

Reproducing Spectral Reflectances from Tristimulus Colors

H. Otsu, M. Yamamoto, and T. Hachisuka

The University of Tokyo, Japan
 {hotsu, myamamoto}@graphics.ci.i.u-tokyo.ac.jp, thachisuka@siggraph.org




Figure 1: Rendering using reconstructed spectral reflectances. We rendered a scene with Lambertian surfaces illuminated by a black-body illuminant at 5500 K using a spectral rendering system. For spectral reflectances, we use the method by Smits [Smi99], the method by Meng et al. [MSHD15], and our method to reconstruct them from tristimulus colors. Tristimulus colors are from the corresponding reference measured spectra [KPJ06]. The bottom row visualizes errors from the reference rendering using the measured spectra. The existing methods suffer from color shifts in indirect illumination due to the difference between the reconstructed spectra and the measured spectra. Our method faithfully reproduces measured spectra only from tristimulus colors, which has been considered challenging in computer graphics.

Abstract

Physically based rendering systems often support spectral rendering to simulate light transport in the real world. Material representations in such simulations need to be defined as spectral distributions. Since commonly available material data are in tristimulus colors, we ideally would like to obtain spectral distributions from tristimulus colors as an input to spectral rendering systems. Reproduction of spectral distributions given tristimulus colors, however, has been considered an ill-posed problem since single tristimulus color corresponds to a set of different spectra due to metamerism. We show how to resolve this problem using a data-driven approach based on measured spectra and propose a practical algorithm that can faithfully reproduce a corresponding spectrum only from the given tristimulus color. The key observation in color science is that a natural measured spectrum is usually well approximated by a weighted sum of a few basis functions. We show how to reformulate conversion of tristimulus colors to spectra via principal component analysis. To improve accuracy of conversion, we propose a greedy clustering algorithm which minimizes reconstruction error. Using precomputation, the runtime computation is just a single matrix multiplication with an input tristimulus color. Numerical experiments show that our method well reproduces the reference measured spectra using only the tristimulus colors as input.

Keywords: spectral rendering, spectral reflectance reconstruction

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

1. Introduction

Spectral rendering simulates light transport considering spectral radiance per each wavelength. The material description in spectral rendering thus needs to be a full spectrum representation instead of a tristimulus representation (e.g., RGB colors). For example, reflectance becomes spectral reflectance defined over all (visible) wavelengths.

This difference causes two major inconvenient aspects for using spectral rendering in practice. Firstly, a full spectrum representation is more expensive in terms of memory footprint than tristimulus values. Unlike tristimulus colors, spectral values are usually needed

at more than three wavelengths. Secondly, obtaining spectral data can be difficult and time consuming. Commonly available measuring devices and image editing systems often support only tristimulus values. It is thus more practical to use tristimulus values as inputs of a spectral rendering system.

This requirement leads to a problem of obtaining a spectrum given a tristimulus color. The main difficulty is that a tristimulus color can correspond to several spectra. This concept is known as *metamerism* in color science and such spectra are called *metamers*.

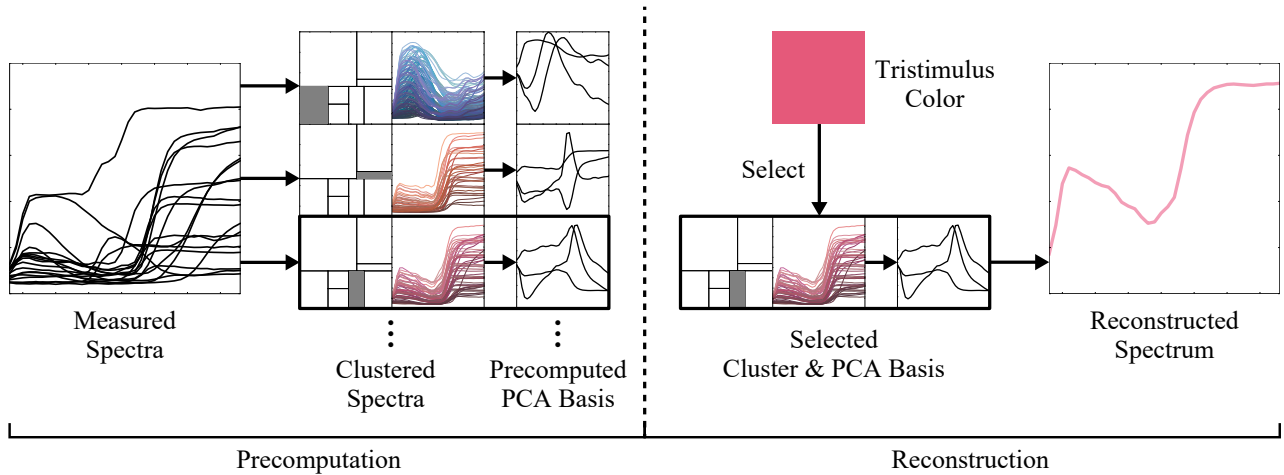


Figure 2: Overview of our method. The algorithm is separated into two phases: the precomputation phase and the runtime phase. In the precomputation phase, measured spectra are separated into several clusters. For each set of clustered spectra, we apply PCA to obtain the basis functions. In the runtime phase, an input tristimulus color is used to determine the corresponding cluster computed in the precomputation phase. Using the basis functions associated with the cluster, we reconstruct a spectrum.

Due to the metamerism, converting a tristimulus color to a spectrum is an ill-posed problem.

Smits [Smi99] showed that smoothness of reconstructed spectra can be used to resolve the metamerism. This heuristic is based on the observation that measured spectra tend to be smooth [Mal86]. A recent work [MSHD15] also uses the same heuristic. Reconstructed spectra with this heuristic, however, have little connection with measured spectra except for being smooth. The difference will show up as visible color shift in spectral rendering due to the presence of interreflections. Iterative multiplications of spectral distributions via interreflections can cause color shift even if reconstructed spectra have the same corresponding tristimulus colors (Fig. 1).

We propose a data-driven approach to reproduce the measured spectra based only on tristimulus colors. The key observation is that principal component analysis (PCA) reveals that typical measured spectra can be accurately reproduced with only a few principle components [Coh64, FB04]. Given measured spectra that satisfy this property, we can reformulate conversion from tristimulus colors to spectra as a well-defined matrix equation by considering the first three components. Unlike the previous methods in computer graphics, our approach does not assume that spectra are smooth. Since we use measured spectra in the conversion process, we are able to reconstruct spectra that match well with actual measured spectra. Unlike similar work in color science [AERN06, HSN16], we use greedy clustering of spectra to directly minimize the difference between the reconstructed spectra and the measured spectra. The runtime algorithm is as simple as one matrix multiplication with an input tristimulus color, which is suitable for spectral rendering. In summary, our contributions are:

- Reformulation of the relationship between tristimulus colors and spectra based on PCA.
- Greedy clustering method which numerically minimizes the reconstruction error.
- Simple runtime algorithm which converts tristimulus colors into spectral reflectances.

