



Efficient Storage and Importance Sampling for Fluorescent Reflectance

Q. Hua,^{1,2}  V. Tázlar,²  A. Fichet^{2,3}  and A. Wilkie² 

¹Saarland University, Saarbrücken, Germany
hua@cg.uni-saarland.de

²Charles University, Prague, Czech Republic
tazlarv@cgg.mff.cuni.cz, alban.fichet@gmx.fr, wilkie@cgg.mff.cuni.cz

³Unity Technologies, San Francisco, CA, USA

Abstract

We propose a technique for efficient storage and importance sampling of fluorescent spectral data. Fluorescence is fully described by a re-radiation matrix, which for a given input wavelength indicates how much energy is re-emitted at other wavelengths. However, such representation has a considerable memory footprint. To significantly reduce memory requirements, we propose the use of Gaussian mixture models for the representation of re-radiation matrices. Instead of the full-resolution matrix, we work with a set of Gaussian parameters that also allow direct importance sampling. Furthermore, if accuracy is of concern, a re-radiation matrix can be used jointly with efficient importance sampling provided by the Gaussian mixture. In this paper, we present our pipeline for efficient storage of bispectral data and provide its extensive evaluation on a large set of bispectral measurements. We show that our method is robust and colour accurate even with its comparably minor memory requirements and that it can be seamlessly integrated into a standard Monte Carlo path tracer.

Keywords: fluorescence, ray tracing, reflectance and shading models, rendering

CCS Concepts: • Computing methodologies → Reflectance modelling

1. Introduction

Traditionally, rendering pipelines mimic the human visual system and rely on RGB or tristimulus values to compute light transport. This approach, however, is insufficient if high colour accuracy or reproduction of spectral effects is required. Resolving the wavelength domain with additional bands, and thus performing *spectral rendering*, allows handling of such cases at the expense of increased computational demands. However, due to the increase in computation power and algorithmic efficiency [WND*14], spectral rendering is becoming increasingly popular even in production [PBC*18]. In this paper, we focus on the reproduction difficulties of one particular spectral effect: fluorescence.

Fluorescence is a phenomenon often found in natural elements and chemical compounds. It has a significant impact on how we perceive an object's colour in terms of hue, saturation and luminance because it transfers a part of non-visible or near-visible light into a different part of the visible range. Therefore, the inclusion of fluorescence can enhance material brightness, making the colour appear

more vivid compared to a pure reflection. While the importance of this effect for appearance modelling in computer graphics has been known for years, it has not been used much in practice, as adding it to graphics workflows—both on the side of user interfaces, as well as within rendering software proper—is a demanding problem.

An issue preventing a more widespread adoption of fluorescence is the significant memory footprint of its representation. To fully describe fluorescent spectra, one typically uses *re-radiation matrices* [Don54], also called *Donaldson matrices*, representing re-emission distributions for a given set of wavelengths. Although explicitly storing these data is manageable for a few fluorescent elements, this does not scale for fine-grained material definitions such as fluorescent textures.

In recent years, there has been an increasing interest in fluorescence in the research community. For example, Jung *et al.* [JWH*19] used fluorescence to enlarge the colour gamut during spectral uplifting in a fashion analogous to the use of optical brighteners in real-world materials. However, due to the lack of




Figure 1: In this scene, we use a near-UV light for illumination. While non-fluorescent materials only reflect shades of the illumination colour (left), fluorescent surfaces also re-emit a portion of the absorbed energy from incident light as a light at additional wavelengths, leading to the colourful appearance of other objects in the scene. To represent fluorescent materials in a renderer, we typically use re-radiation matrices (centre), which have a significant memory overhead. Instead, we propose a more efficient representation of these matrices using Gaussian mixture models (right – eight Gaussians). Although compact, this representation also provides compelling results even with such challenging illumination.

better alternatives, their method had to store a full re-radiation matrix for every texel, which in practice limited its usage to low-resolution input images.

We have proposed a method for efficient storage and importance sampling of re-radiation matrices using a Gaussian mixture model (Figure 1) in Hua *et al.* [HFW21]. This preliminary work was evaluated on three measured fluorescent materials. In this revision, we present an extended evaluation using an additional database of measured materials [GF00]. Further, we compare the previously proposed weighted Expectation–Maximization (EM) method with the weighted Bayesian approach of Gaussian mixture fitting. Finally, we provide additional evaluation metrics to assess the fidelity of the fitted re-radiation matrices.

2. Previous Work

2.1. Bispectral rendering

Although not widely used, the concept of fluorescence in computer graphics is not new. Glassner [Gla95] introduced fluorescence and phosphorescence to the computer graphics community and proposed a method to support those effects in a Whitted ray tracer. Wilkie *et al.* [WTP01] proposed a path tracing system capable of rendering both fluorescence and polarization effects, but they relied on a defined number of spectral bands, the wavelength not being part of the (MC) integration. Mojzík *et al.* [MFW18] made the wavelength domain part of the integrand and adapted Hero Wavelength Spectral Sampling (HWSS) [WND*14] to their method capable of handling fluorescent media in a Monte Carlo (MC) path tracer.

2.2. Colour gamut enlargement

One of the challenges when using a spectral renderer is the requirement of spectrally defined assets. Assets are commonly defined in

a tristimulus space (e.g., RGB, XYZ) and cannot be used directly as the input of a spectral renderer. To allow the usage of RGB assets in a spectral renderer, the original RGB data needs to be ‘uplifted’ to a spectral form. A wide range of work addresses non-fluorescent spectral uplifting, assuming smooth spectra [Smi99, [JH19], or data-driven from measured spectra [OYH18, TWF21]. Recently, Jung *et al.* [JWH*19] proposed an uplifting pipeline that additionally uses fluorescent spectra to enhance the colour gamut. A significant drawback of this bispectral uplifting method is its memory requirement. In the worst case, when dealing with texture uplifting, each pixel needs to be stored as a full re-radiation matrix. However, this is not an intrinsic drawback of the uplifting method, but an indicator that the representation of fluorescent data in rendering technology required improvement.

