

Using Moments to Represent Bounded Signals for Spectral Rendering

CHRISTOPH PETERS, Karlsruhe Institute of Technology, Germany

SEBASTIAN MERZBACH, University of Bonn, Germany

JOHANNES HANIKA, Karlsruhe Institute of Technology and Weta Digital, Germany

CARSTEN DACHSBACHER, Karlsruhe Institute of Technology, Germany



Fig. 1. We store six moments per texel in 8 bytes to represent the spectral albedo textures in this scene. Our novel reconstruction method (green stripes) uses this data to approximate the actual reflectance spectrum. Thus, color reproduction is accurate, even with challenging illuminants such as the fluorescent lamp used here. The interleaving in the upper part, the magnified insets and the insets with distances in CIELAB show how strongly state of the art tristimulus techniques [Jakob and Hanika, 2019] (orange stripes) deviate from ground truth (black stripes).

We present a compact and efficient representation of spectra for accurate rendering using more than three dimensions. While tristimulus color spaces are sufficient for color display, a spectral renderer has to simulate light transport per wavelength. Consequently, emission spectra and surface albedos need to be known at each wavelength. It is practical to store dense samples for emission spectra but for albedo textures, the memory requirements of this approach are unreasonable. Prior works that approximate dense spectra from tristimulus data introduce strong errors under illuminants with sharp peaks and in indirect illumination. We represent spectra by an arbitrary number of Fourier coefficients. However, we do not use a common truncated Fourier series because its ringing could lead to albedos below zero or above one. Instead, we present a novel approach for reconstruction of bounded densities based on the theory of moments. The core of our technique is our bounded maximum entropy spectral estimate. It uses an efficient closed form to compute a smooth signal between zero and one that matches the given Fourier coefficients exactly. Still, a ground truth that localizes all of its mass around a few wavelengths can be reconstructed adequately. Therefore, our representation covers the full gamut of valid reflectances. The resulting textures are compact because each coefficient can be stored in 10 bits. For compatibility with existing tristimulus assets, we implement a mapping from tristimulus color spaces to three Fourier coefficients. Using three coefficients,

Authors' addresses: Christoph Peters, christoph.peters@kit.edu, Karlsruhe Institute of Technology, Am Fasanengarten 5, 76131, Karlsruhe, Germany; Sebastian Merzbach, merzbach@cs.uni-bonn.de, University of Bonn, Endenicher Allee 19a, 53115, Bonn, Germany; Johannes Hanika, hanika@kit.edu, Karlsruhe Institute of Technology, and Weta Digital, Karlsruhe, Germany; Carsten Dachsbacher, dachsbacher@kit.edu, Karlsruhe Institute of Technology, Karlsruhe, Germany.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3306346.3322964>.

our technique gives state of the art results without some of the drawbacks of related work. With four to eight coefficients, our representation is superior to all existing representations. Our focus is on offline rendering but we also demonstrate that the technique is fast enough for real-time rendering.

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

Additional Key Words and Phrases: spectral rendering, reflectance spectra, emission spectra, bounded signals, bounded maximum entropy spectral estimate, bounded MESE, Fourier coefficients, trigonometric moments

ACM Reference Format:

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 (July 2019), 14 pages. <https://doi.org/10.1145/3306346.3322964>

1 INTRODUCTION

Color representation is one of the most fundamental problems in computer graphics. The ubiquitous tristimulus color spaces are adequate for color display, yet insufficient for accurate rendering. A physically based renderer should simulate light transport at each wavelength and thus dense spectra for light emission and surface reflectances need to be known. The accuracy of techniques that infer dense spectra from tristimulus data is limited as the problem is severely underconstrained [Jakob and Hanika, 2019, Meng et al., 2015, Otsu et al., 2018, Smits, 1999]. Especially for scenes with spiky

emission spectra, e.g. due to fluorescent lamps, defining reflectance with tristimulus values gives uncontrollable errors (Figure 1).

To overcome this problem, more information about the spectra must be stored. For emission spectra, dense sampling is viable because they usually do not take a major portion of the scene memory. On the other hand, textures storing reflectance spectra for potentially multi-layered surface albedos take huge amounts of memory, even for tristimulus colors. Access to densely sampled reflectance spectra during rendering is hardly practical. The spectra may be projected into a lower-dimensional linear function space, using e.g. a truncated Fourier series. However, such approaches are prone to ringing, which may lead to albedos outside of the valid range between zero and one. Non-linear representations of spectra in turn are incompatible with linear prefiltering of textures.

We present a novel approach that offers good approximations to reflectance spectra using three to eight scalars per spectrum. The spectra are represented by common Fourier coefficients. Thus, the approach is compatible with linear filtering. However, we do not use a truncated Fourier series for the reconstruction. Instead, we use a non-linear reconstruction originating from the theory of moments. The so-called maximum entropy spectral estimate (MESE) provides a good way to reconstruct positive densities from Fourier coefficients [Burg, 1975, Peters et al., 2015]. It yields useful reconstructions for emission spectra but violates the upper bound of one for reflectance spectra. Therefore, we employ a duality between problems with and without an upper bound [Markoff, 1896] and derive a novel algorithm to compute a bounded MESE for use with reflectance spectra.

The MESE and the bounded MESE always realize the given Fourier coefficients exactly. At the same time, they always obey their respective bounds. The reconstructions are smooth functions but nonetheless they can localize mass strongly in particular points or intervals if the spectra call for it. A constant spectrum is represented just as well as a spectrum made up of several nearly monochromatic components. Hence, the color space covers the full valid gamut. Both reconstructions are computed by efficient closed-form solutions. We also compute error bounds in closed form.

Unlike the Fourier basis, spectra are usually aperiodic. We resolve this issue with an approach similar to the discrete cosine transform. To ensure that information captured by tristimulus color spaces is also captured well by our approach, we apply an optimized warp to spectra before we compute Fourier coefficients. Our representation is compact. Storing each Fourier coefficient in 10 bits is sufficient such that six coefficients only take up eight bytes. Assets that use tristimulus color spaces are converted to three real Fourier coefficients at loading time by means of a lookup table.

Using four to eight coefficients per texel, the accuracy of our approach is superior to all prior work. With six coefficients, even the most unusual illuminants only lead to minor errors in rendered images (Figure 1). The method to convert assets from tristimulus color spaces to three Fourier coefficients yields state of the art results but does not require lookup tables during rendering and still supports linear filtering. Although we focus on offline rendering, our technique is fast enough for use in real-time renderers with stochastic sampling of wavelengths. We provide optimized implementations in C, C++, HLSL and Python.

2 RELATED WORK

Any perceived color arises from a spectrum s mapping a wavelength $\lambda \in \mathbb{R}$ to a density $s(\lambda) \geq 0$. The CIE has specified a standard observer by defining color matching functions $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, $\bar{z}(\lambda)$. The XYZ coefficients

$$X := \int_0^\infty s(\lambda)\bar{x}(\lambda) d\lambda, Y := \int_0^\infty s(\lambda)\bar{y}(\lambda) d\lambda, Z := \int_0^\infty s(\lambda)\bar{z}(\lambda) d\lambda$$

fully describe how a stimulus $s(\lambda)$ is perceived by the standard observer. For all other purposes, the XYZ coefficients provide incomplete information about the spectrum s .

A variety of RGB color spaces is defined in terms of linear transforms of the XYZ color space. Many renderers use the simplification that each of these channels corresponds to exactly one wavelength and only simulate light transport at these three wavelengths. While this model is self-consistent, simple and highly efficient, it fails to capture the complex reality [Fascione et al., 2017].

Spectral densities arise for any non-negative, wavelength-dependent quantity but emission spectra and reflectance spectra are particularly important. Emission spectra describe how much light is emitted at each wavelength. They can take arbitrarily large values and may have sharp peaks. Reflectance spectra are wavelength-dependent albedos and as such they are bounded between zero and one. Furthermore, they tend to be fairly smooth [Wandell, 1987]. It is common practice to factor the albedo out of the BRDF and to texture it. Thus, a compact representation is needed.

Most albedo textures use an RGB color space. A spectral renderer then needs to recover a dense reflectance spectrum from these tristimulus values. The simplest approach is to construct a linear combination of three predefined functions [Glassner, 1989]. Such linear approaches are compatible with any linear color space [Peercy, 1993, Wandell, 1987]. However, a spectrum that obeys the bounds, may violate them after the orthogonal projection into the span of the basis functions (see supplementary document). Smits [1999] addresses this problem with a piecewise linear approach using seven base spectra to cover sRGB.

Meng et al. [2015] strive for smoother spectra by precomputing them for a dense sampling of the chromaticities

$$x := \frac{X}{X + Y + Z} \quad \text{and} \quad y := \frac{Y}{X + Y + Z}.$$

They observe that some XYZ triples cannot be realized by reflectance spectra and propose methods to project them back into the solid of valid reflectances. MacAdam [1935a, 1935b] observes that binary spectra with two discontinuities realize maximal brightness for a given chromaticity and constructs corresponding lookup tables. Otsu et al. [2018] take a piecewise linear approach subdividing the xy chromaticity space into eight clusters. Per cluster they use a principal component analysis of spectra relevant for the scene and map XYZ triples into this basis. Jakob and Hanika [2019] define a three-dimensional non-linear space of functions by concatenating a sigmoid with quadratic polynomials. A look-up table converts XYZ triples to polynomial coefficients at loading time.

All of these techniques work with incomplete information. Errors are inevitable and are particularly noticeable under illuminants with sharp peaks and in indirect illumination.

Work on higher-dimensional representations of reflectance spectra is less extensive. Peercy [1993] proposes to fix an orthonormal basis, possibly by using a principal component analysis, and to perform all spectral computations in terms of linear combinations of this basis. Wandell [1987] proposes a first-order Fourier basis. The issue that projections of valid spectra into the basis may violate the physical bounds is not addressed. In the supplementary document, we prove that the only way to address this issue with linear least squares is to use basis functions with pairwise disjoint support. Indeed, it is a common practice to resort to disjoint box functions covering the range of visible wavelengths although this approach suffers from high memory requirements. For example, Kravchenko et al. [2017] found that eleven wavelength samples are needed for reflectance spectra under certain well-behaved illuminants.

Measured data may be turned into dense spectra using models learned from reflectance databases. Heikkinen et al. [2008] survey methods to turn tristimulus measurements under one or more known illuminants into spectra and propose a way to enforce physical bounds. Park et al. [2007] propose active imaging with a multiplexed LED array to capture input for such techniques at video frame rates. Their reconstruction relies on regularized quadratic programming. Aeschbacher et al. [2017] learn a dictionary of linear maps from RGB to dense spectra similar to Otsu et al. [2018] but with far more clusters. In general, these approaches emphasize quality over efficiency and are more suitable for content generation than rendering.

Despite the open problems, spectral rendering received more attention with the advent of elaborate material models and rendering techniques. An accurate capture and modeling of light material interactions is, for example, required for realistic skin rendering [Chen et al., 2015], material aging [Kimmel et al., 2013], wave optics effects from surface microstructures [Yan et al., 2018], thin films [Belcour and Barla, 2017] or diffraction gratings and holographic surfaces [Toisoul et al., 2018, Toisoul and Ghosh, 2017a]. Bergner et al. [2009] design spectra to produce given color impressions for different combinations of illuminants and surfaces.