2. Related Work

Reconstruction of a spectrum from various information is a major topic in color science. The problem setting varies according to the information available for reconstruction. We focus on related work only on reconstruction from tristimulus colors.

Optimization Approaches. As presented in the pioneering work by Smits [Smi99], conversion from a tristimulus color to a spectrum can be formulated as a numerical optimization problem. This formulation tries to find a spectrum that optimizes a certain cost function such as the difference between converted tristimulus values and input tristimulus values. Smits formulated this problem as a linear optimization of discretized spectra. This method tries to find the smoothest possible spectrum that retains the input tristimulus color which is based on the observation that natural spectra are often smooth [Mal86]. Dupont [Dup02] tested various optimization methods for finding a metameric spectrum similar to the method proposed by Smits, and concluded that the Hawkyard method [Haw93] is among the best.

One issue of this approach is that we need to perform a heavy optimization process for a given input tristimulus color, which is too costly for certain applications. Smits [Smi99] thus proposed to use precomputed spectra only on some primal colors and to rely on linear interpolation for other colors at runtime. Meng et al. [MSHD15] later reported that we in fact need more precomputed samples to cover the XYZ tristimulus color space as input. Our method is free from such interpolation and intrinsically supports inputs from the XYZ color space. Another issue is that reconstructed spectra with this approach can be very different from measured spectra, if the smoothness heuristic does not hold. We avoid using smoothness heuristics via a data-driven approach.

Data-driven Approaches. Since measured spectra data are publicly available, one can also analyze those data and reformulate the conversion process as a problem of finding the relationship between tristimulus colors and measured spectra. Such approaches usually




Figure 3: Approximation (red) of a measured spectrum (dotted) via direct PCA. Although the approximated spectrum is generally similar to the original spectrum, the corresponding colors (left: original, right: reconstructed) are visually different.

express spectra as a weighted sum of basis functions, especially based on PCA. Cohen [Coh64] is the first to perform PCA on measured spectra, and he found that spectral reflectances can be accurately represented by a small number of basis functions. Since then, a weighted sum of basis functions is widely used as a compact representation of measured spectra [PHJ89, FB04, TB05] for the purpose of compressing measured data.

Building upon this idea, if one can find a well-defined relationship between tristimulus colors and the weights of basis functions, we are able to reconstruct spectra from tristimulus colors. While a direct application of PCA for this purpose is appealing, research in color science concluded that it does not work well. Marimont et al. [MW92] thus proposed to use a different set of basis functions including the effect of sensor responses by changing the optimization process in PCA. Other researchers suggested to use variants of PCA such as weighted PCA [AAA08] and adaptive NMF [AA10]. Our method is also built upon PCA, but we propose to cluster spectra in the input color space before performing PCA. Unlike similar work in color science [AERN06, ZX08], our clustering algorithm directly minimizes reconstruction errors at runtime. As far as we know, we are the first to take this data-driven approach in computer graphics.

Abed et al. [AAA09] proposed to use a set of measured spectra directly and interpolate them in the input color space for reconstruction. Some other methods suggest to use non-linear mapping between a tristimulus color and a spectrum such as a feed-forward neural network [SW02], a radial basis function network [NPB14], and nonlinear PCA with an auto-associative neural network [BAPA13]. While these methods improve the accuracy of reconstruction, the tristimulus colors of the reconstructed spectra are not guaranteed to match with the input tristimulus color (called *round-trip error*). Our formulation theoretically guarantees zero round-trip errors (Eq. 9), which is important for applications in computer graphics.

Representation of Spectral Distributions. In addition to discretized samples over wavelengths, one could represent a spectral distribution as a set of basis functions such as polynomials [RF91, SFDC01], Fourier basis [Pee93], Gaussian quadrature [Mey88], or wavelets [CW05]. Since our goal is not to find a representation of spectral distributions, but to generate spectral distributions themselves, these works are orthogonal to ours.

3. Overview

Our goal is to reconstruct a spectral reflectance from a tristimulus color. We determined the following design criteria for this goal.




Figure 4: Accumulated contribution with respect to the number of components for direct PCA applied to the entire dataset (left) and the per-cluster contributions (right) with our clustered PCA. These plots essentially show how accurately the entire dataset is reproduced for each number of components.

1. *Reproduction of spectra:* The reconstructed spectra should match the measured spectra as well as possible.
2. *Recovery of tristimulus colors:* The reconstructed spectrum should represent the input color.
3. *Fast reconstruction:* The runtime computational cost of reconstruction should not be too high.

Note that achieving the criterion #1 does not necessarily mean satisfying the criterion #2 automatically. For example, Fig. 3 shows L2-error minimized reconstruction of spectra with three bases by PCA, but their resulting tristimulus colors are generally different from the corresponding input colors. We thus need to account for these points separately.

Fig. 2 shows the overview of our method. In the precomputation phase, we first cluster a set of measured spectra into several disjoint subsets. The clustering algorithm is designed to minimize the difference between the reconstructed spectra and the measured spectra (criterion #1). For each cluster, we precompute a set of basis functions based on PCA over the measured spectra within the cluster. A direct application of PCA to our problem, however, is impossible since we do not know the original measured spectra at runtime to compute weights for the basis functions. We thus introduce a practical approximation to calculate those weights only from input tristimulus colors. In order to alleviate additional error introduced by this approximation, we also propose to use a clustered PCA with greedy minimization of reconstruction error.

In the reconstruction phase, we first select the corresponding cluster based on the given tristimulus color. We then reconstruct a spectrum as a weighted sum of the precomputed basis functions associated with the selected cluster. This formulation guarantees the recovery of the input tristimulus color after reconstruction (criterion #2). The conversion process is just a matrix-vector multiplication, which is fast enough for many applications (criterion #3).

4. Method

4.1. Tristimulus Colors

A color visible to the human eye can be represented by a spectrum $S(\lambda)$ defined for the wavelength λ . We use the range $\Lambda = [380, 730]$ nm in the following. In order to describe various aspects of the human vision system, many color spaces have been developed. Among them, the CIE XYZ color space is often used as the reference color space among other color spaces. The conversion




Figure 5: Measured spectra plotted with the corresponding colors (top) and the corresponding PCA basis functions (bottom) for each cluster. The thick, normal, and thin lines show the first, second, and third principal components respectively. Our clustering method can group spectra with similar shapes and the basis functions are indeed very different for each cluster.

process from a spectrum $S(\lambda)$ to tristimulus values $\mathbf{c} = [XYZ]^T$ is

$$X = \int_{\Lambda} S(\lambda)\bar{x}(\lambda)d\lambda \quad Y = \int_{\Lambda} S(\lambda)\bar{y}(\lambda)d\lambda \quad Z = \int_{\Lambda} S(\lambda)\bar{z}(\lambda)d\lambda, \quad (1)$$

where \bar{x} , \bar{y} , and \bar{z} are the color matching functions that describe chromatic responses of the standard observer. While our framework is not limited to this choice, we used the analytical approximation of the color matching functions by Wyman et al. [WSS13] for the CIE 1931 standard observer. The RGB color space we used is the sRGB color space (Rec.709) [Uni02].

