

A New In-Camera Imaging Model for Color Computer Vision and its Application

Seon Joo Kim, *Member, IEEE*, Hai Ting Lin, *Student Member, IEEE*, Zheng Lu, *Student Member, IEEE*, Sabine Süsstrunk, *Senior Member, IEEE*, Stephen Lin, *Member, IEEE*, and Michael S. Brown, *Member, IEEE*

Abstract—We present a study of the in-camera image processing through an extensive analysis of more than 10,000 images from over 30 cameras. The goal of this work is to investigate if image values can be transformed to physically meaningful values, and if so, when and how this can be done. From our analysis, we found a major limitation of the imaging model employed in conventional radiometric calibration methods and propose a new in-camera imaging model that fits well with today’s cameras. With the new model, we present associated calibration procedures that allow us to convert sRGB images back to their original CCD RAW responses in a manner that is significantly more accurate than any existing methods. Additionally, we show how this new imaging model can be used to build an image correction application that converts an sRGB input image captured with the wrong camera settings to an sRGB output image that would have been recorded under the correct settings of a specific camera.

Index Terms—Radiometric calibration, in-camera image processing, gamut mapping, white balance.

1 INTRODUCTION

MANY computer vision algorithms assume that cameras are accurate light measuring devices which capture images that are directly related to the actual scene radiance. Representative algorithms include photometric stereo, shape from shading, color constancy, intrinsic image computation, and high dynamic range imaging. Digital cameras, however, are much more than light measuring devices; the imaging pipelines used in digital cameras are well known to be nonlinear. Moreover, the primary goal of many cameras is to create visually pleasing pictures rather than capturing accurate physical descriptions of the scene.

In this paper, we present a study of the in-camera image processing through an extensive analysis of an image database collected by capturing images of scenes under different conditions with over 30 commercial cameras. The ultimate goal is to investigate if image values can be transformed to physically meaningful values and if so, when and how this can be done. From our analysis, we found a glaring limitation in the conventional imaging model employed to determine the nonlinearities in the imaging pipeline (i.e. radiometric calibration). In particular, the conventional radiometric models assume that the irradiance (RAW) to image intensity (sRGB) transformation is attributed to a single nonlinear tone-mapping step. However, this tone-mapping step alone is inadequate to describe saturated colors. As a result, such color values are often mis-interpreted

by the conventional radiometric calibration methods.

In our analysis, we found that the color mapping component which includes gamut mapping [1] has been missing in previous models of imaging pipeline. In this paper, we describe how to introduce this step into the imaging pipeline, together with calibration procedures to estimate the associated parameters for a given camera model. This allows us to model the full transformation from RAW to sRGB with much more accuracy than demonstrated by prior radiometric calibration techniques.

In addition, we demonstrate how our new imaging pipeline model can be used to develop a system that converts an sRGB input image captured with the wrong settings to an sRGB output image that would have been recorded under different and correct camera settings. In essence, our model allows us not only to undo the onboard image processing, but also to refinish an image in a camera-specific manner, producing a result that would appear almost identical to the sRGB output that the camera would have produced with the new settings. For example, given a JPEG image (sRGB) taken with a Canon EOS-1D under a certain white balance and picture style¹, we can reproduce this photograph as it would appear from the same camera but with a different white balance and picture style. To our knowledge, this is the first system capable of providing such camera-specific refinishing abilities. Moreover, with minor modifications to our approach, we can even allow the user to produce a photograph using another camera’s settings.

Preliminary findings reported in this paper appeared in

- S. J. Kim is with SUNY Korea. E-mail: seonjookim@snykorea.ac.kr
- H. Lin, Z. Lu, and M. S. Brown are with National University of Singapore. E-mail: linhait, luzheng, brown@comp.nus.edu.sg
- S. Süsstrunk is with EPFL. E-mail: sabine.susstrunk@epfl.ch
- S. Lin is with Microsoft Asia. E-mail: stevelin@microsoft.com

1. Picture style refers to the photofinishing feature of Canon cameras to produce optimized pictures under specific scenes, such as portrait and landscape. Other camera manufacturers offer similar photofinishing styles, e.g. Nikon’s “Image Optimizer” and Sony’s “Creative Style”. For simplicity, we collectively refer to these functions as *picture style*.

[2]. This prior work provides the initial analysis of the missing gamut mapping step in the in-camera imaging pipeline. In addition to presenting the new application, this paper extends the work in [2] in several ways. We present a more detailed imaging model that factors the in-camera color transformation into different operations which need to be separately considered for flexible modeling. We also present a more robust technique for radial basis function computation to accurately model the gamut mapping in cameras. We also expand the applicability of color gamut mapping calibration, from only cameras that provide RAW images as in [2] to any camera at all, including point-and-shoot models that are popular among consumers, and we also introduce a method to transform colors between different cameras under arbitrary settings.

The remainder of the paper is organized as follows: we first discuss related work in Section 2. We then explain how our database was collected and describe significant observations from this database in Section 3. We introduce a new in-camera imaging model in Section 4 and the associated calibration procedures in Section 5. An application of our new framework is developed in Section 6 and the experimental results are shown in Section 7. We conclude with a discussion about our findings and future work in Section 8.

2 PRELIMINARIES AND RELATED WORK

Radiometric calibration is an area in computer vision in which the goal is to compute the camera response function (f) that maps the amount of light collected by each CCD pixel (irradiance e) to pixel intensities (I) in the output image:

$$I_{\mathbf{x}} = f(e_{\mathbf{x}}), \quad (1)$$

where \mathbf{x} is the pixel location. Eq. 1 can be extended to deal with color as follows [3]:

$$\begin{bmatrix} I_{rx} \\ I_{gx} \\ I_{bx} \end{bmatrix} = \begin{bmatrix} f_r(e_{rx}) \\ f_g(e_{gx}) \\ f_b(e_{bx}) \end{bmatrix}, \quad (2)$$

$$\mathbf{e}_x = \begin{bmatrix} e_{rx} \\ e_{gx} \\ e_{bx} \end{bmatrix} = \mathbf{T}\mathbf{E}_x. \quad (3)$$

\mathbf{T} is a 3×3 transformation that captures both the transformation from the camera's color space (\mathbf{E}_x) to sRGB (\mathbf{e}_x) and white balancing.

The radiometric mapping f is almost always nonlinear due to the design factors built into digital cameras for a variety of reasons, including compressing the dynamic range of the imaged scene (tone-mapping), accounting for nonlinearities in display systems (gamma correction), mimicking the response of films, or to create aesthetic effects [4], [5]. When the response function f is known, the image intensities can be inverted back to relative scene radiance values enabling physics-based photometric analysis of the scene.

2.1 Radiometric Calibration

Conventional radiometric calibration algorithms rely on multiple images of a scene taken with different exposures. Assuming constant radiance, which implies constant illumination, a change in intensity is explained by a change in exposure. Given a pair of images with an exposure ratio of k' , the response function f is computed by solving the following equation from intensity values (I , I') at corresponding points:

$$\frac{f^{-1}(I'_x)}{f^{-1}(I_x)} = k'. \quad (4)$$

The main difference among various calibration methods is the model used to represent the response function. The existing models for a radiometric response function include the gamma curve [6], polynomial [7], non-parametric [8], and PCA based model [9]. Other than the work in [10] where the color was explained as having the same response function for all the channels but with different exposure level per channel, most methods do not deal with color and compute the response function independently per channel.

While different radiometric calibration methods vary in either how the response function is modeled and/or computed, all methods share a common view that it is a fixed property of a given camera model. In fact, this view was exploited to compute the radiometric response function by applying statistical analysis on images downloaded from the web in [11]. One exception is the work in [12] where the response function was modeled differently per image using a probabilistic approach. Another exception is the recent work in [3] where the goal was to provide an analysis of the factors that contribute to the color output of a camera for internet color vision. They proposed a 24-parameter model to explain the imaging pipeline and the parameters are iteratively computed using available RAW data. Through their analysis, they suggest that the color rendering function f is scene-dependent. They go further to suggest that fixed nonlinearities per channel/camera as used in traditional radiometric calibration are often inadequate.

Before moving forward, it is important to clarify the issue of scene dependency of the in-camera imaging process. If the process is scene dependent as mentioned in [3], traditional radiometric calibration would be inadequate and the only option would be to use single-image based radiometric calibration methods [13], [5]. While the single image calibration algorithms are conceptually the best choice, they are sometimes unstable because they rely on edge regions, which are sensitive to noise and may go through further processing such as sharpening onboard the camera.

There are generally two color rendering strategies with regards to how digital cameras convert CCD RAW responses to the final output: the *photofinishing* model and the *slide or photographic reproduction* model [1]. In the "photofinishing" model, the imaging pipeline varies (possibly in a spatially varying manner) to produce a visually pleasing image. The auto-mode in cameras will trigger the photofinishing model as well as the optimizers such as Dynamic Lighting Optimizer on the Canon EOS550D

and D-Range Optimizer in Sony α -200. The photographic reproduction model, on the other hand, uses fixed color rendering. For most high-end cameras, it is possible to set the camera in this photographic mode by turning the camera settings to manual and turning off all settings pertaining to photofinishing, such as Dynamic Lighting Optimizer. For the remainder of this paper it is assumed that the algorithms discussed are intended to work in the photographic reproduction mode.

2.2 White Balancing and Color Transfer

White balance and color transfer are applications affected by the imaging model of a camera. White balancing, or computational color constancy, seeks to estimate the illumination color of a scene and remove its effects in the image [14]. Simply put, a white balancing algorithm aims to make a white object look white in the image regardless of the scene illumination. Computational color constancy is inherently an ill-posed problem. As a result, color constancy algorithms rely on assumptions about the illumination and/or scene surfaces, such as the grey-world assumption [15] and retinex theory [16]. Other recent color constancy methods exploit geometric models of color space [17], [18] and statistical analysis [19], [20] to recover illumination color.

