

Spectral Sensitivity Estimation Without a Camera

Grigory Solomatov, *Member, IEEE*, and Derya Akkaynak, *Member, IEEE*

Abstract—A number of problems in computer vision and related fields would be mitigated if camera spectral sensitivities were known. As consumer cameras are not designed for high-precision visual tasks, manufacturers do not disclose spectral sensitivities. Their estimation requires a costly optical setup, which triggered researchers to come up with numerous indirect methods that aim to lower cost and complexity by using color targets. However, the use of color targets gives rise to new complications that make the estimation more difficult, and consequently, there currently exists no simple, low-cost, robust *go-to* method for spectral sensitivity estimation that non-specialized research labs can adopt. Furthermore, even if not limited by hardware or cost, researchers frequently work with imagery from multiple cameras that they do not have in their possession. To provide a practical solution to this problem, we propose a framework for spectral sensitivity estimation that not only does not require *any* hardware (including a color target), but also does not require physical access to the camera itself. Similar to other work, we formulate an optimization problem that minimizes a two-term objective function: a camera-specific term from a system of equations, and a universal term that bounds the solution space. Different than other work, we utilize publicly available high-quality calibration data to construct both terms. We use the colorimetric mapping matrices provided by the Adobe DNG Converter to formulate the camera-specific system of equations, and constrain the solutions using an autoencoder trained on a database of ground-truth curves. On average, we achieve reconstruction errors as low as those that can arise due to manufacturing imperfections between two copies of the same camera. We provide our code and predicted sensitivities for 1,000+ cameras at <https://github.com/COLOR-Lab-Eilat/Spectral-sensitivity-estimation>, and discuss which tasks can become trivial when camera responses are available.

Index Terms—Adobe DNG SDK, colorimetric mapping, color constancy, reflectance recovery, illumination estimation

1 Introduction

CONSUMER digital cameras (i.e., point-and-shoot, DSLR, mirrorless cameras and smartphones) are not tools designed for scientific imaging, i.e., they are not scientific light-measuring devices [1–3]. Yet their outputs—digital photos and videos—constitute major sources of data for image processing, colorimetry, computational photography, computer vision, and machine learning. Typically, research utilizing consumer camera imagery focuses on the development of filters or algorithms that alter the visual appearance of images [4–7], or on the understanding of scene content [8–10] and structure [11, 12] with downstream goals like recognizing, tracking, or counting objects. For many of these goals, successfully recovering scene reflectance and/or illumination is key (and often the main goal itself), but these tasks are complicated by the fact that consumer cameras do not capture colors in a standardized way [13, 14].

Each camera captures colors in its own, device-specific color space. This is because manufacturers make different design choices for each sensor, which materialize within the limitations inherent to silicon-based semiconductor technology [18]. As a result, each make and model of camera is characterized by its unique set of spectral sensitivities that define how that camera responds to light. Effectively, this means that different cameras will register a different set of RGB values for a given scene—even if the scene is imaged under identical conditions [13, 19]. If the camera’s color bias could be eliminated, the recovery of scene reflectance (e.g., for material

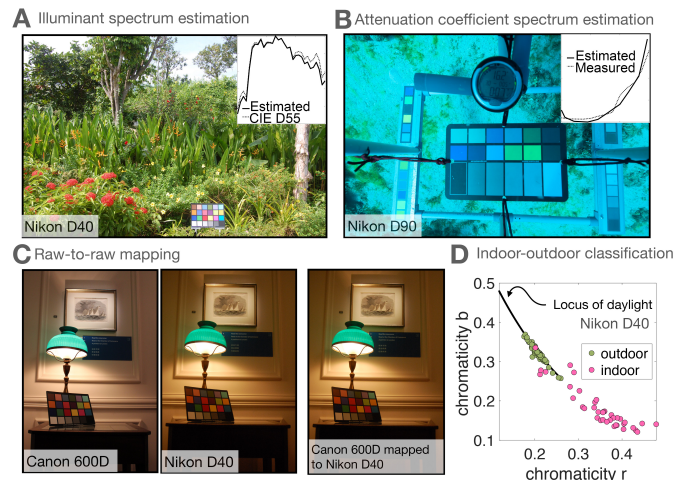


Fig. 1. Examples of tasks that can become possible or easier with known camera response. **A)** The spectrum of a uniformly lit outdoor scene can be recovered using a color chart. Image from [13]. **B)** Underwater, when the imaging distance is small, the diffuse downwelling attenuation coefficient can be recovered using a color chart. Data from [15]. **C)** When mapping one camera’s response to another, the tedious pairwise imaging and alignment [13] or the costly supervised learning steps [16] can be entirely omitted. Data from [17]. **D)** The locus of a camera’s daylight chromaticities can be used to classify whether the image was taken indoors or outdoors. Images from [17]. These results were obtained by implementing previously published methods in a way that incorporates the spectral sensitivities of the relevant camera. We did not have physical access to any of the cameras used; see Supplementary Materials for details.

• G. Solomatov and D. Akkaynak are with the Hatter Department of Marine Technologies, Leon H. Charney School of Marine Sciences, University of Haifa, and the Interuniversity Institute for Marine Sciences in Eilat
E-mail: gsolomat@campus.haifa.ac.il

property estimation [20–22], spectral super-resolution [23, 24], etc.) and/or scene illumination (e.g., for computational color constancy [17, 25, 26], image relighting [27], data augmentation [28–30] etc.) could be achieved simpler and faster, and

with more consistent results across cameras. Figure 1 gives some examples of tasks that could become possible or easier if camera spectral sensitivities were known.

Unfortunately, however, manufacturers of consumer cameras rarely make spectral sensitivity curves available, and the burden falls on the researchers to derive them. Their derivation (detailed in Section 2) requires the stimulation of the camera’s sensor with monochromatic light across the entire wavelength range to which it is sensitive. Such light can be obtained through the use of a monochromator or a spectrally-tuneable light source, but these optical instruments are generally too specialized and expensive to be found in most research labs. A seemingly lower cost and complexity alternative is to use photos of calibrated color targets, but with this indirect option comes a hidden cost: the implicit requirement that the illumination under which the color target is imaged is *also* known. In practice, this either means that the light source is calibrated, or its spectrum is measured with a spectrophotometer; necessitating, again, that the user has access to non-standard hardware (Table 1).

Even when the light spectrum is known, photos of color charts do not guarantee a “good enough” estimate of curves because there are too many unknowns in camera spectral sensitivities and not enough color chart patches whose reflectances are linearly independent [31–33]. The ill-posed nature of the problem introduces new challenges due to the need to constrain the solutions to a “realistic” space (Table 1). Furthermore, the commonly used image formation model does not take into account the optoelectronics of the camera [34], making it difficult to relate the captured RGB values to the predicted ones. Finally, even if not limited by hardware or cost, it is not uncommon that researchers work with imagery captured by cameras that they actually do not have in their possession. Consequently, the missing knowledge regarding the spectral responses of consumer cameras hampers the otherwise fruitful utilization of a consumer electronic device in academic research.

Our own research in underwater computer vision has been held back because we did not have the spectral sensitivities of the ever-evolving set of consumer cameras that produce the imagery we work with, and this motivated us to search for a solution that did not require: 1) any specialized hardware, 2) a color chart photo taken by the camera, or 3) physical access to the camera itself. Unlike how it may seem at first glance, this is not an impossible task—one simply has to be aware of the wealth of publicly available data regarding colorimetric mapping of consumer digital cameras. Here, we describe a method that satisfies these requirements with reconstruction errors that are, on average, as low as those that can arise due to manufacturing imperfections between two copies of the same camera.

Our key observation is the following: Among the tags specific to Adobe’s Digital Negative (DNG) format [35], are color transformation matrices prefixed *ColorMatrix*. These 3x3 matrices map camera-specific colors to the standard CIE XYZ color space under a known illuminant. Calibrated by the manufacturer, Adobe, or third-party companies and appended to image metadata by the Adobe DNG Converter, each DNG file includes at least one *ColorMatrix* tag, whereas

frequently, two matrices per camera can be found¹. In the common case when two matrices are available per camera, they are generally calibrated under illuminant A (tungsten, CCT≈2856K) and D65 (daylight, CCT≈6504K). We use these matrices to construct a linear system of equations which contains the sought spectral sensitivities in its solution space; and this solution space can be further constrained using an autoencoder trained on database of ground-truth spectral sensitivities from the literature. Using this novel approach we make the following contributions:

- 1) a zero-cost method with which to obtain reasonable estimates of camera spectral sensitivities, and
- 2) predicted spectral sensitivities of more than 1,000 makes and models of cameras that are currently in the market.

In addition, we provide a thorough survey of indirect spectral sensitivity estimation methods that have evolved over the last three decades, corresponding with the advent of consumer cameras as prevalent research tools.

2 Background and related work

The gold standard of spectral sensitivity estimation is the use of a monochromator setup in an optically sealed environment [34, 36–38]. A monochromator generates light in narrowband segments, which is then shone directly onto the sensor being calibrated or photographed from a calibration tile with uniform reflectance. At the same time as the light shines on the sensor, its spectrum is measured using a spectrophotometer, whose signal is then used as the reference for recovering the camera’s response from its captured RGB values. Such a setup is described in detail in [36].

