Workshop on Material Appearance Modeling (2020)
H. Rushmeier and R. Klein (Editors)

Improving Spectral Upsampling with Fluorescence

Lars König[1] and Alisa Jung[1] and Carsten Dachsbacher[1]

[1]Karlsruhe Institute of Technology (KIT)

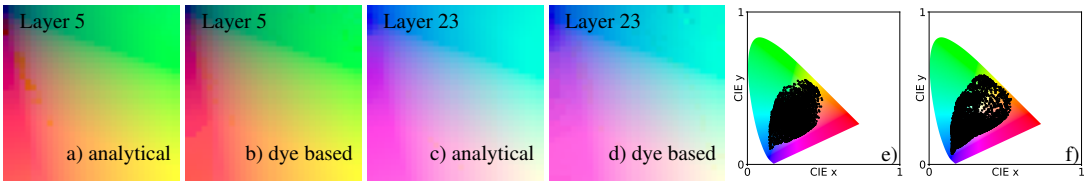

Figure 1: *Comparison of the analytical (a, c, e) and dye-based (b, d, f) fluorescence model without interpolation. The cube grid layers for $B = 0.161$ (a, b) and $B = 0.742$ (c, d) reveal artifacts in the dye-based model. e, f): Chromaticity of all colors represented in a 32^3 coefficient cube with a maximum slope of 0.01, only showing colors with brightness of $X + Y + Z \geq 1.5$.*

Abstract

Modern photorealistic rendering simulates spectral behaviour of light. Since many assets are still created in different RGB color spaces, spectral upsampling of the RGB colors to a spectral representation is required to use them in a spectral renderer. Limiting the upsampled spectra to physically valid and natural, i.e. smooth, spectra results in a more realistic image, but decreases the size of the gamut of colors that can be recreated.
In order to upsample wide gamut color spaces with colors outside the gamut of physically valid reflectance spectra, a previous approach added fluorescence to create accurate and physically valid representations. We extend this approach to increase the realism and accuarcy while considering memory and computation time.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

1. Introduction

Considering the spectral properties of light is crucial for physically accurate light transport simulation. Due to the physically motivated limitations of spectral representations, the gamut of colors of individual reflectance spectra is limited. To support a wider gamut of representable colors, while remaining physically accurate, *fluorescence* can be included. Building on the models of [JH19] and [JWH*19], this work examines possible improvements regarding the realism, accuracy and quality of *spectral upsampling with fluorescence*.

2. Related Work

In 1935, MacAdam [Mac35] discovered that the brightest and most saturated colors possible are represented by box-shaped reflectance spectra. Following work aimed at reproducing the smoother behaviour of physically realistic spectra. We build on the approach by [JH19] who use an analytical function space for reflectances r:

$$\begin{aligned} r(\lambda) &= S(c_0\lambda^2 + c_1\lambda + c_2) \\ S(x) &= 0.5 + x/2\sqrt{(1+x^2)} \end{aligned} \tag{1}$$

with three paramters (c_0, c_1, c_2) that are optimized to match the resulting spectrum to a target color.

As all previous methods create physically valid reflectances bound between 0 and 1, they can only reproduce colors inside the gamut of solid reflectances limited by MacAdam's box spectra, but fail at more saturated colors near the border of wide-gamut RGB spaces such as ACEScg. [JWH*19] extend [JH19] by adding fluorescence for such colors. They add three coefficients to analytically model flu-

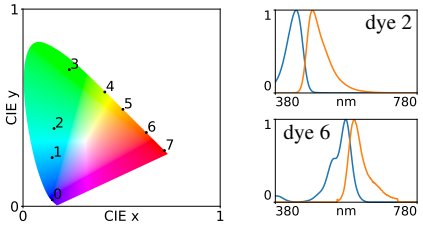

Figure 2: *Left: Chromaticity of 8 fluorescent dyes without non-fluorescent reflection. (Dyes 1 and 2: [AT20], others: [Sci]). Right: absorption (blue) and emission of dye 2 and 6.*

orescent absorption and emission spectra as b-splines with emission peak wavelength λ_e, stokes shift s and a parameter $c \in [0,1]$ for mixing fluorescence and reflectance.

This paper extends the work of [JWH*19]. We modify the parameterization (c_0, c_1, c_2) of the underlying function space (equation 1) for non-fluorescent reflectances to improve interpolation (section 4.3), replace the analytical fluorescence model (λ_e, s, c) with a set of real-world fluorescence measurements (section 4.2) and analyze how the color error metric affects the optimization process (section 4.1).