4.2. Representation of Measured Spectra

We represent a spectrum as a weighted sum of a fixed set of basis functions via PCA. Similar to previous work, we first discretize a spectrum $S(\lambda)$ into n fixed-width bins. The width of each bin in our experiments is $\Delta = 10$ nm ($n = 36$). We denote this discretized spectrum as an n dimensional vector $\mathbf{s} \in \mathbb{R}^n$. Running PCA over a set of measured spectral reflectance data gives us a set of basis functions $\{\mathbf{b}_j\}_{j=1,\dots,n}$. Each measured spectrum under these new basis functions then becomes

$$\mathbf{s} = \sum_{j=1}^n w_j \mathbf{b}_j, \quad (2)$$

where $w_j = \mathbf{b}_j \cdot \mathbf{s}$ is the weight for the j th basis function. Research in color science shows that each measured spectrum can be well approximated by a few basis functions [Coh64, Mal86, FB04].

This observation is confirmed by the accumulated contribution in Fig. 4. In PCA, the accumulated contribution is utilized to assess how much of the information is kept with a given number of components. This value for k -th component is defined as a cumulative sum $\sum_{i=1}^k r_i$ of the explained variance ratio $r_i = \sigma_i / \sum_{j=1}^n \sigma_j$ where σ_j is the variance of the data around j -th principal component. The accumulated contributions quickly approach to one only with a few basis functions regardless of clustering. We thus use only three basis functions in the following.

This choice also means that our method is not applicable to the cases where measured spectra cannot be accurately reconstructed just by three bases. While reproduction of such general measured spectra from tristimulus colors remains challenging, being a data-driven approach, our method allows us to take advantage

of measured spectra unlike the previous approaches in computer graphics [Smi99, MSHD15].

By truncating the sum to three components, we obtain a compact representation of a spectrum \mathbf{s} via three weights w_1, w_2, w_3 using the common three basis functions $\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3$ as

$$\mathbf{s} \approx w_1 \mathbf{b}_1 + w_2 \mathbf{b}_2 + w_3 \mathbf{b}_3. \quad (3)$$

This direct approximation, however, may not recover the input tristimulus color. In other words, the approximated spectrum and the original spectrum \mathbf{s} may not be metameric spectra. Fig. 3 shows an example of this problem: while the approximation is similar to the original spectrum, it represents a different tristimulus color.

This property is not suitable for rendering because we eventually convert the obtained spectrum to tristimulus colors to display final images and the difference might be visually noticeable. We introduce a novel formulation that achieves a compact presentation and the recovery of the input color at the same time. As such, our work is not about merely having a compact representation of measured spectra for given tristimulus colors.

4.3. Reconstruction of Spectra via PCA

We can circumvent the above issue by considering a full conversion process between spectra and tristimulus colors. We start with the relationship between \mathbf{c} and \mathbf{s} as in Eq. 1. By discretizing the color matching functions as n dimensional vectors $\mathbf{x}, \mathbf{y}, \mathbf{z}$, we have

$$X \approx \Delta(\mathbf{x} \cdot \mathbf{s}) \quad Y \approx \Delta(\mathbf{y} \cdot \mathbf{s}) \quad Z \approx \Delta(\mathbf{z} \cdot \mathbf{s}). \quad (4)$$

Instead of representing a spectrum directly with basis functions as Eq. 3, we approximate $\mathbf{s} - \mu$ with basis functions where μ is the average of the measured spectra:

$$\mathbf{s} \approx \tilde{\mathbf{s}} = w_1 \mathbf{b}_1 + w_2 \mathbf{b}_2 + w_3 \mathbf{b}_3 + \mu. \quad (5)$$

This modification is often adopted in PCA [MLBB08], and also used in color science [JMW*64, Bri02, WB04]. Therefore we can approximate the corresponding tristimulus color \mathbf{c} as

$$\mathbf{c} \approx \Delta \begin{bmatrix} \mathbf{x}^T \\ \mathbf{y}^T \\ \mathbf{z}^T \end{bmatrix} \tilde{\mathbf{s}} = \Delta \underbrace{\begin{bmatrix} \mathbf{x}^T \\ \mathbf{y}^T \\ \mathbf{z}^T \end{bmatrix} [\mathbf{b}_1 \mathbf{b}_2 \mathbf{b}_3]}_M \begin{bmatrix} w_1 \\ w_2 \\ w_3 \end{bmatrix} + \mathbf{c}_\mu = M \begin{bmatrix} w_1 \\ w_2 \\ w_3 \end{bmatrix} + \mathbf{c}_\mu, \quad (6)$$




Figure 6: Visualization of the precomputed clusters and measured spectra in the xy -plane. The color of each point is the converted RGB color. The subdivision is done for the entire xy plane, but we zoomed into the region where the measured spectra exist.

where $\mathbf{c}_\mu = \Delta [\mathbf{xyz}]^T \mu$. The matrix M is a 3×3 matrix where each element is defined by a dot product of a basis function and a color matching function. The matrix is a constant given the basis functions.

While each weight is given by $w_j = \mathbf{b}_j \cdot (\mathbf{s} - \mu)$, it is impossible to obtain weights using this definition since the input is \mathbf{c} , not \mathbf{s} . We thus use the (pseudo) inverse of M to approximate the weights as

$$\begin{bmatrix} w_1 \\ w_2 \\ w_3 \end{bmatrix} \approx M^{-1}(\mathbf{c} - \mathbf{c}_\mu) = \begin{bmatrix} \hat{w}_1 \\ \hat{w}_2 \\ \hat{w}_3 \end{bmatrix} \quad (7)$$

and obtain the reconstructed spectrum \mathbf{r} as

$$\mathbf{s} \approx \tilde{\mathbf{s}} \approx \mathbf{r} = \hat{w}_1 \mathbf{b}_1 + \hat{w}_2 \mathbf{b}_2 + \hat{w}_3 \mathbf{b}_3 + \mu. \quad (8)$$

