Controlled Spectral Uplifting for Indirect-Light-Metamerism

MARK VAN DE RUIT, Delft University of Technology, Netherlands
ELMAR EISEMANN, Delft University of Technology, Netherlands

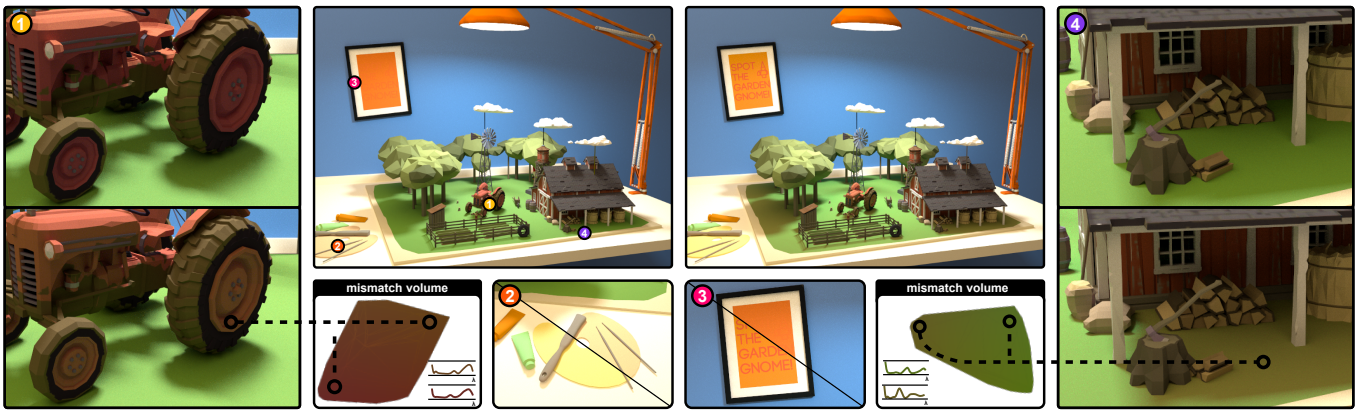

Fig. 1. Spectral renders of a scene lit by multiple illuminants ($D65$, $FL2$). We demonstrate several spectral modifications. *1, 4*; metamer mismatching under indirect illumination, for which mismatch volumes indicate valid color options at a surface. *2, 3;* constrained metamer mismatching under direct illumination.

Spectral rendering has received increasing attention in recent years. Yet, solutions to define spectral reflectances are mostly limited to uplifting techniques which deterministically augment existing RGB inputs. Only recently has uplifting been able to ensure a certain surface appearance under direct illuminants. Yet, prior work in this area limits artist expressiveness and is not well suited for designing the appearance of a scene, as indirect illumination is ignored entirely.

We present an uplifting technique with fine-grained spectral appearance control under direct and indirect illumination, even enabling the placement of spectral constraints in a specific scene. Our approach allows for a flexible authoring process, and solves for the resulting spectra efficiently. Additionally, we show that our method's memory overhead during rendering is kept small, by introducing a compact spectral texture format.

CCS Concepts: • **Computing methodologies** → **Reflectance modeling**.

Additional Key Words and Phrases: spectral rendering, reflectance spectra, spectral uplifting, convex hull, linear system, metamer mismatch volume

ACM Reference Format:
Mark van de Ruit and Elmar Eisemann. 2024. Controlled Spectral Uplifting for Indirect-Light-Metamerism. In *SIGGRAPH Asia 2024 Conference Papers (SA Conference Papers '24), December 03–06, 2024, Tokyo, Japan.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3680528.3687698

Authors' Contact Information: Mark van de Ruit, Delft University of Technology, Netherlands, m.vanderuit-1@tudelft.nl; Elmar Eisemann, Delft University of Technology, Netherlands, e.eisemann@tudelft.nl.

SA Conference Papers '24, December 03–06, 2024, Tokyo, Japan

ACM ISBN 979-8-4007-1131-2/24/12
https://doi.org/10.1145/3680528.3687698

1 Introduction

Spectral rendering has started taking a critical role in production. Trichromatic systems approximate radiance and reflectance as RGB values, but spectra are required for accurate color reproduction under different illuminants. However, spectral surface reflectances are challenging to acquire at the scales required for texture data and impose a memory impact during rendering. One option is spectral uplifting, which generates spectra from existing RGB inputs. This process is ill-posed due to *metamerism* (different reflectances producing the same color under a given light source). Uplifting typically establishes a one-to-one mapping, e.g., opting for the smoothest reflectance, and hereby inherits problems of trichromatic rendering, as a produced reflectance might exhibit a wanted appearance under one illuminant only. Therefore, working in a full spectral pipeline is challenging.

Recent work introduced constrained uplifting [Tódová et al. 2021; van de Ruit and Eisemann 2023], focusing on appearance under various illuminants. Yet, in practice, there are limitations, including performance impacts, limited control, and ignored interreflections, which can significantly impact scene appearance (Fig. 1).

Our spectral uplifting extends prior work, while being simple and efficient. It supports reproducing spectral measurements, direct-illumination color constraints, and constraints on color appearance under complex indirect illumination. We uplift input RGB data via a $\mathcal{R}^3$ convex polytope englobing the data with specific spectra on each vertex. Its interior is tessellated into simplices, and interior vertices can be added for fine-grained control. We retrieve the reflectance of an RGB value by localizing the surrounding simplex and performing interpolation between the spectra of the simplex' vertices. We extend prior work on metamer mismatch volumes [Finlayson and Morovic 2005; Logvinenko et al. 2013; Mackiewicz et al. 2019], incorporating estimated light transport to handle indirect illumination. Using a

small basis, we "bake" the uplifting into an efficient texture format, which supports local constraints in parts of the scene, and is practical for rendering. Finally, we make our source code available online[1].

In short, we make the following contributions:

- A controllable spectral uplifting method via constraints;
- A solution for metamer color control under indirect light;
- A compact spectral texture format for rendering.

In the following, we cover related material (Sec. 2) and our method (Sec. 3). We evaluate roundtrip error and authoring (Sec. 4), discuss findings (Sec. 5), and conclude (Sec. 6).

