

Revisiting Radiometric Calibration for Color Computer Vision

Haiting Lin¹ Seon Joo Kim^{1,2} Sabine Süsstrunk³ Michael S. Brown¹

¹National University of Singapore

²UIUC Advanced Digital Sciences Center, Singapore

³École Polytechnique Fédérale de Lausanne

Abstract

We present a study of radiometric calibration and the in-camera imaging process through an extensive analysis of more than 10,000 images from over 30 cameras. The 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 show that the conventional radiometric model fits well for image pixels with low color saturation but begins to degrade as color saturation level increases. This is due to the color mapping process which includes gamut mapping in the in-camera processing that cannot be modeled with conventional methods. To this end, we introduce a new imaging model for radiometric calibration and present an effective calibration scheme that allows us to compensate for the nonlinear color correction to convert non-linear sRGB images to CCD RAW responses.

1. Introduction

Many computer vision algorithms assume that cameras are accurate light measuring devices that 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.

We present a study of radiometric calibration and 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 number of factors that cause instability in the camera response function computation and use the findings to

present a practical radiometric algorithm that enhances the overall accuracy. More importantly, the analysis shows a limitation of the imaging model employed in conventional radiometric calibration methods when dealing with pixels with high color saturation levels. In particular, conventional radiometric models cannot explain the color mapping component which includes gamut mapping [8] in the imaging pipeline. This makes it difficult to convert points with high color saturation to physically meaningful values. To address this limitation, we introduce a new imaging model for radiometric calibration and propose an effective calibration procedure that allows us to compensate for this color correction to convert non-linear sRGB images to CCD's RAW responses.

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. This radiometric mapping 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 [5, 14]. 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.

Related Work. 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 (\mathbf{I}, \mathbf{I}') 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:

$$\frac{f^{-1}(i'_x)}{f^{-1}(i_x)} = k'. \quad (2)$$

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 [15], polynomial [16], non-parametric [3], and PCA based model [4]. 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 [18] where the response function was modeled differently per image using a probabilistic approach. Another exception is the recent work in [2] 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 as follows:

$$\mathbf{I}_x = \begin{bmatrix} i_{rx} \\ i_{gx} \\ i_{bx} \end{bmatrix} = \begin{bmatrix} f_r(e_{rx}) \\ f_g(e_{gx}) \\ f_b(e_{bx}) \end{bmatrix}, \quad (3)$$

$$\mathbf{e}_x = \begin{bmatrix} e_{rx} \\ e_{gx} \\ e_{bx} \end{bmatrix} = \mathbf{T}\mathbf{E}_x. \quad (4)$$

The term \mathbf{E}_x is the irradiance value captured as RAW, \mathbf{T} is a 3×3 transformation, and f is modeled with 5th order polynomial per channel (r , g , and b). The difference in their model compared to the conventional model (Eq. 1) is that they explain color by including the color transformation \mathbf{T} , which is responsible for both the transformation from the camera's color space to sRGB and white balancing. Therefore, the image values can be inverted back to the actual irradiance value at CCD (RAW). Additionally, the function f is explained as including the color rendering process in addition to the compressive nonlinearity. Using the available RAW data from the cameras, they iteratively compute the color transformation \mathbf{T} and the responses f that map the RAW data to the output sRGB image. 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.

Scene Dependency and Camera Settings 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 [2], traditional radiometric calibration would be inadequate and the only option would be to use single-image based radiometric calibration methods [13, 14]. 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 [8]. The digital camera community defines color rendering as the operations that apply the tone/color reproduction aims for the imaging system and change the state of the image from a scene-referred image state to an output-referred image state [9]. Color rendering transforms may include tone and gamut mapping to account for the dynamic range and color gamut of the output color image encoding (e.g., sRGB, Adobe RGB), compensation for differences in the input- and output-viewing conditions, and other color adjustments (e.g., selective hue, saturation) to account for color reproduction preferences of the human observer.

