

# Interactive White Balancing for Camera-Rendered Images

Mahmoud Afifi and Michael S. Brown; York University, Toronto


## Abstract

White balance (WB) is one of the first photo-finishing steps used to render a captured image to its final output. WB is applied to remove the color cast caused by the scene's illumination. Interactive photo-editing software allows users to manually select different regions in a photo as examples of the illumination for WB correction (e.g., clicking on achromatic objects). Such interactive editing is possible only with images saved in a RAW image format. This is because RAW images have no photo-rendering operations applied and photo-editing software is able to apply WB and other photo-finishing procedures to render the final image. Interactively editing WB in camera-rendered images is significantly more challenging. This is because the camera hardware has already applied WB to the image and subsequent nonlinear photo-processing routines. These nonlinear rendering operations make it difficult to change the WB post-capture. The goal of this paper is to allow interactive WB manipulation of camera-rendered images. The proposed method is an extension of our recent work [6] that proposed a post-capture method for WB correction based on nonlinear color-mapping functions. Here, we introduce a new framework that links the nonlinear color-mapping functions directly to user-selected colors to enable interactive WB manipulation. This new framework is also more efficient in terms of memory and run-time (99% reduction in memory and  $3\times$  speed-up). Lastly, we describe how our framework can leverage a simple illumination estimation method (i.e., gray-world) to perform auto-WB correction that is on a par with the WB correction results in [6].

## 1. Introduction

A number of photo-finishing steps are applied onboard a camera to render the RAW sensor image to the final output images. One of the earliest and most critical steps in this rendering process is the white-balance (WB) correction that is applied to remove color casts caused by the scene illumination. The challenge for performing WB correctly is estimating the scene's illumination from a captured RAW image. Illumination estimation is performed by the camera's auto-white-balance (AWB) module that estimates a 3D color vector representing the color cast caused by the illumination. The WB correction itself is a simple diagonal matrix operation applied to scale each RGB color channel based on the estimated illumination vector. If WB is applied incorrectly, or based on an estimated illumination that the user does not prefer, the image will have a noticeable color cast that is difficult to adjust post-capture. This difficulty arises because several nonlinear photo-finishing operations are applied on the camera after the WB operation.

There is photo-editing software (e.g., Adobe Lightroom [28], Skylum [2], Affinity Photo [1]) that enables interactive WB manipulation. Instead of applying AWB, these methods allow the user to select a pixel's RGB values in the image of achromatic scene materials to serve as the estimated illumination color vector. In some scenes, there may be more than one illuminant present and the users can choose which illumination they prefer to correct (see Fig. 1). This interactive WB editing, how-

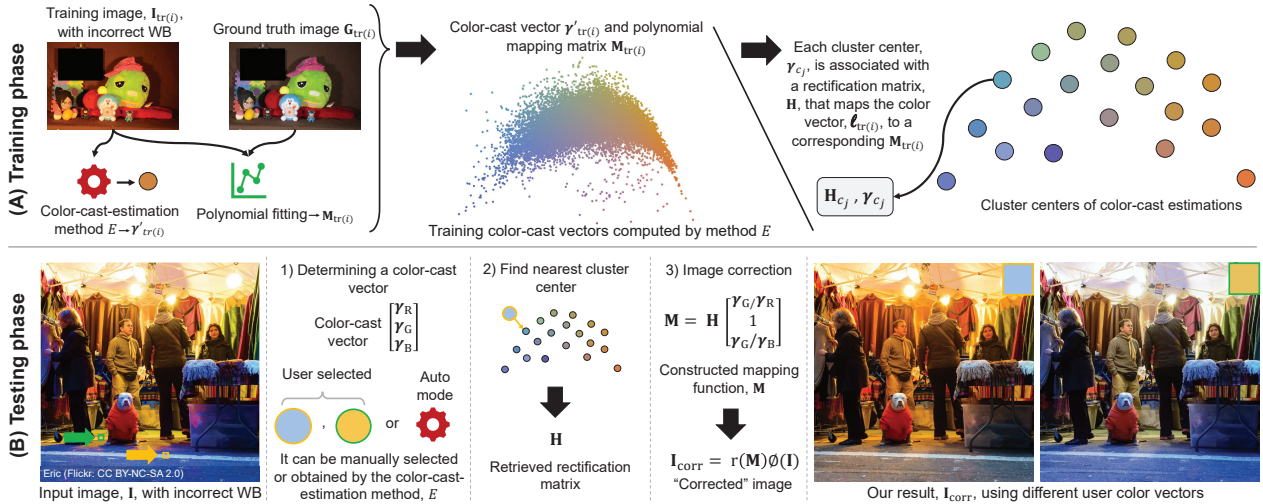


**Figure 1.** (A) A camera-rendered image with the wrong WB applied. The user selects two reference colors in the scene representing achromatic scene materials. (B) The results correction using the conventional diagonal WB correction matrix. Due to the nonlinear camera-rendering, the conventional approach is not sufficient to correct the WB. (C) Our method's AWB on (A). (D) Our method's results using the user-supplied reference colors.

