on Computer Graphics*. SCCG ’10. Budmerice, Slovakia: Association for Computing Machinery, 2010, 77–84. ISBN: 9781450305587. DOI: [10.1145/1925059.1925074](https://doi.org/10.1145/1925059.1925074) 8.
- [GF00] GONZALEZ, SERGIO and FAIRCHILD, MARK. “Evaluation of Bispectral Spectrophotometry for Accurate Colorimetry of Printing Materials”. *Proceedings of Color Imaging Conference*. 2000, 14–23 2, 5.
- [Gla95] GLASSNER, ANDREW S. “A Model for Fluorescence and Phosphorescence”. *Photorealistic Rendering Techniques*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1995, 60–70 2.
- [HHA*10] HULLIN, MATTHIAS B., HANIKA, JOHANNES, AJDIN, BORIS, et al. “Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions”. *ACM Trans. Graph. (Proc. SIGGRAPH 2010)* 29.4 (2010), 97:1–97:7 1, 2, 8.
- [ILW23] ISER, TOMÁŠ, LACHIVER, LOÏC, and WILKIE, ALEXANDER. “Affordable method for measuring fluorescence using Gaussian distributions and bounded MESE”. *Opt. Express* 31.15 (July 2023), 24347–24362. DOI: [10.1364/OE.495459](https://doi.org/10.1364/OE.495459) 2.
- [JA18] JARABO, ADRIAN and ARELLANO, VICTOR. “Bidirectional rendering of vector light transport”. *Computer Graphics Forum*. Vol. 37. 6. Wiley Online Library. 2018, 96–105 4.
- [JHD20] JUNG, ALISA, HANIKA, JOHANNES, and DACHSBACHER, CARSTEN. “Spectral mollification for bidirectional fluorescence”. *Computer Graphics Forum*. Vol. 39. 2. Wiley Online Library. 2020, 373–384 2.
- [JHMD18] JUNG, ALISA, HANIKA, JOHANNES, MARSCHNER, STEVE, and DACHSBACHER, CARSTEN. “A simple diffuse fluorescent BRRDF model”. *Proceedings of the Eurographics 2018 Workshop on Material Appearance Modeling*. 2018, 15–18 2.
- [JWH*19] JUNG, A., WILKIE, A., HANIKA, J., et al. “Wide Gamut Spectral Upsampling with Fluorescence”. *Computer Graphics Forum* 38.4 (2019), 87–96. DOI: [10.1111/cgf.13773](https://doi.org/10.1111/cgf.13773) 1, 2.

- [LF20] LANGLANDS, ANDERS and FASCIONE, LUCA. “PhysLight: An End-to-End Pipeline for Scene-Referred Lighting”. *ACM SIGGRAPH 2020 Talks*. SIGGRAPH ’20. Virtual Event, USA: Association for Computing Machinery, 2020. ISBN: 9781450379717. DOI: [10.1145/3388767.3407368](https://doi.org/10.1145/3388767.3407368).
- [MFW18] MOJŽÍK, M., FICHET, A., and WILKIE, A. “Handling Fluorescence in a Uni-directional Spectral Path Tracer”. *Computer Graphics Forum* 37.4 (2018), 77–94. DOI: [10.1111/cgf.13477](https://doi.org/10.1111/cgf.13477) 2, 5.
- [Tob*18] TOBLER R. Wilkie, A. et al. *The Advanced Rendering Toolkit*. 2018. URL: <https://cgg.mff.cuni.cz/ART/1>.
- [TS67] TORRANCE, K.E. and SPARROW, E.M. “Theory for Off-Specular Reflection from Roughened Surfaces”. *Journal of the Optical Society of America (JOSA)* 57.9 (Sept. 1967), 1105–1114 2.
- [WTP01] WILKIE, ALEXANDER, TOBLER, ROBERT F., and PURGATHOFER, WERNER. “Combined Rendering of Polarization and Fluorescence Effects”. *Rendering Techniques 2001*. Vienna: Springer Vienna, 2001, 197–204. ISBN: 978-3-7091-6242-2 2.