The intent of the “photofinishing” color rendering operations is to produce a pleasing image that is not solely dependent on the exposure received by the image sensor. In this model, the imaging pipeline varies the color rendering based on the captured scene, possibly in a spatially varying manner. Some examples of camera settings that enable photofinishing are Dynamic Lighting Optimizer on the Canon EOS550D and D-Range Optimizer in Sony α -200. Different photofinishing methods can also be associated with “scene modes”, e.g. Portrait, Landscape, Night-time, etc. For images produced using such scene dependent processing it is very difficult to convert image values to physically meaningful values.

The photographic reproduction model, on the other hand, uses fixed color rendering. This mode is intended to allow the digital camera to mimic traditional film cameras and targets professional photographers [8]. For most high-end cameras, it is possible to set the camera in this photographic mode by turning the camera settings to manual and turn off all settings pertaining to photofinishing. However, for cheaper “point-and-shoot” cameras, it should be noted that this may not be possible. The implications of this are discussed in Section 8, and for the remaining of this paper it is assumed that the algorithms discussed are intended to work in the photographic reproduction mode.

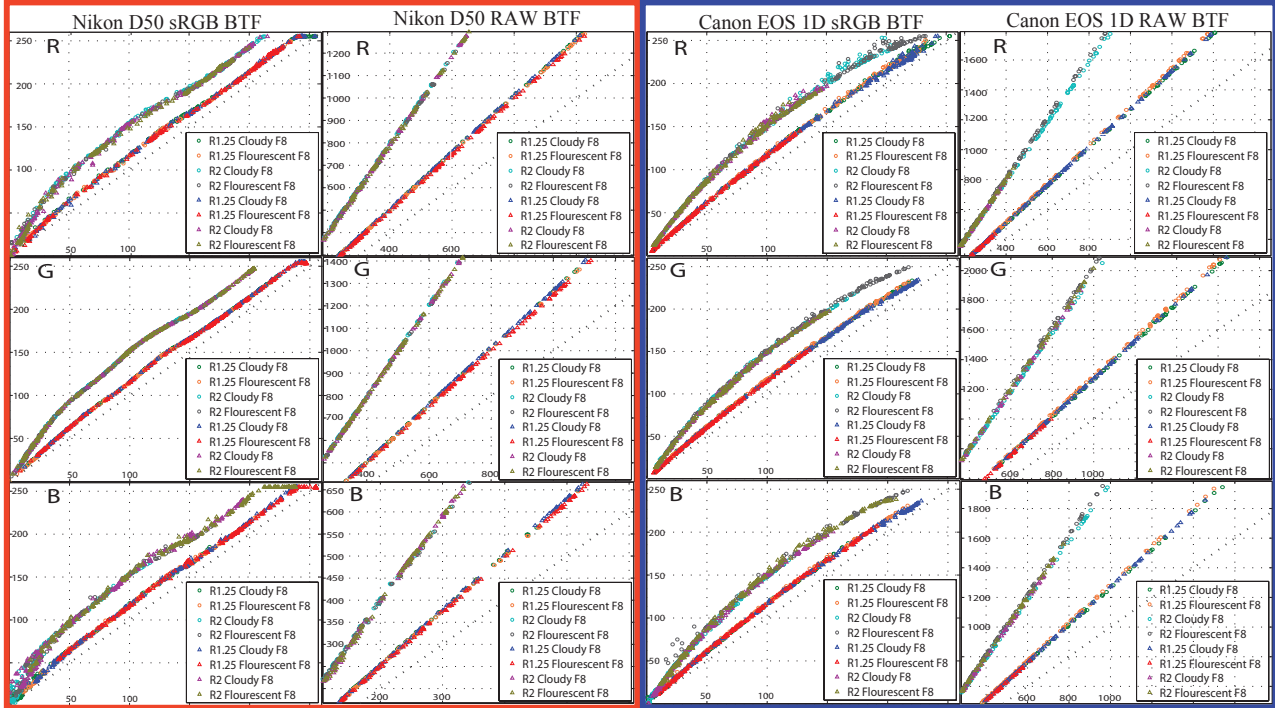


Figure 1. Brightness transfer functions for Nikon D50 and Canon EOS-1D. Each plot includes several BTFs with different exposure ratios (1.25 and 2.0), different lighting environments (\circ : outdoors, \triangle : indoors), and different white balance settings (cloudy and fluorescent). The key observation 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.

