

Rank-Based Radiometric Calibration

Han Gong, Graham D. Finlayson[^], and Maryam M. Darrodi

University of East Anglia, UK

E-mail: h.gong@uea.ac.uk

Robert B. Fisher

University of Edinburgh, UK

Abstract. Raw images are more useful than JPEG images for machine vision algorithms and professional photographers because raw images preserve a linear relation between pixel values and the light measured from the scene. A camera is radiometrically calibrated if there is a computational model which can predict how the raw image is mapped to the corresponding rendered image (e.g., JPEGs) and vice versa. Our method makes use of the observation that the rank order of pixel values is mostly preserved post-color correction. We show that this observation is the key for getting a compact and robust radiometric calibration model. Since our method requires fewer variables, it can be solved for using less calibration data. An additional advantage is that we can derive the camera pipeline from a single pair of raw–JPEG images. Experiments demonstrate that our method delivers state-of-the-art results (especially for the most interesting conversion from JPEG to raw). © 2018 Society for Imaging Science and Technology. [DOI: 10.2352/J.ImagingSci.Technol.2018.62.5.050404]

1. INTRODUCTION

Many computer vision algorithms (e.g., photometric stereo [1], photometric invariants [2], shadow removal [3, 4], and color constancy [5]) rely on the assumption that the captured RGBs in images are linearly related to the actual scene radiance. However, mostly, the actual output of a digital camera imaging pipeline is necessarily non-linear in order to produce perceptually pleasing photos as opposed to their physically meaningful counterparts. In this paper, we present a compact rank-based (RB) radiometric calibration method which solves for the bidirectional mappings between the camera’s raw responses and the rendered RGBs produced by digital cameras.

There is prior art in this field, which models the pipeline with a large number of parameters (up to several thousands [6]). This means that a large corpus of data is required to uncover the pipeline and there is at least tacitly the premise that the underlying pipeline is quite complex. The key insight in our approach is that post-color correction (a 3×3 matrix correction), the linear corrected raw RGBs are to the greatest extent of the same rank order as the final rendered RGBs. Based on this insight, we develop a

compact RB radiometric calibration model that models the camera pipeline with significantly fewer parameters and concomitantly needs much less training data.

In Figure 1, we illustrate a typical image reproduction pipeline which is representative of many cameras [7]. An exemplar raw image, Fig. 1(a), is mapped by a 3×3 color correction matrix to give a color corrected image (Fig. 1(b)). The color correction matrix implements several processing steps (e.g., illumination correction [8, 9], display RGB mapping [10], and color preference adjustments [8]). It is well known that a display device cannot show all captured image colors and some RGBs will fall outside the RGB cube after mapping (e.g., the pixels marked in light purple in Fig. 1(b)). Gamut mapping is therefore required, e.g., [6, 7, 11], to move the colors back inside the cube as shown in Fig. 1(c). Finally, the gamut mapped image is tone mapped to arrive at the final rendered output [7–9] shown in Fig. 1(d). Tone mapping accounts for the display non-linearity [10], dynamic range compression, and some aspects of preference [12].

In general, the camera color processing pipeline can be written as Eq. (1).

$$\underline{P} = \underbrace{f(\Upsilon(\underline{M}\underline{\rho}))}_{(1a)} = \underbrace{\Upsilon(f(\underline{M}\underline{\rho}))}_{(1b)} \approx \underbrace{\text{LUT}(\underline{\rho})}_{(1c)} \quad (1)$$

where $\underline{\rho}$ denotes a camera raw and \underline{P} refers to its rendered RGB counterpart. The 3×3 correction matrix, gamut mapping and tone mapping are respectively denoted by the matrix \underline{M} and the functions $\Upsilon()$ and $f()$. The function $f()$ can implement a single or three per-channel tone curves. Since gamut mapping only implements a small change in comparison with color and tone mapping steps, the order of gamut mapping and tone mapping may be switched (Eqs. (1b) and (1c)), a property that we exploit in this paper. Equally, we can also merge three processing steps into one and directly solve for a 3-D LUT (Look-Up-Table) that maps raw to rendered counterparts. This LUT interpolation function is denoted by $\text{LUT}()$ [13] in Eq. (1c). Readers may refer to the top row of Fig. 1 to link each mathematical function to our example processed image.

In radiometric calibration, given a set of $\underline{\rho}$ and \underline{P} , we solve for the parametrized pipeline parts (e.g., \underline{M} , $\Upsilon()$, $f()$ and $\text{LUT}()$). A disadvantage of the current best performing methods is that a great deal of data may be required to

[^]IS&T Member.

Received Apr. 12, 2018; accepted for publication Aug. 2, 2018; published online Sep. 5, 2018. Associate Editor: Dileepan Joseph.