Appendix A: Alternative bases

We have explored a couple choices of alternative basis functions besides the XYZ CMFs.

First, we have considered adding a single basis U to capture reradiation from the ultraviolet, focusing on the UVA range (*i.e.*, [315 – 400] nm) which accounts for approximately 95% of the ultraviolet radiation reaching the Earth’s surface. To this end, we build a Partition of Unity (PU) over the [300 – 800] nm range, using 5 second-order B-Spline bases. We obtained satisfying results with the first of these B-Splines as our additional UV basis U : it is monotonically-decreasing, with $U > 0.5$ for wavelengths below 400 nm. Yet other options could be considered, for instance when knowing specifics of material reradiation in the UV range.

Second, we have experimented with the use of more than 3 bases in the visual range. Our approach is to leave the Z CMF basis unchanged, but split the large Y CMF basis in two, and the bimodal X CMF basis in three. Starting with the Y basis, we split the corresponding CMF to get $\bar{y} = \bar{y}_1 + \bar{y}_2$, with:

$$\bar{y}_1(\lambda) = \bar{y}(\lambda)\mathcal{S}(\lambda; \mu_y, \sigma_y) \quad (17)$$

$$\bar{y}_2(\lambda) = \bar{y}(\lambda)[1 - \mathcal{S}(\lambda; \mu_y, \sigma_y)], \quad (18)$$

where \mathcal{S} is a smoothstep function centered on μ_y and scaled by σ_y . It is given by $\mathcal{S}(\lambda; \mu_y, \sigma_y) = 3\lambda_y^2 - 2\lambda_y^3$ with $\lambda_y = \frac{1}{2} \left[\frac{\lambda - \mu_y}{\sigma_y} + 1 \right]_0^1$ clipped to the [0, 1] range. We use $\mu_y = 570$ nm and $\sigma_y = 60$ nm.

The splitting of the X basis is similar, except we separate it in three $\bar{x} = \bar{x}_1 + \bar{x}_2 + \bar{x}_3$, with \bar{x}_1 (resp. $\bar{x}_2 + \bar{x}_3$) corresponding to the mode below (resp. above) 500nm. Formally, we have:

$$\bar{x}_1(\lambda) = \bar{x}(\lambda)[1 - \mathcal{S}(\lambda; 500, 2)] \quad (19)$$

$$\bar{x}_2(\lambda) = \bar{x}(\lambda)\mathcal{S}(\lambda; 500, 2)\mathcal{S}(\lambda; \mu_x, \sigma_x) \quad (20)$$

$$\bar{x}_3(\lambda) = \bar{x}(\lambda)\mathcal{S}(\lambda; 500, 2)[1 - \mathcal{S}(\lambda; \mu_x, \sigma_x)]. \quad (21)$$

We use $\mu_x = 590$ nm and $\sigma_x = 60$ nm.

Adding the U basis for ultraviolet reradiation yields $K = 7$ bases.

We obtain the final color by multiplying \mathbf{c}_o by:

$$T = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \quad (22)$$

The advantage of splitting XYZ basis functions to get finer bases is that they map to the final XYZ color through T without incurring further approximations. In general, we suggest that any set of basis functions that may be linearly combined (through T) to form the XYZ CMFs is an adequate choice.

Note that when rendering with forward path tracing, this requires transporting 7×7 reduced matrices denoted $\mathbf{P}_7^{\text{Ours}}$. An experiment we leave to future work would be to transport 4×4 reduced matrices $\mathbf{P}_4^{\text{Ours}}$, but use a 4×7 matrix $\mathbf{P}_{7 \rightarrow 4}^{\text{Ours}}$ when connecting with the light source. The latter matrix is then given by $\mathbf{P}_{7 \rightarrow 4}^{\text{Ours}} = T \times \mathbf{P}_7^{\text{Ours}}$.