3. Data Collection

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. We additionally use the database in [2] which includes over 1000 images from 35 cameras. Cameras from most of the major manufacturers are included as shown in Fig. 4. 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.

4. Observations

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 [4]. These pairs can be arranged into a plot that represents the *brightness transfer function* (BTF [10]), which can be formulated from Eq. 2 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. 4, 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 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 speed¹, and it is well known that the shutter speeds of cameras may be imprecise [7]. 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 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 [8, 9, 17]. 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 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

¹We use shutter speed to control exposure because changing the aperture could result in spatial variation of irradiance due to vignetting.

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. Additionally, 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.

5. Radiometric Calibration Algorithm

In this section, we describe practical steps to make radiometric calibration more robust by taking into account the observations and findings from the previous section. The overall procedure follows conventional radiometric calibration methods which operate in the log domain, and we use the PCA based model of camera response functions introduced in [5]:

$$g(I_x) = \log(f^{-1}(I_x)) = g_0(I_x) + \sum_{n=1}^M h_n(I_x) c_n \quad (7)$$

where g_0 is the mean response function and h_n 's are the PCA basis functions of the response curve in the log domain. Given multiple images with different exposures, the response function can be computed linearly by putting the model in Eq. 7 into the following equation, which is the log version of Eq. 2:

$$g(I'_x) - g(I_x) = K', \quad (8)$$

where $g = \log f^{-1}$ and $K' = \log k'$.

A common practice for solving Eq. 8 is to read the exposure values from the camera's meta data and use it for K' . However, actual image recordings are not always consistent with the camera's meta data as discussed in the previous section. To deal with this issue, capturing images both in sRGB and RAW format is recommended if the camera supports RAW capture. With RAW images, the exposure ratio can be directly computed by dividing the raw values between two images. If RAW images are unavailable, K' in Eq. 8 can be treated as an unknown variable and be solved together with the log response g as in [10]. Notice that the solution using the method in [10] is determined up to an unknown scale in the log domain, meaning that the computed response function f^{-1} is known up to an unknown exponential γ . To address this ambiguity, we fixed the exposure ratio of an image pair using the exposure values from the

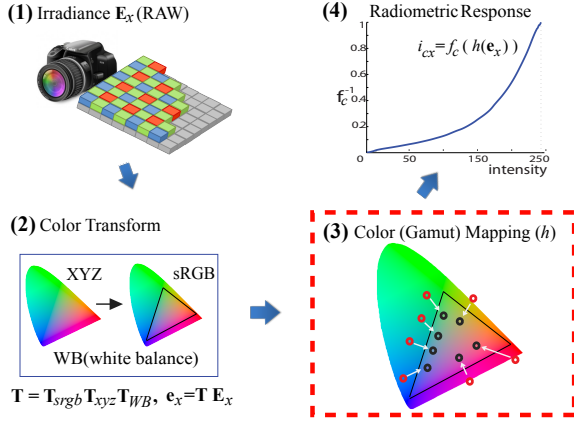


Figure 2. A new radiometric model: the color (gamut) mapping process h [9] is added to overcome the limitation of the conventional model.

camera. We took multiple shots of each image to deal with the instability of the shutter speeds.

The outliers mentioned in Section 4 also need to be addressed for the accurate radiometric calibration. In some cameras and scenes, the number of outliers is small and the effect of outliers on the estimation of the response curve is minimal. However, we have found that the outliers can have a serious effect in some cameras and scenes, so the outliers need to be detected and discarded from the response estimation. First, points with 0 or 255 in any of the channels are rejected as outliers. Next, we use the color saturation value to detect the points that lie close to the boundary of the color gamut. We convert RGB to HSV, and points with saturation over a threshold (β) are considered to be outliers. Suggestions on how to best select β to remove outliers is described in Section 7. In the following Section, we describe a calibration scheme to deal with these outliers by modeling the color mapping component of the in-camera processing.