1062-3701/2018/62(5)/050404/8/\$25.00

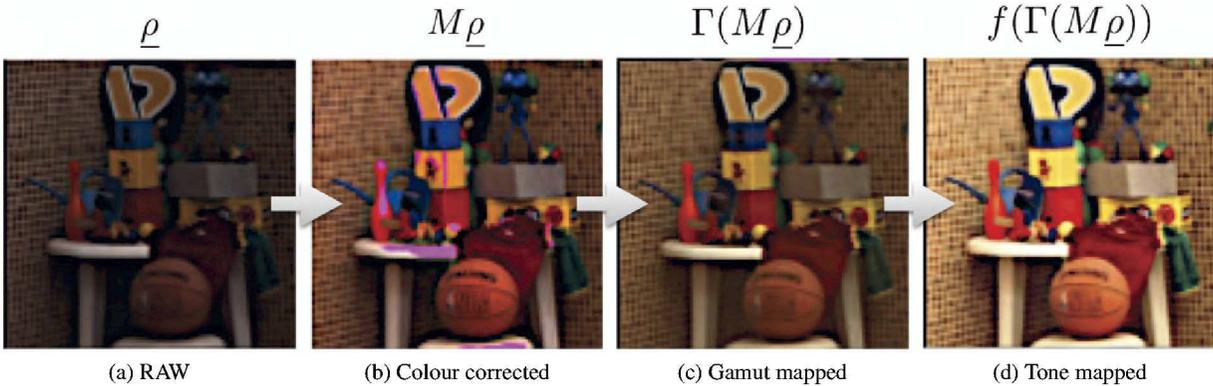


Figure 1. (a) A raw input image is color corrected to give image (b). Non-displayable colors are highlighted in purple pseudo color. Gamut mapping, in step (c), brings colors within gamut. Finally, in (d), a tone mapping step results in the final rendered image. The image is taken from [14].

fit their assumed models. In Eqs. (1a) and (1b), the gamut mapping step could be modeled by 1000s of radial basis functions (RBFs) [6, 7, 13] and in Eq. (1c), the deployed LUT interpolation function could also have several thousand control points.

Our proposed method exploits the simple observation [15] that, assuming the gamut mapping step slightly changes image colors and the tone curves are always monotonically increasing, we expect mostly the rank ordering of the rendered \underline{p} to be the same as $\underline{\rho}$ multiplied by the correction matrix \mathbf{M} . Suppose that two rendered (JPEG) responses—in the 1st color channel—are denoted by P_1^a and P_1^b and that $P_1^a > P_1^b$. The rank order of two corresponding raw red channel measurements post-color correction is written as $M_1 \underline{\rho}^a > M_1 \underline{\rho}^b$ (where M_1 denotes the first row of \mathbf{M} and $\underline{\rho}^a$ and $\underline{\rho}^b$ are a pair of raw RGBs). This implies that $M_1(\underline{\rho}^a - \underline{\rho}^b) > 0$ which defines a half-space constraint.

The row vector M_1 can be considered as a point in 3-space and this inequality (ranking constraint) forces the point to be located in only one half of 3-space. Because we have multiple pixels, each pair of pixels (2 raw and 2 JPEG RGBs) generates a half-space constraint and intersecting all these constraints delimits the region in which M_1 must lie. Our experiments demonstrate that a small number of patches suffices to estimate \mathbf{M} accurately.

Once we have \mathbf{M} we then find the best rank-preserving tone curves $f()$. At this stage, only using \mathbf{M} and $f()$ we have a good approximation of the pipeline. Indeed, we argue that our construction of \mathbf{M} and $f()$ also incorporates, to a first order, gamut mapping. Now we adopt (Eq. (1b)) and find a 125-parameter per-channel LUT interpolation to reduce any remaining errors due to gamut mapping (higher order terms).

Below, we review radiometric calibration paying special attention to methods which adopt Eq. (1). Rank-based radiometric calibration is then described and is shown to provide leading performance on a public dataset. Finally, an application of one-shot radiometric calibration without the access to raw is shown.

2. RELATED WORK

Using the pipeline form of Eq. (1b), Chakrabarti et al. [6] first solve for \mathbf{M} and $f()$ in iteration and then solve directly for $\Gamma()$. In their approach, $f()$ is constrained to be a 7th order increasing polynomial. They model $\Gamma()$ by the RBF method of [7] where several thousands of RBFs are potentially used. A restriction of the above calibration is presented in [9] where the gamut mapping $\Gamma()$ is ignored. This less general model works tolerably well on many real pairs of raw and rendered images and this is a point we will return to later in this paper. In either version ([6] or [9]), the coupled nature of the minimization indicates that a global minimum is not guaranteed. Therefore, a random start point search is implemented to find a better set of parameters.

Kim et al. [7] solve for the pipeline in the form of Eq. (1a) and makes additional assumptions to decouple the optimization. They assume that images of the same scene are captured with respect to two or more exposures and their $\Gamma()$ is a multi-thousand set of RBFs. Regarding solving for $f()$, Debevec et al. [16] showed how relating corresponding pixels under known exposure differences suffices to solve for $f()$ (assuming there is no gamut mapping step). Importantly, in [7], it was argued that for the set of desaturated pixels (i.e., raws far from the RGB cube boundary), the gamut mapping step has little or no effect and can be ignored. Relative to this assumption, $f()$ can be solved using the Debevec method. Given $f()$ the color correction matrix \mathbf{M} can be found (again using desaturated pixels).

We point out that for most off-shelf capture devices (e.g., for most mobile phones), manual exposure control is usually unavailable and the requirement of multiple exposures is impractical. We also note that, in [7], the adopted gamut mapping RBF network requires a large number of parameters and thus a large corpus of data [6, 7].

