Spectral Reconstruction with Uncertainty Quantification via Differentiable Rendering and Null-Space Sampling

MENGQI (MANDY) XIA, Yale University, USA BAI XUE, Yale University, USA RACHEL LIANG, Yale University, USA HOLLY RUSHMEIER, Yale University, USA

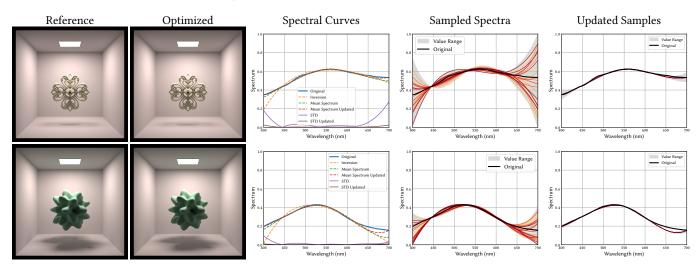


Fig. 1. This figure demonstrates reflectance spectrum reconstruction from an RGB image. We first obtain a candidate spectrum through differentiable rendering, then apply our null-space sampling method to generate multiple spectra that produce the same RGB color. From these samples, we compute both the mean spectrum and the spectral standard deviation across wavelengths. Our sampling method enables uncertainty quantification in the spectral reconstruction and guides additional measurements that significantly reduce uncertainty and improve reconstruction accuracy.

Spectral information plays a crucial role in many domains, including remote sensing, cultural heritage analysis, food inspection, and material appearance modeling. Spectral measurements, such as hyperspectral imaging, provide a powerful means of acquiring this information but often require expensive equipment and time-consuming capture procedures.

We propose a new method for recovering spectral information from multispectral images using differentiable rendering, which naturally incorporates 3D geometry and light transport. However, the inverse problem is ill-posed: conventional pipelines produce a single spectrum that may differ significantly from the ground truth. To address this ambiguity, we introduce a spectral upsampling framework based on null-space sampling, which generates multiple candidate spectra consistent with the input multi-band image. This enables uncertainty quantification across wavelengths and informs the design of additional measurements to improve reconstruction. We

Authors' addresses: Mengqi (Mandy) Xia, Yale University, USA, mengqi.xia@yale.edu; Bai Xue, Yale University, USA, bai.xue@yale.edu; Rachel Liang, Yale University, USA, rachel.liang@yale.edu; Holly Rushmeier, Yale University, New Haven, USA, holly.rushmeier@yale.edu.

Please use nonacm option or ACM Engage class to enable CC licenses
This work is licensed under a Creative Commons Attribution 4.0 International License.

SA Conference Papers '25, December 15–18, 2025, Hong Kong, Hong Kong

2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2137-3/2025/12

https://doi.org/10.1145/3757377.3763918

also demonstrate how to incorporate interreflections into the algorithm to enhance reconstruction accuracy.

We validate our method on synthetic scenes using real-world spectral data and RGB renderings, and demonstrate its effectiveness in physical experiments. Our approach not only avoids the cost and complexity of hyperspectral imaging, but also significantly accelerates the reconstruction process compared to brute-force methods that treat each wavelength independently. Moreover, it supports spectral material authoring by generating diverse, physically plausible spectra from a single RGB input, enabling greater flexibility and artistic control in spectral rendering.

CCS Concepts: • Computing methodologies → Rendering.

ACM Reference Format:

Mengqi (Mandy) Xia, Bai Xue, Rachel Liang, and Holly Rushmeier. 2025. Spectral Reconstruction with Uncertainty Quantification via Differentiable Rendering and Null-Space Sampling. In SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25), December 15–18, 2025, Hong Kong, Hong Kong. ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3757377. 3763918

1 INTRODUCTION

Spectral information is crucial for studying material substances, with wide-ranging applications across science and engineering, including analytical chemistry [Harris 2010], remote sensing [Chang

2013], cultural heritage preservation [Kim et al. 2014], and biomedical research [Offerhaus et al. 2019]. In computer graphics, spectral rendering has proven essential for achieving higher realism and is increasingly adopted in both research and production renderers [Weidlich et al. 2021]. Representing materials with spectral reflectance allows their appearance to be more accurately preserved.

Acquiring spectral properties of real-world materials is challenging. Point-wise spectroscopy works well for spatially uniform materials but falls short in scenarios with spatial variation. In such cases, hyperspectral imaging (HSI) becomes essential, capturing hundreds of contiguous spectral bands that encode both spectral and spatial information. However, HSI systems are expensive, slow to operate, and require complex calibration procedures.

To obtain spectral information more efficiently, a recent trend in the computer vision community is to infer spectra from RGB images. This line of research generally falls into two categories: prior-based methods and deep learning-based methods. Prior-based approaches rely on smaller datasets and use known spectral response functions (SRFs) along with hand-crafted priors that capture the statistical properties of the data. In contrast, data-driven methods often achieve higher accuracy but require large amounts of training data to learn the underlying mapping. We refer readers to [Zhang et al. 2022] for a comprehensive survey on this topic.

Spectral information extraction methods from the vision and HSI communities operate only at the image level and often lack physical interpretability, as illumination is entangled with material appearance and global illumination effects are not properly handled. In contrast, we propose a novel approach that leverages differentiable rendering to reconstruct reflectance spectra from multispectral or RGB images. As a physically grounded framework, differentiable rendering naturally accounts for global illumination and enables the joint reconstruction of shape, material properties, and other scene parameters. Moreover, interreflections embedded in the rendering process help disambiguate the spectral reconstruction problem.

Reconstructing spectra from multispectral input is inherently ill-posed, as a single multi-band observation can correspond to infinitely many spectral solutions—a phenomenon known as metamerism. Existing differentiable rendering frameworks yield only one estimate, which may still deviate from the ground truth despite global illumination constraints. To address this, we propose a null-space sampling method that generates a family of plausible spectra, enabling uncertainty quantification and guiding additional measurements to improve reconstruction.

Our method is versatile, supporting the integration of light spectra, additional priors, and interreflection constraints—all of which improve reconstruction. Since reflectance spectra are typically smooth, with higher-order details exhibiting lower amplitudes, we use B-spline functions for compact representation. To avoid ill-posedness while capturing fine details, we introduce an adaptive basis refinement strategy: starting with a coarse basis to estimate the primary structure, we then refine the residual using a finer basis.

We validate our method on synthetic scenes using real-world spectral data and RGB renderings, and further demonstrate its effectiveness through physical experiments. Compared to hyperspectral imaging, our approach avoids complex acquisition and significantly

speeds up reconstruction over brute-force methods that optimize each wavelength independently.

Concretely, our main contributions are:

- A constrained null-space sampling method that generates multiple plausible reflectance spectra from differentiable rendering outputs, enabling uncertainty quantification and guiding additional measurements.
- Enhancement techniques that incorporate light spectra, interreflections, and adaptive basis refinement to improve reconstruction accuracy.
- A method which can be used for spectral upsampling, providing greater flexibility and artistic control in spectral rendering.