3. Background

Color Difference Metrics. When working with colors, we sometimes need to measure their difference. For colors in an RGB or the CIE XYZ color space, a simple metric is the Euclidean distance in this space. However, such distances do not relate to the human perception of color difference. An alternative is the CIE ΔE^*_{76} color distance metric defined as the Euclidean distance in the CIE LAB color space, which was designed to reflect color differences as perceived by the human eye. Still, the CIE ΔE^*_{76} metric has non-uniformities in some parts of the visible spectrum. The CIE ΔE^*_{00} metric [SWD05] defined in the CIE LCh color space resolves this.

Fluorescence. The phenomenon of fluroescence describes light that is absorbed and then instantaneously emitted at different wavelengths [Lak13]. The distribution of absorbed and emitted wavelengths is described by the absorption and emission spectrum a and e, examples for which are shown in Figure 2. These spectra can be combined with a non-fluorescent reflectance spectrum r to construct the fluorescent BRDF used by [JWH*19]. They model a and e analytically based on b-splines, which we replace with measurements of fluorescent dyes in section 4.2, while for r we use the sigmoid function space proposed by [JH19] (equation 1).

4. Our Work

Optimizing the parameters to fit the model's appearance to a given RGB triple is relatively slow, so we precompute them in a lookup-table based on a 3D regular grid inside the RGB color cube as in [JWH*19] using the Ceres solver [AMO].

Figure 3: *Part of the blue $= 0.7$ layer of the coefficient cube. Left: based on the Euclidean distance in the RGB color space, 32 iterations. Right: based on the ΔE^*_{00} color distance metric, 4 iterations.*

4.1. Perception Based Color Distance Metric

In order to optimize parameters (e.g. $[c_0, c_1, c_2, \lambda_e, s, c]$ in [JWH*19]) to minimize the difference between a target color and the color represented by them, we need a meaningful metric for this difference. The original approach uses the Euclidean distance in RGB space, which approaches zero as the colors become more similar, but ignores the perceived error.

We employ the ΔE^*_{76} and ΔE^*_{00} color difference metrics to aid the optimization in convergence for colors that can be represented perfectly, and in finding better approximations for colors that cannot be represented by the analytical model.

4.2. Fluorescent Dyes

The objective of spectral rendering is to generate a realistic image, so fluorescence should be modeled as realistically as possible. We therefore replace the analytical model in [JWH*19] with measurements from fluorescent substances. We chose eight fluorescent dyes from [AT20] and [Sci] with absorption peak wavelengths ranging from 390 to 633 nm and sort them based on their emission spectrum's chromaticity as shown in Figure 2.

Our model combines two adjacent dyes with an interpolation weight $w \in [0,1]$ to achieve a greater range of colors. As before, we need three coefficients to model fluorescence: The mixing coefficient c, interpolation weight w and index idx of the first dye. If memory is an issue this model may be easier to compress as idx can only take on eight values.

For simplicity we interpolate the absorption and emission spectra of the two dyes. Thus the combined absorbed energy is reradiated according to the combined emission spectra. It would be more accurate to calculate absorption and emission for both dyes separately and interpolate the result.

One drawback of this model is a more complicated optimization, since the index idx can only take on discrete values. We use the brute-force solution of iterating over all eight pairs of adjacent dyes and optimize the remaining coefficients to find the best fitting pair and coefficients.

Another challenge is interpolation between colors, as adjacent entries in the precomputed coefficient cube can have different dye indices. A possible solution would be a BRDF that includes all fluorescent dyes referenced by the coefficients of all corners of a cube grid cell. Using this approach would allow to use only three fluorescence coefficients per grid point, while still being able to interpolate them.

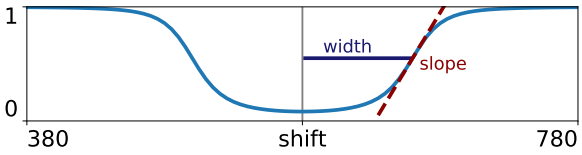

Figure 4: *Parabola features shift, width and slope.*

4.3. Alternative Interpolation

As it is infeasible to precompute coefficients for every single color in an RGB color space, we need to interpolate the upsampling model inbetween the precomputed grid points. It is thus important to define coefficients in a way that allows for a meaningful interpolation, such that the result is as close as possible to a true interpolation of the corresponding spectra.

This is particularly important at the border between coefficients for spectra with and without fluorescence, where the reflectance spectra differ more. The reflectance model (equation 1) is highly non-linear for linearly changing coefficients c_0, c_1, c_2. We present a new parameterization for this function space to improve interpolation in such cases (Figure 5).

