

Optimizing Illuminant Estimation in Dual-Exposure HDR Imaging

Mahmoud Afifi[✉], Zhenhua Hu[✉], and Liang Liang

Google

{mafifi, zhenhuahu, annyliang}@google.com

Abstract. High dynamic range (HDR) imaging involves capturing a series of frames of the same scene, each with different exposure settings, to broaden the dynamic range of light. This can be achieved through burst capturing or using staggered HDR sensors that capture long and short exposures simultaneously in the camera image signal processor (ISP). Within camera ISP pipeline, illuminant estimation is a crucial step aiming to estimate the color of the global illuminant in the scene. This estimation is used in camera ISP white-balance module to remove undesirable color cast in the final image. Despite the multiple frames captured in the HDR pipeline, conventional illuminant estimation methods often rely only on a single frame of the scene. In this paper, we explore leveraging information from frames captured with different exposure times. Specifically, we introduce a simple feature extracted from dual-exposure images to guide illuminant estimators, referred to as the dual-exposure feature (DEF). To validate the efficiency of DEF, we employed two illuminant estimators using the proposed DEF: 1) a multilayer perceptron network (MLP), referred to as exposure-based MLP (EMLP), and 2) a modified version of the convolutional color constancy (CCC) to integrate our DEF, that we call ECCC. Both EMLP and ECCC achieve promising results, in some cases surpassing prior methods that require hundreds of thousands or millions of parameters, with only a few hundred parameters for EMLP and a few thousand parameters for ECCC.

Keywords: Computational color constancy · Illuminant estimation · HDR imaging

1 Introduction and Related Work

Camera image signal processor (ISP) comprises several modules, each dedicated to enhancing specific aspects of the quality of captured raw images by the camera sensor [16]. One key component among these modules is the white-balance module, which aims to eliminate undesirable color casts introduced by the combination of scene lighting and camera sensitivity. To achieve this, the auto white-balance module runs an illuminant estimator onboard the camera ISP to estimate the RGB color of the illuminant, under the assumption that a single global illuminant illuminates the scene for simplicity [30, 44]. Image white balancing can thus be described as follows:




Fig. 1: Conventional illuminant estimators often rely on a single frame for illuminant color estimation. Although the HDR camera pipeline includes at least two frames of the same scene, conventional methods usually consider only a single frame (either with auto exposure, shown in the dashed green line, merged frame, or one frame of the burst capture) [12, 16, 44, 62]. This paper proposes leveraging information from dual-exposure capturing in multi-exposure HDR imaging to enhance illuminant estimation in camera pipelines. Our method uses frames captured at long and short exposures, available in multi-exposure bursts [28] or staggered HDR sensors [52, 58]. Achieving comparable or superior results, our method employs lightweight models (~ 300 – 6000 parameters) compared to those using hundreds of thousands or millions of parameters. In this figure and the following figures, all raw images have the gamma operator applied to aid visualization, and all sRGB images are rendered using the HDR+ camera pipeline [44].

$$I_{\text{WB}} = \text{diag} \left(\left[\frac{\hat{\ell}_G}{\hat{\ell}_R}, 1, \frac{\hat{\ell}_G}{\hat{\ell}_B} \right]^T \right) I, \quad (1)$$

$$[\hat{\ell}_R, \hat{\ell}_G, \hat{\ell}_B]^T = f(I), \quad (2)$$

where I and I_{WB} are $3 \times k$ RGB colors of raw image and white-balanced raw image, respectively, with k refers to the total number of pixels in the image, $\text{diag}(\cdot)$ creates a 3×3 diagonal matrix from a given vector, $\hat{\ell} \in \mathbb{R}^3$ refers to the estimated illuminant color vector, T refers to vector transpose, and f is an illuminant estimator function.

Modern camera ISPs capture several frames of each scene with different exposure times to enhance final image dynamic range [28, 38]. This can be achieved through burst capturing, involving rapid capture with varying exposure times in quick succession, or by using staggered high dynamic range (sHDR) sensors that simultaneously capture long and short exposures [52, 58]. The captured images are then combined, incorporating multiple exposures of the same scene at different levels, resulting in a greater dynamic range than what would be possible with a single image [43, 49, 51, 55]. This HDR imaging camera pipeline, therefore, includes more than a single frame of the captured scene, and internal camera ISP modules, such as the auto white-balance module, can access these additional frames. While such additional frames may have beneficial information to help illuminant estimators, conventional illuminant estimation methods often rely on a single frame (e.g., [9, 36, 41, 47, 54, 63, 64]) in both traditional single-frame camera pipelines [39] and multi-frame HDR ISPs [44].

Such single-frame illuminant estimators can be categorized into: 1) statistical methods that rely on statistics computed from input raw image information (e.g., colors, edges) [13, 14, 25, 31, 50, 56, 57] and 2) learning-based methods that learn, from a set of training images labeled with ground-truth illuminant colors, to map from input raw color information to the corresponding illuminant color [6, 8, 10, 11, 24, 36, 42, 45, 47, 48, 54, 63, 64]. While the latter category typically achieves better results than the statistical methods, most learning-based methods are camera-dependent, meaning that they require domain adaptation or fine-tuning when deployed on new cameras to achieve similar accuracy on cameras used for training due to the influence of the camera response function on both raw image colors and ground-truth colors [3–5, 35, 53, 59].

A limited number of attempts have proposed learning-based methods that go beyond the single-frame input scheme. For example, the cross-camera convolutional color constancy (C5) [4] suggests utilizing additional unlabeled images captured by the testing camera, in addition to the primary single-frame image of the scene being captured to improve the generalization for cameras that were not included in the training phase. Xing et al. [60] proposed the use of a depth map, captured by a time-of-flight (ToF) sensor, along with the primary raw image to predict illuminant color in the scene by leveraging the geometry information obtained from depth map. Abdelhamed et al., [2] proposed leveraging the presence of two cameras in most modern mobile phone devices. Assuming dual streaming from both cameras, they derived a feature within the framework of chromagenic color constancy theory [21–23] to enhance the accuracy of illuminant estimation, yielding promising results.

Our method, in contrast to the majority of prior work, adopts the strategy of benefiting from multiple frames available in the camera ISP (similar to [2]). However, unlike [2, 60], which requires streaming from dual cameras and may lead to impractical high power consumption, our method relies on two frames captured of the same scene under different exposure settings, already present in the HDR imaging pipeline, to estimate the illuminant color of the captured scene (see Fig. 1). Specifically, we utilize: 1) a frame captured with short exposure time, that we refer to as short-exposure image (I_s), and 2) another frame of the same scene captured with long exposure time, that we refer to as long-exposure image (I_l). Having dual-exposure images in the HDR imaging pipeline is feasible, making our method practical for most HDR imaging pipelines. Accordingly, Eq. 2 can be modified to include our dual-exposure input images as follows:

$$[\hat{\ell}_R, \hat{\ell}_G, \hat{\ell}_B]^T = f(I_l, I_s). \quad (3)$$

Our objective in this work is to design the function f in a manner that enables the effective utilization of the additional information provided by the dual-exposure images, I_l and I_s .

