

Sensor-Independent Illumination Estimation for DNN Models

Mahmoud Afifi¹

<http://www.cse.yorku.ca/~mafifi/>

Lassonde School of Engineering¹

York University

Toronto, Canada

Michael S. Brown^{1,2}

<http://www.cse.yorku.ca/~mbrown/>

Samsung AI Center (SAIC)²

Samsung Research

Toronto, Canada

Abstract

While modern deep neural networks (DNNs) achieve state-of-the-art results for illuminant estimation, it is currently necessary to train a separate DNN for each type of camera sensor. This means when a camera manufacturer uses a new sensor, it is necessary to retrain an existing DNN model with training images captured by the new sensor. This paper addresses this problem by introducing a novel sensor-independent illuminant estimation framework. Our method learns a sensor-independent *working space* that can be used to canonicalize the RGB values of any arbitrary camera sensor. Our learned space retains the linear property of the original sensor raw-RGB space and allows unseen camera sensors to be used on a single DNN model trained on this working space. We demonstrate the effectiveness of this approach on several different camera sensors and show it provides performance on par with state-of-the-art methods that were trained per sensor.

1 Introduction and Motivation

Color constancy is the constant appearance of object colors under different illumination conditions [26]. Human vision has this illumination adaption ability to recognize the same object colors under different scene lighting [89]. Camera sensors, however, do not have this ability and as a result, computational color constancy is required to be applied onboard the camera. In a photography context, this procedure is typically termed *white balance*. The key challenge for computational color constancy is the ability to estimate a camera sensor's RGB response to the scene's illumination. Illumination estimation, or auto white balance (AWB), is a fundamental procedure applied onboard all cameras and is critical in ensuring the correct interpretation of scene colors.

Computational color constancy can be described in terms of the physical image formation process as follows. Let $\mathbf{I} = \{\mathbf{I}_r, \mathbf{I}_g, \mathbf{I}_b\}$ denote an image captured in the linear raw-RGB space. The value of each color channel $c = \{R, G, B\}$ for a pixel located at x in \mathbf{I} is given by the following equation [9]:

$$\mathbf{I}_c(x) = \int_{\gamma} \rho(x, \lambda) R(x, \lambda) S_c(\lambda) d\lambda, \quad (1)$$




Figure 1: A scene captured by two different camera sensors results in different ground truth illuminants due to different camera sensor responses. We learn a device-independent *working space* that reduces the difference between ground truth illuminants of the same scenes.

where γ is the visible light spectrum (approximately 380nm to 780nm), $\rho(\cdot)$ is the illuminant spectral power distribution, $R(\cdot)$ is the captured scene’s spectral reflectance properties, and $S(\cdot)$ is the camera sensor response function at wavelength λ . The problem can be simplified by assuming a single uniform illuminant in the scene as follows:

$$\mathbf{I}_c = \ell_c \times \mathbf{R}_c, \quad (2)$$

where ℓ_c is the scene illuminant value of color channel c . A standard approach to this problem is to use a linear model (i.e., a 3×3 diagonal matrix) such that $\ell_R = \ell_G = \ell_B$ (i.e., white illuminant).

Typically, ℓ is unknown and should be defined to obtain the true objects’ body reflectance values \mathbf{R} in the input image \mathbf{I} . The value of ℓ is specific to the camera sensor response function $S(\cdot)$, meaning that the same scene captured by different camera sensors results in different values of ℓ . Fig. 1 shows an example.

Illuminant estimation methods aim to estimate the value ℓ from the sensor’s raw-RGB image. Recently, deep neural network (DNN) methods have demonstrated state-of-the-art results for the illuminant estimation task. These approaches, however, need to train the DNN model per camera sensor. This is a significant drawback. When a camera manufacturer decides to use a new sensor, the DNN model will need to be retrained on a new image dataset captured by the new sensor. Collecting such datasets with the corresponding ground-truth illuminant raw-RGB values is a tedious process. As a result, many AWB algorithms deployed on cameras still rely on simple statistical-based methods even though the accuracy is not comparable to those obtained by the learning-based methods [9].

Contribution In this paper, we introduce a sensor-independent learning framework for illuminant estimation. The idea is similar to the color space conversion process applied on-board cameras that maps the sensor-specific RGB values to a perceptual-based color space – namely, CIE XYZ. The color space conversion process estimates a color space transform (CST) matrix to map white-balanced sensor-specific raw-RGB images to CIE XYZ [65, 94]. This process is applied onboard cameras *after* the illuminant estimation and white-balance step, and relies on the estimated scene illuminant to compute the CST matrix [10]. Our solution, however, is to learn a new space that is used *before* the illuminant estimation step. Specifically, we design a novel unsupervised deep learning framework that learns how to map each input image, captured by arbitrary camera sensor, to a non-perceptual sensor-independent *working space*. Mapping input images to this space allows us to train our model using training sets captured by different camera sensors, achieving good accuracy and generalizing well for unseen camera sensors, as shown in Fig. 2.




Figure 2: (A) Traditional learning-based illuminant estimation methods train or fine-tune a model per camera sensor. (B) Our method can be trained on images captured by different camera sensors and generalizes well for unseen camera sensors. Shown images are rendered in the sRGB color space by the camera imaging pipeline in [65] to aid visualization.

2 Related Work

We discuss two areas of related work: (i) illumination estimation and (ii) mapping camera sensor responses.