2 RELATED WORK

Our work combines spectral reconstruction and differentiable rendering and we will review related work from these two perspectives.

2.1 Spectral Reconstruction and Spectral Uplifting

Image-based Spectral Reconstruction. Many methods for recovering lost spectral information efficiently from RGB images in computer vision rely on representation learning [Aeschbacher et al. 2017; Akhtar and Mian 2018; Arad and Ben-Shahar 2016; Duan et al. 2024; Geng et al. 2019; Jia et al. 2017], where a spectral dictionary is built from HSI data and projected into RGB space using SRFs. The resulting RGB dictionary is then used to reconstruct HSI from RGB inputs. Recent advances in deep learning have further improved reconstruction quality [Alvarez-Gila et al. 2017; Can and Timofte 2018; Li et al. 2020; Stiebel et al. 2018; Xiong et al. 2017]. While these data-driven approaches achieve high fidelity, they require large HSI training datasets and lack physical interpretability. Another related problem in HSI is extracting the pure spectral signatures of objects from hyperspectral inputs [Chang 2013, p. 13], [Xue et al. 2017; Zhao et al. 2014]. All these methods operate solely on 2D images, whereas our approach incorporates 3D geometry through differentiable rendering.

Spectral Uplifting in Computer Graphics. Spectral rendering has been adopted in research and production renderers [Fascione et al. 2018; Jakob et al. 2022; Pharr et al. 2023; Weidlich et al. 2022]. It is also essential to render wave optical effects [Steinberg 2023; Xia et al. 2023b, 2020, 2023a; Yan et al. 2018; Yu et al. 2023]. Using spectral renderers requires converting RGB assets into spectral representations. Early spectral uplifting work includes pioneering studies by MacAdam[1935], Peercy[1993], and Smits[1999]. Recent efforts have explored compact spectral representations [Jakob and Hanika 2019; Meng et al. 2015; Otsu et al. 2018; Peters et al. 2019]. Tódová et al. [2022] introduced constrained uplifting, focusing on appearance under varying illuminants, while van de Ruit and Eisemann [2024] addressed interactions between multiple light sources and indirect illumination. Belcour et al. [2023] proposed a technique that generates a set of spectral solutions from an RGB input. While their approach is the closest to ours, our method generalizes to multispectral input beyond RGB and avoids failure cases where their greedy algorithm produces no valid solution.

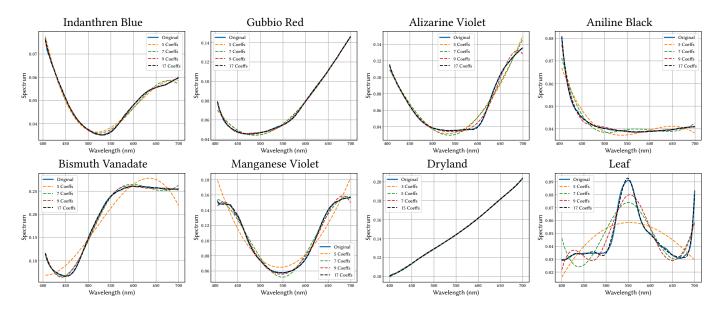


Fig. 2. We demonstrate the ability of B-splines to represent real-world spectra by fitting paints of various colors, as well as dry land and leaf spectra, using different numbers of coefficients. The data are from Kremer Pigments [Deborah 2022] and EcoSIS plant and forest spectra [Wagner et al. 2018].

2.2 Differentiable Rendering

Differentiable rendering is a powerful approach for jointly reconstructing shape, material properties, and other scene parameters while accounting for full light transport. Early works [Bangaru et al. 2020; Li et al. 2018; Loubet et al. 2019; Zhang et al. 2019, 2021] focused on computing derivatives in various contexts. Subsequent research [Nimier-David et al. 2020; Vicini et al. 2021; Xu et al. 2023; Yan et al. 2022; Yu et al. 2023; Zhang et al. 2023] has concentrated on improving the efficiency of derivative estimation. One challenge in differentiable rendering, and optimization problems in general, is exploring multiple optima in the solution landscape. Our work addresses this issue in the specific context of spectral reconstruction from multispectral images.

3 METHOD

Our goal is to reconstruct the reflectance spectrum or spectra $S(\lambda)$ of one or more materials from multispectral image captures. We first review the spectrum-to-multiband projection and B-spline representation, and then introduce our spectral reconstruction algorithm based on null-space sampling.

Spectral Integration and B-Spline Representation

Assume we are given a multispectral imaging system with K spectral bands. The response v_i for each band i is computed as

$$v_i = \int_{\Lambda} S(\lambda) M_i(\lambda) d\lambda, \quad i = 1, ..., K$$
 (1)

where $M_i(\lambda)$ is the sensitivity function for the *i*th band. For RGB renderings and captures, M refers to the color matching functions and the camera sensitivity functions respectively, and K = 3. The wavelength range Λ lies within the visible spectrum, where we use 400-700nm in practice.

We represent the spectrum using B-splines because they are naturally smooth and allow the reflectance spectrum to be easily constrained to the range [0, 1] by bounding the B-spline coefficients. Specifically, we express the spectrum as a weighted sum of N basis functions $B_i(\lambda)$ defined over the wavelength range Λ with nonnegative coefficients c_i :

$$S(\lambda) = \sum_{j=1}^{N} c_j B_j(\lambda), \text{ where } \sum_{j=1}^{N} B_j(\lambda) = 1, \quad \forall \lambda \in \Lambda.$$
 (2)

The partition of unity property of B-Spline functions ensures that

$$\forall j \in [1, N], c_j \in [0, 1] \implies S(\lambda) = \sum_{i=1}^{N} c_j B_j(\lambda) \in [0, 1].$$
 (3)

So we only need to constrain the coefficients to produce a physically plausible reflectance spectrum.

Spectral Reconstruction and Null-Space Sampling

To reconstruct the reflectance spectrum, we first use differentiable rendering to iteratively update the basis coefficients c_i , converging to a plausible spectral solution consistent with the multiband image input. However, spectrum reconstruction from multiband data is inherently ill-posed. For instance, RGB input provides only three channels, while accurately representing real-world spectra often requires five or more coefficients (Figure 2). As a result, the solution obtained through differentiable rendering can deviate significantly from the ground truth.

