DOI: 10.1111/cgf.15252
Pacific Graphics 2024
R. Chen, T. Ritschel, and E. Whiting
(Guest Editors)

COMPUTER GRAPHICS *forum*
*Volume 43 (2024), Number 7*

# Color-Accurate Camera Capture with Multispectral Illumination and Multiple Exposures

H. Gao[1] R. K. Mantiuk[1] G. D. Finlayson[2]

[1] University of Cambridge, UK
[2] University of East Anglia, UK

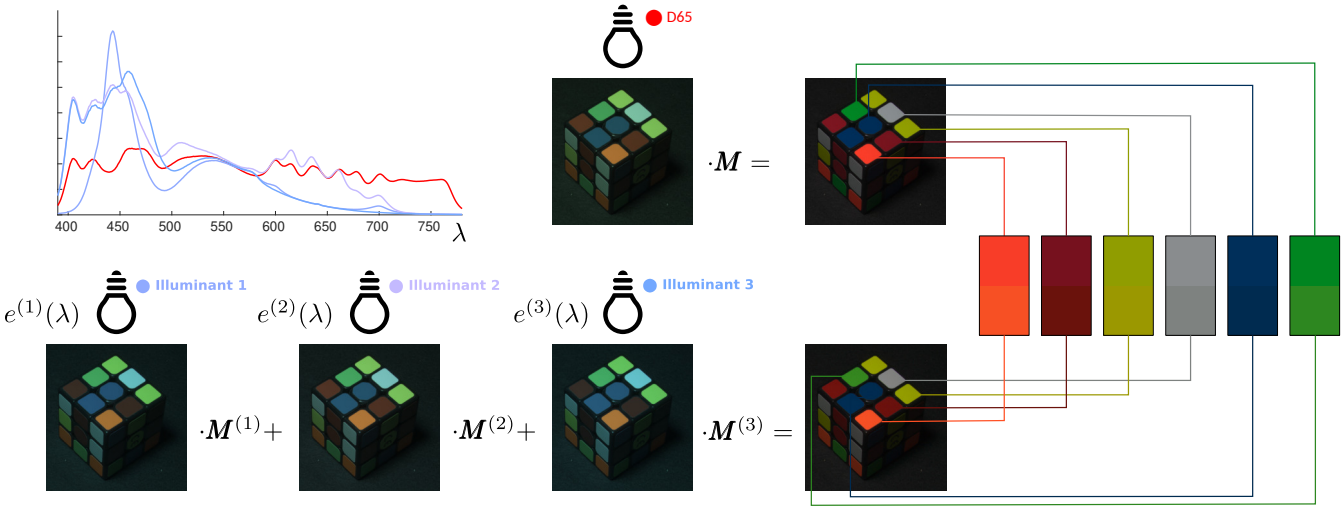


**Figure 1:** *To improve the color accuracy of a camera, we capture an object in multiple exposures (between 1 and 3), each time illuminating the object with a different light spectrum. The spectra are optimized in such a way that a linear combination of captured RGB values can be transformed into the CIE XYZ tristimulus values. In this example, the color differences are subtle but noticeable. When the Rubik's cube is captured under a D65 illuminant and transformed to the standard color space (BT.709), orange is mapped to a bright reddish color. Our optimized illuminants can correctly reproduce the original colors of the Rubik's cube. All colors in the paper are reproduced for the sRGB color space.*

**Abstract**

*Cameras cannot capture the same colors as those seen by the human eye because the eye and the cameras' sensors differ in their spectral sensitivity. To obtain a plausible approximation of perceived colors, the camera's Image Signal Processor (ISP) employs a color correction step. However, even advanced color correction methods cannot solve this underdetermined problem, and visible color inaccuracies are always present. Here, we explore an approach in which we can capture accurate colors with a regular camera by optimizing the spectral composition of the illuminant and capturing one or more exposures. We jointly optimize for the signal-to-noise ratio and for the color accuracy irrespective of the spectral composition of the scene. One or more images captured under controlled multispectral illuminants are then converted into a color-accurate image as seen under the standard illuminant of D65. Our optimization allows us to reduce the color error by 20–60% (in terms of CIEDE 2000), depending on the number of exposures and camera type. The method can be used in applications in which illumination can be controlled, and high colour accuracy is required, such as product photography or with a multispectral camera flash. The code is available at* `https://github.com/gfxdisp/multispectral_color_correction`.

**CCS Concepts**

• ***Computing methodologies*** → ***Computational photography;*** *Image processing; Perception;*

## 1. Introduction

Even professional photographic cameras are inherently color-inaccurate. This is because the camera's color filter design must balance two opposite requirements — color accuracy, which benefits from spectral transmission akin to the eye's cone sensitivities, and camera sensitivity, which benefits from wide spectral transmission resulting in color that can be transformed into a standard color space without amplifying noise. Because of that, the design of color filters in cameras is often a compromise between these two opposing goals. This poses a serious problem in applications that require accurate colors, such as material (BRDF/BTF) acquisition or product photography.

The requirements for accurate color capture are well understood and known as the Luther condition [Ive15, VL27] — camera spectral sensitivity functions (SSFs) must be linear combinations of color matching functions (CMFs) of the human eye, such as CIE XYZ 1931 or CIE XYZ/LMS 2006. If this condition is met, there exists a unique $3 \times 3$ matrix that can transform sensor triplets to cone responses (XYZ or LMS tristimulus values) for any light spectra. However, the Luther condition does not hold for typical cameras, and therefore, color errors are unavoidable for any $3 \times 3$ color-correction matrix. Color accuracy can be improved either by the better spectral design of camera color filters [IQRB01, Har04, NA06, XX16, FZ20] or by finding more robust functions that map camera RGB responses to XYZ tristimulus color values [Hun93, HLR01, CW02, FMH15]. The former approach requires custom filters that could be difficult to produce and will reduce camera sensitivity. The latter approach can optimize colors only for a limited range of spectral reflectance samples (e.g. 24 patches of the Color Checker) and can never generalize to arbitrary spectral reflectances.

In this work, we follow the third approach, pioneered by Zhu and Finlayson [ZF22], where instead of modifying a camera spectral sensitivity, we optimize the spectral composition of a multispectral light source illuminating the scene, so that accurate colors can be captured irrespectively of the material spectral reflectances. This work is distinct from the previous work in that (a) we explicitly optimize for both camera sensitivity (signal-to-noise ratio) and color accuracy; and (b) we further improve the color accuracy by capturing multiple exposures, each using a different illuminant for each exposure. We demonstrate that such an approach can drastically reduce color errors and improve camera sensitivity. The method is applicable to reflective materials that do not emit or re-emit light (fluorescence), as emission breaks the assumptions of our approach.

## 2. Related Work

In this section, we review the previous works that improve color correction transformation and those that design or optimize better filters or illumination.

### 2.1. Color Correction Transformation

Cameras record device-dependent sensor triplets that need to be transformed to a device-independent color space such as CIE XYZ (before mapping to one of the standard color spaces). Wandell [Wan87] has shown that image sensor values can be synthesized using linear models of the surface reflectance and light spectral power distribution. If the Luther condition is met, a $3 \times 3$ linear color-correction (LCC) matrix could accurately map the sensor triplets to XYZ tristimulus values. In addition, Drew and Funt [DF92] demonstrated that if the surface reflectance spectra can be represented by a 3-dimensional linear model, then a unique $3 \times 3$ matrix exists that could perfectly transform RGB values to XYZ values under any illumination.

Unfortunately, neither the Luther condition nor the 3-dimensional reflectance model is satisfied in real situations, so color error between the ground truth and predicted color values cannot be entirely removed. However, LCC is still favorable in cameras' Image Signal Processing (ISP) modules due to its simplicity and exposure-invariance [KFMA23].

Despite the above advantages, LCC may produce large color errors for some reflectance spectra. To reduce the color error of LCC, some extensions have been proposed, most notably the polynomial [HLR01] and root-polynomial [FMH15] color correction. Polynomial color correction (PCC) incorporates a multi-variate polynomial (typically 2nd degree) without constant terms to map sensor values to the tristimulus XYZ values. Although PCC can significantly reduce the mapping error, it does not adhere to the exposure-invariance property and may produce a large error when the exposure is changed. As an alternative that is exposure-invariant, Finlayson et al. [FMH15] proposed the root-polynomial color correction (RPCC). For the $2^{nd}$ degree case, the basis are $R$, $G$, $B$, $\sqrt{RG}$, $\sqrt{GB}$, $\sqrt{BR}$. Importantly, the root-polynomial expansion was also shown to generate an independent basis, i.e., given enough terms, it can always be used to approximate data to within any desired accuracy (being mindful, of course, that a high order fit is unlikely to generalize to unseen data).