While a monochromator setup provides a *direct* way of measuring spectral sensitivity curves and yields results with high accuracy, it consists of specialized hardware that can be prohibitively expensive, or simply inaccessible, for many research labs. Thus, over the last 30 years since consumer digital cameras became commonplace tools in research labs, numerous *indirect* methods, which aim to estimate a camera’s spectral response at a lower cost or complexity than the monochromator method, have been proposed (Table 1).

Despite the apparent diversity in techniques for indirect spectral sensitivity reconstruction, their core strategy rarely deviates from the same underlying principles, and our method is no different. At the outset, the sought sensitivity functions are described as a solution to a system of equations, typically linear or quadratic, which generally turns out to be underdetermined in the sense of admitting infinitely many solutions. Additional constraints such as non-negativity, smoothness and unimodality that are then assumed to hold for a broader family of spectral sensitivities are consequently employed with the aim of discarding spurious and “unrealistic” candidate solutions. Finally, the approximate sensitivity functions are obtained by solving an optimization problem whose objective function is comprised of two terms: one that originates from the system of equations and is specific to the camera under consideration, and one that corresponds to additional assumptions made about the space of functions that presumably

1. Since the release of DNG version 1.6.0.0, a third color matrix is also specified for some cameras [35].

TABLE 1

A summary of indirect spectral sensitivity estimation methods developed for digital cameras in the last 30 years and their requirements.

Reference	Approach	Specialized hardware required?	Camera required?	Chart required?	Light known?
Hubel et al. (1994)[39]	Linear regression, rank deficient pseudoinverse, Wiener method	Tungsten light, color filters, spectrophotometer	Yes	Yes	Yes
Hardeberg et al. (1998)[40]	Pseudoinverse, rank deficient pseudoinverse	No	Yes	Yes	Yes
Finlayson et al. (1998)[41]	Constrained linear regression	No	Yes	Yes	Yes
Barnard & Funt (1999)[42]	Constrained linear regression	Multiple illuminants, color filters, spectroradiometer	Yes	Yes	Yes
Dyas (2000)[43]	SVD + regularization	No	Yes	Yes	Yes
Carvalho et al. (2001)[44]	Learning with generalized cross validation function	No	Yes	Yes	Yes
Barnard & Funt (2002)[25]	Constrained linear regression	Multiple illuminants, color filters, spectroradiometer	Yes	Yes	Yes
Alsam & Finlayson (2002)[31]	Feasible set	No	Yes	Yes	Yes
Quan et al. (2003)[45]	Iterative multiscale bases	No	Yes	Yes	Yes
DiCarlo et al. (2004)[46]	Emissive target	LEDs, opto-electronics	Yes	Emissive	Yes
Ebner (2007)[47]	Evolution strategy	Spectrometer	Yes	Yes	Yes
Mauer & Wueller (2009)[48]	Narrowband interference filters	Filters and mount	Yes	No	Yes
Zhao et al.(2009)[49]	Optimum basis functions	Spectrometer	Yes	Yes	Yes
Rump et al. (2011)[50]	Color specularity & light nonuniformity + optimization	Spectrometer	Yes	Yes	Yes
Pike (2011)[51]	Parameterization of [41] + UV	Spectrometer, UV target	Yes	Yes	Yes
Bedarkar (2011)[52]	Optimization	Blackbody furnace	Yes	No	Planck's law
Han et al. (2012)[53]	Fourier bases	Fluorescent chart	Yes	Yes	No
Jiang et al. (2013)[54]	PCA	None	Yes	Yes	Yes/daylight
Prasad et al. (2013)[55]	Optimization	Multiple illuminants	Yes	Yes	No
Kawakami et al. (2013)[56]	Sky turbidity, need clear sky & camera-sun direction + PCA	No	Yes	No	No
Bongiorno et al. (2013)[57]	Linear variable edge filter (LVEF)	LVEF, smooth light, spectrometer, Spectralon	Yes	No	Yes
Huynh & Robles-Kelly (2014)[58]	Optimization	No	Yes	Yes	No
Macdonald (2015) [59]	Emissive target + transmission filters	Spectrometer, emissive target, color filters	Yes	Emissive	Yes
Bartczak et al. (2015) [60]	Emissive target	Spectrometer, spectrally-tuneable light, diffuser	Yes	Emissive	Yes
Finlayson et al. (2016)[61]	Rank-based + sine bases	No	Yes	Yes	Yes
Manakov (2016)[62]	Band-pass filters	Color filters, integrating sphere, spectrometer	Yes	No	Yes
Zhang et al. (2017)[63]	Regularized least squares	Hyperspectral sensor	Yes	No	No
Chaji et al. (2018)[64]	Neural network	No	Yes	Yes	Daylight
Karge (2018)[65]	Diffraction image + least squares	Grating, light, radiometer	Yes	Yes	Yes
Burggraaff et al. (2019)[1]	Pocket spectrometer + solar irradiance model	Spectrometer, halogen lamp	Yes	No	Yes
Zhu et al. (2020)[66]	Orthogonal test + optimization	Spectral irradiance colorimeter, LCD display	Yes	Yes	Yes
Toivonen & Klami (2020) [67]	Diffraction imaging + optimization	Diffuser, diffraction grating, color filters	Yes	Yes	Yes
Ji et al. (2021) [68]	Compressive sensing	Xenon-arc lamp	Yes	Yes	Yes
Tominaga et al. (2021)[69]	PCA & color patches	Incandescent light, Munsell white paper, Spectralon	Yes	Yes	Yes
Ma et al. (2021)[70]	Constrained least squares + sine bases	No	Yes	Yes	Yes
Xu et al. (2022)[71]	Rank-based	Multiple illuminants	Yes	Yes	Yes
Fan & Ronnier (2023)[72]	Emissive target + PCA	OLED display, radiometer	Yes	Emissive	Yes
Ours	Color matrices + autoencoder	No	No	No	No

contains the sought sensitivities. We will refer to these terms as the *specific term* and the *universal term*, respectively; the former depending on the data available from the experimental setup, and the latter depending on assumptions that may or may not hold true, since nobody really knows what spectral sensitivities of future cameras could look like (although we can probably expect them to become more colorimetric).

There is an inherent interplay between these two terms in the sense that the universal term compensates for the lack of information in the specific term. Indeed, using a monochromator yields nearly perfect information for the specific term, rendering the universal term unnecessary. Conversely, if no experimental data is available from the camera under consideration, then the best approximation one could hope for would be some kind of “averaging” of all known spectral sensitivities, which constitutes an extremely strong assumption on the space of candidate functions.² With this trade-off in mind, we set out to devise a method that achieves the following:

- 1) the experimental setup required by the specific term is minimal; and
- 2) the universal term is as unassuming as possible while ensuring 1).

We propose a method that attains 1) by constructing the specific term using publicly calibration matrices, which does not even require the user to own the camera in question, while 2) is achieved by training an autoencoder on a database of known spectral sensitivities.

2.1 The image formation model and its challenges

In this subsection, we discuss the physical image formation model that gives spectral sensitivity functions their meaning; the discretized version of this model is by far the most common choice for constructing the specific term in the optimization problem used for spectral sensitivity reconstruction.

For any wavelength λ , we denote the sensitivity of color channel $k \in \{\mathbf{r}, \mathbf{g}, \mathbf{b}\}$ in a device d to light of that wavelength by $S_d^{(k)}(\lambda)$. If a target object with reflectance $R(\lambda)$ is photographed by d under some illuminant $L(\lambda)$, then the resulting pixel intensity for channel k is given by

$$I_d^{(k)} = \int_0^\infty R(\lambda)L(\lambda)S_d^{(k)}(\lambda)d\lambda. \quad (1)$$

The sensitivity $S_d^{(k)}$ is typically zero outside of the visible spectrum 400–700nm, in which case the above integral only needs to be considered on this interval. Doing so makes it possible to discretize the above image formation model over n uniformly spaced values of λ as follows:

$$I_d^{(k)} \approx \sum_{i=1}^n R(\lambda_i)L(\lambda_i)S_d^{(k)}(\lambda_i)\Delta\lambda, \quad (2)$$

2. The particular kind of “averaging” would depend on the targeted error metric. We computed such an “average” with respect to the *relative full-scale error* (see Section 4), minimizing the average of this error across our dataset. This yielded 21% median error, which is much worse than the 8% achieved by our method. This clearly demonstrates that color matrices do contain a lot of useful information about spectral sensitivities.

where

$$\Delta\lambda = \frac{700 - 400}{n - 1} \text{nm} = \frac{300}{n - 1} \text{nm} \quad \text{and} \\ \lambda_i = 400 \text{nm} + (i - 1)\Delta\lambda \quad \text{for } i = 1, \dots, n.$$