We propose a spectral upsampling technique based on null-space sampling that generates multiple candidate spectra and quantifies uncertainty across wavelengths. Spectral integration is expressed in matrix form using a B-spline representation. Assume we sample *P* points in the wavelength range. Let $\mathbf{B} \in \mathbb{R}^{P \times N}$ denote the matrix that maps the coefficient vector **c** to the sampled spectrum **S**, and let $\mathbf{M} \in \mathbb{R}^{K \times P}$ be the matrix projecting the full spectrum to the multi-band oberservation **y**. Thus we have:

$$\mathbf{B}_{i,j} = B_j(\lambda_i), \quad i = 1, \dots, P, \ j = 1, \dots, N,$$

$$\mathbf{M}_{i,j} = M_i(\lambda_j)S(\lambda_j)\Delta\lambda, \quad i = 1, \dots, K, \ j = 1, \dots, P,$$

$$\mathbf{S} = \mathbf{Bc},$$

$$\mathbf{y} = \mathbf{MS} = \mathbf{MBc} = \mathbf{Ac},$$

$$(4)$$

where $\Delta \lambda$ is the wavelength step size and A = MB.

Given the multi-band value y^* , obtained from the spectral solution via differentiable rendering, we aim to solve $Ac = y^*$ subject to the constraint $c \in [0, 1]^N$. When the system is underdetermined, there are infinitely many solutions. The entire space of coefficient vectors that produce the same y^* can be expressed as

$$\mathbf{c} = \mathbf{c}_0 + \mathbf{N}\mathbf{z},\tag{5}$$

where \mathbf{c}_0 is a particular solution, $\mathbf{N} \in \mathbb{R}^{N \times D}$ spans the null space of \mathbf{A} , capturing the degrees of freedom that leave $\mathbf{A}\mathbf{c}$ unchanged, and $\mathbf{z} \in \mathbb{R}^D$ contains the free parameters with D = N - K. The solution \mathbf{c} must satisfy $\mathbf{c} \in [0,1]^N$.

We achieve this by sampling feasible vectors z such that the reconstructed spectrum coefficients $c = c_0 + Nz$ stay within the bounds $[0,1]^N$. The full algorithm proceeds in the following steps:

- (1) Acquire a specific solution c_0 : Either use the coefficients obtained from differentiable rendering as c_0 , or run a linear least-squares solve to determine c_0 .
- (2) **Compute valid ranges for each dimension of z:** For each dimension i of z, compute the lower and upper bounds that ensure all entries of $c = c_0 + Nz$ remain between 0 and 1. This is done element-wise by solving for intervals where

$$c_i = c_{0,i} + (Nz)_i \in [0, 1].$$

The bounds are clipped to a reasonable range (e.g., $[-10^3, 10^3]$) to avoid numerical instability.

- (3) **Sample uniformly within bounds:** After computing the valid intervals $[z_{\min}, z_{\max}]$ for each dimension, generate a set of **z** samples uniformly within those bounds.
- (4) Filter Infeasible Samples:

For each sampled z, compute the corresponding $\mathbf{c} = \mathbf{c}_0 + Nz$. Retain only those samples where all components of \mathbf{c} lie within [0,1].

Quantify Uncertainty. After we acquire the z samples, we can compute the B-Spline coefficients and each of the sampled spectra S^i :

$$\mathbf{c}^{i} = \mathbf{c}_{0} + \mathbf{N}\mathbf{z}^{i}, \quad \mathbf{S}^{i} = \mathbf{B}\mathbf{c}^{i}. \tag{6}$$

Then we can compute the mean spectrum $\bar{S}(\lambda)$ and standard deviation $\sigma(\lambda)$ of the sampled spectra as functions of λ :

$$\bar{\mathbf{S}}(\lambda) = \frac{1}{T} \sum_{i=1}^{T} \mathbf{S}^{i}(\lambda), \quad \sigma^{2}(\lambda) = \frac{1}{T} \sum_{i=1}^{T} \left(\mathbf{S}^{i} - \bar{\mathbf{S}}(\lambda) \right)^{2}$$
 (7)

The standard deviation in Equation 7 quantifies the reconstruction uncertainty. In our experiments, both metrics converge after approximately 1,000 samples.

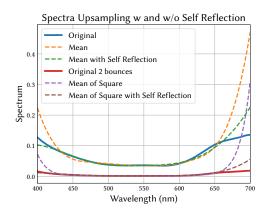


Fig. 3. The ground truth spectra are the blue and red curves. Considering self-reflection (the green and brown curves) improves the accuracy of the sampled spectra.

3.3 Enhancements to the Baseline Algorithm

We will describe three enhancements to the baseline algorithm to further improve the reconstruction accuracy.

3.3.1 Incorporating the Light Source Spectrum. Real captures are often conducted in a controlled environment where the light spectrum $L(\lambda)$ is known. To incorporate this, we update the computation of \mathbf{M} as follows:

$$\mathbf{M}_{i,j} = M_i(\lambda_j)S(\lambda_j)L(\lambda_j)\Delta\lambda, \quad i = 1, \dots, K, \ j = 1, \dots, P,$$
 (8) which in turn changes **A**.

3.3.2 Incorporating Additional Constraints. Reconstruction accuracy can be further improved by incorporating additional constraints. For example, we can guide supplementary measurements at specific wavelengths where the baseline upsampling results exhibit high uncertainty. To achieve this, we define a loss function at the additional wavelengths as

$$\mathcal{L} = \sum_{i} \|S(\lambda_i) - S^*(\lambda_i)\|^2, \tag{9}$$

where $S^*(\lambda_i)$ denotes additional info at wavelength λ_i . We then use this loss to rank the sampled spectra and select the top K candidates, where K is a user-defined parameter. In the last columns of Figure 1 and Figure 4, we show that adding measurements at 400 nm and 700 nm — where the uncertainty is highest — significantly improves the spectral reconstruction.

Another effective way to enhance reconstruction accuracy is to account for interreflections, either within a single object or between multiple objects. We define a loss function based on the multispectral responses \mathbf{y}_i^* obtained under interreflection, and select the top candidates by minimizing this loss:

$$\mathcal{L} = \sum_{i} \|\mathbf{y} - \mathbf{y}_{i}^{*}\|^{2}.$$
 (10)

The responses \mathbf{y}_i^* are computed by multiplying the spectral outputs of differentiable rendering and then converting the result to RGB. If we use only the RGB values derived from the spectral output of

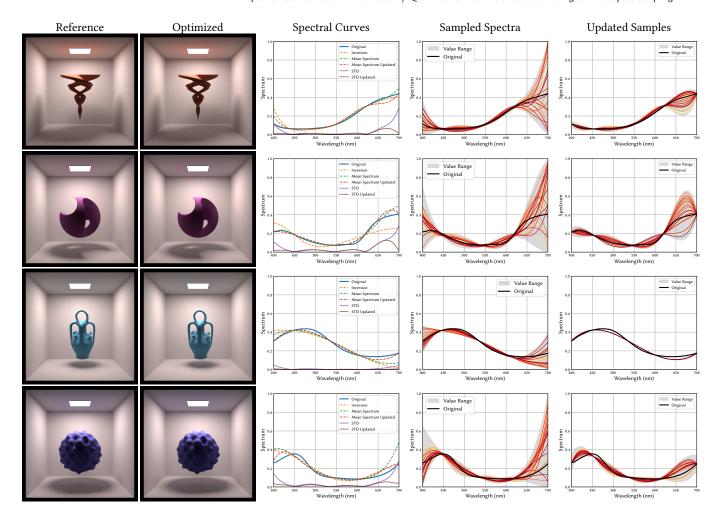


Fig. 4. We evaluate single reflectance spectrum reconstruction on spectral examples with different characteristics. We first obtain a candidate spectrum through inverse rendering. We then generate multiple spectra that produce the same RGB color using our method, and compute both the mean spectrum and the spectral standard deviation (STD) across wavelengths. This sampling process enables uncertainty quantification in the spectral reconstruction and guides additional captures to improve accuracy.

differentiable rendering, the upsampled spectra may fail to preserve the interreflection effects. Figure 3 shows that incorporating interreflection improves the accuracy of the sampled spectra, and the mean spectrum becomes significantly closer to the ground truth.