The white balancing methods described above may not be satisfactory if they are applied directly to nonlinear sRGB images because the image values are nonlinearly transformed from the linear irradiance (RAW) values. By modeling the nonlinearities in the imaging process, our method performs white balancing in the correct domain. Our work does not propose a white balancing algorithm itself, but learns what a particular camera does for white balancing and applies this later to change the white balance of an image. Besides white balancing, our work also deals with other color transformations that occur in cameras, e.g. Canon's Picture Style, which makes a picture more vivid, neutral, or colorimetrically faithful to actual colors under standard daylight conditions.

Another topic related to this paper is color transfer, in which the colors of an image are modified to match the color characteristics of another [21]. Color transfer has been used for various purposes, which include adjusting color to enhance the harmony among the colors of a photograph [22], transferring the look of a model photograph through tone management [23], transferring illumination using webcam data [24], interactive appearance editing by model-based navigation [25], and transferring the color of an image to enhance a desired color theme [26]. Our work may be viewed as a form of color transfer, but differs from conventional color transfer techniques in that it specifically aims to model the color transformations of a given camera under different settings and apply them for the purpose of color correction.

3 DATA COLLECTION AND OBSERVATION

For the analysis, we collected more than 10,000 images from 31 cameras ranging from DSLR cameras to point-

and-shoot cameras. Images were taken in manual mode under different settings including white balance, aperture, and shutter speed. The images were also collected under different lighting conditions: indoor lighting and/or outdoor cloudy condition. Images are captured three times under the same condition to check the shutter speed consistency. RAW images are also saved if the camera supports RAW and the RAW files are rendered using the software *dcraw*². We additionally use the database in [3] which includes over 1000 images from 35 cameras. Cameras from most of the major manufacturers are included as shown in Fig. 6. Though the cameras used for data collection are not uniformly distributed among manufacturers, they reflect the reality of certain manufacturers being more popular than others.

For cameras with RAW support, both sRGB and RAW data are recorded. The target objects for our dataset are two Macbeth ColorChecker charts, specifically a 24-patch chart and a 140-patch chart. There are several reasons why these color charts were used for our analysis. First, since the patches are arranged in a regular grid pattern, we can automatically extract colors from different patches with simple registration. Also, measurements from different pixels within a patch can be averaged to reduce the impact of image noise on the analysis. Finally, these color charts include a broad spectrum of colors and different levels of gray, which facilitate radiometric calibration and color analysis.

Using the conventional radiometric model, pairs of intensity measurements at corresponding patches in two differently exposed images constitute all the necessary information to recover the radiometric response function of a camera [9]. These pairs can be arranged into a plot that represents the *brightness transfer function* (BTF [10]), which can be formulated from Eq. 4 as

$$I'_x = \tau_k(I_x) = f(k'f^{-1}(I_x)), \quad (5)$$

where τ_k is the BTF, f is the response function, and k' is the exposure ratio. The BTF describes how image intensity changes with respect to an exposure change under a given response function. If the response function is a fixed property of a camera and the model in Eq. 1 is valid, the BTF should be the same for all pairs of images that share the same exposure ratio regardless of other camera settings and lighting conditions. Notice that even if we consider the color transformation in Eq. 3, the BTFs should still remain the same for the same exposure ratio as long as the color transformation remains unchanged between images, i.e.:

$$\frac{f^{-1}(I'_{cx})}{f^{-1}(I_{cx})} = k' \frac{\mathbf{t}'_c \mathbf{E}_x}{\mathbf{t}_c \mathbf{E}_x} = k' \text{ if } \mathbf{t}_c = \mathbf{t}'_c. \quad (6)$$

In the above equation, \mathbf{t}_c is a row of the color transformation \mathbf{T} that corresponds to the color channel c .

To validate the model in Eq. 1 and the assumption that the response f is a fixed property of a camera, we compare

² <http://www.cybercom.net/~dcoffin/dcrw/>. We used the command `dcraw -v -D -4 -T`.

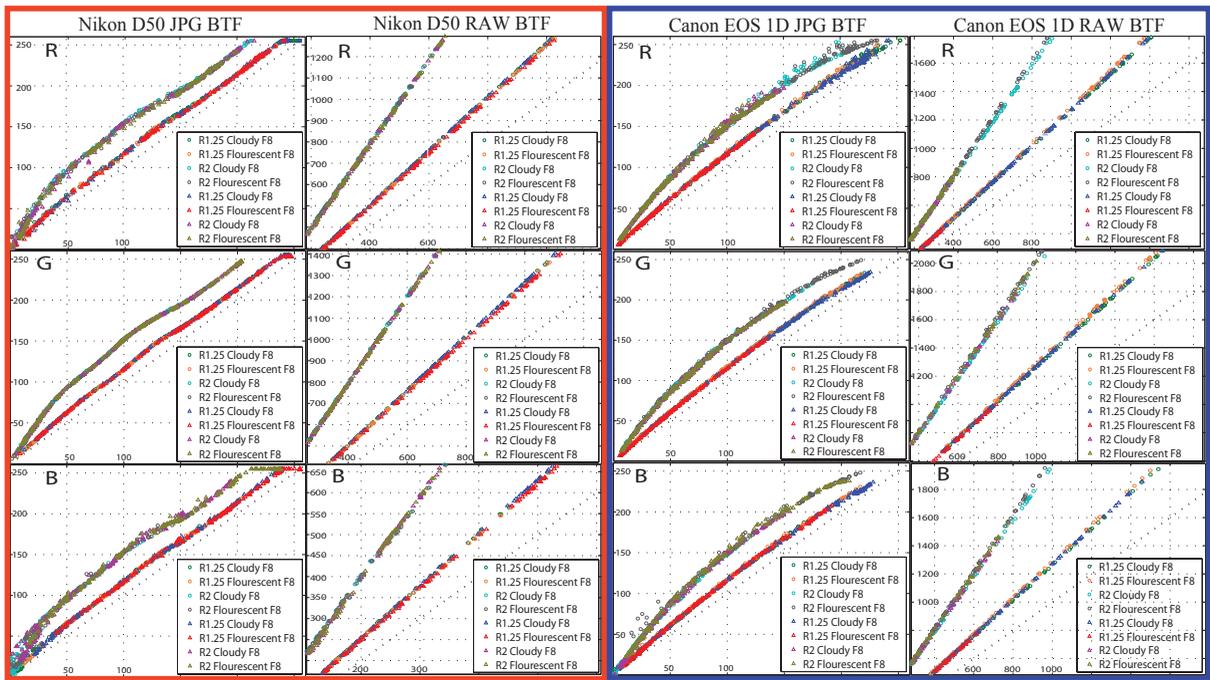


Fig. 1: Brightness transfer functions for Nikon D50 and Canon EOS-1D. Each plot includes several BTFs with different exposure ratios (1.25 and 2.0), different lighting environments (\circ : outdoors, \triangle : indoors), and different white balance settings (cloudy and fluorescent). The key observation from these plots is that the BTFs of sRGB images with the same exposure ratio exhibit a consistent form aside from outliers and small shifts. For better viewing, please zoom the electronic PDF.

the BTFs of different cameras under different settings. Representative examples from two cameras are shown in Fig. 1 for clarity. In the figure, each point represents the change in brightness for a given patch between the image pair.

Through our analysis of the database, we made several key observations, which can also be observed in Fig. 1. The BTFs of a given camera and exposure ratio exhibit a consistent shape up to slight shifts and a small number of measurement outliers. BTFs recorded in the green channel are generally more stable than in the other channels and have a smaller amount of outliers. Also, the appearance of shifts and outliers tends to increase with larger exposure ratios.

The shifts can be explained with the inconsistency of the shutter speed. In our experiments, we control the exposure by changing the shutter speed³, and it is well known that the shutter speeds of cameras may be imprecise [27]. In particular, we have found that shutter speeds of cameras with high shutter-usage count tend to be less accurate, as observed from measurement inconsistency over repeated image captures under the same setting. We should note that we can rule out the illumination change as a cause because of our illumination monitoring and the consistent BTFs measured by other cameras under the same conditions. As these shifts also exist in raw image BTFs, onboard camera processing can also be ruled out.

We found that some outliers, though having intensity

3. We use shutter speed to control exposure because changing the aperture could result in spatial variation of irradiance due to vignetting and depth-of-focus.

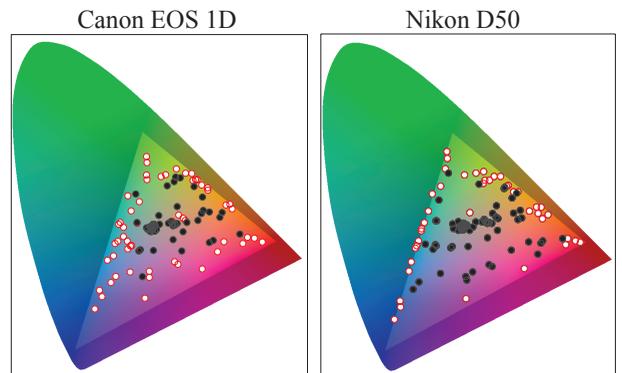


Fig. 2: Positions of color points in the sRGB chromaticity gamut. Inliers (filled with black) are surrounded by outliers (filled with white). Outliers (as observed in Fig. 1) are color points with high saturation levels, and lie close to the boundary of the sRGB gamut.

values well within the dynamic range of the given color channel, have a 0 or 255 intensity value in at least one of the other channels. These clipped values at the ends of the dynamic range do not accurately represent the true scene irradiance. One significant reason for outliers observed is that when a camera's color range extends beyond that of the sRGB gamut, gamut mapping is needed to convert colors from outside the sRGB gamut to within the gamut for the purpose of sRGB representation [1], [28], [29]. As seen in Fig. 2, we can observe the vast majority of outliers in our dataset have high color saturation levels and lie close to the boundary of the sRGB color gamut. This gamut mapping essentially produces a change in color for points