Rewriting (2) in matrix form yields:

$$\mathbf{I}_d \approx \mathbf{R}\mathbf{L}\mathbf{S}_d, \quad (3)$$

where

$$\mathbf{I}_d = \begin{bmatrix} I_d^{(\mathbf{r})} & I_d^{(\mathbf{g})} & I_d^{(\mathbf{b})} \end{bmatrix} \in \mathbb{R}^{1 \times 3}, \\ \mathbf{R} = \begin{bmatrix} R(\lambda_1) & \dots & R(\lambda_n) \end{bmatrix} \in \mathbb{R}^{1 \times n}, \\ \mathbf{S}_d = \begin{bmatrix} S_d^{(\mathbf{r})}(\lambda_1) & S_d^{(\mathbf{g})}(\lambda_1) & S_d^{(\mathbf{b})}(\lambda_1) \\ \vdots & \vdots & \vdots \\ S_d^{(\mathbf{r})}(\lambda_n) & S_d^{(\mathbf{g})}(\lambda_n) & S_d^{(\mathbf{b})}(\lambda_n) \end{bmatrix} \in \mathbb{R}^{n \times 3} \quad \text{and} \\ \mathbf{L} = \begin{bmatrix} L(\lambda_1) & 0 & \dots & 0 \\ 0 & L(\lambda_2) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & L(\lambda_n) \end{bmatrix} \Delta\lambda.$$

There is of course no reason to restrict this setting to a single target object with a single reflectance function $R(\lambda)$. In the general case of m objects, we simply get $\mathbf{R} \in \mathbb{R}^{m \times n}$ and $\mathbf{I}_d \in \mathbb{R}^{m \times 3}$, i.e., one row per target object. Finally, even more equations can be obtained by taking multiple photographs under different illuminants.

In the context of spectral sensitivity reconstruction, this suggests an experimental setup where the camera under consideration is used to take photographs of objects with known reflectances under known light conditions. Indeed, if n linearly independent equations are obtained this way, then this setup is essentially equivalent to having a monochromator, as we can then just solve Eq. (3) for \mathbf{S}_d . Unfortunately, this conceptually simple approach suffers from a few practical issues:

- 1) The pixel intensities \mathbf{I}_d in the photographs will always be subject to noise due to the idealized nature of the underlying image formation model in Eq. (1) that omits the optoelectronics of the camera [34, 40], not to mention the discretization error introduced in Eq. (2).
 - 2) The illuminant \mathbf{L} under which the pixel intensities \mathbf{I}_d are obtained is typically not known; and controlling \mathbf{L} often requires specialized hardware (see Table 1).
 - 3) Knowing the reflectances \mathbf{R} requires a special object like a calibrated color chart in the photograph(s).
 - 4) The linear system does not admit a unique solution as long as $m < n$. Although it is in principle possible to obtain $m = n$ target objects with known linearly independent reflectances, no cheap and convenient options seem to be available on the market today. The common color charts used in photography, for example—even those containing a large number of different colors—are printed using just a handful of unique pigments, which essentially limits the possible number of linearly independent reflectances to the number of pigments used.³
3. The surface on which the colors are printed could contribute with one reflectance as well.

In our spectral sensitivity reconstruction method, we rely on an alternative, previously overlooked linear system (see Section 3.1) which sidesteps all but the last of these challenges. In order to tame the multitude of candidate solutions arising from the fourth challenge, we employ an autoencoder trained on a database of known spectral sensitivities, which is used to measure how “realistic” any given candidate solution is.

3 Proposed method

In this section we give a detailed description of our method for indirect spectral sensitivity reconstruction. We begin with a brief introduction to color matrices as well as the linear system of equations that can be obtained from them.

3.1 Color matrices as a source for equations

In 2004, Adobe introduced the Digital Negative (DNG), an open-source image file format for storing RAW image information from images created with different cameras [35]. Each DNG file includes at least one *ColorMatrix* tag, whereas frequently, two matrices per camera can be found. These matrices map camera colors to the CIE Standard Observer, and are calibrated under illuminants with sufficiently far apart correlated color temperatures (CCT), such as illuminant A (tungsten, CCT≈2856K) and D65 (daylight, CCT≈6504K).

For any device d and illuminant $\mathbf{L} \in \mathbb{R}^{n \times n}$, one defines the color matrix $C_d(\mathbf{L}) \in \mathbb{R}^{3 \times 3}$ as the least-squares solution to the matrix equation

$$\mathbf{L}\mathbf{S}_{\text{XYZ}}C_d(\mathbf{L}) = \mathbf{L}\mathbf{S}_d,$$

where XYZ refers to the CIE XYZ standard observer. Letting $\mathbf{A}^+ = (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \in \mathbb{R}^{t \times s}$ denote the pseudoinverse of any full-rank matrix $\mathbf{A} \in \mathbb{R}^{s \times t}$, we can write

$$C_d(\mathbf{L}) = (\mathbf{L}\mathbf{S}_{\text{XYZ}})^+ \mathbf{L}\mathbf{S}_d. \quad (4)$$

For multiple illuminants $\mathbf{L}_1, \dots, \mathbf{L}_\ell$, we can extend (4) to

$$\begin{bmatrix} C_d(\mathbf{L}_1) \\ \vdots \\ C_d(\mathbf{L}_\ell) \end{bmatrix} = \begin{bmatrix} (\mathbf{L}_1\mathbf{S}_{\text{XYZ}})^+ \mathbf{L}_1 \\ \vdots \\ (\mathbf{L}_\ell\mathbf{S}_{\text{XYZ}})^+ \mathbf{L}_\ell \end{bmatrix} \mathbf{S}_d, \quad (5)$$

which will for us serve as the substitute for (3). Notice that (5) yields 9 equations for each known calibration matrix, amounting to 9ℓ equations in total. In practice, calibration matrices are made available only for illuminants A and D65, i.e., $\ell = 2$. It is worth noting that the equations in Eq. (5) are far more practical than those in Eq. (3). This is because the former can be constructed entirely from publicly available data, while the latter requires one to obtain photographs of objects with known reflectances under a controlled or estimated illuminant.

Adobe matrices can be acquired as follows: For a given RAW image, the available color calibration matrices will be appended to the image metadata when the image is processed through the Adobe DNG Converter. These matrices can then be obtained via a viewer, e.g., *ExifTool* [73]. This option still does not necessitate having physical access to the camera in question, as example RAW images from many cameras can be freely found online (see [74] for an excellent source). We obtained the color matrices for 1,013 cameras for which the Adobe DNG Converter version 12.4.0.555 supports, using the code provided by [75]. Some statistics for these cameras are shown in Fig. 2 for manufacturers who had ≥ 5 cameras in the dataset.

3.2 Optimization problem

To begin with, assume that we describe the sought spectral sensitivities $\mathbf{S}_d \in \mathbb{R}^{n \times 3}$ as a solution to a linear system $\mathbf{A}\mathbf{S}_d = \mathbf{B}$ for some known matrices $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{B} \in \mathbb{R}^{m \times 3}$. For example, this linear system could be (3), or (5), or a combination of the two. Independently of which linear system it is, we can assume that is underdetermined, i.e., $m < n$; otherwise, we can just solve the system and there is nothing more to do. To address the insufficient number of equations at our disposal, we follow the usual practice of inventing a function F that maps any potential spectral sensitivity candidate $\mathbf{S} \in \mathbb{R}^{n \times 3}$ to a single real number that is supposed to indicate how “realistic” this candidate is. For example, if we assume that spectral sensitivities are generally rather smooth, then we could choose $F(\mathbf{S}) = \|\mathbf{D}\mathbf{S}\|_2$, where $\mathbf{D} \in \mathbb{R}^{n \times n}$ is the second derivative matrix and $\|\cdot\|_2$ is the Frobenius norm; variants of this were explored in e.g., [25] and [58].

Being equipped with the aforementioned system of equations as well as the function F , we combine the two in the formulation of the final optimization problem

$$\operatorname{argmin}_{\mathbf{S} \in \Omega} \underbrace{D(\mathbf{A}\mathbf{S}, \mathbf{B})}_{\text{specific term}} + \underbrace{F(\mathbf{S})}_{\text{universal term}}, \quad (6)$$

where $\Omega \subseteq \mathbb{R}^{n \times 3}$ is the feasible domain and D is some distance function on $\mathbb{R}^{m \times 3}$. In a sense, the whole field of indirect spectral sensitivity reconstruction can be viewed as the search for the right combination of the two terms in Eq. (6), hoping that an optimization algorithm would converge to a solution not too far off from the sought \mathbf{S}_d . Before proceeding to our take on this framework, we ought to mention that the system of equations in the specific term does not have to be linear; an example of a quadratic system can be found in [41]. Furthermore, the two terms can of course be summed in a weighted manner—we have omitted the weight coefficients for a cleaner presentation.