2 Background

Color Theory. While light spectra consist of various wavelengths, human color perception builds mainly upon three sensors. In consequence, colors are typically equated to a sensor trio observing the spectral distribution of light [Wyszecki and Stiles 1982], which are described by observer functions. Denoting three observer functions as e, and given an illuminant distribution i, we express the response to surface reflectance r under direct illumination as

$$\Phi(r) = \int_{\Lambda} e(\lambda)\, r(\lambda)\, i(\lambda)\, d\lambda, \tag{1}$$

where Λ describes typically the visible spectrum (e.g. $360 - 830$ nm). This is formalized in the CIE XYZ standard observers and derived color spaces such as sRGB. A *color system* combining observer and illuminant describes a linear transformation $\Phi(r) : \mathcal{X} \rightarrow \mathcal{R}^3$. Of particular interest is the valid region of responses $\{r \in \mathcal{X} \mid \Phi(r) \neq 0\}$, which forms a convex region called the *object color solid.*

The illuminant i describes radiant energy over wavelength, while reflectance r describes a surface's efficacy in reflecting said energy. Illuminants are positively unbounded and vary in shapes dependent on the underlying processes. Reflectances are $[0, 1]$-bounded and are generally low-banded in the visible spectrum. This property holds for most pigments, but not for structural colors [Maloney 1986]. We, as most related work, restrict ourselves to smooth reflectances.

Metamer Mismatching. We briefly cover metamerism and metamer mismatching; for extended overviews, please refer to [Finlayson and Morovic 2005; Logvinenko et al. 2013]. Given some color system, consider the problem of inverting Equation 1, i.e. $\Phi^{-1}(\Phi(r)) = r$. It is ill-posed; as it is typically a set of reflectances:

$$\Phi^{-1}(\Phi(r)) = \{\, r' \in \mathcal{X} \mid \Phi(r) = \Phi(r') \,\},$$

i.e. many reflectances achieve a particular color signal. This *metamer set* is convex. Metamers produce the same signal under Φ, but in a secondary color system Ψ (differing in observer or illuminant) they can produce different responses. Mapping the set to Ψ, we find a non-singleton color solid called a *metamer mismatch region.* We describe the sampling-based method of Mackiewicz et al. [2019] to find mismatch region boundaries, as we extend this method to incorporate indirect illumination (Sec. 3.3).

Consider mapping $\mathcal{T} : \mathcal{X} \rightarrow \mathcal{R}^6$, $\mathcal{T}(r) = (\phi, \psi)$, where $\phi = \Phi(r)$, $\psi = \Psi(r)$ form color signals in two color systems. For given signal ϕ, the set of *mismatched signals* under Ψ is a cross-section of $\mathcal{T}$, i.e.:

$$\mathcal{M}(\phi, \Phi, \Psi) = \{\, \psi \in R^3 \mid (\phi, \psi) \in \mathcal{T} \,\},$$

where we simplify notations by identifying $\mathcal{T}$ with its image. While this region's interior can be complex, boundary spectra are unique step-wise functions consisting only of transitions between zeroes and ones. These *optimal spectra* are extrema [Logvinenko 2009; Schrödinger 1920]. Finding the mismatch boundary $\delta\mathcal{M}$ thus reduces to extremizing spectra mapped to $\mathcal{T}$, subject to $\Phi(r) = \phi$. In practice, one spherically samples unit vector $\hat{u} \in \mathcal{R}^6$, and projects the color system spectra in $\mathcal{T}$ along $\hat{u}$. In the discrete case, reflectance r is represented as a k-dimensional vector, Φ, Ψ are $k \times 3$ color system matrices, and we can solve for the boundary with a linear program:

$$\max_{r \in \Phi^{-1}(\phi)} \left(\begin{bmatrix} \Phi \\ \Psi \end{bmatrix} \hat{u} \right)^{\mathsf{T}} r \tag{2}$$

This produces a discrete reflectance on the boundary $\delta\mathcal{M}(\phi, \Phi, \Psi)$.

Spectral Uplifting. A recent overview of the scope of spectral uplifting is given by Weidlich et al. [2021]. Uplifting is having a known color signal ϕ and then finding a metamer $r \in \Phi^{-1}(\phi)$. Across different methods, three criteria are identifiable: smoothness of produced reflectances, boundedness, and roundtrip error.

The earliest approaches to spectral uplifting are now of mostly theoretical interest, as they produce blocky distributions or break boundedness [MacAdam 1935; Smits 1999]. The later approach of Meng et al. [2015] precomputes spectra on a grid spanning the xy-chromaticity plane, uplifting colors through interpolation of these spectra. It produces smooth reflectances, but requires scaling for values above or below the plane, introducing errors. Otsu et al. [2018] address this by clustering sets of spectra into a spatial hierarchy covering the xy-chromaticty plane. Inside clusters, the authors store localized bases for spectrum recovery. While efficient, this approach introduces discontinuities between clusters, uplifting gradients poorly. More recently, van de Ruit and Eisemann [2023] precompute spectra on a convex polytope enclosing an input texture, ensuring correct interpolation, if such a polytope can be found.

A separate class of techniques started with Jakob and Hanika [2019], who demonstrate a low-dimensional parameterization of a sigmoidal function space, for which they precompute coefficients across a three-dimensional color lookup table. Subsequent function reconstruction is inherently fast, and produces smooth reflectances with minimal roundtrip error. This approach has been extended to handle out-of-gamut spectra [Jung et al. 2019; König et al. 2020]. Afterwards, a more complex, Fourier-space approach is introduced by Peters et al. [2019], addressing roundtrip issues of the sigmoidal.

As previously stated, most methods establish one-to-one mappings, associating a specific metamer with a color signal. Two techniques differ in this aspect. Tódová et al [2021; 2022] extend the Fourier-space representation [Peters et al. 2019], seeding that method's coefficient table s.t. specific measured spectra are reproduced during uplifting. Van de Ruit and Eisemann [2023] uplift based on artist-specified color constraints. While their approach enables targeting of specific color behavior, control is limited. We detail these limitations later (Sec. 3.1), as we extend their work. Our solution is the first to integrate indirect illumination constraints.

[1] https://graphics.tudelft.nl/indirect_uplifting

3 Method

We now present our constrained spectral uplifting. We first cover our method's foundation (Sec. 3.1), followed by the basic uplifting procedure (Sec. 3.2). We then derive a color system to constrain the uplifting under indirect illumination (Sec. 3.3), and finally specify a practical texture format for rendering (Sec. 3.4).

3.1 Foundation

A color system (Equation 1) describes a linear transformation, preserving convexity. Let $r_1, \ldots, r_n$ be n reflectances with corresponding mappings $\Phi(r_1), \ldots, \Phi(r_n)$. If we combine these mappings using convex weights $w_1, \ldots, w_n \;:\; \forall_i w_i \geq 0 \cap \sum_i w_i = 1$, we observe:

$$\sum w_i \Phi(r_i) = \Phi\left(\sum w_i r_i\right), \tag{3}$$

i.e., the color signal of linearly combined reflectances equals the linear combination of the corresponding color signals. When interpolating two metamers, the result itself is a metamer - but it also holds for arbitrary reflectances. This principle is employed in most prior work. Van de Ruit and Eisemann [2023] note that, minimally, interpolation of $\mathcal{R}^3$ color signals must occur within a 3-simplex, as planar methods unavoidably struggle with error [Mallett and Yuksel 2019; Meng et al. 2015; Otsu et al. 2018]. However, instead of a simplex, they solve for a complex polytope enclosing an input RGB texture. Uplifting then reduces to a linear mixture of the polytope's uplifted vertices. The authors enable targeted uplifting by constraining the vertex spectra, but this is limited in effect; vertices necessarily lie on the polytope, away from the input, and therefore constraints never affect or modify the input's uplifting directly. In contrast, our method allows for interior constraints, which can directly affect the uplifting of specific RGB inputs.