In [13], it was shown that it is possible to ignore the underlying structure of the color processing pipeline and directly solve for the best 3-D surjective function—implemented as a LUT interpolation function that maps the raws to rendered RGBs Eq. (1c). Finally, in [17], a method is presented for solving for $f()$ by examining the

edge distribution in an image. This method has the advantage that the method works for a single image (without multiple exposures) but the method is sensitive to processing steps such as image sharpening which is used extensively in mobile phone image processing.

3. THE RANK-BASED METHOD

In this paper, we are interested in calibrating in the most general circumstances when the amount of training data is modest and there is only a single calibration image. Assuming that the rank order of intensities are almost preserved after the camera processing steps, we present a method that solves for an accurate rank-preserving camera pipeline model. This paper extends our previously published RB method [18] with more details and experiments.

To make the RB method work we need to assume that the gamut mapping step $\Gamma()$ only changes color slightly. In fact our assumption is more nuanced. We assume that—to a first order—gamut mapping can mostly be implemented as an affine transform and that this affine transform can be folded into the color correction matrix \mathbf{M} and the monotonically increasing tone mapping functions $f()$.

3.1 Gamut Mapping as An Affine Transform

After color correction, some colors are mapped outside the color cube and become non-displayable. To address this, gamut mapping in Eq. (1b) is applied. A Taylor expansion to model $\Gamma()$ around a point \underline{a} inside the gamut is used:

$$\Gamma(M\rho) \approx \Gamma(\underline{a}) + J(\underline{a})(M\rho - \underline{a}) \quad (2)$$

where J is the 3×3 Jacobian (matrix of derivatives of Γ). Not only does Eq. (2) show that, to a first approximation, gamut mapping is an affine transform it is also one of the gamut mapping algorithms proposed in [11]. We solve for the affine mapping that maps all pixel colors into the unit RGB space cube.

$$\min_{T, \underline{o}} \sum_i \|TM\rho_i + \underline{o} - M\rho_i\|^2 \text{ s.t. } \underline{0} \leq TM\rho_i + \underline{o} \leq \underline{1}. \quad (3)$$

In Eq. (3), T and \underline{o} are respectively a 3×3 matrix and 3×1 offset vector defining the affine gamut mapping algorithm, i is the index of an input RGB vector. The 3-vectors of 0s and 1s are denoted by $\underline{0}$ and $\underline{1}$. Eq. (3) is solved directly by Quadratic Programming [19]. The gamut mapping shown in Fig. 1(c) is the result of solving Eq. (3).

Here, we make two important remarks about affine gamut mapping: (1) Gamut mapping and color correction combined can be represented by the single affine transform: 3×3 matrix TM and offset \underline{o} ; (2) It follows that the RB method presented in the next section will actually solve for TM . The offset term can be incorporated directly in $f()$.

Our hypothesis is that the part of gamut mapping that is not described by an affine transform will be small and the remaining error can be modeled with a function that has fewer parameters (100s in contrast to the prior art 1000s).

3.2 Rank-Based Estimation for Color Correction

Denote the k th row of \mathbf{M} as M_k . We assume that given two color corrected raws, $M_k\rho^a$ and $M_k\rho^b$ the rank order is the same as for the corresponding rendered RGBs:

$$P_k^a > P_k^b \Rightarrow M_k\rho^a > M_k\rho^b \Rightarrow M_k(\rho^a - \rho^b) > 0. \quad (4)$$

Defining the difference vector $\underline{d}^j = \rho^a - \rho^b$:

$$M_k \underline{d}^j > 0 \quad (5)$$

where it is understood the superscript j denotes the difference vector from the j th of $\binom{n}{2}$ pairs of image pixel values (n is the total number of image pixels). Suppose that we have a vector M_k where Eq. (5) holds, then the inequality cannot be true for $-M_k$. That is, Eq. (5) defines a half-plane constraint [15, 20]. The vector \underline{d}^j is perpendicular to the half-plane: any M_k less than 90° to \underline{d}^j is a possible solution. Given multiple difference vectors we have multiple half-plane constraints which taken together delimit a region in 3-space where M_k must lie. Denoting the half-plane as $\mathcal{H}(\underline{d}^j)$, M_k must satisfy:

$$M_k \in \bigcap_j \mathcal{H}(\underline{d}^j). \quad (6)$$

The intersection in Eq. (6) defines an unbounded cone, anchored at the origin, in 3-D space. Clearly, if M_k is in the intersection region defined by Eq. (6), αM_k (where α is a positive scalar) is another solution. Using ranking we solve for each row of \mathbf{M} up to an unknown scalar multiplier.

Let us visualize the computation of M_k using ranking. Without loss of generality let us assume that $M_{k,3} = 1$. We rewrite Eq. (5) as

$$M_{k,1}d_1^j + M_{k,2}d_2^j + d_3^j > 0. \quad (7)$$

If $[a \ b \ c]$ is a solution to Eq. (6), then $[a/c \ b/c \ c/c]$ for Eq. (7) is also true since $M_{k,1} = a/c$ and $M_{k,2} = b/c$. Solutions for $[M_{k,1}, M_{k,2}]$ lie on one side of the line, i.e., the 3-D half-space constraints map directly to a 2-D half-plane constraint. Alternatively, if we consider the whole set of intersections, the cone in 3-D defined by Eq. (6) maps to a 2-D convex region [21]. Denoting half-planes as $\mathcal{P}(\underline{d}^j)$ we, equivalently, solve for

$$[M_{k,1}, M_{k,2}] \in \bigcap_j \mathcal{P}(\underline{d}^j). \quad (8)$$

The intersection problem of Eq. (8) is easily visualized. In Figure 2(a) we show the intersection of 4 half-plane constraints and indicate the solution set where M_k must lie.