2.1 Illumination Estimation

As previously discussed, illuminant estimation is the key routine that makes up a camera’s AWB function. Illuminant estimation methods aim to estimate the illumination in the imaged scene directly from a raw-RGB image without a known achromatic reference scene patch. We categorize the illuminant estimation methods into two categories, which are: (i) sensor-independent methods and (ii) sensor-dependent methods.

Sensor-Independent Methods These methods operate using statistics from an image’s color distribution and spatial layout to estimate the scene illuminant. Representative statistical-based methods include: Gray-World [14], White-Patch [13], Shades-of-Gray [20], Gray-Edges, and PCA-based Bright-and-Dark Colors [16]. These methods are fast and easy to implement; however, their results are not always satisfactory.

Sensor-Dependent Methods Learning-based models outperform statistical-based methods by training sensor-specific models on training examples provided with the labeled images with ground-truth illumination obtained from physical charts placed in the scene with achromatic reference patches. These training images are captured with the sensor make and model being trained. Representative examples include Bayesian-based methods [12, 29, 45], gamut-based methods [23, 25, 30], exemplar-based methods [5, 30, 34], bias-correction methods [0, 19, 27], and, more recently, DNN methods [0, 8, 33, 38, 40, 46], including few-shot learning [40]. The obvious drawback of these methods is that they do not generalize well for arbitrary camera sensors without retraining/fine-tuning on samples captured by testing camera sensor. Our learning method, however, is explicitly designed to be sensor-independent and generalizes well for unseen camera sensors without the need to retrain/tune our model.




Figure 3: Our proposed method consists of two networks: (i) a sensor mapping network and (ii) an illuminant estimation network. Our networks are trained jointly in an end-to-end manner to learn an image-specific mapping matrix (resulting from the sensor mapping network) and scene illuminant in the learned space (resulting from the illuminant estimation network). The final estimated illuminant is produced by mapping the result illuminant from our learned space to the input image’s camera-specific raw space.

2.2 Mapping Camera Sensor Responses

Another research topic related to our work is mapping camera raw-RGB sensor responses to a perceptual color space. This process is applied onboard digital cameras to map the captured sensor-specific raw-RGB image to a standard device-independent “canonical” space (e.g., CIE XYZ) [65, 44]. Usually this conversion is performed using a 3×3 matrix and requires an accurate estimation of the scene illuminant [15]. It is important to note that this mapping to CIE XYZ requires that white-balance procedure first be applied. As a result, it is not possible to use CIE XYZ as the canonical color space to perform illumination estimation.

Work by Nguyen et al. [40] studied several transformations to map responses from a source camera sensor to a target camera sensor, instead of mapping to a perceptual space. In their study, a color rendition reference chart is captured by both source and target camera sensors in order to compute the raw-to-raw mapping function. Learning a mapping transformation between responses of two different sensors is also adapted in [27]. While our approach is similar in this goal, the work in [27, 40] has no mechanism to map an unseen sensor to a canonical working space without explicit calibration.

3 Proposed Method

Fig. 3 provides an overview of our framework. Our method accepts thumbnail (150×150 pixels) linear raw-RGB images, captured by an arbitrary camera sensor, and estimates scene illuminant RGB vectors in the same space of input images.

We rely on color distribution of input thumbnail image I to estimate an image-specific transformation matrix that maps the input image to our working space. This mapping allows us to accept images captured by different sensors and estimate scene illuminant values in the original space of input images.

We begin with the formulation of our problem followed by a detailed description of our framework components and the training process. Note that we will assume input raw-RGB images are represented as $3 \times n$ matrices, where $n = 150 \times 150$ is the total number of pixels in the thumbnail image and the three rows represent the R, G, and B values.

3.1 Problem Formulation

We propose to work in a new learned space for illumination estimation. This space is sensor-independent and retains the linear property of the original raw-RGB space. To that end, we introduce a learnable 3×3 matrix \mathbf{M} that maps an input image \mathbf{I} from its original sensor-specific space to a new working space. We can reformulate Eq. 2 as follows:

$$\mathbf{M}^{-1}\mathbf{M}\mathbf{I} = \text{diag}(\mathbf{M}^{-1}\mathbf{M}\boldsymbol{\ell}), \quad (3)$$

where $\text{diag}(\cdot)$ is a diagonal matrix and \mathbf{M} is a learned matrix that maps arbitrary sensor responses to a sensor-independent space.

Given a mapped image $\mathbf{I}_m = \mathbf{M}\mathbf{I}$ in our learned space, we aim to estimate the mapped vector $\boldsymbol{\ell}_m = \mathbf{M}\boldsymbol{\ell}$ that represents the scene illumination values of \mathbf{I}_m in the new space. The original scene illuminant (represented in the original sensor raw-RGB space) can be reconstructed by the following equation:

$$\boldsymbol{\ell} = \mathbf{M}^{-1}\boldsymbol{\ell}_m. \quad (4)$$

3.2 RGB-uv Histogram Block

Prior work has shown that the illumination estimation problem is related primarily to the image’s color distribution [14, 16]. Accordingly, we use the image’s color distribution as an input for our method. Representing the image using a full 3D RGB histogram requires significant amounts of memory – for example, a 256^3 RGB histogram requires more than 16 million entries. Down-sampling the histogram – for example, to 64-bins – still requires a considerable amount of memory.