We use three features of a reflectance spectrum as coefficients: the *shift* (i.e. the center wavelength), the *maximum slope* and the *width* of the parabola measured between the maximum slope and shift as shown in Figure 4. Note that the maximum slope value is defined as the maximum absolute slope of a parabola, but *slope* itself is positive or negative to indicate a u-shaped or n-shaped parabola respectively.

These coefficients are related to certain color features. The shift corresponds to the chromaticity. The width and the slope simultaneously affect saturation and brightness, which corresponds to the fact that physically valid reflectance spectra cannot appear arbitrarily bright *and* saturated. The color's brightness also depends on the shift, as the human eye is less sensitive to wavelengths near the border of the visible range.

To create a spectrum from the three features *shift*, *width* and *slope* (z, w and s), we use algorithm 1 to transform them to coefficients (c_0, c_1, c_2) before evaluating equation 1.

Algorithm 1 (*shift* z, *width* w, *slope* s) $\rightarrow (c_0, c_1, c_2)$

$$t \leftarrow (|s| \cdot w + (s^2 \cdot w^2 + 1/9)^{1/2})/(2 \cdot |s| \cdot w)$$
$$c_0 \leftarrow s \cdot t^{3/2}/w$$
$$c_1 \leftarrow -2 \cdot c_0 \cdot z$$
$$c_2 \leftarrow c_0 \cdot z^2 + s \cdot w \cdot (5 \cdot t^{3/2} - 6 \cdot t^{1/2})$$

Challenges. This model struggles with representing unsaturated, gray colors. We can specify colors that are perceived as gray by creating an almost horizontal line using a large width and small slope. However, a given gray color can in fact be represented by infinitely many combinations of *width* and inversely scaled *slope* values, while varying the *shift* causes little noticeable effect. This can cause artifacts when interpolating between these extreme values (see Figure 8).

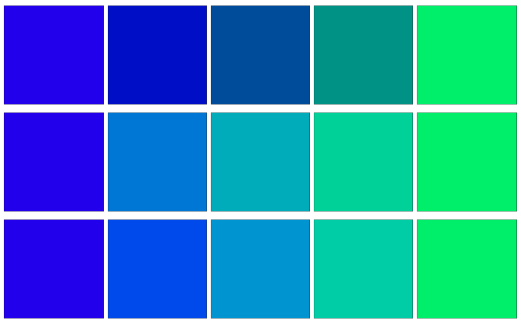

Figure 5: *Linear interpolation of the parabola coefficients (top), the color spectrum itself (center) and the feature based coefficients (bottom) with weights* 0, 0.25, 0.5, 0.75, 1.

Another problem is the distinction between u-shaped and n-shaped spectra, which are caused by positive and negative *slope* values respectively, as spectra with slope values near 0 behave differently depending on the sign. For typical widths, small positive slopes produce bright colors, while small negative slopes produces dark colors. This causes artifacts when interpolating from a u-shaped to an n-shaped parabola.

We avoid this issue with a hybrid interpolation approach, where we transform our feature based coefficients to the original ones (c_0, c_1, c_2) before interpolating such cases.

5. Evaluation

In this chapter we visualize layers of an optimized 32^3-coefficient grid in the ACEScg color space, based on the technique proposed by [JWH*19]. We use the ACES primaries #1 (AP1) for optimization and analysis, but interpret the resulting triples as sRGB for visualization. Reflectance spectra are limited to smooth spectra with a maximum slope of 0.01. Ultraviolet light is included for the optimization of the fluorescent coefficients. The cube grid layer visualizations represent a layer of the cube grid for a fixed value of the blue component. We use the CIE D65 illuminant as light source and the CIE 1931 2 degree standard observer functions. For visualization purposes we plot fluorescent emission spectra (Figure 2, orange) scaled to 1.

5.1. Perception-Based Color Distance Metric

Using the ΔE^*_{76} color distance metric leads to a significantly smoother color transition through the coefficient cube as shown in Figures 3 and 6, but also introduces artifacts in the red/green part of the coefficient cube when used with too few iterations. Using the computation heavier ΔE^*_{00} color

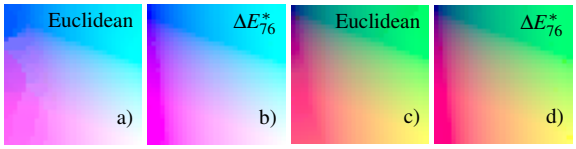

Figure 6: *Cube layers 31 ($B = 1$; a, b) and 8 ($B = 0.258$; c, d) for the Euclidean distance in RGB (a, c) and the ΔE^*_{76} metric (b, d), both with 4 iterations.*

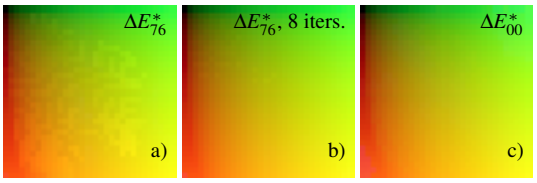

Figure 7: *Cube layer 0 ($B = 0$) for the ΔE^*_{76} metric and 4 iterations (a) has artifacts in the red/green region, which mostly disappear with 8 iterations (b) or the ΔE^*_{00} metric (c).*

distance metric or increasing the number of iterations which propagate the coefficients to adjacent grid points in the coefficient cube removes the artifacts as shown in Figure 7.