We solve for M_k one sensor channel at a time. Due to noise or small deviations in real camera data, it is likely that no common intersection can be found that satisfies every half-plane's constraint. To solve this problem, we generate 100,000 unit length vectors that are uniformly distributed on the surface of the unit sphere [22], which is visualized in Fig. 2(b). With respect to this sampling, the furthest distance between any point and its nearest neighbor is less than 1.15° . Therefore, the orientation of the rows of \mathbf{M} is found to this

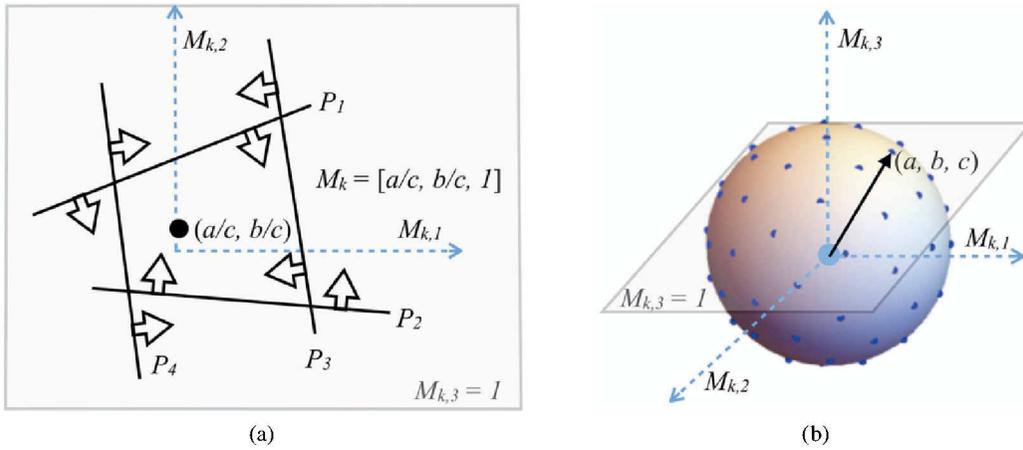


Figure 2. (a) The region where four half-plane constraints intersect delimits the region where $[M_{k,1}, M_{k,2}]$ must lie where the black point is a feasible solution. (b) On a unit sphere, each vector represented by the origin and a blue surface point is a probe for a possible solution (e.g., the black arrow). All 3-D points and constraints are projected to a 2-D plane $M_{k,3} = 1$.

accuracy. For each point on the sphere (i.e., a possible row of M_k), we count how many half-space constraints are satisfied. The point on the unit sphere that has the largest number of satisfying half-plane constraints—or the median of multiple points if there is a tie—defines M_k .

To maintain a reasonable computational cost, we have to be careful not to generate too many half-planes. We simply select 200 random RGB-JPEG pairs for estimating M . Since it is not guaranteed that these 200 random pairs are the optimum selections, we practically generate 25 random sets of 200 RGB-JPEG pairs and select the “best luck” set which gives the lowest training error. Note that the other pipeline components introduced later are also applied to evaluate the training error.

Overall, we find the M that places all the corresponding raw and rendered image RGBs in the most similar rank order. That is, if we plot the mapped raw red responses, for example, against the rendered red JPEG corresponding values then the graph should be a monotonically increasing function. How well a monotonically increasing function fits our data can be used to judge the efficacy of each M .

3.3 Rank-Preserving Optimization of Tone Curves

We now solve for the optimal per-channel tone curves which map color corrected raws to corresponding rendered RGBs. Let us denote the i th color corrected raw and rendered RGB pixel pairs for the k th channel as $(M_k \rho_{k,i}, P_{k,i})$. The k th-channel rank-preserving tone curve $f_k()$ is optimized as a 7th order monotonic and smooth polynomial function as follows:

$$\min_{f_k} \sum_i \|f_k(M_k \rho_{k,i}) - P_{k,i}\|^2 + \lambda \int_t \|f_k''(t)\|^2 dt \quad \text{s.t. } f_k'() \geq 0. \quad (9)$$

where the first term is for data fitness, the second term is for curve smoothness and λ is a small weight (e.g., 10^{-5}). The constraint enforces monotonicity. This polynomial fitting is

solved by Quadratic Programming [19]. Note that these 7th order polynomials already include the affine gamut mapping offset parameters described previously. In this paper, we further denote the combination of all 3-channel mappings $f_{1-3}()$ as $f()$.

