Off the Planckian Locus: Using 2D Chromaticity to Improve In-Camera Color

SaiKiran Tedla*[1,3], Joshua Little*[1]
Hakki C. Karaimer[2], and Michael S. Brown[1]

[1]York University [2]AI Center-Toronto, Samsung Electronics [3]University of Toronto
{tedlasai, jelittle, mbrown}@yorku.ca {hakki.k}@samsung.ca

Abstract. Traditional in-camera colorimetric mapping relies on correlated color temperature (CCT)–based interpolation between pre-calibrated transforms optimized for Planckian illuminants such as CIE A and D65. However, modern lighting technologies such as LEDs can deviate substantially from the Planckian locus, exposing the limitations of relying on conventional one-dimensional CCT for illumination characterization. This paper demonstrates that transitioning from 1D CCT (on the Planckian locus) to a 2D chromaticity space (off the Planckian locus) improves colorimetric accuracy across various mapping approaches. In addition, we replace conventional CCT interpolation with a lightweight multi-layer perceptron (MLP) that leverages 2D chromaticity features for robust colorimetric mapping under non-Planckian illuminants. A lightbox-based calibration procedure incorporating representative LED sources is used to train our MLP. Validated across diverse LED lighting, our method reduces angular reproduction error by 22% on average in LED-lit scenes, maintains backward compatibility with traditional illuminants, accommodates multi-illuminant scenes, and supports real-time in-camera deployment with negligible additional computational cost. Code and data can be found on the project webpage: `ccmmlp.github.io`

Keywords: Color Constancy · Camera Pipeline

1 Introduction

Digital cameras perform a wide range of in-camera processing routines to convert sensor-specific RGB responses into the final output image [15]. A critical component of the in-camera processing pipeline is the colorimetric mapping stage, which performs illumination correction followed by a color correction stage, also known as color space transformation (CST) [24, 25], to map sensor values to a standard color space such as CIE XYZ [14]. The accuracy of this colorimetric mapping impacts color reproduction quality across different lighting conditions.

The current in-camera colorimetric mapping method was developed when early digital cameras primarily captured images illuminated by conventional light sources (e.g., incandescent, fluorescent, and natural/outdoor light). The

* Equal contribution.

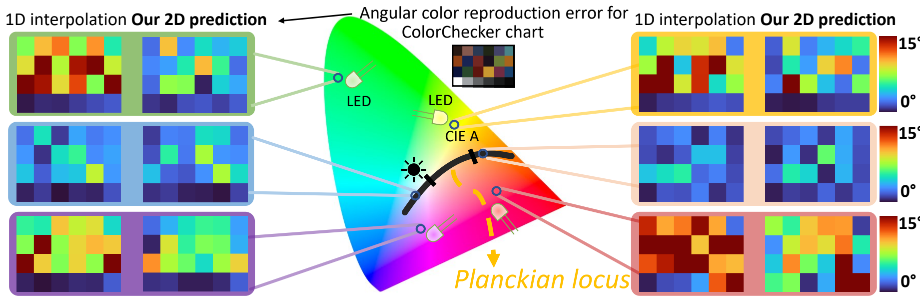

Fig. 1: 1D Planckian interpolation vs. our 2D chromaticity-based prediction. Traditional 1D methods perform adequately for light sources whose CCTs lie close to the Planckian locus but are less accurate for light sources (namely, LEDs) that lie off the Planckian locus. Our approach leverages 2D chromaticity coordinates for more robust color reproduction across all illuminant types. Heat maps show angular reproduction error, from 0° (dark blue) to 15° (dark red), demonstrating comparable performance along the Planckian curve but substantial improvements in LED-dominated regions where existing methods break down.

perceived colors of these early light sources' spectra are conveniently correlated with blackbody radiators—idealized physical bodies whose spectral distributions depend solely on temperature according to Planck's law [51]. In the CIE xy chromaticity space [14], these blackbody sources form the Planckian locus, as shown in Fig. 1. Correlated color temperature (CCT) provides a simple 1D representation for characterizing such conventional light sources. The standard procedure for color correction involves computing the color correction matrix (CCM) using a one-dimensional CCT-based interpolation of pre-calibrated matrices along the Planckian locus. Typically, two pre-calibrated illuminants are used (1) daylight (6500 K) and (2) incandescent light (CIE A).