2.3. Bispectral measurements

Bispectral reflectance measurements are not as widespread as BRDF measurements because their acquisition is a lengthy process requiring expensive equipment. Additionally, with the limited number of renderers supporting fluorescence, there is little incentive to take bispectral material measurements for computer graphics applications. However, few exceptions are noteworthy. Gonzalez and Fairchild [GF00] offer a database of re-radiation matrices of numerous materials, mainly papers and inks. They used a Labsphere BFC-450 bispectral fluorescence colorimeter for the measurements. We use this database to consolidate our earlier published evaluation on the three re-radiation matrices from Labsphere distributed in ART [ART18] measured with the same device.

2.4. Bispectral material models

Wilkie *et al.* [WWLP06] proposed one of the early bispectral models adapted to current rendering techniques. This model

uses a layered BRDF with a diffuse fluorescent component. Hullin *et al.* [HHA*10] proposed an efficient acquisition setup to capture bispectral bidirectional reflectance and re-radiation distribution functions (BRRDF). They guided their acquisition by a principal component analysis to lower the acquisition time. Based on this work, Jung *et al.* [JHMD18] derived a new BRRDF. They exploited Kasha's rule to represent their distribution with three 1D distributions (absorption and emission spectra and non-fluorescent reflectance) and two ratios (emitted to absorbed energy and re-radiation to non-fluorescent reflectance).

2.5. Gaussian representations in rendering

Gaussian distributions are often used in rendering practice, as their sufficiently large combination can approximate distributions of any shape. Jakob *et al.* [JRJ11] represent radiance using a Gaussian mixture and applied it to the rendering of volumetric media. Herholz *et al.* [HEV*16] and Vorba *et al.* [VKŠ*14] employ Gaussian mixtures to represent importance sampling distributions for path guiding. Herholz *et al.* [HES*18] used Gaussian mixtures to sample analytic BRDF models and measured BRDF data.

3. Background

In this section, we review the core concepts used in the remainder of this article. We explain the fluorescent effect, its properties and the typical representation of fluorescent data in a bispectral renderer. We also review the basic notation of Gaussian mixture models (GMMs), which we employ to represent re-radiation matrices.


3.1. Fluorescence

Fluorescence is an effect where a *fluorophore* absorbs energy from incident photons, causing excitation of its molecules which, during their return to the ground state, may re-emit part of this energy as photons with lower energy. In contrast to phosphorescence, where emission can transpire over a long time period, fluorescence occurs in a few nanoseconds (10^{-9} to 10^{-7} s) and, as such, is considered to be instantaneous in the context of computer graphics.

Fluorescence produces noticeable visual effects and is commonly used to increase the brightness and saturation of pigments and dyes. This works by shifting light from the barely visible ultra-violet region to longer wavelengths, where the human eye is more sensitive.

A fluorophore can be characterized by its absorption and re-emission spectrum (Figure 2a). The first defines which incident wavelengths are absorbed and lead to re-emission events, with magnitudes representing the total intensity of the re-emissions. The second describes the amount of re-emission across all incident wavelengths.

According to Kasha's rule, the re-emission spectrum shape can be considered independent of the specific incident wavelength that caused the excitation (re-emission intensity scales). The difference between the spectral positions of the band maxima of absorption and re-emission arising from the same electronic transition is called a *Stokes shift* [MW97]. In practice, this shift usually corresponds to



(a) Fluorophore characterised by its absorption and re-emission spectrum. The offset between the two spectra is called the Stokes shift.

(b) Fluorophore characterised by its re-radiation matrix. The diagonal ($\lambda_i = \lambda_o$) was removed to emphasise the fluorescent effect.

Figure 2: (a) If Kasha's rule holds, the shape of a re-emission spectrum is the same regardless of the incident excitation wavelength, and absorption with re-emission spectrum form a sufficient representation of a fluorophore. (b) Otherwise, a re-radiation matrix providing precise information about the re-emission spectrum given an incident wavelength has to be used at the cost of its significant memory footprint.

a difference between the rightmost local absorption maximum and the re-emission maximum.

However, there are exceptions to Kasha's rule, leading to the need for re-radiation matrices (Figure 2b), which characterize, for each incident wavelength (λ_i), a distribution of re-emission wavelengths ($\lambda_i \neq \lambda_o$) and surface reflectance ($\lambda_i = \lambda_o$).

In the context of path tracing, a normalized re-emission spectrum (column excluding $\lambda_i = \lambda_o$) can be seen as a PDF of the wavelength shift of an irradiation λ_i and a normalized absorption spectrum (row excluding $\lambda_i = \lambda_o$) can be seen as a PDF of absorption of a wavelength λ_i for a given outgoing radiation λ_o .

To preserve energy conservation, the sum of each column of the re-radiation matrix never exceeds 1. The resulting value corresponds to the total albedo of the material at a given incident wavelength, that is, the sum of surface reflectance and fluorescence re-emission.

But, while re-radiation matrix is more flexible and suited for computer graphics applications, it has a significant memory footprint.

3.2. Gaussian mixture model

A GMM is a parametric probabilistic model representing a dataset as a linear superposition of Gaussian distributions. This linear superposition is represented by

$$p(\mathbf{x}) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x} \mid \mu_k, \Sigma_k) \quad (1)$$

where:

$p(\mathbf{x})$ is the density of mixture of Gaussians,
 $\mathcal{N}(\mathbf{x} \mid \mu_k, \Sigma_k)$ is a single Gaussian density,
 π_k is the mixing coefficient (weight) of each Gaussian,
 K is the number of Gaussians.

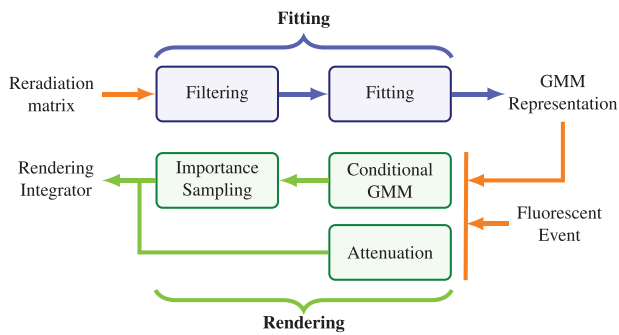


Figure 3: Fluorescence rendering pipeline employing our technique. We model a re-radiation with a GMM in the fitting phase, and then during the rendering phase, we importance sample wavelength shifting events according to the GMM.