Any renderer can be extended to use continuous spectra rather than tristimulus values but the increased complexity motivated research into efficient integrators. Hero wavelength spectral sampling [Radziszewski et al., 2009, Wilkie et al., 2014] reduces color noise by sampling paths based on a chosen hero wavelength, plus additionally evaluating the path for a fixed number of equidistant wavelengths. Ray differentials [Elek et al., 2014] and gradient sampling [Petitjean et al., 2018] have been extended to the spectral domain to improve path sampling. Light transport in spectral participating media has a large potential for improvements over naive rendering; recent examples include distance tracking in fluorescent materials [Mojzík et al., 2018] and spectral tracking [Kutz et al., 2017], a technique to alter free-path distributions such that the fluctuation of path throughputs and estimation variance is reduced. In real-time rendering, spectral rendering is important for faithful reproduction of surface shading under complex illuminants or for rendering complex materials such as diffraction gratings [Toisoul and Ghosh, 2017b].

Meanwhile, a branch of graphics research has used the theory of moments to efficiently reconstruct densities from power moments or Fourier coefficients. This approach has led to new solutions for

shadow map filtering [Peters and Klein, 2015, Peters et al., 2017], order-independent transparency [Münstermann et al., 2018, Sharpe, 2018] and transient imaging [Peters et al., 2015]. The work on transient imaging [Peters et al., 2015] uses the MESE [Burg, 1975], which reconstructs a smooth density from given Fourier coefficients. The density is always positive and if the ground truth signal is sufficiently sparse, the reconstructed density has sharp peaks at the same locations.

3 RECONSTRUCTION OF DENSITIES

In this section, we review the MESE [Burg, 1975, Peters et al., 2015] and simultaneously develop our novel bounded MESE. The MESE is a suitable representation for general spectral densities whereas the bounded MESE is designed for reflectance spectra. We begin with fundamental definitions (Section 3.1), establish a duality between bounded and unbounded problems (Sections 3.2 and 3.3) and then introduce and analyze the two reconstructions (Sections 3.4 to 3.8). Figure 2 gives an overview of how these concepts relate. Our application to spectral rendering is discussed in Section 4.

3.1 Unbounded and Bounded Moment Problems

Consider 2π -periodic signals $d(\varphi), g(\varphi)$ where $d(\varphi) \geq 0$ is a density and $g(\varphi) \in [0, 1]$ is a bounded density for all phases $\varphi \in \mathbb{R}$. Later, these signals will represent emission spectra and reflectance spectra, respectively. Methods to turn aperiodic, wavelength-dependent signals into periodic, phase-dependent signals are discussed in Section 4.1.

We represent these signals by $m + 1 \in \mathbb{N}$ complex Fourier coefficients. The Fourier basis written as vectorial function is

$$c(\varphi) := \frac{1}{2\pi} (\exp(-ij\varphi))_{j=0}^m \in \mathbb{C}^{m+1}.$$

Then the Fourier coefficients of d and g are

$$\gamma := \int_{-\pi}^{\pi} d(\varphi) c(\varphi) d\varphi \in \mathbb{C}^{m+1}, \quad (1)$$

$$c := \int_{-\pi}^{\pi} g(\varphi) c(\varphi) d\varphi \in \mathbb{C}^{m+1}. \quad (2)$$

Since $d(\varphi) \geq 0$ is a density, γ may also be viewed as the vector of trigonometric moments¹ representing the finite measure $d(\varphi) d\varphi$. Similarly, c is the vector of trigonometric $(0, 1)$ -moments representing g [Krein and Nudel'man, 1977, p. 246].

The theory of moments is concerned with the characterization of densities that realize given moments. In particular, it offers ways to reconstruct such densities. While a truncated Fourier series would exhibit ringing and may take negative values or values exceeding one, the non-linear reconstructions offered by the theory of moments are guaranteed to obey the original bounds. Speaking in terms of Fourier series, these reconstructions do not simply set all unknown Fourier coefficients to zero. They have energy across the entire frequency range.

¹Burg [1975] and Peters et al. [2015] define trigonometric moments as $2\pi\bar{\gamma}$ but we prefer to align the definition with the definition of Fourier coefficients.

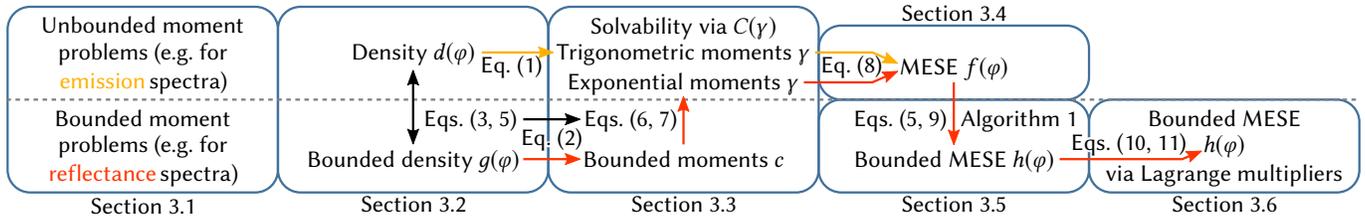


Fig. 2. An overview of the derivation of the bounded MESE. Reconstruction of emission spectra uses the MESE directly (yellow arrows) whereas reflectance spectra are reconstructed through the duality between bounded and unbounded problems giving rise to the bounded MESE (red arrows).

3.2 Reducing Bounded to Unbounded Problems

Solutions to trigonometric moment problems for unbounded densities have been used in graphics and other fields before [Burg, 1975, Peters et al., 2015]. There are closed-form solutions to obtain a smooth density representing given trigonometric moments γ . No such method exists for the reconstruction of smooth densities that are bounded above by one. However, Markov [Markoff, 1896] developed a duality between bounded and unbounded moment problems and Kreĭn and Nudel'man [1977] continued his work. Their duality allows us to transfer results on unbounded moment problems to the bounded case.

It is formulated in terms of the Herglotz transform. A density $d(\varphi) \geq 0$ is transformed into a signal on the open unit disk by

$$\mathcal{H}[d](z) := \frac{1}{2\pi} \int_{-\pi}^{\pi} \frac{\exp(i\varphi) + z}{\exp(i\varphi) - z} d(\varphi) d\varphi \in \mathbb{C},$$

where $z \in \mathbb{C}$ with $|z| < 1$. Then for any bounded density $g(\varphi) \in [0, 1]$, there exists a unique density² $d(\varphi) \geq 0$ and a constant $\alpha \in \mathbb{R}$ such that

$$\exp\left(\frac{\pi i}{2} \mathcal{H}[2g - 1](z)\right) = i\alpha + 2\pi \mathcal{H}[d](z) \quad (3)$$

for all $z \in \mathbb{C}$ with $|z| < 1$ [Kreĭn and Nudel'man, 1977, p. 245 f.]. Thus, a detour to functions in the unit disk gives us a one-to-one correspondence between bounded and unbounded densities.

We can compute g from moments c by first computing the corresponding d (Sections 3.3 and 3.4) and then solving Equation (3) for g . In the supplementary document we prove that for all $\varphi \in \mathbb{R}$

$$\lim_{z \rightarrow \exp(i\varphi)} \Re \mathcal{H}[g](z) = g(\varphi), \quad (4)$$

i.e. the original density can be recovered from the real part of the Herglotz transform. Hence, we take the limit on both sides in Equation (3) to recover g from $\mathcal{H}[d](z)$:

$$\begin{aligned} \mathcal{H}[2g - 1](z) &= -\frac{2}{\pi} i \log(i\alpha + 2\pi \mathcal{H}[d](z)) \\ \stackrel{\text{Eq. (4)}}{\Rightarrow} (2g - 1)(\varphi) &= \frac{2}{\pi} \lim_{z \rightarrow \exp(i\varphi)} \Re(-i \log(i\alpha + 2\pi \mathcal{H}[d](z))) \\ \Rightarrow g(\varphi) &= \frac{1}{\pi} \arg\left(i\alpha + \lim_{z \rightarrow \exp(i\varphi)} 2\pi \mathcal{H}[d](z)\right) + \frac{1}{2} \end{aligned} \quad (5)$$

This formulation reveals how our moment-based approach is able to guarantee $g(\varphi) \in [0, 1]$. It computes the argument of a complex

²Strictly speaking, $d(\varphi) d\varphi$ should be replaced by a finite measure that may include Dirac- δ components. See Section 3.7.

number with non-negative real part (by Equation (4)). This angle in $[-\frac{\pi}{2}, \frac{\pi}{2}]$ is then mapped to the interval $[0, 1]$ linearly.

3.3 Solvability

Before we introduce methods to reconstruct densities from their trigonometric moments, we want to establish when such solutions exist. In doing so, we encounter important building blocks for our reconstruction algorithms. One of them is the so-called Toeplitz matrix. Let trigonometric moments for negative indices $j \in \{-m, \dots, -1\}$ be given by the complex conjugate $\gamma_j := \overline{\gamma_{-j}}$. We denote the Toeplitz matrix by

$$C(\gamma) := \left(\frac{\gamma_{j-k}}{2\pi} \right)_{j,k=0}^m = \frac{1}{2\pi} \begin{pmatrix} \gamma_0 & \overline{\gamma_1} & \cdots & \overline{\gamma_m} \\ \gamma_1 & \gamma_0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \overline{\gamma_1} \\ \gamma_m & \cdots & \gamma_1 & \gamma_0 \end{pmatrix} \in \mathbb{C}^{(m+1) \times (m+1)}.$$

Each of the diagonals is filled with one of the trigonometric moments. The Toeplitz matrix makes it easy to decide solvability:

Proposition 1. *There exists a density $d(\varphi) > 0$ with moments*

$$\int_{-\pi}^{\pi} d(\varphi) \mathfrak{c}(\varphi) d\varphi = \gamma$$

if and only if $C(\gamma)$ is positive definite. There exists a finite measure utilizing Dirac- δ distributions that represents γ if and only if $C(\gamma)$ is positive semi-definite.