3.2 Tessellated Color System

Convex reconstruction of a polytope interior avoids roundtrip error, which is why we also use this as our foundation. However, we select a polytope that describes the color-system boundary, which encloses all color inputs. Given the polytope vertices, we can construct a Delaunay tessellation, which results in a set of 3-simplices, each associated with four vertices. Each vertex contains a spectrum producing the color encoded by the vertex position under the illuminant of the color system. We detail spectrum generation further below. Uplifting a color input then reduces to (a) localizing the enclosing simplex for the color within the tessellation, and (b) a barycentric interpolation of the simplex' associated spectra (Equation 3). Further, we can constrain the interior by inserting vertices into the tessellation, which then affect the uplifting. Fig. 2 shows an overview of the uplifting procedure.

Boundary spectra. Reflectances on the color-system boundary are extrema, being the MacAdam limits. Interpolation between such saturated distributions results in correct but physically implausible spectra. We therefore find a smaller boundary formed by smooth spectra as outlined below.

Like prior methods, we use a weighted PCA basis to construct low-banded spectra [Cohen 1964; Fairman and Brill 2004; Otsu et al. 2018; Parkkinen et al. 1989; Tzeng and Berns 2005; van de Ruit and Eisemann 2023]. We use a dataset of $\sim 41M$ measured reflectances ($400 - 700$ nm) from a variety of materials [Zhang et al. 2016]. Let reflectance r be a discretized k-dimensional vector. We apply PCA to the dataset and retain the first m principal components as a $k \times m$ matrix B. Then we represent a reflectance as $r = Bw$, where w are m-dimensional coefficients. Though $m = 3$ suffices to reproduce most colors [Cohen 1964], we retain $m > 3$ as in prior work [van de Ruit and Eisemann 2023]. Otherwise, if the system were fully determined, we would eliminate the ability to output metamers.

To generate a boundary within this basis, we employ the method of Mackiewicz et al. [2019]. We spherically sample a unit vector $\hat{u} \in \mathcal{R}^3$, and project the $k \times 3$ discretized color system Φ onto $\hat{u}$. Maximizing this projection necessarily results in a position on the system boundary. Expressed in our basis, this becomes:

$$\max_w (\Phi\hat{u})^\top Bw, \quad \text{with } \forall_i 0 \leq (Bw)_i \leq 1. \tag{4}$$

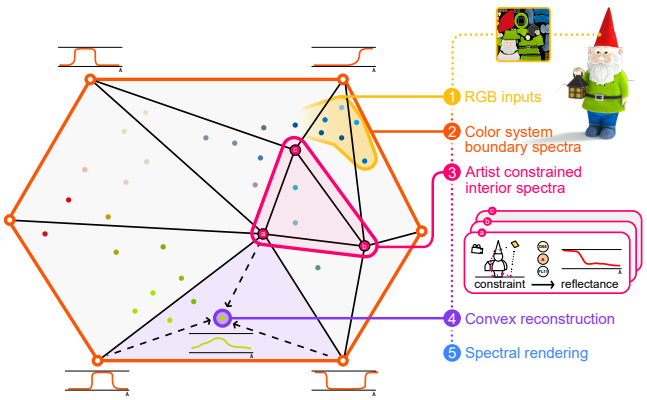

Fig. 2. Uplifting prodecure. We treat input colors as $\mathcal{R}^3$ positions and sample an enclosing boundary on the input's color system, while artist-specified constraints then specify spectra in the interior. Boundary and interior spectra are connected using a Delaunay tessellation, and inputs are then uplifted using barycentric interpolations inside the tessellation's simplices.

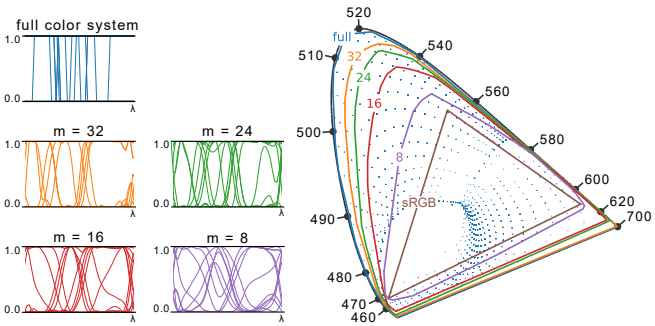

Fig. 3. Color system support. We show sample boundary spectra and the effective support of our basis with $m = 32, 24, 16, 8$ components. We further show support of a discrete system ($k = 64$) with sampled MacAdam limits, as well as the sRGB gamut, which is enclosed for all choices of m.

As PCA acts as a dimensionality reduction, resultant spectra become low-banded, enabling their use for interpolation. Fig. 3 illustrates color system support and boundary spectra for choices of m.

Interior spectra. In the polytope interior, we allow user-specified constraints. Each provides a color signal - in effect a vertex position - and an associated reflectance, uplifted or directly specified. We detail three constraint types across this and prior work.

1. Measurement: Given a concrete spectral reflectance r, we project it into the PCA basis, which gives a w, minimizing $||Bw - r||$, and then insert vertex $\Phi(Bw)$. Similar to Tódová et al. [2021], we can reproduce the representations of spectral measurements. We investigate full spectral texture reproduction in Sec. 4.2.

2. Direct color: Given n color constraints $\{\psi_1, \ldots, \psi_n\}$ under color systems $\{\Psi_1, \ldots, \Psi_n\}$, the linear program

$$\min_{w} \; ||Bw||, \text{ with } \forall_j \Psi_j Bw = \psi_j \text{ and } \forall_i 0 \leq (Bw)_i \leq 1 \qquad (5)$$

produces coefficients for a constraint-satisfying metamer, which we insert as vertex $\Phi(Bw)$. Following van de Ruit and Eisemann [2023], we restrict inputs to the intersection of relevant mismatch regions.

3. Indirect color: Given a scene with observer, and an observed surface position with reflectance r, we enable constraining observed color ψ at this position under an indirect color system $\Psi(r)$. We derive this color system in the following (Sec. 3.3).