ever, is possible only for photos saved in a RAW image format. This is because RAW images have no photo-finishing applied—instead, the photo-editing software mimics the onboard camera rendering using the user-supplied parameters.

The goal of this paper is to allow interactive WB manipulation for camera-rendered images. As previously mentioned, WB manipulation in camera-rendered images is challenging due to the nonlinear operations applied by the camera hardware. In recent work [6], we showed that even when an exact achromatic reference scene point is known in a camera-rendered image, the conventional diagonal correction method cannot sufficiently remove the color casts caused by WB errors (see Fig. 1). To address this issue, [6] proposed an effective nonlinear polynomial color correction function in lieu of the conventional diagonal WB matrix. The work in [6] used a histogram feature computed from an input image to determine which nonlinear correction function to use to correct a camera-rendered image that had the wrong WB applied. The work in [6], however, provided no mechanism to link the nonlinear WB color-mapping functions to the user's selected pixel values.

**Contribution** We build upon the idea in [6] and propose a new framework that allows interactive WB editing in camera-rendered images. Our approach works by associating color-cast vectors with rectification functions that output nonlinear color-mapping functions to correct the camera-rendered image's WB. This allows the user to supply a color vector from the image that results in a nonlinear color mapping to modify the image's WB based on the specified color vector. Our framework requires only 1% of the memory used in [6] and runs  $3\times$  faster, allowing interactive functionality (see Fig. 1). In addition, we show that our framework can be used to perform AWB correction for camera-rendered images that were incorrectly white-balanced with results on a par with [6].



**Figure 2.** Overview of our proposed method. (A) [Training Phase] For each image in the training dataset, we estimate a color-cast vector – shown here projected into the  $r_g$  chromaticity space to aid visualization – using an off-the-shelf-illumination estimation algorithm  $E$  (e.g., gray-world [15]). A nonlinear color-mapping function that maps this training image to a correctly white-balanced image is also computed. The color-cast vectors of each training image are clustered. A rectification function is computed for each cluster that returns a color-mapping function based on a color-cast vector. (B) [Testing Phase] When applying our method, the user either manually provides a color-cast vector or uses the method  $E$  to predict a color-cast vector. Using this color-cast vector, the most similar cluster in the training data is found and its rectification function is used to compute a mapping function  $M$  that is applied to correct the image.

## 2. Related Work

We first discuss related work targeting auto WB correction in Sec. 2.1, followed by work related to user-interactive WB correction approaches in Sec. 2.2.

### 2.1. Auto White-Balance Correction

Cameras perform WB correction to remove color casts caused by scene illumination. This procedure is applied to the linear scene-referred images (i.e., RAW sensor images) and is either performed in a fully automated mode or based on a manually selected WB pre-set (e.g., Daylight, Shade, Tungsten, Fluorescent). In the AWB mode, cameras employ illuminant estimation algorithms to estimate the scene illumination color; representative examples of illuminant estimation methods include [4, 11–18, 20–22, 24, 25, 27, 29, 31–33]. After determining the scene illuminant color, a linear diagonal correction matrix is used to remove the effect of the scene illuminant color [23]. Afterwards, cameras apply a set of operations to render the RAW values into the final display-referred colors represented in one of the standard display-referred color spaces—for example, the standard RGB (sRGB) space [8].

Once the image is rendered it is difficult to modify the image’s WB to either correct an improperly applied WB or adjust the image for aesthetic purposes [5, 6]. This difficulty is due to the nonlinear camera-specific rendering procedures applied after the WB correction. As a consequence, utilizing such linear correction for camera-rendered image white balancing mostly fails to produce satisfactory results, as discussed in [6].