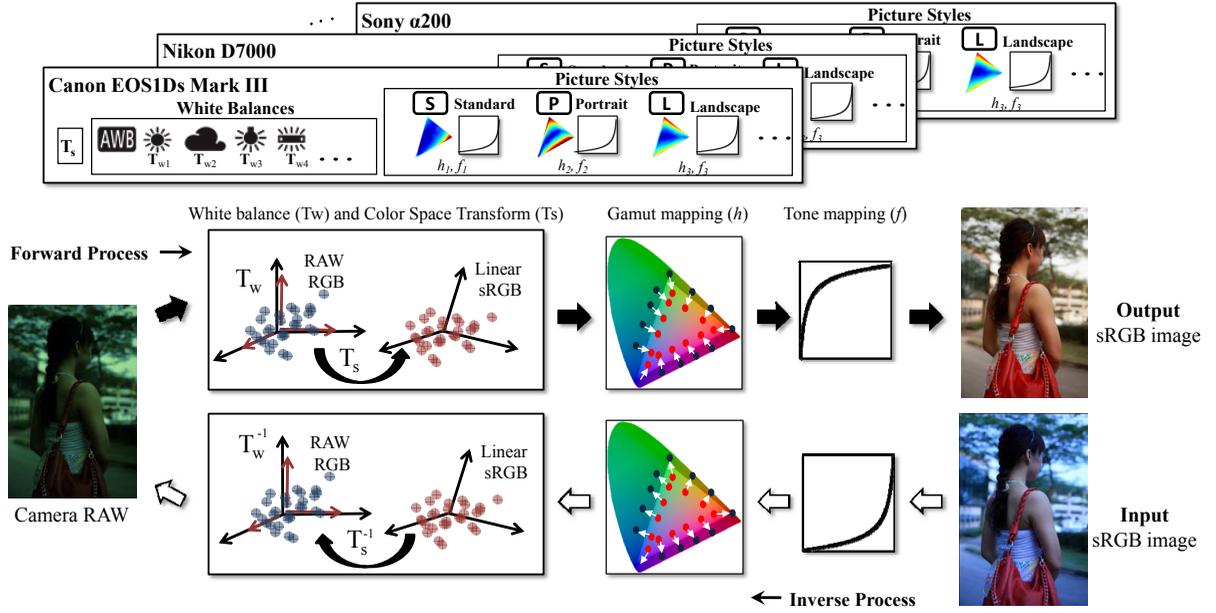


Fig. 3: Overview of the new imaging model of Eq. 7 and its application. The parameters of the imaging process for different cameras and settings including the white balance and the picture style are calibrated using training images (Section 5). An sRGB image under a certain setting can be transformed to RAW through reverse imaging, and then to another sRGB image under a different setting through forward imaging using the corresponding parameters (Section 6).

outside the sRGB gamut, and if out-of-gamut colors are shifted in different ways between different exposures, the color transformation becomes different ($\mathbf{T} \neq \mathbf{T}'$ in Eq. 6) between the two images. Thus these points become outliers positioned off from the BTF. This effect of gamut mapping becomes more significant with larger exposure ratios, since the out-of-gamut colors need a greater displacement in color space to move into the sRGB gamut.

To summarize, the observations imply that factors such as shutter speed inaccuracies and gamut mapping have to be considered to compute the radiometric response function accurately. Most importantly, the observations show that less saturated colors can be modeled with the conventional radiometric model (Eq. 1) and be linearized accurately. However, it is shown that the conventional model has an essential limitation in representing the nonlinear color mapping in the imaging pipeline and highly saturated colors will not be linearized accurately with the model in Eq. 1.

4 IN-CAMERA IMAGING MODEL

Based on our observation, we introduce the following model for the imaging pipeline inside digital cameras, which is illustrated in Fig. 3.

$$\begin{bmatrix} I_{rx} \\ I_{gx} \\ I_{bx} \end{bmatrix} = \begin{bmatrix} f_r(e_{rx}) \\ f_g(e_{gx}) \\ f_b(e_{bx}) \end{bmatrix}, \text{ where} \\ \begin{bmatrix} e_{rx} \\ e_{gx} \\ e_{bx} \end{bmatrix} = h(\mathbf{T}_s \mathbf{T}_w \mathbf{E}_x). \quad (7)$$

$\mathbf{E}_x = [E_{rx}, E_{gx}, E_{bx}]^T$ is the irradiance, which can be recorded as a RAW image in certain digital cameras⁴. In our model, the RAW values are first white balanced by a 3×3 diagonal matrix \mathbf{T}_w . Then the white balanced raw values, defined in the camera's color space, are transformed to the linear sRGB space by a 3×3 matrix \mathbf{T}_s . Having the linear transformation decomposed to \mathbf{T}_w and \mathbf{T}_s allows more flexibility in designing applications compared to having a single transformation that combines both factors. Notice that the white balance \mathbf{T}_w could actually be applied at a different stage to the same effect, e.g. after the color space transformation \mathbf{T}_s and after the function h in Equation 7. We place white balancing as the first operator in the imaging pipeline based on empirical data from our experiments: in all cameras that we tested, the order in Eq. 7 yielded the best results. Next, the nonlinear color gamut function $h: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is applied and then the final image in the nonlinear sRGB space is computed with the camera response function f .

A noticeable difference between the new model in Eq. 7 and the conventional model in Eq. 1 is the addition of color transformations, especially the nonlinear color gamut mapping function h . In digital cameras, both tone mapping and gamut mapping are employed to transform the colorimetry of the source image to one that produces a visually pleasing image on the actual reproduction medium [1]. The tone mapping by the camera response function f aims to compress the dynamic range of the luminance recorded from the imaged scene. The gamut mapping (h) acts on the color itself and brings the colors that are outside the

4. We assume that the raw value \mathbf{E}_x is demosaicked (i.e. the color filter array values are interpolated) and is linearly related to the actual irradiance as shown in [3]. The RAW BTFs in Fig. 1 also show the linearity of \mathbf{E}_x

sRGB gamut to within the gamut. That is, when a camera's color range extends beyond that of the sRGB gamut, gamut mapping converts colors outside the sRGB gamut to inside the gamut for the purpose of sRGB representation. The gamut mapping process is usually nonlinear with greater transformation of more highly saturated colors near and beyond the boundary of the gamut. In addition to color range compression, gamut mapping may also include different transformations for some specific color ranges, e.g. to make the sky more blue and to make skin tone more vivid. By incorporating this nonlinear color mapping h in the pipeline, our model in Eq. 7 describes the in-camera imaging more accurately than the conventional model. Note that in our new model, the color space transformation \mathbf{T}_s is fixed per camera model, the white balance parameter \mathbf{T}_w is fixed per white balance setting of a specific camera, and the response function f and the color mapping h are fixed per picture style of a camera (Fig. 3).

5 CALIBRATION

Converting a given sRGB image to its RAW representation requires knowledge of the model parameter values in Eq. 7, namely of f , \mathbf{T}_s , \mathbf{T}_w , and h . To calibrate these values, we assume that we are given a number of training images taken by the camera under varied settings with different exposures, white balance, and picture styles. We also assume that the RAW images associated with these training images are provided as well. For each camera model, we compute the color space transform \mathbf{T}_s , a matrix \mathbf{T}_w for each white balance preset, and f and h per picture style.

5.1 Camera Response Function Estimation

We first compute the camera response function f from a set of images taken with varying exposures. At first glance, using a conventional radiometric calibration procedure does not look feasible due to the presence of h in Eq. 7. However, for color points (\mathbf{p}) that satisfy $h(\mathbf{p}) = \mathbf{p}$, the following equation holds between a pair of image intensity values varied by the exposure ratio k :

$$f_c^{-1}(I_{c2})/f_c^{-1}(I_{c1}) = k. \quad (8)$$

Eq. 8 represents the basic principle of traditional radiometric calibration methods, and any of them can be used to compute f . In this paper, we use the method in [4] which is based on a PCA model of camera responses.

The key is then to find a set of points (\mathbf{p}) that satisfy $h(\mathbf{p}) = \mathbf{p}$. In other words, we need to find points that do not get transformed by the gamut mapping function. Since the main purpose of the gamut mapping is to bring the color points outside the gamut into the inside, we assume that colors with low saturation are not significantly transformed by the gamut mapping. Therefore, we only use points with a saturation value (S in HSV color space) below a threshold (β) to compute the response function. Additionally, points with 0 or 255 in any of the color channels are rejected as outliers.

5.2 Color Transformation Matrix Estimation

After computing the camera response function f , we can convert the image values to linearized sRGB values. Then the linear color transformation matrices \mathbf{T}_w and \mathbf{T}_s are computed also by using the points with low color saturation that are not affected by the gamut mapping. The white balance matrix \mathbf{T}_w is a diagonal matrix defined per white balance setting, and the color space transformation matrix \mathbf{T}_s , which aligns the camera's color space with the sRGB space, is defined per camera. We compute the \mathbf{T}_w 's and \mathbf{T}_s that minimize the following error function:

$$\sum_{i=1}^M \sum_{j=1}^N \|\mathbf{T}_s^{-1} \mathbf{X}_{ij} - \mathbf{T}_w \mathbf{E}_{ij}\|^2, \quad (9)$$

where M is the number of white balance settings, N is the number of color points used for estimation, \mathbf{X}_{ij} is the linearized sRGB values computed from the inverse response functions ($\mathbf{X}_{ij} = [f_r^{-1}(I_{r,ij}), f_g^{-1}(I_{g,ij}), f_b^{-1}(I_{b,ij})]^T$), and \mathbf{E}_{ij} denotes the RAW image values ($\mathbf{E}_{ij} = [E_{r,ij}, E_{g,ij}, E_{b,ij}]^T$) that correspond to \mathbf{X}_{ij} .

We note that few camera models such as the Canon EOS-1D and Nikon 200D provide the white balancing scale factors for each channel (\mathbf{T}_w) in its image metadata (EXIF). For those cameras, we can compute the color space transformation \mathbf{T}_s just by incorporating \mathbf{T}_w from this metadata into Eq. 9.

