



Affordable method for measuring fluorescence using Gaussian distributions and bounded MESE

TOMÁŠ ISER, *  LOÏC LACHIVER, AND ALEXANDER WILKIE

Charles University, Faculty of Mathematics and Physics, Malostranské nám. 25, Prague, 11800, Czech Republic

*tomas@cgg.mff.cuni.cz

Abstract: We present an accurate and low-cost method for measuring fluorescence in materials. Our method outputs an estimate of the material's *Donaldson matrix*, which is a commonly used two-dimensional spectral characterization of its fluorescence and reflectance properties. To find the estimate, only a few measurements of the material's reflectance under a few illuminants are needed, which we demonstrate using low-cost optical components. Internally, our algorithm is based on representing each Donaldson matrix with a *multivariate Gaussian mixture model* and its diagonal with a *bounded MESE* (maximum entropy spectral estimate). It parametrizes and constrains the estimate in a robust and simple way, allowing the use of gradient-descent optimization. We evaluate our algorithm on a combination of real and synthetic data, and four examples of distinct optical components. We reach significantly lower errors than the current state of the art on the exact same inputs, our estimates do not suffer from artifacts such as oscillations of the spectra, and they are stable and robust.

© 2023 Optica Publishing Group under the terms of the [Optica Open Access Publishing Agreement](#)

1. Introduction

When a material is illuminated, a fraction of the incoming photons is reflected. The ratio between the reflected and incoming photons is called the material's *surface reflectance*. In general, the reflectance varies per wavelength, which gives materials their various colors. Materials may also absorb and re-emit a fraction of the incoming photons at a different wavelength than they originally had, which is called *wavelength shifting*. For instance, minerals such as fluorite or calcite will glow in color even when illuminated only with invisible ultraviolet light. This effect is called *fluorescence*, and it is immediate and observable with bare eyes.

Fluorescence is common in both natural (minerals, corals, tissues) and synthetic (pigments in papers, textiles, plastics) objects. Distinct *fluorophores* (chemicals) can also be combined in a single object. Due to the re-emission of incoming photons into longer wavelengths, fluorescent materials can appear brighter and more saturated than their non-fluorescent counterparts, and their color or reflectance may significantly change under various illuminants (Fig. 1). This has been studied in computer graphics to allow accurate visualizations and color rendering of such objects [1–7], but also in remote sensing of vegetation [8], and in fluorescence imaging in biology and medicine, including diagnostics in animals and humans [9], plants [10], and coral reefs [11].

All of these methods rely on measuring the fluorescence in the materials. A simple solution is shining monochromatic light on the sample, measuring the whole reflected spectrum, and repeating this for each illumination wavelength, e.g., in 10 nm steps (Sec. 2.1; [12]; partially illustrated in Fig. 1(bc)). It has several downsides: it needs repeating for each wavelength, it requires an expensive monochromator with a powerful broadband light source, and the monochromatic light is weak leading to even weaker fluorescent emission and substantial measurement noise.

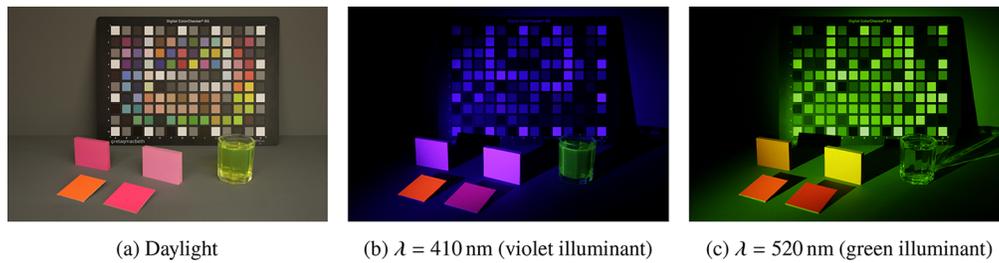


Fig. 1. Photographs of fluorescent (sticky notes, liquid detergent) and non-fluorescent (color chart) materials under different illuminants. The monochrome illuminants reveal which objects are fluorescent: they emit different wavelengths than the illuminants had.

Alternative solutions are based on *sparse measurements* (Sec. 2.2): instead of scanning all combinations, only a few illuminants are used and the reflectances may be measured by a sensor with a coarse resolution. The fluorescence is then estimated by solving an *inverse problem*, i.e., by numerically finding which fluorescence could have resulted in the observed reflectances. Some of these approaches still require expensive equipment such as programmable filters [13] or high-frequency spectral illuminants [14], only assume RGB data [15], or can inherently only support one fluorophore in each material [16–20]. The most important work in our context is by Blasinski et al. [21], who not only summarize the previous approaches, but also unify and generalize several of them into a new cohesive framework. Their method can estimate fluorescence, even from multiple fluorophores simultaneously, using any arbitrary illuminant spectra and sensor spectral sensitivities, which makes it compatible with low-cost optical setups.

However, their estimates can be significantly improved, which we show both quantitatively and qualitatively in our evaluations, where our new method outperforms their results on the exact same inputs. For example, their estimated spectra often suffer from oscillations and ripples: a flat spectrum is reconstructed as a wavy spectrum oscillating around the ground truth value, or a single peak is reconstructed as multiple peaks instead. Similar problems are common in methods that represent the spectra with basis functions that were built from fluorophore and reflectance datasets [17–21]. In our method, we represent fluorescence in a different way, using Gaussian distributions and maximum entropy spectral estimates, which we picked to ensure that the estimated fluorescence is accurate, including fluorescent peaks and potential flat spectra.

Our contribution We develop a robust algorithm for accurately estimating fluorescence and reflectance from sparse measurements acquirable with low-cost optical setups. Our algorithm is based on solving an inverse problem. While several previous approaches use a dataset of reflectances and fluorophores to reduce the dimensionality of the problem [17–21], we instead represent the fluorescence with a *multivariate Gaussian mixture model* and the reflectance with a *bounded MESE* (maximum entropy spectral estimate), which, to our knowledge, is their first application in such measurements. We show that our approach is simple, yet accurate and robust, and even though the inverse problem is non-convex, it behaves well with our chosen gradient-descent optimization strategy. We demonstrate how our method adapts to optical setups of different equipment costs. We evaluate all on a combination of real and synthetic data and compare to state of the art, showing consistently lower errors and higher-quality spectra.

2. Problem statement and prior work

The reflectance from fluorescent materials can be mathematically modeled (Sec. 2.1) and the measurement problem can then be written and solved as an inverse problem (Secs. 2.2 and 2.3).

2.1. Modeling fluorescence with the Donaldson matrix

Assuming fixed illumination and observation angles, when a fluorescent material is illuminated by a light source with spectrum $I(\lambda_i)$, then the reflected intensity $r(\lambda_o)$ is (Fig. 2):

$$r(\lambda_o) = \int \Phi(\lambda_i, \lambda_o) I(\lambda_i) d\lambda_i, \quad (1)$$

where $\Phi(\lambda_i, \lambda_o)$ is a two-dimensional function jointly describing the material's reflectance and fluorescence, and λ_i and λ_o are incoming (excitation) and outgoing (reflected, emitted) wavelengths, respectively. For practical purposes, we discretize the spectral dimension into N_i incoming and N_o outgoing wavelengths, and then $I \in \mathbf{R}^{N_i}$ and $r \in \mathbf{R}^{N_o}$ become vectors, $\Phi \in \mathbf{R}^{N_o \times N_i}$ becomes a matrix, and Eq. (1) becomes a matrix-vector multiplication (Fig. 2):

$$r = \Phi \cdot I. \quad (2)$$

Finding the exact values of the Φ matrix for the given material is our goal. The matrix is often called the *Donaldson matrix* after Donaldson [12], who measured the matrix values directly using monochromatic light (as introduced in Sec. 1). Note that the matrix exists also for non-fluorescent materials: it is simply zero everywhere except for the diagonal ($\lambda_i = \lambda_o$), which represents pure reflectance. As Fig. 2 shows, it is natural to separate the diagonal Φ_{diag} (*pure reflectance*) and the off-diagonal Φ_{fluo} (*pure fluorescence*), and write $\Phi = \Phi_{\text{diag}} + \Phi_{\text{fluo}}$.