Recently, we showed in [6] that correcting WB for camera-rendered images can be achieved using nonlinear color-mapping functions or via the nonlinearity provided in deep learning (DL) networks [5]. Our work is inspired by the effectiveness of these nonlinear color-mapping functions to adjust WB in camera-rendered images. However, in contrast to [6], our method provides a mechanism to link these mapping functions to color vectors found in the images. This enables the interactive editing ability that is the focus of this paper.

### 2.2. User-Interactive White-Balance Correction

Prior work dedicated to post-capture interactive WB correction does allow interactive WB modification, but only in the form of a color temperature slider [5, 7]. For example, Afifi et al. [7] proposed embedding metadata in camera-rendered images at capture time that allowed for post-capture WB modification. In their method, a color temperature slider is used to interpolate between different pre-defined color temperatures. Later, Afifi and Brown [5] introduced a deep-learning (DL) framework for WB correction and modification. The DL framework “re-renders” the input image with two different indoor/outdoor WB settings and allows the user to manipulate a temperature slider to interpolate between the two indoor/outdoor WB images.

Abdelhamed et al. [3] showed through a set of user studies that the color temperature slider is less preferred compared to directly marking achromatic reference points present in the scene. This approach allows the user to mark a known achromatic reference point that can represent current scene illumination color and was used directly by the camera’s hardware to render the image. Such manual selection cannot properly work when the image is rendered with strong color casts caused by camera WB errors, as demonstrated in [6]. In contrast to existing approaches for post-capture user-interactive WB correction, our method allows the user to select any achromatic reference point to change the WB setting of the camera-rendered image.

## 3. Methodology

Given an input image,  $I$ , that is rendered with an incorrect or undesired WB setting, the goal is to generate a new output image,  $I_{\text{corr}}$ , that represents the input  $I$  as it would appear if re-rendered with a new (presumably correct or desired) WB setting. We will refer to the target “ground truth” white-balanced image as  $G$ . In the remaining part of this paper, each image is represented as  $3 \times N$  matrix of the R, G, B triplets, where  $N$  is the total number of pixels in the image.

A traditional diagonal-based solution relies on determining a 3D vector  $\boldsymbol{\gamma} = [\boldsymbol{\gamma}^{(R)}, \boldsymbol{\gamma}^{(G)}, \boldsymbol{\gamma}^{(B)}]^T$  that represents the scene illuminant color. Since in our application, we are applying a cor-



**Figure 3.** This figure shows examples from the *Rendered WB dataset* [6] used in order to generate our training rectification functions. (A) Three examples of the same scene rendered with different incorrect WB settings. (B) The ground truth white-balanced image.

rection not of the scene illumination but instead to a color cast present in the camera-rendered image with the wrong WB applied, we will refer to  $\boldsymbol{\gamma}$  as a color-cast vector instead of an illumination vector. In our case, this color-cast vector is provided manually by the user or based on an algorithm when in AWB mode. Given a color-cast vector, the image is assumed to be corrected using the following equation:

$$\mathbf{I}_{\text{corr (diag)}} = \text{diag}(\boldsymbol{\ell}) \mathbf{I}, \quad (1)$$

where  $\text{diag}(\cdot)$  is a  $3 \times 3$  diagonal matrix of the *color-cast-correction* vector,  $\boldsymbol{\ell}$ . The vector  $\boldsymbol{\ell} = [\boldsymbol{\gamma}(\text{G})/\boldsymbol{\gamma}(\text{R}), 1, \boldsymbol{\gamma}(\text{G})/\boldsymbol{\gamma}(\text{B})]^\top$  and represents a simple modification of the color-cast vector.

Due to the nonlinearity applied to camera-rendered images, this simple scaling operation cannot properly correct WB errors. In [6], the diagonal correction matrix was replaced with a polynomial color-mapping function that could deal with the nonlinearities in the input  $\mathbf{I}$ . This function was computed in the following form:

$$\mathbf{I}_{\text{corr (poly)}} = r(\mathbf{M}) \phi(\mathbf{I}), \quad (2)$$

where  $\phi$  is a kernel mapping function [26], s.t.  $\phi: \phi([\text{R}, \text{G}, \text{B}]^\top) \rightarrow [\text{R}, \text{G}, \text{B}, \text{RG}, \text{RB}, \text{GB}, \text{R}^2, \text{G}^2, \text{B}^2, \text{RGB}, 1]^\top$ ,  $\mathbf{M} \in \mathbb{R}^{33}$  is a vectorized form of the polynomial mapping matrix, and  $r(\cdot)$  is a reshaping function that constructs the  $3 \times 11$  matrix from the vectorized version,  $\mathbf{M}$ . This polynomial matrix can be computed in a closed-form via standard least squares methods.