3.4 Gamut Correction Step

As argued previously, we propose that $f(M\rho)$ has the expressive power to implement color correction, tone correction and gamut mapping (to the first order in a Taylor expansion). However, we wish to add a further gamut mapping step for the higher order terms. But, since our hypothesis is that much of the gamut mapping will have been accounted for we are going to adopt a simple small parameter solution. Further, since this additional correction is going to be carried out at the end of the process, we adopt Eq. (1b). Specifically, we find a $5 \times 5 \times 5$ LUT interpolation function by using lattice regression [23] that minimizes $\min_{LUT} \sum_i \|LUT(g(f(M\rho_i))) - \underline{P}_i\|^2$ where $g()$ is a non-linear function that stretches highlights. We found empirically there was an advantage in deploying more LUT resolution in the highlight region where gamut mapping is created. We implemented this not by changing the sampling structure of the LUT control points (which is uniform) but by stretching our data, by applying the function $g()$ shown in Figure 3. The function $g()$ is fixed for all our experiments.

3.5 Rank-Based Recovery of Raw

Suppose we wish to map rendered RGBs to raws. Using the previously described method, M has already been solved in the RAW-to-JPEG forward estimation phrase. Now, in a least-squares optimal way, we use the same polynomial fitting method Eq. (9) to find f^{-1} by optimizing $\min_{f^{-1}} \sum_i \|f^{-1}(\underline{P}_i) - M\rho_i\|$. Finally, we solve for the backward LUT by optimizing $\min_{LUT} \sum_i \|LUT(g(M^{-1}f^{-1}(\underline{P}_i))) - \underline{P}_i\|$ where the LUT

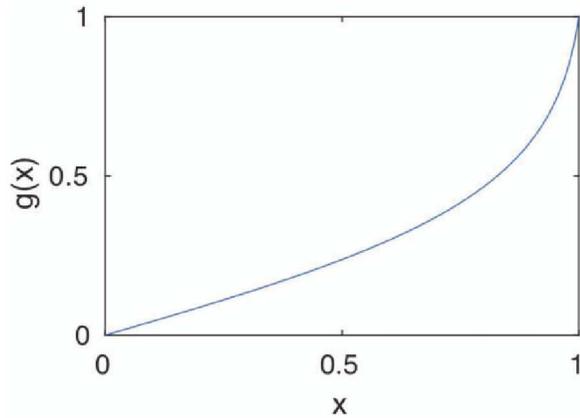


Figure 3. Plot of a non-linear function $g(x)$ where $g(0) = 0$ and $g(1) = 1$.

interpolation function is fitted by a $5 \times 5 \times 5$ lattice regression [23].

3.6 Parameter Counting

Assuming we solve for three independent tone curves then our method requires 9 (for M) + 8×3 (for $f(\cdot)$) + 125×3 (the LUT for $I(\cdot)$) = 408 parameters which is significantly less (even an order of magnitude less) than [6, 7, 13].

4. EVALUATION

Our evaluation is based on two challenging datasets from Harvard [6] and NUS [7]. The Harvard dataset [6] contains the RAW/JPEG intensity pairs of 140 color checker patches viewed under multiple viewing conditions. The color chart is captured by eight cameras and under 16 illuminants across many different exposures. Compared with the Harvard dataset, the NUS dataset contains large data captured with a 24-patch color checker and 31 camera sensors though its capture conditions (i.e., light and exposure) are relatively limited.

We carried out the same experiment described in [6, 7]. We are interested in validating whether our method, with much reduced number of parameters, can produce similar or even better results compared with the state-of-the-art [6]. We evaluate both RAW-to-JPEG and JPEG-to-RAW. The Harvard dataset [6] captures a sort of “worst-case” viewing conditions. Normally, when we capture a picture there is a single prevailing illuminant color. In the Harvard dataset, all camera processing parameters are turned off and the same reflectances are viewed under multiple colored lights. As Forsyth observed [24], the reddest red camera response cannot be observed under a blue light, and he exploited this observation to solve for the color of the light. In real imaging conditions, the greenest green and the bluest blue do not typically appear at the same time. A pipeline that suffices for the combinations of all lights and all surfaces is unlikely to be needed. This means the prior art pipelines are probably more complex than they need to be. As described in [6], for each camera, we estimate the parameters of a calibration model using different subsets of the available RAW-JPEG pairs. For