In this article, we model re-radiation matrices as mixtures of two-dimensional multivariate Gaussian distributions:

$$\mathcal{N}(\mathbf{x} \mid \mu, \Sigma) = \frac{1}{2\pi \sqrt{|\Sigma|}} e^{-\frac{1}{2}(\mathbf{x}-\mu)^T \Sigma^{-1}(\mathbf{x}-\mu)} \quad (2)$$

where:

μ is the mean,
 Σ is the covariance matrix.

4. Method Overview

Our technique consists of two related but distinct operations that form an efficient fluorescence rendering pipeline (Figure 3):

1. Fitting, where the re-radiation matrices are fitted to GMMs (5).
2. Rendering, where the GMM representation is used to importance sample wavelength-shifting events (6).

The pipeline starts with a fitting preprocessing step. Here, the re-radiation matrices representing fluorophores are filtered to remove invalid values and fitted to GMMs. During the rendering phase, we use these GMM representations for fluorescent wavelength-shifting events, where a conditional GMM is importance sampled. Depending on the path direction, we either sample based on absorption or re-emission wavelengths.

Using a GMM to represent a bispectral reflectance drastically reduces its memory footprint in the renderer, as instead of relying on a tabulated re-radiation matrix, we only need to store a few parameters of the mixture. This reduction is even more pronounced when the rendering technique relies on random sampling, as additional tabulated CDFs are necessary for efficient sampling of re-radiation matrices.

5. Fitting Phase

In the fitting phase, we apply parametric density estimation to fit a re-radiation matrix into a GMM. Later, we show that we can accurately re-construct the re-radiation matrix information from the GMM on the fly.

5.1. Dataset filtering

Some measurements in the dataset include a significant amount of noise, potentially lowering the accuracy of the fit. To mitigate this problem, we filtered the re-radiation matrices before the fitting phase.

Although not ideal, we zeroed values under the diagonal ($\lambda_i > \lambda_o$) and negative values above it, since these values in our dataset correspond to measurement errors. We do not apply any further data filtering during the fitting phase. A further improvement would be possible with a proper characterization of the acquisition device's noise ratio.

5.2. Fitting algorithms

In our previous publication, we used the implementation of weighted EM provided by the Pomegranate library [Sch22]. In this revision, we instead use weighted EM in the better-behaving scikit-learn [PVG*11, sci21] library and compare its fitting accuracy with the weighted Bayesian fitting algorithm.

5.2.1. Weighted expectation-maximization

EM is a two-stage iterative algorithm suitable for the finding of mixture model parameters. The E-step (expectation) performs an estimation of the expected log-likelihood for a complete dataset. The M-step (maximization) maximizes the expected complete log-likelihood. The E-step and M-step are run iteratively until the expectation converges to the target. As a detailed introduction of EM is beyond the scope of our work, we review only a sketch of EM. For a detailed study on EM, please refer to Bishop [Bis06]. Vorba *et al.* [VKŠ*14] also showcase an adaptation of EM to the rendering context.

The EM algorithm works with a set of observation samples and computes a mixture matching its distribution. But, the re-radiation matrix instead provides weights of samples of pairs (λ_i, λ_o) . We could use a large set of samples distributed according to these weights, but then the accuracy of the representation would depend on the size of the sample set, whereas bigger sample sets would lead to higher memory consumption and a slower fitting process.

A more suitable approach is to use the weighted EM algorithm [GAPFH16, VKŠ*14] where all non-zero samples are used as observation outcomes and weighted according to their respective values.

5.2.2. Variational Bayesian inference

Variational inference follows the principle of the EM algorithm. It is a two-stage iterative algorithm that maximizes a posterior probability (MAP) instead of maximizing local likelihood (ML) as in the EM. Apart from updating the approximations in EM, the model evidence in variational inference also includes a pre-defined lowest prior distribution. By maximizing their posteriors using Bayesian inference in each iteration, the variational inference model gives more stable results compared to the EM, which can produce singularities and exhibit overfitting [Bis06]. Implementation of weighted

variational Bayesian inference in our pipeline directly follows the derivation of weighted EM [HFW21].

5.2.3. Optimal number of Gaussians

Choosing an optimal number of Gaussians for the EM is a non-trivial problem, as with a higher number of Gaussians, we may face overfitting issues. Bayesian information criterion (BIC) can mitigate overfitting by the introduction of a penalty for the number of parameters of the model.

Variational Bayesian inference can take advantage of the predefined prior distribution, such as the Dirichlet process (used in the scikit-learn implementation as default) to select the optimal number of clusters based on the fitted data.

We evaluate our fitting on different numbers of Gaussians for EM as the selection may depend on the use case. Conversely, we let Bayesian inference use a lower number of Gaussians than the specified maximum if it opts to do so for improved accuracy of the fit.

5.3. Re-radiation re-construction

We can re-construct the re-radiation matrix from the GMM found previously. Gaussian mixture gives a probability density function, and so its integral is unity. Because the re-radiation integral does not hold this property, the resulting Gaussian mixture has to be scaled to re-construct the re-radiation later on:

$$\Phi(\lambda_i, \lambda_o) = S \cdot p(\lambda_i, \lambda_o) \quad (3)$$

where:

- S is the scaling factor,
- $\Phi(\lambda_i, \lambda_o)$ is the re-radiation function,
- $p(\lambda_i, \lambda_o)$ is the probability density function of the GMM defined in Equation (1),
- λ_i, λ_o are the wavelengths of a given absorption and re-emission event.

There are different strategies to retrieve the scaling factor S . We originally proposed two approaches:

5.3.1. Error minimisation

The first strategy consists of a minimisation process. S is set as a parameter to minimise the error between the measured re-radiation and the re-constructed mixture.

$$S = \arg \min_S \sum_{\Lambda_i} \sum_{\Lambda_o} \|\Phi_{\text{measured}}(\lambda_i, \lambda_o) - S \cdot p(\lambda_i, \lambda_o)\| \quad (4)$$

The choice of a norm function has little influence on the result when dealing with fluorescence — the data have a low dynamic range. By definition, this strategy reduces the average re-construction error.