6. Image to Irradiance (RAW)

It was shown in Section 4 that the existing radiometric model (Eq. 1) cannot represent the nonlinear color mapping (gamut mapping) effectively by simply having a different response function per color channel. Therefore, points with high color saturation cannot be mapped back to physical values as well as neutral colors. To overcome the limitation of the conventional model in Eq. 1, we introduce a new radiometric imaging model as follows:

$$i_{cx} = f_c(h(\mathbf{T}\mathbf{E}_x)), \quad (9)$$

where c represents the color channel, f is the conventional radiometric response function responsible for the tone-compression, $h: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is the color mapping function, \mathbf{T} is a 3×3 matrix which includes the white balance and the transformation from camera color space to linear

sRGB, and \mathbf{E}_x is the irradiance (RAW). Fig. 2 shows a diagram of this radiometric model including the color mapping.

Because the nonlinear color mapping process h is specific to each camera manufacturer and can be drastically different, it is difficult to design a parametric model for this process. Instead, we use scatter point interpolation via radial basis functions (RBF) to estimate this nonlinear mapping as:

$$h^{-1}(\mathbf{e}) = \sum_{i=1}^N w_i \|\mathbf{e} - \mathbf{e}_i\|_2, \quad (10)$$

where \mathbf{e} represents a linearized sRGB color point and \mathbf{e}_i represents a RBF control point. The weights (w_i) are computed from a set of selected sRGB-RAW control point pairs in the form of $\mathbf{e}_i \rightarrow \mathbf{T}\mathbf{E}_i$, where the $\mathbf{T}\mathbf{E}_i$ is the corresponding RAW value that has been corrected by \mathbf{T} . For more information on RBF, readers are referred to [1]. This inverse color mapping $h^{-1}(\mathbf{e})$ is essentially a 3D warping that reverses the gamut mapping process which enables more accurate reconstruction of a RAW image from a given sRGB image.

We pre-calibrate the functions f_c^{-1} and h^{-1} per camera and the transformation \mathbf{T} per white balance setting. The response functions (f_c^{-1}) are computed as described in Section 5 using a number of differently exposed images. With the response function computed, the transformation \mathbf{T} is then computed using Eq. 4 from a number of sRGB-RAW pairs. Finally, the color mapping function (h^{-1}) which should map the linearized image values to the RAW values transformed by \mathbf{T} is computed from a number of sRGB-RAW pairs with various white balance settings (typically 1500 samples are used to define h^{-1}). After the pre-calibration, a new image which is in the non-linear sRGB space can be mapped to the RAW by

$$\begin{bmatrix} E_{rx} \\ E_{gx} \\ E_{bx} \end{bmatrix} = \mathbf{T}^{-1} \cdot h^{-1} \left(\begin{bmatrix} f_r^{-1}(i_{rx}) \\ f_g^{-1}(i_{gx}) \\ f_b^{-1}(i_{bx}) \end{bmatrix} \right). \quad (11)$$

7. Experiments

Radiometric response function estimation. We first compare the performance of the practical algorithm (Section 5) against the conventional approach [5] upon which we have built our algorithm. Fig. 3 shows an example of the outliers detected by our algorithm and the response functions recovered by the two methods. There is a significant difference in the estimations and the proposed algorithm clearly outperforms on the linearization results.

A selected few response functions computed using our algorithm for some cameras in our database are shown in Fig. 4. Note that the responses differ from the gamma curve ($\gamma = 2.2$) commonly used for linearization in some color




Figure 4. Radiometric response functions for a set of cameras in the database and mean linearization errors (δ in Eq. 12) 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.