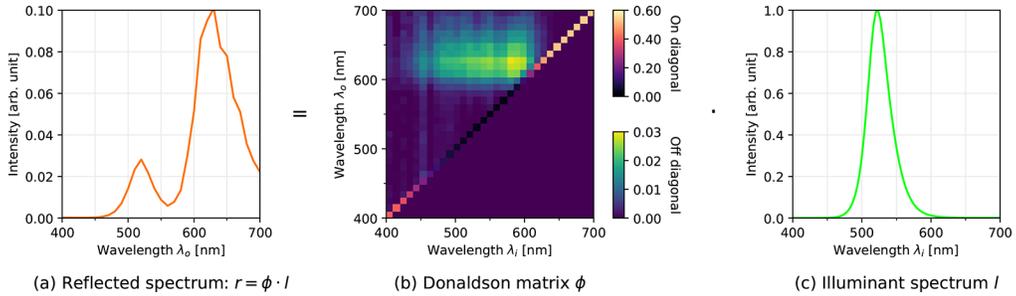


Fig. 2. Illustration of Eqs. (1) and (2). The reflected spectrum (a) is computed from the Donaldson matrix (b) and illuminant spectrum (c). While the illuminant is green, the reflection is orange with two peaks, which results from fluorescence.

Sometimes, assumptions are made to simplify the shape of Φ [7,12,21]. First, we expect the emitted photons to have longer wavelengths than the excitation, hence $\Phi(\lambda_i, \lambda_o) = 0$ for $\lambda_i > \lambda_o$, or equivalently, the matrix is zero below the diagonal. Second, by following Kasha's rule, Φ_{fluo} would be a separable function $\Phi_{\text{fluo}}(\lambda_i, \lambda_o) = \phi_{\text{em}}(\lambda_i) \cdot \phi_{\text{ex}}(\lambda_o)$ with one-dimensional emission $\phi_{\text{em}}(\lambda_i)$ and excitation $\phi_{\text{ex}}(\lambda_o)$ spectra [21]. But this does not hold in general, especially for materials with multiple fluorophores, so we do not assume such separability in this publication.

2.2. Sparse measurements

As measuring the Donaldson matrix Φ directly is expensive, slow, and suffers from noise, we focus on the idea of estimating Φ from only *sparse measurements* (Fig. 3) that are typically faster to acquire and do not always rely on expensive components [13–21]. In general, such a measurement setup consists of:

- a set of n_i illuminants with spectra $I_1(\lambda), \dots, I_{n_i}(\lambda)$, and
- a detector of the reflected light, which can detect n_s spectral channels with spectral sensitivities $s_1(\lambda), \dots, s_{n_s}(\lambda)$.

In our method, we impose no implicit restrictions. The detector can be an RGB camera with only 3 channels, as well as a spectrometer with a sub-nanometer resolution and hundreds of channels. The illuminants can be individual LEDs, but also a single broadband lamp with switchable filters.

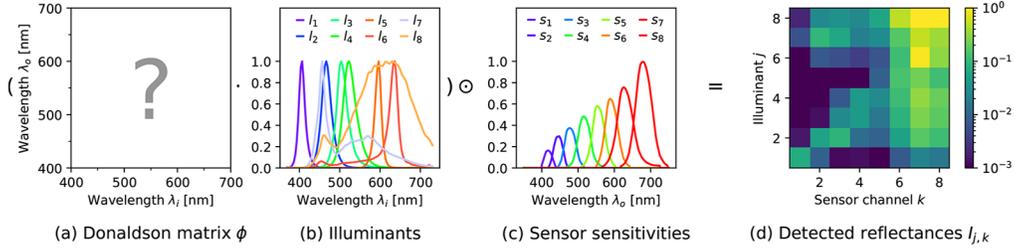


Fig. 3. Illustration of Eq. (3) (sparse measurements). In this example, the Donaldson matrix (a) is unknown, and the sparse measurements are based on 8 illuminants (b) and an 8-channel sensor (c), resulting in 8×8 detected intensity values $I_{j,k}$ (d).

By discretizing the problem and using the notation above, we can write that when the fluorescent sample is illuminated by the j -th illuminant, then the k -th channel of the detector should detect the reflected intensity of (Fig. 3):

$$(\Phi \cdot l_j) \odot s_k, \quad (3)$$

where \odot is an element-wise multiplication (Hadamard product). In order to estimate an unknown Φ , we can hence illuminate the sample once with each illuminant and detect the reflected intensities per channel, which gives $n_l \times n_s$ measurements denoted as $I_{j,k}$. Estimating the Donaldson matrix $\hat{\Phi}$ then becomes an inverse optimization problem, for which we chose the Euclidean distance (ℓ^2 norm) between the actual measurements and their estimation:

$$\hat{\Phi} = \underset{\Phi}{\operatorname{argmin}} \sqrt{\sum_j^{n_l} \sum_k^{n_s} (I_{j,k} - (\Phi \cdot l_j) \odot s_k)^2}. \quad (4)$$

In general, this problem is significantly underdetermined and has infinitely many solutions, most of which are not even physically plausible. Therefore, a good method needs not only to minimize the error, but mainly converge to a realistic and plausible solution.

2.3. Previous approaches finding a solution of Eq. (4)

To simplify the space of possible solutions of Eq. (4), many approaches use basis functions derived from a database of a priori known reflectances and fluorophores and essentially interpolating between them [17–21]. As we explained in Sec. 1 and show in Sec. 4.5, fitting the spectra onto the small linear bases results in artifacts such as a flat spectrum being reconstructed with oscillations.

From the broader perspective, Blasinski et al. [21] categorized the existing sparse approaches into two groups: *bispectral separation* [14,16–19] and *computational separation* [13,15,20,21]. Using our notation from Sec. 2.2, we could say that bispectral separation methods rely on high-resolution spectral detectors, meaning $n_s > 30$. Such tailored methods can operate under a very small number of illuminants n_l , e.g., two broadband illuminants [16], two high-frequency illuminants [14], or even just one spiky illuminant [18]. On the other hand, the computational separation methods employ more complex algorithms to allow cheaper detectors, e.g., just a simple RGB camera with two broadband illuminants [15] or a set of narrowband illuminants [20]. Using this terminology, our method would certainly fall into computational separation.

Out of all of these methods, only these by Suo et al. [13], Blasinski et al. [21], and ours are general enough to resolve materials with more than one fluorophore, mainly because they do not

build on the assumption of Kasha's rule and separability (Sec. 2.1). The algorithm of Blasinski et al. reaches higher accuracy than Suo et al., it employs and unifies the concepts from the other methods, and it is the most recent method. For the purpose of comparisons, we refer to it as the *state of the art*.

3. Estimating fluorescence with Gaussian distributions and bounded MESE

We design a sparse measurement algorithm that finds a solution to Eq. (4). It ensures that the Donaldson matrix is smooth and plausible by parametrizing it via Gaussian mixtures and bounded MESE, and it finds the final estimate using gradient-descent optimization.