5.3.2. Integral ratio

The second strategy consists of defining S as the ratio between the integral of the original dataset and the GMM.

$$S = \frac{\int \int_{\Lambda} \Phi_{\text{measured}}(\lambda_i, \lambda_o) d\lambda_i d\lambda_o}{\int \int_{\Lambda} p(\lambda_i, \lambda_o) d\lambda_i d\lambda_o} \quad (5)$$

This method ensures that the re-constructed re-radiation has the same albedo as the original measured one.

Although this technique does not reduce the measurable re-construction error, its application provides better rendering results because our eyes are more sensitive to brightness variation than to a slight chromaticity shift.

5.3.3. Strategy selection

The strategy for scaling factor computation depends on the target application. The integral ratio is better suited for perceptual applications, such as rendering with tone mapping. If the accuracy of the spectral re-construction of the re-radiation is needed, then error minimisation is a better choice.

In this publication, we apply integral ratio scaling as it is more closely related to the intended application of our method.

6. Rendering Phase

When a fluorescent event occurs during the rendering phase, we first need to sample an in-shifting or an out-shifting wavelength depending on the ray direction. Then, we need to evaluate the attenuation given a set of λ_i and λ_o wavelengths. The sampling phase relies on conditional probability: either λ_i or λ_o is known, and we have to sample a λ_o or a λ_i according to the direction of the ray, that is, evaluate an in- or out-shifting event.

6.1. Conditional GMM

To sample a random wavelength shift after a fluorescent event, we need to ‘sample a slice’ on the multivariate Gaussian mixture. With a fixed input wavelength λ_i , we construct the conditional parameters for the resulting Gaussian mixture and generate the random variable as the output wavelength λ_o .

Our model is made of:

$$\mu_k = \begin{pmatrix} \mu_{k_a} \\ \mu_{k_b} \end{pmatrix}, \Sigma_k = \begin{pmatrix} \sigma_{k_{aa}} & \sigma_{k_{ab}} \\ \sigma_{k_{ba}} & \sigma_{k_{bb}} \end{pmatrix} \quad (6)$$

where μ_k and Σ_k are the GMM parameters of the k^{th} Gaussian.

We define T , the inversion of the k^{th} covariance matrix as:

$$T_k = \Sigma_k^{-1} = \begin{pmatrix} \tau_{k_{aa}} & \tau_{k_{ab}} \\ \tau_{k_{ba}} & \tau_{k_{bb}} \end{pmatrix}. \quad (7)$$

The conditional probability that, given λ_i , we shift to λ_o in component k is:

$$p_k(\lambda_i \rightarrow \lambda_o) = \mathcal{N}(\lambda_o | \mu_{k_{b|a}}, \tau_{k_{bb}}^{-1}) \quad (8)$$

$$\text{with : } \mu_{k_{b|a}} = \mu_{k_b} - \tau_{k_{bb}}^{-1} \tau_{k_{ba}} (\lambda_i - \mu_{k_a}). \quad (9)$$

Then, the probability for the entire conditional GMM from a given λ_i to λ_o is:

$$p(\lambda_i \rightarrow \lambda_o) = \sum_{k=1}^K \pi_k^{(\lambda_i)} p_k(\lambda_i \rightarrow \lambda_o) \quad (10)$$

$$\text{with : } \pi_k^{(\lambda_i)} = \frac{\pi_k \mathcal{N}(\lambda_o | \mu_{k_a}, \sigma_{k_{aa}})}{\sum_k \mathcal{N}(\lambda_o | \mu_{k_a}, \sigma_{k_{aa}})}. \quad (11)$$

Finally, we generate the random variable based on the new conditional GMM model. First, we choose which Gaussian component we will sample using the conditional mixing coefficients $\pi_k^{(\lambda_i)}$, and then we generate the sample following the normal distribution of the chosen k^{th} Gaussian:

$$\mu_k^{(\lambda_i)} = \mu_{k_{b|a}}, \sigma_k^{(\lambda_i)} = \tau_{k_{bb}}^{-1}. \quad (12)$$

We only discuss sampling on the light path, but the same approach can be symmetrically applied on an eye path.

6.2. Importance sampling

When a ray hits a fluorophore, two events can occur:

- the ray is reflected with the same wavelength as the incident ray — no fluorescent event,
- the ray is absorbed and its energy re-radiated to a different wavelength — fluorescent event.

We first determine if a fluorescent event occurs based on the ratio between reflectance and fluorescence for a given incident wavelength. If a fluorescent event occurs, we use importance sampling to select a Gaussian distribution from the mixture and a specific wavelength from the distribution (Algorithm 1).

Figure 4 shows an example of a conditional PDF of outgoing wavelengths of a specific incident wavelength λ_i . Given an incident wavelength, we can generate this distribution from GMM and use it to importance sample the re-emitted wavelength λ_o .

The probability of transition between absorption at wavelength λ_i and re-emission at wavelength λ_o is

$$p_{\text{shift}}(\lambda_i \rightarrow \lambda_o) = \begin{cases} \frac{\Phi(\lambda_i, \lambda_i)}{r_t(\lambda_i)} & \text{if } \lambda_i = \lambda_o \\ \left(1 - \frac{\Phi(\lambda_i, \lambda_i)}{r_t(\lambda_i)}\right) \cdot \left(\sum_{k=1}^K \pi_k^{(\lambda_i)} \mathcal{N}(\lambda_o | \mu_k^{(\lambda_i)}, \sigma_k^{(\lambda_i)})\right) \end{cases}$$

$$r_t(\lambda_i) = \int_{\Lambda} \Phi(\lambda_i, \lambda_r) d\lambda_r \quad (13)$$

where:

$\Phi(\lambda_i, \lambda_o)$ is the bispectral reflectance,

$r_t(\lambda_i)$ is the total reflected and re-emitted energy given λ_i .