Remind that an approximated spectrum $\tilde{\mathbf{s}}$ via PCA may not recover the input tristimulus color \mathbf{c} . The reconstructed spectrum \mathbf{r} , however, converts back to \mathbf{c} by construction since we have

$$\Delta \begin{bmatrix} \mathbf{x}^T \\ \mathbf{y}^T \\ \mathbf{z}^T \end{bmatrix} \mathbf{r} = \Delta \underbrace{\begin{bmatrix} \mathbf{x}^T \\ \mathbf{y}^T \\ \mathbf{z}^T \end{bmatrix} [\mathbf{b}_1 \mathbf{b}_2 \mathbf{b}_3]}_M \begin{bmatrix} \hat{w}_1 \\ \hat{w}_2 \\ \hat{w}_3 \end{bmatrix} + \mathbf{c}_\mu = M \underbrace{M^{-1}(\mathbf{c} - \mathbf{c}_\mu)}_{[\hat{w}_1 \hat{w}_2 \hat{w}_3]^T} + \mathbf{c}_\mu = \mathbf{c}. \quad (9)$$

Therefore the reconstructed spectrum \mathbf{r} and the original spectrum \mathbf{s} are metamers by construction in this formulation. Since we can precompute \mathbf{b}_1 , \mathbf{b}_2 , \mathbf{b}_3 , M , and M^{-1} , the runtime conversion is just one matrix vector multiplication $M^{-1}(\mathbf{c} - \mathbf{c}_\mu)$ as in Eq. 7. This formulation of conversion from tristimulus colors to spectra no longer relies on smoothness heuristics.

4.4. Hierarchical Chromatic Clustering

While the reconstruction approach above works, Eq. 8 reveals that this approach involves two different approximations: $\mathbf{s} \approx \tilde{\mathbf{s}}$ and $\tilde{\mathbf{s}} \approx \mathbf{r}$. Even though PCA minimizes the L2 error $\|\mathbf{s} - \tilde{\mathbf{s}}\|^2$, due to another approximation $\tilde{\mathbf{s}} \approx \mathbf{r}$, the reconstructed spectrum \mathbf{r} does not necessarily well approximate the original spectrum \mathbf{s} .




Figure 7: Relationship between the number of clusters and the reconstruction error defined in Eq. 13. We can observe that the reconstruction error decreases as the number of clusters increases, eventually becomes flat. Based on this observation, we empirically fixed the number of clusters to eight.

We address this issue by directly minimizing the L2 reconstruction error $\|\mathbf{s} - \mathbf{r}\|^2$. Our idea is inspired by that natural spectra with similar colors tend to be similar [GBNHAR98]. Related work in color science [AERN06, ZX08] suggest that similar approximation via PCA on such spectra indeed improves the reconstruction accuracy, although they did not facilitate the clustering scheme to directly minimize the reconstruction error. Building upon this idea, we cluster the spectra in the xy plane of the CIE xyY color space such that the resulting L2 reconstruction error $\|\mathbf{s} - \mathbf{r}\|^2$ over all the spectra is numerically minimized.

To be concrete, we iteratively split the measured spectra to construct a kD-tree on the xy plane. Each split tries to minimize the L2 reconstruction error $\|\mathbf{s} - \mathbf{r}\|^2$ in a greedy manner. This procedure is inspired by a top-down kD-tree construction with SAH [Hav00]; our approach tries to minimize the L2 reconstruction error, while the top-down kD-tree construction tries to minimize SAH.

Splitting Spectra. We explain the process of splitting a set of spectra $S \equiv \{\mathbf{s} \in \mathbb{R}^n\}$ into two subsets S_1 and S_2 . We determine the splitting plane perpendicular to each axis and compute the splitting position so that the sum of the L2 reconstruction error for each subset is minimized. Given the splitting axis $p \in \{x, y\}$ and the splitting position $v \in [0, 1]$, the subsets S_1 and S_2 can be written as

$$S_1(p, v) = \{\mathbf{s} \in S \mid \mathbf{c} \cdot \mathbf{p} \geq v\}, \quad (10)$$

$$S_2(p, v) = S \setminus S_1(p, v), \quad (11)$$

where $\mathbf{c} \cdot \mathbf{p}$ returns x or y coordinates of the measured spectrum \mathbf{s} in the xyY color space according to the axis p . Using this definition, the splitting axis and position is written as

$$(i^*, v^*) = \underset{i, v}{\operatorname{argmin}} \left(\Delta E(S_1(i, v)) + \Delta E(S_2(i, v)) \right). \quad (12)$$

$\Delta E(S)$ is the sum of the reconstruction error for a set of spectra S :

$$\Delta E(S) = \sum_{\mathbf{s} \in S} \|\mathbf{s} - \mathbf{r}\|^2, \quad (13)$$

where \mathbf{r} is obtained by using the approach described in the last section with PCA over the subset S . We note that the splitting axis

and position is only determined by the optimization process, thus our approach does not depend on the starting condition.

At each leaf node, we store the basis functions, the conversion matrix, and the mean values of the corresponding subset of spectra. At runtime, we traverse the tree using the input tristimulus color and select the corresponding leaf node to reconstruct the spectrum. Since the reconstruction process still recovers input tristimulus colors, a spectrum and a color has a bijective mapping. A set of reconstructed spectra at each leaf node minimizes L2 errors as in Eq. 13.

5. Results

Setup. We used the two datasets of the matte Munsell color chips [KPJ06]. The dataset I is a collection of 1269 measurements ranging from 380 nm to 800 nm with 1 nm interval. The dataset II is a collection of 1250 measurements ranging from 400 nm to 700 nm with 5 nm interval. The precomputation and the evaluation are conducted respectively on the dataset I and II. The range of the vertical axis is $[0, 1]$ for all the plots in this section unless otherwise specified. The range of the horizontal axis is $[380, 730]$ nm or $[400, 700]$ nm respectively for the results using the dataset I or II. We integrated our method into a spectral rendering system based on bidirectional path tracing [LW93, VG94]. All images are rendered with equal time (1 hour) on the same environment. We conducted all the experiments on a machine with an Intel Core i7-5960X at 3.0 GHz using 16 threads. Both the rendered images and visualized colors are in the sRGB color space under the D65 illuminant. The algorithm itself is independent of a choice of the color space and the illuminant. We plan to distribute an implementation of our method.

Clustering of Measured Spectra. Fig. 5 shows the measured spectra in each cluster for the dataset I. Fig. 6 visualizes the distribution of measured spectra and the clusters in the xy-plane for the precomputation using the same dataset. We used eight clusters so that the amount of precomputed data is roughly equal to that of the method by Smits [Smi99]. We also experimentally verified that the reconstruction error does not decrease much for the number of clusters more than eight (Fig. 7).

We can observe that our algorithm successfully clusters spectra that have similar shapes with different overall magnitudes. Ignoring the magnitudes is desirable in our case since the reconstruction via PCA will take care of overall scaling. Fig. 5 also shows the corresponding three PCA basis functions at the bottom. Different clusters have quite different basis functions, validating the use of clustered PCA instead of PCA over an entire dataset.