PROOF. The proofs for the positive definite case [Peters et al., 2015, Theorem 1] and the positive semi-definite case [Kreĭn and Nudel'man, 1977, p. 65, p. 78] are in the literature. \square

This result is powerful but it does not tell us when the vector of trigonometric $(0, 1)$ -moments c may be represented by a density $g(\varphi) \in [0, 1]$. For example, $c_0 > 1$ certainly means that such a representation cannot exist. In general, we should not use any method for unbounded problems with the trigonometric $(0, 1)$ -moments c directly. We have to use the duality from the previous section to convert them to trigonometric moments γ :

Proposition 2. *Suppose that $g(\varphi) \in [0, 1]$ and $d(\varphi) \geq 0$ are related by Equation (3) and c and γ are the respective trigonometric moments. Then*

$$\gamma_0 = 2\Re \gamma'_0, \quad \alpha = 4\pi \Im \gamma'_0 \quad \text{with} \quad \gamma'_0 := \frac{1}{4\pi} \exp\left(\pi i \left(c_0 - \frac{1}{2}\right)\right) \quad (6)$$

and for all $l \in \{1, \dots, m\}$

$$\gamma_l = \frac{2\pi i}{l} \left(l\gamma'_0 c_l + \sum_{j=1}^{l-1} (l-j)\gamma_j c_{l-j} \right). \quad (7)$$

PROOF. See the supplementary document. The technique employed here is attributed to Roger Barnard [Gustafsson and Putinar, 2017, p. 12]. \square

Equations (6) and (7) define a linear recurrence that allows us to compute γ from c in time $O(m^2)$. Owing to the exponential in Equation (3), the moments γ computed in this manner are referred to as exponential moments [Gustafsson and Putinar, 2017, p. 11]. Among other things, they let us decide whether it is possible at all to reconstruct a density $g(\varphi) \in [0, 1]$ from the trigonometric $(0, 1)$ -moments c . We simply analyze the Toeplitz matrix for the exponential moments γ [Krein and Nudel'man, 1977, p. 246]:

Proposition 3. *Given $c \in \mathbb{C}^{m+1}$ with $c_0 \in [0, 1]$, let $\gamma \in \mathbb{C}^{m+1}$ as in Equations (6) and (7). There exists a function $g(\varphi) \in [0, 1]$ with*

$$\int_{-\pi}^{\pi} g(\varphi) c(\varphi) d\varphi = c$$

if and only if the Toeplitz matrix $C(\gamma)$ is positive semi-definite.

Thus, the question of solvability is settled. We just verify that the Toeplitz matrix $C(\gamma)$ is positive definite. Along the way we have learned that Equations (6) and (7) map trigonometric $(0, 1)$ -moments c defining a bounded moment problem to exponential moments γ defining a dual unbounded moment problem.

3.4 Maximum Entropy Spectral Estimate

We now review the MESE. It was introduced by Burg [1975] and has been used in graphics before [Peters et al., 2015]. The MESE is a suitable tool for the reconstruction of emission spectra. Though, we are primarily interested in its dual, the bounded MESE, which we develop in the next section. We have already made all the necessary preparations to define the MESE $f(\varphi)$:

Theorem 4. *Let trigonometric moments $\gamma \in \mathbb{C}^{m+1}$ with positive definite $C(\gamma)$ be given. Let $e_0 := (1, 0, \dots, 0)^T \in \mathbb{C}^{m+1}$. Then*

$$f(\varphi) := \frac{1}{2\pi} \frac{e_0^* C^{-1}(\gamma) e_0}{|e_0^* C^{-1}(\gamma) c(\varphi)|^2} \quad (8)$$

is a positive density satisfying

$$\int_{-\pi}^{\pi} f(\varphi) c(\varphi) d\varphi = \gamma.$$

Among all such densities, it is the global maximum of the Burg entropy

$$\int_{-\pi}^{\pi} \log f(\varphi) d\varphi.$$

PROOF. See [Burg, 1975, p. 8 ff.] or [Peters et al., 2015, Theorem 1]. \square

The simplicity of this result is remarkable. Just like a truncated Fourier series, the MESE f matches all of the given Fourier coefficients γ exactly. However, it is guaranteed to be positive and by Proposition 3, it accomplishes this goal whenever it is possible at

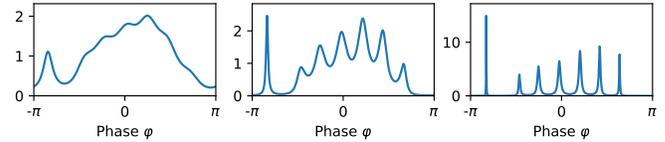


Fig. 3. Examples of the MESE for $m = 7$. The trigonometric moments $\gamma_1, \dots, \gamma_m \in \mathbb{C}$ are the same for all three plots but γ_0 is set to 1.06, 0.79 and 0.75 to approach a sparse reconstruction.

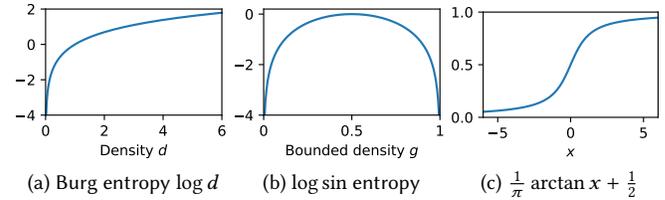


Fig. 4. The entropies, which are optimized by the MESE and bounded MESE, and the scaled arctangent that is used with the Lagrange multipliers.

all. Figure 3 demonstrates how it reconstructs smooth densities as well as densities that localize mass in just a few points. The Burg entropy (Figure 4a) punishes small densities heavily and rewards large densities moderately. Intuitively, the optimal solution in terms of the Burg entropy is relatively smooth and uses moderately large densities as much as possible.

The most costly part of the evaluation of the MESE is the solution of a system of linear equations for $C^{-1}(\gamma)e_0$. In general, this step would take time $O(m^3)$ but thanks to the special structure of the Toeplitz matrix, Levinson's algorithm solves the problem in time $O(m^2)$ [Peters et al., 2015].

3.5 Bounded Maximum Entropy Spectral Estimate

We are now ready to describe the core of our reconstruction for reflectance spectra: The bounded MESE. Our strategy is to consider the MESE f for the exponential moments γ . Then we apply the duality in Equation (5) to transform it into the bounded MESE. As we combine these concepts, there is a challenging problem. Equation (5) requires us to compute the Herglotz transform of the MESE

$$\mathcal{H}[f](z) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \frac{\exp(i\varphi) + z}{\exp(i\varphi) - z} f(\varphi) d\varphi.$$

We solve this problem in the Fourier domain. The infinite sequence of Fourier coefficients for the MESE f obeys a linear recurrence. The Fourier coefficients of the other factor are given by a simple closed form. Then the Herglotz transform equals the inner product of the two Fourier transforms. We solve this infinite series using a generalization of the geometric series to matrices. For details, please refer to the supplementary document. The whole procedure is implemented by Algorithm 1.

Assembling all of the pieces, we obtain the bounded MESE $h(\varphi)$:

Theorem 5. *Algorithm 1 is correct and takes time $O(m)$.*

Given trigonometric $(0, 1)$ -moments $c \in \mathbb{C}^{m+1}$ with $c_0 \in [0, 1]$, let $\gamma \in \mathbb{C}^{m+1}$ and γ'_0 be exponential moments as in Equations (6) and (7).

Algorithm 1 Herglotz transform of the MESE.

Input: Trigonometric moments $\gamma \in \mathbb{C}^{m+1}$ such that $C(\gamma)$ is positive definite, $q := C^{-1}(\gamma)e_0$ and $z \in \mathbb{C}$ with $|z| \leq 1$.

Output: $\mathcal{H}[f](z) \in \mathbb{C}$ for f as in Equation (8).

$p_m := q_0$

For $l = m - 1, \dots, 0$:

$$p_l := q_{m-l} + p_{l+1}z^{-1}$$

Return $\gamma_0 + \frac{2}{p_0} \sum_{k=1}^m p_k \gamma_k$

Suppose that $C(\gamma)$ is positive definite and f is the MESE. Then

$$h(\varphi) := \frac{1}{\pi} \arg(2i\Im\gamma'_0 + \mathcal{H}[f](\exp(i\varphi))) + \frac{1}{2} \quad (9)$$

is a smooth density in $[0, 1]$ satisfying

$$\int_{-\pi}^{\pi} h(\varphi) \mathbf{c}(\varphi) \, d\varphi = c.$$

Among all such densities, it is the global maximum of the log sin entropy

$$\int_{-\pi}^{\pi} \log \sin(\pi h(\varphi)) \, d\varphi.$$

PROOF. Equation (9) arises from Equation (5) through substitution of α and by exploiting that Algorithm 1 accepts inputs with $|z| = 1$. The entire proof is given in the supplementary document. \square

Each property of the bounded MESE is dual to a property of the MESE. It matches all given Fourier coefficients with a smooth density while obeying the bounds and it accomplishes this whenever it is possible at all. As shown in Figure 4b, the log sin entropy punishes densities near zero or one heavily while a density of 0.5 is optimal. From a physical point of view, this is a desirable behavior. Materials that absorb or reflect all of the light at some wavelength are rare. It also benefits the convergence of renderers because spectra without extreme values have lower variance. Nonetheless, the bounded MESE is able to get arbitrarily close to zero or one if the situation demands it (Figure 5).

Summary of the Algorithm. Initially, we use the recurrence in Equations (6) and (7) to transform the trigonometric $(0, 1)$ -moments \mathbf{c} into exponential moments γ (time $O(m^2)$). As for the MESE, we use Levinson's algorithm to compute $q = C^{-1}(\gamma)e_0$ in time $O(m^2)$ [Peters et al., 2015]. Finally, Algorithm 1 and Equation (9) evaluate the bounded MESE at any phase $\varphi \in \mathbb{R}$ in time $O(m)$. The overall cost of evaluating the bounded MESE is roughly twice as high as for the MESE.

3.6 Evaluation at Many Phases

When we evaluate the bounded MESE at multiple phases $\varphi \in \mathbb{R}$, we only need to compute the exponential moments γ and $q = C^{-1}(\gamma)e_0$ once. Then the time per phase φ is $O(m)$ but with a moderately large constant. For evaluation at many wavelengths, the following reformulation provides a more efficient solution:

Proposition 6. *Suppose we are in the setting of Theorem 5. Let $q := C^{-1}(\gamma)e_0$ and for all $j \in \{1, \dots, m\}$ let $\gamma'_j := \gamma_j$. We define*

Lagrange multipliers for all $l \in \{0, \dots, m\}$:

$$\lambda_l := \frac{1}{\pi i q_0} \sum_{k=0}^{m-l} \gamma'_k \sum_{j=0}^{m-k-l} \overline{q_{j+k+l}} q_j \in \mathbb{C}. \quad (10)$$

Then the bounded MESE as defined in Equation (9) is given by

$$h(\varphi) = \frac{1}{\pi} \arctan \left(\Re \lambda_0 + 2\Re \sum_{l=1}^m \lambda_l \exp(-il\varphi) \right) + \frac{1}{2}. \quad (11)$$

PROOF. See the supplementary document. \square

All sums involving q in Equation (10) are entries of the autocorrelation of the sequence q_0, \dots, q_m , which is trivially computed in time $O(m^2)$. Computation of the Lagrange multipliers $\lambda_0, \dots, \lambda_m$ takes time $O(m^2)$ in total. Once they are available, evaluation of the bounded MESE is very efficient. We only have to evaluate the Fourier series in Equation (11) and apply the arctangent. Except for the arctangent, the cost per phase φ is the same as for a common truncated Fourier series. However, our solution never violates the bounds between zero and one through ringing.

This reformulation also offers another way to understand the behavior of the bounded MESE h . It is a real-valued Fourier series whose values are subsequently forced into the allowable range by the arctangent (Figure 4c). If the Fourier series takes moderate values, h will be smooth and stay away from the bounds. If it has large magnitudes, h may transition from zero to one almost instantaneously. This class of functions resembles those used by Jakob and Hanika [2019] and Heikkinen et al. [2008] but ours has arbitrary dimension and the Lagrange multipliers $\lambda_0, \dots, \lambda_m \in \mathbb{C}$ are computed efficiently from a linear representation of the spectrum.