Our method relies on the RGB-uv histogram feature used in [14]. This feature represents the image color distribution in the log of chromaticity space [18]. Unlike the original RGB-uv feature, we use two learnable parameters to control the contribution of each color channel in the generated histogram and the smoothness of histogram bins. Specifically, the RGB-uv histogram block represents the color distribution of an image \mathbf{I} as a three-layer histogram $\mathbf{H}(\mathbf{I})$ represented as an $m \times m \times 3$ tensor. The produced histogram $\mathbf{H}(\mathbf{I})$ is parameterized by uv and computed as follows:

$$\begin{aligned} \mathbf{I}_y(i) &= \sqrt{\mathbf{I}_{R(i)}^2 + \mathbf{I}_{G(i)}^2 + \mathbf{I}_{B(i)}^2}, \\ \mathbf{I}_{u1(i)} &= \log \left(\frac{\mathbf{I}_{R(i)}}{\mathbf{I}_{G(i)}} + \varepsilon \right), \mathbf{I}_{v1(i)} = \log \left(\frac{\mathbf{I}_{R(i)}}{\mathbf{I}_{B(i)}} + \varepsilon \right), \\ \mathbf{I}_{u2(i)} &= \log \left(\frac{\mathbf{I}_{G(i)}}{\mathbf{I}_{R(i)}} + \varepsilon \right), \mathbf{I}_{v2(i)} = \log \left(\frac{\mathbf{I}_{G(i)}}{\mathbf{I}_{B(i)}} + \varepsilon \right), \\ \mathbf{I}_{u3(i)} &= \log \left(\frac{\mathbf{I}_{B(i)}}{\mathbf{I}_{R(i)}} + \varepsilon \right), \mathbf{I}_{v3(i)} = \log \left(\frac{\mathbf{I}_{B(i)}}{\mathbf{I}_{G(i)}} + \varepsilon \right), \end{aligned} \quad (5)$$

$$\mathbf{H}(\mathbf{I})_{(u,v,c)} = \left(s_c \sum_i \mathbf{I}_{y(i)} \exp(-|\mathbf{I}_{uc(i)} - u|/\sigma_c^2) \exp(-|\mathbf{I}_{vc(i)} - v|/\sigma_c^2) \right)^{1/2},$$

where $i = \{1, \dots, n\}$, $c \in \{1, 2, 3\}$ represents each color channel in \mathbf{H} , ε is a small positive constant added for numerical stability, and s_c and σ_c are learnable scale and fall-off parameters, respectively. The scale factor s_c controls the contribution of each layer in our histogram,

while the fall-off factor σ_c controls the smoothness of the histogram’s bins of each layer. The values of these parameters (i.e., s_c and σ_c) are learned during the training phase.

3.3 Network Architecture

As shown in Fig. 3, our framework consists of two networks: (i) a sensor mapping network and (ii) an illuminant estimation network. The input to each network is the RGB- uv histogram feature produced by our histogram block. The sensor mapping network accepts an RGB- uv histogram of a thumbnail raw-RGB image \mathbf{I} in its original sensor space, while the illuminant estimation network accepts RGB- uv histograms of the mapped image \mathbf{I}_m to our learned space. In our implementation, we use $m = 61$ and each histogram feature is represented by a $61 \times 61 \times 3$ tensor.

We use a simple network architecture for each network. Specifically, each network consists of three conv/ReLU layers followed by a fully connected (fc) layer. The kernel size and stride step used in each conv layer are shown in Fig. 3.

In the sensor mapping network, the last fc layer has nine neurons. The output vector \mathbf{v} of this fc layer is reshaped to construct a 3×3 matrix \mathbf{V} , which is used to build \mathcal{M} as described in the following equation:

$$\mathcal{M} = \frac{1}{\|\mathbf{V}\|_1 + \varepsilon} |\mathbf{V}|, \quad (6)$$

where $|\cdot|$ is the modulus (absolute magnitude), $\|\cdot\|_1$ is the matrix 1-norm, and ε is added for numerical stability. The modulus step is necessary to avoid negative values in the mapped image \mathbf{I}_m , while the normalization step is used to avoid having extremely large values in \mathbf{I}_m . Note the values of \mathcal{M} are image-specific, meaning that its values are produced based on the input image’s color distribution in the original raw-RGB space.

There are three neurons in the last fc layer of the illuminant estimation network to produce illuminant vector $\hat{\ell}_m$ of the mapped image \mathbf{I}_m . Note that the estimated vector $\hat{\ell}_m$ represents the scene illuminant in our learned space. The final result is obtained by mapping $\hat{\ell}_m$ back to the original space of \mathbf{I} using Eq. 4.

3.4 Training

We jointly train our sensor mapping and illuminant estimation networks in an end-to-end manner using the adaptive moment estimation (Adam) optimizer [66] with a decay rate of gradient moving average $\beta_1 = 0.85$, a decay rate of squared gradient moving average $\beta_2 = 0.99$, and a mini-batch with eight observations at each iteration. We initialized both network weights with Xavier initialization [62]. The learning rate was set to 10^{-5} and decayed every five epochs.

We adopt the recovery angular error (referred to as the angular error) as our loss function [2]. The angular error is computed between the ground truth illuminant ℓ and our estimated illuminant $\hat{\ell}_m$ after mapping it to the original raw-RGB space of training image \mathbf{I} . The loss function can be described by the following equation:

$$L(\hat{\ell}_m, \mathcal{M}) = \cos^{-1} \left(\frac{\ell \cdot (\mathcal{M}^{-1} \hat{\ell}_m)}{\|\ell\| \|\mathcal{M}^{-1} \hat{\ell}_m\|} \right), \quad (7)$$

where $\|\cdot\|$ is the Euclidean norm, and (\cdot) is the vector dot-product.