Though this nonlinear color mapping achieves better results compared to the diagonal-based correction, it lacks a correlation to a color-cast vector. In the following section, we describe how to efficiently associate these color-mapping functions with color vectors that can be used for WB correction in a camera-rendered image  $\mathbf{I}$ . Fig. 2 provides an overview of our framework.

### 3.1. Training Phase

The first step of our method is to have a large number of training examples that exhibits a wide range of WB errors in camera-rendered images. We used the *Rendered WB dataset* proposed in [6]. This dataset contains  $\sim 65,000$  pairs of improperly white-balanced camera-rendered (sRGB) images and their corresponding ground truth white-balanced sRGB images. All images have WB settings applied to the original raw-RGB image followed by an emulation of in-camera nonlinear rendering

operations via Adobe Lightroom [28] to get the final camera-rendered image. This dataset consists of two sets: (i) the training set, referred to as Set 1, and (ii) the testing set, referred to as Set 2. Fig. 3 shows example images taken from the *Rendered WB dataset* [6].

We used all training images in Set 1 to compute the correction functions  $\mathbf{M}$  described in Eq. 2. For each pair,  $i$ , of an improperly white-balanced image,  $\mathbf{I}_{\text{tr}(i)}$ , and the corresponding ground truth image,  $\mathbf{G}_{\text{tr}(i)}$ , we compute our 33-dimensional vectorized polynomial mapping matrix  $\mathbf{M}_{\text{tr}(i)}$  as follows:

$$\arg \min_{\mathbf{M}_{\text{tr}(i)}} \left\| r(\mathbf{M}_{\text{tr}(i)}) \phi(\mathbf{I}_{\text{tr}(i)}) - \mathbf{G}_{\text{tr}(i)} \right\|_{\text{F}}, \quad (3)$$

where  $\|\cdot\|_{\text{F}}$  is the Frobenius norm. Afterwards, each training image,  $\mathbf{I}_{\text{tr}(i)}$ , is associated with a color-cast vector  $\boldsymbol{\gamma}'_{\text{tr}(i)}$ , which will be later used to compute a color-cast-correction vector  $\boldsymbol{\ell}_{\text{tr}(i)}$ . This color-cast vector can be computed using any off-the-shelf illuminant estimation algorithm,  $E: E(\mathbf{I}_{\text{tr}(i)}) \rightarrow \boldsymbol{\gamma}'_{\text{tr}(i)}$ .

We cluster the training data based on their color-cast vectors into  $k$  clusters. In our experiments, we used k-means++ [9] with a cosine similarity distance metric and set  $k$  to 50. Each cluster, noted as  $\mathbf{c}$ , also has a number of  $\mathbf{M}_i$  mapping functions associated with it, where  $i \in \mathbf{c}$ . Instead of storing all of these mapping functions, we derive a single mapping function, termed a *rectification function*, that can estimate the color-mapping function based on the polynomial matrix  $\mathbf{M}$ . This is described in the following section.

### 3.2. Rectification Function

Our rectification function is inspired by a bias-correction method proposed to rectify scene illuminant estimation errors [20]. Specifically, we propose a rectification function,  $\mathbf{H}$ , that maps a color-cast-correction vector,  $\boldsymbol{\ell}$ , directly to a nonlinear correction matrix as follows:

$$\mathbf{M} = \mathbf{H} \boldsymbol{\ell}, \quad (4)$$

where  $\mathbf{M}$  is a  $33 \times 1$  vectorized polynomial matrix computed to map the colors of an incorrectly white-balanced image,  $\mathbf{I}$ , into the corresponding colors of the correctly white-balanced image,  $\mathbf{G}$ ,  $\mathbf{H}$  is our  $33 \times 3$  rectification matrix, and  $\boldsymbol{\ell}$  is the color-cast-correction vector computed from the color-cast vector as described earlier. This type of correction function was used by Finlayson [20] to correct biases made by illumination estimation algorithms. In [20], an estimated illumination vector would be mapped to a new illumination vector that was closer to the ground truth based on the training dataset. In our case, the function  $\mathbf{H}$  maps  $\boldsymbol{\ell}$  to its corresponding matrix  $\mathbf{M}$ , allowing us to connect a color-cast vector directly to a mapping function.