Other color correction methods include 3D Look-Up-Tables [Hun93], hue-plane preserving color correction [AH05, AC16] and neural networks [CW02, Xin11, MM21, KFMA22]. Hung [Hun93] divided the color gamut into many tetrahedrons and performed non-linear tetrahedral interpolation. A $33 \times 33 \times 33$ look-up table was proposed to predict an output color triplet. Anderson and Hardeberg [AH05] splitted the device color gamut into multiple hue-angle delimited subregions defined by the white point and two adjacent chromatic training set colors and applied the same number of white-point-preserving $3 \times 3$ matrices to the individual subregions to perform the color correction. Anderson and Connah [AC16] further improved this method by imposing continuous transitions between the subregions and optimizing for a better white-point-preserving matrices. Cheung and Westland [CW02] compared a 2-layer multi-layer perceptron (MLP) and $2^{nd}$ degree PCC and claimed that they achieve approximately the same performance. MacDonald and Mayer [MM21] demonstrated that 3-layer MLP could reduce the error below $1\Delta E$, which is a just noticeable difference [HP11], on approximately 50% of the real materials dataset over a wide range of daylight illuminations.

### 2.2. Designing Color Filters and Special Illumination

To find the optimal mapping from the sensor triplets to the XYZ tristimulus values, researchers have been trying to find optimal

color filters to gather more color information so that higher accuracy transformation from sensor triplets to tristimulus values could be obtained. Vrhel and Trussell [VT94] built a mathematical framework of selecting $P$ color filters under $K$ viewing illuminants using priori information about the spectral reflectances. Vora and Trussell [VT97] incorporated constraints of physical realizability into the filter optimization and designed a set of filters, which may be fabricated in practice. Over the years, many works [IQRB01, Har04, NA06, XX16, FZ20, Vrh20] continued to make progress in this direction by producing filters with higher color accuracy. Instead of designing new spectral sensitivities from scratch, Finlayson and Zhu [FZ20] formulated a filter optimization framework to find a single color filter that is placed in front of the camera. The filter is optimized so that the product of the filter and camera spectral sensitivity functions (SSFs) is more linearly related to the XYZ color matching functions (CMFs) or the filtered sensor triplets are more linearly related to the target tristimulus values from a data-driven approach. To further improve the drawbacks of the optimized filter, such as filter smoothness and limited transmission, constraints were incorporated in the framework [FZ20]. Although the numerical solution presented in their work indeed achieves a much smaller color error, the filter was not actually manufactured, and whether the filter could be physically reproduced accurately still remains unclear.

However, optimizing for color filters is not the only way to achieve smaller color errors. Zhu and Finlayson [ZF22] tackled the problem from the illumination perspective by creating the color filtering effect with a spectrally tunable lighting system that consists of 10 LED channels with adjustable intensities. Their goal was to find the illumination that made the camera more colorimetric in terms of the Luther condition. Like their work on filter design, the theoretical reduction in measurement error was significant, resulting in half the measurement error. Experiments on real results yielded an error reduction of about 1/3 across viewing conditions. Our method builds on their approach and adds two distinct advancements: a regularization term that minimizes noise in the resulting images and an extension to multiple exposures, which further improve color accuracy.

## 3. Method

We first outline the illumination optimization technique of [ZF22] (Section 3.1), which we extend to use multiple exposures, each with a distinct optimized illuminant (Section 3.2). To ensure that the solution is unique, we introduce a regularization on the color correction matrix (Section 3.3). Then, we formulate the cost function that maximizes the signal-to-noise ratio of captured images (Section 3.4).

### 3.1. Color-Optimal Illuminant

If we had a camera with RGB spectral sensitivity $q(\lambda)$ that meets the Luther condition, we would be able to optimize for a $3\times3$ color-correction matrix $\boldsymbol{M}^{\mathrm{L}}$ such that the linear combination of camera sensitivities closely resembles color matching functions $p(\lambda)$:

$$\underset{\boldsymbol{M}^{\mathrm{L}}}{\arg\min} \sum_{d\in\{X,Y,Z\}} \int_{\lambda} \left( \sum_{c\in\{R,G,B\}} q_c(\lambda)\, m^{\mathrm{L}}_{c,d} - p_d(\lambda) \right)^2 d\lambda. \quad (1)$$

**Table 1:** *Notation and symbols used in the text.*

| Symbol | Description |
|---|---|
| $w_l$ | weight for LED channel $l$ of multispectral light source |
| $b_l(\lambda)$ | emission spectrum of LED channel $l$ |
| $t(\lambda)$ | target light spectrum |
| $e(\lambda)$ | optimized light spectrum |
| $r(\lambda)$ | surface reflectance spectrum |
| $q_c(\lambda)$ | camera spectral sensitivity for primary $c \in \{R,G,B\}$ |
| $p_d(\lambda)$ | CIE XYZ color matching functions, $d \in \{X,Y,Z\}$ |
| $\boldsymbol{M}$ | $3\times3$ color-correction matrix with elements $m_{c,d}$ |
| $\boldsymbol{v}$ | a row vector with RAW linear RGB color values |
| $\boldsymbol{u}$ | a row vector with CIE 1931 XYZ tristimulus values |

In this work, $p(\lambda)$ will represent CIE 1931 color matching functions. The notation used throughout the paper is listed in Table 1.

The problem is that most cameras do not meet the Luther condition, and the linear combination of camera spectral sensitivities does not match the CIE XYZ color matching functions. We cannot easily modify camera spectral sensitivity $q(\lambda)$. However, Zhu and Finlayson [ZF22] noted that we can find the spectrum of an illuminant $e(\lambda)$ for the scene that modulates $q(\lambda)$ in such a way that its linear combination can reproduce color matching functions under a target illuminant $t(\lambda)$ (e.g., D65):

$$\underset{\boldsymbol{M},e(\lambda)}{\arg\min} \sum_{d\in\{X,Y,Z\}} \int_{\lambda} \left( e(\lambda) \sum_{c\in\{R,G,B\}} q_c(\lambda)\, m_{c,d} - t(\lambda) p_d(\lambda) \right)^2 d\lambda$$
$$\text{s.t.} \quad 0 \le e(\lambda) \le e_{\max}. \quad (2)$$

The crux of the method is to illuminate a scene with an artificial (optimized) illuminant $e(\lambda)$ so that the captured RGB values can be later transformed to color-accurate XYZ tristimulus values, as seen under the target illuminant $t(\lambda)$ (e.g., D65).

It is easy to show that when we find suitable $e(\lambda)$, our camera will capture accurate colors regardless of the reflectance spectrum in the scene. If we have a surface with the spectral reflectance $r(\lambda)$, illuminated by D65 light $t(\lambda)$, its CIE XYZ tristimulus responses are:

$$u_d = \int_{\lambda} r(\lambda)\, t(\lambda)\, p_d(\lambda)\, d\lambda, \quad d\in\{X,Y,Z\}. \quad (3)$$

If a camera captures the same surface but under the optimized illuminant $e(\lambda)$, the captured RGB values are:

$$v_c = \int_{\lambda} r(\lambda)\, e(\lambda)\, q_c(\lambda)\, d\lambda, \quad c\in\{R,G,B\}. \quad (4)$$

If the optimization from Eq. (2) is successful, then:

$$e(\lambda) \sum_{c\in\{R,G,B\}} q_c(\lambda)\, m_{c,d} \approx t(\lambda) p_d(\lambda), \quad d\in\{X,Y,Z\} \quad (5)$$

$$r(\lambda)\, e(\lambda) \sum_{c\in\{R,G,B\}} q_c(\lambda)\, m_{c,d} \approx r(\lambda)\, t(\lambda) p_d(\lambda), \quad d\in\{X,Y,Z\} \quad (6)$$

and, from Eq. (3) and Eq. (4), we get:

$$\boldsymbol{v}\boldsymbol{M} \approx \boldsymbol{u}, \quad (7)$$

regardless of the surface reflectance spectrum $r(\lambda)$.

Because we cannot generate an arbitrary illuminant $e(\lambda)$ with our spectrally tunable lights, we are restricted to a space of spectra that is formed by the linear combination of the spectral channels $b_l(\lambda)$ of our multispectral light:

$$e(\lambda) = \sum_{l=1}^{L} w_l\, b_l(\lambda) \tag{8}$$

where $w_l$ are weights of the LED channels. Therefore, in practice, in Eq. (2) we optimize for weights $w_l$ rather than for $e(\lambda)$.

### 3.2. Multiple Illuminants

The optimization from Eq. (2) was originally proposed for a single illuminant. Here, we show that the color accuracy and signal-to-noise ratio can be significantly improved by taking multiple exposures, each under a different illuminant. In that case, our optimization problem becomes:

$$\begin{aligned}&\underset{\boldsymbol{M}^{(k)},\boldsymbol{w}^{(k)}}{\arg\min}\\ &\sum_{d\in\{X,Y,Z\}}\int_{\lambda}\left(\sum_{k=1}^{K}\left(e^{(k)}(\lambda)\sum_{c\in\{R,G,B\}}q_c(\lambda)\,m_{c,d}^{(k)}\right)-t(\lambda)p_d(\lambda)\right)^2 d\lambda\\ &\text{s.t.}\quad 0\le w_l^{(k)}\le 1\quad k=1,\dots,K,\end{aligned} \tag{9}$$

where $K$ is the number of exposures/illuminants, $e^{(k)}(\lambda)$ is $k$-th illuminant given by Eq. (8), $\boldsymbol{w}^{(k)}$ are LED channel weights for $k$-th illuminant, and $\boldsymbol{M}^{(k)}$ is the color-correction matrix for $k$-th exposure.