5.3 Color Gamut Mapping Function Estimation

With the camera response function f and the linear color transformations of \mathbf{T}_w and \mathbf{T}_s computed, the last step in our calibration procedure is to solve for the color gamut mapping function h in Eq. 7. The gamut mapping function is a key element in defining the color characteristics of a camera. This nonlinear mapping can be vastly different among cameras as shown in Fig. 4, making colors in one camera more vivid and colors in another camera look softer. Designing a single parametric model that can describe the gamut mapping functions on different camera models is challenging. We therefore opted for a nonparametric approach to model the gamut mapping function based on scatter point interpolation using radial basis functions (RBFs).

Among several variants of RBFs, we adopt the following form [30], [31] to model the inverse gamut mapping function h^{-1} :

$$h^{-1}(\mathbf{X}) = p(\mathbf{X}) + \sum_{i=1}^N \lambda_i \phi(\|\mathbf{X} - \mathbf{X}_i\|) \quad (10)$$

where $\mathbf{X} = [f_r^{-1}(I_r), f_g^{-1}(I_g), f_b^{-1}(I_b)]^T$, color points \mathbf{X}_i are the control (or center) points of the RBF, and N is the number of control points. The λ_i 's are the weights for the basis function ϕ , and we chose $\phi(r) = r$ as the basis function. For $p(\mathbf{X})$, we set it as $p(\mathbf{X}) = \mathbf{c}^T \tilde{\mathbf{X}}$ where $\mathbf{c} = [c_1, c_2, c_3, c_4]^T$ and $\tilde{\mathbf{X}} = [1, \mathbf{X}^T]^T$.

With data from the given sRGB-RAW image pairs and the computed matrices \mathbf{T}_w and \mathbf{T}_s , the corresponding

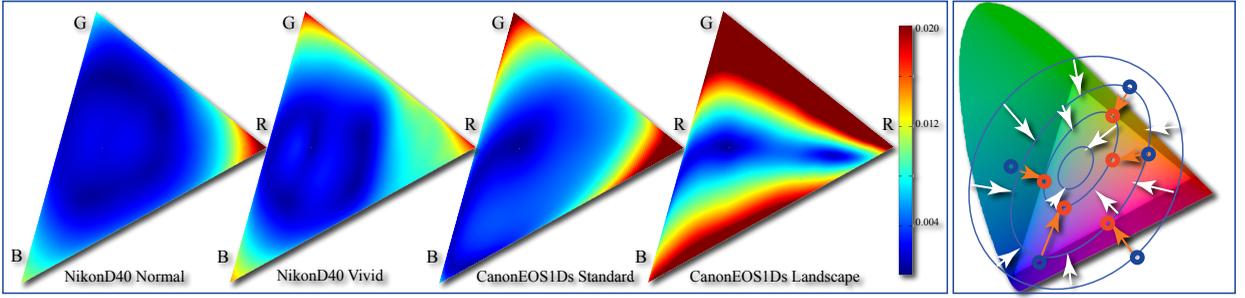


Fig. 4: (Left) Gamut mapping functions (h) display large variations depending on the mode and the camera. The colors in the map indicate the color displacement magnitude of the gamut mapping at a specific color ($\| [r, g, b]^T - h([r, g, b]^T) \|^2$). Interesting to note is that our estimate of the “Landscape” mode of Canon’s picture style matches the description by Canon: “Landscape expresses hues from green to blue more vividly than the Standard settings”. (Right) The gamut mapping function is estimated with scatter point interpolation via radial basis functions. The rings and white arrows show the interpolated color mapping function h and the colored dots and arrows indicate control points.

instance of a control point \mathbf{X}'_i is given by $\mathbf{X}'_i = h^{-1}(\mathbf{X}_i) = \mathbf{T}_s \mathbf{T}_w \mathbf{E}_i$, where \mathbf{E}_i is the RAW value of the control point. Note that all points regardless of their saturation values are used in this stage in contrast to the previous steps where only points with low saturation were used to compute for f and \mathbf{T} 's. From a set of control point pairs $(\mathbf{X}_i, \mathbf{X}'_i)$, the parameters of the RBF in Eq. 10, $\boldsymbol{\lambda} = [\lambda_1, \lambda_2, \dots, \lambda_N]^T$ and \mathbf{c} , are computed as follows [31]:

$$\begin{pmatrix} \mathbf{D} - 8N\pi\rho\mathbf{I} & \tilde{\mathbf{P}} \\ \tilde{\mathbf{P}}^T & \mathbf{0}_{4 \times 4} \end{pmatrix} \begin{pmatrix} \boldsymbol{\lambda} \\ \mathbf{c} \end{pmatrix} = \begin{pmatrix} \mathbf{P}' \\ \mathbf{0}_{4 \times 3} \end{pmatrix}, \quad (11)$$

where \mathbf{D} is an $N \times N$ matrix with $D_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|$, $\tilde{\mathbf{P}}$ is an $N \times 4$ matrix with the i -th row being $\tilde{\mathbf{X}}_i^T$, and \mathbf{P}' is an $N \times 3$ matrix with the i -th row as $\mathbf{X}'_i{}^T$. The parameter ρ balances smoothness of the RBF against fidelity to the data.

With the computed parameters, the inverse gamut mapping at any point ($h^{-1}(\mathbf{X})$) is evaluated by Eq. 10 (Fig. 4). The overall performance of the RBF relies on the selection of the control points. While we could use all possible points from the training data as control points, this would be inefficient since the evaluation time grows with the number of control points. Additionally, a larger number of control points could also lead to over-fitting. We instead use a greedy algorithm similar to the one used in [31] to select a small subset of control points from a large number of available points that maintains the desired accuracy. The number of control points used in this work varies from 3000 to 5000. As previously mentioned, the gamut mapping function h is computed per camera picture style and the training data set for each picture style contains images taken from all the white balance settings. Having data from different white balance settings is necessary to have the color points well distributed throughout the color space in the training data. In most of our experiments, we use 70 image pairs per picture style for the training: 7 different white balance settings with 10 RAW-sRGB pairs per each setting.

5.4 Calibrating Cameras without RAW support

Thus far, computing the color transformations \mathbf{T}_w , \mathbf{T}_s , and h relied on having the associated RAW image for each sRGB image in the training set. However, there are many cameras that do not provide RAW images, especially point-and-shoot cameras. Therefore, a calibration scheme for cameras without RAW support is necessary to broaden the applicability of our work.

For those cameras without RAW support, we use a RAW image of the same scene from another camera as a reference. In this case, Eq. 7 changes to:

$$\begin{bmatrix} e_{rx} \\ e_{gx} \\ e_{bx} \end{bmatrix} = h(\mathbf{T}_s \mathbf{T}_w \mathbf{T}_c \mathbf{E}'_x). \quad (12)$$

The 3×3 matrix \mathbf{T}_c is a transformation that approximates the transformation between the color space of two different cameras. \mathbf{E}'_x contains the RAW values given by the other camera. For cameras without RAW images, the different color transformations are combined into one transformation ($\mathbf{T}_{z_i} = \mathbf{T}_s \mathbf{T}_w \mathbf{T}_c$), which is computed as the one that minimizes the following error:

$$\sum_{i=1}^M \sum_{j=1}^N \|\mathbf{T}_{z_i}^{-1} \mathbf{X}_{ij} - \mathbf{E}'_{ij}\|^2. \quad (13)$$

After computing the \mathbf{T}_{z_i} 's, the gamut mapping function h is computed just as explained in Section 5.3. While an image of a camera cannot be converted back to its own RAW image with this approach, it can still be transformed to sRGB images with different settings as described in the next section.

6 CAMERA SPECIFIC IMAGE TRANSFORMATION

One benefit of our calibration procedure is that we inherently have camera-specific photofinishing information pertaining to different white-balance and picture style settings. This allows us to not only revert an sRGB photo back to RAW, but we can reapply the processing pipeline to refinish

a photo. An excellent example of when this is useful is when a photo has been taken under the wrong settings.

Fig. 3 illustrates the procedure for transferring colors between different settings. Given an input image \mathbf{I} taken under a white balance w_i and a picture style p_i , a new image \mathbf{I}' under a new white balance w_o and a new picture style p_o can be generated with k times the exposure of the original as follows:

$$\begin{bmatrix} I'_r \\ I'_g \\ I'_b \end{bmatrix} = \begin{bmatrix} f_{r,p_o}(e'_{rx}) \\ f_{g,p_o}(e'_{gx}) \\ f_{b,p_o}(e'_{bx}) \end{bmatrix}, \text{ where} \quad (14)$$

$$\begin{bmatrix} e'_{rx} \\ e'_{gx} \\ e'_{bx} \end{bmatrix} = h_{p_o} \left(\mathbf{T}_s \mathbf{T}_{w_o} k \mathbf{T}_{w_i}^{-1} \mathbf{T}_s^{-1} h_{p_i}^{-1} \left(\begin{bmatrix} f_{r,p_i}^{-1}(I_{rx}) \\ f_{g,p_i}^{-1}(I_{gx}) \\ f_{b,p_i}^{-1}(I_{bx}) \end{bmatrix} \right) \right).$$

6.1 Manual Mode

Frequently the wrong settings that ruin a photo are manually set by the user, in many cases by mistake. For those images taken under a manual mode, the input settings for white balance (w_i), picture style (p_i), and exposure can all be read from the EXIF data of the input image. The user then only has to select the exposure and the choices for white balance and picture style among the camera presets to correct an image. This correction procedure is intuitive because the user chooses the output settings just as one would do when using a camera.

We also provide a feature which enables the user to change the white balance setting of the output in a continuous manner rather than just selecting from preset options. The white balance parameters (diagonal elements of \mathbf{T}_w) are associated with color temperature and thus can be ordered, e.g. tungsten (3200K), fluorescent (4000K), daylight (5200K), and cloudy (6000K). The output white balance parameters (\mathbf{T}_{w_o}) in Eq. 14 could then be computed by linear interpolation with respect to either a user supplied color temperature or user scrolling between preset white balance settings.