Having presented the general optimization framework for spectral sensitivity reconstruction, let us introduce our proposed realization of it. For the specific term, we rely on the linear system in Eq. (5) based on color matrices, choosing

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix} = \begin{bmatrix} (\mathbf{L}_A\mathbf{S}_{\text{XYZ}})^+ \mathbf{L}_A \\ (\mathbf{L}_{D65}\mathbf{S}_{\text{XYZ}})^+ \mathbf{L}_{D65} \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \end{bmatrix} = \begin{bmatrix} C_d(\mathbf{L}_A) \\ C_d(\mathbf{L}_{D65}) \end{bmatrix},$$

where \mathbf{L}_A and \mathbf{L}_{D65} denote the illuminants A and D65 respectively (see Section 3.1). For the feasible domain and the distance function we respectively use

$$\Omega = \{\mathbf{S} \mid \mathbf{S} \neq \mathbf{0} \text{ has no negative entries}\} \quad \text{and}$$

$$D(\mathbf{A}\mathbf{S}, \mathbf{B}) = \sum_{i=1}^2 \angle(\mathbf{A}_i\mathbf{S}, \mathbf{B}_i), \quad \text{where}$$

$$\angle(\mathbf{U}, \mathbf{V}) \triangleq \arccos\left(\frac{\mathbf{U} \cdot \mathbf{V}}{\|\mathbf{U}\|_2 \|\mathbf{V}\|_2}\right),$$

and $\mathbf{U} \cdot \mathbf{V}$ denotes the dot product of \mathbf{U} and \mathbf{V} seen as vectors. Finally, our universal term takes the form

$$F(\mathbf{S}) = \angle(\mathbf{S}, A_{\mathbf{w}}(\mathbf{S})), \quad (7)$$

where $A_{\mathbf{w}}$ is an autoencoder with weights \mathbf{w} trained on a database $\mathcal{S} \subset \mathbb{R}^{n \times 3}$ of known spectral sensitivities. Details about the chosen architecture of A , as well as the training process, can be found in Section 3.3.




Fig. 2. We extracted color matrices for 1,013 digital cameras from the Adobe DNG Converter and manually compiled metadata for each camera including camera type (e.g., DSLR, compact, mirrorless, cell phone), sensor type (CMOS or CCD), sensor size (in mm), and release year. Manufacturers with fewer than 5 cameras were omitted. See Supplementary Materials for details.

3.3 Autoencoder

Letting $\mathcal{S} \subset \mathbb{R}^{n \times 3}$ denote a database of ground-truth spectral sensitivities, the weights \mathbf{w} of our autoencoder $A_{\mathbf{w}}$ are obtained as a solution to the optimization problem

$$\operatorname{argmin}_{\mathbf{w}} \frac{1}{|\mathcal{S}|} \sum_{\mathbf{S} \in \mathcal{S}} \Delta(\mathbf{S}, A_{\mathbf{w}}(\mathbf{S})), \quad (8)$$

where for any $\mathbf{U} = [\mathbf{U}_r, \mathbf{U}_g, \mathbf{U}_b]$ and $\mathbf{V} = [\mathbf{V}_r, \mathbf{V}_g, \mathbf{V}_b]$ in $\mathbb{R}^{n \times 3}$

$$\Delta(\mathbf{U}, \mathbf{V}) = \left\| \left[\frac{\mathbf{U}_k - \mathbf{V}_k}{\|\mathbf{U}_k\|_2} \right]_{k=r,g,b} \right\|_2.$$

The choice for Δ might come across as somewhat peculiar, but it simply turns out to work better than $\angle(\cdot, \cdot)$, which would be more consistent with our choice of the universal term in Eq. (7). Another choice could be $\Delta(\mathbf{U}, \mathbf{V}) = \|\mathbf{U} - \mathbf{V}\|_2$, which is essentially what we are doing, except that the scaling by $\|\mathbf{U}_k\|_2^{-1}$ is introduced to make $A_{\mathbf{w}}$ assign roughly equal importance to the three color channels.

The description thus far constitutes the conceptual core of our training process, but we have omitted two important details, both related to the small size of our database \mathcal{S} (see Table 4). First and foremost, it has shown to be absolutely crucial for our method to use regularization in order to avoid dramatic overfitting; in particular, we use weight decay and dropout (for parameter choices see Table 3). Secondly, the stability of our method seems to improve with addition of modest data augmentation. During training, every time we draw a ground-truth sensitivity \mathbf{S} from \mathcal{S} , we randomly scale and “roll” its channels by applying the map $\mathbf{S} \mapsto \mathbf{GSH}$ where $\mathbf{H} \in \mathbb{R}^{3 \times 3}$ is a diagonal matrix whose diagonal entries are drawn uniformly and independently from the interval $[h, 1]$ for a real-valued parameter $0 < h \leq 1$. The matrix $\mathbf{G} \in \mathbb{R}^{n \times n}$ is a special kind of random permutation matrix whose i -th column is 0 in every entry except for the index

$$(i - 1 + u \bmod n) + 1,$$

where the value is 1; here u is a random integer drawn uniformly and independently for every column of \mathbf{G} from the set $\{-g, \dots, g\}$, where $0 \leq g < n$ is an integer-valued parameter. The two parameters h and g are chosen by the user and decide the level of data augmentation. Intuitively, h decides by how much we can randomly scale the channel peaks up and down, while g decides by how much we can shift them left and right. Our choice for these parameters can be

found in Table 3, while the neural network architecture of the autoencoder $A_{\mathbf{w}}$ is presented in Table 2. Finally, we ought to mention that all ground-truth sensitivities $\mathbf{S} \in \mathcal{S}$ were normalized by their maximal value (before augmentation).

TABLE 2

Architecture of the autoencoder $A_{\mathbf{w}}$. Information travels from top to bottom. In the table $N = 3n$, and p indicates that dropout with the specified retention probability was applied after the corresponding layer during training.

Layer type	Inputs	Outputs	Activation	
Linear ($p = 0.2$)	N	$4N$	ReLU	} Encoder
Linear ($p = 0.5$)	$4N$	$2N$	ReLU	
Linear	$2N$	6	None	
Linear	6	$2N$	ReLU	} Decoder
Linear ($p = 0.5$)	$2N$	$4N$	ReLU	
Linear	$4N$	N	None	

TABLE 3

Training details for weights \mathbf{w} in (8). Explanations for channel scaling and rolling (data augmentation) can be found in Section 3.3.

Parameter	Value
Framework	PyTorch v1.11.0
Optimizer	SGD
Channel scaling (h)	$2 \cdot 10^{-1}$
Channel rolling (g)	2
Learning rate (lr)	10^{-1}
Momentum	$5 \cdot 10^{-1}$
Weight decay	10^{-4}
Scheduler	ReduceLRonPlateau
Scheduler decay	$5 \cdot 10^{-1}$
Scheduler patience	$2 \cdot 10^3$
Stop criterion	$lr < 10^{-5}$

3.4 Ground-truth database

The database \mathcal{S} consists of 51 spectral sensitivity functions that we compiled from the literature, which were reported to have been obtained directly using monochromatic light (Table 4). If the authors did not make their data available in a table, we digitized them from plots using DataThief [76].

We excluded some cameras from the sources listed. Specifically, we excluded all cameras for which Adobe DNG Converter version 12.4.0.555 did not provide a *ColorMatrix* tag. In some cases, the presented curves were scaled by the authors so as to make the maximal value in each channel equal to

one—these were excluded from our dataset due to the across-channel nature of our approach. If a camera’s curves showed sensitivity to a wider range than the VIS (i.e., UV or IR filters were removed), we omitted that camera. Finally, we also excluded discontinued cameras with atypical curves, such as the Kodak DCS series (a poignant outcome given their pioneering role in the history of digital photography).

Among the resulting set of 51, we had 9 cameras whose sensitivities were derived by more than one source (Fig. 3). The relative full-scale errors between these duplicates range between $\approx 2.4 - 17\%$, which are likely due to measurement errors and uncertainties; this can be quite large for some wavelengths, as demonstrated by Darrodi et al. [36]. It is also possible, however, that two copies of the same make and model camera have quantifiably different spectral sensitivities due to the imperfections of the manufacturing process. This indeed seems to be the case for two Nikon Coolpix 5700 cameras that were reportedly purchased simultaneously and characterized using the same experimental setup by Stevens et al. [77] (note that Adobe does not provide matrices for this camera but we retained it in our database).

TABLE 4

Our compiled ground-truth spectral sensitivity database of 51 cameras. All original sources report estimation using a monochromator, except for those that used a liquid crystal modulator⁺, diffraction grating^{*}, and narrowband filters[†] to obtain monochromatic light.