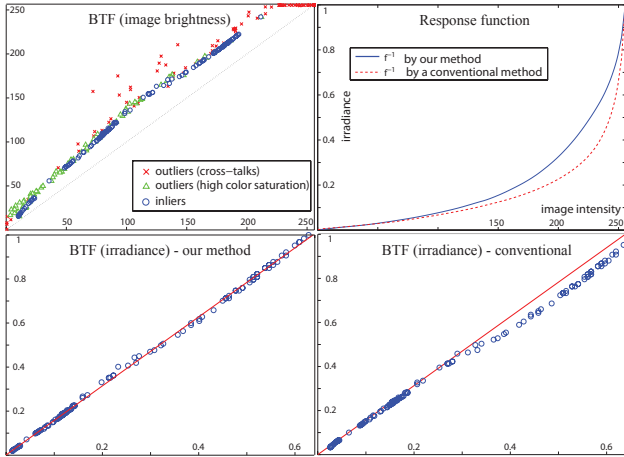


Figure 3. A BTF, estimated response function, and linearization results for the blue channel of Nikon D40 using our practical algorithm and a conventional method [5]. Using the practical algorithm, outliers can be effectively removed for more accurate calibration.

vision work. The response functions for the cameras in the database, as well as data and the software for exploring the data, are all available online at the author’s project webpage. For a quantitative evaluation of the response estimation, we use the following measure per channel to gauge the accuracy of linearization from Eq. 2:

$$\delta_c = \sqrt{\frac{\sum_{n=1}^N \sum_{x \in A} \|k'_n f^{-1}(i_{cx}^n) - f^{-1}(i_{cx}^{n'})\|^2}{N|A|}}, \quad (12)$$

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 com-

puted 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. 4 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. 4. If we exclude outliers from the computation, δ converges almost to 0 in many cameras. So the δ in Fig. 4 is related to the amount of outliers, or the degree of color mapping in 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. For this reason, we had to rely mostly on gray patches to reliably compute the response functions for the cameras with high δ , using a low threshold ($\beta = 0.4$) for determining saturation outliers in Section 5. A higher threshold ($\beta > 0.7$) was used for most other cameras with low δ ’s.

Color mapping function estimation and converting image to RAW. Next, we evaluate the new radiometric model (Eq. 9) and the method for converting images to RAW responses. The 3D color mapping functions (h) for the Nikon D50 and the Canon EOS-1D are shown as slices in Fig. 5 (a). The colors on the map in Fig. 5 (a) encode the magnitude of 3D color warping at a color point e , $\|e - h^{-1}(e)\|^2$. The results confirm the existence of the 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 algorithm for converting image values to RAW responses described in Section 6 is shown in Fig. 5 (b). In the figure, we compare the results from three different techniques given a number of sRGB-RAW image pairs. The first method is the implementation of the algorithm from [2] where f (5th order polynomial) and T are computed iteratively. The second method computes the

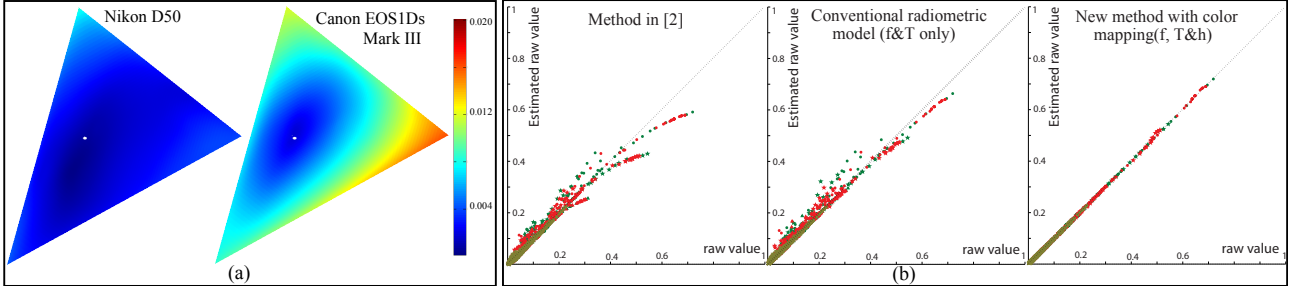


Figure 5. (a) Estimated color mapping function h for Nikon D50 and Canon EOS-1D. The maps shown here are slices of the 3D functions. It can be seen that the gamut mapping was only applied to the points near the boundaries in Nikon D50 whereas the gamut mapping influences the points well within the gamut and the degree of the mapping is more severe in Canon EOS-1D. (b) Performance of mapping image values to RAW values (Canon EOS-1D) with different techniques: using the technique in [2], using f and T from Section 6 without h , and the new method with h . Using our new model, images can be mapped back to RAW accurately.