3.1. Overview

Input The inputs of our algorithm are sparse measurements of the material following Fig. 3 and Sec. 2.2. The inputs include the spectra $l_1(\lambda), \dots, l_{n_l}(\lambda)$ of the n_l illuminants, the spectral sensitivities $s_1(\lambda), \dots, s_{n_s}(\lambda)$ of the detector's n_s spectral channels, and the $n_l \times n_s$ detected values of the material's reflectances corresponding to $I_{j,k}$ in Eq. (4) and Fig. 3(d).

Output The output of our algorithm is an estimate of the material's Donaldson matrix Φ . Because it is represented parametrically in our algorithm, it has a small memory footprint, and a fully continuous signal Φ can be easily reconstructed with any spectral resolution.

Processing We split the estimate $\Phi = \Phi_{\text{diag}} + \Phi_{\text{fluo}}$ into the off-diagonal elements Φ_{fluo} (Fig. 4(ab)) and the diagonal Φ_{diag} (Fig. 4(c)). The fluorescence Φ_{fluo} is represented by a multivariate Gaussian mixture model (GMM), which consists of individual two-dimensional Gaussian distributions parametrized by their intensity, mean, and covariance (Sec. 3.2). The diagonal Φ_{diag} is represented by a small set of trigonometric moments that define a continuous signal given by the bounded maximum entropy spectral estimate (bounded MESE) (Sec. 3.3). Using Eq. (4) then allows estimating the whole matrix via gradient-descent optimization (Sec. 3.4).

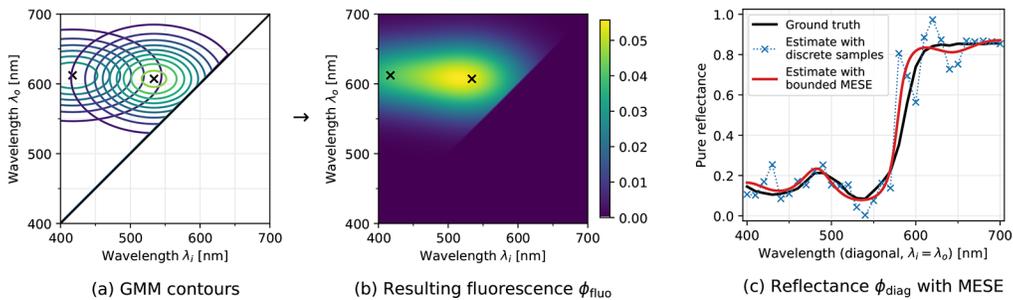


Fig. 4. Illustration of our parametric model. The Gaussian mixture model (GMM) consists of independent distributions. Their contours and mean values (\times symbols) are visualized in (a), and their superposition Φ_{fluo} in (b). The pure reflectance Φ_{diag} (c) is visualized with its discrete samples (blue) together with the final bounded MESE (red).

3.2. Fluorescence as a Gaussian mixture

A two-dimensional Gaussian mixture model (GMM) is a linear superposition of Gaussian distributions (Fig. 4(ab)). Following Hua et al. [7], we define it as a weighted sum of n_g

individual Gaussians:

$$f(\mathbf{x}) = \sum_{m=1}^{n_g} w_m \mathcal{N}(\mathbf{x} \mid \mu_m, \Sigma_m), \quad (5)$$

where $\mathcal{N}(\mathbf{x} \mid \mu, \Sigma): \mathbf{R}^2 \rightarrow \mathbf{R}$ is a two-dimensional Gaussian distribution:

$$\mathcal{N}(\mathbf{x} \mid \mu, \Sigma) = \left(2\pi\sqrt{|\Sigma|}\right)^{-1} \exp\left(-\frac{1}{2}(\mathbf{x} - \mu)^T \Sigma^{-1}(\mathbf{x} - \mu)\right), \quad (6)$$

and $w_m > 0$, $\mu_m \in \mathbf{R}^2$, and $\Sigma_m \in \mathbf{R}^{2 \times 2}$ are the weight, mean, and covariance matrix of the m -th distribution in the mixture. In our context, the GMM and Eq. (6) are interpreted to be in the wavelength domain, meaning $\mathbf{x} = (\lambda_i, \lambda_o)$ is a vector of the incoming and outgoing wavelengths, and $\mu = (\mu_i, \mu_o)$ represents the mean incoming and outgoing wavelengths. The pure fluorescence of the Donaldson matrix is then defined as $\Phi_{\text{fluo}}(\lambda_i, \lambda_o) = f(\mathbf{x})$.

This concept follows the recent publication of Hua et al. [7], who worked on compressing fluorescent textures in photorealistic rendering. They observed that the fluorescence distribution's shape resembles Gaussian distributions and it can be modeled as a sum of a few Gaussians. We use this observation in the new context of measurements as it allows using only a few parameters to control the fluorescence while also ensuring that the result will be physically plausible to a certain degree. Furthermore, the estimate naturally supports multiple fluorophores, it is continuous and can be evaluated at any wavelength resolution, and it is directly compatible with efficient importance sampling in Monte Carlo rendering [7]. We observed that the GMM is a smooth, differentiable function that behaves well in the gradient-descent optimization setting.

3.3. Diagonal as a bounded MESE

We represent the diagonal Φ_{diag} as a one-dimensional bounded maximum entropy spectral estimate (bounded MESE, Fig. 4(c)). Following Peters et al. [22], we know that a bounded 2π -periodic signal $g(\varphi) \in [0, 1]$ can be represented by $m + 1$ complex Fourier coefficients $\mathbf{c} = (c_0, \dots, c_m)$, which can also be viewed as the signal's trigonometric moments:

$$\mathbf{c} = \int_{-\pi}^{\pi} g(\varphi) \mathcal{F}(\varphi) d\varphi \in \mathbf{C}^{m+1}, \quad (7)$$

where $\mathcal{F}(\varphi)$ is the Fourier basis:

$$\mathcal{F}(\varphi) = \frac{1}{2\pi} (\exp(-ij\varphi))_{j=0}^m \in \mathbf{C}^{m+1}. \quad (8)$$

With the whole signal $g(\varphi)$ reduced to only $m + 1$ complex coefficients, we also need a way to solve the inverse problem: finding a signal that corresponds to the given coefficients. Broadly speaking, MESE [22,23] is the result of such an inverse process: it is an estimated signal that best corresponds to the given coefficients \mathbf{c} , and by "best" we mean that it maximizes Burg entropy, so the estimate is a relatively smooth signal. This is a memory compression mechanism, as a whole curve is represented by only a few numbers, but it also acts as a smoothing operator, because after computing the first few moments of a noisy signal, its MESE will be a smoother signal (Fig. 4(c), compare the blue samples to the red curve).

Spectral reflectances, in our case $\Phi_{\text{diag}}(\lambda) \in [0, 1]$, are generally not periodic signals and they are additionally bounded due to energy conservation. Hence, in our algorithm, we follow Peters et al.'s method [22] for computing bounded MESE and transforming them to reflectance spectra by mirroring and wavelength-mapping the signals. Such transformed moments also lose their imaginary part, so they become real instead of complex numbers. The meaningful ranges of the moment values are $c_0 \in [0, 1]$ and $c_j \in [-\frac{1}{\pi}, \frac{1}{\pi}]$ for $j \in \{1, \dots, m\}$. We confirm Peters et al.'s observation that 8 moments ($m = 7$) are sufficient for reflectance spectra.

3.4. Estimating the Donaldson matrix via gradient descent

The GMM and bounded MESE models from Secs. 3.2 and 3.3 together form a *parametric model* that describes a continuous Donaldson matrix Φ with only a handful of parameters (Fig. 4). Plugging our model into Eq. (4) leads to an inverse problem, in which these model parameters are the unknown variables that we need to solve for.