Camera make & model	Source
Canon 1Ds II	Akkaynak et al. [78]
Panasonic DMCL X5, Sony NEX7	Berra et al. [79]
Canon 1Ds II, Nikon D3	Brady & Legge [80]
Samsung Galaxy S8	Burggraaff et al. [1]
Canon 10D, Nikon D70	Huynh & Robles-Kelley [81] ⁺
Canon 1D III, Canon 20D, Canon 300D, Canon 40D, Canon 500D, Canon 50D, Canon 5D II, Canon 600D, Canon 60D, Nikon D200, Nikon D3, Nikon D300s, Nikon D3X, Nikon D40, Nikon D50, Nikon D5100, Nikon D700, Nikon D80, Nikon D90, Olympus EPL2, Pentax K5, Pentax Q, Sony NEX5N	Jiang et al. [54]
Canon 10D, Canon 5D II, Canon 5D, Nikon D1x, Nikon D70	Kawakami et al. [56]
Canon 400D	Lebourgeois et al. [82] [*]
Canon 40D, Leica M8, Nikon D200, Nikon D700, Panasonic DMCL-X3	Mauer & Wueller [48] [†]
Nikon 5700 (2 copies)	Stevens et al. [77]
iPhone 8, iPhone X, iPhone 11, iPhone 12 Pro Max, Samsung Galaxy S7 Edge, Samsung Galaxy S9, Samsung Galaxy S20	Tominaga et al. [69]

4 Results

When it comes to indirect methods for spectral sensitivity reconstruction, it is generally difficult—if not impossible—to formulate any concrete theoretical guarantees about reconstruction accuracy. The main reason for this is that all contemporary models for the space of spectral sensitivity functions substitute the lacking theoretical understanding of

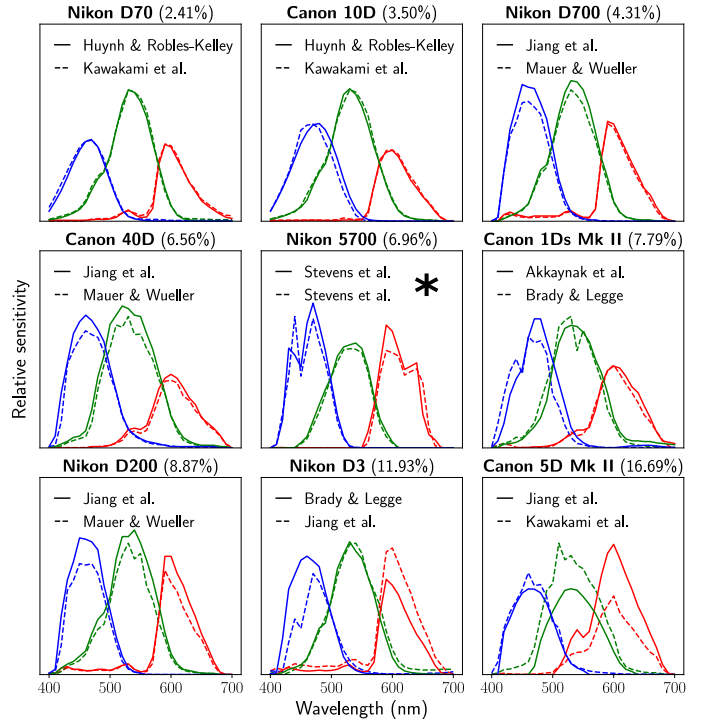


Fig. 3. Cameras in our database for which two ground-truth derivations exist. Percentages indicate relative full-scale error between estimations reported in two different publications for each camera. These differences probably arise from measurement errors and uncertainties associated with the experimental setup employed by each party. Note, however, that the case for the two Nikon 5700 Coolpix cameras (*) is slightly different. These cameras were reported to be purchased at the same time and their spectral sensitivity measurements were made with the same experimental setup [77]. Thus, potentially, the 6.96% error between their spectral sensitivities primarily represents the error resulting from the imperfections of the manufacturing process and/or the variability of the sensor materials.

this space with convenient assumptions that may or may not hold in practice. In the case of our method, it is assumed that this space can be learned by an autoencoder with a particular neural network architecture, which doesn’t make theoretical analysis any easier. In acknowledgment of this, we proceed with the standard approach of assessing the effectiveness of our method by empirical validation.

Our validation procedure is a form of leave-one-out validation and can be summarised as follows: For every device d represented in our database $\mathcal{S} \subset \mathbb{R}^{n \times 3}$, we train an autoencoder $A^{(d)}$ on $\mathcal{S} \setminus \{S_d\}$. We then proceed by solving the optimization problem discussed in Section 3.2, i.e.,

$$\operatorname{argmin}_{\mathbf{S}} \alpha \sum_{i=1}^2 \mathcal{L}(\mathbf{A}_i \mathbf{S}, \mathbf{B}_i) + \beta \mathcal{L}(\mathbf{S}, A^{(d)}(\mathbf{S})), \text{ where}$$

$$\begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix} = \begin{bmatrix} (\mathbf{L}_A \mathbf{S}_{XYZ})^+ \mathbf{L}_A \\ (\mathbf{L}_{D65} \mathbf{S}_{XYZ})^+ \mathbf{L}_{D65} \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \end{bmatrix} = \begin{bmatrix} C_d(\mathbf{L}_A) \\ C_d(\mathbf{L}_{D65}) \end{bmatrix},$$

and comparing the found solution $\hat{\mathbf{S}}_d$ to the ground-truth \mathbf{S}_d . Here, $\alpha, \beta \geq 0$ are just weight coefficients; their values, as well as other optimization parameters can be found in Table 5. The comparison is done similarly to [66] with respect to the average *relative full-scale error* across the three color channels,

i.e.,

$$\text{RE}(\widehat{\mathbf{S}}_d, \mathbf{S}_d) = \frac{1}{3} \sum_{k \in \{r, g, b\}} \frac{\text{RMSE}(\widehat{\mathbf{S}}_d^{[*], k}, \mathbf{S}_d^{[*], k})}{\max \mathbf{S}_d^{[*], k}}, \quad \text{where}$$

$$\text{RMSE}(\widehat{\mathbf{S}}_d^{[*], k}, \mathbf{S}_d^{[*], k}) = \sqrt{\frac{1}{n} \sum_{i=1}^n (\widehat{\mathbf{S}}_d^{[i], k} - \mathbf{S}_d^{[i], k})^2}$$

is the root mean square error; here the superscript $[i, k]$ means taking the entry in the i -th row and the k -th column, while $[*], k$ refers to the k -th column in its entirety. As mentioned, our database \mathcal{S} contains a few duplicate “ground-truth” sensitivities for the same camera. Whenever this is the case, we leave out *all* of these duplicates when training the autoencoder during the validation (but errors are still included). The resulting error histogram can be found in Fig. 4 and our best, median, and worst results are shown in Fig. 5. We also performed validation where we left out one brand at a time—the results were similar to those in Fig. 4 (e.g. 7.08% median error across all Canon cameras). The rest of the results are in the Supplementary Material.

It is difficult to compare this error range to other methods, as there is no consensus in the literature about which error metric should be used, and many of the common error metrics are scale-dependent. One value with which we can compare is the reported 8.54% average relative full-scale error across channels from Zhu et al. [66] for Canon EOS 600D; this is slightly worse than our error of 6.87% with respect to the ground-truth from Jiang et al. [54]. The method of Zhu et al. [66] required a camera and specialized hardware (Table 1).

We wondered if reconstruction accuracy would improve if we also used a color chart photo for a given camera. We used the outdoor images from the NUS dataset [17] all of which contained a ColorChecker, and included the equations from the image formation model (3) into the specific term of our optimization problem. We approximated the light by the CIE standard daylight model, parameterized by CCT, which we treated as an additional unknown to be solved for in our optimization problem. Our reconstruction accuracy did not improve by any significant amount ($< 0.5\%$). Most likely, this was because the color chart offered little new information that was not captured by the color matrices.




Fig. 4. **A)** Histogram of relative full-scale errors observed during leave-one-out validation of our method. The two ground-truths of Nikon 5700 from Stevens et al. [77] were excluded from the predictions as well as from training data due to missing color matrices. **B)** Breakdown of errors per color channel—the median errors are 9.3%, 6.3% and 8.3% for red, green and blue channels, respectively.

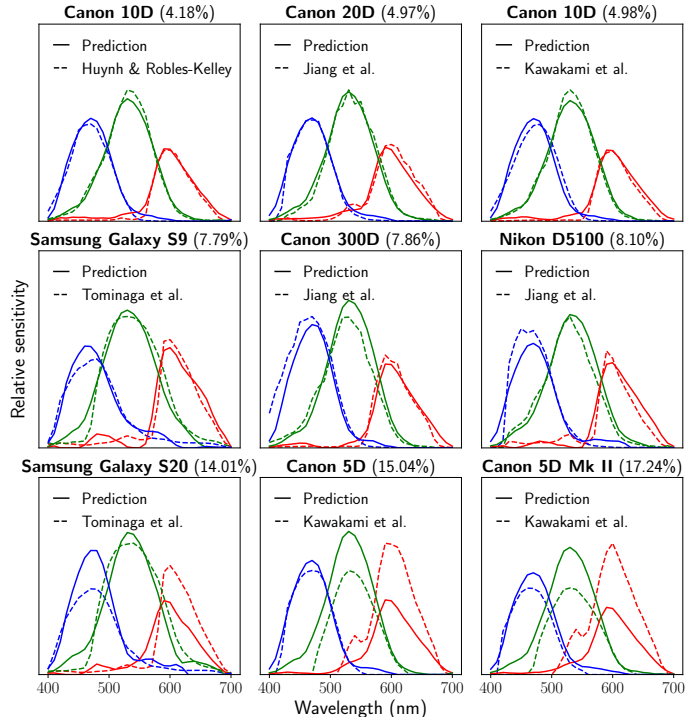


Fig. 5. Our best (first row), median (middle row), and worst (bottom row) results for the 51 cameras in our ground-truth set. The rest of the results for the ground-truth dataset, as well as predictions for 1,000+ cameras are given in Supplementary Material.