Working from Eq. 4, we compute a rectification matrix for each cluster, denoted as  $\mathbf{H}_{\mathbf{c}(j)}$  for cluster  $j$ . Let  $n$  be the number of training examples belonging to each cluster  $\mathbf{c}(j)$ , s.t.  $j \in [1, 2, \dots, k]$ . For each cluster, we minimize the following equation to compute its rectification matrix:

$$\arg \min_{\mathbf{H}_{\mathbf{c}(j)}} \left\| \sum_{i \in \mathbf{c}(j)} \mathbf{H}_{\mathbf{c}(j)} \boldsymbol{\ell}_i - \mathbf{M}_i \right\|_{\text{F}}, \quad (5)$$

where  $\boldsymbol{\ell}_i$  is a  $3 \times 1$  color-cast-correction vector of the  $i^{\text{th}}$  training example in cluster  $\mathbf{c}(j)$ ,  $\mathbf{M}_i \in \mathbb{R}^{33}$  is a vectorized polynomial matrix associated with the  $i^{\text{th}}$  training example in cluster  $\mathbf{c}(j)$ , and  $\mathbf{H}_{\mathbf{c}(j)}$  is a  $33 \times 3$  rectification matrix assigned to cluster  $\mathbf{c}(j)$ .



Eq. 5 essentially estimates a single  $\mathbf{H}$  per cluster that minimizes the error over all  $\ell_i$  associated with this cluster.

After this procedure, each color-cast cluster is now represented by the mean color-cast vector,  $\boldsymbol{\gamma}_{c_j}$ , of all color-cast vectors that belong to it. For each cluster, we store the corresponding rectification matrix,  $\mathbf{H}_{c(j)}$ , to be used in the testing phase. This model requires only 0.04 MB of memory to encode the 50 rectification functions, compared to  $\sim 25$  MB required in [6]. This represents a  $\sim 99\%$  reduction in memory requirements.

### 3.3. Testing Phase

Our method can easily be used in an AWB mode. To do this, given an input image  $\mathbf{I}$ , the same illuminant estimation algorithm,  $E$ , used in the training phase to compute the illuminant vector,  $\boldsymbol{\gamma}'$ , is applied. Afterwards, we search among the pre-computed color-cast cluster centers to find the closest cluster  $c(h)$  to our testing color vector. Then, the rectification function of the closest cluster is retrieved and used along with the computed color-cast-correction vector,  $\ell'$ , to correct the testing image as described in the following equation:

$$\mathbf{I}_{\text{CORR}} = r \left( \mathbf{H}_{c(h)} \ell' \right) \phi(\mathbf{I}). \quad (6)$$

In the user-interactive mode, we use the color-cast vector supplied by the user of some achromatic reference point in the image instead of the estimated color vector  $\boldsymbol{\gamma}'$ . Our method requires  $\sim 0.5$  seconds to correct a 12-mega-pixel image on an Intel<sup>®</sup> Xeon<sup>®</sup> E5-1607 @ 3.10 GHz machine, compared to 1.5 seconds required in [6] using a similar machine. Our approach can significantly improve the results of diagonal-based methods (e.g., [15, 21, 27, 32]) to correct improperly white-balanced images.

## 4. Experimental Results

We evaluate our method extensively through quantitative and qualitative comparisons with existing solutions for AWB and user-interactive WB correction. Sec. 4.1 provides quantitative evaluation of our method, while qualitative comparisons are provided in Sec. 4.2.

### 4.1. Quantitative Evaluation

We evaluated our AWB correction results using  $\sim 13,000$  testing images from the Rendered WB testing set (Set 2) [6] and the rendered version of the Cube+ dataset [6, 10]. As mentioned in Sec. 3, any off-the-shelf illumination estimation method  $E$  can be used. We tested different illuminant estimation methods in our AWB framework. Specifically, we utilized the following methods: the gray-world (GW) [15], the shades-of-gray (SoG) [21], and the FC4 [27] methods. In each experiment, we use the training data of the Rendered WB dataset [6] to compute our rectification functions, as described in Sec. 3.2. For each of the illuminant estimation methods, we compare our results with the diagonal correction with and without linearizing the testing images. The linearizing process was performed using the standard de-gamma linearization operation [8, 19]. We include this linearization process in our comparisons as it is a common misconception that a simple gamma linearization can remove the nonlinearity applied by cameras.