Reconstructed Spectra. Fig. 8 and Fig. 9 compare several measured spectra [KPJ06] and reconstructed spectra by three different methods; the method by Smits [Smi99] (blue line), the method by Meng et al. [MSHD15] (green line), and ours (red line). The reference RGB color is shown in the top left of each plot. We can observe that the reconstructed spectra with our method match well with the measured spectra in many cases. The two previous methods, while producing metameric spectra like ours, result in very different spectra. It is because real-world spectra are not simply characterized as smooth in general, and a data-driven approach like ours is necessary to capture more complex characteristics of real-world

Dataset	Method	Δe Mean	Δe Med.	Δe Max.
I	Smits	0.1904	0.0819	1.8439
	Meng et al.	0.1633	0.0710	1.1410
	PCA only	0.0388	0.0156	0.9744
	Ours	0.0138	0.0053	0.9875
II	Smits	0.1297	0.0633	1.1825
	Meng et al.	0.1566	0.0985	1.0305
	PCA only	0.0024	0.0014	0.0214
	Ours	0.0022	0.0012	0.0220

Table 1: Reconstruction errors of the four reconstruction methods: Smits [Smi99], Meng et al. [MSHD15], PCA without clustering, and the proposed method. We show the mean, the median, and the maximum of the reconstruction errors for each dataset computed from the reference spectra. In almost all cases, our method is the most accurate.

spectra. Our method is the first to achieve such results in computer graphics. Tab. 1 summarizes the reconstruction errors for each reconstruction method. Δe is the L2 reconstruction error for each reference spectrum $\|\mathbf{s} - \mathbf{r}\|^2$. Our method outperforms the other methods in terms of most metrics in the table.

We note that the purpose of our approach is to achieve a compact representation of the spectrum based on the specific data and the accurate recovery of the input color at the same time. In this sense, because we are not focus on the reconstruction of the general reflectance spectra, our approach can be considered as data compression rather than learning. In other words, we do not focus on applying the precomputed data to reconstruct the features in the different datasets. This explains why, in our experiments, we evaluated the reconstruction methods only with the precomputed data obtained from the same dataset, instead of applying the precomputed data obtained from the different dataset.

Fig. 10 highlights the importance of clustering, illustrating the reconstructed spectra with and without clustering for the dataset I. For the results without clustering, we just apply PCA over all the measured spectra and use the same set of basis functions for any input. While based on the same formulation via PCA (Eq. 8), clustering dramatically improves the accuracy of reconstruction. We found that similar approaches in color science [JMW*64, Bri02, WB04, MLBB08] are not significantly different from our PCA without clustering in terms of error since they do not directly minimize reconstruction error.

Rendering with Reconstructed Spectra. We tested the importance of reproducing of measured spectra in spectral rendering. The scene shown in Fig. 1 is rendered with black-body illuminants of 5500 K. All the materials except for the light source are Lambertian with spectral reflectances. We also render a less naive scene in Fig. 11. The scene is illuminated by black-body illuminants of 4000 K and 5500 K. Several light sources casts indirect lighting to the red wall. In both scenes, the top row shows the resulting images using spectral reflectances produced by the methods of Smits [Smi99], Meng et al. [MSHD15], and ours. The bottom row visualizes per-pixel error from the reference rendering using the measured spectra. We also show the error values (rRMSE) cor-

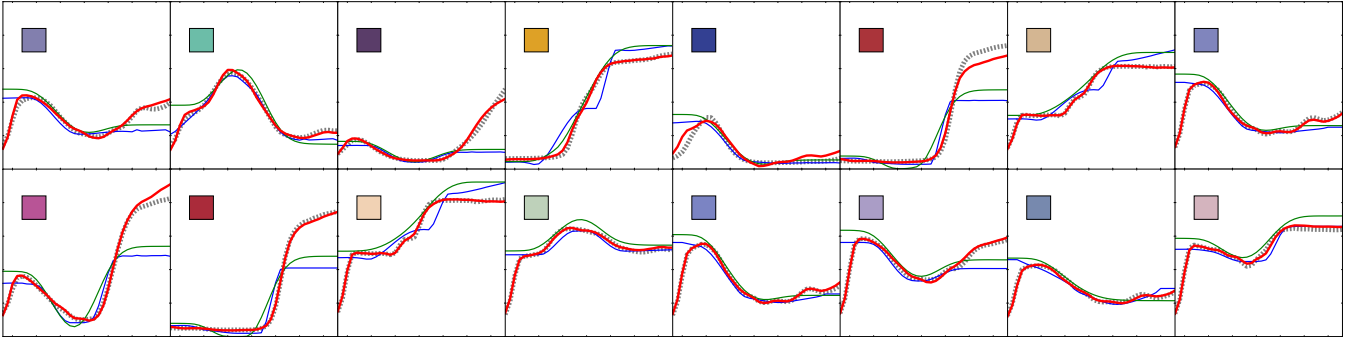


Figure 8: Comparison of the reference spectrum (dotted) and the reconstructed spectra for the dataset I with the method by Smits [Smi99] (blue), the method by Meng et al. [MSHD15] (green), and the proposed approach (red) as well as the RGB color represented by the reference spectrum. All reconstructed spectra are converted from the tristimulus colors of the reference spectra.

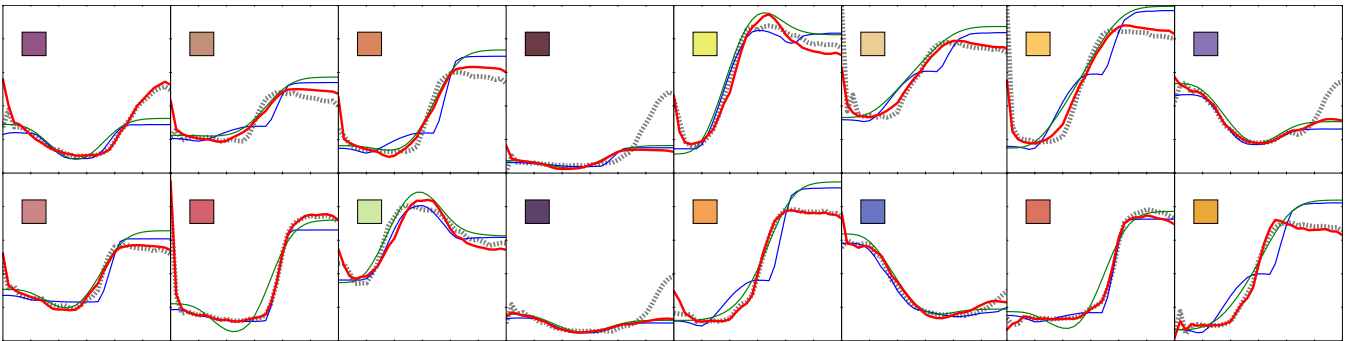


Figure 9: Comparison of the reference and reconstructed spectra for the dataset II in similar configuration to Fig. 8.

responding to the error images. The color map next to the error images is linear according to the error values because we compute the errors *before* tone mapping. We determine the range of the color map between zero and the threshold near the maximum error among the images. The tristimulus color used for reconstruction is obtained from the corresponding measured spectrum. We can observe that the final color in the rendered images with the previous reconstruction methods suffer from color shifts especially for the indirect illumination. This result implies that different metameric spectra can in fact lead to visible differences because the reflectances are multiplied by many times. This observation underlines the importance of accurate reproduction of measured spectra.