This problem has significantly fewer parameters and is more constrained than just naïvely discretizing the whole matrix, which in our early experiments did not lead to any meaningful results. For example, a Donaldson matrix in the wavelength range of 400 to 700 nm discretized with a 10 nm resolution has 496 free parameters (the whole matrix has $31 \cdot 31 = 961$ elements, out of which 31 are the diagonal, 465 are the upper left triangle, and 465 are zero), whereas our parametric model has almost ten times fewer parameters.

Our parametric model is non-linear and non-convex, since a sum of multiple Gaussian distributions can have more local extrema. Fortunately, it has well-defined derivatives, which allows using numerical optimization algorithms based on *gradient descent* [24]. We take inspiration from the field of deep learning, where non-convex optimizations are often solved by adaptive step sizes. Specifically, we noticed that the adaptive moment estimation algorithm *Adam* [25] converges well to estimates with high accuracy.

Optimized parameters In order to minimize Eq. (4), the optimization algorithm estimates the following parameters:

- for each Gaussian distribution, its weight $w > 0$, mean wavelength $\mu \in [400, 700]^2$, and covariance matrix $\Sigma \in \mathbf{R}^{2 \times 2}$ (we used only two non-zero elements to prevent rotations of the distribution), in total 5 real numbers per distribution,
- for the diagonal, the trigonometric moments $c_0 \in [0, 1]$ and $c_j \in [-\frac{1}{\pi}, \frac{1}{\pi}]$ for $j \in \{1, \dots, m\}$ defining the bounded MESE curve, in our case using 8 moments ($m = 7$).

While estimating the trigonometric moments directly via the optimization algorithm is possible and leads to valid results, in practice, it is computationally expensive and also hinders the accuracy. Hence, during the optimization, we instead discretize the diagonal with a 10 nm resolution, leading to 31 parameters (visualized by the blue crosses in Fig. 4(c)), and we let the optimizer estimate them individually with a simple penalization to keep them in the bounded $[0, 1]$ range. Only then, we compute the 8 trigonometric moments corresponding to the discretized diagonal, which finally gives a continuous smooth bounded MESE curve (red curve in Fig. 4(c)).

Initial estimates For the optimization to converge, each parameter needs to be initialized within a valid and meaningful range. Experimentally, we verified that the specific initialization does not play such an important role. We initialized the diagonal as a constant function of $0.5 = 50\%$ reflectance, and each Gaussian distribution with a covariance matrix $\Sigma = \{\{4000, 0\}, \{0, 4000\}\}$ to make their support large enough to reach from the edges to the center, weight $w = 100$, and their mean wavelengths covering the whole triangle, making sure every corner is covered, i.e., for 6 Gaussian distributions: (400, 400), (550, 550), (700, 700), (550, 400), (700, 400), and (700, 550) nm.

Iterations and step sizes To ensure fast convergence, we ran each optimization in three loops with decreasing step sizes. Hence, the Adam algorithm is executed 3 times per Donaldson matrix: first for 100 iterations with a step size of 0.10; second for 200 iterations and 0.05; and third for 800 iterations and 0.02. The second and third loops are initialized with the parameters that yielded the lowest error in the previous loop. Furthermore, the step sizes need to be scaled for each parameter separately, because different parameters have different ranges, e.g., the reflectance is only valid between 0 and 1, whereas the Gaussian covariance matrix can reach large values. The exact values are not critical, because Adam is an adaptive algorithm. In our case, we multiplied the reflectance's step size by 0.05, and each Gaussian parameter's step size by 200, e.g., the actual step size of the Gaussian mean wavelength in the first loop would be $0.10 \cdot 200 = 20$.

3.5. Algorithm implementation

The algorithm itself was implemented in Python and the source codes together with accompanying data are attached in [Code 1](#), Ref. [26]. Most of the implementation is straightforward application of the equations in the paper, with one exception, which is the optimization itself. The gradient-descent algorithm requires computing the gradients (partial derivatives) of Eq. (4) w.r.t. the parameters of the Gaussian distributions and the diagonal. We achieved that using the `jax` auto-differentiation framework [27]. With such a naive implementation, the run-time per material is a couple of minutes, but since it can be trivially executed in parallel, estimating a batch of 32 materials on a 32-core CPU takes the same time as a single material. For comparison, the implementation of Blasinski et al. [21,28] is faster, taking a couple of seconds per material, but that is for the price of lower accuracy.

4. Evaluation

We evaluate our method on a combination of synthetic data and real measurements. We first introduce examples of optical setups compatible with our method (Sec. 4.1), then our synthetic dataset (Sec. 4.2), and then our real measurements (Sec. 4.3). The evaluation itself is both quantitative, based on commonly used error metrics (Sec. 4.4), and qualitative, based on visual examinations of the matrices, spectra, and predictive image rendering (Sec. 4.5).

4.1. Optical setups

We introduce two sets of illuminants (LEDs and color filters) and two kinds of detectors (a spectrometer and a multispectral sensor chip) to demonstrate the variability and compatibility of our method. These illuminants and detectors can be used interchangeably, which gives $2 \cdot 2 = 4$ different examples of compatible measurement setups. We illustrate some of them in Fig. 5 and their spectra in Fig. 6.

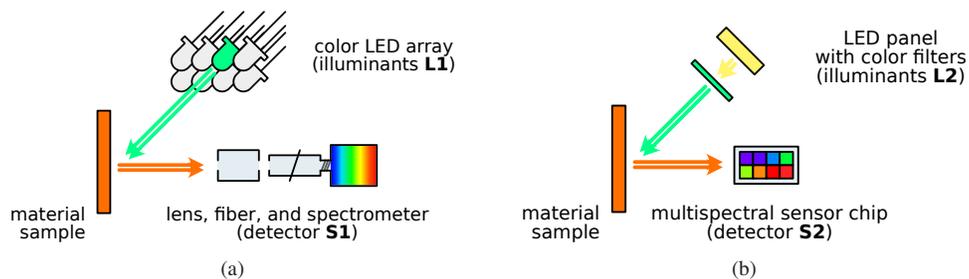


Fig. 5. Examples of optical setups with different optical components compatible with our algorithm. The abbreviations L1, L2, S1, and S2 are from Sec. 4.1.

Measurement geometry Figure 5 shows what is often called a $45^\circ/0^\circ$ measurement: a diffuse material sample is illuminated under approximately 45° , while the detector is placed orthogonally to the sample to avoid specular reflections. To measure how the material performs under varying angles, our method could in principle be extended to a fully goniophotometric setup, e.g., with the detector on a rotating arm. Our algorithm would then reconstruct one Donaldson matrix per each angle separately, and these results could be interpolated if needed.