- 3.3.3 Adaptive Basis Refinement. We observe that higher-order variations in real-world spectra typically have much smaller amplitudes than lower-order components. Naively increasing the number of B-spline coefficients can make the reconstruction problem more ill-posed and often results in higher uncertainty. To address this, we propose an adaptive basis refinement strategy: we first estimate the coefficients using a coarse set of basis functions. Then given
 - A set of coarse coefficients $\mathbf{c}_c \in \mathbb{R}^{N_c}$

 - A coarse basis matrix $\mathbf{B}_c \in \mathbb{R}^{P \times N_c}$ A fine basis matrix $\mathbf{B}_f \in \mathbb{R}^{P \times N_f}$ with $N_f > N_c$
 - A measurement matrix $\mathbf{A}_f = \mathbf{M}\mathbf{B}_f \in \mathbb{R}^{K \times N_f}$

• Function-space Gram matrix $G = \mathbf{B}_f^{\mathsf{T}} \mathbf{B}_f$

We want to lift the coarse solution to a fine one and explore additional degrees of freedom. This is achieved as follows:

(1) Coefficient Lifting: Lift the coarse solution into the fine basis:

$$\mathbf{c}_{f0} = \mathbf{T}\mathbf{c}_c,\tag{11}$$

where

$$\mathbf{T} = \arg\min_{\mathbf{T}} \|\mathbf{B}_f \mathbf{T} - \mathbf{B}_c\|_F^2 \Rightarrow \mathbf{T} = \mathbf{B}_f^{\dagger} \mathbf{B}_c, \tag{12}$$

and \mathbf{B}_f^{\dagger} is the Moore–Penrose pseudoinverse of \mathbf{B}_f .

(2) **Residual Subspace:** Choose any G-orthonormal basis Q_c of the lifted coarse span, so $Q_c^{\top}GQ_c = I$ and span $(Q_c) = I$ $\operatorname{span}(T)$. Complete it to a G-orthonormal basis $[Q_c \ Q_{\perp}]$

$$[\mathbf{Q}_c \ \mathbf{Q}_\perp]^\top \mathbf{G} [\mathbf{Q}_c \ \mathbf{Q}_\perp] = \mathbf{I}. \tag{13}$$

(3) **Refined Nullspace**: Restrict the measurement operator to the residual subspace compute the null:

$$S = A_f Q_{\perp} \in \mathbb{R}^{K \times r}, \qquad N_{\perp} = \text{null}(S).$$
 (14)

The valid refinement directions in coefficient space are

$$N_{\text{residual}} = Q_{\perp} N_{\perp} \in \mathbb{R}^{N_f \times d}, \quad d = \dim \text{null}(S).$$
 (15)

(4) **Final Solution:** Construct the full solution by adding a residual perturbation to the lifted coarse solution:

$$\mathbf{c}_f = \mathbf{c}_{f0} + \mathbf{N}_{\text{residual}}\mathbf{z},\tag{16}$$

where **z** parametrizes the residual degrees of freedom.

Figure 5 demonstrates that lifting the coarse solution to a denser basis, followed by sampling of the residual degrees of freedom, improves upon the results obtained by directly sampling in the fine-basis representation.

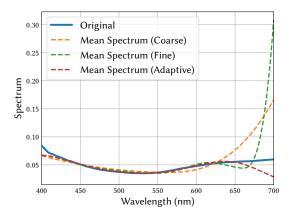


Fig. 5. Adaptive basis refinement can improve on the results obtained by directly using a fine basis representation.

3.4 Comparison with Previous Work

Belcour et al. [2023] proposed a method for generating multiple spectral solutions from RGB input. Compared to their approach, our method offers two main advantages: (1) we support general multispectral inputs beyond RGB, whereas they rely on chromaticity and luminance conversion and is therefore limited to RGB; (2) we use a general null-space sampling approach, while they use a greedy algorithm to iteratively sample each B-spline coefficient, which often fails to produce valid solutions unless the number of coefficients is small. In our experiments, their method often struggles to produce solutions when fitting real spectra with 7 or more coefficients. For example, their method fails to generate any solutions for upsampling the RGB values corresponding to the spectrum in Figure 8 when using 7 or more coefficients.

4 RESULTS

We first validate B-spline fitting on real-world spectral reflectance data, then perform spectral reconstruction using synthetic scenes in Mitsuba 3. Finally, we demonstrate the effectiveness of our algorithm through physical experiments.

4.1 B-Spline Fitting

To validate that spectral reflectance data can be accurately represented using B-splines, we fit known spectra from Kremer Pigments [Deborah 2022] and EcoSIS plant and forest spectra [Wagner et al. 2018]. We used cubic B-splines with a uniform knot vector and endpoint multiplicity of 4 to ensure clamped boundaries, and varied the number of coefficients to assess the accuracy of the fitting (Figure 2).

As expected, increasing the number of control points improved the fitting accuracy. Smoother spectra, such as those from the paint dataset, could be approximated with fewer coefficients, while more complex natural spectra from EcoSIS and certain fluorescent pigments required greater flexibility for accurate interpolation. The quadratic representation used in [Jakob and Hanika 2019]—the default spectral upsampling method in Mitsuba 3—is often inadequate for representing real-world spectra.

While not comprehensive, this validation demonstrates the ability of our B-Spline approach to fit broad classes of material reflectance spectra.

4.2 Spectral Reconstruction Using Synthetic Scenes

Single Spectrum Reconstruction. In Figures 1 and 4, we evaluate single-spectrum reconstruction across a variety of spectra with different characteristics. We first generate the reference image using the ground-truth spectrum. We then obtain a candidate spectrum through differentiable rendering, where each spectrum is represented using either 5 or 7 coefficients, with the optimization initialized from zero coefficients (corresponding to a black appearance). The blue curve shows the ground-truth spectrum, while the orange curve represents the spectrum obtained through differentiable rendering.

We then apply our null-space sampling method to generate spectra (30k–80k valid samples out of 100k trials) that produce the same RGB color. We compute both the mean spectrum (green curve) and the spectral standard deviation (purple curve) across wavelengths, enabling uncertainty quantification in spectral reconstruction.