Summary of the Algorithm. As in Section 3.5, we first compute exponential moments γ using Equations (6) and (7) and then use Levinson's algorithm to compute $q = C^{-1}(\gamma)e_0$. Next, we compute entries of the autocorrelation $\sum_{j=0}^{m-k} \overline{q_{j+k}} q_j$ for all $k \in \{0, \dots, m\}$. The results let us compute Lagrange multipliers $\lambda_0, \dots, \lambda_m \in \mathbb{C}$ with Equation (10). Up until this point, every step took time $O(m^2)$. With the Lagrange multipliers, Equation (11) lets us evaluate the bounded MESE $h(\varphi)$ for any $\varphi \in \mathbb{R}$ in time $O(m)$.

The supplementary code provides implementations for all variants of the bounded MESE (algorithms in Sections 3.5 and 3.6) in C, C++, HLSL and Python.

3.7 Boundary Cases

The MESE is known to be suitable for smooth densities as well as sparse densities [Peters and Klein, 2015]. Figure 3 gives an example. The reconstruction to the left is smooth. As we remove a uniform component from the signal by reducing γ_0 , it develops more and more pronounced peaks. In the limit, the peaks approach m Dirac- δ distributions. This limit case allows a perfect reconstruction because no other distribution realizes the trigonometric moments γ . It corresponds to a Toeplitz matrix that is positive semi-definite but singular (cf. Proposition 1) [Peters et al., 2015, p. 6]. The ability to reconstruct smooth and sparse densities is valuable for emission spectra because they may be entirely smooth or may have sharp peaks (Figure 9).

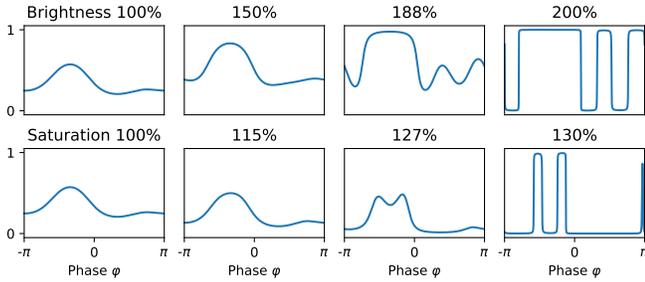


Fig. 5. Examples of the bounded MESE for $m = 3$. In the top row, we use the moments tc , i.e. we increase brightness. In the bottom row, the moments are $(1 - t)e_0 + tc$, i.e. we increase saturation. In the limit, we approach the brightest and most saturated spectra that are theoretically possible.

For the bounded MESE, we get analogue properties which have a useful interpretation in terms of color representation. MacAdam [1935a, 1935b] sought reflectance spectra of maximal brightness for a given chromaticity x, y and learned that such spectra are binary with only two discontinuities. If we think of the wavelength range as periodic, these MacAdam spectra take a value of one on an interval and zero elsewhere. At any fixed brightness, they trace out the largest gamut that can be covered by reflectance spectra.

The same construction is possible with our trigonometric $(0, 1)$ -moments c (Figure 5). We consider $\frac{c}{c_0} \in \mathbb{C}^{m+1}$ as a higher-dimensional chromaticity. It is not limited by human perception but provides a more complete description of the reflectance spectrum. We can change the overall brightness of the spectrum without changing the chromaticity by considering tc with $t \geq 1$. As we increase t , the initial effects are similar to common scaling of the spectrum (Figure 5 top). However, scaling would quickly violate the bound one.

Our bounded MESE instead approaches spectra that assign full brightness to m intervals and no brightness outside these intervals. We think of these spectra as generalized MacAdam spectra with m intervals instead of one. In the limit, this reconstruction is the only one possible and once more the Toeplitz matrix $C(\gamma)$ is positive semi-definite but singular [Kreĭn and Nudel'man, 1977, p. 78, 246]. Increasing the scaling t beyond this limit case means that the Toeplitz matrix $C(\gamma)$ has a negative eigenvalue and in this case, there is no matching reflectance spectrum for tc by Proposition 3. Thus, the bounded MESE gets arbitrarily close to the brightest spectra for a given chromaticity $\frac{c}{c_0}$.

In a similar spirit, we manipulate the saturation of a spectrum by moving away from the constant white spectrum of maximal brightness. This spectrum has moments $e_0 \in \mathbb{C}^{m+1}$, so we reconstruct from the moments $(1 - t)e_0 + tc$ for $t \geq 1$. Once more, the initial effect is similar to a linear change but as we approach the boundary case of a singular Toeplitz matrix, the spectra approach generalized MacAdam spectra (Figure 5 bottom). Still greater values of t have no interpretation in terms of reflectance spectra. The bounded MESE gets arbitrarily close to the most saturated version of an initial spectrum given by moments c .

While generalized MacAdam spectra, just like common MacAdam spectra, are not physically plausible, it is valuable to have a color

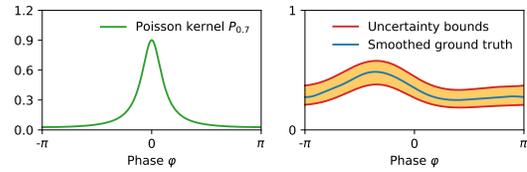


Fig. 6. An example of uncertainty bounds with $r = 0.7$ and $m = 4$. Using solely the knowledge of the bounds $g(\varphi) \in [0, 1]$ and the moments c , we are certain that $P_r * g$ lies between the shown uncertainty bounds.

space large enough to include them. As observed by MacAdam [1935a] bright and saturated colors have artistic value and dye chemists have been pushing the limits in this regard. Our bounded MESE provides the first higher-dimensional color space to get arbitrarily close to these idealized colors in rendered images. Of course, we are still limited by the gamut of the display device.

3.8 Uncertainty Bounds

There is another interesting result on the MESE that carries over to the bounded MESE. Using only the knowledge of the stored trigonometric moments γ and the fact that the ground truth does not take negative values, we can compute sharp bounds for a smoothed version of the signal [Karlsson and Georgiou, 2013, Kovalishina and Potapov, 1982, Peters and Klein, 2015]. Thus, we get quality guarantees.

An analogue result holds for the bounded MESE. Smoothing the signal is necessary because for moments $c \in \mathbb{C}^{m+1}$ with positive definite $C(\gamma)$, we can always construct a bounded density with value zero or one at any given phase $\varphi \in \mathbb{R}$ [Kreĭn and Nudel'man, 1977, p. 254 ff.]. By smoothing the signal first, the bounds become robust to small perturbations and sharp enough to be useful. This smoothing is done through convolution with a Poisson kernel

$$P_r(\varphi) := \frac{1}{2\pi} \frac{1 - |r|^2}{|1 - r \exp(i\varphi)|^2},$$

where $r \in [0, 1)$ controls the sharpness of the kernel. The Poisson kernel is closely related to the Herglotz transform and therefore this particular approach enables computation of the bounds in closed form. Figure 6 shows an example. For further results and the complete formulation of the uncertainty bounds, please refer to the supplementary document.

4 SPECTRAL RENDERING

Now that the reconstruction algorithms are available, we investigate their use for the storage and reconstruction of emission and reflectance spectra. We begin by defining the mapping between wavelengths λ and phases φ (Section 4.1). Then we analyze the impact of quantization errors in the trigonometric $(0, 1)$ -moments c (Section 4.2). To support assets authored using tristimulus color spaces, we convert such data to our representation (Section 4.3). Finally, we discuss the integration into renders (Section 4.4).

4.1 Mapping Wavelengths to Phases

Unlike the Fourier basis c , typical emission and reflectance spectra are aperiodic. This poses a problem because our reconstructions will

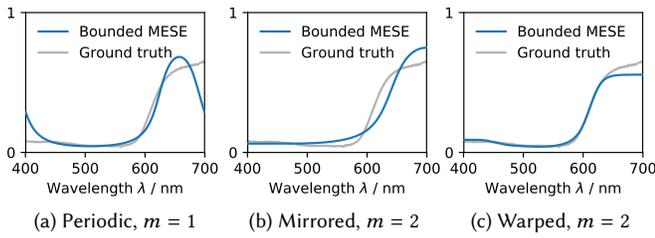


Fig. 7. Reconstructions with the bounded MESE using different methods to map wavelengths to phases. In all cases, the spectrum is stored by three real scalars.

be periodic. If we just map wavelengths $\lambda \in [\lambda_{\min}, \lambda_{\max}]$ to phases linearly by setting

$$\varphi = 2\pi \frac{\lambda - \lambda_{\min}}{\lambda_{\max} - \lambda_{\min}} - \pi \in [-\pi, \pi],$$

we get strong artifacts at the boundaries and moderate distortions throughout the entire range of wavelengths (Figure 7a).

We overcome this problem using the same approach as in the discrete cosine transform. The mapping to phases only uses negative values

$$\varphi = \pi \frac{\lambda - \lambda_{\min}}{\lambda_{\max} - \lambda_{\min}} - \pi \in [-\pi, 0]. \quad (12)$$

Once we have mapped a reflectance spectrum to the phase domain in this manner, we mirror it by defining $g(\varphi) = g(-\varphi)$ for all $\varphi \in [0, \pi]$. The moments c are then computed for this mirrored signal. The reconstruction will be even and we only use the part for $\varphi \in [-\pi, 0]$. Thus, aperiodic spectra can be reconstructed adequately (Figure 7b).

At first sight, this approach may seem wasteful since we drop half of the reconstruction. However, the moments for the even signal are real numbers. Thus, storing them takes $m + 1$ real scalars rather than the $2m + 1$ real scalars needed for complex moments. We need twice as many moments for signals of similar complexity but the storage cost is roughly the same. Only the computation becomes more expensive but remains affordable.

Additionally, we account for the greater perceptual importance of some wavelength ranges. While slight changes are perceivable near 550 nm, light is hardly visible at 700 nm. When using a small number of coefficients, we would like to focus accuracy on important regions. We accomplish this weighting by means of a differentiable, bijective function $\varphi : [\lambda_{\min}, \lambda_{\max}] \rightarrow [-\pi, 0]$. The moments for the warped and mirrored signal $g(\varphi)$ are

$$c = 2\Re \int_{-\pi}^0 g(\varphi)c(\varphi) d\varphi = 2 \int_{\lambda_{\min}}^{\lambda_{\max}} g(\varphi(\lambda))\Re c(\varphi(\lambda))\varphi'(\lambda) d\lambda.$$

In the extreme case $m = 2$, it would be best if $\Re c(\varphi(\lambda))\varphi'(\lambda)$ resembles the color matching functions $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, $\bar{z}(\lambda)$. Thus, we set up a non-convex optimization that minimizes the distance between the function space spanned by \bar{x} , \bar{y} , \bar{z} and the three component functions of $\Re c(\varphi(\lambda))\varphi'(\lambda)$. Our objective function is the matrix 2-norm of the difference between the two orthogonal projectors onto the

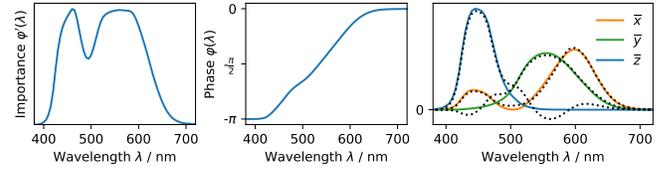


Fig. 8. We optimize the derivative $\varphi'(\lambda)$ (left) of the mapping from wavelengths to phases $\varphi(\lambda)$ (middle). The objective is to make the space spanned by the three warped and weighted Fourier basis functions $\Re c(\varphi(\lambda))\varphi'(\lambda)$ similar to the space spanned by \bar{x} , \bar{y} , \bar{z} . The dotted lines (right) show squares fits to \bar{x} , \bar{y} , \bar{z} in the optimized space.

discretized function spaces [Golub and Van Loan, 2012, p. 82]. Positivity constraints enforce $\varphi'(\lambda) > 0$. The optimization procedure and the resulting warp $\varphi(\lambda)$ are part of the supplementary code.