We also show rendered images with different number of bounces using the three different methods (Fig. 12). Similarly, we also show the error images obtained by the comparison with the references using the corresponding measured spectra for each path length. We can observe the similar tendency that our method can generate the better results.

Fig. 13 shows how the number of clusters affects the rendered images. The bottom line shows the error compared to the reference obtained with the measured spectra corresponding to the input tristimulus color. Due to the inaccurate reconstruction, the rendered image with a small number of clusters suffers from the color shift compared to the reference. We can observe that the color shift can be alleviated by adopting sufficient number of clusters.

6. Limitations

Reconstruction in Cluster Boundaries. Fig. 14 illustrates the change of the reconstructed spectra in the cluster boundaries. The left-most plot shows the xy plane with cluster boundaries. The sampled points on each of lines (a) - (c) correspond to the set of the spectra in the right three sets of spectra. Fig. 15 shows some rendered images for the closest two samples in the lines (b) and (c) within the difference of 10^{-5} in the value of x or y. Since our clustering algorithm does not explicitly consider the continuity across the boundaries, we found that the shapes of the spectra do not necessarily vary smoothly across the boundaries. For example, while the cases (a) and (b) show smooth changes, the case (c) shows some discontinuous changes in the shape of the spectrum. For the case (c), the visual difference in the rendered images (Fig. 15, bottom) is small, but noticeable for interreflection.

Reconstruction of Saturated Colors. To obtain physically plausible results in spectral rendering, spectral reflectances must be in the range of $[0, 1]$ for every wavelength. Our formulation using PCA does not guarantee this property. That is, the reconstructed spectrum can be less than zero or more than one at certain wavelengths. We thus need to clamp the output spectra to the range of $[0, 1]$ to guarantee this property. Clamped spectra, however, may not convert back to the input tristimulus color. Fig. 16 illustrates some examples of the case. We reconstructed the spectrum from the saturated colors, constraining each component of the RGB color to




Figure 10: Comparison of the reconstructed spectra with (red) and without (blue) clustering for the dataset I. Similar to Fig. 8, we used the measured spectra for the reference (dotted), and reconstructed them. We can observe that clustering improves the accuracy of reconstruction, especially for longer wavelengths.




Figure 11: Rendering using reconstructed spectral reflectances with a less naive scene. The scene contains Lambertian surfaces and illuminated by black-body illuminants at 4000K and 5500K. Several light sources are placed behind the sofa. Similar to Fig. 1, we render the images with three different methods and visualize the error images as well as the corresponding error values, compared to the reference obtained with the measured spectra.

Dataset	Method	Δe Mean	Δe Med.	Δe Max.
Glossy	Smits	0.0050	0.0023	0.0523
	Meng et al.	0.0061	0.0028	0.0484
	PCA only	0.0014	0.0006	0.0441
	Ours	0.0006	0.0002	0.0455
Natural Colors	Smits	0.0057	0.0036	0.0643
	Meng et al.	0.0054	0.0037	0.0602
	PCA only	0.0027	0.0017	0.0285
	Ours	0.0021	0.0013	0.0223

Table 2: Reconstruction errors of the four reconstruction methods similar to Tab. 1 for the two different datasets: glossy and natural colors [KPJ06].

either zero or one. The values of the spectra reconstructed from such colors tend to be outside the range of $[0, 1]$. Meng et al. [MSHD15] pointed out that some colors are impossible to be represented within the range of $[0, 1]$, and suggested some adjustment techniques to produce physically plausible spectra in such cases. Since our formulation intrinsically supports the XYZ color space, we can use their technique.

We also visualize the ratio of the successful reconstruction in the range of $[0, 1]$ in the xy -plane (Fig. 17). The region enclosed by the solid line shows the range of visible chromacities and the region enclosed by the dotted line shows the contour line of the measured spectra for the dataset I. The color map represents the ratio of the successful reconstruction. The color is assigned for each bin representing a small part in the xy -plane and the ratio is estimated by computing the number of successful and unsuccessful reconstructions for each bin. The red color means all reconstructed

spectra are in the range of $[0, 1]$ and the blue color represents the spectra are outside the range of $[0, 1]$, or the visible region. We can observe that the reconstruction succeeds in a large part of the domain defined by the measurements yet fails in the outside of the domain.

The maximum error in Tab. 1 shows that there are cases where our method still fails to reproduce measured spectra. We do not claim that our method is capable of reproducing *any* natural spectra since some natural spectra can have very different characteristics than typical ones [WS00]. Our method is certainly not perfect and metamerism may still occur if we consider all possible spectra. In the paper, we tested our method with the measured spectra of Munsell Matte color chips. While those spectra certainly do not represent *all* the natural spectra, our method is still applicable when measured spectra has a few degree of freedoms as was confirmed firmly in color science [Coh64, Mal86, FB04] for different datasets. Our method thus can be seen as a general approach to find a compact and reversible (in the sense that it recovers the input colors) representation of a given measured spectra.

For example, if measured spectra for various metals are given and we use the input color to specify the color of a metal, our method reproduces spectra of various metals, *not*, Munsell Matte color chips. Tab. 2 shows such examples where the reconstruction errors are calculated for the two different datasets: glossy and natural colors [KPJ06]. Our method indeed generates more accurate results than the other existing methods even in these examples.

Reconstruction of General Spectra. Our method does not support spectra with intrinsic distributions such as the spectra measured from the material with the fluorescence or the structural colors. For such spectra, principle component analysis might not lead to an efficient approximation. Our method does not directly support emission