For the sake of completeness, we also compare our results with the recent nonlinear method for post-capture WB correction [6]. In this section, we refer to Afifi et al.'s method as the *WB-sRGB* method. Table 1 shows the first, second, and third quantile

and the mean of the error values obtained by each method. We followed the evaluation metrics used in [6], which are: (i) mean square error (MSE), (ii) mean angular error (MAE), and (iii)  $\Delta E$  2000 [30].

As shown in Table 1, our rectification function significantly improves the results of diagonal-based methods and achieves results on a par with the WB-sRGB method on both testing sets. As can be seen from the results, our method reduces the MSE by  $\sim 55\%$ , the MAE by  $\sim 32\%$ , and the  $\Delta E$  2000 by  $\sim 35\%$  on average compared to the gamma linearization process, which reduces the MSE by  $\sim 15\%$ , the MAE by  $\sim 5\%$ , and the  $\Delta E$  2000 by  $\sim 5\%$  on average.

### 4.2. Qualitative Evaluation

We qualitatively evaluated our method against different methods, including a commercial photo-editing software, for auto and user-interactive WB correction. Fig. 4 shows a comparison between our results and the diagonal correction of two illuminant estimation methods—namely, the GW method [15] and the quasi-unsupervised color constancy method [13]. As shown, our AWB is superior to diagonal WB and achieves similar results to the recent WB-sRGB method [6].

As previously mentioned, in contrast to the WB-sRGB method [6], our method allows interactive correction to improve the results by manually adjusting the color-cast color. Fig. 5 shows that this manual adjustment feature produces arguably visually superior results compared to the WB-sRGB method [6].

We further compared our method against Adobe Lightroom, as it is one of the most common photo-editing software programs that provide the same manual WB correction feature. As can be seen in Fig. 6, our method produces perceptibly superior results in comparison with Adobe Lightroom's results for both auto and manual WB correction.

## 5. Conclusion

We have introduced an interactive WB method for use on camera-rendered images that allows the user to directly specify scene points to be used for white balancing. Previously, this type of interaction could be performed only by photo-editing software operating on RAW images. We have enabled this feature for camera-rendered images that already have a white-balance correction applied as well as additional photo-finishing. Our method works by efficiently computing nonlinear color-correction mappings based on user-supplied color-cast vectors directly from the camera-rendered image. We also showed how our method can easily perform auto-WB correction in a camera-rendered image. Our method enables a new photo-editing feature for color manipulation.

## Acknowledgements


This study was funded in part by the Canada First Research Excellence Fund for the Vision: Science to Applications (VISTA) programme and an NSERC Discovery Grant.

## References

- [1] Affinity Photo – Professional image editing software. <https://affinity.serif.com>. Accessed: 2020-04-26.
- [2] Skylum. <https://skylum.com>. Accessed: 2020-04-26.
- [3] A. Abdelhamed, I. S. MacKenzie, and M. S. Brown. Mark-white: An improved interactive white-balance method for smartphone cameras. In *Graphics Interface*, 2019.



**Figure 4.** Qualitative results on the Rendered WB dataset [6] and the rendered version of the Cube+ dataset [6, 10]. (A) Input images. (B) Diagonal correction using the gray-world (GW) method [15]. (C) Diagonal correction using the quasi-supervised color constancy (quasi-U CC) method [13]. (D) Results of the WB-sRGB method [6]. (E) Results of applying our rectification function to initial estimation of GW. (F) Ground truth images.



**Figure 5.** Our method allows the user to manually select achromatic points in order to improve the results. (A) Input images. (B) Results of the gray-world (GW) method [15]. (C) Results of the quasi-supervised color constancy (quasi-U CC) method [13]. (D) Results of the WB-sRGB method [6]. (E) Our results using the GW's estimated illuminant colors. (F) Our results using manually selected achromatic reference points.