Calibration The Donaldson matrix is a unitless distribution that describes ratios of the reflected and re-emitted energy, so the exact units of $l_1(\lambda), \dots, l_{n_l}(\lambda)$ and $s_1(\lambda), \dots, s_{n_s}(\lambda)$ are arbitrary. It is critical to normalize the values with respect to a single value to ensure that the estimated ratios in the Donaldson matrix are meaningful. For example, Fig. 6(bd) are normalized to the peak of the strongest spectrum. To calibrate the illuminants, we used Spectralon, which has a

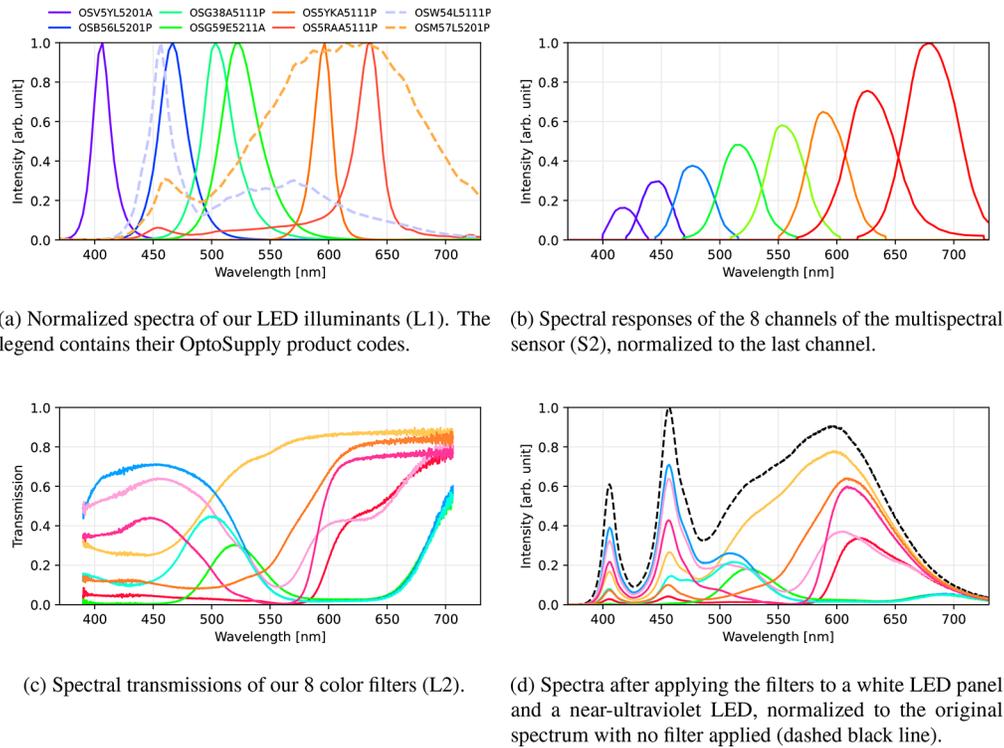


Fig. 6. Measured spectra of our exemplar optical setups from Sec. 4.1 and Fig. 5.

guaranteed 99% diffuse reflectance in the whole visible spectrum. By illuminating the Spectralon with each of the illuminants, the illuminant spectra are obtained from the known reflectance of 99%. It is also important that the detector's spectral sensitivity is calibrated beforehand, such that the detected peaks have correct intensities. Many spectrometers and detectors come factory pre-calibrated, or their calibration is possible with a stabilized light source of a known spectrum.

Examples of illuminants (L1, L2) The goal is to have a set of illuminants distinctly covering the intended spectral range, in our case the visible range from around 400 to 700 nanometers. We demonstrate two affordable options. The first option (denoted **L1**) is *light-emitting diodes* (LEDs), which are readily available with various spectra. In our example (measured in Fig. 6), we use 6 narrowband and 2 broadband (warm white and cold white) LEDs from OptoSupply with the cost of about 0.50 U.S. dollars per piece. The second option (denoted **L2**) is placing different *color filters* in front of a white illuminant. In our example (measured in Fig. 6), we purchased colored foils from a photography equipment store (about 25 U.S. dollars) and we placed them in front of a 200 W white LED panel (about 140 U.S. dollars). Note that it was necessary to add an additional near-ultraviolet LED (about 0.50 U.S. dollars), otherwise the spectrum would not cover the wavelengths around 400 nm. The main difference between L1 and L2 is that L1 relies on multiple LEDs where each is electronically controlled, whereas in L2, the filters can be exchanged quickly by hand.

Examples of detectors (S1, S2) An optimal spectral detector in our case can differentiate between reflected light throughout the whole spectral range. The most straightforward option (denoted **S1**) is a *spectrometer*, which can easily have a sub-nanometer resolution and hundreds of spectral channels. An entry-level USB spectrometer (Ocean Insight USB-650) can be purchased for around 2000 U.S. dollars. A significantly more low-cost solution (denoted **S2**) relies on a

multi-spectral sensor, which is a small chip containing a few pixels, each with a different spectral sensitivity. In our example (measured in Fig. 6), we used the 8 narrowband channels of the AMS AS7341 sensor (about 25 U.S. dollars).

Cost comparison In Table 1, we summarize the aforementioned rough estimated costs of the individual optical setups L1-S1, L1-S2, L2-S1, and L2-S2. We also compare these costs to the reference setup that is used for ground truth measurements, which includes a monochromator, a broadband Xenon light source including its power supply, and a spectrometer. Because the manufacturer of our reference setup does not exist anymore, we instead decided to use the publicly available price lists of Newport Corporation (newport.com, sections “CS130B 1/8m Configured Monochromator” and “Low Power Xenon (Xe) Research Light Sources”) to estimate the prices for the monochromator and light source.

Table 1. Rough estimated costs for the individual optical setups.

	Reference setup	L1-S1	L1-S2	L2-S1	L2-S2
Illuminants	Monochromator: 8,500 USD	4 USD	4 USD	165 USD	165 USD
	Xe light: 10,000 USD				
Detector	2,000 USD	2,000 USD	25 USD	2,000 USD	25 USD
Total	20,500 USD	2,004 USD	29 USD	2,165 USD	190 USD

4.2. Synthetic dataset and simulated measurements

For evaluating fluorescence estimation methods, it is common to simulate measurements on synthetic datasets of Donaldson matrices [14,16,21], i.e., instead of performing an actual measurement, the optical setup is only simulated to evaluate the method’s performance. Such an approach has many benefits: the ground truth matrices are precisely known, the simulations are efficient, and the synthetic datasets can contain many materials, so the evaluation is statistically meaningful compared to real measurements on only a few isolated samples.

Base dataset We use a dataset of real materials that were measured by Gonzales [29], who directly acquired the Donaldson matrices using a specialized bi-spectral spectrophotometer (Labsphere BFC-450), which has two monochromators measuring the ground truth as shown by Donaldson [12] (Sec. 2.1). These ground truth matrices are available in a 10 nm resolution and they cover our range of interest of 400 to 700 nm. As the matrices were acquired with a real setup, they already contain significant noise, hence no synthetic noise was added. We picked 32 samples with strong fluorescence, including color pigments, polymer clays, golf balls, textiles, papers, and index cards, all showing various reflectances and levels of fluorescence. Out of these samples, 28 contained a single fluorophore, and we judged 4 to have two fluorophores.

Enlarging the dataset To enlarge the dataset, we further created 56 synthetic Donaldson matrices containing two fluorophores. First, we sorted the original 28 matrices based on their emission peaks (to prevent entirely overlapping spectra), and then we summed the Donaldson matrices while preserving the original diagonals (reflectances). Physically, this would be roughly equivalent to overlaying two transparent fluorescent slides with different fluorophores, which was done by Blasinski et al. [21].

Simulating the measurements Because we know the parameters of the optical setups from Sec. 4.1, mainly the illuminant and sensor spectra I_j, s_k (Fig. 6), and we know the ground truth matrix Φ from the synthetic dataset, we can use Eq. (3) to compute $I_{j,k}$ (the reflectances that the sensor detects). These values are then used as the sparse inputs to the actual fluorescence estimation method. The method then outputs the estimated matrix $\hat{\Phi}$. The similarity between $\hat{\Phi}$ and Φ is then evaluated (Secs. 4.4 and 4.5). A perfect estimation method would output $\hat{\Phi} = \Phi$. The data and the algorithm for these simulations are part of the attachment Code 1, Ref. [26].