Contribution In this paper, we present a feature inspired by the chromagenic color constancy theory [21–23], termed the dual-exposure feature (DEF), that is derived from images captured with both short exposure time (I_s) and long

exposure time (I_l). DEF leverages the variations in chromatic information between these dual-exposure images, providing valuable guidance for illuminant estimator methods. To assess its effectiveness, we trained a lightweight multi-layer perceptron network (MLP) for illuminant estimation that utilizes solely our DEF as input, departing from the conventional approach of using actual RGB values of the captured scene colors. Additionally, we explored the integration of DEF into an established color constancy framework, specifically the convolutional color constancy (CCC) [4, 8, 9, 37], which we refer to as exposure-based CCC (ECCC). The experimental results on a multi-exposure dataset, collected to evaluate our work empirically, show that these models – namely, EMLP and ECCC – achieve promising results with a reasonable number of parameters – 354 learnable parameters for EMLP and 6,156 learnable parameters for ECCC. This outperformance across diverse evaluation metrics is observed when comparing to prior methods that require significantly higher number of parameters, ranging from tens to hundreds of thousands, or even millions.

2 Illuminant-Linked Dual Exposure Feature

To develop an efficient illuminant estimator that benefits from both I_l and I_s , we introduce a compact feature that aims to capture the correlation between these dual-exposure images. Our feature is inspired by the chromagenic color constancy theory [21–23]. Specifically, we explore the analogy between long and short exposure images, I_l and I_s , and aligned images captured by two cameras [2] under the chromagenic color constancy theory, leading us to develop our dual-exposure feature (DEF). To begin, we review the chromagenic color constancy theory under the Lambertian reflectance model with a single illuminant assumption. The captured raw image I can mathematically be described by:

$$I_{\rho}^{(y)} = \int_{\gamma} S_{\rho}(x)D(x)R(y, x) dx + z, \quad (4)$$

where $S(\cdot)$, $D(\cdot)$, and $R(\cdot)$ represent the camera response function (typically represented by camera sensitivity, infrared cut-off filter, and spectral lens transmission), the spectral power distribution of light, and scene reflectance, respectively. Here, x refers to a wavelength within the visible range γ , y refers to the pixel location in image I , $\rho \in \{R, G, B\}$ refers to the color channel, and z denotes the undesired noise, typically represented by signal-dependent and additive components [1]. According to the chromagenic color constancy theory [22], if we capture two images of the same scene with a specially chosen chromagenic filter, Q , applied between image captures, the linear color matrix that maps between those captured images is *unique* to the illuminant color present in this scene. That is, given an image, I , that is captured by the main camera, and a filtered image, I_f , that can be described as:

$$I_{f\rho}^{(y)} = \int_{\gamma} S_{\rho}(x)Q_{\rho}(x)D(x)R(y, x) dx + z, \quad (5)$$

the 3×3 color matrix C_c that maps between the colors of I and I_f is indicative of the scene illuminant. This color matrix can be computed by minimizing the following equation:

$$\arg \min_{C_c} \|C_c I_f - I\|_F, \quad (6)$$

where $\|\cdot\|_F$ is the Frobenius norm. The closed-form solution for Eq. 6 can be obtained using the pseudoinverse (i.e., $C_c = II_f^\dagger$). While theoretically validated, the chromagenic filter conditions required to obtain a *unique* mapping matrix per illuminant are challenging to meet in practice. Empirically, Finlayson et al., [23] found that most filters exhibit a reasonable level of *correlation* between the computed matrix and the illuminant color, excluding neutral filters, which consistently result in a scaling relationship between the unfiltered image, I , and the filtered image, I_f . Relaxing the conditions to include *normal* colored filter, Abdelhamed et al., [2] proposed using two cameras, with the second camera serving as the main camera after applying a colored filter—i.e., $S(\cdot)Q(\cdot)$. Surprisingly, even when chromagenic filter conditions are not met, using a normal colored filter [23] or another camera with a different response function [2] still shows a correlation between the computed matrix C_c and the illuminant color to some extent, achieving promising results.

We argue that this correlation arises due to the variations, or what can be considered a form of “distortion”, in the colors captured by the second (filtered) camera when compared to the original colors captured by the main camera. This color distortion varies based on the interplay between the scene irradiance, $D(\cdot)R(\cdot)$, and the camera response functions – namely, the main camera’s $S(\cdot)$ and the filtered/second camera’s $S(\cdot)Q(\cdot)$ in Eqs. 4 and 5. In the context of machine learning (ML) models, the camera response functions remain fixed across the entire dataset. Consequently, ML models, such as the one proposed in [2], learn the correlation between the ground-truth illuminant color and the scene irradiance through the matrix C_c that represents the level of “distortion” between the colors of captured scene images.

In a dual-exposure setup, we have a long-exposure image, I_l , and a short-exposure image, I_s , both capturing the same scene. The extent of color distortion in each image varies based on the scene irradiance; for instance, I_s may exhibit a higher level of color distortion and noise than I_l under suboptimal lighting conditions (e.g., indoor light), while I_l may have a higher level of color distortion and over-saturated colors than I_s in well-lit scenes (e.g., outdoor light). This difference in color distortion arises because the non-linear camera response function receives a different number of photons to form each image colors based on the exposure time and scene irradiance [32]. As a result, I_l and I_s can exhibit color variations that differ across individual color channels [15], that are somewhat akin to those caused by a color filter (though to a lesser extent). Even within the same color channel, different levels of color differences between I_l and I_s can be observed spatially due to the interaction between the non-linear camera response function of that channel and the spatially varying scene irradiance (see

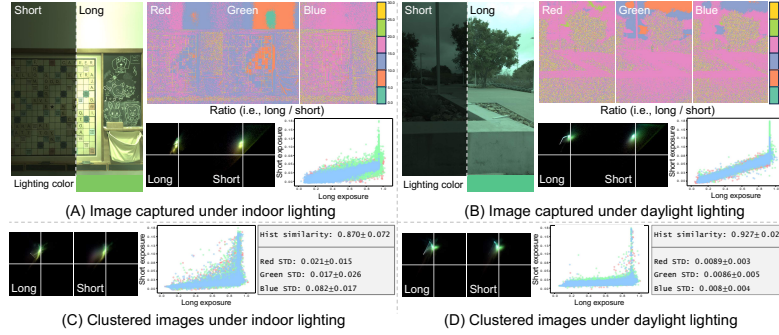


Fig. 2: Cameras perceive different amount of photons when capturing the same scene with different exposure times. Images taken with both long and short exposure times exhibit variations in each channel due to the camera response function and scene irradiance. Additionally, spatial variations, in each color channel, can be observed based on object reflectance, as the interplay of lighting, object reflectance, and the camera response function leads to different outcomes. (A) and (B) show raw images of scenes captured under indoor and outdoor lighting, respectively. In (C) and (D), we present the average rg -chromaticity histogram and aggregated red, green, and blue pixel values from 25 images sharing similar lighting conditions in (A) and (B), respectively.