### 3.3. Color-Correction Matrix Regularizer

The optimization from Eq. (9) is still under-constrained. Assuming the LED weights $w_l$ do not reach the maximum intensity for all the elements, if we multiply $\boldsymbol{w}$ by a factor, we could get the same objective function value if we divide $\boldsymbol{M}$ by the same factor. A small $\boldsymbol{w}$ is not ideal because a dim light requires either a longer exposure time or a higher gain (ISO). At the same time, a $\boldsymbol{M}$ with large elements is also undesirable as they amplify sensor values and thus noise. To address this problem, we introduce a regularization term:

$$\epsilon_{\mathrm{R}} = \beta\sum_{k=1}^{K}\sum_{d\in\{X,Y,Z\}}\sum_{c\in\{R,G,B\}}\left(m_{c,d}^{(k)}\right)^2 \tag{10}$$

which biases color-correction matrices towards smaller magnitudes. Although this regularizer improves the color correction numerical stability, it constrains the capacity of the color-correction matrix towards a more accurate color correction. $\beta$ is the importance of this term, and we explore the effect of $\beta$ in Section 4.2. The regularization term $\epsilon_{\mathrm{R}}$ is then added to the objective function in Eq. (2) or Eq. (9).

### 3.4. Signal-to-Noise Ratio

An illuminant that is optimized for higher color accuracy may result in lower signal-to-noise ratio. Here, we want to find the best balance between the two — the aspect that has not be addressed in the previous works.

We want to maximize the SNR of the resulting image in the target color space (XYZ, after transformation via color-correction matrix). Let $U_d$ be random variables explaining pixel values in that target space, separately for each channel $d$. Then, the signal-to-noise ratio for each channel is given by:

$$r_d(\boldsymbol{v}) = \log_{10}\frac{\mathbb{E}^2(U_d)}{\mathbb{V}(U_d)}\quad d\in\{X,Y,Z\}, \tag{11}$$

where $\mathbb{E}$ is the expected value and $\mathbb{V}$ is the variance operator. $\boldsymbol{v}$ is the sensor RAW RGB triplet, explained below.

A camera's noise can be well modeled in the camera's native color space (linear RGB, or RAW values) as the sum of the Poisson distribution, explaining the photon shot noise, and the normal distribution, explaining the readout and ADC noise [FTKE08, HZM20]:

$$V_c \sim a_c\,\mathrm{Pois}(\psi_c) + \mathcal{N}(0,\sigma_c^2)\quad c\in\{R,G,B\}, \tag{12}$$

where per-color channel parameter $a_c$ and variance $\sigma_c^2$ vary between cameras and across the camera's gain settings. In our experiments, we use the noise parameters measured by Hanji et al. [HZM20]. $\psi_c$ is the effective number of photons reaching the sensor. Since $\mathbb{E}(V_c) = a_c\psi_c = v_c$, we get $\psi_c = {}^{v_c}/_{a_c}$, where $v_c$ is the RAW sensor value (RGB). When the camera's native RGB color values are transformed to the target XYZ color space via a color-correction matrix $\boldsymbol{M}$, the noise distribution will be similarly transformed:

$$\begin{aligned}\mathbb{E}(U_d) &= m_{R,d}\,\mathbb{E}(V_R) + m_{G,d}\,\mathbb{E}(V_G) + m_{B,d}\,\mathbb{E}(V_B)\\ &= m_{R,d}a_R\psi_R + m_{G,d}a_G\psi_G + m_{B,d}a_B\psi_B,\\ \mathbb{V}(U_d) &= m_{R,d}^2\,\mathbb{V}(V_R) + m_{G,d}^2\,\mathbb{V}(V_G) + m_{B,d}^2\,\mathbb{V}(V_B)\\ &= m_{R,d}^2(a_R^2\psi_R+\sigma_R^2) + m_{G,d}^2(a_G^2\psi_G+\sigma_G^2) + m_{B,d}^2(a_B^2\psi_B+\sigma_B^2),\end{aligned} \tag{13}$$

where $m_{c,d}$ are the individual components of the color-correction matrix $\boldsymbol{M}$. Eq. (13) brings a few interesting observations. First, the noise in the target color space becomes correlated across color channels; all color channels in the input color space contribute to the noise in the target color space. Second, the elements of the color-correction matrix $m_{c,d}$ can be both positive and negative. When $m_{c,d}$ is negative, it will reduce signal ($\mathbb{E}(U_d)$), but it will increase the noise $\mathbb{V}(U_d)$ because of the squares in the equation. Therefore, the optimization will bias the solution towards positive coefficients in the color-correction matrix.

With the noise properly modeled, we can introduce the signal-to-noise regularization term into our optimization:

$$\epsilon_{\mathrm{SNR}} = -\frac{\gamma}{N}\sum_{i=1}^{N} r(\boldsymbol{v}_i) \tag{14}$$

where $\gamma$ is the importance of this term (explored in Section 4.2). Note that the amount of noise depends on the sensor RAW values. Therefore, $\boldsymbol{v}_i$ represents the triplet of RAW RGB sensor values, drawn from the unit RGB color cube sampled in the range $[0.0001, 1]$ with the logarithmic spacing. $N$ is the number of color

samples. The regularization term $\epsilon_{\text{SNR}}$ is then added together with $\epsilon_{\text{SNR}}$ from Eq. (10) to the objective function in Eq. (2) or Eq. (9).

### 3.5. Optimization

We use MATLAB's interior-point non-linear constrained optimization method (`fmincon` function) to solve for the optimal LED channel weights and color-correction matrix. We initialize all the parameters with random values between 0 to 1. To avoid local minima, we ran the optimization using 10 different seeds and picked the one that resulted in the smallest values of the objective function. The optimization for three illuminants takes around 48 seconds on Intel Core i7-7700 CPU. Note that this optimization needs to be done only once per camera.

### 3.6. Fine-tuning of color correction matrix

The optimization problem in Eq. (9) jointly solves for the illuminants $e^{(k)}(\lambda)$ and color correction matrices $\boldsymbol{M}^{(k)}$, assuming that the measured camera spectral sensitivities and the optimized illuminants are accurate. However, small inaccuracies in camera sensitivity measurements or deviations in light emitted from the spectrally tunable LEDs may result in degraded color accuracy. To compensate for such inaccuracies, we introduce an optional fine-tuning step in which we refit the color correction matrices $\boldsymbol{M}^{(k)}$. For each patch in a color checker, we capture RGB pixel values under the optimized illuminants $e^{(k)}(\lambda)$ and then measure the corresponding XYZ under the target illuminant $t(\lambda)$ (D65). Then, we optimize:

$$\arg\min_{\boldsymbol{M}} \| \sum_{k=1}^{K} \boldsymbol{v}^{(k)} \boldsymbol{M}^{(k)} - \boldsymbol{u} \|_2 . \tag{15}$$

Such optimization can compensate for the inaccuracies in measurements, but it can also overfit to the color checker spectra. Therefore, we compare results with and without this additional step in the results section below.

## 4. Experiments

### 4.1. Hardware setup

#### 4.1.1. Spectrally Tunable Light Module

The spectrally tunable light module iQ-LED (Image Engineering) consists of 41 high-power LEDs separated into 20 channels on a $10 \times 10$ cm board. Each LED channel peaks at different wavelength and its intensity is tunable. By adjusting the intensity of individual LED channels, we could generate our desired light spectrum such as D65 or other specially designed light spectra.

In order to mix the light produced by individual LEDs located at different positions on the circuit, we made an elbow-shaped wooden case with iQ-LED lights directed towards the elbow (see Figure 2). The interior of the elbow was lined with white paper. We put the light source on top of a small desk so that objects under the desk are illuminated properly, as shown in Figure 2b. To measure the light spectrum, we put a patch of white diffuse reflectance standard (Labsphere's Spectralon has uniform spectral reflectance of about 91%) under the desk and position our spectroradiometer

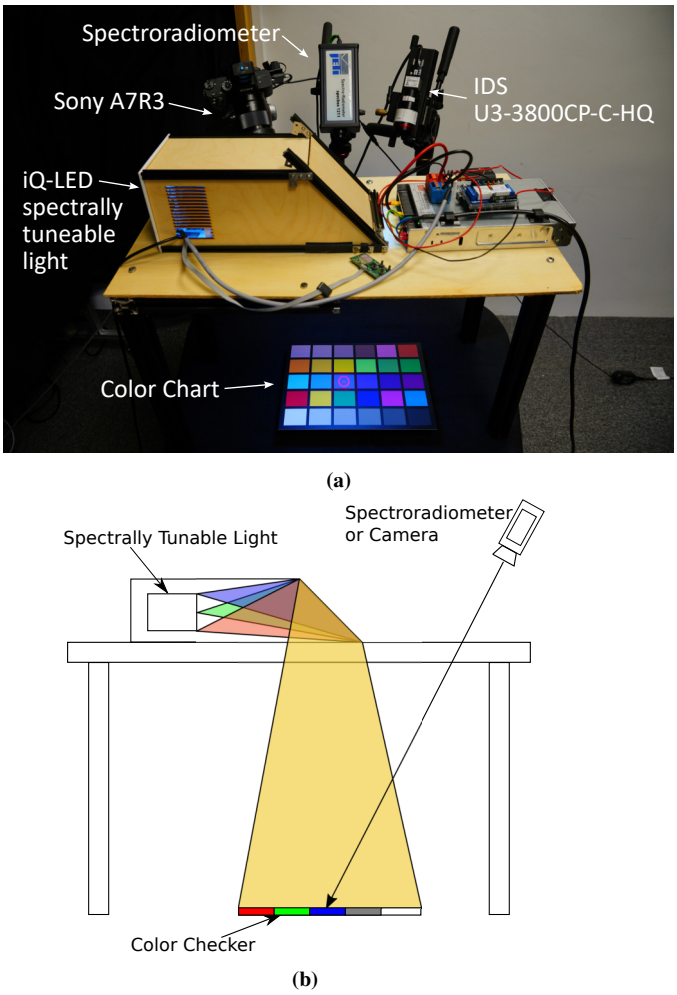


**Figure 2:** *A photograph (a) and a diagram (b) of our image capture setup with the multispectral light, color checker, two cameras, and a spectroradiometer. To avoid inconsistency in illumination, the color checker was manually shifted to capture the color at the same spatial location (indicated by the red laser pointer of the spectroradiometer).*