Although both ΔE^* metrics require an additional step to convert the color from the initial RGB color space to the CIE L*a*b* color space, overall convergence of the optimization was faster, decreasing from 5h 59m (Euclidean distance in RGB) to 3h 18m (ΔE^*_{76}) for otherwise equal parameters and 4 iterations. Figure 7 shows that remaining artifacts are removed best by switching to the ΔE^*_{00} metric with 4 iterations in 3h 20min, though they also disappear in 5h 36min when increasing iterations for the ΔE^*_{76} metric from 4 to 8.

5.2. Fluorescent Dyes

Using real-world measurements improves the realism of the fluorescence model, although some approximations are still required in order to limit memory and computation resources. Figure 1 shows that with our eight dyes from Figure 2 we achieve a similar gamut as the analytical model by [JWH*19]. Note that the gamut could be of a different shape or size for a different set of fluorescent dyes.

Including this fluorescence model in the coefficient cube optimizer proposed by [JWH*19] shows promising results, though some artifacts remain and computation time is higher due to the brute force optimization over dye pairs. Figure 1 compares our result to the analytical fluorescence model.

5.3. Alternative Interpolation

Figure 5 shows that interpolating the feature based coefficients (shift, width and slope) produces more perceptually meaningful results than interpolating the original coefficients (c_0, c_1, c_2). Our approach works well with bright, saturated colors, as long as all colors have the same sign of the slope value, but produces artifacts for flipping slope signs.

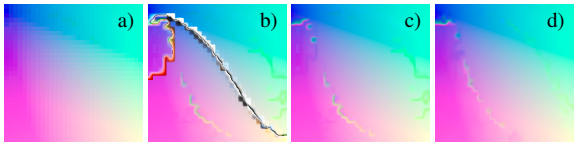

Figure 8: *Cube layer for $B = 0.7$ and feature based coefficients with no (a), feature based (b), hybrid for different slope signs (c) and old (c_0, c_1, c_2) (d) interpolation. Interpolating slopes with different signs leads to distortions.*

Figure 8 visualizes a slice of an optimized coefficient cube. As discussed in section 4.3, interpolating our coefficients leads to heavy artifacts, which can be avoided by the presented hybrid interpolation approach. Note that the discontinuities in the left image without interpolation are at the same position as the distortions in the right image.

6. Conclusion

We proposed several modifications to the approach of [JWH*19]. The ΔE^* color distance metrics improved the overall quality and convergence speed and were easy to implement. Using real-world measurements from fluorescent substances produced artifacts in the optimization process, but provides a more realistic model for simulating fluorescence in a spectral renderer. Our feature based coefficients for reflectance spectra interpolate colors more meaningfully, but require additional work to arrive at a stable solution.

References

[AMO] AGARWAL S., MIERLE K., OTHERS: Ceres solver. http://ceres-solver.org. 2

[AT20] ATTO-TEC: Fluorescence label - atto, 2020. URL: https://www.atto-tec.com/. 2

[JH19] JAKOB W., HANIKA J.: A low-dimensional function space for efficient spectral upsampling. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library. 1, 2

[JWH*19] JUNG A., WILKIE A., HANIKA J., JAKOB W., DACHSBACHER C.: Wide gamut spectral upsampling with fluorescence. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 87–96. 1, 2, 3, 4

[Lak13] LAKOWICZ J. R.: *Principles of fluorescence spectroscopy*, third edition ed. Springer Science & Business Media, 2013, ch. 1. 2

[Mac35] MACADAM D. L.: Maximum visual efficiency of colored materials. *J. Opt. Soc. Am. 25*, 11 (Nov. 1935), 361–367. doi:10.1364/JOSA.25.000361. 1

[Sci] SCIENTIFIC T. F.: Fluorescence spectraviewer. accessed 2020-02-29. URL: https://tinyurl.com/nrza7qq. 2

[SWD05] SHARMA G., WU W., DALAL E. N.: The ciede2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application 30*, 1 (2005), 21–30. 2