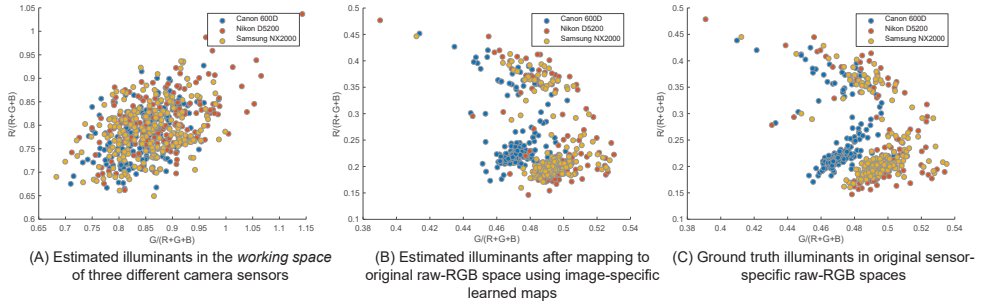


Figure 4: Raw-RGB images capture the same set of scenes using three different cameras taken from the NUS 8-Cameras dataset [16]. (A) Estimated illuminants resulted from the illuminant estimation network in our learned *working space*. (B) Estimated illuminants after mapping to the original raw-RGB space. This mapping is performed by multiplying each illuminant vector by the inverse of the learned image-specific mapping matrix (resulting from the sensor mapping network). (C) Corresponding ground truth illuminants in the original raw-RGB space of each image.

As the values of \mathbf{M} are produced by the sensor mapping network, there is a possibility of producing a singular matrix output. In this case, we add small offset $\mathcal{N}(0, 1) \times 10^{-4}$ to each parameter in \mathbf{M} to make it invertible.

At the end of the training process, our framework learns an image-specific matrix \mathbf{M} that maps an input image taken by an arbitrary sensor to the learned space. Fig. 4 shows an example of three different camera responses capturing the same set of scenes. As shown in Fig. 4-(A), the estimated illuminants of these sensors are bounded in the learned space. These illuminants are mapped back to the original raw-RGB sensor space of the corresponding input images using Eq. 4. As shown in Fig. 4-(B) and Fig. 4-(C), our final estimated illuminants are close to the ground truth illuminants of each camera sensor.

4 Experimental Results

In our experiments, we used all cameras of three different datasets, which are: (i) NUS 8-Camera [16], (ii) Gehler-Shi [29], and (iii) Cube+ [8] datasets. In total, we have 4,014 raw-RGB images captured by 11 different camera sensors.

We followed the leave-one-out cross-validation scheme to evaluate our method. Specifically, we excluded all images captured by one camera for testing and trained a model with the remaining images. This process was repeated for all cameras. We also tested our method on the Cube dataset. In this experiment, we used a trained model on images from the NUS and Gehler-Shi datasets, and excluded all images from the Cube+ dataset. The calibration objects (i.e., X-Rite color chart or SpyderCUBE) were masked out in both training and testing processes.

Unlike results reported by existing learning methods which use three-fold cross-validation for evaluation, our reported results were obtained by models that were *not* trained on any example of the testing camera sensor.

In Tables 1–2, the mean, median, best 25%, and the worst 25% of the angular error between our estimated illuminants and ground truth are reported. The best 25% and worst 25% are the mean of the smallest 25% angular error values and the mean of the highest 25% angu-

Table 1: Angular errors on the NUS 8-Cameras [16] and Gehler-Shi [29] datasets. Methods highlighted in gray are trained/tuned for each camera sensor (i.e., sensor-specific models). The lowest errors are highlighted in yellow.