6.2 Auto White Balance Mode

In some cases, one may not like a photograph taken under the camera's auto white balance mode and wish to correct it. The problem with auto-mode images is that it is difficult to recover the specific settings of the camera from the EXIF data. Therefore, we cannot determine which white balance (w_i) and picture style (p_i) to use for Eq. 14. For the auto-mode case, we rely on user assistance to convert an image to another setting. The user can either set the input or the output setting as he wishes and then tune the other settings until he is satisfied with the final output image. The user can choose any of the available picture styles for the camera and change the white balance setting in a continuous manner using interpolation as explained previously.

6.3 Camera-to-Camera Transfer

Thus far, we have described how to transform an image to another image under different settings but from the same camera. We can extend our framework to transfer color between different cameras and their settings. One can imagine such a feature being useful for many applications. For example, it could be used to compare color differences between cameras to guide a person planning to purchase a new camera. It can also be used to align colors of images from different cameras to create seamless mosaics and texture maps.

While the information on sensor spectral sensitivity of the cameras is necessary to accurately compute camera-to-camera color transfers, we approximate the color space transformation between the color spaces of two cameras by a 3×3 matrix \mathbf{T}_c . The matrix \mathbf{T}_c is computed using two aligned RAW images ($\mathbf{E}_1, \mathbf{E}_2$) of the same scene, one for each camera:

$$\mathbf{T}_c = \underset{\mathbf{T}}{\operatorname{argmin}} \sum_x \|\mathbf{E}_{2x} - \mathbf{T}\mathbf{E}_{1x}\|_2 \quad (15)$$

Then the color transfer between cameras is achieved similar to Eq. 14:

$$\begin{bmatrix} e'_{rx} \\ e'_{gx} \\ e'_{bx} \end{bmatrix} = h_{p_o} \left(\mathbf{T}_s \mathbf{T}_{w_o} \mathbf{T}_c k \mathbf{T}_{w_i}^{-1} \mathbf{T}_s^{-1} h_{p_i}^{-1} \mathbf{z} \right), \quad (16)$$

$$\mathbf{z} = \begin{bmatrix} f_{r,p_i}^{-1}(I_{rx}) \\ f_{g,p_i}^{-1}(I_{gx}) \\ f_{b,p_i}^{-1}(I_{bx}) \end{bmatrix}.$$

The transfer matrix \mathbf{T}_c between two cameras can also be computed via transformations to a reference camera: $\mathbf{T}_{c,1 \rightarrow 3} = \mathbf{T}_{c,1 \rightarrow 2} \mathbf{T}_{c,2 \rightarrow 3}$. Note that \mathbf{T}_c is inherently included in \mathbf{T}_z in Section 5.4 and transferring color between cameras that do not support RAW is not a problem.

7 EXPERIMENTAL RESULTS

7.1 Radiometric Response Function Estimation

We first compare the performance of the response function estimation (Section 5.1) against the conventional approach [4] upon which we have built our algorithm.

Fig. 5 shows an example of the outliers detected by our algorithm and the response functions recovered by the two methods. Note that the only difference between the two methods is the existence of the outlier removal procedure. There is a significant difference in the estimations and the proposed algorithm for removing the outliers clearly outperforms on the linearization results.

A few selected inverse response functions computed using our algorithm for some cameras in our database are shown in Fig. 6. Note that the responses differ from the sRGB gamma curve commonly used for linearization in some color vision work. For a quantitative evaluation of the response estimation, we use the following measure per channel to gauge the accuracy of linearization from Eq. 4:

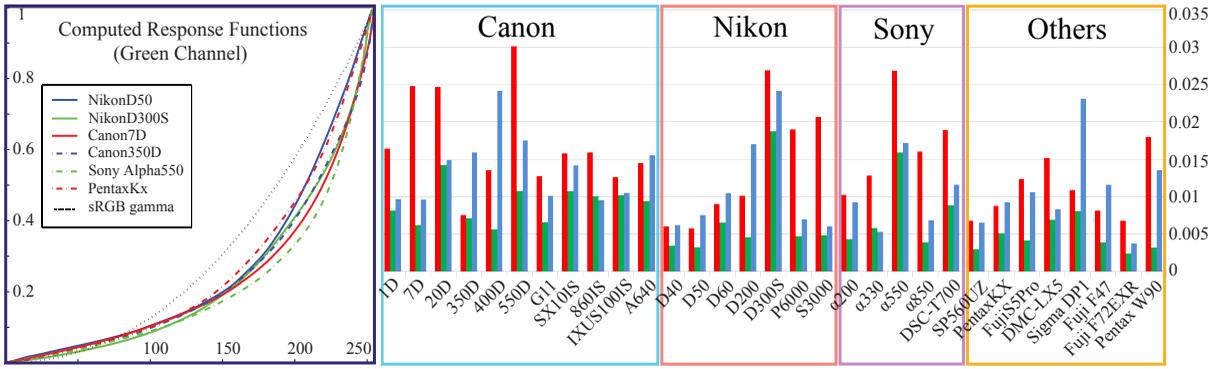


Fig. 6: Inverse radiometric response functions for a set of cameras in the database and mean linearization errors (δ in Eq. 17) for all cameras in the database. High errors indicate the degree of color mapping in the in-camera image processing. Images from cameras with high errors will not be linearized accurately with the conventional radiometric model and calibration, hence a new imaging model is necessary. Using the color mapping (h), the linearization errors converge very close to zero (the maximum among the cameras in the database above is 0.004). Using the The bar colors indicate different color channels.

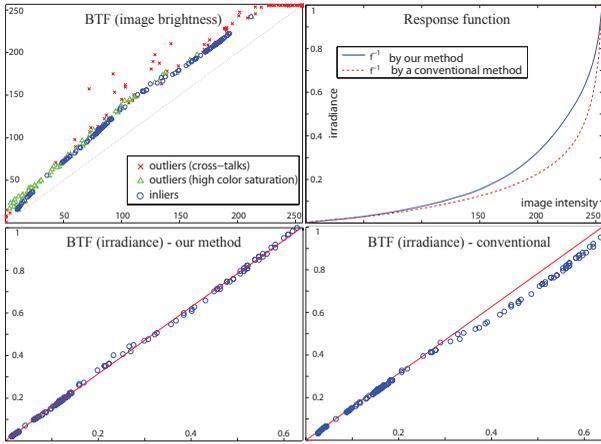


Fig. 5: A BTF, estimated response function, and linearization results for the blue channel of Nikon D40 using our radiometric calibration algorithm with outliers removed and a conventional method [4]. With our method, the outliers are effectively removed for more accurate calibration.

$$\delta_c = \sqrt{\frac{\sum_{n=1}^N \sum_{x \in A} \|k'_n f^{-1}(i_{cx}^n) - f^{-1}(i'_{cx})\|^2}{N|A|}}, \quad (17)$$

where N is the number of image pairs, A is the set of all image points, and $|A|$ is the size of the set A . To compute δ for each camera, we use all available sets of images in the database for the particular camera, not just the ones used for calibration. This is to verify that a response function computed under a specific condition can be used to accurately linearize images captured under different settings such as the lighting condition and the white balance setting.

Fig. 6 plots the δ 's for all cameras in the database. We can see that for many cameras in our database, the image can be linearized very well with an average error of less than 1%. Note that outliers were included for the statistics in Fig. 6. If we exclude outliers from the computation, δ converges almost to 0 in many cameras. So the δ in Fig. 6 is related to the amount of outliers, or the degree of color mapping h in

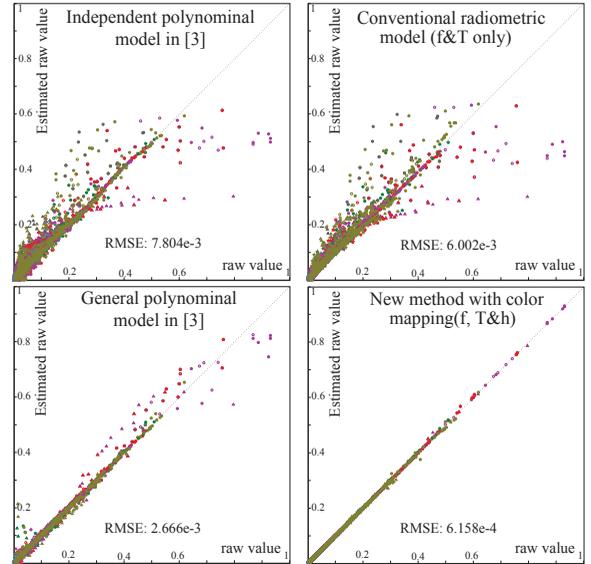


Fig. 7: Performance of mapping image values to RAW values (Canon EOS-1D) with different techniques: using the technique in [3] with the independent polynomial model per channel, using f and T in Eq. 7 without h , the all-channel 3D polynomial model in [3], and the new method with h . Using our new model, images can be mapped back to RAW accurately.

the in-camera image processing. For the cameras with high δ 's, the gamut mapping is applied to points well within the sRGB gamut as opposed to other cameras where it applies only to points close to the boundary of the gamut.