3.3 Indirect Color System

While prior constraint types enable uplifting control, they only account for mismatching in a linear form. Yet, in complex scenes, illuminant-induced mismatching can occur due to complex light transport (Fig. 4). To control this effect, we derive a convex formulation of a non-linear color system, w.r.t. a constraint reflectance r at some surface position in a scene. In the following, we estimate the indirect light transport from this point and then factor out r. We then use this factorization to formulate a maximization that finds the metamer mismatch boundary under indirect illumination.

Path-integral formulation. For an environment with a particular constrained reflectance r at a given surface point, we define a color system over a per-wavelength incident radiance I_r and observer e:

$$\Psi(r) = \int_\Lambda e(\lambda)\, I_r(\lambda)\, d\lambda. \qquad (6)$$

The radiance measure can be expressed using a path-integral formulation of light transport [Veach and Guibas 1995]:

$$I_r(\lambda) = \int_\Omega f_r(\bar{x}, \lambda)\, d\mu\bar{x} \qquad (7)$$

where μ forms a measure over samples in path domain Ω, and $f_r : (\Omega, \Lambda) \to \mathcal{R}$ denotes the per-wavelength measurement contribution along a path of length n, $\bar{x} = \{x_1, \ldots, x_n\}$. This contribution describes light throughput (geometric terms, cosine attenuation, bidirectional reflectance, illuminant) along the path. Let $\{r_1, \ldots, r_{n-1}\}$ be underlying reflectances at path vertices not on a light source. We factor out these reflectances using a secondary function $f_0 : \Omega \to \mathcal{R}$,

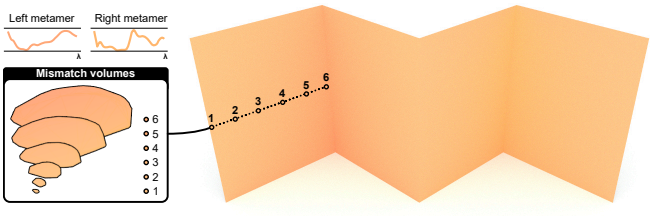

Fig. 4. Two folded patches on a white plane under $D65$. While metameric in direct light, these concave patches differ due to interreflections. At equidistant points on a patch, we show generated indirect mismatch volumes.

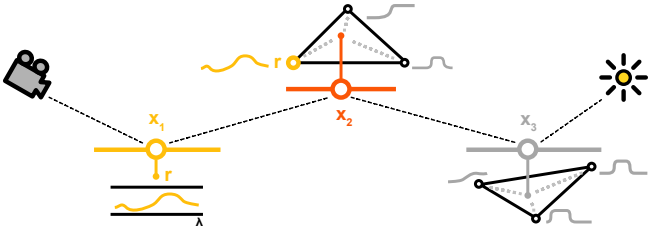

Fig. 5. Formulation. We show a light path $\{x_1, x_2, x_3\}$, and constraint reflectance r at surface x_1. Further reflectances are expressed as convex combinations of spectra in simplices from our tessellation, which can contain r.

which is independent of reflectances ($\forall_{r \in \mathcal{X}}\, r(\lambda) = 1$). We express

$$f_r(\bar{x}, \lambda) = f_0(\bar{x}) \prod_{i=1}^{n-1} r_i(\lambda), \qquad (8)$$

which implies that BRDFs have to allow for associative rearrangement of surface reflectances.

Following this, we can express path reflectance r_i w.r.t. constraint reflectance r. Recall that uplifting is the convex combination of four spectra, associated with vertices of a 3-simplex. As r is one vertex in our tessellated color system, each encountered spectrum r_i on surfaces along the path either uses r or not. Fig. 5 illustrates an example. Next, we denote simplex spectra and associated convex weights as $(s_1, a_1), \ldots, (s_4, a_4)$, such that $r_i = \sum_{j=1}^{4} s_j a_j$. Then:

$$r_i = \begin{cases} a_k r + \sum_{j \neq k}^{4} s_j a_j & (\exists_k\, s_k = r) \\ 0r \;\; + \sum_{j=1}^{4} s_j a_j & (else) \end{cases} \quad \to \quad r_i = a_i r + w_i$$

where $a_i \geq 0$ is either zero or a convex weight, and w_i sums the weighted remainder spectra. This enables expansion of Equation 8:

$$\begin{aligned} f_0(\bar{x}) \prod_{i=1}^{n-1} (a_i r + w_i) = f_0(\bar{x}) (\;& r^{n-1}(a_1 a_2 \cdots a_{n-1}) \\ + \;& r^{n-2}(w_1 a_2 \cdots a_{n-1} + a_1 w_2 \cdots a_{n-1}) \\ + \;& r^{n-3}(w_1 w_2 a_3 \cdots a_{n-1} + \cdots) \\ & \vdots \\ + \;& r^{0}(w_1 w_2 \cdots w_{n-1})) \end{aligned} \qquad (9)$$

For a power $b < n$, we define t_b as the coefficient of r^b in Equation 9. The contribution then simplifies to a truncated power series:

$$f_r(\bar{x}, \lambda) = f_0(\bar{x}) \sum_{b=0}^{n-1} t_b(\bar{x}, \lambda)\ r^b(\lambda)$$

We can now revisit the path-integral formulation (Equation 7). Given the complexity of light-transport problems, we apply Monte Carlo integration, introducing the estimator

$$I_r(\lambda) \approx \hat{I}_r(\lambda) = \frac{1}{N} \sum_{i=1}^{N} \frac{f_r(\bar{x}_i, \lambda)}{p(\bar{x}_i)},$$

converging to the expected value as $N \to \infty$. Here $p(\bar{x}_i)$ is the probability density of sample $\bar{x}_i$. Given a known, finite length n across paths - noted as approximation - the estimator becomes

$$\begin{aligned}
\hat{I}_r(\lambda) &= \frac{1}{N} \sum_{i=1}^{N} \frac{f_0(\bar{x}_i)}{p(\bar{x}_i)} \sum_{b=0}^{n-1} t_b(\bar{x}_i, \lambda)\ r^b(\lambda) \\
&= \sum_{b=0}^{n-1} \left(\frac{1}{N} \sum_{i=1}^{N} \frac{f_0(\bar{x}_i)\ t_b(\bar{x}_i, \lambda)}{p(\bar{x}_i)} \right) r^b(\lambda) \qquad (10) \\
&= \sum_{b=0}^{n-1} c_b(\lambda)\ r^b(\lambda) \ \text{, where } \ c_b(\lambda) := \frac{1}{N} \sum_{i=1}^{N} \frac{f_0(\bar{x}_i)\ t_b(\bar{x}_i, \lambda)}{p(\bar{x}_i)},
\end{aligned}$$

producing a simplified expression. In practice, we can estimate coefficient spectra $c_0, \ldots, c_{n-1}$ by accumulating incident radiance along N paths and factoring out b interreflections of constraint reflectance r. Given these coefficients, we next employ Equation 10 to express light transport under mismatching constraint reflectances. This follows prior numerical approximations of interreflections treating surfaces as a finite number of patches [Nayar et al. 1991].