Based on the computed uncertainty, we introduce two additional measurements at 400nm and 700nm. Uncertainty is typically highest at the ends of the wavelength range, where the sensitivity functions are low and spectral variations have minimal impact on the final RGB values. Using single-wavelength reconstruction at these points, we filter the sample set by selecting the top 10k candidates that best match the additional values. In the last two columns of the figure, we plot the first 30 samples before and after introducing these measurements, with the gray region indicating the range covered by all selected samples. We observe a significant reduction in uncertainty, and the mean spectrum becomes much more accurate. The updated mean spectrum and standard deviation are shown as red and brown curves in the first column.

Joint Spectrum and Shape Reconstruction. Figure 6 demonstrates the joint reconstruction of shape and spectrum. Starting from a black sphere, differentiable rendering successfully recovers both the shape and appearance, matching the observed RGB color. We apply the large-step method [Nicolet et al. 2021] along with a coarse-to-fine geometry refinement strategy. Our spectral

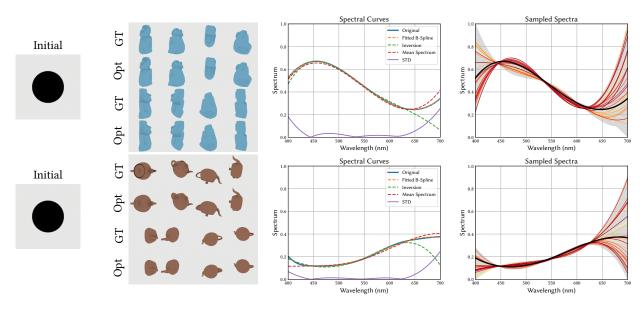


Fig. 6. This figure shows joint shape and spectral reconstruction. Differentiable rendering reconstructs both the shape and appearance with the same RGB color. Our spectral sampling algorithm quantifies the uncertainty in the spectral reconstruction.

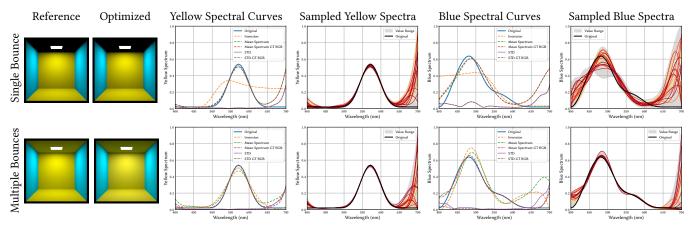


Fig. 7. Interreflection leads more accurate spectral reconstruction and less uncertainty of the blue and yellow walls. This example further illustrates that more accurate RGB information improves sampling accuracy.

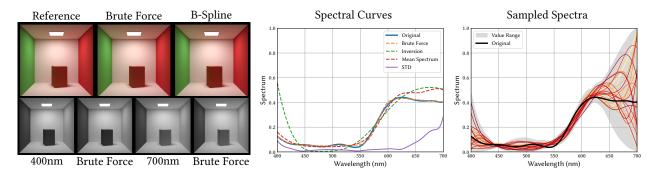


Fig. 8. We perform brute-force spectral reconstruction on the orange box in the Cornell box scene. The brute-force method achieves good accuracy but takes 7.1 times longer to run.

sampling algorithm further quantifies the uncertainty in the reconstructed spectrum. In the middle column, the orange curve shows a B-spline fit directly to the ground-truth spectrum, illustrating the ideal reconstruction achievable with the current representation.

Spectral Reconstruction with Interreflection. Figure 7 illustrates the influence of interreflection on spectral reconstruction. We recover the spectra of the blue and yellow walls in the Cornell Box scene. The first row considers only direct illumination; in this case, the spectral solution obtained from differentiable rendering (orange curve) deviates significantly from the ground truth (blue curve). In contrast, the mean spectrum of the samples generated by our algorithm is much closer to the ground truth than the inverse rendering result. Note that in this test, the initial coefficient \mathbf{c}_0 used in Equation 5 and for sampling the spectra was computed by a linear least-squares solve, rather than taken directly from the inverse rendering, which results in more accurate spectral estimations.

The second row shows reconstruction with global illumination enabled, where the spectral solution from inverse rendering is substantially more accurate compared to the direct-illumination case. In the spectral curve plots, we also show the mean spectrum of the samples generated using the ground-truth RGB color of each wall, highlighting that sampling accuracy is influenced by the errors in the inverse rendering result.

Comparison with Brute-Force Reconstruction. The brute-force method for reconstructing reflectance spectra while accounting for 3D geometry and light transport involves capturing hyperspectral images and applying differentiable rendering to reconstruct each wavelength independently. We simulate this process by first performing forward rendering using the ground-truth spectrum and the monowavelength variant in Mitsuba 3, then optimizing the reflectance value at each wavelength for the box

Figure 8 compares the results of the brute-force method with our B-spline-based method, which reconstructs the spectrum from RGB input. On the left, we show the single-wavelength reconstruction results at 400nm and 700nm using the brute-force method, along with the corresponding RGB images produced by both the brute-force and B-spline methods.

In the spectral curve plot, we observe that the mean spectrum (red curve) of the samples generated from the B-spline-based solution achieves higher accuracy than the direct inverse rendering result (green curve). Although the brute-force spectral solution (orange curve) is more accurate, it requires 8.7 minutes to compute on an RTX 4090 GPU, and is 7.1 times longer than the B-spline-based method (1.2 minutes), which uses Mitsuba's spectral mode with the sampling method from [Radziszewski et al. 2009]. Alternatively, the hero wavelength sampling method [Wilkie et al. 2014] can be used. The sampling step in our method takes only 0.001 seconds.

Comparison with Image Based Methods. Many methods for recovering spectral information efficiently from RGB images in computer vision rely solely on 2D images and do not take global illumination and 3D geometry fully into account. We evaluated the NTIRE 2022 Spectral Reconstruction Challenge winning method [Cai et al. 2022] on the filigree example in the teaser. Their method outputs

a hyperspectral image, from which we extract the spectra corresponding to the filigree. The mean error between our reconstructed spectra and the ground truth is 0.0810 (over 90k spectra), compared to 0.4033 (over 11k spectra) for their method. Figure 9 shows that the mean spectrum and spectral samples from their method differ significantly from the ground-truth spectrum. It is worth noting that image-based methods do not disentangle lighting from material properties — making this white light setting effectively their easiest test case.

4.3 Physical Experiments

For physical validation, we constructed a Cornell Box with 16-inch interior dimensions and applied Valspar® paints to all interior surfaces. The ground-truth spectra of the painted surfaces, objects, and the light source (an iPhone flashlight) were measured using an ASD FieldSpec Pro spectroradiometer (Analytical Spectral Devices, Inc., USA).