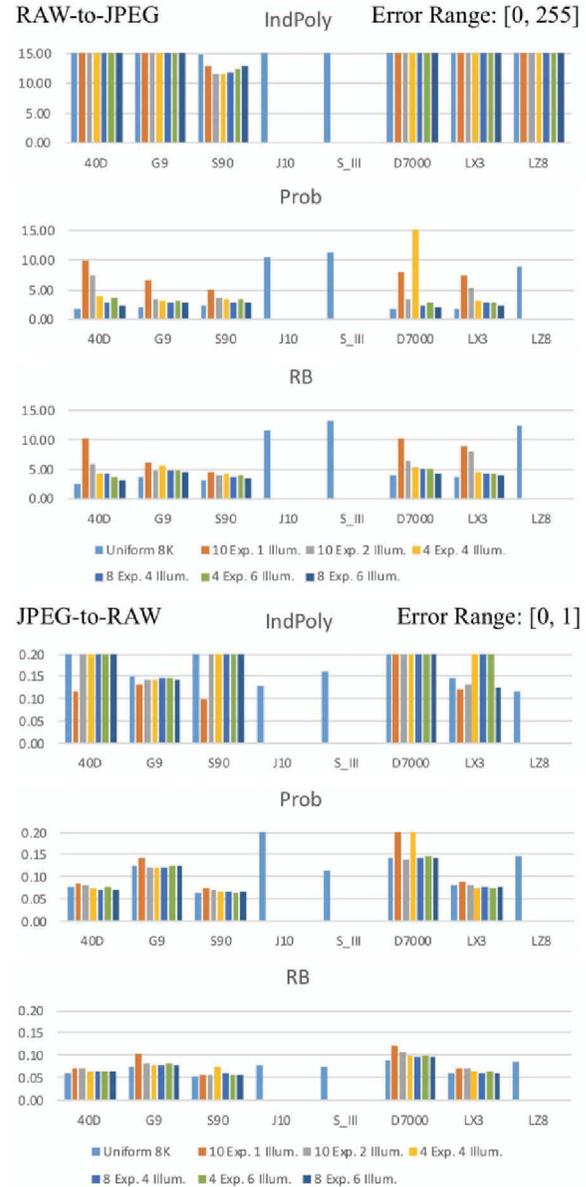


Figure 4. Test RMSE bar chart of the Harvard dataset [6] for IndPoly [9], Prob [6], and our RB [18]. The figure shows RMSE between ground truth and prediction for bidirectional raw and JPEG mappings. “Exp.” and “Illum.” are respectively short for “Exposure” and “Illuminant.” Each bin color indicates a capture condition. The horizontal and vertical axes indicate camera sensor and calibration error, respectively. The displayed forward and backward errors are clipped at 15.0 and 0.20, respectively.

each subset and a selected camera, the root mean squared error (RMSE) between the prediction and ground truth is validated by using all available RAW-JPEG pairs.

Figure 4 (top half) shows the raw-to-JPEG mapping error plot (where pixel intensities are coded as integers in the interval $[0, 255]$). In both forward and backward tests, our RB method [18] is significantly better than the independent polynomial method (IndPoly) [9]. IndPoly is a simple model which only contains three per-channel tone mapping (or

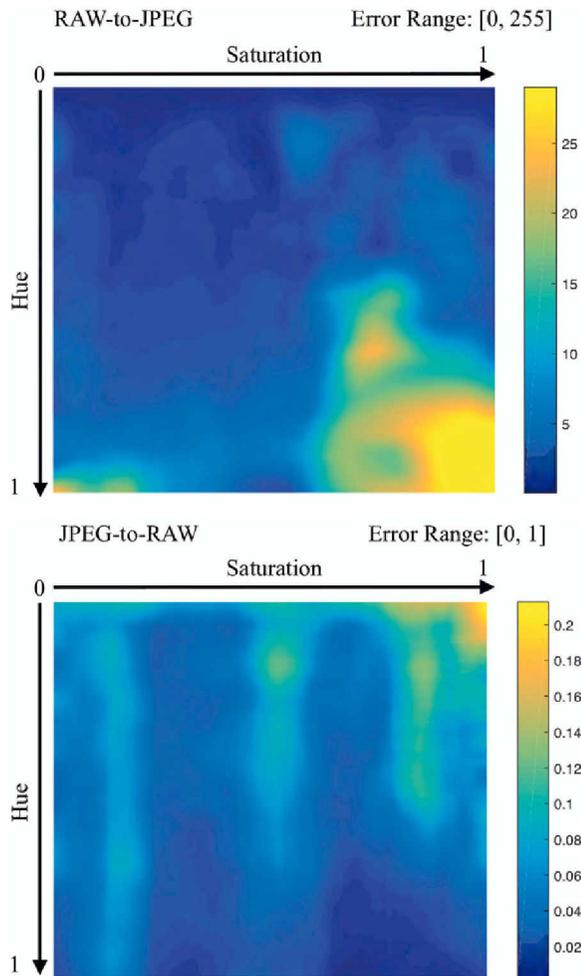


Figure 6. Distributions of forward and backward mapping RMSE of our method over the input hue and saturation gamut. It is tested with the “1 illuminant + 10 exposures” set from the Harvard dataset [6].

Table I. RMSE of our RB method between ground truth and prediction for bidirectional RAW and JPEG conversions. The results are based on 50 repeated tests.

(a) RAW-to-JPEG	40D	G9	S90	D7000	LX3
Mean	10.52	8.85	4.92	14.19	9.74
Std	1.65	2.54	0.41	3.14	1.00
(b) JPEG-to-RAW	40D	G9	S90	D7000	LX3
Mean	0.073	0.106	0.058	0.124	0.071
Std	0.002	0.003	0.001	0.004	0.001

6. CALIBRATION ERROR DISTRIBUTION

To better understand how well the proposed algorithm performs with respect to input hue and saturation, we visualize the distributions of forward and backward mapping errors (RMSE) over the hue and saturation gamut (HSV color space [25]) which is shown in Figure 6. The distributions are generated for the common capture condition set—“1 illuminant + 10 exposures”—as the examples. Specifically,