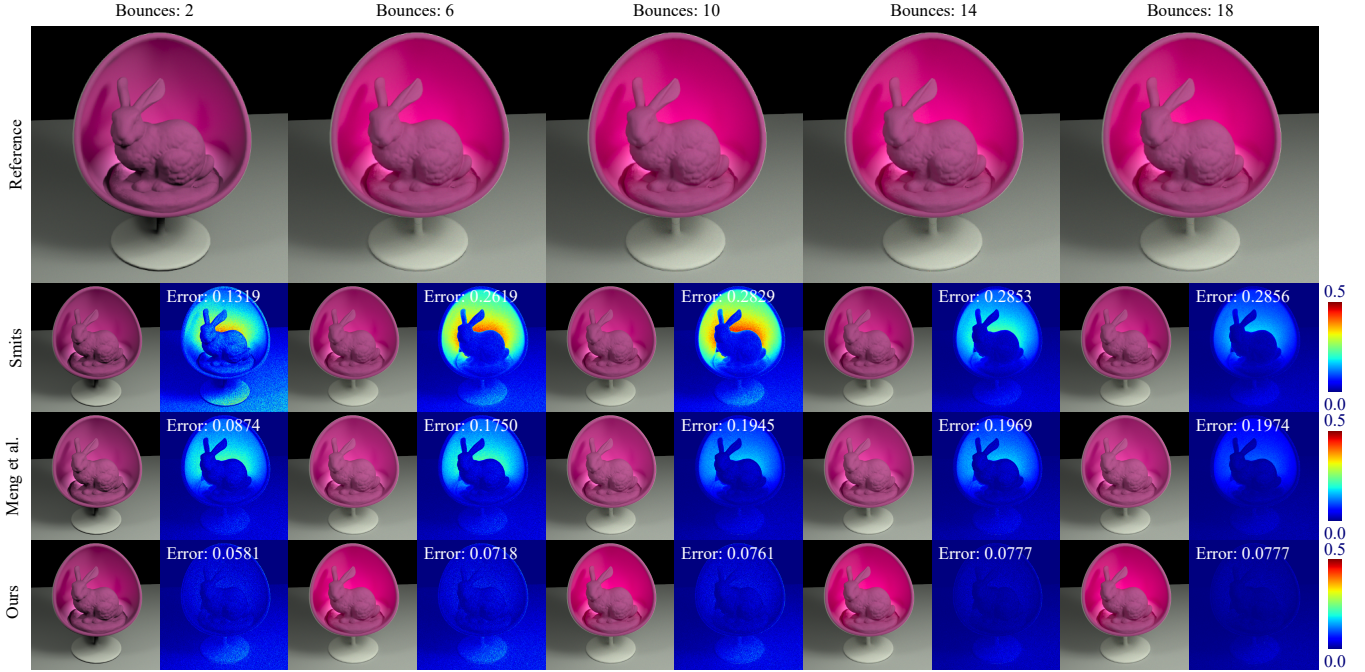


Figure 12: Rendering with different number of bounces using the spectra converted from the same tristimulus color with the three different methods: the method by Smits [Smi99], the method by Meng et al. [MSHD15], and our method. As well as the rendered images we also show error images and the error values corresponding to the images. We note that, because we use different references for each column, the error images and the error values are irrelevant with respect to the different number of path length, although we used the same error scale for the visualization.

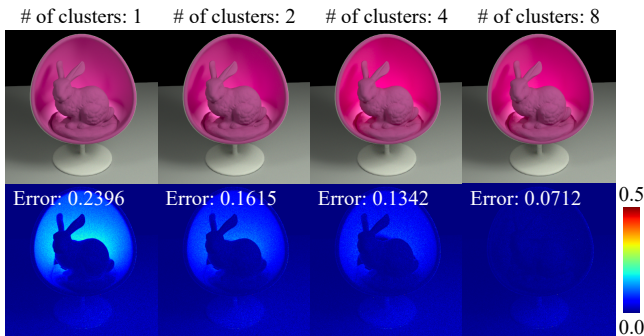


Figure 13: Rendered images using the reconstructed spectra from the multiple number of clusters, converted from the same tristimulus color. The top row shows the rendered images using the spectrum with the corresponding number of clusters. The bottom row shows the pixel-wised error compared to the reference rendered with the original spectrum corresponding to the tristimulus color, as well as the error values.




Figure 14: Some sets of spectra across cluster boundaries. The left-most plot is the xy plane with three different lines across the cluster boundaries. The right three figures are corresponding reconstructed spectra to the samples points on the lines. We assigned the different line types (solid, dashes) for the spectra in the different clusters.

7. Conclusion

We introduced a novel method to reproduce spectral reflectances from tristimulus colors. Our key idea is to use prior knowledge on the actual measured spectra as studied in color science. Unlike existing methods in computer graphics, the use of measured spectra allows us to formulate the conversion process without smoothness heuristics. Our formulation also clarifies that a similar formulation in color science would have uncontrolled approximation error. In order to address this problem, we proposed a clustering technique to directly minimize the approximation error. At runtime, an input tristimulus color is converted to a spectrum using the precomputed basis functions and the conversion matrix to weights for those basis functions. Our experiments demonstrate that the proposed method can faithfully reproduce measured spectra without relying on exis-

spectra either. One technique to use emission spectra within our method is scaling down the input emission tristimulus color into $[0, 1]^3$, converting it to an emission spectrum using our method, then scaling up the resulting spectrum. This technique, however, is not guaranteed at all to reproduce a measured emission spectrum with the same emission tristimulus color.




Figure 15: Rendered images using the reconstructed spectra with the two consecutive colors for the lines (b) and (c) in Fig. 15. We selected the consecutive colors so that the differences of x or y values are within 10^{-5} . We rendered the scene containing indirect illumination with the spectral reflectances using two reconstructed spectra from the two colors. The right-most images are the difference of two images.

ting heuristics or heavy optimizations at runtime. We expect that our method can be used for many different applications with spectral data. In addition, since our work introduces the mapping between RGB values a certain set of measured spectra, it might lead to more accurate color reproduction in printing, given RGB images.

Acknowledgment

We would like to thank Hajime Uchimura at Polyphony Digital Inc. for his insightful comments on the draft of this paper. We thank the blendswap.com artist Wig42 for the *living room* scene.