RAW just from f and T , which are computed as described in Section 6 without the color mapping function h . Finally, the third method computes RAW from Eq. 11 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.

Finally, we show the results of applying the calibration results to convert images of real scenes back to RAW responses for various scenes and cameras in Fig. 6. 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 T and the ground truth RAW images are used only for the evaluation purposes. Using the new model and the algorithm introduced in Section 6, we can accurately convert the image values to RAW values even for the highly saturated colors in the scene.

8. Conclusion and Future Work

In this paper, we presented a study of radiometric calibration and the in-camera image processing through an extensive analysis of a large image database. One of the key contributions of this paper is bringing the color (gamut) mapping in the in-camera image processing to light to overcome the limitations of the conventional radiometric model and calibration methods. By considering the color mapping in the imaging process, we could compute the radiometric response function more accurately. Moreover, by introducing the color mapping function to the radiometric model and the algorithm to estimate it, we could convert any given image to RAW accurately using a pre-calibration scheme.

There are several directions that we need to explore more in the future. We currently rely on a pre-calibration scheme, where different mappings (f , h , and T) are first computed given a number of sRGB-RAW pairs and are used later to transform images to RAW images. While this is an effective

scheme for cameras with RAW support, we cannot use this scheme for cameras that do not provide RAW images. Further investigation is needed to explore the possibility of extending the work to cameras without RAW support.

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. In the future, we plan to investigate to see 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 [2, 6, 11, 12] research in the future.

Acknowledgement

We thank Dr. Steve Lin for his help and valuable discussions. This work was supported in part by the NUS Young Investigator Award, R-252-000-379- 101.

References

- [1] M. D. Buhmann. *Radial Basis Functions: Theory and Implementations*. Cambridge University Press, 2003. 5
- [2] A. Chakrabarti, D. Scharstein, and T. Zickler. An empirical camera model for internet color vision. In *Proc. British Machine Vision Conference*, 2009. 2, 3, 6, 7
- [3] P. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Proc. ACM SIGGRAPH*, pages 369–378, 1997. 2
- [4] M. Grossberg and S. Nayar. Determining the camera response from images: What is knowable? *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 25(11):1455–1467, 2003. 2, 3
- [5] M. Grossberg and S. Nayar. Modeling the space of camera response functions. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 26(10):1272–1282, 2004. 1, 4, 5, 6
- [6] T. Haber, C. Fuchs, P. Bekaert, H.-P. Seidel, M. Goesele, and H. Lensch. Relighting objects from image collections.




Figure 6. 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. 5 (a)). More examples can be found in the supplementary materials.

- In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008. 7
- [7] P. D. Hiscocks. Measuring camera shutter speed, 2010. <http://www.syscompdesign.com/AppNotes/shutter-cal.pdf>. 4
- [8] J. Holm, I. Tastl, L. Hanlon, and P. Hubel. Color processing for digital photography. In P. Green and L. MacDonald, editors, *Colour Engineering: Achieving Device Independent Colour*, pages 79 – 220. Wiley, 2002. 1, 2, 4
- [9] 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. 2, 4, 5
- [10] S. J. Kim and M. Pollefeys. Robust radiometric calibration and vignetting correction. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 30(4):562–576, 2008. 2, 3, 4
- [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*, pages 74–86, 2008. 2, 7
- [12] 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*, 28(5), 2009. 7
- [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*, pages 938–945, 2004. 2
- [14] 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*, pages 66–73, 2005. 1, 2
- [15] 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*, pages 422–428, 1995. 2
- [16] T. Mitsunaga and S. Nayar. Radiometric self-calibration. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 374–380, 1999. 2
- [17] J. Morovic and M. R. Luo. The fundamentals of gamut mapping: A survey. *Journal of Imaging Science and Technology*, 45(3):283–290, 2001. 4
- [18] 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*, pages 173–180, 2004. 2