Mismatch volume boundary. We employ Equation 10 to specify an estimate of the indirect color system (Equation 6), which becomes

$$\Psi(r) \approx \hat{\Psi}(r) = \int_\Lambda e(\lambda)\ \hat{I}_i(\lambda)\ d\lambda = \int_\Lambda e(\lambda) \sum_{b=0}^{n-1} (c_b(\lambda)\ r^b(\lambda))\ d\lambda,$$

which forms a non-linear system. We build upon the method of Mackiewicz et al. [2019] to determine mismatch boundary $\delta\mathcal{M}$. First, we discretize the color systems ($3 \times k$), specifically the uplifting's color system Φ with known signal ϕ, and define discretized indirect color system spectra as $\tilde{\Psi}_0, \ldots, \tilde{\Psi}_{n-1}$, where $\Psi_b = e \circ c_b$; here $\circ$ denotes a component-wise multiplication. We then spherically sample a unit vector $\hat{u} \in \mathcal{R}^6$, along which we project color system spectra. As with the linear form (Sec. 2), maximizing in the direction of this projection results in a position on the region's boundary. We denote $\hat{u} := (\alpha, \beta)$, where $\alpha, \beta \in \mathcal{R}^3$, forming

$$\max_{r \in \Phi^{-1}(\phi)} (\Phi\alpha)^\mathsf{T} r + \sum_{b=0}^{n-1} (\tilde{\Psi}_b \beta)^\mathsf{T} r^b \qquad (11)$$

which adopted to our basis becomes

$$\begin{aligned}
&\max_w\ (\Phi\alpha)^\mathsf{T}(Bw) + \sum_{b=0}^{n-1} (\tilde{\Psi}_b \beta)^\mathsf{T} (Bw)^b \\
&\text{with } \ \Phi^\mathsf{T}(Bw) = \phi \text{ and } \forall_i 0 \le (Bw)_i \le 1.
\end{aligned}$$

We show mismatch volumes generated over a patch with varying interreflections in Fig. 4. Note that this problem is convex as the solution spans r, which is $[0, 1]$-bounded, while $b \ge 0$ and $\forall_b c_b \ge 0$. As our basis uses negative components, convexity does not hold. In practice, we may find mismatch boundaries that are interior to the exact boundary, which is sufficient in practice.

Indirect reflectance generation. We next generate a metamer in the mismatch volume. Given indirect color constraint ψ and discretized color system spectra $\hat{\Psi}_1, \ldots, \hat{\Psi}_{n-1}$, we extend Equation 5, solving:

$$\begin{aligned}
&\min_w \|Bw\|,\ with\ \forall_i 0 \le (Bw)_i \le 1 \\
&\text{with } \phi = \Phi^T(Bw) \text{ and } \psi = \textstyle\sum_b^{n-1} \hat{\Psi}_b^T\ (Bw)^b.
\end{aligned}$$

Note that, as the mismatch volume we present to the user to specify ϕ is a color solid, we can avoid this minimization. Instead, we tessellate this solid in $\mathcal{R}^3$, localize ψ's enclosing simplex within the tessellation, and perform barycentric interpolation of the simplex' vertex coefficients. This w encodes the wanted metamer Bw.

3.4 Texture format

Our uplifting is a reconstruction, localized to a simplex in the tessellation. A direct approach to storing an uplifted input, is the index to its enclosing simplex, together with three interior convex weights, as the fourth can be deduced ($\sum a_i = 1$). Given that weights are in $[0, 1]$ and the number of simplices in the tessellation is low, a low-bit RGBA texture typically suffices. However, users may define conflicting constraints, as we discuss in Sec. 5, which implies additional operations during rendering.

Instead, we preprocess and "bake" our reflectances, representing each reflectance in the orthonormal basis B. We already rely on coefficients in this basis throughout our entire pipeline, including the reflectances stored in vertices of the tessellation. Hence, to encode an uplifted reflectance, we find the enclosing simplex and use the interpolated basis coefficients. This representation is compact and conversion is easily parallelized. We scale our basis to a $[-1, 1]$ boundary on the projection of PCA inputs; recovered coefficients are thus $[0, 1]$-bounded. This enables a fixed-point representation; in practice, we store 128 bits per pixel, packing 8, 12 or 16 coefficients at 16, 10 or 8 bits respectively. Note that texture filtering is applied to coefficients after unpacking; the filtered result is then used for uplifting. We evaluate variants of our representation in Sec. 4.1.

4 Evaluation

In the following, we discuss implementation, evaluate uplifting quality (Sec. 4.1) and spectral texture recovery (Sec. 4.2). Afterwards, we demonstrate the indirect color system (Sec. 4.3).

Our implementation relies on a sequential quadratic programming [Kraft 1994] algorithm in the *NLopt* framework [Johnson 2007] for the constrained optimization and *Qhull* [Barber et al. 2013] for the Delaunay tessellations. Uplifting and rendering use *OpenGL*, and we employ continuous wavelength hero sampling [Wilkie et al. 2014]. During uplifting, discrete spectra use $k = 64$ bins to handle high-frequency illuminants. For color solid sampling, we use 128 spherical samples in all cases. Our method supports any spectral range, but uses $400 - 700$ nm due to the dataset underlying our basis.

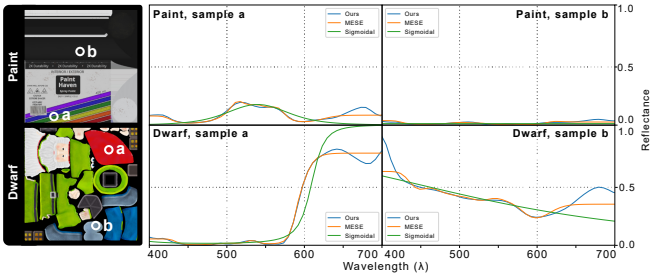

Fig. 6. Unconstrained reflectances. We show sample spectra from unconstrained uplifting (*Fig.* 11) for our method ($m = 12$), the *sigmoidal* [Jakob and Hanika 2019] and the bounded *MESE* [Peters et al. 2019]. All spectra accurately reproduce the RGB input under $D65$.