Algorithm 1. Random wavelength shifting. Given a λ_i , we either have no fluorescent event, that is, $\lambda_i = \lambda_o$ or, we sample a λ_o after a fluorescent event.

```

input :  $\lambda_i$ 
output:  $\lambda_o$  and pdf :=  $p_{\text{shift}}(\lambda_i \rightarrow \lambda_o)$ 
 $\xi_1 \leftarrow \mathcal{U}(0, 1)$ ;
pdiag  $\leftarrow \frac{\Phi(\lambda_i, \lambda_i)}{r_t(\lambda_i)}$ ;
if  $\xi_1 \leq p_{\text{diag}}$  then
  // Sample the main diagonal
   $\lambda_o \leftarrow \lambda_i$ ;
  pdf  $\leftarrow p_{\text{diag}}$ ;
else
  // Wavelength shifting event
  /* Select one Gaussian from the mixture */
   $\xi_2 \leftarrow \mathcal{U}(0, 1)$ ;
   $s \leftarrow 0$ ;
  for  $k \leftarrow 1$  to  $K$  do
     $s \leftarrow s + \pi_k^{(\lambda_i)}$ ;
    if  $s \leq \xi_2$  then
      | break;
    end
  end
  /* Conditional sampling with the rejection method */
  do
    |  $\lambda_o \leftarrow \mathcal{N}(\mu_k^{(\lambda_i)}, \sigma_k^{(\lambda_i)})$ ;
  while  $\lambda_o \leq \lambda_i$ ;
  // Eqn. 13
  pdf  $\leftarrow (1 - p_{\text{diag}}) \cdot \left(\sum_{k=1}^K \pi_k^{(\lambda_i)} \mathcal{N}(\lambda_o | \mu_k^{(\lambda_i)}, \sigma_k^{(\lambda_i)})\right)$ ;
end

```




Figure 4: Comparison of conditional PDFs derived from a re-radiation matrix and its corresponding GMM representation. The figure shows the expected distribution of the re-emission wavelengths λ_o when the HERPICH A material is illuminated at $\lambda_i = 440$ nm. When a fluorescent event occurs, the outgoing wavelength λ_o is sampled according to this conditional PDF. Notice the presence of measurement noise lowering the highest values of the PDF derived from the raw re-radiation matrix after normalization compared to the PDF derived from the fitted GMM.

In our renderer, we use HWSS [WND*14]. The HWSS technique multiplexes multiple wavelengths in a MC sample. It differs from a simple multiplexing approach by using a ‘Hero’ component, which is used to make all directional decisions. This technique drastically reduces colour noise in contexts with a wavelength dependence on the directional decisions, such as in a participating medium. In the case of fluorescence, we allow each of the wavelengths in the HWSS vector to be independently importance sampled. If done directly, we can find non-zero samples with a zero probability, which is incorrect for a MC integrator. Mojzík *et al.* [MFW18] introduced a balance term to avoid this artefact. We use the same term in our implementation:

$$p_j(\lambda_{i_j} \rightarrow \lambda_{o_j}) = \prod_{k=1, k \neq j}^N p_{\text{shift}}(\lambda_{i_k} \rightarrow \lambda_{o_k}) \quad (14)$$

where:

$p_j(\lambda_{i_j} \rightarrow \lambda_{o_j})$ is the probability of j th HWSS sample,
 $p_{\text{shift}}(\lambda_i \rightarrow \lambda_o)$ is the probability of shifting from λ_i to λ_o ,
 N is the size of the HWSS vector.

6.3. Implementation details

To evaluate the re-radiation, we store the mean, covariance matrix and weight for each Gaussian of the mixture. We also need to store the scaling factor S . The determinant and inverse of the covariance matrix are pre-computed to speed up the evaluation of the mixture during rendering. The diagonal is stored in its original tabulated form without any alteration.

To efficiently perform conditional wavelength shifting, we also pre-compute two additional tabulated values to retrieve $\pi_k^{(\lambda_i)}$ or $\pi_k^{(\lambda_o)}$: the sum of rows and the sum of columns of the Gaussian mixture. In our implementation, we use a 1-nm sampling rate.

Although additional pre-computed values increase the total memory footprint during rendering, it is still much smaller than the memory required when using tabulated CDFs, which are necessary for the use of re-radiation matrices.

We use rejection sampling to guarantee that absorption wavelengths do not exceed re-emission wavelengths. It proved sufficient in our case, as we observed only a small number of rejected samples. Alternatively, truncated distributions could be used.

The implementation of our method is available on Gitlab.

7. Results and Discussion

Our fluorophore representation consists of the original re-radiation matrix diagonal representing the reflectance spectrum and the scaling factor S with the GMM representing the fluorescent effect. We need to store seven values for each Gaussian of the mixture:

- a mixing coefficient π (1 element),
- a mean vector μ (2 elements),
- a covariance matrix Σ (2×2 elements).

7.1. Dataset

In our previous publication [HFW21], we evaluated our method on three measured re-radiation matrices distributed in the assets of the ART rendering toolkit — originally provided by Labsphere Inc. In this revision, we provide a more in-depth evaluation on a much larger bispectral dataset. We additionally use the dataset from RIT [GF00], which contains 144 measured re-radiation matrices of different materials.

All measurements in the resulting dataset come from the same device: Labsphere BFC-450. The materials were sampled at 10-nm precision with an absorption wavelength ranging from 300 to 780 nm and a re-emission ranging from 380 to 780 nm, resulting in 49×41 re-radiation matrices.

One material measurement (IXCAXORA) was discarded from our evaluations. Its re-radiation matrix appears to be incorrect due to human error. We still include this sample in the supplemental.

In this section, we provide the results of the evaluation on the whole dataset but directly showcase only its small sub-set. Detailed results of our method on all samples are available in the supplemental.

7.2. Illuminants

To evaluate the accuracy of our representation, we compare the results of our fitted model with the original tabulated data. We base our accuracy evaluation on the CIE 2000 Delta E (ΔE_{00}^*) between a rendering of the reference data and the corresponding fits observed under a large set of illuminants.

We use two sets of illuminants and distinguish the results of each group:

- **Monochromatic illuminants:** We evaluate the bispectral reflectance of each fluorophore under monochromatic illumination from 300 to 780 nm to cover the full range of incident wavelengths in the dataset. The performance of our model is calculated as an average ΔE_{00}^* of all monochromatic illuminants.
- **CIE standard illuminants:** We also evaluate the colour reproduction accuracy of our model under the CIE standard illuminants, which cover often used illumination conditions in rendering practice.

Any measured differences between our model and pure re-radiation matrices are caused by our compact representation of the fluorescence, as the reflectance spectrum is not modified.