NUS 8-Cameras Dataset					Gehler-Shi Dataset				
Method	Mean	Med.	Best 25%	Worst 25%	Method	Mean	Med.	Best 25%	Worst 25%
White-Patch [16]	9.91	7.44	1.44	21.27	White-Patch [16]	7.55	5.68	1.45	16.12
Pixel-based Gamut [16]	5.27	4.26	1.28	11.16	Edge-based Gamut [16]	6.52	5.04	5.43	13.58
Grey-world (GW) [16]	4.59	3.46	1.16	9.85	Grey-world (GW) [16]	6.36	6.28	2.33	10.58
Edge-based Gamut [16]	4.40	3.30	0.99	9.83	1st-order Gray-Edge [16]	5.33	4.52	1.86	10.03
Shades-of-Gray [16]	3.67	2.94	0.98	7.75	2nd-order Gray-Edge [16]	5.13	4.44	2.11	9.26
Bayesian [16]	3.50	2.36	0.78	8.02	Shades-of-Gray [16]	4.93	4.01	1.14	10.20
Local Surface Reflectance [16]	3.45	2.51	0.98	7.32	Bayesian [16]	4.82	3.46	1.26	10.49
2nd-order Gray-Edge [16]	3.36	2.70	0.89	7.14	Pixels-based Gamut [16]	4.20	2.33	0.50	10.72
1st-order Gray-Edge [16]	3.35	2.58	0.79	7.18	Quasi-unsupervised [16]	3.46	2.23	-	-
Quasi-unsupervised [16]	3.00	2.25	-	-	PCA-based B/W Colors [16]	3.52	2.14	0.50	8.74
Corrected-Moment [16]	2.95	2.05	0.59	6.89	NetColorChecker [16]	3.10	2.30	-	-
PCA-based B/W Colors [16]	2.93	2.33	0.78	6.13	Grayscale Index [16]	3.07	1.87	0.43	7.62
Grayscale Index [16]	2.91	1.97	0.56	6.67	Meta-AWB w 20 tuning images [16]	3.00	2.02	0.58	7.17
Color Dog [16]	2.83	1.77	0.48	7.04	Quasi-unsupervised [16] (tuned)	2.91	1.98	-	-
APAP using GW [16]	2.40	1.76	0.55	5.42	Corrected-Moment [16]	2.86	2.04	0.70	6.34
Conv Color Constancy [16]	2.38	1.69	0.45	5.85	APAP using GW [16]	2.76	2.02	0.53	6.21
Effective Regression Tree [16]	2.36	1.59	0.49	5.54	Bianco et al.'s CNN [16]	2.63	1.98	0.72	3.90
WB-sRGB (modified for raw-RGB) [16]	2.26	1.60	0.48	5.21	Effective Regression Tree [16]	2.42	1.65	0.38	5.87
Deep Specialized Net [16]	2.24	1.46	0.48	6.08	WB-sRGB (modified for raw-RGB) [16]	2.07	1.38	0.29	5.11
Meta-AWB w 20 tuning images [16]	2.23	1.49	0.49	5.20	Fast Fourier - thumb, 2 channels [16]	2.01	1.13	0.30	5.14
SqueezeNet-FC4 [16]	2.23	1.57	0.47	5.15	Conv Color Constancy [16]	1.95	1.22	0.35	4.76
AlexNet-FC4 [16]	2.12	1.53	0.48	4.78	Deep Specialized Net [16]	1.90	1.12	0.31	4.84
Fast Fourier - thumb, 2 channels [16]	2.06	1.39	0.39	4.80	Fast Fourier - full, 4 channels [16]	1.78	0.96	0.29	4.62
Fast Fourier - full, 4 channels [16]	1.99	1.31	0.35	4.75	AlexNet-FC4 [16]	1.77	1.11	0.34	4.29
Quasi-unsupervised (tuned) [16]	1.97	1.41	-	-	SqueezeNet-FC4 [16]	1.65	1.18	0.38	3.78
Avg. result for sensor-independent	4.26	3.25	0.99	9.43	Avg. result for sensor-independent	5.10	4.03	1.91	10.77
Avg. result for sensor-dependent	2.40	1.64	0.50	5.75	Avg. result for sensor-dependent	2.62	1.75	0.50	5.95
Sensor-independent (Ours)	2.05	1.50	0.52	4.48	Sensor-independent (Ours)	2.77	1.93	0.55	6.53

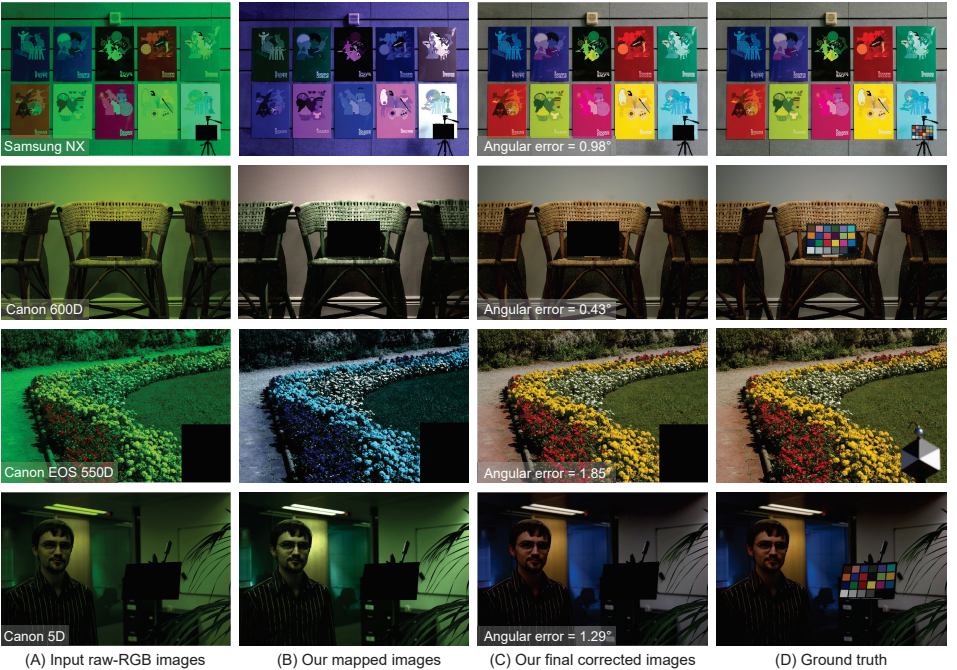


Figure 5: Qualitative results of our method. (A) Input raw-RGB images. (B) After mapping images in (A) to the learned space. (C) After correcting images in (A) based on our estimated illuminants. (D) Corrected by ground truth illuminants. Shown images are rendered in the sRGB color space by the camera imaging pipeline in [35] to aid visualization.

Table 2: Angular errors on the Cube and Cube+ datasets [8]. Methods highlighted in gray are trained/tuned for each camera sensor (i.e., sensor-specific models). The lowest errors are highlighted in yellow.