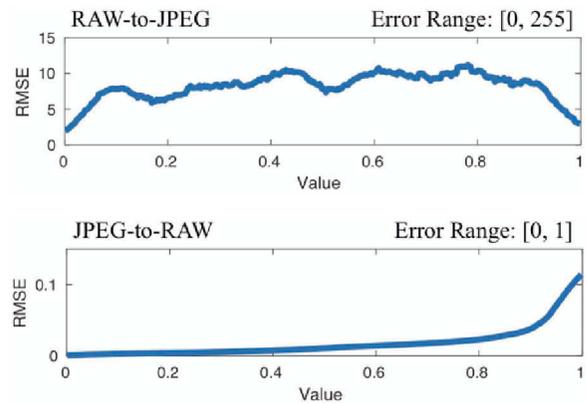


Figure 7. Distributions of forward and backward mapping RMSE of our method with respect to input lightness (value channel of the HSV color space). It is tested with the “10 Exps. and 1 illum.” set from the Harvard dataset [6].

the complete distribution data is interpolated for a 400×400 uniform grid. This grid (as an image) is then filtered by a 41×41 Gaussian kernel with a 20 standard deviation.

In addition, in Figure 7, we also plot the RMSE distributions for the same capture condition—“1 illuminant + 10 exposures”—with respect to lightness (i.e., value channel of the HSV color space). Similarly, the complete distribution data are interpolated for a 400-tick uniform 1-D space. This interpolated 1-D space data is then filtered by a 41×1 Gaussian kernel with a 20 standard deviation.

Overall, higher forward and backward errors are observed when saturation is high and hue is close to 0 or 1 (i.e., reddish colors). Value (of the HSV color space) does not seem to have a great impact for the forward errors although lower forward errors are found near both clipping boundaries of 0 and 1.

7. CALIBRATION WITH SMALL NUMBERS OF PARAMETERS

We wished to visually validate our claim that we can calibrate with few parameters. We took 4 RAW+JPEG pairs (for different cameras) from [9]. We then uniformly selected 140 corresponding pixels from the RAW and JPEG. We solved for all the 408 parameters in our RB method. We then applied our model to the rest of the image. The result of this experiment for four images (JPEG-to-RAW) is shown in Figure 8.

8. CONCLUSION

In this paper we have shown how the rank order of image responses is a powerful tool for solving for the individual steps in a camera processing pipeline (color correction, gamut, and tone mapping). A simple ranking argument, relating color corrected raws to corresponding rendered RGBs suffices to solve for the color correction matrix. Then, the rank-preserving tone map is found and, finally, a simple gamut correction step is derived. Compared with the prior art, our RB method requires the fewest assumptions

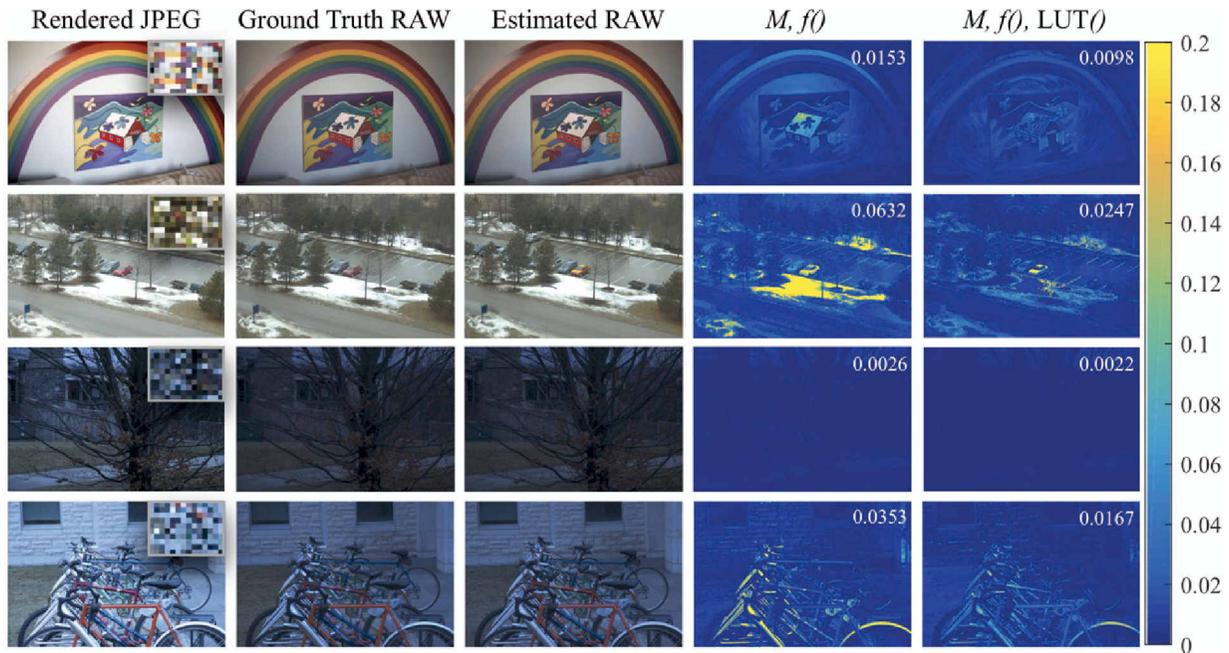


Figure 8. Visualization of one-shot radiometric calibration through a simulated 140-patch color checker, shown at the top-right corner of each Rendered JPEG image. The error maps in the 4th and 5th columns respectively visualize the per pixel RMSE for our RB method with & without the gamut mapping LUT interpolation. The RMSE of each whole image is shown at the top-right corner of each error map. All raw images are shown with a 0.5 gamma.

and delivers state-of-the-art radiometric calibration results. Experiments also show that excellent calibration is possible given a single image exposure and limited color diversity (e.g., a color chart).