Sensor Sensitivity and Illumination Calibration. The default sensitivity function used in Mitsuba for converting spectral to RGB is the CIE-1913 color matching function. To account for the real

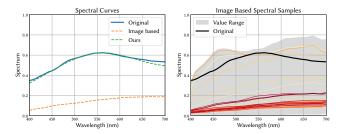


Fig. 9. This figure compares the mean spectra of [Cai et al. 2022] and our method; it also shows the spectral samples from their method.

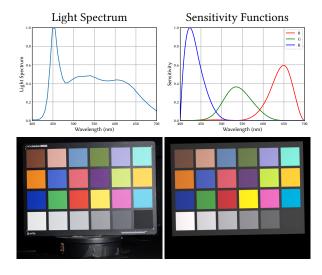


Fig. 10. This figure shows the measured light spectrum and the fitted sensitivity functions (both normalized), as well as the captured ColorChecker chart and the rendered one.

sensitivty function of the Sony a7 II camera we used, we fitted customized sensitivity functions. Specifically, we took a photo of the ColorChecker (Macbeth) Chart, illuminated using the same light source and captured using the same camera. We acquired the spectra of the color grids from [Pascale 2023] where they report the average data derived from measurements of 30 charts.

We assume that all the grids have diffuse BRDF. For each grid, we integrate the outgoing radiance $E(\lambda)$ over a patch A:

$$E(\lambda) = \int_{A} L_{i}(\lambda) \frac{\rho(\lambda)}{\pi} \frac{\cos \theta_{i}}{r^{2}} dA, \tag{17}$$

where $L_i(\lambda)$ is emitted radiance from the light, θ_i is the incident angle, and *r* is the distance from the light source to the integrated point. We extract the average RGB color of the patch from the captured photo.

We estimate the spectral sensitivity functions of a 3-channel sensor using B-Spline basis functions. Let $\mathbf{B} \in \mathbb{R}^{P \times N}$ be a matrix of basis functions sampled over P wavelengths. There are 24 color grids with known spectra and let $S_i \in \mathbb{R}^P$ be the *i*-th reflectance spectrum associated with a ground truth RGB value \mathbf{y}_{i}^{*} . In practice we set P = 400 and we solve for 10 B-Spline coefficients for each channel. The predicted RGB values are:

$$\hat{r}_i = \mathbf{S}_i^{\mathsf{T}} \mathbf{B} \mathbf{c}_R, \quad \hat{g}_i = \mathbf{S}_i^{\mathsf{T}} \mathbf{B} \mathbf{c}_G, \quad \hat{b}_i = \mathbf{S}_i^{\mathsf{T}} \mathbf{B} \mathbf{c}_B$$
 (18)

We solve for the weight vectors \mathbf{c}_R , \mathbf{c}_G , \mathbf{c}_B by minimizing the total squared error:

$$\min_{\mathbf{c}_R, \mathbf{c}_G, \mathbf{c}_B} \sum_{i=1}^{N} \left\| \begin{bmatrix} \hat{r}_i \\ \hat{g}_i \\ \hat{b}_i \end{bmatrix} - \begin{bmatrix} r_i \\ g_i \\ b_i \end{bmatrix} \right\|^2, \tag{19}$$

subject to boundary conditions:

$$\mathbf{b}_{400}^{\mathsf{T}} \mathbf{c}_Q = 0, \quad \mathbf{b}_{700}^{\mathsf{T}} \mathbf{c}_Q = 0, \quad \text{for } Q \in \{R, G, B\}$$

These constraints ensure that the sensor sensitivities are zero at 400nm and 700nm, enforcing realistic falloff near spectral boundaries.

The spectral irradiance of the light source was measured. To convert it into emitted radiance for use in the rendering scene, we performed differentiable rendering to optimize a scale factor for the light. Figure 10 shows the normalized light spectrum, the fitted sensitivity functions, and both the captured and rendered ColorChecker chart.

Results. Figure 11 presents the results of the physical experiments. In each of the three setups, a cube painted in a different color is placed inside the Cornell box. We jointly reconstruct the spectra of the walls and the cube. The first three rows show the captured photos, renderings after spectral optimization, forward renderings using measured spectra, and spectral loss during optimization. Renderings with ground-truth spectra show slight mismatches with the photos, providing a visualization of the cumulative effect of measurement errors in the spectral reflectances, source emission, and sensitivity function. The bottom three rows display the spectral curves and sampled spectra for the colored walls and cubes. Our method enables uncertainty quantification, and the mean spectrum for the green wall reconstruction is significantly more accurate than that from inverse rendering within the interior range of wavelengths. Note that measurement errors reduce the accuracy of the inversion compared to synthetic results. Since the upsampling step assumes the differentiable rendering spectrum to be RGB-faithful, these errors also propagate into the upsampling results. This highlights the particular importance of uncertainty quantification in real-world scenarios.

DISCUSSION AND CONCLUSION

This paper presents a novel approach that leverages differentiable rendering to reconstruct reflectance spectra from RGB and general multispectral images. By naturally accounting for light transport and geometry, our method recovers physically plausible spectra. To address the inherent ambiguity in mapping multiband observations to full spectra, we introduce a null-space sampling technique that generates plausible spectral candidates and quantifies uncertainty across wavelengths. We further propose enhancement techniques to improve reconstruction accuracy. The method is validated on both synthetic scenes and real-world experiments. Compared to traditional hyperspectral imaging, our approach eliminates complex acquisition procedures and significantly accelerates reconstruction compared to brute-force methods. It can also be applied to spectral upsampling in spectral rendering.

Limitations and Future Work. Our work opens new avenues for follow-up research at the intersection of differentiable rendering and spectral reconstruction. Future directions include reconstructing illumination spectra, handling non-smooth spectra, and optimizing reconstruction setups to improve spectral recovery. Knowing the appropriate number of coefficients in advance improves reconstruction and reduces uncertainty, whereas we currently assume this value. Finally, our method operates as a post-processing step after inverse rendering; an interesting direction is to integrate it directly into the differentiable rendering pipeline to enable a more principled exploration of multiple solutions during optimization.

ACKNOWLEDGMENTS

We would like to thank Nadia Zikiou for helping with measuring the ground truth spectra for the physical experiment. This work was supported by the National Science Foundation under grant 2303328.

REFERENCES

Jonas Aeschbacher, Jiqing Wu, and Radu Timofte. 2017. In defense of shallow learned spectral reconstruction from RGB images. In Proceedings of the IEEE International Conference on Computer Vision Workshops. 471-479.

Naveed Akhtar and Ajmal Mian. 2018. Hyperspectral recovery from RGB images using Gaussian processes. IEEE transactions on pattern analysis and machine intelligence 42, 1 (2018), 100-113.