- [4] M. Afifi and M. S. Brown. Sensor-independent illumination estimation for dnn models. In *BMVC*, 2019.
- [5] M. Afifi and M. S. Brown. Deep white-balance editing. In *CVPR*, 2020.
- [6] M. Afifi, B. Price, S. Cohen, and M. S. Brown. When color constancy goes wrong: Correcting improperly white-balanced images. In *CVPR*, 2019.
- [7] M. Afifi, A. Punnappurath, A. Abdelhamed, H. C. Karaimer, A. Abuolaim, , and M. S. Brown. Color temperature tuning: Allowing accurate post-capture white-balance editing. In *Color Imaging Conference*, 2019.
- [8] M. Anderson, R. Motta, S. Chandrasekar, and M. Stokes. Proposal for a standard default color space for the internet - sRGB. In *Color and Imaging Conference*, pages 238–245, 1996.
- [9] D. Arthur and S. Vassilvitskii. K-means++: The advantages of careful seeding. In *Annual ACM-SIAM Symposium on Discrete Algorithms*, 2007.
- [10] N. Banić and S. Lončarić. Unsupervised learning for color constancy. *arXiv preprint arXiv:1712.00436*, 2017.
- [11] J. T. Barron. Convolutional color constancy. In *ICCV*, 2015.
- [12] J. T. Barron and Y.-T. Tsai. Fast fourier color constancy. In *CVPR*, 2017.
- [13] S. Bianco and C. Cusano. Quasi-supervised color constancy. In *CVPR*, 2019.
- [14] S. Bianco, C. Cusano, and R. Schettini. Color constancy using CNNs. In *CVPR Workshops*, 2015.
- [15] G. Buchsbaum. A spatial processor model for object colour perception. *Journal of the Franklin Institute*, 310(1):1–26, 1980.
- [16] V. C. Cardei and B. Funt. Committee-based color constancy. In *Color and Imaging Conference*, 1999.
- [17] V. C. Cardei, B. Funt, and K. Barnard. Estimating the scene illumination chromaticity by using a neural network. *Journal of the Optical Society of America A (JOSA A)*, 19(12):2374–2386, 2002.
- [18] D. Cheng, B. Price, S. Cohen, and M. S. Brown. Effective learning-based illuminant estimation using simple features. In *CVPR*, 2015.
- [19] M. Ebner. *Color Constancy*, volume 6. John Wiley & Sons, 2007.
- [20] G. D. Finlayson. Corrected-moment illuminant estimation. In *ICCV*, 2013.
- [21] G. D. Finlayson and E. Trezzi. Shades of gray and colour constancy. In *Color and Imaging Conference*, 2004.
- [22] P. V. Gehler, C. Rother, A. Blake, T. Minka, and T. Sharp. Bayesian color constancy revisited. In *CVPR*, 2008.
- [23] A. Gijsenij, T. Gevers, and J. Van De Weijer. Computational color constancy: Survey and experiments. *IEEE Transactions on Image Processing*, 20(9):2475–2489, 2011.
- [24] A. Gijsenij, T. Gevers, and J. Van De Weijer. Improving color constancy by photometric edge weighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(5):918–929, 2012.
- [25] D. Hernandez-Juarez, S. Parisot, B. Busam, A. Leonardis, G. Slabaugh, and S. McDonagh. A multi-hypothesis approach to color constancy. *arXiv preprint arXiv:2002.12896*, 2020.
- [26] G. Hong, M. R. Luo, and P. A. Rhodes. A study of digital camera colorimetric characterisation based on polynomial

Table 1. Quantitative results on the Rendered WB testing set (Set 2) [6] and the rendered version of the Cube+ dataset [6, 10]. We applied our rectification function (RF) to the estimated illuminants of different methods. The term “linearized” refers to applying the standard gamma linearization [8, 19] to images before estimating and correcting images. The terms Q1, Q2, and Q3 denote the first, second (median), and third quartile, respectively. The terms MSE and MAE stand for mean square error and mean angular error, respectively. The best results are highlighted in yellow and boldface. As shown, we achieve on par results with the WB-sRGB method [6] by applying our post-processing rectification to different illuminant estimation methods. In contrast to the WB-sRGB method [6], our method allows the user to interactively adjust the results by manually changing the selected achromatic reference point.