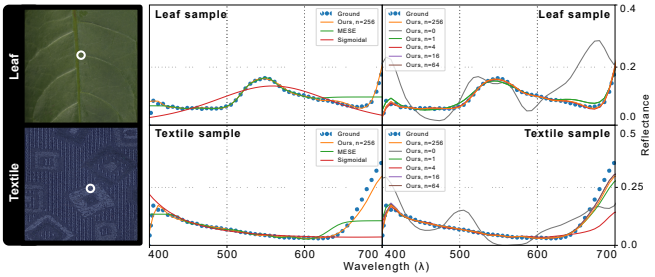

Fig. 7. Constrained reflectances. We show sample spectra from constrained spectral texture reconstruction (Fig. 12) for our method ($m = 12$), the *sigmoidal* [Jakob and Hanika 2019] and the bounded *MESE* [Peters et al. 2019]. Our uplifting is constrained with n random samples from the input texture.

A supplemental video shows real-time spectral modifications on a *RTX 3070*. When a user first places an indirect constraint in a scene, we trace *65K* paths GPU-side, and reduce to a power series CPU-side, which takes under a second. Texture baking occurs whenever a relevant constraint is edited, and takes 12.5 ms for a 4K texture.

Throughout the evaluation, we compare with the *sigmoidal* uplifting [Jakob and Hanika 2019] and the Fourier-space bounded *MESE* [Peters et al. 2019]. For the latter, we select $m = 12$ coefficients (11 moments). This method has high accuracy on RGB roundtrip with fewer coefficients; however, it proved insufficient to reproduce complex metamers on mismatch boundaries. Note that the method's reconstruction time is quadratic w.r.t. the choice of m. In the authors' implementation, we use their *mirrored, warped* configuration.

We evaluate color accuracy for the different methods using CIE ΔE_{00} color difference [Sharma et al. 2005]. Note the following thresholds: $\Delta E \leq 1$ implies no perceptible difference; $\Delta E \in (1, 2]$ implies a very close match; $\Delta E > 2$ implies close but visible mismatching.

4.1 Reconstruction of RGB data

We first evaluate unconstrained uplifting of RGB input data. We measure color reproduction of *BabelColor Average* patches [BabelColor Company 2019] under $D65$. Prior methods achieve sufficient accuracy here; ideally we produce the same roundtrip error or less. For our method, we test bases with $m = 8, 12, 16$ principal components at 32 bit, and low-bit *packed* variants storing 16, 10, 8 bits per coefficient. For the bounded *MESE*, we likewise test a variant storing 10 bits per coefficient, using the same code as our method.

Fig. 10 shows roundtrip results for all methods and Table 1 lists mean and maximum ΔE_{00}. While roundtrip of full-precision variants of our method consistently improves on prior methods, none produces perceptible error. The exception is our method's $m = 16$ packed variant, which visibly mismatches darker colors due to the low bitrate. As packed variants are intended for practical rendering we discard $m = 16$ variants in the following.

Fig. 11 shows unconstrained uplifting of RGB textures, achieving similar error. All methods uplift correctly, though in packed variants our method avoids mismatching, while the bounded *MESE* struggles with darker colors. We further show uplifted spectra from these textures (Fig. 6); evidently, the basis can introduce an oscillating behavior near spectral range boundaries, compared to the smoother outputs of prior methods. This is explained by the use of warped coefficients, which was presented by the authors of bounded *MESE*; i.e. greater precision is given to the centre of the visible range.

4.2 Reconstruction of spectral data

We next demonstrate recovery of hyperspectral textures. Given their size, such data is typically impractical for rendering. We approximate the textures in a compact format. We repeat a prior experiment [van de Ruit and Eisemann 2023], fitting textures from the *HyTexila* dataset [Khan et al. 2018]; each 1024^2 texture stores 186 spectral channels over $400 - 1000$ nm ($744MB$). We then sample $n = 0, 1, 4, 16, 64, 256$ spectra from the texture, inserting these in our tessellation as *measurement* constraints. As n increases, we expect our method to increasingly resemble the input texture. We could fit all pixels into our basis directly instead. However, as spectral texture capture is challenging, we test whether a smaller input suffices. For the bounded *MESE*, we fit per pixel, though we note that the work of Tódová et al. [2021; 2022] enables a compact fitting of this method.

Fig. 12 compares outputs under standard illuminants $D65$, $FL11$, and $LED - RGB1$. All unconstrained methods correctly handle recovery under $D65$, but visibly mismatch under one of the other illuminants. Given any number of constraints, our method's output strongly improves. For $n \geq 16$, all outputs achieve mean $\Delta E_{00} \leq 1.2$, outperforming prior methods. We show several output spectra for all methods and sample counts in Fig. 7.

4.3 Reconstruction of indirect color constraints

Finally, we evaluate indirect color constraining (Sec. 3.3). We set up a simple scenario; a neutral-gray surface, illuminated by a constant $D65$ environment. The surface consists of flat and folded parts, perpendicular to an orthogonal camera. We constrain the flat part to simply reproduce the input RGB color. We then constrain the center of the fold, where interreflections occur and metameric mismatching is possible. We generate a mismatch volume for the constraint, and select constraint values that we subsequently apply to the surface.

Fig. 8 displays results; for each position, we render with the $m = 12$, *packed* variant of our method. All variants of the surface correctly reproduce the intended colors; roundtrip error remains below perceptible limits. Error on the flat patch (mean $\Delta E_{00} = 0.34$) is partially due to the low-bit representation used for rendering.

Table 1. Perceptual error metrics. We list mean and maximum CIE LAB ΔE_{00} for unconstrained RGB uplifting (Fig. 11). We compare output for our method ($m = 8, 12, 16$), low-bit variants, the *sigmoidal* [Jakob and Hanika 2019] and the bounded *MESE* [Peters et al. 2019].