Fig. 2). In Fig. 2, we show two examples of dual-exposure images captured under different lighting conditions. As can be seen, the differences between the colors in I_l and I_s under each lighting condition exhibit variations, which are observable in the histogram similarity (here we use the Bhattacharyya distance similarity metric) and variations of the ratios in each color channel between I_l and I_s . That is, the correlation between I_l and I_s is *not* always a proportional scaling, as in the case of a neutral filter. We observed similar patterns as shown in Fig. 2 when studying examples from the two-camera dataset in [2] (see supplemental materials).

Based on this discussion, we propose to compute the color matrix C_c to map between the rgb -chromaticity values (i.e., $[R/\kappa, G/\kappa, B/\kappa]^T$, with $\kappa = R+G+B$) of I_s and I_l . The reason for not using the RGB triples, similar to chromagenic color constancy, is that we aim to reduce the influence of intensity differences between I_l and I_s when computing C_c . In addition, we compute the covariance matrix, C_v , of the ratio image $X \in \mathbb{R}^{3 \times k}$, where $X_p^{(y)} = I_{s_p}^{(y)} / (I_{l_p}^{(y)} + \epsilon)$, and ϵ is a small number added for numerical stability. Computing C_v is performed as described in Eq. 7 to measure the variance in each color channel between I_s and I_l and the joint variability across channels.

$$C_v(X) = \mathbb{E}[(X - \mathbb{E}[X])(X - \mathbb{E}[X])^T], \quad (7)$$

where $\mathbb{E}(\cdot)$ is the expected value (mean) of the matrix. Both C_c and C_v form our dual-exposure feature (DEF) that represents the differences in color distortion in I_l and I_s . Our DEF is represented as a vector $\in \mathbb{R}^{15}$, where we exclude

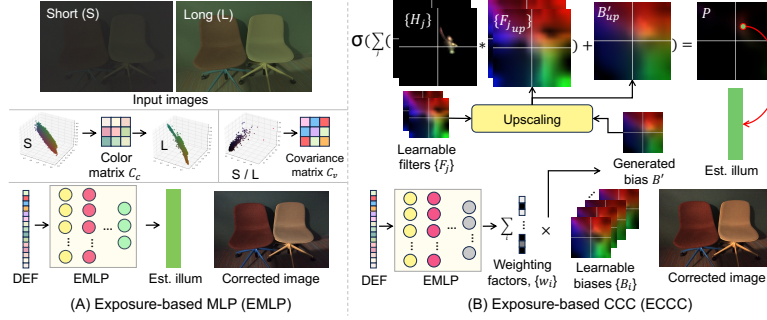


Fig. 3: We present an illuminant-related dual-exposure feature (DEF), derived from a pair of images captured with short and long exposures. Using DEF, we deploy a simple multilayer perceptron network (MLP) with only 354 parameters, referred to as exposure-based MLP or EMLP, for illuminant estimation, as shown in (A). We further explore the integration of DEF into the CCC framework, as shown in (B), by dynamically generating bias map based on DEF. We denote this modified CCC framework as exposure-based CCC or ECCC.

the redundant values in C_v over the symmetric positions. To evaluate DEF’s effectiveness as a clue for scene illuminant colors, we employed a lightweight MLP for scene illuminant estimation (Fig. 3-A), that we call exposure-based MLP (EMLP). EMLP relies solely on our DEF as input. It comprises an input fully connected (fc) layer with 9 output neurons, followed by leaky ReLU (LReLU) [61], two hidden fc layers, each with 9 output neurons with LReLU in between, and an output fc layer with three neurons. EMLP achieves results comparable to complex models with thousands or millions of parameters (refer to Tables 1 and 2), while maintaining a lightweight design with just a few hundred parameters. Consequently, DEF is empirically shown as a valuable feature providing strong insights into the illuminant scene color.

3 Integration with Convolutional Color Constancy

In this section, we incorporate the proposed DEF into one of the most established frameworks for illuminant estimation. Specifically, we introduce modifications to the convolutional color constancy (CCC) framework [4, 8, 9, 37] to leverage the benefits of the DEF, which we refer to as exposure-based CCC or ECCC for short. It is important to note that ECCC *does not* introduce a new CCC method; instead, it serves as an illustration of how the DEF can be seamlessly integrated into existing, well-established illuminant estimation frameworks.

Let’s begin with a brief description of the CCC [8, 9]. Given a single raw image, the CCC operates by learning one or more convolutional 2D filters, denoted as $\{F_j\}$, which convolve over the 2D histogram(s), $\{H_j\}$, of the image colors in the uv color space (i.e., the log of G/R and G/B chroma values of pixel colors) [8, 20]. This convolutional operation can be accelerated by FFTs as

proposed in [9]. Afterwards, a 2D bias map, B , is added, followed by a softmax operation, σ , to compute the “probability” map, P , of the illuminant bin in the uv 2D histogram space. This simplified version of the CCC can be described as follows:

$$P = \sigma \left(\sum_j (F_j * H_j) + B \right). \quad (8)$$

In CCC [8,9], the filters and bias are learned across the entire training dataset. This means it consists of a single bias and one or more filters (the number depends on the histograms used, either a single histogram for image colors or two histograms, including the edge color histogram). Later, C5 [4] proposes a hypernetwork that dynamically generates filters and bias based on the input raw image and additional images taken from the same camera, aiming to improve the generalization of CCC across cameras. While cross-camera color constancy is out of the scope of this paper, we propose to use our DEF to dynamically “generate” a bias map based on the input image. That is, we learn a bank of biases, $\{B_i\}$, where $i \in \{1, \dots, n\}$, such that we linearly interpolate between them based on blending weights emitted from an MLP network that processes our DEF. In this way, the DEF controls the bias of the CCC model, acting as “candidate” illuminant priors within the uv space for the input image (see Fig. 3-B). To implement this, we need to have multiple bias maps, which definitely will lead to an impractical increase in model size. Thus, we propose to use a downsampled size ($1/4$ of histogram size) for the learnable biases, $\{B_i\}$. Furthermore, we propose to feed two histograms (i.e., $j \in \{l, s\}$) of both I_l and I_s , denoted as H_l and H_s , respectively, into the model. This leads to the learning of two downsampled filters, $\{F_j\}$, corresponding to the histograms of the long-exposure image (I_l) and the short-exposure image (I_s). Learning small-sized filters and biases has the following benefits. First, we can learn many biases within an affordable model size (e.g., $n = 20$ requires $\sim 6K$ parameters, while FFCC [9] with a single bias requires $\sim 12K$ parameters). Second, it implicitly produces smooth filters and biases, which are desirable to avoid overfitting [4,9]. With this modification, Eq. 8 can be rewritten as follows:

$$P = \sigma \left(\sum_j (\uparrow (F_j) * H_j) + B'_{\text{up}} \right), \quad (9)$$

$$B'_{\text{up}} = \uparrow \left(\sum_{i=1}^n w_i B_i \right). \quad (10)$$

where $\uparrow(\cdot)$ refers to upscaling through bilinear interpolation, $[w_1, \dots, w_n]^T$ is a weighting vector produced by a lightweight MLP (similar to EMLP, but with n output neurons) that processes our DEF. After computing P , the estimated illuminant can be obtained by:




Fig. 4: Examples from the dataset used in this work. For each scene, we captured the scene with a gray calibration object placed in the scene to obtain the ground-truth illuminant (A) and captured the scene using different exposure settings without the gray object (B-H). The terms ‘short / e ’ (C-E) and ‘long $\times e$ ’ (F-H) refer to multiplying and dividing auto exposure time by a factor e , respectively. The first image in (A) is displayed in sRGB, while the rest are shown in raw RGB space.

$$\hat{\ell}_u = \sum_{u,v} uP(u,v), \quad \hat{\ell}_v = \sum_{u,v} vP(u,v), \quad (11)$$

$$\hat{\ell} = \left[\exp\left(-\hat{\ell}_u\right)/q, 1/q, \exp\left(-\hat{\ell}_v\right)/q \right]^T, \quad (12)$$

$$q = \sqrt{\exp\left(-\hat{\ell}_u\right)^2 + \exp\left(-\hat{\ell}_v\right)^2 + 1}. \quad (13)$$

4 Experiments

4.1 Data

To the best of our knowledge, the existing HDR multi-exposure datasets (e.g., [17, 33]) do not provide ground-truth illuminant colors for training and evaluation of our method. Thus, we compiled a dataset of scenes captured with both auto exposure and various multiple-exposure settings for each scene using Pixel 7 Pro camera (see Fig. 4). Images captured with auto exposure are used for training (in the case of learning-based methods) and evaluating other methods, while those captured with multiple exposure settings are employed for validating our method. Our set comprises 558 scenes, each captured using auto exposure and six additional exposure settings: $\times 2$, $\times 4$, $\times 8$, $1/2$, $1/4$, and $1/8$, indicating adjustments to the original exposure time. Specifically, $\times 2$ signifies doubling the original exposure time, $\times 4$ and $\times 8$ denote quadrupling and octupling, respectively, while $1/2$, $1/4$, and $1/8$ represent dividing the original exposure time by 2, 4, and 8, respectively. These adjustments are dynamically computed as ratios from the initial auto exposure, following the common practice in multi-exposure

HDR imaging on smartphone cameras [28, 38]. The camera was fixed on a tripod to minimize potential misalignment between images captured by different exposure settings. In total, we captured 3,906 raw images (558×7 captures of each scene). The dataset is organized as follows: 83 and 86 scenes were randomly selected for validation and testing, respectively, with the remaining 387 scenes were used for training. Additionally, each scene was captured with a gray calibration object to obtain the ground-truth illuminant. The images with the gray calibration object are not included in any of the training, validation, nor testing sets.

During training, we optionally augment training data using chromatic adaptation. Specifically, we employ clustering on the training data using our DEF as query feature. We use K-means [34] with L2 to create 80 clusters. Subsequently, we augment the training data by generating three additional images for each sample. For each augmented image, we randomly select an illuminant color from the cluster to which the original image belongs. We then apply Von Kries transform [18] to remove the original illuminant from the image, using the actual ground truth, and apply the newly selected illuminant color to the image using the diagonal scaling operator. This new illuminant serves as the ground truth for the augmented image. As our method is designed to handle two images, namely I_s and I_l , we use, without loss of generality, a single exposure factor, e , for both images, such that I_s was captured with $1/e$ of the auto-exposure time, while I_l was captured with $\times e$ of the auto-exposure time, for simplicity. In Sec. 4.3, we present an ablation study on the impact of different values of e on our final results.

4.2 Training

We trained EMLP and ECCC using the Adam optimizer [40] with a mini-batch size set to 32 for 1000 and 200 epochs, respectively. For EMLP, the learning rate was 10^{-3} , while ECCC was trained with an incremental mini-batch size (similar to [4]) with a learning rate of 5×10^{-3} with a cosine annealing schedule [46], and the weight decay (i.e., L2 regularization penalty) was set to 10^{-5} . For ECCC, the filter weights were initialized to zeros, and the biases were initialized to n 2D histograms of training ground-truth illuminant colors after clustering into n clusters using our DEF as a query feature with K-means. The n histograms of training ground-truth illuminants were first processed by morphological dilation using a 3×3 diamond-shaped structuring element before being used as initial values for our learnable biases. This initialization improves the results, as it assists in establishing reasonable “candidate” illuminant priors linked to the corresponding DEFs (see Sec. 4.3 for an ablation study). Both models were trained using the angular error, $L(\cdot)$, between the predicted illuminant color $\hat{\ell}$ and the ground-truth illuminant ℓ , as described in the following equation [19]:

$$L(\ell, \hat{\ell}) = \cos^{-1} \left(\frac{\ell \cdot \hat{\ell}}{\|\ell\| \|\hat{\ell}\|} \right), \quad (14)$$

where $\|\cdot\|$ denotes the Euclidean norm, and (\cdot) represents the vector dot product. For ECCC, we further added two smoothness loss terms to encourage smoothness in the learned filters and biases, similar to [4]. See supplemental materials for more details.

4.3 Results

We conducted a comprehensive comparison of our proposed methods against various existing techniques, which include: training-free statistical methods [13, 14, 25, 50], camera-independent learning-based methods [4, 5], and camera-dependent learning-based methods [6, 9, 29, 41, 42, 54, 63, 64]. For the camera-dependent learning-based methods, we trained each model on our data using the provided code and recommended parameters by the respective authors. Meanwhile, for sensor-independent learning methods, we utilized the provided pre-trained models on other datasets (e.g., [7, 14, 27]).