(JETI Specbos 1211) approximately 60 cm away from the Spectralon. Since the reflectance spectrum of the Spectralon is approximated flat within the visible spectrum, we can obtain the illuminant spectral power distribution (SPD) by measuring the reflected spectrum on the Spectralon. To measure and record the LED channel response, we illuminate one channel at a time and record the individual channel response at full intensity. As plotted in Figure 3, the peak wavelength of all 20 LED channels ranges from 405 nm to 810 nm. In this paper, we work on the wavelength range between 390 nm to 780 nm, which covers the human visible spectrum, so we only use the first 18 LED channels to generate our desired light spectrum.

The channel responses of iQ-LED, shown in Figure 3, are unfortunately not perfectly linear, and therefore, a linear combination of

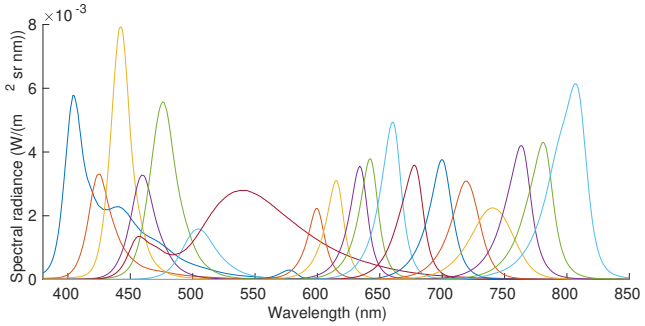


**Figure 3:** *20 channel responses of iQ-LED spectrally tunable light.*

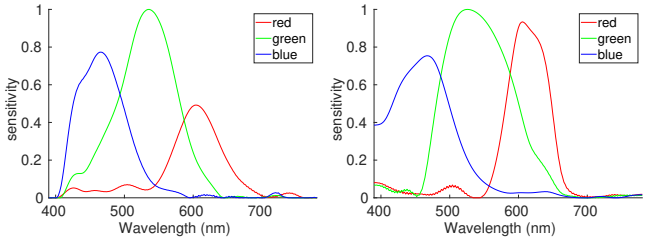


**Figure 4:** *Calculated camera spectral sensitivity functions for Sony A7R3 (left) and IDS U3-3800CP-C-HQ (right).*

channels may not give the desired light spectrum. For that reason, we had to optimize channel driving values with spectroradiometer in the loop to obtain the desired light spectra. The starting point of such an optimization is found by optimizing channel values $\boldsymbol{w}$ to match the desired spectrum $t(\lambda)$ and minimize color difference:

$$\arg\min_{\boldsymbol{w}} \int_{\lambda} (e(\lambda)-t(\lambda))^2 + \alpha \Delta E_{00}(e(\lambda),t(\lambda)) \quad \text{s.t.} \quad 0 \le \boldsymbol{w} \le 1 \tag{16}$$

where $t(\lambda)$ is given in Eq. (8), and $\Delta E_{00}$ is CIEDE2000 color difference [LCR01].

Despite our best attempt to mix individual channels of iQ-LED using the elbow construction, the resulting light could be still spatially non-uniform (have slightly different spectra on the illuminated surface). To improve the accuracy, all our measurements are performed at a single spatial location (see the red laser point in Figure 2a). That is, to measure different patches of a color checker, we manually shifted the color checker.

#### 4.1.2. Color Chart

To validate the effectiveness of our method, we need a reflectance dataset where we can associate the numerical integration values to pixel values in the real image. We use the 30-color Preferred Memory Colour Chart—PMCC (THOUSLITE), which consists of 18 "memory" colors, 6 unitary hue colors and 6 neutral scale colors. With the spectrally tunable light illumination on top of the PMCC, the reflected spectrum of each color patch is the product of the illuminant SPD shining on the patch and its surface reflectance spectrum, so the surface reflectance spectrum of each color patch could be calculated by dividing the reflected spectrum by the illuminant SPD. We first put the Spectralon at the position of the red laser point and illuminate the scene using our approximated D65. After the spectroradiometer measurement, we replace the Spectralon with the PMCC and measure the color patch reflected spectrum one by one. By dividing the D65 spectroradiometer measurement from the spectroradiometer measurement of individual color patch, we can easily obtain the reflectance of each color patch.

#### 4.1.3. Cameras

Our technique requires the knowledge of camera spectral sensitivity functions (SSFs), and we will discuss our approach to estimating the SSFs $q_c(\lambda)$ in the following part.

As we use vector-matrix representations in the optimization, the $1 \times 3$ vector of sensor value triplet $v_c$ defined in Eq. (3) can be calculated as the element-wise vector product

$$v_c = e \cdot s \cdot q_c, \quad c \in \{r,g,b\} \tag{17}$$

where $q_c$ is the camera SSF vector with red, green and blue; $s$ is the vector of the surface reflectance spectrum; $e$ is the vector of illuminant SPD. According to Eq. 17, we could construct a matrix equation for solving $q_c$ by iterating a 3-step procedure for $n$ known illuminants onto one color patch: shining one illuminant onto the color patch, capturing the scene with the target camera, and taking the image mean value of the central area on the color patch as $v_c$. By placing each of the 30 color patches under the fixed position, we can accumulate $30n$ equations for solving the SSFs. Let $D$ denote the matrix where we put the element-wise product of each illuminant SPD $e_i$ and color patch reflectance spectrum $s_j$ in each column, so

$$D = [e_1 \cdot s_1, e_1 \cdot s_2, ...e_1 \cdot s_{30}, e_2 \cdot s_1, ...e_n \cdot s_{30}]. \tag{18}$$

For channel $c$ ($c \in \{r,g,b\}$), the sensor value vector $\nu$ can be constructed from the mean value of each cropped color patch arranged in the same order as $D$

$$\nu = [\nu_{1,1}, \nu_{1,2}, ...\nu_{1,30}, \nu_{2,1}, ...\nu_{n,30}] \tag{19}$$

where $\nu_{i,j}$ corresponds to the image value of color patch $s_j$ under illuminant SPD $e_i$. The optimization equation for single channel SSF $\mu$ can be formulated as the summation of a data term $\|D^{\mathsf{T}}\mu - \nu^c\|_2^2$ and a smoothness term $\|\nabla\mu\|_2^2$ weighted by $\alpha$

$$\arg\min_{\mu} \|D^{\mathsf{T}}\mu - \nu^{\mathsf{T}}\|_2^2 + \alpha\|\nabla\mu\|_2^2 \quad \text{s.t.} \quad 0 \le \mu \le 1 \tag{20}$$

where the SSF vector $\mu$ should be greater than or equal to 0 as the sensor does not produce any negative response. When we plug $\nu$ for red, green and blue values into the equation, we could obtain the SSFs for all three channels.

To construct matrix $D$ in Eq. (20), we generated 24 illuminants, namely 18 single channel LED responses and 6 typical illuminants such as D65. Sequential quadratic programming is employed to solve our objective equation and we set $\alpha$ to 0.001. In our experiments, we used two cameras: Sony A7R3 photographic camera with a Sony 35 mm/F1.8 lens and IDS U3-3800CP-C-HQ computer vision camera with Kowa 25 mm/F1.8 lens (LM25FC24M). The estimated SSFs between 390 nm to 780 nm of the two cameras are shown in Figure 4.