7.2 Color Mapping Function Estimation

Next, we evaluate the performance of the color mapping function (h) estimation and the overall accuracy of the new imaging model (Eq. 7). The 3D color mapping functions (h) for the Nikon D40 and the Canon EOS-1D are shown as slices in Fig. 4. The results confirm the existence of gamut mapping in the in-camera imaging process and the need to include the color mapping function in the radiometric model. The performance of our new imaging model and its calibration procedures for converting image values to RAW

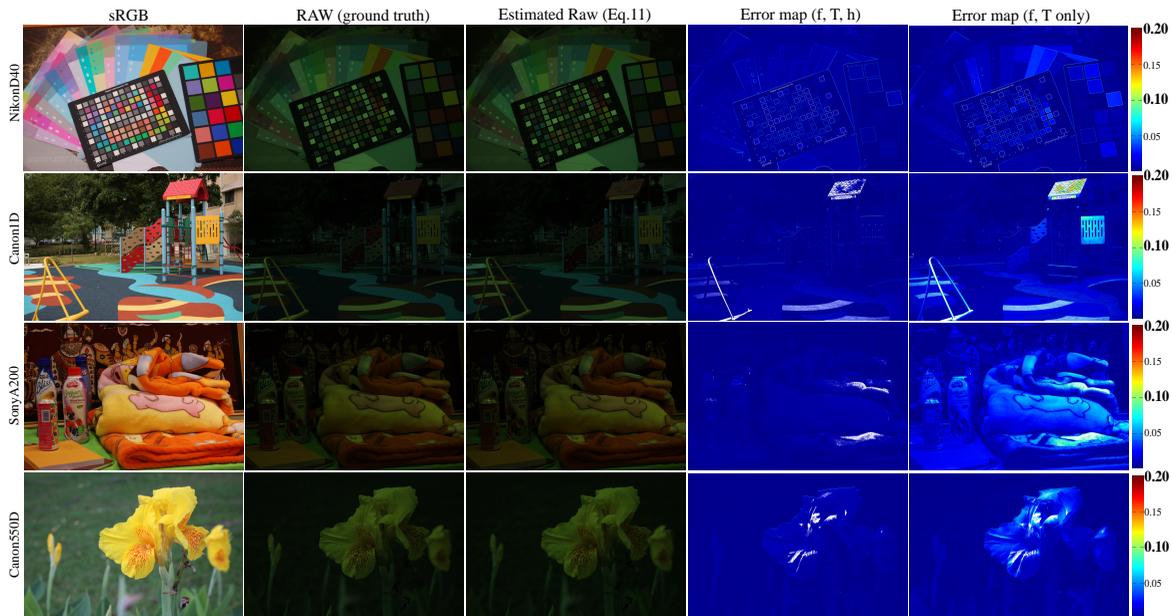


Fig. 8: Mapping images to RAW. Our method for mapping images to RAWs works well for various cameras and scenes. The white points on the difference maps indicate pixels with a value of 255 in any of the channels which are impossible to linearize. The RMSE's for the new method and the conventional method from the top to the bottom are (0.006, 0.008), (0.004, 0.010), (0.003, 0.017), and (0.003, 0.007). Notice that the errors are high in edges due to demosaicing. For Nikon cameras, the difference in performance between the new method and the conventional method is not as big as other cameras because the effect of the gamut mapping is not as severe as the others (see Fig. 7 (a)).

is shown in Fig. 7. In the figure, we compare the results from four different techniques given a number of sRGB-RAW image pairs. The first method is the implementation of the algorithm from [3] where f is modeled as a 6th order polynomial per channel. The second method computes the RAW just from f and \mathbf{T} , which are computed as described in Section 5 without the color mapping function h . Next, we computed f as a 3D polynomial function as described in [3]. Finally, the last method computes RAW from Eq. 7 with the color mapping function included. As can be seen, the image can be mapped backed to RAW accurately by including the color mapping function in the radiometric model and approximating the mapping function with radial basis functions. In addition, our results show that in-camera color manipulation introduces nonlinearities that cannot be sufficiently modeled by a 3D polynomial function [3].

Fig. 8 shows the results of applying the calibration results to convert images of real scenes back to RAW responses for various scenes and cameras. The estimates of RAW images are compared with the ground truth RAW images. Note that the estimates are purely from the pre-calibrated values of f , h , and \mathbf{T} and the ground truth RAW images are used only for evaluation purposes. Using the new model and the calibration algorithms introduced in Section 5, we can accurately convert the image values to RAW values even for the highly saturated colors in the scene.

7.3 Refinishing Examples

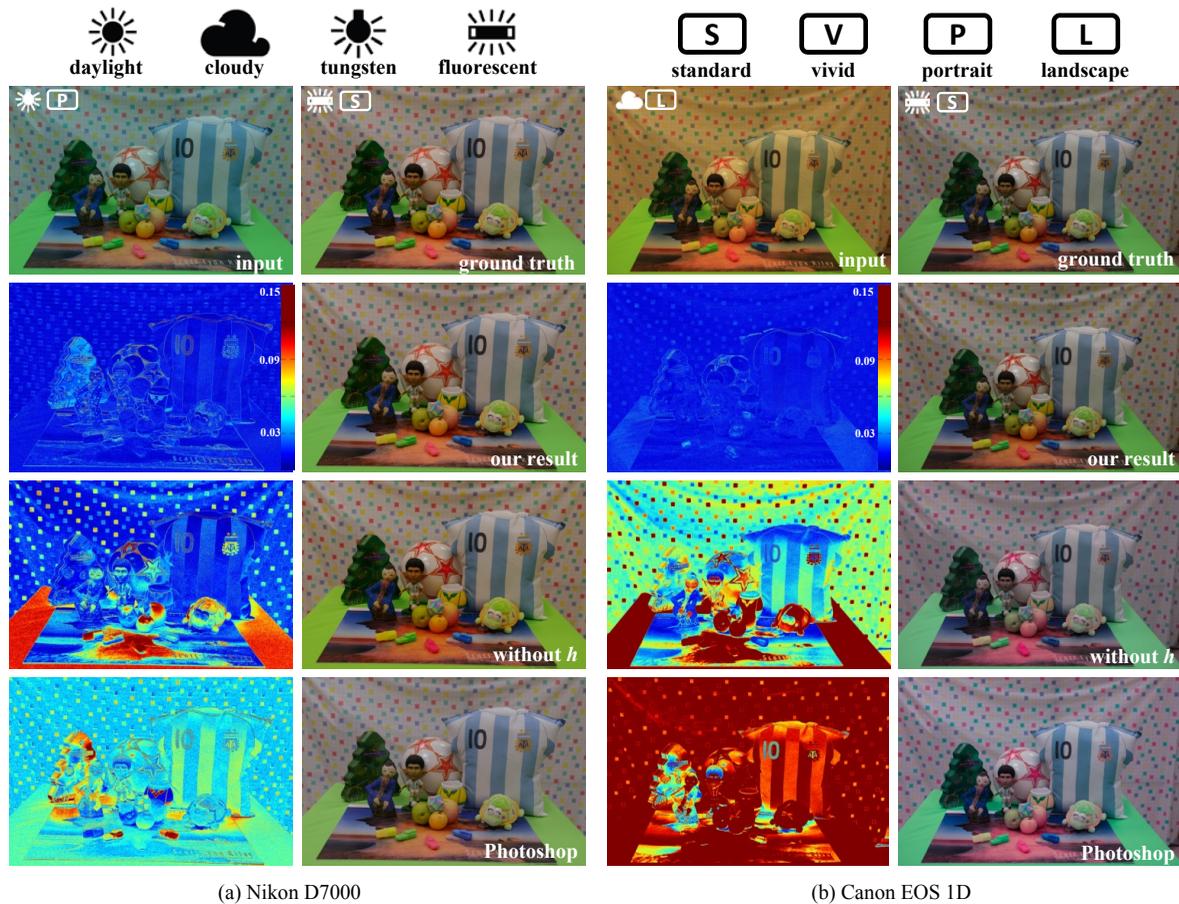
Here we show the ability of our approach to refinish photographs using the extracted parameter settings. For the sake of comparisons, we compare our method with

Photoshop and a version of our method without gamut mapping (no h in Eq. 14) in Fig. 9.

For the Photoshop results, we use the Camera RAW utility and choose the best result either from the auto white balancing feature or the semi-auto feature in which we chose a point in the image to be white. As can be seen from the error maps, our photo refinishing technique can transfer colors between different settings accurately, therefore provide a practical method to correct undesired visual errors in photographs taken with the wrong camera settings. Meanwhile, the other two methods have difficulties dealing with the nonlinearities in the imaging process. This is especially visible in the Canon's Landscape mode which is shown to have greater nonlinearity in Fig. 4. More examples for different cameras are shown in Fig. 10, and results for point-and-shoot cameras that do not support RAW (Section 5.4) are shown in Fig. 11. Both cameras used for Fig. 11 were calibrated using a RAW image from a Canon EOS1-D, and the results closely approximate the ground truth.

We should note that correction using Photoshop can yield visually satisfactory results as can be seen from the second example of Fig. 11. However, the quality of Photoshop results is unreliable since it can vary greatly depending on scene (e.g. distribution of color, especially white and gray colors) and camera settings. Furthermore, we have found that Photoshop almost never reproduces accurate camera specific images as our technique does.

Next, we show examples of transferring colors between cameras. In Fig. 12, three images of a scene were taken each with different cameras, namely a Canon EOS-1D,



(a) Nikon D7000

(b) Canon EOS 1D

Fig. 9: Comparisons of different methods for correcting input images taken under inappropriate settings (WB, picture style). Our photo refining technique recovers images that are very close to the images from the cameras themselves (ground truth) while the technique without consideration of gamut mapping h and the Photoshop method do not effectively deal with the nonlinearities in the imaging process. The scale for the error maps is the same for all the error maps shown. The RMSE's for the new method, the conventional method, and the Photoshop are (a) (0.02, 0.05, 0.06) and (b) (0.02, 0.1, 0.18).

Sony α -200, and Nikon D40. All the cameras were under fluorescent white balance and the standard picture style. Images from these cameras exhibit differences in color, notably the color of the face and the balloon in the middle. The second and the third rows of Fig. 12 are the results of transferring camera colors to the Nikon and Sony, respectively. As can be seen, the colors from different cameras can be transferred and matched accurately using our framework (Section 6.3).

As the last example, we show the result of photofinishing an image taken under auto white balance. As mentioned earlier, when a photograph is taken under the auto mode, the input white balance T_{wi} is unknown and the system relies on user provided information on the unknown parameter. In Fig. 13, the user presumes that the image was taken under “daylight” and the system produces images under different white balance settings. In the end, using a slightly warmer color temperature than daylight provides an image more satisfying than the auto white balanced image (Fig. 13-d).