For the sake of conducting realistic experiments, we assessed all methods, including ours, on 384×256 images—a suitable size for evaluating illuminant estimation methods designed for embedded hardware devices; such devices may utilize even smaller image sizes in white-balance modules [9]. For CCC methods [4, 9], including the ECCC, 64×64 histograms were used. The ECCC incorporates the proposed DEF into the CCC by generating an interpolated bias map based on the DEF. This can be seen as a spatial case analogous to C5 [4], where C5 utilizes a neural network to emit bias and convolution filters based on additional images taken by the same camera to enhance generalization across cameras. To validate our modification against C5, we trained different versions of C5 not aimed at improving cross-camera generalization (as our method also focuses on a single camera). We refer to these modified versions as follows:

- **C5 (model A):** We use the histogram of the averaged short and long exposure images without additional images. That is, the C5 network uses a single histogram to generate CCC model parameters.
- **C5 (model B):** We use histograms of both images taken by long and short exposures as input to the C5 network to generate CCC model parameters. The CCC mode is then applied to the histogram of the averaged long-short images.
- **C5 (models C and D):** Both Model C and Model D use histograms of both images taken by long and short exposures to feed the C5 network and generate CCC model that is then applied to histogram of short exposure image (model C) and long exposure image (model D).

The results are given in Tables 1 and 2 on the validation and testing sets, respectively. We report the mean, median, tri-mean, best 25%, worst 25%, worst 5%, and max angular errors in each set. We also report the total number of parameters for other methods, including ours. Additionally, we present the results of a set of ablation studies conducted to examine the impact of the exposure factor e in both EMLP and ECCC. We also studied the impact of different color

Table 1: Angular errors on the validation set. \diamond and \star denote camera-independent models and training-free statistical methods, respectively. Methods are listed chronologically by publication year, with the top and second-best results highlighted in yellow and red. We present the results of camera-dependent trained models for various versions of C5 [4], identified as models A, B, C, and D (see main text for more details). Our results are reported using EMLP, ECCC, and the ensemble model (EMLP + ECCC). Ablation studies are included using different mapping matrices (CM: 3×3 color mapping matrix, TM: 3×4 affine transformation matrix, HM: 3×3 homography matrix), different color representations (rgb: raw RGB triplet as used in [2, 21–23, 26], *rg/rgb*-chroma: *rg*-chromaticity/*rgb*-chromaticity of the raw RGB triplet), different exposure ratios ($e \in \{2, 4, 8\}$), different input image sizes and histograms (s indicates 48×32 input images, s^2 indicates 48×32 input images and 32×32 histograms), using different input histograms for ECCC of both long and short exposure images, average image, and long/short image, different number of learned biases (n), and results with (w/) and without (w/o) augmentation, w/ and w/o the covariance matrix’s parameters (for EMLP), w/ and w/o DEF and bias initialization (BI) (for ECCC).

Method	Validation set						Params
	Mean	Med.	Tri.	Best 25%	Worst 25%	Worst 5%	
Grayworld \star [13]	5.54	3.13	4.11	0.97	12.83	19.45	28.82
Shades of Gray \star [25]	7.16	6.21	6.43	0.99	15.06	18.39	19.21
PCA \star [14]	6.13	4.53	5.02	0.91	13.66	18.28	20.33
Ganutt (pixels) [29]	7.53	7.38	7.03	2.04	13.96	18.34	21.01
Ganutt (edges) [29]	7.03	6.06	6.16	1.49	14.54	19.51	21.98
PECC [9]	3.77	1.73	2.32	0.54	9.56	16.08	28.99
Gray Index \star [50]	6.00	3.62	4.28	0.58	15.44	24.91	31.72
APAP-LUT [6]	4.84	3.30	3.87	1.12	10.84	16.92	27.53
SHF \star [5]	4.70	4.04	4.23	1.35	9.21	13.27	16.24
BiCCP [41]	4.84	3.30	3.87	1.12	10.84	16.92	27.53
C1 [63]	3.91	3.26	3.24	1.06	8.50	12.76	17.35
CWCC [42]	4.49	3.35	3.48	1.61	9.21	13.31	16.11
CS \star [4]	4.01	2.81	3.14	1.08	8.51	12.47	21.81
CS (model A) [4]	3.93	2.29	2.90	1.01	9.40	13.33	18.75
CS (model B) [4]	4.40	2.84	3.24	0.88	10.52	15.23	20.55
CS (model C) [4]	4.25	2.67	3.19	0.93	10.12	14.57	19.82
CS (model D) [4]	4.51	2.94	3.38	0.87	10.74	16.01	20.74
TLCC [54]	4.16	2.72	3.15	0.81	9.35	14.05	21.39
PCC [64]	4.61	4.19	4.12	1.75	8.50	12.06	17.86
EMLP ($e = 8$, <i>rgb</i> -chroma, HM, w/o cov)	3.86	4.26	4.96	1.36	12.5	16.56	20.14
EMLP ($e = 8$, <i>rgb</i> -chroma, TM, w/o cov)	4.78	3.94	4.27	1.14	9.57	13.20	16.54
EMLP ($e = 8$, rgb, CM, w/o cov)	4.54	3.53	3.87	1.26	9.51	13.74	19.22
EMLP ($e = 8$, <i>rg</i> -chroma, CM, w/o cov)	4.81	3.98	4.22	1.29	9.32	12.75	14.99
EMLP ($e = 2$, <i>rgb</i> -chroma, CM, w/o cov)	4.33	3.33	3.57	1.11	9.44	13.60	17.14
EMLP ($e = 4$, <i>rgb</i> -chroma, CM, w/o cov)	4.16	3.56	3.77	1.00	8.38	12.16	16.79
EMLP ($e = 8$, <i>rgb</i> -chroma, CM, w/o cov)	4.07	3.51	3.64	1.23	7.77	10.38	13.29
EMLP ($e = 8$, <i>rgb</i> -chroma, CM, w/ cov)	3.77	2.89	2.92	0.86	8.59	11.48	12.71
EMLP+ ($e = 8$, <i>rgb</i> -chroma, CM, w/ aug, w/ cov)	3.83	2.82	3.27	0.88	8.23	11.29	12.81
EMLP ($e = 8$, <i>rgb</i> -chroma, CM, w/ aug, w/ cov)	3.52	2.43	2.85	0.86	8.68	11.32	13.11
ECCC ($e = 8$, $n = 5$, w/ DEF, both)	3.69	2.78	3.11	0.79	7.83	10.62	13.27
ECCC ($e = 8$, $n = 10$, w/ DEF, both)	3.66	2.69	2.92	0.95	8.03	11.06	12.32
ECCC ($e = 8$, $n = 15$, w/ DEF, both)	3.58	2.71	2.91	0.91	7.96	11.85	13.51
ECCC ($e = 2$, $n = 20$, w/ DEF, both)	3.91	2.99	3.23	1.07	8.33	11.81	15.23
ECCC ($e = 4$, $n = 20$, w/ DEF, both)	4.28	3.19	3.36	0.99	9.63	13.72	14.72
ECCC ($e = 8$, $n = 20$, w/o DEF, both)	4.02	3.23	3.69	1.22	9.17	12.14	14.47
ECCC ($e = 8$, $n = 20$, w/o BI, both)	3.99	3.13	3.51	0.97	8.53	11.41	13.66
ECCC ($e = 8$, $n = 20$, w/ DEF, avg)	4.03	2.97	3.34	1.03	8.68	12.09	14.71
ECCC ($e = 8$, $n = 20$, w/ DEF, short)	3.71	2.74	2.93	1.00	8.07	12.20	15.43
ECCC ($e = 8$, $n = 20$, w/ DEF, long)	3.66	2.58	2.85	0.79	8.20	11.34	12.22
ECCC ($e = 8$, $n = 20$, <i>rgb</i> -chroma, w/ DEF, both)	3.79	2.73	3.23	0.82	8.41	11.63	15.44
ECCC ($e = 8$, $n = 20$, w/ DEF, both, w/ aug)	3.61	2.97	3.14	0.79	7.37	10.54	11.41
ECCC+ ($e = 8$, $n = 20$, w/ DEF, both)	3.69	2.62	3.07	0.9	8.12	10.95	15.1
ECCC+* ($e = 8$, $n = 20$, w/ DEF, both)	4.04	3.34	3.50	0.66	8.75	11.28	12.96
ECCC ($e = 8$, $n = 20$, w/ DEF, both)	3.56	2.50	2.79	0.87	7.93	11.32	13.47
EMLP + ECCC	3.24	2.37	2.53	0.84	7.11	10.64	12.06