4.3. Real measurements

Apart from the synthetic measurements, we also performed real measurements. Specifically, we measured fluorescent papers of different colors and fluorophores. To obtain ground truth measurements of the same materials, we relied on a tunable monochromator by AMKO GmbH, with a spectral range from 350 nm to 780 nm. We illuminated the sample in 10 nm steps, while the reflected light was being measured by a spectrometer. From these data, the ground truth Donaldson matrix was reconstructed [12]. When obtaining the ground truth matrices, the material needs to be illuminated under the same measurement geometry as when obtaining the sparse measurements. We solved this by using optical fiber to direct the light from the monochromator and from the LEDs into the exact same spot on the material samples. Each measurement was then performed once with the monochromator (for the ground truth Donaldson matrix), and once with the LED spectra (for the Donaldson matrix estimated using our method). Our experiments confirmed that our method works in real conditions, and we discuss the accuracy of the estimated matrices of real and synthetic data together in Sec. 4.5.

4.4. Quantitative evaluation

We evaluate the accuracy of our method on the various low-cost optical setups from Sec. 4.1. We focus on the synthetic dataset as it allows for drawing statistically meaningful conclusions. The exact same inputs were also evaluated using `fiToolbox` [28], an open-source implementation of the state-of-the-art method by Blasinski et al. [21]. This allows us to directly compare the estimated matrices between our method, the state of the art, and ground truth.

Root-mean-square error A commonly used quantitative metric is the *root-mean-square error* (RMSE) of the estimated matrices compared to the ground truths. Since the diagonal Φ_{diag} and fluorescence Φ_{fluo} can be separated, we evaluate their errors independently, which gives more insight into the algorithms' behaviors. We compare the RMSE on all combinations of the exemplar optical setups from Sec. 4.1, namely L1-S1, L1-S2, L2-S1, and L2-S2, which vary by their illuminants (L) and sensors (S). The resulting errors are visualized in Fig. 7 using standard boxplots, which show the medians, quartiles, minimum and maximum errors, and outliers, all separately for the four optical setups, and for one and two fluorophores.

Following Fig. 7, we can conclude that our algorithm reaches *lower errors* than the state of the art, and it is also more stable with *fewer outliers*. The biggest difference is visible in the L1-S1 setup on the single fluorophore dataset, where our fluorescence median error is *three times lower* than Blasinski et al. Notice that our algorithm's error significantly decreases with better optical setups (L1-S1 vs. L2-S2, the worst setup), whereas Blasinski et al. reaches similar errors in most setups. In the original publication by Blasinski et al. [21], Fig. 4, we can see that for a small number of illuminants, increasing the sensor resolution (camera filters) above some point does not improve the accuracy anymore, whereas our algorithm still takes advantage of the extra resolution to lower the estimation error. The only test in which our algorithm underperformed was the most difficult L2-S2 setup on the two-fluorophores dataset, where Blasinski et al. benefitted from their pre-learned basis functions.

Color accuracy Another important aspect is the *color accuracy under narrowband illumination* provided by our estimates (Fig. 8) to ensure that our method is suitable for predictive rendering of such scenarios in computer graphics. Given the estimated Donaldson matrices, we evaluated the error in predicting the color that is reflected from the fluorescent materials when illuminated by varying monochromatic illuminants. We calculated the color difference using the standardized perceptual metric CIEDE2000 (denoted ΔE_{00} , [30]). Following Fig. 8, our color predictions are consistently and significantly more accurate than in the estimates of Blasinski et al. The highest differences can be seen for low wavelengths, because most fluorophores in our dataset

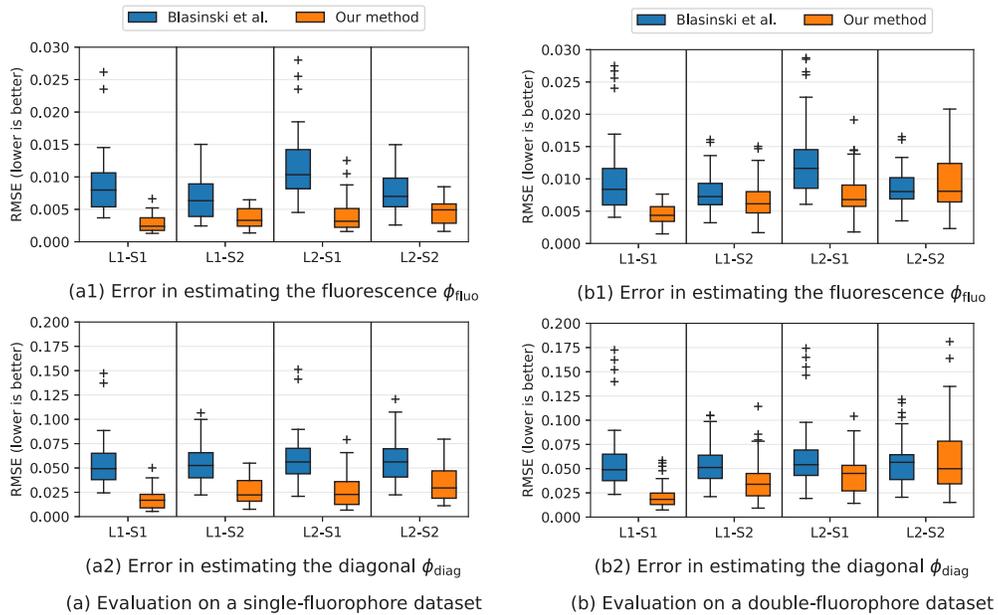


Fig. 7. Quantitative evaluation comparing the root mean square errors (RMSE) of Donaldson matrices estimated with our method (orange) and the state of the art (blue, Blasinski et al. [21]). The boxes show quartiles, horizontal lines are median errors, whiskers minimum and maximum errors, and the plus symbols are outliers. The L1-S1, L1-S2, L2-S1, and L2-S2 correspond to different optical setups per Sec. 4.1.

are triggered with a blue illuminant. Towards red wavelengths, the reflected color becomes independent of the fluorescence, it is given mostly by the pure reflectance (diagonal), and hence the overall color difference is lower.

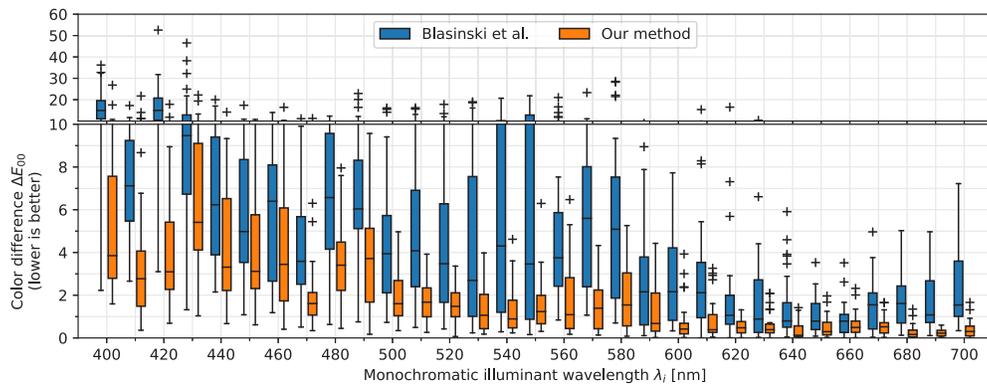


Fig. 8. Quantitative evaluation measuring the color difference (CIEDE2000, ΔE_{00} [30], lower is better) between the ground truth materials and our estimates (orange boxplots), and the state-of-the-art estimates (blue boxplots, Blasinski et al. [21]). The vertical axis shows how much the predicted reflected color is wrong, given the material was monochromatically illuminated with the wavelength on the horizontal axis. The dataset in this plot is the single-fluorophore, L1-S1 optical setup.

4.5. Qualitative evaluation

While the previous section focused on objective quantitative metrics using the synthetic dataset, we now perform a qualitative examination of all our data to understand the behavior of the estimation methods in various contexts.