Figure 8 visualizes the optimized warp and how it accomplishes the goal of approximating the color matching functions with the Fourier basis. We have discretized this warp using 95 equidistant samples. Thus, it can be evaluated with a single lookup using linear interpolation. The reconstruction in Figure 7c shows how the warp improves reconstruction quality in important wavelength ranges while diminishing it in others.

We could repeat the optimization outlined above for $m > 2$ but that would reward solutions where all moments carry the same information as the XYZ coefficients. Thus, it is preferable to use the same warp for all values of m . For $m \geq 5$, the moments capture entire reflectance spectra well without the warp. Results with the warp or with Equation (12) are similar. For emission spectra, we use more coefficients and refrain from using the warp.

When computing moments for a sampled spectrum, the sample locations should be transformed to phases first. It is important to use an integration scheme that respects the bounds. For reflectance spectra, the sample locations must not be idealized as Dirac- δ pulses. Instead the samples should be interpolated with linear interpolation. The supplementary document has formulas to compute the moments for such interpolated signals. Higher-order integration schemes must not be used because the underlying polynomial approximations may violate the bounds. If data is unavailable for part of the range of phases, the nearest available value should be repeated or the range should be shortened. Treating parts with unknown value as zero introduces unnecessary complexity to the signal.

4.2 Quantization

Once we have computed the moments, we need to store them. Since each texel of a texture may hold one reflectance spectrum, a compact representation is important and it is worthwhile to investigate the impact of quantization errors.

The zeroth moment c_0 is in $[0, 1]$. The other moments take a minimal value when the reflectance spectrum is one wherever the basis function is negative and zero elsewhere. Solving integrals, we find that $\Re c_j, \Im c_j \in [-\frac{1}{\pi}, \frac{1}{\pi}]$ for $j \in \{1, \dots, m\}$. There is no good reason to have greater precision for values of a small magnitude. Thus, we favor fixed-point quantization over floating-point.

Table 1. Metrics for the error introduced to the bounded MESE with warping by fixed-point quantization of the moments c . For each of the 3076 spectra to reconstruct, we compute the root-mean-square error (RMSE). Then we list the arithmetic mean of these RMSEs and their maximum. Additionally, we report the maximal error encountered at a single wavelength throughout the experiment.

m	Mean RMSE		Max. RMSE		Max. error	
	10 bits	16 bits	10 bits	16 bits	10 bits	16 bits
2	$4 \cdot 10^{-3}$	$7 \cdot 10^{-5}$	$3 \cdot 10^{-2}$	$3 \cdot 10^{-4}$	$5 \cdot 10^{-2}$	$4 \cdot 10^{-4}$
3	$5 \cdot 10^{-3}$	$8 \cdot 10^{-5}$	$4 \cdot 10^{-2}$	$3 \cdot 10^{-4}$	$7 \cdot 10^{-2}$	$7 \cdot 10^{-4}$
5	$6 \cdot 10^{-3}$	$1 \cdot 10^{-4}$	$2 \cdot 10^{-2}$	$6 \cdot 10^{-4}$	$5 \cdot 10^{-2}$	$1 \cdot 10^{-3}$
7	$7 \cdot 10^{-3}$	$1 \cdot 10^{-4}$	$1 \cdot 10^{-1}$	$7 \cdot 10^{-4}$	$3 \cdot 10^{-1}$	$1 \cdot 10^{-3}$

To determine sufficient bit counts for real-world spectra, we have evaluated our bounded MESE on a database of 3076 measured reflectance spectra (see Section 5.1). Table 1 summarizes the results. The mean errors are on a similar order of magnitude as the error introduced to the moments. The error grows as more quantized moments are used. Though, even with $m = 7$ and 10-bit quantization, the mean errors are still acceptable.

Near the boundary cases discussed in Section 3.7, errors are stronger. For example, a small increase in saturation in Figure 5 leads to a substantially different spectrum. When these cases are of particular interest, 16-bit quantization should be used. In these cases, rounding errors may even lead to vectors of moments that make reconstruction impossible in the sense of Proposition 3. Then the vector of moments c should be biased towards the spectrum of constant value 0.5, i.e. replaced by $(1 - \alpha)c + \alpha 0.5e_0$ [Münstermann et al., 2018]. A moment bias $\alpha > 0$ on the order of magnitude of the rounding error guarantees valid results. Though, this procedure is not needed for typical reflectance spectra and we do not use it in our experiments.

In practice, memory alignment is important and thus the storage per spectrum should be a multiple of 4 bytes. Therefore, we recommend to use 10-bit quantization if the number of coefficients is divisible by three and 16-bit quantization otherwise. Then we need 4 bytes for three coefficients, 8 bytes for four or six coefficients and 16 bytes for eight coefficients.

4.3 Conversion of Tristimulus Data

Our main motivation is to provide a compact representation for spectral reflectance data. Though, the vast majority of production assets uses tristimulus color spaces. A single representation that works for spectral data as well as tristimulus data would be useful. Thus, we strive to convert tristimulus data to three real moments with 10-bit quantization.

The approach is simple. We use the bounded MESE to reconstruct reflectance spectra for all 1024^3 triples of moments (discarding the ones that are invalid according to Proposition 3). For each spectrum, we integrate the XYZ coefficients. The XYZ coefficients are converted to the perceptually uniform CIELAB color space and entered into an acceleration structure for nearest-neighbor queries. Next we sample a three-dimensional grid of the used tristimulus color space. For example, we take 256^3 samples for sRGB. For each cell,

we look up the nearest neighbor and store the corresponding triple of moments. Since distances are defined using CIELAB, tristimulus colors that cannot be realized by a reflectance spectrum are mapped back into the solid of valid reflectances using the minimal distortion method of Meng et al. [2015].

Such a lookup table of size 256^3 takes 64 MiB. It would be costly to deal with such a big lookup table during rendering but we only use it at loading time. Each asset is converted once and then the tristimulus data can be discarded. If the tristimulus data is quantized with more than 8 bits, interpolation in the lookup table is a sensible operation because the moments are defined linearly.

To enable construction of such lookup tables for arbitrary color spaces, we provide the C++ code of our construction method.

4.4 Use in Renderers

We have implemented our technique in an offline renderer and a real-time renderer. The offline renderer is a path tracer with hero wavelength spectral sampling [Wilkie et al., 2014] using eight wavelengths. Reflectance spectra are reconstructed using the bounded MESE. They are stored in textures and texture sampling uses trilinear filtering with mipmapping. This simple linear filtering is possible thanks to the linear definition of the moments. We only have to avoid filters with negative filter weights because they may invalidate vectors of moments.

The real-time renderer uses Direct3D 12 through the Falcor rendering framework [Benty et al., 2018]. It is a deferred renderer storing three real moments in the G-buffer. Shading supports direct illumination only, albeit with ray traced shadows. To obtain a low variance with sample counts affordable in real time, we sample the emission spectrum exactly by precomputing an inverted CDF. The corresponding lookup table holds the phase at which the bounded MESE needs to be evaluated and the corresponding values of \bar{x} , \bar{y} , \bar{z} . In the spirit of Hero wavelength spectral sampling, we only use one stratified random number per pixel and sample the other wavelengths using equidistant random numbers.

5 RESULTS

In the following, we evaluate our techniques in comparison to related work [Jakob and Hanika, 2019, MacAdam, 1935a, Meng et al., 2015, Otsu et al., 2018, Smits, 1999]. We begin with an overview of the spectral data that we use for this evaluation (Section 5.1). Then we briefly analyze the suitability of the MESE for emission spectra (Section 5.2) and refractive indices (Section 5.3). Furthermore, we present detailed experiments concerning the quality of the bounded MESE for reconstruction of reflectance spectra (Section 5.4). Finally, we demonstrate that real-time rendering with our approach is feasible (Section 5.5) and analyze run times (Section 5.6).

5.1 Sources of Spectral Data

The Lamp Spectral Power Distribution Database³ (LSPDD) contains densely sampled spectra for a diverse selection of 254 commercially available lamps. We use their data for most of our experiments. The exception is the spectrum of a sodium vapor lamp. For individual reflectance spectra, we use the X-Rite color checker and a database

³<http://galileo.graphyics.cegepsheerbrooke.qc.ca/app/en/lamps>

of 3052 spectra measured at the University of Eastern Finland⁴. This database includes matte and glossy color chips manufactured by Munsell as well as natural spectra from flowers, leaves and other colorful plants.

Additionally, we have measured spectral svBRDFs using the X-Rite TAC7-Scanner⁵. This off-the-shelf hardware samples the reflectance for each point on the probe using 29 white LEDs and four cameras. Five LEDs are behind filter wheels with ten narrow-band filters. The proprietary post-processing pipeline first uses this data to estimate HDR images for each pair of light and view directions at 31 equidistant wavelengths from 400 nm to 700 nm. Then parameters of the Geisler-Moroder Ward BRDF [Geisler-Moroder and Dür, 2010] with Schlick Fresnel-term are fitted.

The diffuse and specular albedo maps that are output by the proprietary svBRDF fitting currently use an RGB color space. Therefore, we have trained spectral A+ [Aeschbacher et al., 2017] with pairs of spectra from the intermediate HDR images and RGB triples from the albedo maps. The technique learns a mapping from RGB to spectra that is specific to a single measured probe. We use it to obtain svBRDFs with albedo maps sampled at 31 wavelengths. More details of this approach will be described in a subsequent publication. We use these textures in Figure 1. The rendering with tristimulus textures [Jakob and Hanika, 2019] has substantial error. Errors of our approach with six moments are hardly perceivable in spite of the challenging illuminant.

5.2 Reconstruction of Emission Spectra

Our efforts are focused on reflectance spectra because we believe that compact representations are less important for emission spectra. Nonetheless, there are instances where emission spectra take noteworthy amounts of memory, e.g. when dealing with light probes. Since we have already introduced the MESE, it is interesting to analyze its capability to represent emission spectra.