TABLE 5
Optimization parameters used to produce Fig. 4.

Parameter	Value
Framework	PyTorch v1.11.0
Optimizer	SGD
Specific term coeff. (α)	10^2
Universal term coeff. (β)	$2 \cdot 10^{-1}$
Learning rate (lr)	$8 \cdot 10^{-4}$
Scheduler	ReduceLROnPlateau
Scheduler decay	$5 \cdot 10^{-1}$
Scheduler patience	$2 \cdot 10^3$
Stop criterion	$lr < 4 \cdot 10^{-4}$

5 Discussion

We introduced a zero-cost framework with which to reasonably estimate the spectral sensitivities of consumer cameras. Our reconstruction errors ranged from 4 – 17% with a median of $\approx 8\%$ (Fig. 4), which is comparable to the errors that might arise due to sensor fabrication flaws or irregularities when two copies of the same are being manufactured (see Nikon 7500 data in Fig. 3). Since our framework relies on publicly available information, our current results can be improved as more, and higher-quality data become available, for instance, when more ground-truth curves are published by researchers, or color matrices for more illuminants are derived by Adobe or others. We also provided predictions of camera responses for 1,000+ cameras for the use of the scientific community.

A lot of value can be added to scientific research with the knowledge of the spectral sensitivities of the camera at hand. This information can enable objectivity and repeatability of color capture, better colorimetric mapping, and create the

possibility of linking RGB values to spectral quantities in the physical world. We demonstrated some of these cases in Fig. 1 where we modified the implementations of previously published methods to accommodate the spectral sensitivities of the cameras used. Moreover, having the spectral sensitivities reduces the number of unknowns in a given problem by $3n$ (where n represents the number of chosen wavelength steps), thereby simplifying solutions in numerous scenarios.

For example, **illuminant spectrum estimation** can become a simpler problem with a more accurate solution that recovers the *spectrum* of the physical illuminant, not just the camera-specific RGB white point (Fig. 1A). Additionally, as suggested by Cheng et al. [17], algorithms that use specular and shadow pixels to estimate the illuminant through distinguishing between specularities and bright surfaces, and shades and dull surfaces, can be refined.

Similarly, **reflectance spectrum estimation** also becomes simpler, especially outdoors where light can be represented using the one-parameter CIE daylight model. Underwater, image formation is complex, but in the cases where haze can be neglected and Eq. (1) applies, the spectrum of the **diffuse downwelling attenuation coefficient** can be recovered using a color chart and the CIE daylight model (Fig. 1B). It should be kept in mind, however, that metamerism and the low-dimensionality of spectral reflectances on most color charts still make the estimation of spectral quantities challenging (see Supplementary Material).

Having access to the spectral sensitivities of *many* cameras, not all of which may be at hand, can also add a lot of value to science. For example, realistic training and/or augmentation data can be generated for **image colorization**, **jpg-to-raw** [83], and **raw-to-raw** [13] mapping problems. **Spectral super-resolution**, a task which obtains hyperspectral images from a single RGB image and generally requires camera curves to be known [84], will be possible to apply to any consumer camera image. Similarly, many cameras will suddenly become candidates for **hyper-spectral super-resolution**, the task of fusing a hyperspectral image with low spatial resolution with an RGB image of high spatial resolution, which also requires camera responses to be known [56, 85, 86]. It will also be trivial to follow the method of Oh et al. [23] to build a Do-it-Yourself hyperspectral imaging system by combining several different consumer cameras, since the dataset we release contains the peak wavelengths of (almost) all the cameras in today’s market.

Perhaps the most curious application of our predicted camera responses is **time travel rephotography**, which is the projection of faded, black-and-white photos into the space of modern-day high resolution sensors [87]. It will be possible to re-make iconic photographs, like that of Abraham Lincoln, in the color space of any modern-day digital camera, simulating that camera’s in-camera photofinishing workflow using the pipeline from [88], and even comparing how each camera in the market would have portrayed the notable persona.

While we are excited about all the capabilities our work will enable, we acknowledge that our method has limitations. First of all, we noticed large differences in the reconstruction errors on the same camera due to the *source* of derivation (Fig. 3), which means that “ground-truth” information should be interpreted with caution. For example, our reconstruction for Canon 5D Mk II yielded only 6.20% error relative to

the ground-truth from Jiang et al. [54], but 17.24% error relative to ground-truth from Kawakami et al. [56], which is actually our worst result (see Fig. 6). Does the sensor in this Canon model show such a large variation? Or was Kawakami’s derivation flawed? We cannot answer, and we caution the reader to be aware of potential measurement errors and uncertainties when using publicly available data. Also, we assume, as does the Adobe DNG Converter, that every copy of a given make and model camera has identical spectral sensitivities. The case of Canon 5D Mk II could be one example where this assumption fails. However, we believe that in practice differences between sensors of identical cameras are small, and that the Nikon 5700 example in Fig. 3 is more typical of this situation, so our general approach is justified.

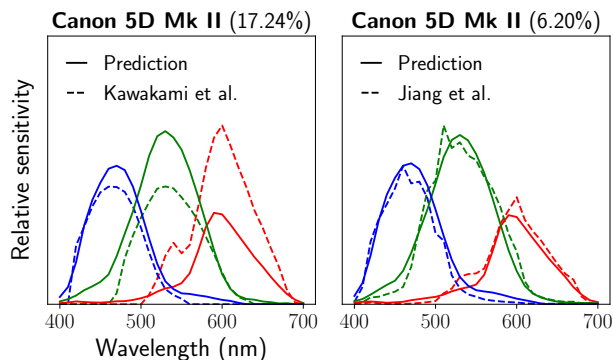


Fig. 6. Leave-one-out validation results for two different “ground-truths” of Canon 5D Mk II. The large discrepancy in reconstruction error could suggest and experimental error in Kawakami et al., but further investigation is needed for a definitive conclusion.

Lastly, we did not have ground-truth data for most of the 1,000+ cameras we presented, because if such data were available, we would not need to write this paper. Despite the lack of ground-truth, we have confidence in the peak wavelengths and the general shapes of the reconstructed curves due to the way Adobe’s color matrices are derived. Even without ground-truth, this information can already be useful, serving as a feasible initial guess for problems where camera responses are unknown, or potentially guiding the purchase of a camera if sensitivity for certain wavelengths are desired.

Acknowledgments

This research was funded by Israel Science Foundation grant #1055/22 and the Schmidt Marine Technology Partners. We thank Olena Kovalenko for assistance with data analysis and visualization, and four anonymous reviewers for their helpful comments.


References

- [1] O. Burggraaff, N. Schmidt, J. Zamorano, K. Pauly, S. Pascual, C. Tapia, E. Spyarakos, and F. Snik, “Standardized spectral and radiometric calibration of consumer cameras,” *Optics express*, vol. 27, no. 14, pp. 19075–19101, 2019.
- [2] D. Akkaynak, T. Treibitz, B. Xiao, U. A. Gürkan, J. J. Allen, U. Demirci, and R. T. Hanlon, “Use of commercial off-the-shelf digital cameras for scientific data acquisition