References

- [AA10] AMIRSHAHI S. H., AMIRSHAHI S. A.: Adaptive non-negative bases for reconstruction of spectral data from colorimetric information. *Optical Review* 17, 6 (2010), 562–569. 3
- [AAA08] AGAHIAN F., AMIRSHAHI S. A., AMIRSHAHI S. H.: Reconstruction of reflectance spectra using weighted principal component analysis. *Color Res. Appl.* 33, 5 (2008), 360–371. 3
- [AAA09] ABED F. M., AMIRSHAHI S. H., ABED M. R. M.: Reconstruction of reflectance data using an interpolation technique. *J. Opt. Soc. Am. A* 26, 3 (2009). 3
- [AERN06] AYALA F., ECHAVARRI F., RENET P., NEGUERUELA A. I.: Use of three tristimulus values from surface reflectance spectra to calculate the principal components to reconstruct these spectra by using only three eigenvector. *J. Opt. Soc. Am. A* 23, 8 (2006), 2020–2026. 2, 3, 5
- [BAPA13] BARAKZEHI M., AMIRSHAHI S. H., PEYVANDI S., AFJEH M. G.: Reconstruction of total radiance spectra of fluorescent samples by means of nonlinear principal component analysis. *J. Opt. Soc. Am. A* 30, 9 (2013), 1862–1870. 3
- [Bri02] BRILL M. H.: A non-PC look at principal components. *Color Res. Appl.* 28, 1 (2002). 4, 6
- [Coh64] COHEN J.: Dependency of the spectral reflectance curves of the munsell color chips. *Psychonomic Science* 1, 1 (1964), 369–370. 2, 3, 4, 8
- [CW05] CHERN J. R., WANG C. M.: A novel progressive refinement algorithm for full spectral rendering. *Real-Time Imaging* 11, 2 (2005), 117–127. 3
- [Dup02] DUPONT D.: Study of the reconstruction of reflectance curves based on tristimulus values: Comparison of methods of optimization. *Color Res. Appl.* 27, 2 (2002), 88–99. 2
- [FB04] FAIRMAN H. S., BRILL M. H.: The principal components of reflectances. *Color Res. Appl.* 29, 2 (2004), 104–110. 2, 3, 4, 8
- [GBNHAR98] GARCÍA-BELTRÁN A., NIEVES J. L., HERNÁNDEZ-ANDRÉS J., ROMERO J.: Linear bases for spectral reflectance functions of acrylic paints. *Color Res. Appl.* 23, 1 (1998). 5
- [Hav00] HAVRAN V.: *Heuristic Ray Shooting Algorithms*. Ph.d. thesis, Department of Computer Science and Engineering, Faculty of Electrical Engineering, Czech Technical University in Prague, November 2000. 5
- [Haw93] HAWKYARD C.: Synthetic reflectance curves by additive mixing. *Journal of the Society of Dyers and Colourists* 109, 10 (1993), 323–329. 2
- [HSN16] HAJIPOUR A., SHAMS-NATERI A.: Effect of classification by competitive neural network on reconstruction of reflectance spectra using principal component analysis. *Color Res. Appl.* (2016). 2
- [JMW*64] JUDD D. B., MACADAM D. L., WYSZECKI G., BUDDE H. W., CONDIT H. R., HENDERSON S. T., SIMONDS J. L.: Spectral distribution of typical daylight as a function of correlated color temperature. *J. Opt. Soc. Am. A* 54, 8 (1964), 1031–1040. 4, 6
- [KPJ06] KOHONEN O., PARKKINEN J., JAASKELAINEN T.: Databases for spectral color science. *Color Res. Appl.* 31, 5 (2006). 1, 6, 8
- [LW93] LAFORTUNE E., WILLEMS Y. D.: Bi-directional path-tracing. In *Proceedings of Third International Conference on Computational Graphics and Visualization Techniques (Compugraphics '93)* (Alvor, Portugal, 1993), pp. 145–153. 6
- [Mal86] MALONEY L. T.: Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *J. Opt. Soc. Am. A* 3, 10 (1986), 1673–1683. 2, 4, 8
- [Mey88] MEYER G. W.: Wavelength selection for synthetic image generation. *Comput. Vision Graph. Image Process.* 41, 1 (1988), 57–79. 3
- [MLBB08] MIRANDA A. A., LE BORGNE Y.-A., BONTEMPI G.: New routes from minimal approximation error to principal components. *Neural Processing Letters* 27, 3 (2008), 197–207. 4, 6
- [MSHD15] MENG J., SIMON F., HANIKA J., DACHSBACHER C.: Physically meaningful rendering using tristimulus colours. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)* 34, 4 (2015), 31–40. 1, 2, 4, 6, 7, 8, 9
- [MW92] MARIMONT D. H., WANDELL B. A.: Linear models of surface and illuminant spectra. *J. Opt. Soc. Am. A* 9, 11 (1992), 1905–1913. 3
- [NPB14] NGUYEN R. M. H., PRASAD D. K., BROWN M. S.: Training-based spectral reconstruction from a single rgb image. In *Computer Vision - ECCV 2014* (Zurich, Switzerland, 2014), pp. 186–201. 3
- [Pee93] PEERCY M. S.: Linear color representations for full speed spectral rendering. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques* (1993), SIGGRAPH '93, pp. 191–198. 3
- [PHJ89] PARKKINEN J. P. S., HALLIKAINEN J., JAASKELAINEN T.: Characteristic spectra of munsell colors. *J. Opt. Soc. Am. A* 6, 2 (1989), 318–322. 3
- [RF91] RASO M., FOURNIER A.: A piecewise polynomial approach to shading using spectral distributions. In *Graphics Interface* (Calgary, Alberta, Canada, 1991), pp. 40–46. 3
- [SFD01] SUN Y., FRACCHIA F. D., DREW M. S., CALVERT T. W.: A spectrally based framework for realistic image synthesis. *The Visual Computer* 17, 7 (2001), 429–444. 3
- [Smi99] SMITS B.: An rgb-to-spectrum conversion for reflectances. *Journal of Graphics Tools* 4, 4 (1999), 11–22. 1, 2, 4, 6, 7, 9
- [SW02] SHARMA G., WANG S.: Spectrum recovery from colorimetric data for color reproductions. *Color Imaging: Device-Independent Color, Color Hardcopy, and Applications VII. Proc. SPIE 4663* (2002), 8–14. 3




Figure 16: Comparison of the reconstructed spectra for the saturated colors with the method by Smits (blue), the method by Meng et al. (green), and our approach (red). We reconstructed the spectrum from the eight combinations of the saturated colors whose components of the color are equal to one or zero. The dotted lines illustrate the range of $[0, 1]$. The color plate in each plot shows the color converted from the reference spectrum.




Figure 17: Visualization of the ratio of the reconstructed spectra in the range of $[0, 1]$ in the xy -plane. The solid line shows the region of the visible spectra and the dotted line shows the contour line of the measured spectra in the dataset I. We separated the xy -plane into a set of small regions and assigned the color for each region according to the estimated ratio of the spectra in the range.

- [TB05] TZENG D.-Y., BERNS R. S.: A review of principal component analysis and its applications to color technology. *Color Res. Appl.* 30, 2 (2005), 84–98. 3
- [Uni02] UNION I. T.: Recommendation ITU-R BT.709: Parameter values for the hdtv standards for production and international programme exchange, 2002. 4
- [VG94] VEACH E., GUIBAS L. J.: Bidirectional estimator for light transport. *Proc. Eurographics Symposium on Rendering* (1994), 147–162. 6
- [WB04] WORTHEY J. A., BRILL M. H.: Principal components applied to modeling: Dealing with the mean vector. *Color Res. Appl.* 29, 4 (2004). 4, 6
- [WS00] WYSZECKI G., STILES W. S.: *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd ed. Wiley-Interscience, 2000. 8
- [WSS13] WYMAN C., SLOAN P.-P., SHIRLEY P.: Simple analytic approximations to the CIE XYZ color matching functions. *Journal of Computer Graphics Techniques* 2, 2 (2013), 1–11. 4
- [ZX08] ZHANG X., XU H.: Reconstructing spectral reflectance by dividing spectral space and extending the principal components in principal component analysis. *J. Opt. Soc. Am. A* 25, 2 (2008), 371–378. 3, 5