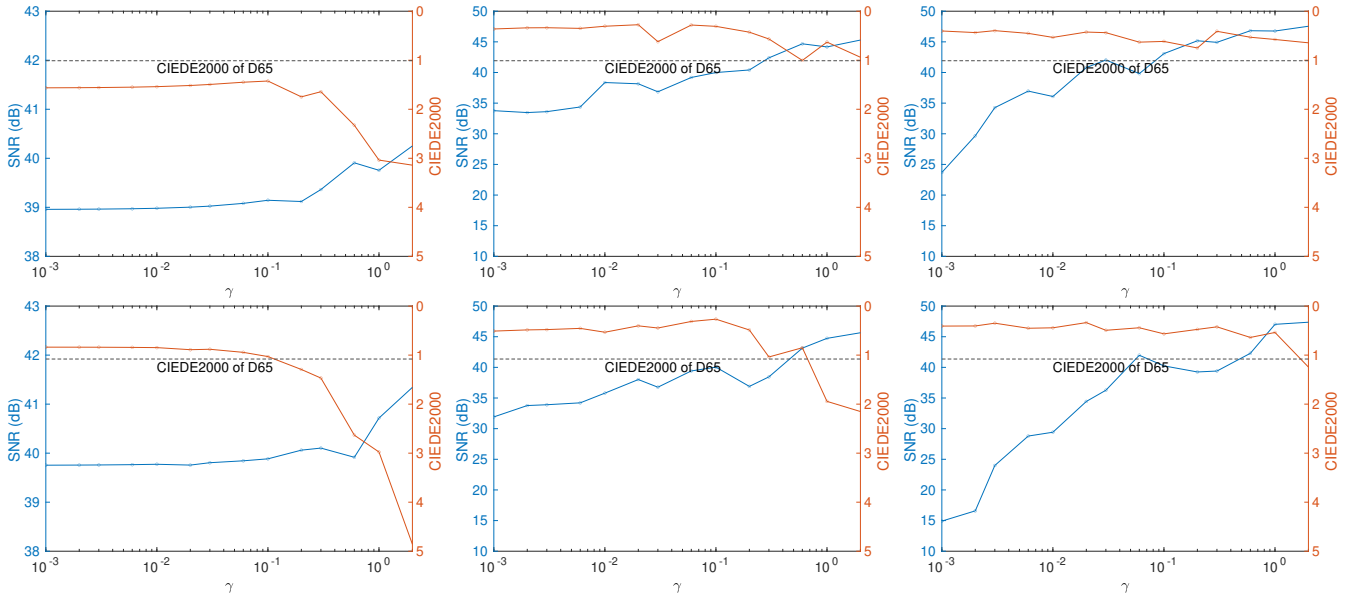


**(a)** *1 illuminant* **(b)** *2 illuminants* **(c)** *3 illuminants*

**Figure 5:** *SNR and CIEDE2000 values of Sony A7R3 (first row) and IDS U3-3800CP-C-HQ (second row) across the range of the regularization term γ. The higher γ, the larger the influence of the SNR term is. The plots correspond to the cases with 1, 2, and 3 illuminants.*

### 4.2. Regularization of Hyper-Parameters

To find the optimal illuminants, we need to select the right values for our regularization hyper-parameters γ (Eq. (14), SNR regularizer) and β (Eq. (10), color-correction matrix regularizer). We explore the range of hyper-parameters by finding the optimum solution for a range of their values. To find suitable γ, we set $\beta = 0$ and plot SNR and $\Delta E$ for the two cameras in Figure 5. The plots show a trade-off between SNR and $\Delta E$, with different γ needed for a different number of illuminants. The plots also show that the SNR regularization is essential when taking multiple exposures, as the lack of such regularization results in very low SNR (see the left-most point in Figure 5-c). We select $\gamma = 0.1$ for both cameras as a compromise.

To find suitable β, we plot SNR and $\Delta E$ of the optimal solution for the two cameras similarly and put the figure in Figure 6. The plots show that when β is too large, $\Delta E$ increases rapidly. But when it is too small, we do not gain the benefit of increased SNR with multiple exposures. We choose $\beta = 1.0$ for Sony A7R3 and $\beta = 0.2$ for the IDS camera.

### 4.3. Optimal Illuminants and Transformed Camera SSFs

The optimized illuminants for Sony A7R3 and the IDS camera are shown in the top row of Figure 8 and Figure 9. The bottom row of the corresponding figures shows how well the camera SSFs, modulated by the illuminants and transformed by the color-correction matrix, align with the CIE 1931 color matching functions. The match becomes better as we increase the number of illuminants.

### 4.4. Validation on Color Chart

We test the color accuracy and SNR on the PMCC in four different configurations: in simulation, to eliminate all potential sources of inaccuracy and estimate the lower limit of the color accuracy, and based on actual images captured with our cameras. Then, each of those is tested with and without the optional color correction fine-tuning step (Section 3.6). Our simulation assumes that camera SSFs from Figure 4, measured spectral reflectances of the color patches, and the optimized illuminants are accurate. The simulated camera values are computed according to Eq. (3). The real-image results are based on RGB pixel values read from the camera's RAW files. We used the Python rawpy library for decoding the RAW files (before demosaicing) for the A7R3 camera and used the APIs to obtain RAW images from the IDS camera. The black level was subtracted from the RAW values in the A7R3 images. The results for both cameras are listed in Table 2, and the comparison of captured colors is shown in Figure 12.

First, we can observe that the real-image and simulated color error values are very close to each other, confirming the accuracy of our measurements. The only exception here is the 1-illuminant real-image case without the color correction fine-tuning step for the IDS camera (2.38 average $\Delta E_{00}$ error). This error was likely caused by the inaccuracy in the estimated IDS SSFs and it was corrected by the fine-tuning step.

Second, we can observe that the fine-tuning step improved the color accuracy in almost all cases and, therefore, we recommend it for the pipeline. Our further analysis focuses on the results with the fine-tuning step. Note that this step was not part of the method proposed in [ZF22].

Third, we can observe an improvement in color accuracy as we

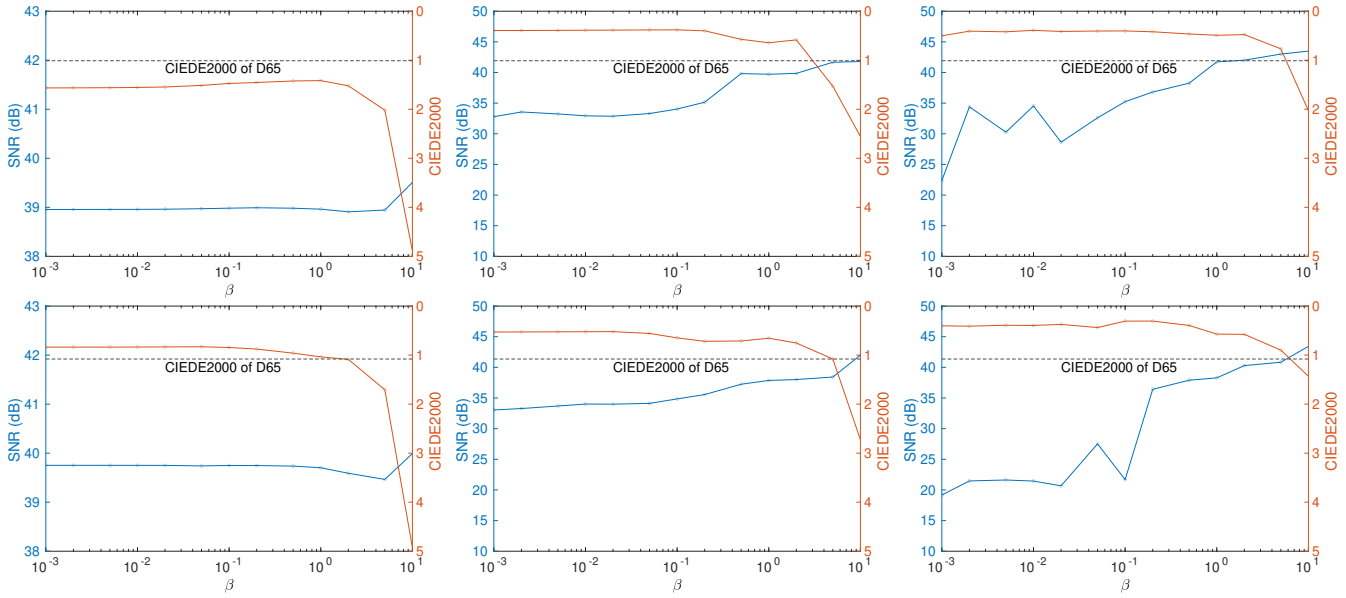


**(a)** *1 illuminant* **(b)** *2 illuminants* **(c)** *3 illuminants*

**Figure 6:** *SNR and ΔE values of Sony A7R3 (first row) and IDS U3-3800CP-C-HQ (second row) across the value of the regularization term β. The higher β, the smaller is the norm of the color-correction matrix. The plots correspond to the cases with 1, 2 and 3 illuminants.*

switch from the standard D65 illuminant, to 1, 2 or 3 optimized illuminants. The average errors are very small, often much below one CIEDE2000 unit. However, in practical applications, we are even more interested in the worst-case scenario, here represented as the maximum error and listed in parenthesis. We can reduce the maximum error for Sony camera (with the fine-tuning step) from 2.46 to 0.67 $\Delta E_{00}$ units, and from 3.17 to 1.91 $\Delta E_{00}$ units for the IDS camera. This is a very substantial reduction of error that is rarely achieved with RGB cameras. We can also notice diminishing gains as we switch from 2 to 3 illuminants, suggesting that further improvements are difficult to achieve. The average errors are comparable across both cameras, however, the maximum errors for real images are smaller for Sony A7R3.

To demonstrate that the technique generalizes to other illuminants (different from D65), we show in Table 4 simulation results for illuminants D50, A and LED-RGB1 whose SPDs are plotted in Figure 7. The results indicate that both illuminant A and LED-RGB1 pose a challenge for accurate color reproduction, especially for Sony A7R3 camera. The large color errors for these illuminants can be reduced by a factor of 2 or 3 when taking multiple exposures under optimized illuminants.