Cube Dataset Method	Mean	Med.	Best 25%	Worst 25%	Cube+ Dataset Method	Mean	Med.	Best 25%	Worst 25%
White-Patch [8]	6.58	4.48	1.18	15.23	White-Patch [8]	9.69	7.48	1.72	20.49
Grey-world (GW) [8]	3.75	2.91	0.69	8.18	Grey-world (GW) [8]	7.71	4.29	1.01	20.19
Shades-of-Gray [8]	2.58	1.79	0.38	6.19	Color Dog [8]	3.32	1.19	0.22	10.22
2nd-order Gray-Edge [8]	2.49	1.60	0.49	6.00	Shades-of-Gray [8]	2.59	1.73	0.46	6.19
1st-order Gray-Edge [8]	2.45	1.58	0.48	5.89	2nd-order Gray-Edge [8]	2.50	1.59	0.48	6.08
APAP using GW [8]	1.55	1.02	0.28	3.74	1st-order Gray-Edge [8]	2.41	1.52	0.45	5.89
Color Dog [8]	1.50	0.81	0.27	3.86	APAP using GW [8]	2.01	1.36	0.38	4.71
Meta-AWB (20) [10]	1.74	1.08	0.29	4.28	Color Beaver [8]	1.49	0.77	0.21	3.94
WB-sRGB (modified for raw-RGB) [8]	1.37	0.78	0.19	3.51	WB-sRGB (modified for raw-RGB) [8]	1.32	0.74	0.18	3.43
Avg. result for sensor-independent	3.57	2.47	0.64	8.30	Avg. result for sensor-independent	4.98	3.32	0.82	11.77
Avg. result for sensor-dependent	1.54	0.92	0.26	3.85	Avg. result for sensor-dependent	2.04	1.02	0.25	5.58
Sensor-independent (Ours)	1.98	1.36	0.40	4.64	Sensor-independent (Ours)	2.14	1.44	0.44	5.06

Table 3: Angular and reproduction angular errors [24] on the Cube+ challenge [8]. The methods are sorted by the median of the errors (shown in bold), as ranked in the challenge [8]. Methods highlighted in gray are sensor-specific models. We show our results w/o training on Cube+ dataset. The lowest errors over all methods are highlighted in yellow.

Cube+ challenge (angular error) Method	Mean	Med.	Best 25%	Worst 25%	Cube+ challenge (reproduction error) Method	Mean	Med.	Best 25%	Worst 25%
Grey-world (GW) [8]	4.44	3.50	0.77	9.64	Grey-world (GW) [8]	5.74	4.60	1.12	12.21
1st-order Gray-Edge [8]	3.51	2.3	0.56	8.53	1st-order Gray-Edge [8]	4.57	3.22	0.84	10.75
V Vuk et al., [8]	6	1.96	0.99	18.81	V Vuk et al., [8]	6.87	2.1	1.06	21.82
Y Qian et al., (1) [8]	2.48	1.56	0.44	6.11	Y Qian et al., (1) [8]	6.87	2.09	0.61	8.18
Y Qian et al., (2) [8]	1.84	1.27	0.39	4.41	Y Qian et al., (2) [8]	2.49	1.71	0.52	6
K Chen et al., [8]	1.84	1.27	0.39	4.41	K Chen et al., [8]	2.49	1.69	0.52	6
Y Qian et al., (3) [8]	2.27	1.26	0.39	6.02	Y Qian et al., (3) [8]	2.93	1.64	0.5	7.78
Fast Fourier [8]	2.1	1.23	0.47	5.38	Fast Fourier [8]	2.48	1.59	0.58	7.27
A Savchik et al., [8]	2.05	1.2	0.41	5.24	A Savchik et al., [8]	2.65	1.51	0.5	6.85
WB-sRGB (modified for raw-RGB) [8]	1.83	1.15	0.35	4.6	WB-sRGB (modified for raw-RGB) [8]	2.36	1.47	0.45	5.94
Ours trained w/o Cube+	2.89	1.72	0.71	7.06	Ours trained w/o Cube+	3.97	2.31	0.86	10.07
Ours trained w/ Cube+	2.1	1.23	0.47	5.38	Ours trained w/ Cube+	2.8	1.54	0.58	7.27

lar error values, respectively. We highlight learning methods (i.e., models trained/tuned for the testing sensor) with gray in the shown tables. The reported results are taken from previous papers, except for the recent work in [10], which was proposed for white balancing images saved in the sRGB color space. We modified [10] to work in the raw-RGB space by replacing the training polynomial matrices with the ground truth illuminant vectors. The shown results of [10] were obtained by using training data from a single camera sensor (i.e., sensor-specific) with the following settings: $k = 15$, $\sigma = 0.45$, $m = 91$, and $c = 191$. For the recent work in [10], we include results of the unsupervised and tuned models. Our method performs better than all statistical-based methods and outperforms some sensor-specific learning methods. We obtain results on par with the *sensor-specific* state-of-the-art results in the NUS 8-Camera dataset (Table 1).