Visualized matrices In Figs. 9 and 10, we show a small subset of ground truth Donaldson matrices and their estimates using our algorithm and state of the art. We chose 2 examples from the synthetic one-fluorophore dataset, 1 example from the synthetic two-fluorophores dataset, and 2 examples from real measurements. While one can see that the RMSE is typically lower for our estimates compared to Blasinski et al., the major observation is that the estimates of Blasinski et al. suffer from the aforementioned oscillations.

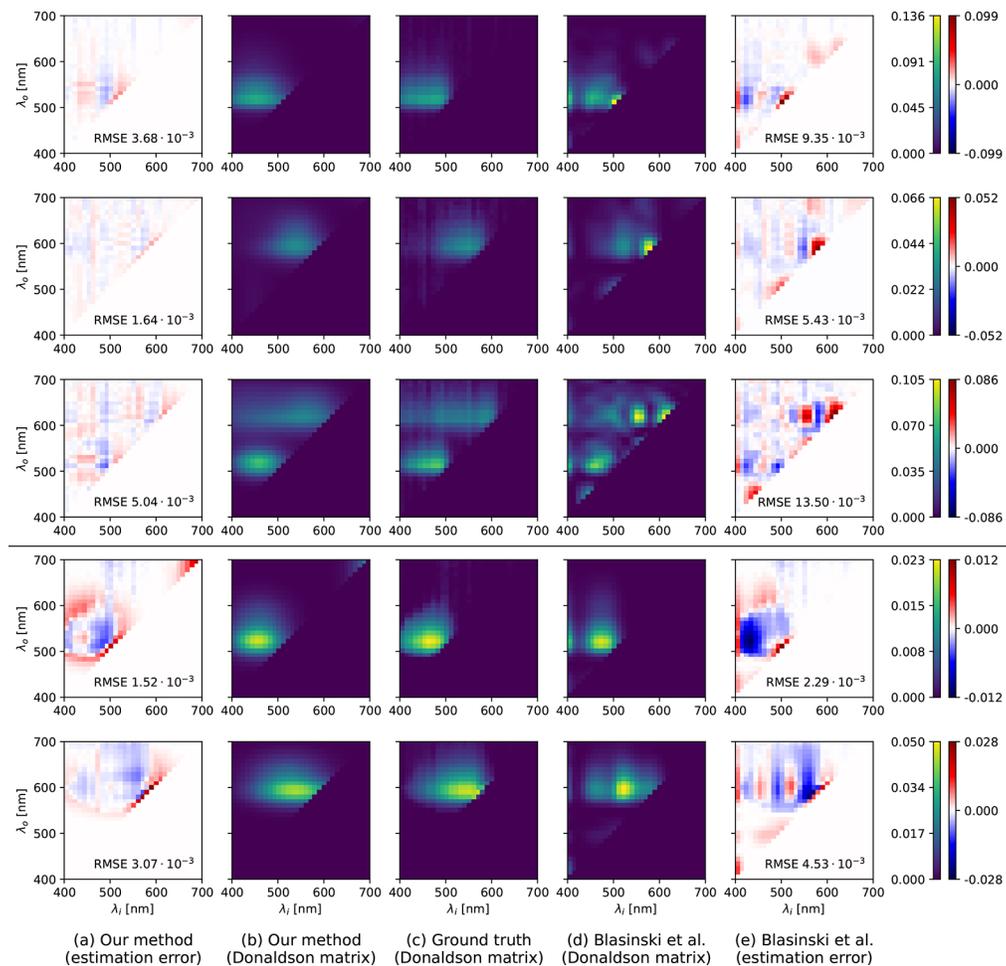


Fig. 9. Examples of Donaldson matrices of five different materials and their estimates using the L1-S1 optical setup, comparing our method (a,b) to ground truth (c) and the state of the art (d,e). The top three measurements are synthetic, the bottom two are real.

For example, in Fig. 9, all the ground truth matrices have either a single fluorescent peak or two peaks. While our estimates manage to reconstruct the number of peaks with a high accuracy, Blasinski et al. often overestimate the number of peaks, and the Donaldson matrices then resemble ripples on the water surface. In Fig. 10, we can see a similar effect in the reflectance

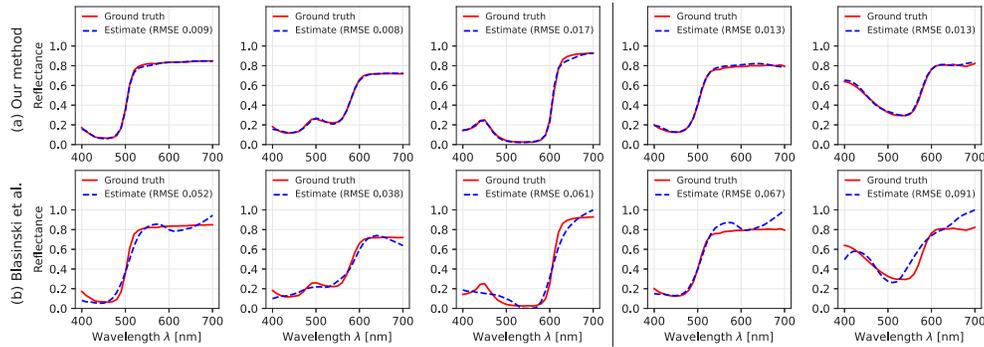


Fig. 10. Examples of ground truth and estimated pure reflectances (diagonals) of the five materials from Fig. 9 under identical conditions, comparing our method (a) to the state of the art (b). The left three measurements are synthetic, the right two are real.

spectra: where the ground truth and our estimates have a relatively flat spectrum, the estimates of Blasinski et al. often oscillate and create waves, which are not present in the ground truth spectra.

Using the estimates in predictive rendering Furthermore, we also present an example of using our estimated Donaldson matrices in the Monte Carlo predictive renderer ART [31], which is open-source and supports fluorescent materials. In Fig. 11, we show a comparison of six images rendered with ground truth matrices, and matrices estimated with our method and by Blasinski et al. [21] for two fluorescent materials lit by a monochromatic illuminant at 560 nm. This is an illustration of Fig. 8: the importance of having a consistently low Δ_E color error over the illuminants to ensure high color accuracy. The renders have been tonemapped with the exact same parameters to ensure that the exposure and contrast do not vary from result to result.

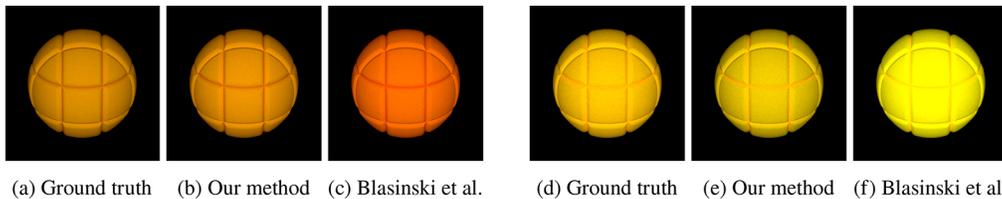


Fig. 11. A fluorescent ball is monochromatically illuminated (560 nm). The ball in each rendered image has a different Donaldson matrix. The balls (a) and (d) use the ground truth matrices from the material dataset, and the balls (b)-(c) and (e)-(f) use matrices estimated with our method, and with Blasinski et al. [21].

5. Conclusion

We have presented a simple and affordable technique for estimating fluorescence in materials. The inputs for our algorithm are easily acquirable with low-cost optical setups, and the output is in the form of a parametrized Donaldson matrix with a small memory footprint, usable in any standard application including predictive and photorealistic rendering. We showed that our estimates have objectively lower errors than the state-of-the-art algorithm on the exact same input data, our algorithm is also more stable and robust with a lower number of outliers, and our estimates are qualitatively better and more suitable for color-accurate predictive rendering.