### 4.5. Validation on the SFU dataset

The validation performed on the PMCC let us conveniently analyze all aspects of the method, but it introduces two problems. First, the PMCC is not representative of the range of reflectance spectra found in the real world. Second, we train and test on the same data, risking overfitting. To overcome these problems, we test the method on the widely-used SFU dataset [BMFC02], consisting of 1995 relectance spectra.

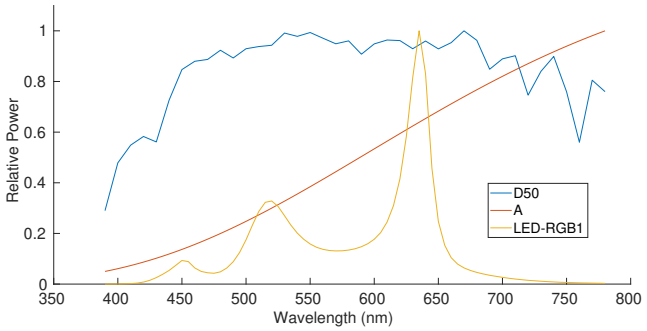


**Figure 7:** *Spectral power distributions (SPDs) of D50, A and LED-RGB1.*

The validation on the SFU dataset is straightforward to perform in simulation, however, as we do not have physical samples of the materials to capture camera images. Therefore, instead, we employ the approach proposed in [ZF22] and use a linear combination of reflectance (and RGB) values of the PMCC as a proxy for that dataset. Zhu and Finlayson [ZF22] observed that the majority of spectra found in the SFU dataset can be well approximated by a linear combination of four reflectances of their color chart:

$$\arg\min_{\alpha} \int_{\lambda} \left( r_{\mathrm{SFU}}(\lambda) - \sum_{i=1}^{25} \alpha_i r_i(\lambda) \right)^2 d\lambda \quad \text{s.t.} \, \|\alpha\|_0 = 4 \,. \qquad (21)$$

where $r_{\mathrm{SFU}}(\lambda)$ is a reflectance spectrum from the SFU dataset, $\alpha$ is a 25-dimensional weight vector and $r_i(\lambda)$ is one of the 25 source spectra from the PMCC. From the 30 source patches in PMCC, we remove 5 gray-scale patches (but keep white) so in total we have

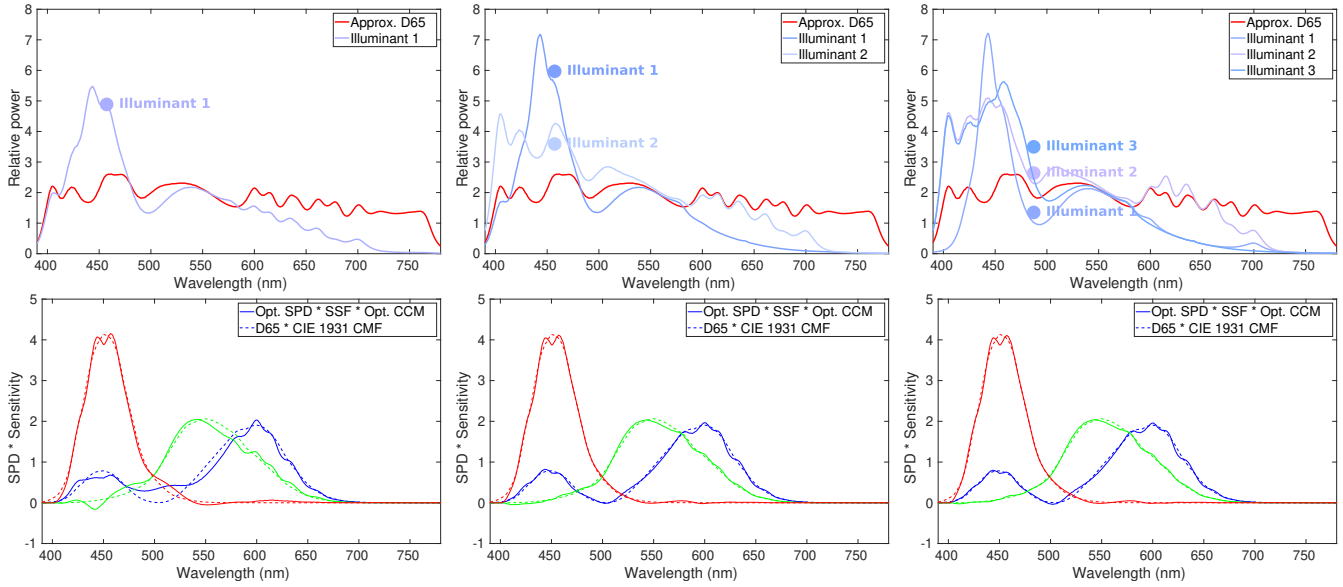


**(a)** *1 illuminant* **(b)** *2 illuminants* **(c)** *3 illuminants*

**Figure 8:** *Top: Optimized illuminant(s) of Sony A7R3 for 1, 2 and 3 exposures. Bottom: The product of the optimized illuminant(s) and camera sensitivity transformed via color-correction matrix vs. the product of D65 and CIE 1931 color matching functions. The closer the solid line to the dashed lines of color matching functions is, the more accurate the color is.*

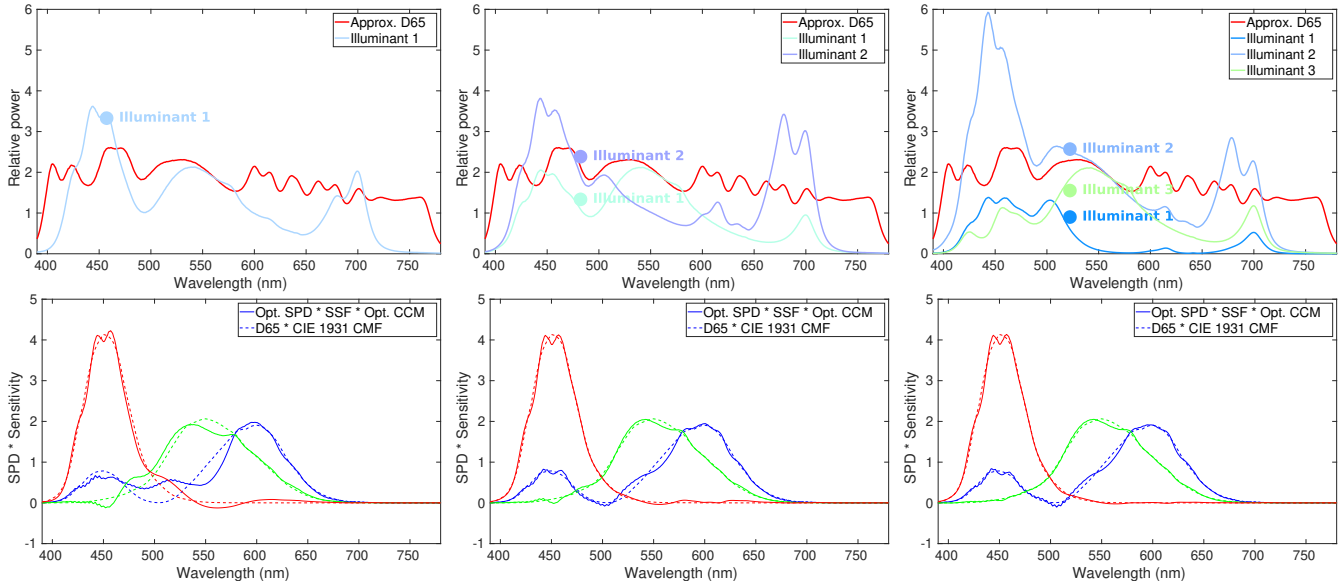


**(a)** *1 illuminant* **(b)** *2 illuminants* **(c)** *3 illuminants*

**Figure 9:** *Top: Optimized illuminant(s) of IDS U3-3800CP-C-HQ for 1, 2 and 3 exposures. Bottom: The product of the optimized illuminant(s) and camera sensitivity transformed via the color-correction matrix vs. the product of D65 and CIE 1931 color matching functions. The closer the solid line to the dashed lines of color matching functions is, the more accurate the color is.*

25 linearly independent source spectra to build our approximation for the SFU dataset. Once we know the linear combination of the reflectance spectra, we can predict the linear RGB values captured by the camera for this reflectance as

$$\boldsymbol{v}_{\mathrm{SFU}} = \sum_{i=1}^{25} \alpha_i \boldsymbol{v}_i(\lambda)\,, \tag{22}$$

where $\boldsymbol{v}_i$ is the linear RGB pixel value captured for the PMCC patch $i$.