7.3. Rendering and colour reproduction

While a MC rendering provides additional information about inter-reflection between different fluorescent materials and self-inter-reflection, the choice of the scene and its setup also limit the evaluation to particular configurations. And so, given the large number of material measurements, we instead predominantly use an analytical rendering method. Specifically, we evaluate the accuracy of the fitting on an analytically calculated interaction of light with a

fluorescent material represented as a re-radiation matrix:

$$C = \int_{\lambda_o} \int_{\lambda_i} \phi(\lambda_i, \lambda_o) L_i(\lambda_i) d\lambda_i \bar{C}(\lambda_o) d\lambda_o \quad (15)$$

where:

C	is the computed outgoing colour,
λ_i, λ_o	are the incident and outgoing wavelengths,
$\phi(\lambda_i, \lambda_o)$	is the re-radiation matrix value for λ_i, λ_o ,
$L_i(\lambda_i)$	is the illuminant radiance at λ_i ,
$\bar{C}(\lambda_o)$	is the colour matching function at λ_o .

This equation is equivalent to a single light bounce configuration on a Lambertian surface where both incident and outgoing directions are aligned with the surface normal. This configuration allows us to assess the quality of the fits without any additional parameters interfering with our evaluations.

The re-radiation matrix corresponding to our fit can be constructed by evaluating the GMM at regular intervals for λ_i, λ_o . This approach allows efficient generation of noise-free results for each material in the dataset for a large variety of illumination types.

7.4. Rendering accuracy

Due to the fluorescent nature of our dataset, illumination conditions have a significant impact on error, making an objective evaluation difficult. Therefore, we evaluate our representation of fluorophores under a wide variety of illuminants visualized as

- **Colour band:** To visualize outgoing colour after interaction of a monochromatic illumination with a fluorophore, we use a ‘colour band’. It consists of three strips representing from top to bottom: outgoing colour based on a re-radiation matrix, outgoing colour based on our fitted model, and visualization of ΔE_{00}^* between these two. The colour bands cover the entire excitation spectrum of the fluorophore (horizontal axis) and allow wavelength-specific comparisons of the measured bispectral reflectance with our fit.
- **Colour patches:** To compare the quality of our fit with the original measured data, we use colour patches to visualize how well a fluorescent material is represented under a specific CIE standard illuminant. The fitted patch is embedded inside of a reference patch. The colour values are divided by the Y value of the used illuminant.

Figure 5 shows the five worst fluorophore fits based on their average ΔE_{00}^* for monochromatic illuminants at four Gaussians using EM. While the colour difference is visible for some monochromatic illuminants, the difference is hardly noticeable for most CIE standard illuminants, especially with a higher number of Gaussians. While four Gaussians would be suitable for most rendering purposes, eight Gaussians already provide high accuracy within a small memory footprint. This also applies to global illumination, as showcased in Figure 1.

This is supported by a comparison of the quality of the GMM fit with respect to the number of Gaussians used (Figure 6). Even though the reconstruction of the original re-radiation at four and two

Gaussians is very rough, it does not translate to big observable differences as our visual system is much more sensitive to luminance discrepancies than to chromaticity variations.

7.5. Average accuracy on the whole dataset

To compare the overall accuracy of the two fitting methods with a varying number of Gaussians, we provide a ΔE_{00}^* averaged across the entire dataset.

7.5.1. Fitting methods comparison

Figures 7b and 8b show a consistent decrease of the ΔE_{00}^* with an increasing number of Gaussians. Bayesian inference performs better with a lower number of Gaussians, but is outperformed by EM from eight Gaussians. This is caused by the Bayesian inference implementation opting to use a lower number of Gaussians for fitting than the allowed maximum.

7.5.2. Model accuracy

Figure 7a shows the average ΔE_{00}^* over the entire dataset for monochromatic illuminants ranging from 300 to 780 nm. The significant error observed in the lower parts of the spectrum can be attributed to limitations of the ΔE_{00}^* involving colours with low luminance. In our case, this part of the spectrum usually represents non-visible light or a negligible amount of re-radiation, as absorption bands of fluorophores are often situated there. Therefore, the practical influence of this error on real-world rendering applications is limited. The monochromatic illumination error is visualized for a few selected fluorophores in Figure 5 and the whole dataset in the supplemental.

In Figure 8a, we provide an evaluation of a more practical application of our approach—rendering with CIE standard illuminants. We show an average of the whole dataset for the ten worst performing illuminants while fitting four Gaussians with EM. On average, we achieve a ΔE_{00}^* lower than 1.1 for two, 0.6 for four and 0.3 for eight Gaussians.

7.6. Memory footprint

While retaining high accuracy, our method is memory efficient, requiring only a fraction of the memory necessary for use of the raw re-radiation matrices.

In the most favourable case, we only have to store the upper triangular part of the re-radiation matrix due to the nature of fluorescent re-radiation being of lower energy than excitation caused by the incident illumination. For our dataset measured at 10 nm precision, it is approximately 1150 values. If maximal precision was required, each matrix could be interpolated to 1 nm precision, resulting in a hundredfold increase in memory footprint. Additionally, to facilitate efficient importance sampling of wavelength-shifting events, it is desirable to compute tabulated CDFs, effectively tripling the total required memory when using re-radiation matrices for rendering.



Figure 5: Comparison of analytical renderings based on our fitted model with EM and re-radiation matrices. Five worst behaving fluorophores on the fitting of four Gaussians based on the average ΔE_{00}^* over the set of monochromatic illuminants from 300 to 780 nm are showcased. The colour bands show the bispectral reflectance of the re-radiation matrix and our corresponding fit under monochromatic illumination along with ΔE_{00}^* for each wavelength. The colour patches show bispectral reflectance of the fitted model embedded inside of a patch rendered with a re-radiation matrix. The worst performing (average ΔE_{00}^*) CIE standard illuminants on the whole dataset are displayed.




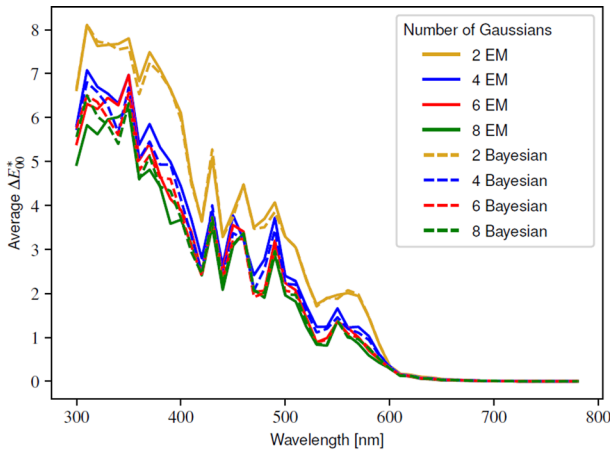
Figure 6: Worst performing fluorophores' re-radiation matrices with absorption and re-radiation spectra and corresponding fits of EM.