- and scene-specific color calibration,” *JOSA A*, vol. 31, no. 2, pp. 312–321, 2014.
- [3] S. J. Kim, H. T. Lin, Z. Lu, S. Süsstrunk, S. Lin, and M. S. Brown, “A new in-camera imaging model for color computer vision and its application,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 12, pp. 2289–2302, 2012.
- [4] S. Paris and F. Durand, “A fast approximation of the bilateral filter using a signal processing approach,” *International journal of computer vision*, vol. 81, pp. 24–52, 2009.
- [5] C. Guo, C. Li, J. Guo, C. C. Loy, J. Hou, S. Kwong, and R. Cong, “Zero-reference deep curve estimation for low-light image enhancement,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 1780–1789.
- [6] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232.
- [7] V. Bychkovsky, S. Paris, E. Chan, and F. Durand, “Learning photographic global tonal adjustment with a database of input/output image pairs,” in *CVPR 2011*. IEEE, 2011, pp. 97–104.
- [8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [9] D. G. Lowe, “Object recognition from local scale-invariant features,” in *Proceedings of the seventh IEEE international conference on computer vision*, vol. 2. Ieee, 1999, pp. 1150–1157.
- [10] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba, “Sun database: Large-scale scene recognition from abbey to zoo,” in *2010 IEEE computer society conference on computer vision and pattern recognition*. IEEE, 2010, pp. 3485–3492.
- [11] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “Nerf: Representing scenes as neural radiance fields for view synthesis,” *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [12] D. C. Lee, M. Hebert, and T. Kanade, “Geometric reasoning for single image structure recovery,” in *2009 IEEE conference on computer vision and pattern recognition*. IEEE, 2009, pp. 2136–2143.
- [13] R. Nguyen, D. K. Prasad, and M. S. Brown, “Raw-to-raw: Mapping between image sensor color responses,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 3398–3405.
- [14] A. Chakrabarti, D. Scharstein, and T. Zickler, “An empirical camera model for internet color vision.” in *BMVC*, vol. 1, no. 2. Citeseer, 2009, p. 4.
- [15] D. Akkaynak, T. Treibitz, T. Shlesinger, Y. Loya, R. Tamir, and D. Iluz, “What is the space of attenuation coefficients in underwater computer vision?” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 4931–4940.
- [16] M. Afifi and A. Abuolaim, “Semi-supervised raw-to-raw mapping,” *arXiv preprint arXiv:2106.13883*, 2021.
- [17] D. Cheng, D. K. Prasad, and M. S. Brown, “Illuminant estimation for color constancy: why spatial-domain methods work and the role of the color distribution,” *JOSA A*, vol. 31, no. 5, pp. 1049–1058, 2014.
- [18] A. Bhandari, A. Kadambi, and R. Raskar, *Computational Imaging*. MIT Press, 2022.
- [19] H. C. Karaimer and M. S. Brown, “Improving color reproduction accuracy on cameras,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 6440–6449.
- [20] A. Munoz, J. I. Echevarria, F. J. Seron, J. Lopez-Moreno, M. Glencross, and D. Gutierrez, “Bssrdf estimation from single images,” in *Computer Graphics Forum*, vol. 30, no. 2. Wiley Online Library, 2011, pp. 455–464.
- [21] P. Debevec, C. Tchou, A. Gardner, T. Hawkins, C. Poullis, J. Stumpfel, A. Jones, N. Yun, P. Einarsson, T. Lundgren *et al.*, “Estimating surface reflectance properties of a complex scene under captured natural illumination,” UNIVERSITY OF SOUTHERN CALIFORNIA LOS ANGELES, Tech. Rep., 2004.
- [22] Y. Yu, P. Debevec, J. Malik, and T. Hawkins, “Inverse global illumination: Recovering reflectance models of real scenes from photographs,” in *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, 1999, pp. 215–224.
- [23] S. W. Oh, M. S. Brown, M. Pollefeys, and S. J. Kim, “Do it yourself hyperspectral imaging with everyday digital cameras,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2461–2469.
- [24] S. Galliani, C. Lanaras, D. Marmanis, E. Baltsavias, and K. Schindler, “Learned spectral super-resolution,” *arXiv preprint arXiv:1703.09470*, 2017.
- [25] K. Barnard and B. Funt, “Camera characterization for color research,” *Color Research & Application*, vol. 27, no. 3, pp. 152–163, 2002.
- [26] Y.-C. Lo, C.-C. Chang, H.-C. Chiu, Y.-H. Huang, C.-P. Chen, Y.-L. Chang, and K. Jou, “Clcc: Contrastive learning for color constancy,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 8053–8063.
- [27] M. El Helou, R. Zhou, S. Süsstrunk, and R. Timofte, “Ntire 2021 depth guided image relighting challenge,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 566–577.
- [28] M. Afifi, J. T. Barron, C. LeGendre, Y.-T. Tsai, and F. Bleibel, “Cross-camera convolutional color constancy,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 1981–1990.
- [29] M. Afifi, A. Abdelhamed, A. Abuolaim, A. Punnapurath, and M. S. Brown, “Cie xyz net: Unprocessing images for low-level computer vision tasks,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 9, pp. 4688–4700, 2021.
- [30] Z. Lou, T. Gevers, N. Hu, M. P. Lucassen *et al.*, “Color constancy by deep learning,” in *BMVC*, 2015, pp. 76–1.
- [31] A. Alsam and G. Finlayson, “Recovering spectral sensitivities with uncertainty,” in *Conference on Colour in Graphics, Imaging, and Vision*, vol. 2002, no. 1. Society for Imaging Science and Technology, 2002, pp. 22–26.
- [32] J. Y. Hardeberg, “On the spectral dimensionality of object colours,” in *Conference on Colour in Graphics,*

- Imaging, and Vision*, vol. 2002, no. 1. Society for Imaging Science and Technology, 2002, pp. 480–485.
- [33] J. P. Parkkinen, J. Hallikainen, and T. Jaaskelainen, “Characteristic spectra of munsell colors,” *JOSA A*, vol. 6, no. 2, pp. 318–322, 1989.
- [34] F. M. Martínez-Verdú, J. Pujol Ramo, P. Capilla Perea *et al.*, “Calculation of the color matching functions of digital cameras from their complete spectral sensitivities,” *Journal of Imaging Science and Technology*, vol. 46, pp. 15–25, 2002.
- [35] “Adobe Digital Negative (DNG) specification, version 1.6.0.0,” https://helpx.adobe.com/content/dam/help/en/photoshop/pdf/dng_spec_1_6_0_0.pdf, accessed: 2023-04-02.
- [36] M. M. Darrodi, G. Finlayson, T. Goodman, and M. Mackiewicz, “Reference data set for camera spectral sensitivity estimation,” *JOSA A*, vol. 32, no. 3, pp. 381–391, 2015.
- [37] L. MacDonald and W. Ji, “Colour characterisation of a high-resolution digital camera,” in *Conference on Colour in Graphics, Imaging, and Vision*, vol. 2002, no. 1. Society for Imaging Science and Technology, 2002, pp. 433–437.
- [38] P. L. Vora, J. E. Farrell, J. D. Tietz, and D. H. Brainard, “Digital color cameras—2—spectral response,” 1997.
- [39] P. M. Hubel, D. Sherman, and J. E. Farrell, “A comparison of methods of sensor spectral sensitivity estimation,” in *Color and Imaging Conference*, vol. 1994, no. 1. Society for Imaging Science and Technology, 1994, pp. 45–48.
- [40] J. Y. Hardeberg, H. Brettel, and F. J. Schmitt, “Spectral characterization of electronic cameras,” in *SYBEN-Broadband European Networks and Electronic Image Capture and Publishing*. International Society for Optics and Photonics, 1998, pp. 100–109.
- [41] G. D. Finlayson, S. Hordley, and P. M. Hubel, “Recovering device sensitivities with quadratic programming,” in *Color and imaging conference*, vol. 1998, no. 1. Society for Imaging Science and Technology, 1998, pp. 90–95.
- [42] K. Barnard and B. V. Funt, “Camera calibration for color research,” in *Human vision and electronic imaging IV*, vol. 3644. SPIE, 1999, pp. 576–585.
- [43] B. Dias, “Robust color sensor response characterization,” in *Color and Imaging Conference*, vol. 2000, no. 1. Society for Imaging Science and Technology, 2000, pp. 144–148.
- [44] P. Carvalho, A. Santos, A. Dourado, and B. Ribeiro, “Learning spectral calibration parameters for color inspection,” in *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, vol. 2. IEEE, 2001, pp. 660–667.
- [45] S. Quan, N. Ohta, and X. Jiang, “Comparative study on sensor spectral sensitivity estimation,” in *Color Imaging VIII: Processing, Hardcopy, and Applications*, vol. 5008. SPIE, 2003, pp. 209–220.
- [46] J. M. DiCarlo, G. E. Montgomery, and S. W. Trovinger, “Emissive chart for imager calibration,” in *Color and Imaging Conference*, vol. 2004, no. 1. Society for Imaging Science and Technology, 2004, pp. 295–301.
- [47] M. Ebner, “Estimating the spectral sensitivity of a digital sensor using calibration targets,” in *Proceedings of the 9th annual conference on Genetic and evolutionary computation*, 2007, pp. 642–649.
- [48] C. Mauer and D. Wueller, “Measuring the spectral response with a set of interference filters,” in *IS&T/SPIE Electronic Imaging*. International Society for Optics and Photonics, 2009, pp. 72 500S–72 500S.
- [49] H. Zhao, R. Kawakami, R. T. Tan, and K. Ikeuchi, “Estimating basis functions for spectral sensitivity of digital cameras,” in *Meeting on Image Recognition and Understanding*, vol. 2009, no. 1, 2009.
- [50] M. Rump, A. Zinke, and R. Klein, “Practical spectral characterization of trichromatic cameras,” in *Proceedings of the 2011 SIGGRAPH Asia Conference*, 2011, pp. 1–10.
- [51] T. W. Pike, “Using digital cameras to investigate animal colouration: estimating sensor sensitivity functions,” *Behavioral Ecology and Sociobiology*, vol. 65, pp. 849–858, 2011.
- [52] R. S. Bedarkar, *Camera spectral sensitivity characterization using a blackbody source*. University of Maryland, College Park, 2011.
- [53] S. Han, Y. Matsushita, I. Sato, T. Okabe, and Y. Sato, “Camera spectral sensitivity estimation from a single image under unknown illumination by using fluorescence,” in *Proc. IEEE CVPR*, 2012, pp. 805–812.
- [54] J. Jiang, D. Liu, J. Gu, and S. Süsstrunk, “What is the space of spectral sensitivity functions for digital color cameras?” in *IEEE Workshop on Applications of Computer Vision (WACV)*, 2013, pp. 168–179.
- [55] D. Prasad, R. Nguyen, and M. Brown, “Quick approximation of camera’s spectral response from casual lighting,” in *Proc. IEEE ICCV Workshops*, 2013, pp. 844–851.
- [56] R. Kawakami, H. Zhao, R. T. Tan, and K. Ikeuchi, “Camera spectral sensitivity and white balance estimation from sky images,” *IJCV*, vol. 105, no. 3, pp. 187–204, 2013.
- [57] D. L. Bongiorno, M. Bryson, D. G. Dansereau, and S. B. Williams, “Spectral characterization of cots rgb cameras using a linear variable edge filter,” in *Digital Photography IX*, vol. 8660. SPIE, 2013, pp. 157–166.
- [58] C. P. Huynh, A. Robles-Kelly *et al.*, “Recovery of spectral sensitivity functions from a colour chart image under unknown spectrally smooth illumination,” in *ICPR*, 2014, pp. 708–713.
- [59] L. W. MacDonald, “Determining camera spectral responsivity with multispectral transmission filters,” in *Color and Imaging Conference*, vol. 23. Society for Imaging Science and Technology, 2015, pp. 12–17.
- [60] P. Bartczak, A. Gebejes, P. Falt, J. Parkkinen, and P. Silfstein, “Led-based spectrally tunable light source for camera characterization,” in *2015 Colour and Visual Computing Symposium (CVCS)*. IEEE, 2015, pp. 1–5.
- [61] G. Finlayson, M. M. Darrodi, and M. Mackiewicz, “Rank-based camera spectral sensitivity estimation,” *JOSA A*, vol. 33, no. 4, pp. 589–599, 2016.
- [62] A. Manakov, “Evaluation of computational radiometric and spectral sensor calibration techniques,” in *Optics, Photonics and Digital Technologies for Imaging Applications IV*, vol. 9896. SPIE, 2016, pp. 130–143.
- [63] L. Zhang, Y. Fu, Y. Zheng, and H. Huang, “Camera spectral sensitivity, illumination and spectral reflectance estimation for a hybrid hyperspectral image capture system,” in *2017 IEEE International Conference on Image*