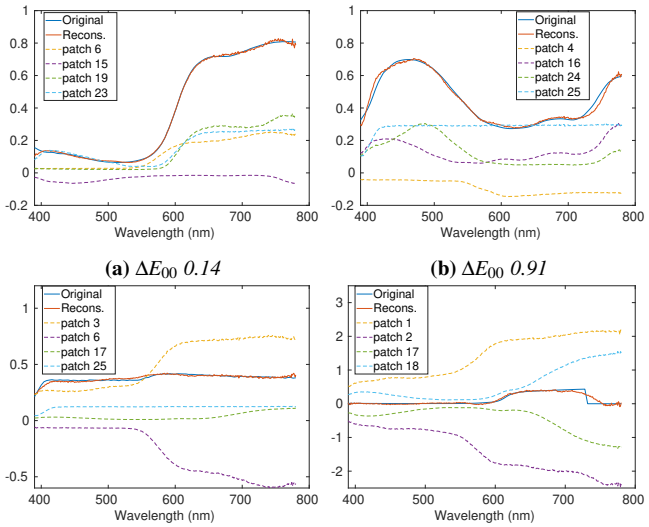


**(a)** $\Delta E_{00}$ *0.14* **(b)** $\Delta E_{00}$ *0.91*

**(c)** $\Delta E_{00}$ *4.90* **(d)** $\Delta E_{00}$ *15.65*

**Figure 10:** *Reflectance in the SFU dataset and their linear reconstruction. The $\Delta E_{00}$ (CIEDE2000) between the original spectra and its reconstruction is listed in the caption.*

According to our spectrum reconstruction results of Eq. (21), the mean error is 1.53% and the maximum error is 2.74%. From all the SFU linear approximations, we select those with reflectance strictly in the range $[0, 1]$ and CIEDE2000 smaller than 1 to guarantee that the linear approximation is physically possible and accurate. In total, 1079 out of 1995 spectra are selected as our SFU validation dataset. Two examples of well-approximated reflectance spectra are plotted in the top row of Figure 10. By contrast, two spectra that cannot be well approximated are shown in the bottom row of Figure 10 — the linear combination either drops below 0 or differs from the original spectra by a large margin. We test the selected 1079 validation spectra using the linear combination of the PMCC linear RGB values and summarize the results as the "Approximated camera" part in Table 3. For completeness, we also report the results for the entire SFU dataset in the "Simulation" part of Table 3, following the simulation approach from Section 4.4.

The results in Table 3 confirm the trends observed in Table 2. We can observe substantial improvement up to 2 optimized illuminants for the Sony camera, and further reduction of error for 3 illuminants for the IDS camera. The mean errors are comparable to those observed for the PMCC. The additional fine-tuning step ("with CC" columns) shows significant improvement, indicating that the optimization does not overfit to the color chart. It should be noted that the table reports the 95th percentile in the parenthesis rather than the maximum error as this dataset is much larger. Those results confirm that the method can generalized to any spectra and achieve very high color accuracy with $\Delta E_{00}$ below 1, even for a computer vision camera. In the top row of Figure 11, we plot examples of two SFU spectra resulting in large color errors when captured under D65. The bottom row of that figure shows that the color correction for our optimized three illuminants results in imperceptible color differences.

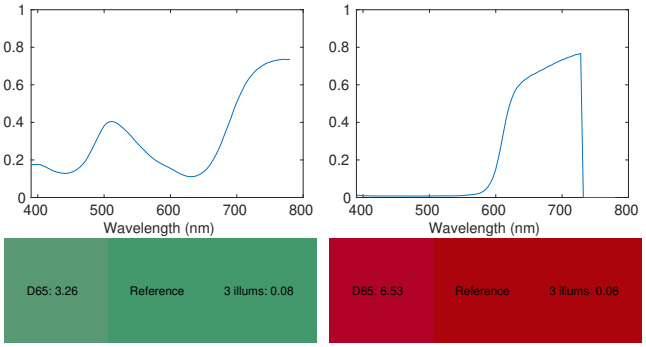


**Figure 11:** *Top: SFU reflectance spectra that result in large color errors when captured with Sony A7R3 under D65. Bottom: The reconstructed colours for D65 illuminant and our optimized triplet of illuminates. The values show $\Delta E_{00}$ (CIEDE2000) error.*

### 4.6. Real-Camera Captures

We also captured a Rubik's cube with Sony A7R3 using the optimized illuminants and D65 as shown in Figure 1. Although we do not have reference colors for those objects, the figures demonstrate that the images produced with optimized illuminants give colors that are close to reality.

## 5. Conclusions

Capturing multiple exposures is widely used to reduce noise (improve camera sensitivity) or expand dynamic range. Here, we show that capturing multiple exposures can also be used to improve the camera's color accuracy. We optimize a spectral composition of light so that when one or more images are captured under the optimized light spectra, the linear combination of captured RGB values can be transformed into CIE XYZ tristimulus color values. This let us radically improve color accuracy ($\Delta E_{00}$ below 1) while maintaining or improving the signal-to-noise ratio and without any modification to the camera. In this project, we employ a spectrally tuneable LED light module (iQ-LED), which offers a practical solution for studio photography or scanning (e.g., material acquisition). However, we foresee that the technique can be adapted to a multi-spectral flash in the future.

## Acknowledgements

We would like to thank Alejandro Sztrajman for his help on the iQ-LED software. We would also like to thank Jinze Sha for his help on building the wooden case of the iQ-LED. This work was partially supported by the EPSRC research grants EP/P007902/1 and EP/S028730/1.

## References

[AC16] ANDERSEN C. F., CONNAH D.: Weighted constrained hue-plane preserving camera characterization. *IEEE Transactions on Image Processing 25*, 9 (2016), 4329–4339. 2

**Table 2:** *Simulation and real image results of PMCC on both cameras. The "without CC" columns correspond to matching spectral responses alone, and "with CC" corresponds to fine-tuning the color correction matrix* $\boldsymbol{M}$ *on the PMCC patches (Section 3.6). The numbers in brackets represent the maximum* $\Delta E_{00}$ *(CIEDE2000).*

| | | A7R3 | | | IDS | | |
|---|---|---|---|---|---|---|---|
| | | without CC ($\Delta E_{00}$) | with CC ($\Delta E_{00}$) | SNR [dB] | without CC ($\Delta E_{00}$) | with CC ($\Delta E_{00}$) | SNR [dB] |
| Simulation | D65 | 2.16 (5.00) | 0.90 (2.72) | N.A. | 1.92 (3.38) | 1.05 (2.54) | N.A. |
| | 1 illuminant | 1.58 (4.19) | 0.51 (1.76) | 39.12 | 1.16 (2.68) | 0.57 (1.85) | 39.91 |
| | 2 illuminants | 0.31 (0.60) | 0.22 (0.53) | 39.34 | 0.84 (2.17) | 0.22 (0.48) | 39.56 |
| | 3 illuminants | 0.25 (0.50) | 0.22 (0.52) | 42.04 | 0.34 (0.91) | 0.22 (0.50) | 42.37 |
| Real image | D65 | 2.13 (5.12) | 0.93 (2.46) | 36.02 | 1.95 (3.74) | 1.19 (3.17) | 31.98 |
| | 1 illuminant | 1.68 (4.30) | 0.65 (1.68) | 35.99 | 2.38 (4.50) | 0.90 (2.57) | 32.11 |
| | 2 illuminants | 1.03 (3.18) | 0.28 (0.62) | 36.33 | 1.54 (2.52) | 0.44 (1.68) | 32.00 |
| | 3 illuminants | 0.96 (2.24) | 0.25 (0.67) | 37.86 | 1.28 (2.49) | 0.39 (1.91) | 31.81 |

**Table 3:** *The* $\Delta E_{00}$ *(CIEDE2000) results for the SFU dataset for both cameras. The 95th percentile* $\Delta E_{00}$ *of the dataset is also listed in the brackets. The results for the 1079 spectra approximated by a linear combination of PMCC patches (see the text) are listed in the left half of the table. The right half shows the simulation results for all spectra in the SFU dataset.*

| | Approximated camera | | | | Simulation | | | |
|---|---|---|---|---|---|---|---|---|
| | A7R3 | | IDS | | A7R3 | | IDS | |
| | without CC | with CC | without CC | with CC | without CC | with CC | without CC | with CC |
| D65 | 2.26 (4.06) | 0.86 (2.19) | 1.93 (3.38) | 1.01 (2.38) | 2.10 (3.83) | 1.06 (3.03) | 2.01 (3.37) | 1.05 (2.57) |
| 1 illuminant | 1.63 (3.59) | 0.73 (1.67) | 2.42 (4.48) | 1.13 (2.66) | 1.70 (3.91) | 0.54 (3.33) | 1.22 (2.57) | 0.63 (3.99) |
| 2 illuminants | 1.00 (2.17) | 0.29 (0.54) | 1.37 (2.28) | 0.39 (0.78) | 0.29 (0.80) | 0.07 (0.73) | 0.66 (1.38) | 0.11 (0.79) |
| 3 illuminants | 1.06 (2.02) | 0.32 (0.83) | 1.22 (2.18) | 0.48 (1.49) | 0.24 (0.62) | 0.06 (0.60) | 0.42 (0.88) | 0.03 (0.72) |