spaces of input images before computing our DEF. Specifically, we used the *rg*-chromaticity and raw RGB values (similar to [2, 21–23, 26]), in addition to our main design that uses *rgb*-chromaticity. We also studied different mapping matrices, in addition to the 3×3 color matrix we used in Sec. 2. Specifically, we examined using the geometric 3×4 affine transformation matrix and the 3×3 homography matrix.

We report results of our EMLP with and without the covariance matrix parameters in addition to the results of ECCC with different numbers of learnable biases, n , and with and without the proposed bias initialization. Furthermore, we

Table 2: Angular errors on the testing set. See caption of Table 1 for abbreviations.

Method	Testing set						
	Mean	Med.	Tri.	[Best 25%]	[Worst 25 %]	Max	
Grayworld* [13]	5.32	3.70	4.51	0.95	11.34	14.41	17.75
Shades of Gray* [25]	6.77	5.77	6.08	1.07	14.24	16.64	17.50
PCA* [14]	5.91	4.99	5.22	1.07	12.62	17.10	19.08
Ganutt (pixels) [29]	7.49	7.25	7.37	2.30	12.96	16.62	19.00
Ganutt (edges) [29]	5.90	4.80	4.93	1.13	12.20	17.61	20.09
FFCC [9]	3.18	1.83	2.28	0.39	7.97	13.13	20.58
Gray Index* [50]	5.18	3.47	4.18	0.60	12.32	16.90	19.67
AFAP-LUT [6]	3.94	2.82	3.32	0.81	8.54	11.90	16.36
SIE* [5]	3.51	2.35	2.74	0.80	7.56	11.06	12.30
BoCF [41]	3.91	3.31	3.40	1.18	7.72	11.24	12.81
C4 [63]	3.79	2.76	3.06	1.05	7.93	11.33	13.39
CWCC [42]	3.71	2.85	3.06	1.11	7.79	11.21	12.83
C5* [4]	3.45	2.93	3.03	0.95	7.00	10.38	14.01
C5 (model A) [4]	3.82	2.33	2.81	0.72	9.08	15.69	20.68
C5 (model B) [4]	3.30	2.05	2.42	0.69	7.58	11.23	14.02
C5 (model C) [4]	3.31	2.01	2.40	0.63	7.67	12.00	14.23
C5 (model D) [4]	3.33	2.20	2.54	0.73	7.53	10.85	13.33
TLCC [54]	3.73	2.89	3.21	0.96	7.54	11.55	14.78
PCC [64]	4.37	3.66	3.64	0.89	9.14	15.17	23.34
EMLP ($e = 2$, w/ aug)	3.64	2.66	2.97	0.88	7.82	11.5	13.62
EMLP ($e = 8$, w/ aug)	3.37	2.34	2.74	0.67	7.24	10.16	14.10
EMLP ($e = 8$, w/ aug)	3.36	2.73	2.84	0.76	7.03	9.53	11.85
ECCC+ ($e = 8$)	3.33	2.92	2.92	0.96	6.56	9.38	10.26
ECCC+* ($e = 8$)	3.41	2.87	3.03	0.81	6.87	9.12	11.03
ECCC ($e = 2$)	2.94	2.10	2.38	0.69	6.31	8.13	9.25
ECCC ($e = 4$)	3.18	2.39	2.60	0.81	6.99	10.01	10.56
ECCC ($e = 8$)	3.00	2.31	2.45	0.74	6.66	9.85	12.23
ECCC ($e = 8$, w/ aug)	2.95	2.09	2.37	0.76	6.06	8.14	9.32
EMLP + ECCC	2.91	2.33	2.44	0.72	6.11	8.38	8.93

show the results of ECCC without utilizing the DEF feature. Lastly, we demonstrate the influence of different sizes of input sizes on the inference accuracy of both EMLP and ECCC. The results show that the best results are obtained with $e = 8$. This choice is intuitively sensible, as a higher exposure factor increases the differences between the dual-exposure images, resulting in greater distinction based on scene lighting conditions. The augmentation is found to be useful for EMLP; however, we did not observe consistent improvement in the case of ECCC. This is likely because EMLP has a limited number of input features, making it more susceptible to overfitting. Thus, augmentation helps in generalization. Conversely, in the case of ECCC, the inclusion of histogram features alongside the small DEF feature suggests that the model may not consistently derive benefits from augmentation. We also noted that in ECCC, using raw RGB colors of dual-exposure images (without chromaticity conversion) tends to yield better results, while in EMLP, the chromaticity conversion tends to enhance performance (see Table 1). Results presented in Table 2 default to *rgb*-chromaticity for EMLP and raw RGB values for ECCC.