- Processing (ICIP)*. IEEE, 2017, pp. 545–545.
- [64] S. Chaji, A. Pourreza, H. Pourreza, and M. Rouhani, “Estimation of the camera spectral sensitivity function using neural learning and architecture,” *JOSA A*, vol. 35, no. 6, pp. 850–858, 2018.
- [65] A. Karge, I. Rieger, B. Eberhardt, and A. Schilling, “Using chromaticity error minimisation for fast camera spectral responsivity measurement,” in *Color and Imaging Conference*, vol. 2018, no. 1. Society for Imaging Science and Technology, 2018, pp. 67–74.
- [66] J. Zhu, X. Xie, N. Liao, Z. Zhang, W. Wu, and L. Lv, “Spectral sensitivity estimation of trichromatic camera based on orthogonal test and window filtering,” *Optics Express*, vol. 28, no. 19, pp. 28 085–28 100, 2020.
- [67] M. E. Toivonen and A. Klami, “Practical camera sensor spectral response and uncertainty estimation,” *Journal of Imaging*, vol. 6, no. 8, p. 79, 2020.
- [68] Y. Ji, Y. Kwak, S. M. Park, and Y. L. Kim, “Compressive recovery of smartphone rgb spectral sensitivity functions,” *Optics Express*, vol. 29, no. 8, pp. 11 947–11 961, 2021.
- [69] S. Tominaga, S. Nishi, and R. Ohtera, “Measurement and estimation of spectral sensitivity functions for mobile phone cameras,” *Sensors*, vol. 21, no. 15, p. 4985, 2021.
- [70] L. Ma, Q. Gao, S. Chen, C. Li, and K. Xiao, “Recovery of camera spectral sensitivity based on multi-objective optimization method,” in *First Optics Frontier Conference*, vol. 11850. SPIE, 2021, pp. 304–309.
- [71] B. Xu, L. Ma, and P. Li, “Rank-based camera spectral sensitivity estimation under multiple illuminations,” in *Second Optics Frontier Conference (OFS 2022)*, vol. 12307. SPIE, 2022, pp. 46–53.
- [72] H. Fan and M. R. Luo, “Camera spectral sensitivity estimation based on a display,” in *Innovative Technologies for Printing and Packaging*. Springer, 2023, pp. 24–30.
- [73] Phil Harvey, *ExifTool*, Kingston, Ontario, Canada, 2016. [Online]. Available: <https://exiftool.org/>
- [74] “Rawsamples.ch,” <http://rawsamples.ch/index.php/en/>, accessed: 2023-04-11.
- [75] “Ilia Sibiriyakov’s github,” <https://github.com/ilia3101/ExtractAdobeCameraMatrices>, accessed: 2023-04-11.
- [76] “DataThief III by B. Tummars, 2006,” <https://datathief.org/>, accessed: 2023-04-11.
- [77] M. Stevens, C. A. Párraga, I. C. Cuthill, J. C. Partridge, and T. S. Troscianko, “Using digital photography to study animal coloration,” *Biological Journal of the Linnean society*, vol. 90, no. 2, pp. 211–237, 2007.
- [78] D. Akkaynak, E. Chan, J. J. Allen, and R. T. Hanlon, “Using spectrometry and photography to study color underwater,” in *Proc. MTS/IEEE OCEANS*, 2011.
- [79] E. Berra, S. Gibson-Poole, A. MacArthur, R. Gaulton, and A. Hamilton, “Estimation of the spectral sensitivity functions of un-modified and modified commercial off-the-shelf digital cameras to enable their use as a multispectral imaging system for uavs,” in *International Conference on Unmanned Aerial Vehicles in Geomatics*. Newcastle University, 2015.
- [80] M. Brady and G. E. Legge, “Camera calibration for natural image studies and vision research,” *JOSA A*, vol. 26, no. 1, pp. 30–42, 2009.
- [81] C. P. Huynh and A. Robles-Kelly, “Comparative colorimetric simulation and evaluation of digital cameras using spectroscopy data,” in *9th Biennial Conference of the Australian Pattern Recognition Society on Digital Image Computing Techniques and Applications (DICTA 2007)*. IEEE, 2007, pp. 309–316.
- [82] V. Lebourgeois, A. Bégué, S. Labbé, B. Mallavan, L. Prévot, and B. Roux, “Can commercial digital cameras be used as multispectral sensors? a crop monitoring test,” *Sensors*, vol. 8, no. 11, pp. 7300–7322, 2008.
- [83] M. V. Conde, R. Timofte, Y. Huang, J. Peng, C. Chen, C. Li, E. Pérez-Pellitero, F. Song, F. Bai, S. Liu *et al.*, “Reversed image signal processing and raw reconstruction. aim 2022 challenge report,” in *European Conference on Computer Vision*. Springer, 2022, pp. 3–26.
- [84] R. M. Nguyen, D. K. Prasad, and M. S. Brown, “Training-based spectral reconstruction from a single rgb image,” in *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part VII 13*. Springer, 2014, pp. 186–201.
- [85] N. Akhtar, F. Shafait, and A. Mian, “Sparse spatio-spectral representation for hyperspectral image super-resolution,” in *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part VII 13*. Springer, 2014, pp. 63–78.
- [86] —, “Hierarchical beta process with gaussian process prior for hyperspectral image super resolution,” in *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*. Springer, 2016, pp. 103–120.
- [87] X. Luo, X. Zhang, P. Yoo, R. Martin-Brualla, J. Lawrence, and S. M. Seitz, “Time-travel rephotography,” *ACM Transactions on Graphics (TOG)*, vol. 40, no. 6, pp. 1–12, 2021.
- [88] H. C. Karaimer and M. S. Brown, “A software platform for manipulating the camera imaging pipeline,” in *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*. Springer, 2016, pp. 429–444.



Grigory Solomatov is a postdoctoral fellow in the Department of Marine Technologies at the University of Haifa and is part of the COLOR Lab at the Interuniversity Institute of Marine Sciences in Eilat. He earned his PhD in mathematics and computer science from the Technical University of Denmark. His current research pivots towards leveraging techniques from optimization and machine learning to advance underwater computer vision and computational photography.



Derya Akkaynak is an assistant professor in the Department of Marine Technologies at the University of Haifa, where she directs the COLOR Lab located at the Interuniversity Institute of Marine Sciences in Eilat. She received her PhD in Mechanical & Oceanographic Engineering from MIT and Woods Hole Oceanographic Institution, MSc in Aeronautics and Astronautics from MIT, and BSc in Aerospace Engineering from METU (Turkey). She studies light and vision in the ocean, at the intersection of which lies color.