	Ours (8)	Ours (12)	Ours (16)	Ours (8P)	Ours (12P)	Ours (16P)	Sigm.	MESE (12)	MESE (12P)
$\Delta E_{00}, \mu$	0.00046	0.00056	0.01024	0.00048	0.11862	0.60964	0.01467	0.05829	0.29709
ΔE_{00}, max	0.00588	0.00725	0.07291	0.01079	0.87168	1.89340	0.03181	0.17747	1.39676

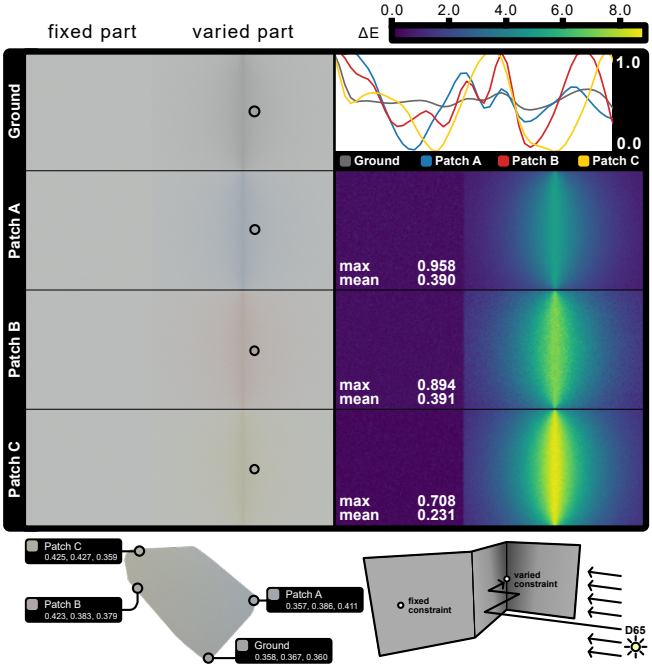

Fig. 8. Indirect color constraining. *Bottom-right.* Scene setup; flat and folded patches under constant $D65$. We constrain the surface at the indicated positions. *Bottom-left.* A mismatch volume showing potential colors under interreflections in the fold. We select four color constraints. *Top.* Renders with the constrained metamers. We show ΔE_{00} for the unchanged flat patch.

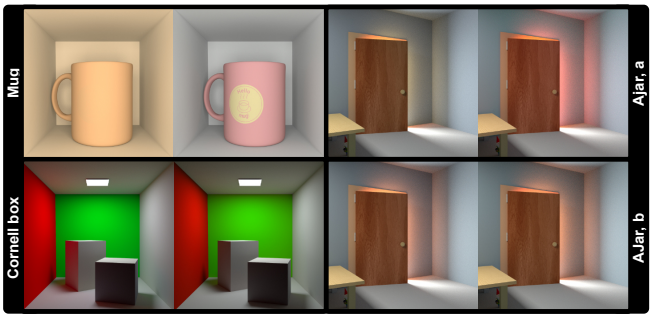

Fig. 9. *Examples. Mug.* We enforce metamerism under $FL2$, hiding the surface texture. *Cornell box.* We constrain the surface to cancel red scattering. *Ajar A.* We constrain wall reflectances, obtaining differing interreflections under $D65$. *Ajar B.* We constrain the door's wood texture.

5 Discussion

Reflectance reconstruction. Our method recovers metamers with excellent roundtrip under $D65$, outperforming or matching prior methods even in the ($m = 12$, packed) variant. While produced reflectances remain smooth, prior methods better target a low-banded output through their smaller basis or warped coefficients, though at the cost of potential mismatching under secondary illuminants. Specifically in Fig. 12, all unconstrained methods mismatch for one or more illuminants. We hence emphasize the necessity of constrainable uplifting, as even a single interior constraint allows us to recover smooth, correctly matching spectra. If no constraints are present, one can even insert smooth spectra into the interior.

Our theory supports specular materials but our implementation is for diffuse surfaces. Their view-independence makes constraint placement intuitive. Implementing general BRDFs is future work.

Scene constraint specification. In our implementation, users can specify constraints by directly clicking in the scene. We then sample the underlying surface color and construct an interior vertex in the tessellation. This process enables significant user control; Fig. 9 demonstrates example scene modifications.

However, as users edit different parts of a scene, they might define conflicting constraints, e.g., two metamers which share the same tessellation vertex and affect the same object. Multiple reflectances per vertex can be supported by using a texture atlas; the user can then guide the process via a brush metaphor. Texture masks encode weights for each constraint, which are blended during uplifting. While this implies that additional operations are needed during uplifting, it is irrelevant in practice because our compact texture encodes a basis representation, which has constant evaluation cost.

Basis function encoding. We employ $m = 12$ basis functions in most results, enabling constrained metamerism. To store basis coefficients in our compact texture, we scale the basis, enabling low-bit representation of $[-1, 1]$ coefficients. Instead of the current basis which encodes a full color system, we can compute a targeted basis for just the output spectra in a texture. This may allow us to improve the precision or size of memory-efficient variants of our texture. Nevertheless, having a different basis per scene does not allow us to easily merge scenes, which is of practical relevance in production.

6 Conclusion

Our solution enables practical, authorable spectral rendering. The tessellated uplifting improves upon prior solutions in expressivity, while remaining compact for rendering. Further, we introduced indirect illuminant constraints for uplifting, which can form a novel artistic tool for designing scene appearance. Our representation is efficient - our implementation reaches interactive frame rates, as

demonstrated in the supplementary video - while offering greater control over uplifting than prior work. Our work opens up new avenues for spectral material design.

In this work, we targeted accurate reflectance representations. In the future, we may focus on perceptual aspects. For one, current representations may be further compressed without visible difference. For another, our work shows that, despite the human visual system's limitations, we can differentiate high-dimensional reflectances in a scene context. Yet, during material design, this context is often lacking. Novel interfaces relying on our solution may one day alleviate this problem.

References

BabelColor Company. 2019. The Colorchecker Pages. Retrieved September 12 2024 from http://www.babelcolor.com/colorchecker.htm

C Bradford Barber, David P Dobkin, and Hannu Huhdanpaa. 2013. Qhull: Quickhull algorithm for computing the convex hull. *Astrophysics Source Code Library* (2013).

Jozef Cohen. 1964. Dependency of the spectral reflectance curves of the Munsell color chips. *Psychonomic science* 1, 1 (1964), 369–370.

Hugh S Fairman and Michael H Brill. 2004. The principal components of reflectances. *Color Research & Application* 29, 2 (2004), 104–110.

Graham D. Finlayson and Peter Morovic. 2005. Metamer sets. *J. Opt. Soc. Am. A* 22, 5 (2005), 810–819. https://doi.org/10.1364/JOSAA.22.000810

Wenzel Jakob and Johannes Hanika. 2019. A Low-Dimensional Function Space for Efficient Spectral Upsampling. *Computer Graphics Forum* (2019).