To obtain good results with this approach, we need to address one issue. Many of the LSPDD emission spectra include long stretches of zero value at both ends. Computing trigonometric moments for such signals directly may lead to undesirable results. The Burg entropy punishes densities near zero heavily (Figure 4a). If the ground truth forces small densities in large regions, the MESE will resolve this conflict by producing a spectrum with sharp peaks. Thus, a spectrum that should resemble the one to the left in Figure 3 may be more similar to the one in the middle. We address this problem by choosing λ_{\min} , λ_{\max} for Equation (12) per spectrum in such a way that we discard up to 0.2% of the overall energy. Alternatively, we could regularize by adding a small uniform component to the trigonometric moment γ_0 (Figure 3).

Figure 9 shows results for a diverse selection of LSPDD emission spectra. We use 16 real Fourier coefficients and additionally store λ_{\min} , λ_{\max} . The reconstructed densities for the smooth spectra in the top row are highly accurate. For the more spiky spectra in the bottom row, the MESE often merges multiple peaks into a single broader peak. However, the cumulative distribution functions show that the overall energy in small intervals is always accurate.

⁴<http://www.uef.fi/web/spectral/-spectral-database>

⁵<https://www.xrite.com/categories/appearance/total-appearance-capture-ecosystem/tac7>

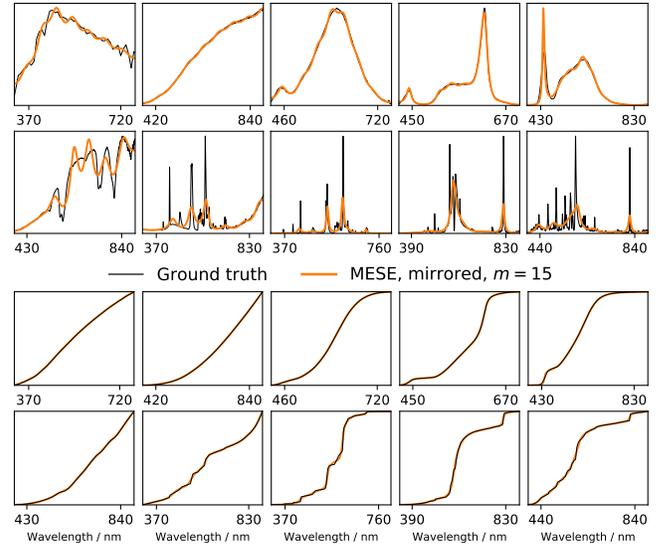


Fig. 9. Reconstructions for a selection of emission spectra using the MESE with mirroring and $m = 15$ (i.e. 16 real moments). We show the spectral densities (top) and the cumulative distribution functions (bottom).

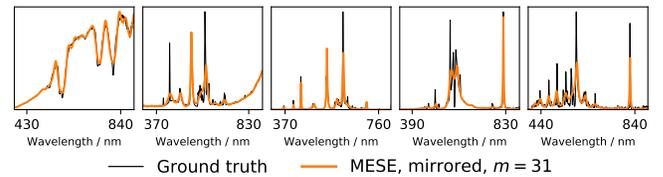


Fig. 10. The MESE with mirroring for difficult emission spectra using twice as many real Fourier coefficients as in Figure 9.

Since most reflectance spectra are smooth, the slight redistribution of energy between nearby wavelengths should not lead to large errors in rendered images. Nonetheless, doubling the number of coefficients yields a far more accurate reconstruction (Figure 10). In general, the benefits of the MESE over simple sampling of the signal taper off at large coefficient counts. Besides the $O(m^2)$ run time of the MESE might become problematic. Still, the MESE remains attractive if the signal of interest is sufficiently sparse.

5.3 Reconstruction of Complex Refractive Indices

Our techniques are applicable to any non-negative, wavelength-dependent signal. Even the complex refractive indices of conductors can be represented by handling the extinction as separate real quantity. Figure 11 uses the MESE with six real coefficients per quantity. The chosen conductors are among the more challenging cases but still the signals are well-behaved and the MESE is highly accurate. Simple linear interpolation of six samples already gives a satisfactory result when computing the Fresnel term for the conductors. Nonetheless, the MESE is far more accurate using the same number of coefficients. Results with more spiky illuminants are similar.

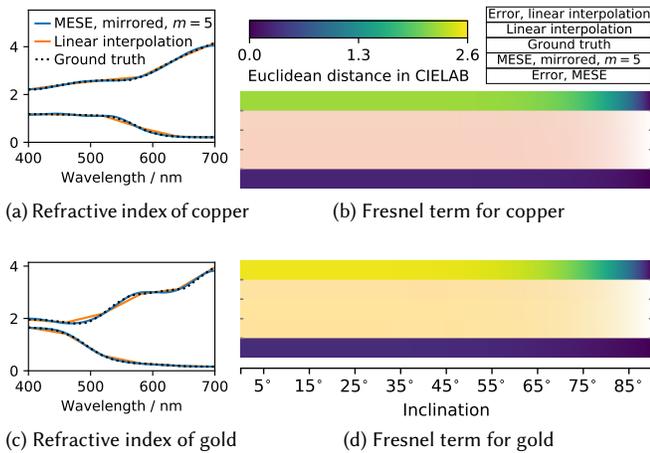


Fig. 11. The wavelength-dependent, complex refractive indices for two conductors. We compare the MESE with six real coefficients per signal to linear interpolation of six samples. The Fresnel term for ideal reflection of illuminant D65 in vacuum is shown to the right.

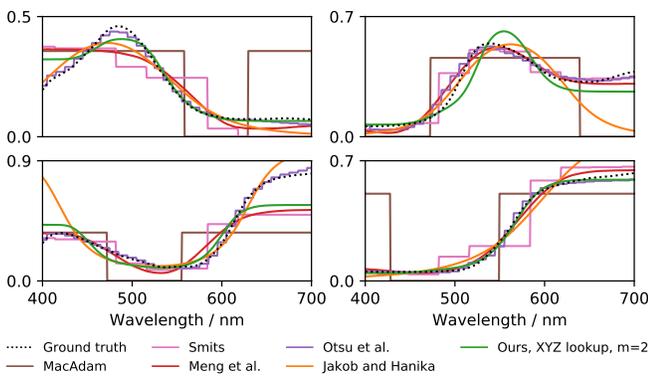


Fig. 12. Reconstructions of the reflectance spectra for four patches of the X-Rite color checker. All techniques only take an XYZ triple as input.

5.4 Reconstruction of Reflectance Spectra

For the remainder of this section, we evaluate our methods for reconstruction of reflectance spectra. We begin by considering plots of reconstructed reflectance spectra. Figure 12 compares all techniques that take a single XYZ triple as input. The two oldest techniques [MacAdam, 1935a, Smits, 1999] deviate from the ground truth substantially. Our technique with the warp and the work of Meng et al. [2015] optimize their reconstruction with regard to perceptual aspects. Therefore, their results are poor in the less important range between 650 and 700 nm but fairly good otherwise. The reconstruction of Jakob and Hanika [2019] is quite far off, especially in the boundary regions. Reconstructions of Otsu et al. [2018] are excellent. However, the color checker spectra used here are similar in nature to the Munsell color chips used for the training of this technique.

Figure 13 compares different variants of our technique. Even the reconstructions from four real moments (i.e. $m = 3$) are already

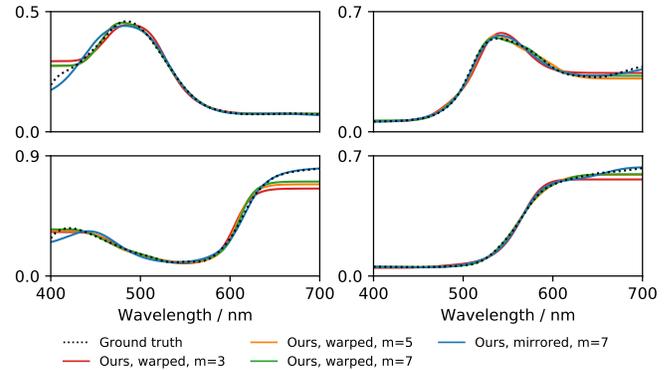


Fig. 13. Reconstructions for four patches of the X-Rite color checker using four, six or eight real moments computed from the ground truth spectrum.

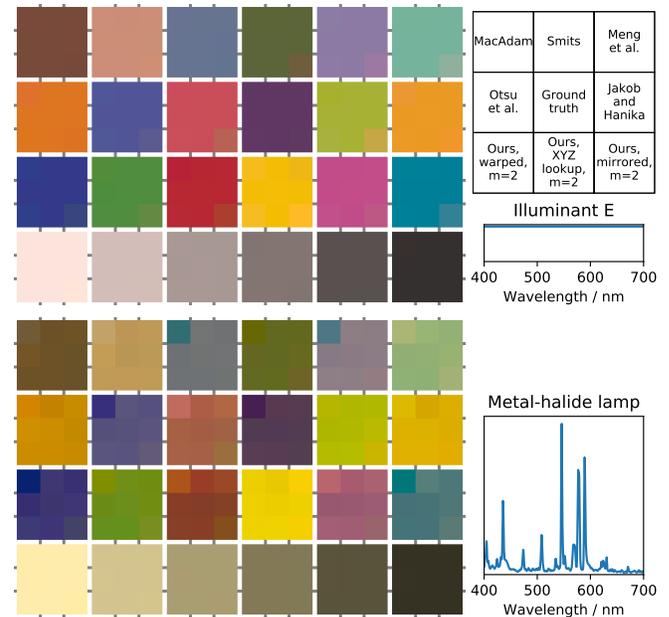


Fig. 14. An X-Rite color checker relit using the emission spectra to the right. Each patch is subdivided into 3×3 squares using the techniques indicated in the legend to reconstruct the spectrum.

close to ground truth in the perceptually important regions. Going up to six real moments (i.e. $m = 5$), the reconstructions become highly accurate. Eight moments (i.e. $m = 7$) only achieve a minor improvement over six moments. Using eight moments without the warp, the reconstruction quality improves in the region between 650 and 700 nm and near 400 nm but degrades everywhere else.

In Figure 14 we use different tristimulus techniques for relighting of a color checker. Six of the techniques use XYZ as input and are supposed to provide a perfect result under illuminant E. Indeed, they do with the exception of Meng et al. [2015] on the yellow patch where the reconstruction exceeds the value of one and is clipped. If we compute three real Fourier coefficients directly from

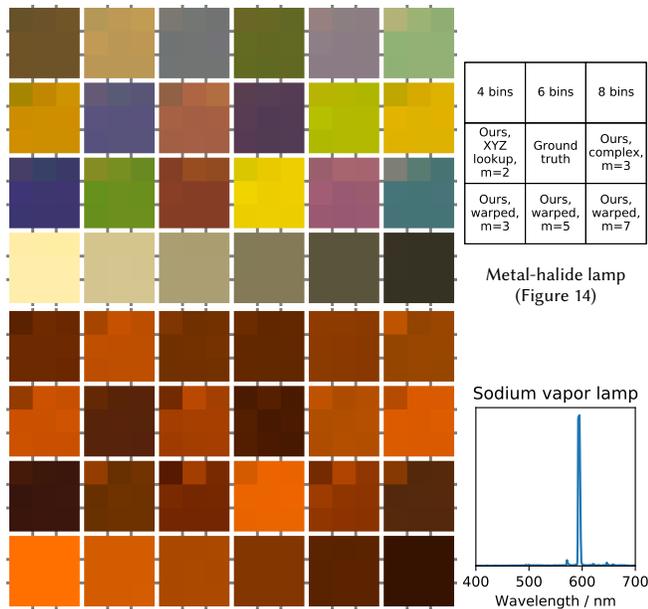


Fig. 15. An X-Rite color checker relit using the emission spectra to the right. The 3×3 squares in each patch use binning of spectra with equidistant bins and different variants of our techniques.

the spectrum, our technique produces minor deviations when using the warp from Section 4.1 but strong deviations without this warp.