From the results presented in Tables 1 and 2, it is evident that our proposed DEF serves as a promising feature for guiding illuminant estimators. This is demonstrated by the performance of our lightweight EMLP model (354 parameters) compared to complex models that rely solely on raw RGB image colors (e.g., TLCC [54] with 32 million parameters). Moreover, incorporating DEF into ECCC reduces the maximum error by over 50% and achieves comparable results across various evaluation metrics when compared to FFCC [9].

Since both models – namely, EMLP and ECCC – have a reasonable number of parameters, we can combine their predictions to create an ensemble model by averaging the predicted illuminant colors from both models. The results of this ensemble model are reported in Tables 1 and 2 under ‘ECCC + EMLP’. The combined efforts of EMLP and ECCC, guided by our DEF, demonstrate promising results, outperforming several state-of-the-art approaches across vari-




Fig. 5: Randomly selected examples from our worst 25% results of ECCC (top two rows are from validation set and remaining rows are from testing set). (A) Input pair of raw images captured with long and short exposures (note that other methods use a single image captured with auto exposure). (B-G) Images corrected with the estimated illuminant by: (B) FFCC [9], (C) C5 [4] (chosen the best results among all variations discussed in Sec. 4.3), (D) TLCC [54], (E) EMLP, (F) ECCC, and (G) Ground-truth illuminant. The estimated illuminant of each method is shown on the right side of each image, along with the angular error written in the top-left corner of the image.

ous evaluation metrics while utilizing considerably fewer parameters. Qualitative examples of our results, randomly selected from the worst 25% of ECCC results, are presented in Fig. 5 alongside results from other methods.

5 Conclusion and Future Work

In this paper, we introduced DEF, a feature derived from dual exposure images to enhance illuminant estimation. Our DEF achieves comparable results with state-of-the-art methods that employ thousands or millions of parameters (e.g., [54, 63]), using only 354 parameters in a straightforward MLP network. We further discuss incorporating the proposed DEF into the established CCC framework, referred to as ECCC. ECCC achieves comparable or better results on different evaluation metrics than classic CCC approaches while remaining lightweight, requiring only 6,156 parameters (50% reduction compared to FFCC [9]), approximately 30 KB of memory, and running in ~ 0.25 milliseconds per image on CPU. Our solution focuses on the single-camera case, intending testing on the same camera used for training. Future work includes studying the stability of this feature across different cameras. We discussed integrating the proposed DEF into the CCC framework. Further exploration may involve examining DEF’s benefits for other illuminant estimation techniques that rely on convNets and raw image pixels as input. Another research direction could involve developing a spatially varying version of DEF, rather than our global feature, to utilize for spatially varying illuminant estimation and image white balancing.

References

1. Abdelhamed, A., Lin, S., Brown, M.S.: A high-quality denoising dataset for smart-phone cameras. In: CVPR (2018)
2. Abdelhamed, A., Punnappurath, A., Brown, M.S.: Leveraging the availability of two cameras for illuminant estimation. In: CVPR (2021)
3. Affi, M., Abuolaim, A.: Semi-supervised raw-to-raw mapping. In: BMVC (2021)
4. Affi, M., Barron, J.T., LeGendre, C., Tsai, Y.T., Bleibel, F.: Cross-camera convolutional color constancy. In: ICCV (2021)
5. Affi, M., Brown, M.S.: Sensor-independent illumination estimation for DNN models. In: BMVC (2019)
6. Affi, M., Punnappurath, A., Finlayson, G., Brown, M.S.: As-projective-as-possible bias correction for illumination estimation algorithms. *Journal of the Optical Society of America A* **36**(1), 71–78 (2019)
7. Banić, N., Koščević, K., Lončarić, S.: Unsupervised learning for color constancy. arXiv preprint arXiv:1712.00436 (2017)
8. Barron, J.T.: Convolutional color constancy. In: ICCV (2015)
9. Barron, J.T., Tsai, Y.T.: Fast Fourier color constancy. In: CVPR (2017)
10. Bianco, S., Schettini, R.: Color constancy using faces. In: CVPR (2012)
11. Brainard, D.H., Freeman, W.T.: Bayesian color constancy. *Journal of the Optical Society of America A* **14**(7), 1393–1411 (1997)
12. Brown, M.: Color processing for digital cameras, chap. 5, pp. 81–98. John Wiley & Sons, Ltd (2023)
13. Buchsbaum, G.: A spatial processor model for object colour perception. *Journal of the Franklin Institute* **310**(1), 1–26 (1980)
14. Cheng, D., Prasad, D.K., Brown, M.S.: 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)
15. Debevec, P.E., Malik, J.: Recovering high dynamic range radiance maps from photographs. In: Annual Conference on Computer Graphics and Interactive Techniques (1997)
16. Delbracio, M., Kelly, D., Brown, M.S., Milanfar, P.: Mobile computational photography: A tour. *Annual Review of Vision Science* **7**, 571–604 (2021)
17. Endo, Y., Kanamori, Y., Mitani, J.: Deep reverse tone mapping. *ACM Transactions on Graphics* **36**(6) (2017)
18. Fairchild, M.D.: Color appearance models. John Wiley & Sons (2013)
19. Finlayson, G.D., Funt, B.V., Barnard, K.: Color constancy under varying illumination. In: ICCV (1995)
20. Finlayson, G.D., Hordley, S.D.: Color constancy at a pixel. *Journal of the Optical Society of America A* **18**(2), 253–264 (2001)
21. Finlayson, G.D., Hordley, S.D.: The chromagenic colour camera and illuminant estimation. In: Color Imaging Conference (2005)
22. Finlayson, G.D., Hordley, S.D., Morovic, P.: Chromagenic colour constancy. In: Congress of the International Colour Association (2005)
23. Finlayson, G.D., Hordley, S.D., Morovic, P.: Colour constancy using the chromagenic constraint. In: CVPR (2005)
24. Finlayson, G.D., Hordley, S.D., Tastl, I.: Gamut constrained illuminant estimation. *International Journal of Computer Vision* **67**(1), 93–109 (2006)
25. Finlayson, G.D., Trezzi, E.: Shades of gray and colour constancy. In: Color and Imaging Conference (2004)