REFERENCES

- A. P. Petrov, "On obtaining shape from color shading," *Color Res. Appl.* **18**, 236–240 (1993).
- T. Gevers, G. D. Finlayson, and R. Schettini, "Color for image indexing and retrieval," *Computer Vision and Image Understanding* (Elsevier, New York, NY, 2004), Vol. 94, pp. 1–3.
- H. Gong and D. Cosker, "User-assisted image shadow removal," *Image Vis. Comput.* **62**, 19–27 (2017).
- H. Gong and D. Cosker, "Interactive removal and ground truth for difficult shadow scenes," *J. Opt. Soc. Am. A* **33**, 1798–1811 (2016).
- K. Barnard, L. Martin, A. Coath, and B. V. Funt, "A comparison of computational color constancy algorithms. ii. Experiments with image data," *IEEE Trans. Image Process.* **11**, 985–996 (2002).
- A. Chakrabarti, Y. Xiong, B. Sun, T. Darrell, D. Scharstein, T. Zickler, and K. Saenko, "Modeling radiometric uncertainty for vision with tone-mapped color images," *IEEE Trans. Pattern Anal. Mach. Intell.* **36**, 2185–2198 (2014).
- S. J. Kim, H. T. Lin, Z. Lu, S. Susstrunk, S. Lin, and M. S. Brown, "A new in-camera imaging model for color computer vision and its application," *IEEE Trans. Pattern Anal. Mach. Intell.* **34**, 2289–2302 (2012).
- R. Ramanath, W. E. Snyder, Y. Yoo, and M. S. Drew, "Color image processing pipeline," *IEEE Signal Process. Mag.* **22**, 34–43 (2005).
- A. Chakrabarti, D. Scharstein, and T. Zickler, "An empirical camera model for internet color vision," *British Machine Vision Conf.* (BMVA, Durham, England, 2009), Vol. 1, p. 4.
- M. Anderson, R. Motta, S. Chandrasekar, and M. Stokes, "Proposal for a standard default color space for the internet:srgb," *Proc. IS&T/SID CIC4: Fourth Color Imaging Conf.* (IS&T, Springfield, VA, 1995), pp. 238–245.
- J. Giesen, E. Schuberth, K. Simon, P. Zolliker, and O. Zweifel, "Image-dependent gamut mapping as optimization problem," *IEEE Trans. Image Process.* **16**, 2401–2410 (2007).
- E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda, "Photographic tone reproduction for digital images," *ACM Trans. Graph.* **21**, 267–276 (2002).
- H. T. Lin, Z. Lu, S. J. Kim, and M. S. Brown, "Nonuniform lattice regression for modeling the camera imaging pipeline," *European Conf. on Computer Vision*, Lecture Notes in Computer Science (Springer, 2012), Vol. 7572, pp. 556–568.
- D. H. Foster, K. Amano, S. M. C. Nascimento, and M. J. Foster, "Frequency of metamerism in natural scenes," *J. Opt. Soc. Am. A* **23**, 2359–2372 (2006).
- G. Finlayson, M. M. Darrodi, and Michal Mackiewicz, "Rank-based camera spectral sensitivity estimation," *J. Opt. Soc. Am. A* **33**, 589–599 (2016).
- P. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," *SIGGRAPH* (ACM, NY, 1997), pp. 369–378.
- S. Lin, J. Gu, S. Yamazaki, and H.-Y. Shum, "Radiometric calibration from a single image," *Proc. IEEE Computer Society Conf. on Computer Vision and Pattern Recognition* (IEEE, Piscataway, NJ, 2004), pp. II–938–II–945.
- H. Gong, G. D. Finlayson, and M. M. Darrodi, "Concise radiometric calibration using the power of ranking," *British Machine Vision Conf.* (BMVA, Durham, England, 2017).
- P. E. Gill, W. Murray, and M. H. Wright, *Practical Optimization* (Academic Press Inc., Harcourt Brace Jovanovich Publishers, London, England, 1981).
- M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars, *Computational Geometry: Algorithms and Applications*, 3rd ed. (Springer, TELOS, Berlin, Germany, 2008).
- G. D. Finlayson, "Color in perspective," *IEEE Trans. Pattern Anal. Mach. Intell.* **18**, 1034–1038 (1996).
- E. A. Rakhmanov, E. B. Saff, and Y. M. Zhou, "Minimal discrete energy on a sphere," *Math. Res. Lett.* **1**, 647–662 (1994).
- E. K. Garcia, R. Arora, and M. R. Gupta, "Optimized regression for efficient function evaluation," *IEEE Trans. Image Process.* **21**, 4128–4140 (2012).
- D. A. Forsyth, "A novel algorithm for color constancy," *Int. J. Comput. Vis.* **5**, 5–35 (1990).
- A. R. Smith, "Color gamut transform pairs," *ACM Siggraph Computer Graphics* **12**, 12–19 (1978).