We also examined our trained models on the Cube+ challenge [8]. This challenge introduced a new testing set of 363 raw-RGB images captured by a Canon EOS 550 D – the same camera model used in the original Cube+ dataset [8]. In our results, we did not include any image from the testing set in the training/validation processes. Instead, we used the same models trained for the evaluation on the other datasets (Tables 1–2). Table 3 shows the angular error and reproduction angular errors [24] obtained by our models and the top-ranked methods that participated in the challenge. Additionally, we show results obtained by other methods [8, 10, 24]. For [8], we show the results of the ensemble model (i.e., averaged estimated illuminant vectors from the three-fold trained models on the Cube+ dataset). We report results of two trained models using our method. The first one was trained without ex-

Table 4: Angular errors on the INTEL-TUT dataset [8]. Methods highlighted in gray are trained/tuned for each camera sensor (i.e., sensor-specific models). The lowest errors are highlighted in yellow.

INTEL-TUT Dataset	Gray-World [8]	Shades-of-Gray [8]	2nd-order Gray-Edge [8]	PCA-based B/W Colors [8]	1st-order Gray-Edge [8]	APAP using GW [8]	WB-sRGB [8] (modified for raw-RGB)	Ours trained on NUS and Cube+	Ours trained on NUS and Gehler-Shi
Mean	4.77	4.99	4.82	4.65	4.62	4.30	1.79	3.76	3.82
Median	3.75	3.63	2.97	3.39	2.84	2.44	0.87	2.75	2.81
Best 25%	0.99	1.08	1.03	0.87	0.94	0.69	0.14	0.81	0.87
Worst 25%	10.29	11.20	11.96	10.75	11.46	11.30	5.08	8.40	8.65

amples from Cube+ camera sensor (i.e., trained on all camera models in NUS and Gehler-Shi datasets). The second model was originally trained to evaluate our method on one camera of the NUS 8-Cameras dataset (i.e., trained on seven out of the eight camera models in NUS 8-Cameras dataset, the Cube+ camera model, and the Gehler-Shi camera models). The latter model is provided to demonstrate the ability of our method to use different camera models beside the target camera model during the training phase. More results of the Cube+ challenge are provided in the supplemental materials.

We further tested our trained models on the INTEL-TUT dataset [8], which includes DSLR and mobile phone cameras that are not included in the NUS 8-Camera, Gehler-Shi, and Cube+ datasets. Table 4 shows the obtained results by the proposed method trained on DSLR cameras from the NUS 8-Camera, Gehler-Shi, and Cube+ datasets.

Finally, we show qualitative examples in Fig. 5. For each example, we show the mapped image I_m in our learned intermediate space. In the shown figure, we rendered the images in the sRGB color space by the camera imaging pipeline in [85] to aid visualization.

5 Conclusion

We have proposed a deep learning method for illuminant estimation. Unlike other learning-based methods, our method is a sensor-independent and can be trained on images captured by different camera sensors. To that end, we have introduced an image-specific learnable mapping matrix that maps an input image to a new sensor-independent space. Our method relies only on color distributions of images to estimate scene illuminants. We adopted a compact color histogram that is dynamically generated by our new RGB-uv histogram block. Our method achieves good results on images captured by new camera sensors that have not been used in the training process.

Acknowledgment 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. Dr. Brown contributed to this article in his personal capacity as a professor at York University. The views expressed (or the conclusions reached) are his own and do not necessarily represent the views of Samsung Research.

References

- [1] Mahmoud Afifi, Brian Price, Scott Cohen, and Michael S Brown. When color constancy goes wrong: Correcting improperly white-balanced images. In *CVPR*, 2019.
- [2] Mahmoud Afifi, Abhijith Punnappurath, Graham Finlayson, and Michael S. Brown.

- As-projective-as-possible bias correction for illumination estimation algorithms. *Journal of the Optical Society of America A*, 36(1):71–78, 2019.
- [3] Çağlar Aytekin, Jarno Nikkanen, and Moncef Gabbouj. A data set for camera-independent color constancy. *IEEE Transactions on Image Processing*, 27(2):530–544, 2018.
- [4] Nikola Banić and Karlo Košćević. Illumination estimation challenge. <https://www.isispa.org/illumination-estimation-challenge>. Accessed: 2019-07-01.
- [5] Nikola Banic and Sven Loncaric. Color dog-guiding the global illumination estimation to better accuracy. In *VISAPP*, 2015.
- [6] Nikola Banić and Sven Lončarić. Unsupervised learning for color constancy. *arXiv preprint arXiv:1712.00436*, 2017.
- [7] Jonathan T Barron. Convolutional color constancy. In *ICCV*, 2015.
- [8] Jonathan T Barron and Yun-Ta Tsai. Fast fourier color constancy. In *CVPR*, 2017.
- [9] Ronen Basri and David W Jacobs. Lambertian reflectance and linear subspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(2):218–233, 2003.
- [10] Simone Bianco and Claudio Cusano. Quasi-unsupervised color constancy. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12212–12221, 2019.
- [11] Simone Bianco, Claudio Cusano, and Raimondo Schettini. Color constancy using cnns. In *CVPR Workshops*, 2015.
- [12] David H Brainard and William T Freeman. Bayesian color constancy. *Journal of the Optical Society of America A*, 14(7):1393–1411, 1997.
- [13] David H Brainard and Brian A Wandell. Analysis of the retinex theory of color vision. *Journal of the Optical Society of America A*, 3(10):1651–1661, 1986.
- [14] Gershon Buchsbaum. A spatial processor model for object colour perception. *Journal of the Franklin Institute*, 310(1):1–26, 1980.
- [15] Hakki Can Karaimer and Michael S Brown. Improving color reproduction accuracy on cameras. In *CVPR*, 2018.
- [16] Dongliang Cheng, Dilip K Prasad, and Michael S Brown. Illuminant estimation for color constancy: Why spatial-domain methods work and the role of the color distribution. *Journal of the Optical Society of America A*, 31(5):1049–1058, 2014.
- [17] Dongliang Cheng, Brian Price, Scott Cohen, and Michael S Brown. Effective learning-based illuminant estimation using simple features. In *CVPR*, 2015.
- [18] Mark S Drew, Graham D Finlayson, and Steven D Hordley. Recovery of chromaticity image free from shadows via illumination invariance. In *ICCV Workshop on Color and Photometric Methods in Computer Vision*, 2003.