With our method, we only need a small set of GMM parameters to both evaluate the re-radiation on the fly and to importance sample wavelength shifting events.

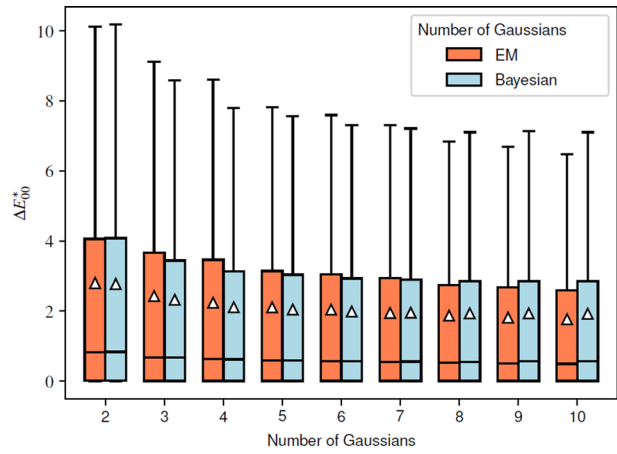
In practice, for a representation using N Gaussians, we need to store $1 + 7 \cdot N$ values. For importance sampling, we need two additional 1D arrays to efficiently evaluate the conditional GMM. This is

much less than the two tabulated 2D CDFs required for a fluorescent material when using a re-radiation matrix. Additionally, if desired, our method can be used solely for importance sampling instead of tabulated CDFs.

The memory efficiency of our method becomes especially pronounced when working with larger numbers of re-radiation

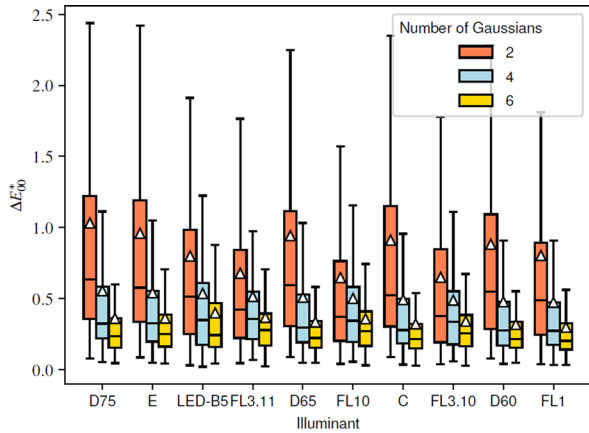


(a) Average ΔE_{00}^* over the whole dataset for each wavelength.




(b) Distribution of ΔE_{00}^* over the whole dataset over all monochromatic illuminants for different number of Gaussians.

Figure 7: Evaluation of the fitting accuracy on monochromatic illuminants. In the box plot, the box extends from first to third quartile, the lines show the median and the triangles show the average value. The whiskers extend from quartile to 1.5 inter-quartile range.



(a) Distribution of ΔE_{00}^* over the whole dataset of the 10 on average worst behaving CIE standard illuminants at 4 Gaussians with EM.



(b) Distribution of ΔE_{00}^* over the whole dataset over all CIE standard illuminants for different number of Gaussians.

Figure 8: Evaluation of the fitting accuracy on CIE standard illuminants. The box extends from first to third quartile, the lines show the median and the triangles show the average value. The whiskers extend from quartile to 1.5 inter-quartile range.

matrices. It makes fluorescent textures where each pixel may have to be represented by a re-radiation matrix usable in rendering practice. This, in turn, facilitates the use of related techniques such as enlargement of the colour gamut by uplifting of RGB textures [JWH*19].

8. Conclusion and Future Work

We presented a method for efficient storage of re-radiation matrices via GMM. We showed that this approach drastically lowers required memory for the rendering of fluorescence while retaining colour reproduction accuracy matching that of the original data. Further, the resulting GMM representation can be directly used in MC-based rendering pipelines, as it supports efficient importance sampling.

We evaluated weighted EM and Bayesian inference fitting approaches on an extensive dataset of measured re-radiation matrices under all CIE standard illuminants and a wide range of monochromatic illuminants. Both fitting methods proved to provide a comparable and highly accurate representation of fluorescent re-radiation with as little as eight Gaussians for most practical rendering use cases.

Different models and fitting approaches could be explored, but considering the rendering accuracy and small memory footprint of the presented solution, the potential gains seem marginal.

We think that the main direction of future work is in the development of a parametric fluorescence model. Such model would not

require tabulated CDFs — further lowering total memory requirements for efficient rendering of fluorophores. More importantly, it would allow a smooth integration of editable fluorescent materials into artistic workflows.

Acknowledgements

We thank the reviewers for their extensive feedback. This work has received funding from Czech Science Foundation Grant 19-07626S, the EU H2020-MSCA-ITN-2020 Grant Number 956585 (PRIME), and the Charles University Grants Number SVV-260588 and GA UK number 240422. The teaser uses the Serapis bust model from McGuire [McG17].

Open access funding enabled and organized by Projekt DEAL.