26. Fredembach, C., Finlayson, G.: The bright-chromagenic algorithm for illuminant estimation. *Journal of Imaging Science and Technology* **52**, 40906:1–40906:11 (2008)
27. Gehler, P.V., Rother, C., Blake, A., Minka, T., Sharp, T.: Bayesian color constancy revisited. In: *CVPR* (2008)
28. Gelfand, N., Adams, A., Park, S.H., Pulli, K.: Multi-exposure imaging on mobile devices. In: *ACM International Conference on Multimedia* (2010)
29. Gijsenij, A., Gevers, T., Van De Weijer, J.: Generalized gamut mapping using image derivative structures for color constancy. *International Journal of Computer Vision* **86**(2), 127–139 (2010)
30. Gijsenij, A., Gevers, T., Van De Weijer, J.: Computational color constancy: Survey and experiments. *IEEE Transactions on Image Processing* **20**(9), 2475–2489 (2011)
31. Gijsenij, A., Gevers, T., Van De Weijer, J.: Improving color constancy by photometric edge weighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **34**(5), 918–929 (2011)
32. Granados, M., Ajdin, B., Wand, M., Theobalt, C., Seidel, H.P., Lensch, H.P.: Optimal HDR reconstruction with linear digital cameras. In: *CVPR* (2010)
33. Hanji, P., Mantiuk, R.K., Eilertsen, G., Hajisharif, S., Unger, J.: Comparison of single image HDR reconstruction methods — the caveats of quality assessment. In: *SIGGRAPH* (2022)
34. Hartigan, J.A., Wong, M.A.: Algorithm as 136: A K-means clustering algorithm. *Journal of the royal statistical society. Series C (applied statistics)* **28**(1), 100–108 (1979)
35. Hernandez-Juarez, D., Parisot, S., Busam, B., Leonardis, A., Slabaugh, G., McDonagh, S.: A multi-hypothesis approach to color constancy. In: *CVPR* (2020)
36. Hu, Y., Wang, B., Lin, S.: FC4: Fully convolutional color constancy with confidence-weighted pooling. In: *CVPR* (2017)
37. Hubel, P.M., Finlayson, G.D., Hordley, S.D.: White point estimation using color by convolution (2007), US Patent 7,200,264
38. Jung, C., Yang, Y., Jiao, L.: High dynamic range imaging on mobile devices using fusion of multiexposure images. *Optical Engineering* **52**(10), 102004–102004 (2013)
39. Karaimer, H.C., Brown, M.S.: A software platform for manipulating the camera imaging pipeline. In: *ECCV* (2016)
40. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
41. Laakom, F., Passalis, N., Raitoharju, J., Nikkanen, J., Tefas, A., Iosifidis, A., Gabbouj, M.: Bag of color features for color constancy. *IEEE Transactions on Image Processing* **29**, 7722–7734 (2020)
42. Laakom, F., Raitoharju, J., Nikkanen, J., Iosifidis, A., Gabbouj, M.: Robust channel-wise illumination estimation. In: *BMVC* (2021)
43. Le, P.H., Le, Q., Nguyen, R., Hua, B.S.: Single-image HDR reconstruction by multi-exposure generation. In: *WACV* (2023)
44. Liba, O., Murthy, K., Tsai, Y.T., Brooks, T., Xue, T., Karnad, N., He, Q., Barron, J.T., Sharlet, D., Geiss, R., et al.: Handheld mobile photography in very low light. *ACM Transactions on Graphics* **38**(6), 164–1 (2019)
45. Lo, Y.C., Chang, C.C., Chiu, H.C., Huang, Y.H., Chen, C.P., Chang, Y.L., Jou, K.: CLCC: Contrastive learning for color constancy. In: *CVPR* (2021)
46. Loshchilov, I., Hutter, F.: SGDR: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983* (2016)
47. Lou, Z., Gevers, T., Hu, N., Lucassen, M.P., et al.: Color constancy by deep learning. In: *BMVC* (2015)

48. Oh, S.W., Kim, S.J.: Approaching the computational color constancy as a classification problem through deep learning. *Pattern Recognition* **61**, 405–416 (2017)
49. Qi, G., Chang, L., Luo, Y., Chen, Y., Zhu, Z., Wang, S.: A precise multi-exposure image fusion method based on low-level features. *Sensors* **20**(6), 1597 (2020)
50. Qian, Y., Kamarainen, J.K., Nikkanen, J., Matas, J.: On finding gray pixels. In: *CVPR* (2019)
51. Shen, R., Cheng, I., Shi, J., Basu, A.: Generalized random walks for fusion of multi-exposure images. *IEEE Transactions on Image Processing* **20**(12), 3634–3646 (2011)
52. Solhusvik, J., Hu, S., Johansson, R., Lin, Z., Ma, S., Mabuchi, K., Manabe, S., Mao, D., Phan, B., Rhodes, H., et al.: A 1392x976 2.8 μm 120db CIS with per-pixel controlled conversion gain. In: *International Image Sensor Workshop* (2017)
53. Solomatov, G., Akkaynak, D.: Spectral sensitivity estimation without a camera. In: *ICCP* (2023)
54. Tang, Y., Kang, X., Li, C., Lin, Z., Ming, A.: Transfer learning for color constancy via statistic perspective. In: *AAAI* (2022)
55. Ulucan, O., Karakaya, D., Turkan, M.: Multi-exposure image fusion based on linear embeddings and watershed masking. *Signal Processing* **178**, 107791 (2021)
56. Ulucan, O., Ulucan, D., Ebner, M.: Block-based color constancy: The deviation of salient pixels. In: *International Conference on Acoustics, Speech and Signal Processing* (2023)
57. Van De Weijer, J., Gevers, T., Gijsenij, A.: Edge-based color constancy. *IEEE Transactions on Image Processing* **16**(9), 2207–2214 (2007)
58. Willassen, T., Solhusvik, J., Johansson, R., Yaghmai, S., Rhodes, H., Manabe, S., Mao, D., Lin, Z., Yang, D., Cellek, O., et al.: A 1280 \times 1080 4.2 μm split-diode pixel HDR sensor in 110 nm BSI CMOS process. In: *International Image Sensor Workshop* (2015)
59. Xiao, J., Gu, S., Zhang, L.: Multi-domain learning for accurate and few-shot color constancy. In: *CVPR* (2020)
60. Xing, X., Qian, Y., Feng, S., Dong, Y., Matas, J.: Point cloud color constancy. In: *CVPR* (2022)
61. Xu, B., Wang, N., Chen, T., Li, M.: Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853* (2015)
62. Yahiaoui, L., Horgan, J., Deegan, B., Yogamani, S., Hughes, C., Denny, P.: Overview and empirical analysis of ISP parameter tuning for visual perception in autonomous driving. *Journal of Imaging* **5**(10), 78 (2019)
63. Yu, H., Chen, K., Wang, K., Qian, Y., Zhang, Z., Jia, K.: Cascading convolutional color constancy. In: *AAAI* (2020)
64. Yue, S., Wei, M.: Color constancy from a pure color view. *Journal of the Optical Society of America A* **40**(3), 602–610 (2023)