-
- [19] Graham D Finlayson. Corrected-moment illuminant estimation. In *ICCV*, 2013.
 - [20] Graham D Finlayson. Colour and illumination in computer vision. *Interface Focus*, 8(4):1–8, 2018.
 - [21] Graham D Finlayson and Elisabetta Trezzi. Shades of gray and colour constancy. In *Color and Imaging Conference*, 2004.
 - [22] Graham D Finlayson, Brian V Funt, and Kobus Barnard. Color constancy under varying illumination. In *ICCV*, 1995.
 - [23] Graham D Finlayson, Steven D Hordley, and Ingeborg Tastl. Gamut constrained illuminant estimation. *International Journal of Computer Vision*, 67(1):93–109, 2006.
 - [24] Graham D Finlayson, Roshanak Zakizadeh, and Arjan Gijsenij. The reproduction angular error for evaluating the performance of illuminant estimation algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(7):1482–1488, 2016.
 - [25] David A Forsyth. A novel algorithm for color constancy. *International Journal of Computer Vision*, 5(1):5–35, 1990.
 - [26] David H Foster. Does colour constancy exist? *Trends in Cognitive Sciences*, 7(10):439–443, 2003.
 - [27] Shao-Bing Gao, Ming Zhang, Chao-Yi Li, and Yong-Jie Li. Improving color constancy by discounting the variation of camera spectral sensitivity. *Journal of the Optical Society of America A*, 34(8):1448–1462, 2017.
 - [28] Shaobing Gao, Wangwang Han, Kaifu Yang, Chaoyi Li, and Yongjie Li. Efficient color constancy with local surface reflectance statistics. In *ECCV*, 2014.
 - [29] Peter V Gehler, Carsten Rother, Andrew Blake, Tom Minka, and Toby Sharp. Bayesian color constancy revisited. In *CVPR*, 2008.
 - [30] Arjan Gijsenij and Theo Gevers. Color constancy using natural image statistics and scene semantics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(4):687–698, 2011.
 - [31] Arjan Gijsenij, Theo Gevers, and Joost Van De Weijer. Generalized gamut mapping using image derivative structures for color constancy. *International Journal of Computer Vision*, 86(2-3):127–139, 2010.
 - [32] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feed-forward neural networks. In *AISTATS*, 2010.
 - [33] Yuanming Hu, Baoyuan Wang, and Stephen Lin. FC4: Fully convolutional color constancy with confidence-weighted pooling. In *CVPR*, 2017.
 - [34] Hamid Reza Vaezi Joze and Mark S Drew. Exemplar-based color constancy and multiple illumination. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(5):860–873, 2014.
 - [35] Hakki Can Karaimer and Michael S Brown. A software platform for manipulating the camera imaging pipeline. In *ECCV*, 2016.

- [36] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [37] Karlo Košćević, Nikola Banić, and Sven Lončarić. Color beaver: Bounding illumination estimations for higher accuracy. In *VISIGRAPP*, 2019.
- [38] Zhongyu Lou, Theo Gevers, Ninghang Hu, Marcel P Lucassen, et al. Color constancy by deep learning. In *BMVC*, 2015.
- [39] Laurence T Maloney. Physics-based approaches to modeling surface color perception. *Color Vision: From Genes to Perception*, pages 387–416, 1999.
- [40] Steven McDonagh, Sarah Parisot, Zhenguo Li, and Gregory Slabaugh. Meta-learning for few-shot camera-adaptive color constancy. *arXiv preprint arXiv:1811.11788*, 2018.
- [41] Rang Nguyen, Dilip K Prasad, and Michael S Brown. Raw-to-raw: Mapping between image sensor color responses. In *CVPR*, 2014.
- [42] Seoung Wug Oh and Seon Joo Kim. Approaching the computational color constancy as a classification problem through deep learning. *Pattern Recognition*, 61:405–416, 2017.
- [43] Yanlin Qian, Jarno Nikkanen, Joni-Kristian Kämäräinen, and Jiri Matas. On finding gray pixels. *arXiv preprint arXiv:1901.03198*, 2019.
- [44] Rajeev Ramanath, Wesley E Snyder, Youngjun Yoo, and Mark S Drew. Color image processing pipeline. *IEEE Signal Processing Magazine*, 22(1):34–43, 2005.
- [45] Charles Rosenberg, Alok Ladsariya, and Tom Minka. Bayesian color constancy with non-gaussian models. In *NIPS*, 2004.
- [46] Wu Shi, Chen Change Loy, and Xiaoou Tang. Deep specialized network for illuminant estimation. In *ECCV*, 2016.
- [47] Joost Van De Weijer, Theo Gevers, and Arjan Gijsenij. Edge-based color constancy. *IEEE Transactions on Image Processing*, 16(9):2207–2214, 2007.