For a more spiky emission spectrum, the differences between the techniques are far greater. The three oldest techniques [MacAdam, 1935a, Meng et al., 2015, Smits, 1999] show major errors on many patches. Our techniques with warp and Otsu et al. [2018] give the best results while Jakob and Hanika [2019] is slightly worse. The drawback of not using our lookup table is less pronounced here but the variant without warp still performs poorly. We recommend to always use the warp for $m < 5$.

In Figure 15 we use challenging emission spectra to test techniques with more than three coefficients. Light from the sodium vapor lamp is almost monochromatic around 590 nm and thus it probes the spectra at a single wavelength. While our XYZ lookup table produces some visible errors, the results with four real moments (i.e. $m = 3$) are very close to ground truth. Careful observation reveals further quality improvements when going to six real moments (i.e. $m = 5$). The improvement going from six to eight real moments (i.e. $m = 7$) is miniscule. We also include results for one variant without mirroring, which stores complex moments with $m = 3$. The seven real scalars needed to store its moments yield a quality comparable to four real moments. Under normal circumstances, we do not believe that the lower computational overhead justifies the reduction in quality (see Section 4.1). The use of wavelength bins is clearly inferior to our approach. Even with eight bins the results are often worse than the ones from the XYZ lookup table.

Figure 16 attests to the importance of accurate spectra for inter-reflections. In spite of the simple illuminant, most tristimulus techniques show strong errors on some of the patches. Especially the method of Jakob and Hanika [2019] performs poorly. Once more,

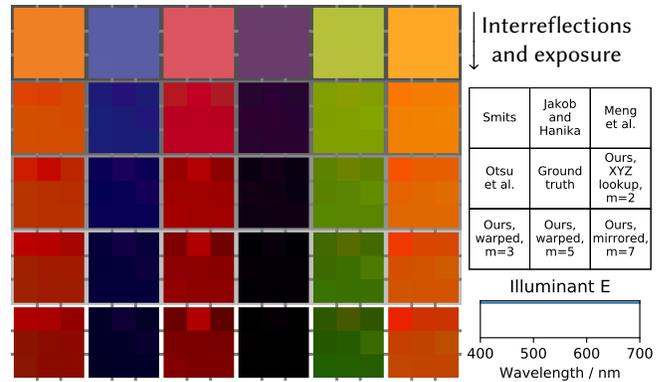


Fig. 16. Reconstructed reflectance for the second row of the X-Rite color checker raised to powers from one (top) to five (bottom) to simulate inter-reflections. The exposure increases by 33% per row. Each patch compares five tristimulus techniques and our techniques with more coefficients.

Table 2. Error metrics for different techniques on the database of 3076 reflectance spectra described in Section 5.1. For each spectrum, we compute the RMSE of the reconstruction from 400 to 700 nm and report the mean and maximal RMSE across all spectra. We also list the mean absolute error.

Technique	Mean error	Mean RMSE	Max. RMSE
MacAdam	$8.8 \cdot 10^{-2}$	$1.2 \cdot 10^{-1}$	$3.6 \cdot 10^{-1}$
Smits	$2.9 \cdot 10^{-2}$	$3.9 \cdot 10^{-2}$	$2.3 \cdot 10^{-1}$
Meng et al.	$2.2 \cdot 10^{-2}$	$3.0 \cdot 10^{-2}$	$4.3 \cdot 10^{-1}$
Otsu et al.	$1.3 \cdot 10^{-2}$	$1.8 \cdot 10^{-2}$	$2.1 \cdot 10^{-1}$
Jakob and Hanika	$3.8 \cdot 10^{-2}$	$6.0 \cdot 10^{-2}$	$2.3 \cdot 10^{-1}$
Ours	XYZ, $m = 2$	$1.8 \cdot 10^{-2}$	$2.6 \cdot 10^{-2}$
	Warped, $m = 2$	$1.8 \cdot 10^{-2}$	$2.5 \cdot 10^{-2}$
	Warped, $m = 3$	$1.2 \cdot 10^{-2}$	$1.9 \cdot 10^{-2}$
	Warped, $m = 5$	$7.6 \cdot 10^{-3}$	$1.4 \cdot 10^{-2}$
	Mirrored, $m = 7$	$5.1 \cdot 10^{-3}$	$8.2 \cdot 10^{-3}$

Otsu et al. [2018] is the best tristimulus technique, closely followed by our technique with the XYZ lookup table. Using up to six coefficients (i.e. $m = 5$) improves quality clearly, whereas eight coefficients (i.e. $m = 7$) have minimal benefit.

Table 2 confirms these findings using the whole reflectance database described in Section 5.1. The technique of Otsu et al. [2018], which was trained on a subset of our test set, performs on par with our technique using four coefficients (i.e. $m = 3$). More coefficients reduce the error as expected. For up to six coefficients (i.e. $m = 5$), the errors compared to ground truth are substantially larger than the errors introduced by 10-bit quantization (Table 1). In this sense, the quantization errors are acceptable.

Figure 17 demonstrates how different techniques behave if we interpolate the coefficients stored in their textures linearly. While the parameters of MacAdam spectra [1935a] change continuously with XYZ, the spectrum itself is discontinuous. Combined with the peaks in the spectrum, discontinuities arise. The technique of Otsu et al. [2018] produces even stronger discontinuities at the boundaries of its clusters. Reconstructions of Jakob and Hanika [2019]

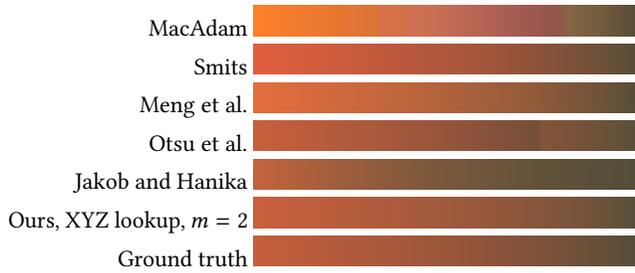


Fig. 17. The impact of linear interpolation of input coefficients for different techniques. In all cases, we perform linear interpolation between the same two colors and light the reconstruction using the metal-halide lamp in Figure 14. For Jakob and Hanika we interpolate the polynomial coefficients, for our technique we interpolate the moments from the lookup table and for all other techniques we interpolate the XYZ triple.



Fig. 18. A rendering from our real-time renderer. All assets have been converted to three real Fourier coefficients at loading time. The renderer samples the metal-halide illuminant from Figure 14 using eight samples per pixel.

change continuously but since the stored polynomial coefficients do not depend upon the spectrum linearly, the change is different from linear interpolation of spectra. Performing the interpolation in XYZ resolves this issue but may require costly lookups during rendering. Interpolation of reconstructed reflectance values is incompatible with hardware-accelerated texture filtering on GPUs. All other techniques, including ours, behave as expected.

5.5 Spectral Real-Time Rendering

Figure 18 shows a rendering of our real-time renderer using converted tristimulus assets for all surfaces. The scene is lit by a challenging illuminant. Nonetheless, the color noise is moderate at eight samples per pixel because we sample proportional to the emission spectrum. Standard temporal and/or spatial reconstruction filters could easily produce noise-free results from this input.

We rendered this image on an NVIDIA RTX 2080 Ti and the total frame time was 0.63 ms at a resolution of 1920×1080. Thus, spectral real-time rendering is entirely feasible with our approach. It is an excellent match for graphics hardware because no instructions execute conditionally, most of the instructions are multiply-add and there are no additional memory reads after reading the moments.

Table 3. Timings for the reconstruction of spectra with different tristimulus techniques. All techniques evaluate the spectrum at 16 wavelengths.

Technique	Time / s	Technique	Time / s
MacAdam	0.48	Otsu et al.	0.94
Smits	0.52	Jakob (AVX)	0.06
Meng et al.	2.0	Ours, $m = 2$, real, Lagr.	0.89

Table 4. Timings for the reconstruction of spectra with different variants of our bounded MESE.

Ours	Time / s							
	Algorithm 1				Lagrange			
	Wavelengths	1	4	8	16	1	4	8
$m = 2$, mirrored	0.2	0.6	0.8	3.7	0.6	0.6	0.8	0.9
$m = 3$, mirrored	0.3	0.8	1.0	4.1	0.8	0.8	1.0	1.1
$m = 5$, mirrored	1.8	2.0	2.0	5.5	2.2	2.3	2.3	2.5
$m = 7$, mirrored	2.8	3.6	7.2	11.8	3.5	3.6	3.6	3.8
$m = 2$, complex	0.5	0.6	0.8	3.8	0.6	0.7	0.8	0.9
$m = 3$, complex	1.3	1.4	1.4	4.6	1.3	1.4	1.4	1.5

5.6 Run Times

We compare CPU run times on a system with an Intel Core i7-8700K having 12 logical cores. To obtain numbers that are not contaminated by other aspects of the rendering, we define a simple benchmark. Each technique has to evaluate the spectrum at multiple wavelengths for each texel of a texture of resolution 4096². We run this benchmark on a single thread and report the total run time.

Table 3 shows the results for all tristimulus techniques. The technique of Jakob and Hanika is by far the fastest. It is designed to be fast to evaluate and additionally it has been implemented with hand-written AVX intrinsics. Among the other techniques, our approach for tristimulus values is in the middle of the field.

Table 4 compares run times of different variants of our technique. We observe that the cost of increasing m is significant, as expected for a technique with time complexity $O(m^2)$. Still, the cost at $m = 5$ or even $m = 7$ is reasonable. The cost per wavelength is only $O(m)$ and thus the increase in cost for evaluating each spectrum at more wavelengths is moderate. To some extent, the compiler successfully applies vectorization but there are instances where it fails, especially for 16 wavelengths. Computing Lagrange multipliers is amortized in the vicinity of eight wavelengths.

Note that performance in our benchmark is limited by compute capability. Since our technique requires no additional memory reads during rendering, its cost may be hidden entirely in bandwidth limited applications.

6 CONCLUSIONS

Our novel representation of spectra offers a practical way to eliminate the high memory and bandwidth requirements of accurate spectral rendering without sacrificing quality. Using only six real moments stored in 8 bytes, scenes are rendered accurately even under the most challenging illuminants. Still, the computational

overhead is low. Among the variants of our technique, there should be an appropriate tradeoff for almost all applications.

These improvements bring spectral rendering closer to widespread adoption and may spark more research on acquisition of spectral data. A common problem of this acquisition is that narrow-band filters lead to long exposure times. In this regard, it could be interesting to take measurements directly in the Fourier domain. It would also be interesting to map colors from the widely used Pantone Matching System to our representation to give countless product designers access to accurate rendering of their colors.