Method	MSE				MAE				$\Delta E$ 2000			
	Mean	Q1	Q2	Q3	Mean	Q1	Q2	Q3	Mean	Q1	Q2	Q3
<b>Testing set of the Rendered WB dataset (Set 2): DSLR and mobile phone cameras (2,881 images)</b>												
GW [15]	500.18	173.69	332.75	615.40	8.89°	5.82°	8.32°	11.33°	10.74	7.92	10.29	13.12
SoG [21]	429.35	147.05	286.84	535.72	9.54°	5.72°	8.85°	12.65°	10.01	7.09	9.85	12.69
FC4 [27]	662.53	304.88	524.42	817.57	8.92°	5.94°	8.03°	10.84°	12.12	8.94	11.79	14.76
GW (linearized) [15]	469.86	163.07	312.28	574.85	8.61°	5.44°	7.94°	10.93°	10.68	7.70	10.13	13.15
SoG (linearized) [21]	393.85	137.21	267.37	497.40	8.96°	5.31°	8.26°	11.97°	9.81	6.87	9.67	12.46
FC4 (linearized) [27]	505.30	142.46	307.77	635.35	10.37°	5.31°	9.26°	14.15°	10.82	7.39	10.64	13.77
GW [15] + our RF	207.13	46.71	111.89	230.10	5.35°	2.89°	4.59°	6.84°	6.74	4.45	6.15	8.45
SoG [21] + our RF	256.10	55.93	132.09	266.61	6.25°	3.28°	5.20°	8.20°	7.27	4.77	6.61	9.05
FC4 [27] + our RF	303.99	50.01	118.88	298.44	6.61°	2.99°	4.99°	8.40°	7.28	4.30	6.22	9.33
WB-sRGB [6]	<b>171.09</b>	<b>37.04</b>	<b>87.04</b>	<b>190.88</b>	<b>4.48°</b>	<b>2.26°</b>	<b>3.64°</b>	<b>5.95°</b>	<b>5.60</b>	<b>3.43</b>	<b>4.90</b>	<b>7.06</b>
<b>Rendered Cube+ dataset with different WB settings (10,242 images)</b>												
GW [15]	312.62	55.16	159.63	358.02	6.85°	3.08°	5.76°	9.70°	9.01	5.35	8.38	12.08
SoG [21]	269.31	21.92	90.37	312.02	6.69°	2.3°	4.63°	9.62°	7.70	3.40	6.38	11.07
FC4 [27]	410.01	79.26	219.05	505.71	6.7°	3.26°	5.45°	8.7°	10.4	6.51	9.73	13.43
GW (linearized) [15]	244.59	32.58	121.42	300.99	6.37°	2.51°	5.13°	9.09°	8.05	4.18	7.25	11.08
SoG (linearized) [21]	275.33	17.16	67.49	309.97	6.66°	2.08°	4.17°	9.58°	7.57	3.00	5.73	11.05
FC4 (linearized) [27]	371.9	79.15	213.41	467.33	6.49°	3.34°	5.59°	8.59°	10.38	6.6	9.76	13.26
GW [15] + our RF	<b>159.88</b>	21.94	54.76	125.02	4.64°	2.12°	3.64°	5.98°	6.2	3.28	5.17	7.45
SoG [21] + our RF	226.83	20.01	58.61	165.03	5.33°	2.1°	3.83°	6.97°	6.61	3.17	5.38	8.56
FC4 [27] + our RF	175.73	<b>17.8</b>	<b>43.65</b>	<b>114.65</b>	4.67°	1.89°	<b>3.10°</b>	5.48°	5.7	<b>2.95</b>	4.63	7.05
WB-sRGB [6]	194.98	27.43	<b>57.08</b>	<b>118.21</b>	<b>4.12°</b>	<b>1.96°</b>	3.17°	<b>5.04°</b>	<b>5.68</b>	3.22	<b>4.61</b>	<b>6.70</b>




Figure 6. Comparison with Adobe Lightroom WB correction. (A) Input image. (B) and (C) Adobe Lightroom’s auto WB and manual WB correction results, respectively. (D) and (E) Our auto and manual WB correction results, respectively.

modelling. *Color Research & Application*, 26(1):76–84, 2001.

[27] Y. Hu, B. Wang, and S. Lin. Fc4: Fully convolutional color constancy with confidence-weighted pooling. In *CVPR*, 2017.

[28] S. Kelby. *The Adobe Photoshop Lightroom 5 Book for Digital Photographers*. Pearson Education, 2014.

[29] Y. Qian, J.-K. Kamarainen, J. Nikkanen, and J. Matas. On finding gray pixels. In *CVPR*, 2019.

[30] G. Sharma, W. Wu, and E. N. Dalal. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application*, 30(1).

[31] W. Shi, C. C. Loy, and X. Tang. Deep specialized network for illuminant estimation. In *ECCV*, 2016.

[32] J. Van De Weijer, T. Gevers, and A. Gijsenij. Edge-based color constancy. *IEEE Transactions on image processing*, 16(9):2207–2214, 2007.

[33] W. Xiong and B. Funt. Estimating illumination chromaticity via support vector regression. *Journal of Imaging Science and Technology*, 50(4):341–348, 2006.