Steven G. Johnson. 2007. The NLopt nonlinear-optimization package. Retrieved September 12 2024 from https://github.com/stevengj/nlopt

A. Jung, A. Wilkie, J. Hanika, W. Jakob, and C. Dachsbacher. 2019. Wide Gamut Spectral Upsampling with Fluorescence. *Computer Graphics Forum* 38, 4 (2019), 87–96. https://doi.org/10.1111/cgf.13773

Haris Ahmad Khan, Sofiane Mihoubi, Benjamin Mathon, Jean-Baptiste Thomas, and Jon Yngve Hardeberg. 2018. HyTexiLa: High Resolution Visible and Near Infrared Hyperspectral Texture Images. *Sensors* 18, 7 (2018). https://doi.org/10.3390/s18072045

Dieter Kraft. 1994. Algorithm 733: TOMP–Fortran modules for optimal control calculations. *ACM Trans. Math. Software* 20 (1994), 262–281. https://doi.org/10.1145/192115.192124

Lars König, Alisa Jung, and Carsten Dachsbacher. 2020. Improving Spectral Upsampling with Fluorescence. In *Workshop on Material Appearance Modeling*. The Eurographics Association. https://doi.org/10.2312/mam.20201139

Alexander D. Logvinenko. 2009. An object-color space. *Journal of Vision* 9, 11 (2009), 5–5. https://doi.org/10.1167/9.11.5

Alexander D Logvinenko, Brian Funt, and Christoph Godau. 2013. Metamer mismatching. *IEEE Transactions on Image Processing* 23, 1 (2013), 34–43. https://doi.org/10.1109/TIP.2013.2283148

David L. MacAdam. 1935. Maximum Visual Efficiency of Colored Materials. *J. Opt. Soc. Am.* 25, 11 (1935), 361–367. https://doi.org/10.1364/JOSA.25.000361

Michal Mackiewicz, Hans Jakob Rivertz, and Graham Finlayson. 2019. Spherical sampling methods for the calculation of metamer mismatch volumes. *J. Opt. Soc. Am. A* 36, 1 (Jan 2019), 96–104. https://doi.org/10.1364/JOSAA.36.000096

Ian Mallett and Cem Yuksel. 2019. Spectral Primary Decomposition for Rendering with sRGB Reflectance. The Eurographics Association. https://doi.org/10.2312/sr.20191216

Laurence T. Maloney. 1986. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *J. Opt. Soc. Am. A* 3, 10 (1986), 1673–1683. https://doi.org/10.1364/JOSAA.3.001673

Johannes Meng, Florian Simon, Johannes Hanika, and Carsten Dachsbacher. 2015. Physically Meaningful Rendering using Tristimulus Colours. *Computer Graphics Forum* 34, 4 (2015), 31–40. https://doi.org/10.1111/cgf.12676

Shree K Nayar, Katsushi Ikeuchi, and Takeo Kanade. 1991. Shape from interreflections. *International Journal of Computer Vision* 6 (1991), 173–195.

H. Otsu, M. Yamamoto, and T. Hachisuka. 2018. Reproducing Spectral Reflectances From Tristimulus Colours. *Computer Graphics Forum* 37, 6 (2018), 370–381. https://doi.org/10.1111/cgf.13332

Jussi PS Parkkinen, Jarmo Hallikainen, and Timo Jaaskelainen. 1989. Characteristic spectra of Munsell colors. *JOSA A* 6, 2 (1989), 318–322.

Christoph Peters, Sebastian Merzbach, Johannes Hanika, and Carsten Dachsbacher. 2019. Using Moments to Represent Bounded Signals for Spectral Rendering. *ACM Trans. Graph.* 38, 4, Article 136 (2019), 14 pages. https://doi.org/10.1145/3306346.3322964

Erwin Schrödinger. 1920. Theorie der Pigmente von grösster Leuchtkraft. *Annal. Der Phys.* 62 (1920), 603–622.

Gaurav Sharma, Wencheng Wu, and Edul N. Dalal. 2005. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application* 30, 1 (2005), 21–30. https://doi.org/10.1002/col.20070

Brian Smits. 1999. An RGB-to-Spectrum Conversion for Reflectances. *Journal of Graphics Tools* 4, 4 (1999), 11–22. https://doi.org/10.1080/10867651.1999.10487511

Di-Yuan Tzeng and Roy S Berns. 2005. A review of principal component analysis and its applications to color technology. *Color Research & Applications* 30, 2 (2005), 84–98.

Lucia Tódová, Alexander Wilkie, and Luca Fascione. 2021. Moment-based Constrained Spectral Uplifting. In *Eurographics Symposium on Rendering*. https://doi.org/10.2312/sr.20211304

L. Tódová, A. Wilkie, and L. Fascione. 2022. Wide Gamut Moment-based Constrained Spectral Uplifting. *Computer Graphics Forum* 41, 6 (2022), 258–272. https://doi.org/10.1111/cgf.14617

Mark van de Ruit and Elmar Eisemann. 2023. Metameric: Spectral Uplifting via Controllable Color Constraints. In *ACM SIGGRAPH 2023 Conference Proceedings*. https://doi.org/10.1145/3588432.3591565

Eric Veach and Leonidas J Guibas. 1995. Optimally combining sampling techniques for Monte Carlo rendering. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*.

Andrea Weidlich, Alex Forsythe, Scott Dyer, Thomas Mansencal, Johannes Hanika, Alexander Wilkie, Luke Emrose, and Anders Langlands. 2021. Spectral Imaging in Production: Course Notes Siggraph 2021. In *ACM SIGGRAPH 2021 Courses (SIGGRAPH '21)*. Article 14, 90 pages. https://doi.org/10.1145/3450508.3464582

A. Wilkie, S. Nawaz, M. Droske, A. Weidlich, and J. Hanika. 2014. Hero Wavelength Spectral Sampling. *Computer Graphics Forum* 33, 4 (2014), 123–131. https://doi.org/10.1111/cgf.12419

G. Wyszecki and W. S. Stiles. 1982. *Color science: Concepts and Methods, Quantative Data and Formulae*. Wiley.

Xiandou Zhang, Brian Funt, and Hamidreza Mirzaei. 2016. Metamer mismatching in practice versus theory. *J. Opt. Soc. Am. A* 33, 3 (Mar 2016), A238–A247. https://doi.org/10.1364/JOSAA.33.00A238

Fig. 10. Unconstrained color uplifting. We uplift color patches [BabelColor Company 2019] and render under $D65$. *Left.* Output for our method ($m = 8, 12, 16$), low-bit variants, the sigmoidal [Jakob and Hanika 2019] and the bounded *MESE* [Peters et al. 2019]. *Right.* ΔE_{00} values are listed for all results.

Fig. 11. Unconstrained texture uplifting. We uplift RGB textures and render under $D65$. We show output for our method ($m = 12$ coefficients), a low-bit packed variant, the *sigmoidal* [Jakob and Hanika 2019] and bounded *MESE* [Peters et al. 2019]. Mean and maximum ΔE_{00} are listed for all outputs.

Fig. 12. Constrained texture recovery. We reproduce a spectral texture dataset [Khan et al. 2018] under three standard illuminants for the *sigmoidal* [Jakob and Hanika 2019], the bounded *MESE* [Peters et al. 2019], and our method ($m = 12$) using $n = 0, \ldots, 256$ constraints. Mean ΔE_{00} is listed for all outputs.