Our system is implemented in C++ and we have two implementations for evaluating the RBF gamut mapping function h (Eq. 10). One implementation evaluates the RBFs of each image on the fly and takes 30 seconds on

average to compute the color transfer, which includes two RBF evaluations, one backward (h^{-1}) and one forward (h). The running time of this implementation can be shortened by using a fast RBF evaluation method as in [32]. The other implementation is based on lookup tables which saves computation time while increasing the amount of memory usage. The color transfer in this paper only depends on the color values (RGB) of each point in the image and is therefore a deterministic process. This allows lookup tables to be built for both the forward and the inverse process by sampling the RAW and the sRGB color space and precomputing the color transfers for each of the sampled points. With the lookup tables, photo refining takes less than a second. More results, as well as the database can be found at www.comp.nus.edu.sg/~brown/radiometric_calibration.

8 DISCUSSION AND FUTURE WORK

We have presented a study of the in-camera image processing through an extensive analysis of a large image database. One of the key contributions of this paper is identifying the need for a color (gamut) mapping step in the in-camera image processing model. The inclusion of this step covers the

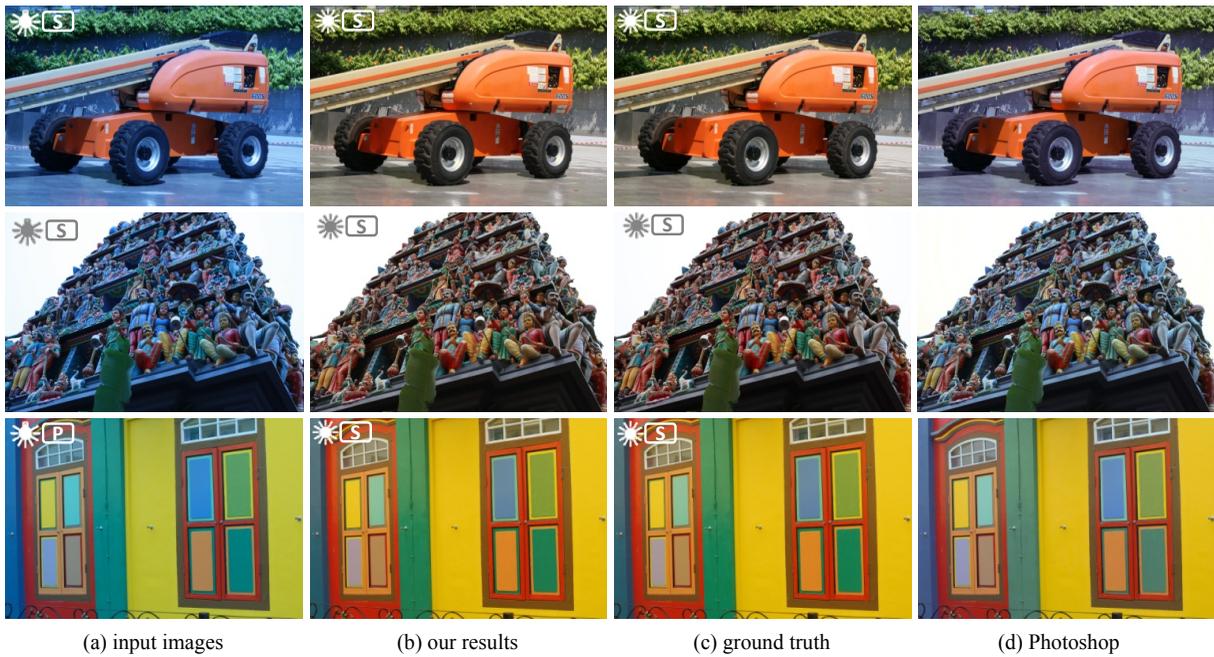


Fig. 10: More examples of our photo refinishing using images from a Sony α -200, Canon EOS-1D, and Nikon D200 (from top to bottom). The ground truth images are actual images from the cameras themselves under the proper settings.

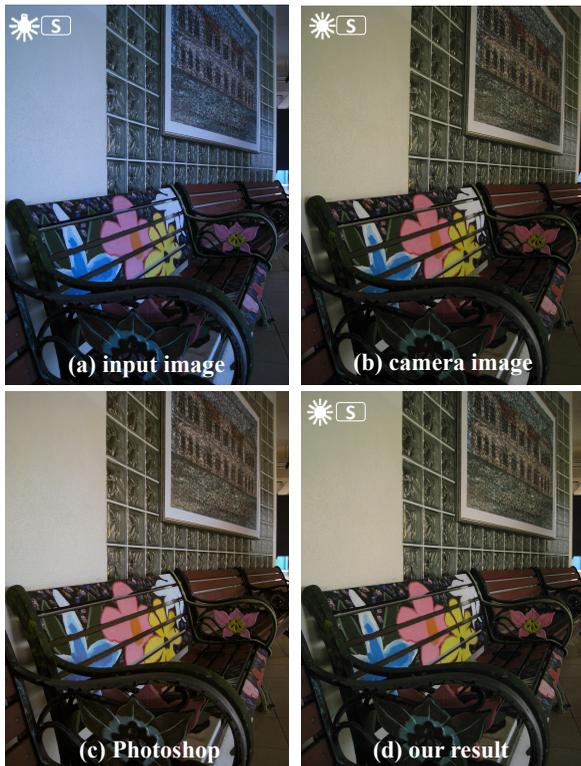


Fig. 11: Photo refinishing result for a camera (Canon IXUS 860IS) without the RAW support (see Section 5.4).

limitations present in the conventional imaging model and calibration methods. By considering color mapping in the imaging process, we can not only compute the radiometric response function more accurate than previous approach, but we can also convert a given sRGB image to RAW

using our calibration scheme. Based on the new model, we further introduced a new framework for refinishing photos, which enables one to correct photographs taken with wrong settings without the associated RAW files.

For the calibration, we relied on a simple assumption that the colors with low saturation are not affected by the gamut mapping. With this assumption, the response function and the linear color space transformations were first computed by using the points filtered by a threshold (β) on the color saturation level. While this simple approach provided satisfying results for our rather controlled dataset (color charts), a more robust approach based on an iterative scheme may reduce the reliance on a hard threshold for more general databases of images.

Note that the color gamut mapping function h may not be invertible depending on the gamut mapping method employed by the camera. For instance, many color points will be mapped to a single color if a camera employs a clipping strategy. However, we rarely observed such instances in our experiments (about 0.2% of total observations). When such instances occurred, we chose the median value as the control point to approximate the inverse.

While we estimate f and h separately during the calibration, one could also consider combining the two functions into a single $\mathbb{R}^3 \rightarrow \mathbb{R}^3$ function that directly maps white balanced RAW values to nonlinear sRGB values. In principle, the radial basis functions should be able to model this, however, in our experiments we obtained better results when we used two separate functions. Our intuition is that this initial linearization of the RGB space by the function f reduces the complexity of the color mapping function h . This allows h to appear smoother and require less control points for the scatter point interpolation. Modeling

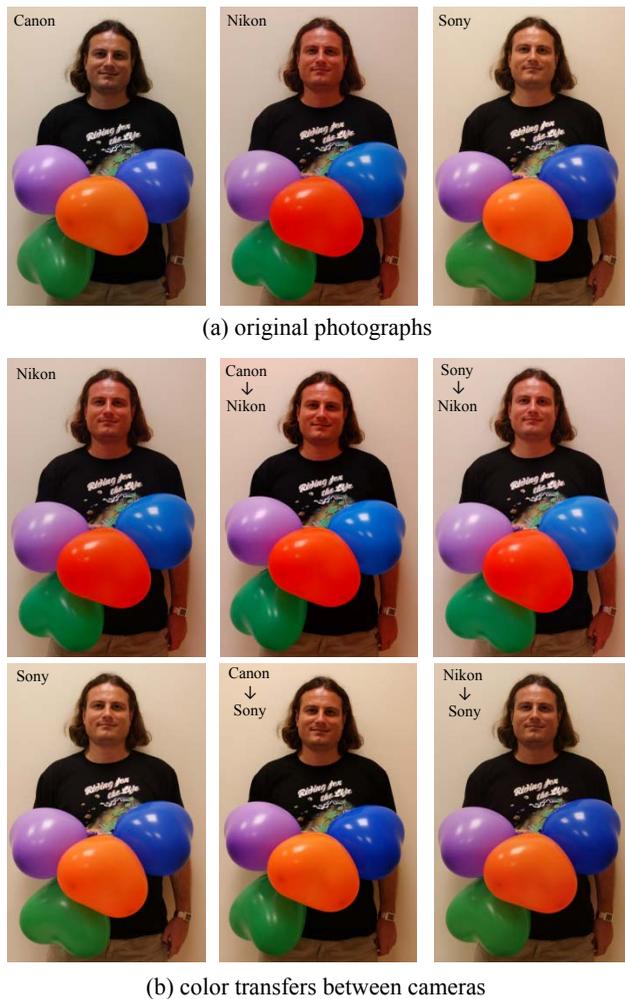


Fig. 12: Transferring colors between cameras. (a) Original photographs from three cameras (Canon EOS-1D, Sony α -200, and Nikon D40) display varying colors. (b) Colors from different cameras can be matched by using the method described in Section 6.3.