Aitor Alvarez-Gila, Joost Van De Weijer, and Estibaliz Garrote. 2017. Adversarial networks for spatial context-aware spectral image reconstruction from RGB. In Proceedings of the IEEE international conference on computer vision workshops. 480-

Boaz Arad and Ohad Ben-Shahar. 2016. Sparse recovery of hyperspectral signal from natural RGB images. In Computer Vision-ECCV 2016: 14th European Conference, Amsterdam, the Netherlands, October 11-14, 2016, Proceedings, Part VII 14. Springer,

Sai Praveen Bangaru, Tzu-Mao Li, and Frédo Durand. 2020. Unbiased warped-area sampling for differentiable rendering. ACM Transactions on Graphics (TOG) 39, 6

Laurent Belcour, Pascal Barla, and Gaël Guennebaud. 2023. One-to-Many Spectral Upsampling of Reflectances and Transmittances. In Computer Graphics Forum, Vol. 42. Wiley Online Library, e14886.

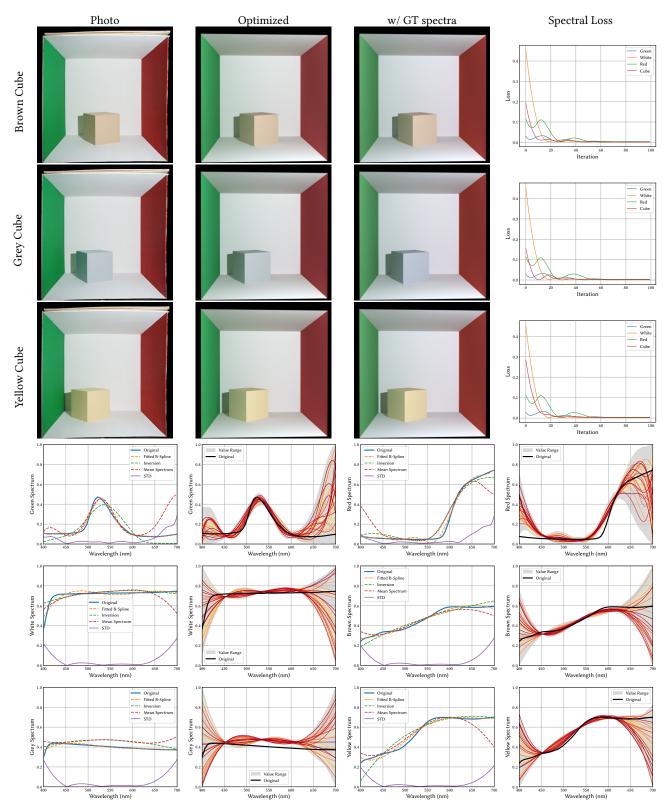


Fig. 11. We constructed a physical Cornell box and performed spectral reconstruction on the green, red, and white walls, as well as the colored cubes. Our method enables uncertainty quantification of the reconstructed spectra. For the green wall, the mean spectrum yields significantly more accurate results within the interior range of wavelengths.

SA Conference Papers '25, December 15–18, 2025, Hong Kong, Hong Kong.

- Yuanhao Cai, Jing Lin, Zudi Lin, Haoqian Wang, Yulun Zhang, Hanspeter Pfister, Radu Timofte, and Luc Van Gool. 2022. Mst++: Multi-stage spectral-wise transformer for efficient spectral reconstruction. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 745-755
- Yigit Baran Can and Radu Timofte. 2018. An efficient CNN for spectral reconstruction from RGB images. arXiv preprint arXiv:1804.04647 (2018).
- Chein-I Chang. 2013. Hyperspectral data processing: algorithm design and analysis. John Wiley & Sons
- H. Deborah. 2022. Hyperspectral Pigment Dataset. In 2022 12th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS). IEEE, Rome, Italy, 1-5. https://doi.org/10.1109/WHISPERS56178.2022.9955067
- Yingtao Duan, Nan Wang, Yifan Zhang, and Chao Song. 2024. Tensor-based sparse representation for hyperspectral image reconstruction using RGB inputs. Mathematics 12. 5 (2024), 708.
- Luca Fascione, Johannes Hanika, Mark Leone, Marc Droske, Jorge Schwarzhaupt, Tomáš Davidovič, Andrea Weidlich, and Johannes Meng. 2018. Manuka: A batch-shading architecture for spectral path tracing in movie production. ACM Transactions on Graphics (TOG) 37, 3 (2018), 1-18.
- Yunhao Geng, Shaohui Mei, Jin Tian, Yifan Zhang, and Qian Du. 2019. Spatial constrained hyperspectral reconstruction from RGB inputs using dictionary representation. In IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium. IEEE, 3169-3172.
- Daniel C Harris, 2010. Quantitative chemical analysis. Macmillan.
- Wenzel Jakob and Johannes Hanika. 2019. A low-dimensional function space for efficient spectral upsampling. In Computer Graphics Forum, Vol. 38. Wiley Online Library, 147-155.
- Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. 2022. Mitsuba 3 renderer. https://mitsuba-renderer.org.
- Yan Jia, Yinqiang Zheng, Lin Gu, Art Subpa-Asa, Antony Lam, Yoichi Sato, and Imari Sato. 2017. From RGB to spectrum for natural scenes via manifold-based mapping. In Proceedings of the IEEE international conference on computer vision, 4705-4713.
- Min H Kim, Holly Rushmeier, John Ffrench, Irma Passeri, and David Tidmarsh. 2014. Hyper3D: 3D graphics software for examining cultural artifacts. Journal on Computing and Cultural Heritage (JOCCH) 7, 3 (2014), 1-19.
- Jiaojiao Li, Chaoxiong Wu, Rui Song, Yunsong Li, and Fei Liu. 2020. Adaptive weighted attention network with camera spectral sensitivity prior for spectral reconstruction from RGB images. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. 462-463.
- Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. 2018. Differentiable monte carlo ray tracing through edge sampling. ACM Transactions on Graphics (TOG) 37, 6 (2018), 1-11.
- Guillaume Loubet, Nicolas Holzschuch, and Wenzel Jakob. 2019. Reparameterizing discontinuous integrands for differentiable rendering. ACM Transactions on Graphics (TOG) 38, 6 (2019), 1-14.
- David L MacAdam. 1935. Maximum visual efficiency of colored materials. Journal of the Optical Society of America 25, 11 (1935), 361-367.
- Johannes Meng, Florian Simon, Johannes Hanika, and Carsten Dachsbacher. 2015. Physically meaningful rendering using tristimulus colours. In Computer Graphics Forum, Vol. 34. Wiley Online Library, 31-40.
- Baptiste Nicolet, Alec Jacobson, and Wenzel Jakob. 2021. Large steps in inverse rendering of geometry. ACM Transactions on Graphics (TOG) 40, 6 (2021), 1-13.
- Merlin Nimier-David, Sébastien Speierer, Benoît Ruiz, and Wenzel Jakob. 2020. Radiative backpropagation: An adjoint method for lightning-fast differentiable rendering. ACM Transactions on Graphics (TOG) 39, 4 (2020), 146-1.
- Herman L Offerhaus, Sarah E Bohndiek, and Andrew R Harvey. 2019. Hyperspectral imaging in biomedical applications. J. Opt 21, 1 (2019), 010202.
- Hisanari Otsu, Masafumi Yamamoto, and Toshiya Hachisuka. 2018. Reproducing spectral reflectances from tristimulus colours. In Computer Graphics Forum, Vol. 37. Wiley Online Library, 370-381.
- Danny Pascale. 2023. ColorChecker® Classic and ColorChecker® Passport 2 Data. https://babelcolor.com/colorchecker-2.htm#CCP2_data. https://babelcolor.com/ colorchecker-2.htm#CCP2_data Accessed: https://babelcolor.com/colorchecker-2.htm#CCP2_data on 2025-04-23.
- 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. 191-198.
- Christoph Peters, Sebastian Merzbach, Johannes Hanika, and Carsten Dachsbacher. 2019. Using moments to represent bounded signals for spectral rendering. ACM Transactions on Graphics (TOG) 38, 4 (2019), 1-14.
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2023. Physically based rendering: From theory to implementation. MIT Press.
- Michal Radziszewski, Krzysztof Boryczko, and Witold Alda. 2009. An Improved Technique for Full Spectral Rendering. 7. WSCG 17, 1-3 (2009), 9-16.
- Brian Smits. 1999. An RGB-to-spectrum conversion for reflectances. Journal of Graphics tools 4, 4 (1999), 11-22.