References

- [ART18] The advanced rendering toolkit. <https://cgg.mff.cuni.cz/ART/> (2018). Accessed 18 October 2022.
- [Bis06] BISHOP C. M.: *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006.
- [Don54] DONALDSON R.: Spectrophotometry of fluorescent pigments. *British Journal of Applied Physics* 5, 6 (June 1954), 210–214.
- [GAPFH16] GEBRU I. D., ALAMEDA-PINEDA X., FORBES F., HORAUD R.: Em algorithms for weighted-data clustering with application to audio-visual scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38, 12 (2016), 2402–2415.
- [GF00] GONZALEZ S., FAIRCHILD M.: Evaluation of bispectral spectrophotometry for accurate colorimetry of printing materials. In *Proceedings of the Color Imaging Conference* (June 2000), pp. 14–23.
- [Gla95] GLASSNER A. S.: A model for fluorescence and phosphorescence. In *Photorealistic Rendering Techniques*. G. Sakas, S. Müller, P. Shirley (Eds.). Springer, Berlin, Heidelberg (1995), pp. 60–70.
- [HES*18] HERHOLZ S., ELEK O., SCHINDEL J., KŘIVÁNEK J., LENSCH H.: A unified manifold framework for efficient BRDF sampling based on parametric mixture model. In *EGSR'18: Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations* (2018), Eurographics Association.
- [HEV*16] HERHOLZ S., ELEK O., VORBA J., LENSCH H., KŘIVÁNEK J.: Product importance sampling for light transport path guiding. *Computer Graphics Forum* 35, 4 (July 2016), 67–77.
- [HFW21] HUA Q., FICHET A., WILKIE A.: A compact representation for fluorescent spectral data. In *Eurographics Symposium on Rendering* (Saarbrücken, Germany, June 2021). <https://hal.archives-ouvertes.fr/hal-03274233>
- [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 Transactions on Graphics* 29, 4 (July 2010). <https://doi.org/10.1145/1778765.1778834>
- [JH19] JAKOB W., HANIKA J.: A low-dimensional function space for efficient spectral upsampling. *Computer Graphics Forum* 38, 2 (May 2019), 147–155. <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13626>
- [JHMD18] JUNG A., HANIKA J., MARSCHNER S., DACHSBACHER C.: A simple diffuse fluorescent BBRDF model. In *Proceedings of the Workshop on Material Appearance Modeling* (2018), The Eurographics Association, pp. 4. <https://diglib.org/handle/10.2312/mam20181193>
- [JRJ11] JAKOB W., REGG C., JAROSZ W.: Progressive expectation-maximization for hierarchical volumetric photon mapping. *Computer Graphics Forum (Proceedings of EGSR)* 30, 4 (June 2011). <https://10.dtwj>
- [JWH*19] JUNG A., WILKIE A., HANIKA J., JAKOB W., DACHSBACHER C.: Wide gamut spectral upsampling with fluorescence. *Computer Graphics Forum* 38, 4 (2019), 87–96.
- [McG17] MCGUIRE M.: Computer graphics archive. <https://casual-effects.com/data> (July 2017). Accessed 18 October 2022.
- [MFW18] MOJZÍK M., FICHET A., WILKIE A.: Handling fluorescence in a uni-directional spectral path tracer. *Computer Graphics Forum* 37, 4 (July 2018), 77–94.
- [MW97] MCNAUGHT A., WILKINSON A.: *IUPAC. Compendium of Chemical Terminology* (2nd edition). Blackwell Scientific Publications, Oxford, 1997.
- [OYH18] OTSU H., YAMAMOTO M., HACHISUKA T.: Reproducing spectral reflectances from tristimulus colours. *Computer Graphics Forum* 37, 6 (2018), 370–381.
- [PBC*18] PHARR M., BURLEY B., CHRISTENSEN P., FAJARDO M., FASCIONE L., KULLA C.: Design and implementation of modern production renderers. In *SIGGRAPH'18: Proceedings of the ACM SIGGRAPH 2018 Panels* (New York, NY, USA, 2018), Association for Computing Machinery. <https://doi.org/10.1145/3209621.3214901>.
- [PVG*11] PEDREGOSA F., VAROQUAUX G., GRAMFORT A., MICHEL V., THIRION B., GRISEL O., BLONDEL M., PRETTENHOFER P., WEISS R., DUBOURG V., VANDERPLAS J., PASSOS A., COURNAPEAU D., BRUCHER M., PERROT M., DUCHESNAY E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [Sch22] SCHREIBER J.: pomegranate. <https://github.com/jmschrei/pomegranate> (2022). Accessed 18 October 2022.
- [sci21] scikit. Gaussian mixture—weighted implementation. <https://github.com/scikit-learn/scikit-learn/pull/17130> (2021). Accessed 18 October 2022.

- [Smi99] SMITS B.: An RGB-to-spectrum conversion for reflectances. *Journal of Graphics Tools* 4, 4 (Dec. 1999), 11–22.
- [TWF21] TÓDOVÁ L., WILKIE A., FASCIONE L.: Moment-based constrained spectral uplifting. In *Proceedings of the Eurographics Symposium on Rendering - DL-only Track* (2021), A. Bousseau and M. McGuire (Eds.), The Eurographics Association. <https://doi.org/10.2312/sr.20211304>
- [VKŠ*14] VORBA J., KARLÍK O., ŠÍK M., RITSCHER T., KŘIVÁNEK J.: On-line learning of parametric mixture models for light transport simulation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2014)* 33, 4 (Aug. 2014), 1–11.
- [WND*14] WILKIE A., NAWAZ S., DROSKE M., WEIDLICH A., HANIKA J.: Hero wavelength spectral sampling. In *EGSR'14: Proceedings of the 25th Eurographics Symposium on Rendering* (Goslar, DEU, 2014), Eurographics Association, pp. 123–131. <https://doi.org/10.1111/cgf.12419>
- [WTP01] WILKIE A., TOBLER R. F., PURGATHOFER W.: Combined rendering of polarization and fluorescence effects. In *Rendering Techniques 2001*. S. J. Gortler and K. Myszkowski (Eds.). Springer, Vienna (2001), pp. 197–204.
- [WWLP06] WILKIE A., WEIDLICH A., LARBOULETTE C., PURGATHOFER W.: A reflectance model for diffuse fluorescent surfaces. In *GRAPHITE'06: Proceedings of the 4th International Conference on Computer Graphics and Interactive Techniques in Australasia and Southeast Asia* (New York, NY, USA, 2006), Association for Computing Machinery, pp. 321–331. <https://doi.org/10.1145/1174429.1174484>
- [XWHM20] XIA M., WALTER B., HERY C., MARSCHNER S.: Gaussian product sampling for rendering layered materials. *Computer Graphics Forum* 39 (2020), 420–435.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Supplemental Data S1 - Supplemental document presenting our whole dataset and fits on all of its materials.

Supplemental Data S2 - Zip file of our repository containing implementation of fitting code.