Finally, the bounded MESE is a novel reconstruction for bounded densities with properties unlike any existing reconstruction. Although we have focused on spectra, we believe that it can be a valuable tool within the graphics community and beyond.

ACKNOWLEDGMENTS

We thank Tobias Zirr and Johannes Meng for constructive input and the authors of the spectral databases described in Section 5.1. Nicholas Hull modeled the arcade in Figure 18 for NVIDIA.

REFERENCES

- Jonas Aeschbacher, Jiqing Wu, and Radu Timofte. 2017. In Defense of Shallow Learned Spectral Reconstruction from RGB Images. In *2017 IEEE International Conference on Computer Vision Workshops (ICCVW)*. IEEE, 471–479. <https://doi.org/10.1109/ICCVW.2017.63>
- Laurent Belcour and Pascal Barla. 2017. A Practical Extension to Microfacet Theory for the Modeling of Varying Iridescence. *ACM Trans. Graph. (Proc. SIGGRAPH)* 36, 4 (2017), 65:1–65:14. <https://doi.org/10.1145/3072959.3073620>
- Nir Benty, Kai-Hwa Yao, Tim Foley, Conor Lavelle, and Chris Wyman. 2018. The Falcor Rendering Framework 3.2. <https://github.com/NVIDIAGameWorks/Falcor>
- Steven Bergner, Mark S. Drew, and Torsten Möller. 2009. A Tool to Create Illuminant and Reflectance Spectra for Light-driven Graphics and Visualization. *ACM Trans. Graph.* 28, 1, Article 5 (2009), 5:1–5:11 pages. <https://doi.org/10.1145/1477926.1477931>
- John Parker Burg. 1975. *Maximum Entropy Spectral Analysis*. Ph.D. dissertation. Stanford University, Department of Geophysics. <http://sepwww.stanford.edu/data/media/public/oldreports/sep06/>
- Tenn F. Chen, Gladimir V. G. Baranoski, Bradley W. Kimmel, and Erik Miranda. 2015. Hyperspectral Modeling of Skin Appearance. *ACM Trans. Graph.* 34, 3, Article 31 (2015), 31:1–31:14 pages. <https://doi.org/10.1145/2701416>
- Oskar Elek, Pablo Bauszat, Tobias Ritschel, Marcus Magnus, and Hans-Peter Seidel. 2014. Spectral Ray Differentials. *Computer Graphics Forum* 33, 4 (2014), 113–122. <https://doi.org/10.1111/cgf.12418>
- Luca Fascione, Johannes Hanika, Marcos Fajardo, Per Christensen, Brent Burley, and Brian Green. 2017. Path Tracing in Production - Part 1: Production Renderers. In *ACM SIGGRAPH 2017 Courses (SIGGRAPH '17)*. ACM, Article 13, 13:1–13:39 pages. <https://doi.org/10.1145/3084873.3084904>
- David Geisler-Moroder and Arne Dür. 2010. A New Ward BRDF Model with Bounded Albedo. *Computer Graphics Forum* 29, 4 (2010), 1391–1398. <https://doi.org/10.1111/j.1467-8659.2010.01735.x>
- Andrew S. Glassner. 1989. How to derive a spectrum from an RGB triplet. *IEEE Computer Graphics and Applications* 9, 4 (1989), 95–99. <https://doi.org/10.1109/38.31468>
- Gene H. Golub and Charles F. Van Loan. 2012. *Matrix Computations, Fourth Edition*. The Johns Hopkins University Press. <https://jhupbooks.press.jhu.edu/content/matrix-computations-0>
- Björn Gustafsson and Mihai Putinar. 2017. *Hyponormal Quantization of Planar Domains, Exponential Transform in Dimension Two*. Lecture Notes in Mathematics, Vol. 2199. Springer International Publishing. <https://doi.org/10.1007/978-3-319-65810-0>
- Ville Heikkinen, Reiner Lenz, Tuija Jetsu, Jussi Parkkinen, Markku Hauta-Kasari, and Timo Jääskeläinen. 2008. Evaluation and unification of some methods for estimating reflectance spectra from RGB images. *J. Opt. Soc. Am. A* 25, 10 (2008), 2444–2458. <https://doi.org/10.1364/JOSAA.25.002444>
- Wenzel Jakob and Johannes Hanika. 2019. A Low-Dimensional Function Space for Efficient Spectral Upsampling. *Computer Graphics Forum* 38, 2 (2019).
- Johan Karlsson and Tryphon T. Georgiou. 2013. Uncertainty Bounds for Spectral Estimation. *IEEE Trans. Automat. Control* 58, 7 (2013), 1659–1673. <https://doi.org/10.1109/TAC.2013.2251775>
- Bradley W. Kimmel, Gladimir V. G. Baranoski, T. F. Chen, Daniel Yim, and Erik Miranda. 2013. Spectral Appearance Changes Induced by Light Exposure. *ACM Trans. Graph.* 32, 1 (2013), 10:1–10:13. <https://doi.org/10.1145/2421636.2421646>
- I. V. Kovalishina and Vladimir Petrovich Potapov. 1982. *Integral representation of Hermitian positive functions*. Hokkaido University, Sapporo, Japan. <https://gso.gbv.de/DB=2.1/PPNSET?PPN=014282208> Private translation by Tsuyoshi Ando of a Russian monograph. Copies were gifted to several libraries.
- Boris Kravchenko, Gladimir V. G. Baranoski, Tenn Francis Chen, Erik Miranda, and Spencer R. Van Leeuwen. 2017. High-fidelity iridal light transport simulations at interactive rates. *Computer Animation and Virtual Worlds* 28, 3-4 (2017). <https://doi.org/10.1002/cav.1755>
- Mark Grigorievich Krein and Adol'f Abramovich Nudel'man. 1977. *The Markov Moment Problem and Extremal Problems*. Translations of Mathematical Monographs, Vol. 50. American Mathematical Society.
- Peter Kutz, Ralf Habel, Yining Karl Li, and Jan Novák. 2017. Spectral and Decomposition Tracking for Rendering Heterogeneous Volumes. *ACM Trans. Graph. (Proc. SIGGRAPH)* 36, 4 (2017). <https://doi.org/10.1145/3072959.3073665>
- David L. MacAdam. 1935a. Maximum Visual Efficiency of Colored Materials. *J. Opt. Soc. Am.* 25, 11 (1935). <https://doi.org/10.1364/JOSA.25.000361>
- David L. MacAdam. 1935b. The Theory of the Maximum Visual Efficiency of Colored Materials. *J. Opt. Soc. Am.* 25, 8 (1935). <https://doi.org/10.1364/JOSA.25.000249>
- André Markoff. 1896. Nouvelles applications des fractions continues. *Math. Ann.* 47, 4 (1896), 579–597. <https://doi.org/10.1007/BF01445804>
- Johannes Meng, Florian Simon, Johannes Hanika, and Carsten Dachsbacher. 2015. Physically Meaningful Rendering using Tristimulus Colours. *Computer Graphics Forum* 34, 4 (2015). <https://doi.org/10.1111/cgf.12676>
- Michal Mojzík, Alban Fichet, and Alexander Wilkie. 2018. Handling Fluorescence in a Uni-directional Spectral Path Tracer. *Computer Graphics Forum* 37, 4 (2018), 77–94. <https://doi.org/10.1111/cgf.13477>
- Cedrick Münstermann, Stefan Krumpfen, Reinhard Klein, and Christoph Peters. 2018. Moment-Based Order-Independent Transparency. *Proc. ACM Comput. Graph. Interact. Tech. (Proc. i3D)* 1, 1 (2018), 7:1–7:20. <https://doi.org/10.1145/3203206>
- Hisanari Otsu, Masafumi Yamamoto, and Toshiya Hachisuka. 2018. Reproducing Spectral Reflectances From Tristimulus Colours. *Computer Graphics Forum* 37, 6 (2018), 370–381. <https://doi.org/10.1111/cgf.13332>
- Jong-Il Park, Moon-Hyun Lee, Michael D. Grossberg, and Shree K. Nayar. 2007. Multi-spectral Imaging Using Multiplexed Illumination. In *IEEE 11th International Conference on Computer Vision*. 1–8. <https://doi.org/10.1109/ICCV.2007.4409090>
- Mark S. Peercy. 1993. Linear Color Representations for Full Speed Spectral Rendering. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '93)*. ACM, 191–198. <https://doi.org/10.1145/166117.166142>
- Christoph Peters, Jonathan Klein, Matthias B. Hullin, and Reinhard Klein. 2015. Solving Trigonometric Moment Problems for Fast Transient Imaging. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 34, 6 (2015). <https://doi.org/10.1145/2816795.2818103>
- Christoph Peters and Reinhard Klein. 2015. Moment Shadow Mapping. In *Proceedings of the 19th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (i3D '15)*. ACM, 7–14. <https://doi.org/10.1145/2699276.2699277>
- Christoph Peters, Cedrick Münstermann, Nico Wetstein, and Reinhard Klein. 2017. Improved Moment Shadow Maps for Translucent Occluders, Soft Shadows and Single Scattering. *Journal of Computer Graphics Techniques (JCGT)* 6, 1 (2017), 17–67. <http://jcg.org/published/0006/01/03/>
- Victor Petitjean, Pablo Bauszat, and Elmar Eisemann. 2018. Spectral Gradient Sampling for Path Tracing. *Computer Graphics Forum* 37, 4 (2018). <https://doi.org/10.1111/cgf.13474>
- Michal Radziszewski, Krzysztof Boryczko, and Witold Alda. 2009. An Improved Technique for Full Spectral Rendering. *Journal of WSCG* 17 (2009), 9–16.
- Brian Sharpe. 2018. Moment Transparency. In *Proceedings of the Conference on High-Performance Graphics (HPG '18)*. ACM. <https://doi.org/10.1145/3231578.3231585>
- 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>
- Antoine Toisoul, Daljit Singh Dhillon, and Abhijeet Ghosh. 2018. Acquiring Spatially Varying Appearance of Printed Holographic Surfaces. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 37, 6 (2018). <https://doi.org/10.1145/3272127.3275077>
- Antoine Toisoul and Abhijeet Ghosh. 2017a. Practical Acquisition and Rendering of Diffraction Effects in Surface Reflectance. *ACM Trans. Graph.* 36, 5 (2017). <https://doi.org/10.1145/3012001>
- Antoine Toisoul and Abhijeet Ghosh. 2017b. Real-time Rendering of Realistic Surface Diffraction with Low Rank Factorisation. In *Proceedings of the 14th European Conference on Visual Media Production*. ACM. <https://doi.org/10.1145/3150165.3150167>
- Brian A. Wandell. 1987. The Synthesis and Analysis of Color Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 9, 1 (1987), 2–13. <https://doi.org/10.1109/TPAMI.1987.4767868>
- Alexander Wilkie, Sehera Nawaz, Marc Droske, Andrea Weidlich, and Johannes Hanika. 2014. Hero Wavelength Spectral Sampling. *Computer Graphics Forum* 33, 4 (2014). <https://doi.org/10.1111/cgf.12419>
- Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering Specular Microgeometry with Wave Optics. *ACM Trans. Graph. (Proc. SIGGRAPH)* 37, 4 (2018), 75:1–75:10. <https://doi.org/10.1145/3197517.3201351>