- Shlomi Steinberg. 2023. Towards Practical Physical Optics Rendering. University of California, Santa Barbara.
- Tarek Stiebel, Simon Koppers, Philipp Seltsam, and Dorit Merhof. 2018. Reconstructing spectral images from RGB-images using a convolutional neural network. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops.
- Lucia Tódová, Alexander Wilkie, and Luca Fascione. 2022. Wide Gamut Moment-based Constrained Spectral Uplifting. In Computer Graphics Forum, Vol. 41. Wiley Online Library, 258-272.
- Mark van de Ruit and Elmar Eisemann. 2024. Controlled Spectral Uplifting for Indirect-Light-Metamerism. In SIGGRAPH Asia 2024 Conference Papers. 1-9.
- Delio Vicini, Sébastien Speierer, and Wenzel Jakob. 2021. Path replay backpropagation: Differentiating light paths using constant memory and linear time. ACM Transactions on Graphics (TOG) 40, 4 (2021), 1-14.
- E. P. Wagner, J. Merz, and P. A. Townsend. 2018. Ecological Spectral Information System: An Open Spectral Library. In AGU Fall Meeting Abstracts, Vol. 2018. Article B41L-2878, B41L-2878 pages.
- Andrea Weidlich, Alex Forsythe, Scott Dyer, Thomas Mansencal, Johannes Hanika, Alexander Wilkie, Luke Emrose, and Anders Langlands. 2021. Spectral imaging in production: course notes Siggraph 2021. In ACM SIGGRAPH 2021 Courses. 1-90.
- Andrea Weidlich, Chloe LeGendre, Carlos Aliaga, Christophe Hery, Jean-Marie Aubry, Jiří Vorba, Daniele Siragusano, and Richard Kirk. 2022. Practical aspects of spectral data in digital content production. In ACM SIGGRAPH 2022 Courses. i-96.
- Alexander Wilkie, Sehera Nawaz, Marc Droske, Andrea Weidlich, and Johannes Hanika. 2014. Hero wavelength spectral sampling. In Computer Graphics Forum, Vol. 33. Wiley Online Library, 123-131.
- Mengqi Xia, Bruce Walter, Christophe Hery, Olivier Maury, Eric Michielssen, and Steve Marschner. 2023b. A practical wave optics reflection model for hair and fur. ACM Transactions on Graphics (TOG) 42, 4 (2023), 1-15,
- Mengqi Xia, Bruce Walter, Eric Michielssen, David Bindel, and Steve Marschner. 2020. A wave optics based fiber scattering model. ACM Transactions on Graphics (TOG) 39. 6 (2020), 1-16.
- Mengqi (Mandy) Xia, Bruce Walter, and Steve Marschner. 2023a. Iridescent Water Droplets Beyond Mie Scattering. Computer Graphics Forum (2023). https://doi.org/ 10.1111/cgf.14893
- Zhiwei Xiong, Zhan Shi, Huiqun Li, Lizhi Wang, Dong Liu, and Feng Wu. 2017. Hscnn: Cnn-based hyperspectral image recovery from spectrally undersampled projections. In Proceedings of the IEEE international conference on computer vision workshops. 518-525
- Peiyu Xu, Sai Bangaru, Tzu-Mao Li, and Shuang Zhao. 2023. Warped-area reparameterization of differential path integrals. ACM Transactions on Graphics (TOG) 42, 6 (2023), 1-18.
- Bai Xue, Shih-Yu Chen, Chunyuan Yu, Yulei Wang, Lin Wang, Meiping Song, Sen Li, and Chein-I Chang. 2017. Kernel automatic target generation process. In 2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS). IEEE, 636-639.
- Kai Yan, Christoph Lassner, Brian Budge, Zhao Dong, and Shuang Zhao. 2022. Efficient estimation of boundary integrals for path-space differentiable rendering. ACM Transactions on Graphics (TOG) 41, 4 (2022), 1-13.
- Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering specular microgeometry with wave optics. ACM Transactions on Graphics (TOG) 37, 4 (2018), 1-10.
- Yunchen Yu, Mengqi Xia, Bruce Walter, Eric Michielssen, and Steve Marschner. 2023. A full-wave reference simulator for computing surface reflectance. ACM Transactions on Graphics (TOG) 42, 4 (2023), 1-17.
- Cheng Zhang, Lifan Wu, Changxi Zheng, Ioannis Gkioulekas, Ravi Ramamoorthi, and Shuang Zhao. 2019. A differential theory of radiative transfer. ACM Transactions on Graphics (TOG) 38, 6 (2019), 1-16.
- Cheng Zhang, Zihan Yu, and Shuang Zhao. 2021. Path-space differentiable rendering of participating media. ACM Transactions on Graphics (TOG) 40, 4 (2021), 1-15.
- Jingang Zhang, Runmu Su, Qiang Fu, Wenqi Ren, Felix Heide, and Yunfeng Nie. 2022. A survey on computational spectral reconstruction methods from RGB to hyperspectral imaging. Scientific reports 12, 1 (2022), 11905.
- Ziyi Zhang, Nicolas Roussel, and Wenzel Jakob. 2023. Projective sampling for differentiable rendering of geometry. ACM Transactions on Graphics (TOG) 42, 6 (2023),
- Liaoying Zhao, Junpeng Zheng, Xiaorun Li, and Lijiao Wang. 2014. Kernel simplex growing algorithm for hyperspectral endmember extraction. Journal of Applied Remote Sensing 8, 1 (2014), 083594-083594.