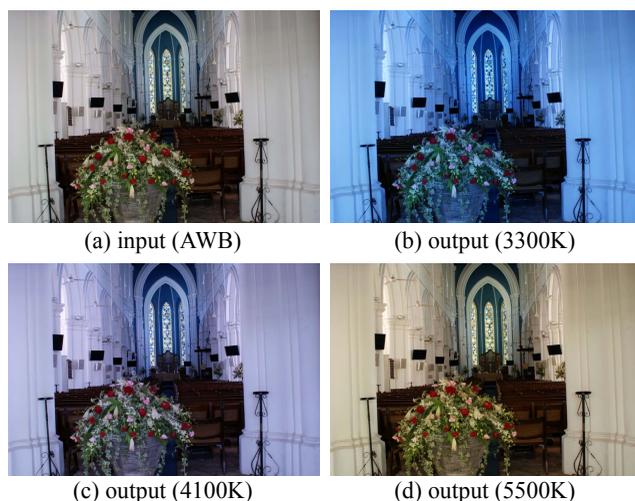


Fig. 13: (a) Input image taken under auto white balance. (b)-(d) The user specifies daylight as the input WB and changes the output WB to different color temperatures. Through this process, our system can produce an output image with warmer whites in comparison to the input image.

f separately from h is also desirable since f can be still computed from multiple images and used for linearization when RAW images are not available.

Recall that the underlying assumption for this work is that cameras are operating under the photographic reproduction mode, which can be achieved by capturing images in the manual mode and turning off features for scene dependent rendering. We did, however, show an instance of dealing with images taken under auto white balance with the help of user-assistance. In the future, we plan to investigate what and how much scene dependent processing is done in images under the photofinishing mode. The analysis on the photofinishing mode together with the analysis done in this paper will suggest a direction for the internet color vision research [3], [33], [11], [24] in the future. We also plan to extend our framework to seek a calibration method that does not require RAW images and to model cameras outside our calibrated database.

Acknowledgement

This work was supported in part by the NUS Young Investigator Award, R-252-000-379- 101.

REFERENCES

- [1] J. Holm, I. Tastl, L. Hanlon, and P. Hubel, "Color processing for digital photography," in *Colour Engineering: Achieving Device Independent Colour*, P. Green and L. MacDonald, Eds. Wiley, 2002, pp. 79 – 220.
- [2] H. Lin, S. J. Kim, S. Süsstrunk, and M. S. Brown, "Revisiting radiometric calibration for color computer vision," in *Proc. International Conference on Computer Vision*, 2011.
- [3] A. Chakrabarti, D. Scharstein, and T. Zickler, "An empirical camera model for internet color vision," in *Proc. British Machine Vision Conference*, 2009.
- [4] M. Grossberg and S. Nayar, "Modeling the space of camera response functions," *IEEE Transaction on Pattern Analysis and Machine Intelligence*, vol. 26, no. 10, pp. 1272–1282, 2004.
- [5] S. Lin and L. Zhang, "Determining the radiometric response function from a single grayscale image," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2005, pp. 66–73.
- [6] S. Mann and R. Picard, "On being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures," in *Proc. IS&T 46th annual conference*, 1995, pp. 422–428.
- [7] T. Mitsunaga and S. Nayar, "Radiometric self-calibration," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 1999, pp. 374–380.
- [8] P. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," in *Proceedings of SIGGRAPH*, 1997, pp. 369–378.
- [9] M. Grossberg and S. Nayar, "Determining the camera response from images: What is knowable?" *IEEE Transaction on Pattern Analysis and Machine Intelligence*, vol. 25, no. 11, pp. 1455–1467, 2003.
- [10] S. J. Kim and M. Pollefeys, "Robust radiometric calibration and vignetting correction," *IEEE Transaction on Pattern Analysis and Machine Intelligence*, vol. 30, no. 4, pp. 562–576, 2008.
- [11] S. Kuthirummal, A. Agarwala, D. Goldman, and S. Nayar, "Priors for large photo collections and what they reveal about cameras," in *Proc. European Conference on Computer Vision*, 2008, pp. 74–86.
- [12] C. Pal, R. Szeliski, M. Uyttendaele, and N. Jovic, "Probability models for high dynamic range imaging," in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition*, 2004, pp. 173–180.
- [13] S. Lin, J. Gu, S. Yamazaki, and H. Shum, "Radiometric calibration from a single image," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2004, pp. 938–945.
- [14] K. Barnard, V. Cardei, and B. Funt, "A comparison of computational color constancy algorithmspart i: Methodology and experiments with synthesized data," *IEEE Transaction on Image Processing*, vol. 11, no. 9, pp. 972–984, 2002.

- [15] G. Buchsbaum, "A spatial processor model for object colour perception," *Journal of The Franklin Institute-engineering and Applied Mathematics*, vol. 310, pp. 1–26, 1980.
- [16] E. Land and J. McCann, "Lightness and retinex theory," *J. Opt. Soc. Am.*, vol. 61, no. 1, pp. 1–11, 1971.
- [17] D. Forsyth, "A novel algorithm for color constancy," *International Journal of Computer Vision*, vol. 5, no. 1, pp. 5–36, 1990.
- [18] G. D. Finlayson and S. D. Hordley, "Improving gamut mapping color constancy," *IEEE Transaction on Image Processing*, vol. 9, no. 10, pp. 1774–1783, 2000.
- [19] D. Brainard and W. Freeman, "Bayesian color constancy," *J. Opt., Soc. Am.*, vol. 14, no. 7, pp. 1393–1411, 1997.
- [20] A. Gijsenij and T. Gevers, "Color constancy using natural image statistics," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2007, pp. 1–8.
- [21] E. Reinhard, M. Ashikhmin, B. Gooch, and P. Shirley, "Color transfer between images," *IEEE Comput. Graph. Appl.*, vol. 21, pp. 34–41, 2001.
- [22] D. Cohen-Or, O. Sorkine, R. Gal, T. Leyvand, and Y.-Q. Xu, "Color harmonization," *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, vol. 25, no. 3, pp. 624–630, 2006.
- [23] S. Bae, S. Paris, and F. Durand, "Two-scale tone management for photographic look," *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, vol. 25, pp. 637–645, 2006.
- [24] J.-F. Lalonde, A. A. Efros, and S. G. Narasimhan, "Webcam clip art: Appearance and illuminant transfer from time-lapse sequences," *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, vol. 28, no. 5, pp. 131:1–131:10, 2009.
- [25] L. Shapira, A. Shamir, and D. Cohen-Or, "Image appearance exploration by model-based navigation," *Comput. Graph. Forum*, vol. 28, no. 2, pp. 629–638, 2009.
- [26] B. Wang, Y. Yu, T.-T. Wong, C. Chen, and Y.-Q. Xu, "Data-driven image color theme enhancement," *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, vol. 29, no. 6, pp. 146:1–146:10, December 2010.
- [27] P. D. Hiscocks, "Measuring camera shutter speed," 2010, <http://www.syscompdesign.com/AppNotes/shutter-cal.pdf>.
- [28] ISO 22028-1:2004, "Photography and graphic technology - extended colour encodings for digital image storage, manipulation and interchange - Part 1: architecture and requirements," *International Organization for Standardization*, 2004.
- [29] J. Morovic and M. R. Luo, "The fundamentals of gamut mapping: A survey," *Journal of Imaging Science and Technology*, vol. 45, no. 3, pp. 283–290, 2001.
- [30] M. D. Buhmann, *Radial Basis Functions: Theory and Implementations*. Cambridge University Press, 2003.
- [31] J. C. Carr, R. K. Beatson, J. B. Cherrie, T. J. Mitchell, W. R. Fright, and B. C. McCallum, "Reconstruction and representation of 3d objects with radial basis functions," in *Proceedings of SIGGRAPH*, 2001, pp. 67–76.
- [32] L. Greengard and V. Rokhlin, "A fast algorithm for particle simulations," *J. Comput. Phys.*, vol. 73, pp. 325–348, 1987.
- [33] T. Haber, C. Fuchs, P. Bekaert, H.-P. Seidel, M. Goesele, and H. Lensch, "Relighting objects from image collections," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2008, pp. 1–8.

Hai Ting Lin received the B. E. degree in computer science from the Renmin University of China in 2008. He is pursuing the PhD degree at the National University of Singapore. His research interests include image processing, computer vision. He is a student member of the IEEE.



Zheng Lu received the B.Comp degree in Computer Science from the National University of Singapore in 2004. He is pursuing the PhD degree at the National University of Singapore. His research interests include computer vision and image/video processing. He is a student member of the IEEE.



Sabine Süsstrunk leads the Images and Visual Representation Group (IVRG) in the School of Computer and Communication Sciences (IC) at EPFL since 1999. Her main research areas are in computational photography, color imaging, image quality metrics, image indexing, and archiving. She has authored and co-authored over 100 publications and holds 5 patents. She was Associate Editor for the IEEE Transactions on Image Processing from 2007-2011. Since 2011, she is Conference Vice-President for IS&T.

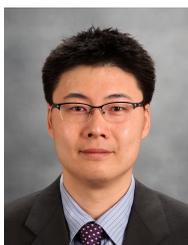


Stephen Lin is a Senior Researcher in the Internet Graphics Group of Microsoft Research Asia. He obtained a B.S.E. from Princeton University and a Ph.D. from the University of Michigan. His research interests include computer vision and computer graphics. Dr. Lin has served as a Program Co-Chair for the IEEE International Conference on Computer Vision 2011 and the Pacific-Rim Symposium on Image and Video Technology 2009, and as a General Co-Chair for



the IEEE Workshop on Color and Photometric Methods in Computer Vision 2003.

Seon Joo Kim received B.S. and M.S. from Yonsei University, Seoul, Korea, in 1997 and 2001. He received Ph.D. in computer science from University of North Carolina at Chapel Hill in 2008. He currently holds a joint appointment as an assistant professor at SUNY Korea and a research scientist at CEWIT Korea. His research interests include computer vision, computer graphics/computational photography, and HCI/visualization.



Michael S. Brown obtained his BS and PhD in Computer Science from the University of Kentucky in 1995 and 2001 respectively. He is currently an associate professor in the School of Computing at the National University of Singapore. Dr. Brown has served as an area chair for CVPR'09, ACCV'10, CVPR'11, ICCV'11, and ECCV'12 and is an associate editor for IEEE TPAMI. His research interests include computer vision, image processing and computer graphics.