**Table 4:** *Simulation results of PMCC on more illuminants. The numbers in brackets represent the maximum* $\Delta E_{00}$ *(CIEDE2000).*

| | A7R3 | | IDS | |
|---|---|---|---|---|
| | without CC | with CC | without CC | with CC |
| D50 | 2.89 (6.25) | 1.16 (3.20) | 2.17 (4.65) | 1.02 (3.29) |
| 1 illuminant | 1.52 (2.80) | 0.72 (2.07) | 1.64 (2.71) | 0.74 (2.05) |
| 2 illuminants | 0.98 (2.06) | 0.66 (1.66) | 0.97 (2.35) | 0.62 (1.59) |
| 3 illuminants | 0.87 (2.07) | 0.62 (1.63) | 0.80 (2.32) | 0.62 (1.63) |
| A | 5.64 (13.15) | 1.60 (4.79) | 3.70 (10.64) | 1.46 (6.49) |
| 1 illuminant | 3.50 (8.19) | 1.27 (3.24) | 4.75 (8.07) | 1.15 (2.14) |
| 2 illuminants | 2.27 (5.33) | 1.19 (2.78) | 2.12 (4.22) | 1.19 (2.68) |
| 3 illuminants | 1.82 (4.35) | 1.18 (2.71) | 1.70 (3.28) | 1.18 (2.70) |
| LED-RGB1 | 7.88 (13.89) | 1.56 (4.89) | 6.80 (13.53) | 1.68 (5.67) |
| 1 illuminant | 7.68 (16.52) | 0.99 (2.54) | 5.00 (9.16) | 1.32 (4.48) |
| 2 illuminants | 3.03 (5.86) | 0.95 (2.59) | 3.19 (5.15) | 0.97 (2.54) |
| 3 illuminants | 2.17 (4.28) | 0.95 (2.50) | 1.91 (3.41) | 0.96 (2.51) |


[AH05] ANDERSEN C. F., HARDEBERG J. Y.: Colorimetric characterization of digital cameras preserving hue planes. In *Color and Imaging Conference* (2005), vol. 13, pp. 141–146. 2

[BMFC02] BARNARD K., MARTIN L., FUNT B., COATH A.: A data set for color research. *Color Research & Application 27*, 3 (2002), 147–151. 8

[CW02] CHEUNG T. L. V., WESTLAND S.: Color camera characterisation using artificial neural networks. In *Color and Imaging Conference* (2002), pp. 117–120. 2

[DF92] DREW M. S., FUNT B. V.: Natural metamers. *CVGIP: Image Understanding 56*, 2 (1992), 139–151. 2

[FMH15] FINLAYSON G. D., MACKIEWICZ M., HURLBERT A.: Color correction using root-polynomial regression. *IEEE Transaction on Image Processing 24*, 5 (2015), 1460–1470. 2

[FTKE08] FOI A., TRIMECHE M., KATKOVNIK V., EGIAZARIAN K.: Practical poissonian-gaussian noise modeling and fitting for single-image raw-data. *IEEE Transaction on Image Processing 17*, 10 (2008), 1737–1754. 4

[FZ20] FINLAYSON G. D., ZHU Y.: Designing color filters that make cameras more colorimetric. *IEEE Transaction on Image Processing 30* (2020), 853–867. 2, 3

[Har04] HARDEBERG J. Y.: Filter selection for multispectral color image acquisition. *Journal of Imaging Science and Technology 48*, 2 (2004), 105–110. 2, 3

[HLR01] HONG G., LUO M. R., RHODES P. A.: A study of digital camera colorimetric characterization based on polynomial modeling. *Color Research & Application 26*, 1 (2001), 76–84. 2

[HP11] HUNT R. W. G., POINTER M. R.: *Measuring colour*. John Wiley & Sons, 2011. 2

[Hun93] HUNG P.-C.: Colorimetric calibration in electronic imaging devices using a look-up-table model and interpolations. *Journal of Electronic imaging 2*, 1 (1993), 53–61. 2

[HZM20] HANJI P., ZHONG F., MANTIUK R. K.: Noise-aware merging of high dynamic range image stacks without camera calibration. In *ECCV Workshops* (2020), pp. 376–391. 4

[IQRB01] IMAI F. H., QUAN S., ROSEN M. R., BERNS R. S.: Digital camera filter design for colorimetric and spectral accuracy. In *Proc. of third international conference on multispectral color science* (2001), pp. 13–16. 2, 3

[Ive15] IVES H. E.: The transformation of color-mixture equations from one system to another. *Journal of the Franklin Institute 180*, 6 (1915), 673–701. 2

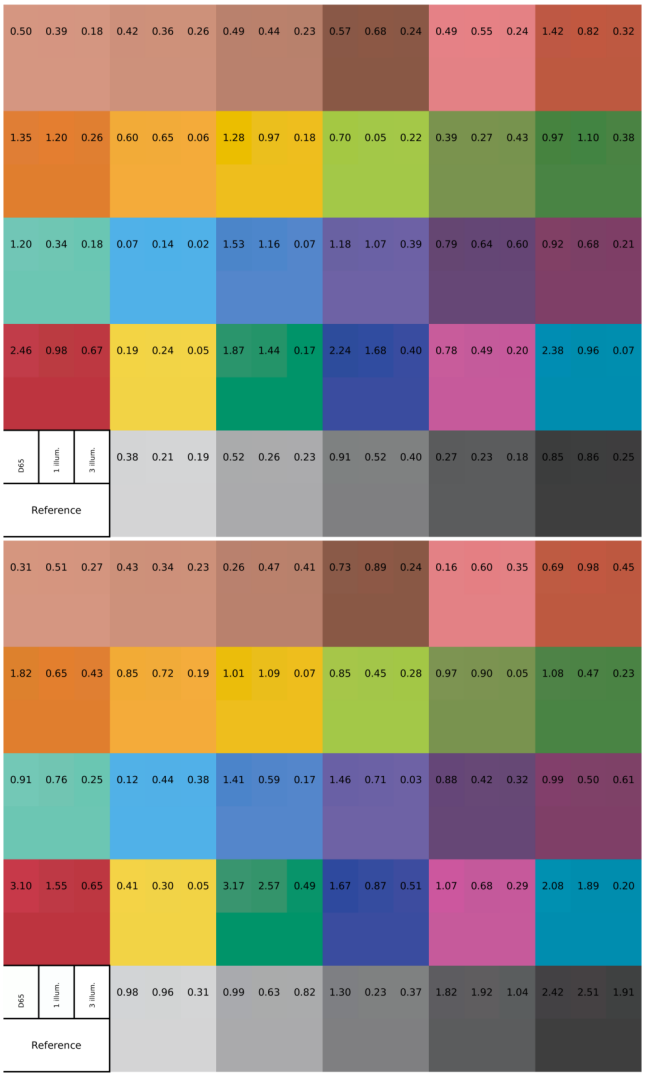


**Figure 12:** *The visualization of 30 patches of the PMCC, captured with both cameras. For each color patch, the bottom part is the reference color as measured with the spectroradiometer. The top half is horizontally split into 3 regions — left: the color captured under the approximated D65; center: color captured under 1 optimized illuminant; right: color obtained after merging information from 3 exposures, each captured under a different illuminant. The results are reported for real-image captures and with the color-correction fine-tuning step. Note that the cyan color (row 4 and column 6) is clipped in the figure because it is outside the sRGB (BT.709) color gamut.*

[KFMA22] KUCUK A., FINLAYSON G. D., MANTIUK R., ASHRAF M.: An exposure invariant neural network for colour correction. In *Color and Imaging Conference* (2022), pp. 176–181. 2

[KFMA23] KUCUK A., FINLAYSON G. D., MANTIUK R., ASHRAF M.: Performance comparison of classical methods and neural networks for colour correction. *Journal of Imaging 9*, 10 (2023), 214. 2

[LCR01] LUO M. R., CUI G., RIGG B.: The development of the cie 2000 colour-difference formula: Ciede2000. *Color Research & Application 26*, 5 (2001), 340–350. 6

[MM21] MACDONALD L., MAYER K.: Camera colour correction using neural networks. In *London Imaging Meeting* (2021), pp. 54–57. 2

[NA06] NG D.-Y., ALLEBACH J. P.: A subspace matching color filter design methodology for a multispectral imaging system. *IEEE Transaction on Image Processing 15*, 9 (2006), 2631–2643. 2, 3

[VL27] VON LUTHER R.: Aus dem gebiet der farbreizmetrik. *Zeitschrift fur Technishe Physik 12* (1927), 540–558. 2

[Vrh20] VRHEL M. J.: Improved camera color accuracy in the presence of noise with a color prefilter. In *Color and Imaging Conference* (2020), vol. 28, pp. 187–192. 3

[VT94] VRHEL M. J., TRUSSELL H. J.: Filter considerations in color correction. *IEEE Transaction on Image Processing 3*, 2 (1994), 147–161. 3

[VT97] VORA P. L., TRUSSELL H. J.: Mathematical methods for the design of color scanning filters. *IEEE Transaction on Image Processing 6*, 2 (1997), 312–320. 3

[Wan87] WANDELL B. A.: The synthesis and analysis of color images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1 (1987), 2–13. 2

[Xin11] XINWU L.: A new color correction model for based on bp neural network. *Advances in Information Sciences and Service Sciences 3*, 5 (2011). 2

[XX16] XU P., XU H.: Filter selection based on light source for multispectral imaging. *Optical Engineering 55*, 7 (2016). 2, 3

[ZF22] ZHU Y., FINLAYSON G. D.: Matched illumination: using light modulation as a proxy for a color filter that makes a camera more colorimetric. *Optics Express 30*, 12 (2022), 6–24. 2, 3, 7, 8