Our paper's core principle was the parametrization of the Donaldson matrix by a Gaussian mixture model (GMM) and bounded maximum entropy spectral estimate (bounded MESE), which can be done in a differentiable way compatible with gradient-descent optimization algorithms.

This makes it compatible with sparse measurements and makes it possible to measure fluorescence accurately without specialized and expensive bi-spectral spectrophotometers.

Future work We can see several minor directions for future work. One improvement would consist of a faster implementation that does not rely on auto-differentiation, and instead computes derivatives using equations derived by hand, possibly accelerated on the GPU. Furthermore, the optimization itself could consist of more adaptive steps, such that only one Gaussian distribution is estimated at first, and more Gaussians are added subsequently to allow faster convergence for materials with only one fluorophore. Finally, our method allows anyone to quickly build a dataset of Donaldson matrices of common materials, which would be useful in computer graphics for predictive and photorealistic rendering.

Funding. Grantová Agentura České Republiky (GAČR-22-22875S); Horizon 2020 Framework Programme (956585); Univerzita Karlova v Praze (SVV-260699).

Disclosures. The authors declare no conflicts of interest.

Data availability. Data underlying the results presented in this paper are available in [Code 1](#), Ref. [26], along with Python source codes used to generate the reported values.

References

1. G. Johnson and M. Fairchild, "Full-spectral color calculations in realistic image synthesis," *IEEE Comput. Graph. Appl.* **19**(4), 47–53 (1999).
2. M. B. Hullin, J. Hanika, B. Ajdin, H.-P. Seidel, J. Kautz, and H. P. A. Lensch, "Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions," *ACM Trans. Graph.* **29**(4), 1–7 (2010).
3. M. Mojzík, A. Fichet, and A. Wilkie, "Handling Fluorescence in a Uni-directional Spectral Path Tracer," *Comput. Graph. Forum* **37**(4), 77–94 (2018).
4. A. Jung, A. Wilkie, J. Hanika, W. Jakob, and C. Dachsbacher, "Wide Gamut Spectral Upsampling with Fluorescence," *Comput. Graph. Forum* **38**, 1–10 (2019).
5. A. Jung, J. Hanika, and C. Dachsbacher, "Spectral Mollification for Bidirectional Fluorescence," *Comput. Graph. Forum* **39**(2), 373–384 (2020).
6. L. König, A. Jung, and C. Dachsbacher, "Improving Spectral Upsampling with Fluorescence," (2020).
7. Q. Hua, V. Tázlar, A. Fichet, and A. Wilkie, "Efficient Storage and Importance Sampling for Fluorescent Reflectance," *Computer Graphics Forum* (2022).
8. G. H. Mohammed, R. Colombo, and E. M. Middleton, *et al.*, "Remote sensing of solar-induced chlorophyll fluorescence (SIF) in vegetation: 50 years of progress," *Remote. Sens. Environ.* **231**, 111177 (2019).
9. P. L. Choyke and H. Kobayashi, "Medical Uses of Fluorescence Imaging: Bringing Disease to Light," *IEEE J. Sel. Top. Quantum Electron.* **18**(3), 1140–1146 (2012).
10. S. Delalieux, A. Auwerkerken, W. W. Verstraeten, B. Somers, R. Valcke, S. Lhermitte, J. Keulemans, and P. Coppin, "Hyperspectral Reflectance and Fluorescence Imaging to Detect Scab Induced Stress in Apple Leaves," *Remote Sens.* **1**(4), 858–874 (2009).
11. M. S. Roth and D. D. Deheyn, "Effects of cold stress and heat stress on coral fluorescence in reef-building corals," *Sci. Rep.* **3**(1), 1421 (2013).
12. R. Donaldson, "Spectrophotometry of fluorescent pigments," *Br. J. Appl. Phys.* **5**(6), 210–214 (1954).
13. J. Suo, L. Bian, F. Chen, and Q. Dai, "Bispectral coding: compressive and high-quality acquisition of fluorescence and reflectance," *Opt. Express* **22**(2), 1697–1712 (2014).
14. Y. Fu, A. Lam, I. Sato, T. Okabe, and Y. Sato, "Separating Reflective and Fluorescent Components Using High Frequency Illumination in the Spectral Domain," in *2013 IEEE International Conference on Computer Vision*, (2013), pp. 457–464. ISSN: 2380–7504.
15. C. Zhang and I. Sato, "Separating reflective and fluorescent components of an image," in *CVPR 2011*, (2011), pp. 185–192. ISSN: 1063–6919.
16. S. Tominaga, K. Hirai, and T. Horiuchi, "Estimation of bispectral Donaldson matrices of fluorescent objects by using two illuminant projections," *J. Opt. Soc. Am. A* **32**(6), 1068–1078 (2015).
17. Y. Zheng, I. Sato, and Y. Sato, "Spectra Estimation of Fluorescent and Reflective Scenes by Using Ordinary Illuminants," in *Computer Vision – ECCV 2014*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, eds. (Springer International Publishing, Cham, 2014), Lecture Notes in Computer Science, pp. 188–202.
18. Y. Zheng, Y. Fu, A. Lam, I. Sato, and Y. Sato, "Separating Fluorescent and Reflective Components by Using a Single Hyperspectral Image," in *2015 IEEE International Conference on Computer Vision (ICCV)*, (2015), pp. 3523–3531. ISSN: 2380–7504.
19. A. Lam and I. Sato, "Spectral Modeling and Relighting of Reflective-Fluorescent Scenes," in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, (2013), pp. 1452–1459. ISSN: 1063–6919.
20. Y. Fu, A. Lam, I. Sato, T. Okabe, and Y. Sato, "Reflectance and Fluorescence Spectral Recovery via Actively Lit RGB Images," *IEEE Trans. Pattern Anal. Mach. Intell.* **38**(7), 1313–1326 (2016).

21. H. Blasinski, J. Farrell, and B. Wandell, "Simultaneous Surface Reflectance and Fluorescence Spectra Estimation," *IEEE Trans. on Image Process.* **29**, 8791–8804 (2020).
22. C. Peters, S. Merzbach, J. Hanika, and C. Dachsbacher, "Using moments to represent bounded signals for spectral rendering," *ACM Trans. Graph.* **38**(4), 1–14 (2019).
23. J. P. Burg, "Maximum Entropy Spectral Analysis," PhD Thesis, Stanford University (1975).
24. A. Antoniou and W.-S. Lu, *Practical Optimization: Algorithms and Engineering Applications*, Texts in Computer Science (Springer US, New York, NY, 2021).
25. D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *ArXiv*, ArXiv:1412.6980 (2017).
26. T. Iser, L. Lachiver, and A. Wilkie, "Source code and data for the fluorescence measurement method using Gaussian distributions and bounded MESE," figshare (2023), <https://doi.org/10.6084/m9.figshare.23567073>.
27. J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, "JAX: composable transformations of Python+NumPy programs," (2018).
28. H. Blasinski, "Fluorescence Imaging Toolbox for Image Systems Engineering Toolbox (fiToolbox for ISET)," (2020). GitHub repository: <https://github.com/hblasins/fiToolbox>.
29. S. Gonzalez, "Evaluation of bispectral spectrophotometry for accurate colorimetry of printing materials," Ph.D. thesis, Rochester Institute of Technology (2000).
30. G. Sharma, W. Wu, and E. N. Dalal, "The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations," *Color Res. Appl.* **30**(1), 21–30 (2005).
31. A. Wilkie and R. F. Tobler, "ART: The Advanced Rendering Toolkit," (2022). URL: <https://cgg.mff.cuni.cz/ART/>.