However, the proliferation of LED technology has fundamentally disrupted this framework. Modern artificial illuminants exhibit chromaticity coordinates that deviate significantly from the Planckian locus, resulting in suboptimal color reproduction when using CCT-based interpolation (Fig. 1). This mismatch highlights the growing disparity between current camera color-processing assumptions and the spectral characteristics of modern illumination.

This paper proposes shifting from a one-dimensional CCT space to a 2D chromaticity space for CCM estimation. Our approach recognizes that modern illuminants require characterization beyond simple CCT, incorporating camera-specific chromaticity coordinates as input features. Moreover, we employ a low-parameter multi-layer perceptron (MLP) to learn optimal interpolation. By training on an expanded dataset including representative LED illuminants alongside traditional sources, our CCM-MLP achieves superior colorimetric mapping that adapts to contemporary lighting spectral diversity while maintaining backward compatibility.

Our key contributions are as follows: (1) CCM-MLP, a lightweight MLP that predicts CCMs in 2D chromaticity space, replacing traditional CCT-based interpolation with learned transforms capable of handling non-Planckian illuminants; (2) adaptation of traditional and baseline methods to operate in the chromaticity domain; (3) extensive validation demonstrating an average 22% improvement in color reproduction accuracy across two smartphones and two DSLRs; (4) a spatially varying extension for multi-illuminant scenarios; and (5) a carefully collected laboratory and in-the-wild dataset captured under non-Planckian illuminants for evaluating our method. This work represents a step toward modernizing in-camera color processing in line with the widespread adoption of LED lighting.

2 Motivation and Related Work

Motivation. The colorimetric mapping stage of the digital camera pipeline (top/middle of Fig. 2 consists of two main steps: white balance (WB) via a diagonal correction and color balance (CB) via a CCM. In the camera pipeline, WB is typically modeled in the 2D raw chromaticity space [24]. Other software, such as Lightroom [2], allows 2D WB control via temperature and tint [41] sliders. While estimated WB values are represented in 2D chromaticity space, the color-balancing stage of the camera pipeline projects 2D chromaticities onto a 1D CCT and utilizes a CCT-based interpolation [1] to derive the CCM (see Supplementary Sec. S5 for details on this procedure).

CCT-based interpolation assumes that conventional light sources lie close to the Planckian locus, however, LEDs and tunable lighting violate the core assumption, as their chromaticity coordinates often deviate significantly from the Planckian locus [5,38]. This divergence arises from inherent LED properties, including discontinuous spectra [26] and narrow-band primaries [45].

When illumination chromaticity coordinates fall far from the Planckian locus, CCT calculations become ill-defined [41], and subsequent CCT-based interpolation leads to errors exceeding just-noticeable color differences [7]. Karaimer and Brown [25] introduced an additional calibration point along the Planckian locus to improve interpolation, an approach later integrated into the Adobe DNG specification [1]. However, because it still relies on CCT-based interpolation, this method cannot represent light sources that fall off the Planckian locus. Other research has explored improved illumination estimation [9, 22, 29, 52] and post-processing strategies [13, 18, 25, 46], yet these approaches do not address the fundamental limitation of CCT-based interpolation in the colorimetric mapping stage. To overcome this, we propose predicting the CCM directly in the 2D chromaticity space using a CCM-MLP, as illustrated in Fig. 2 (bottom).

In the following, we present a short review of relevant areas for in-camera color reproduction.

Color constancy. Early color-correction stages in camera ISP are motivated by the human visual system's color constancy ability to perceive object colors similarly under different illuminations. Computational color constancy generates

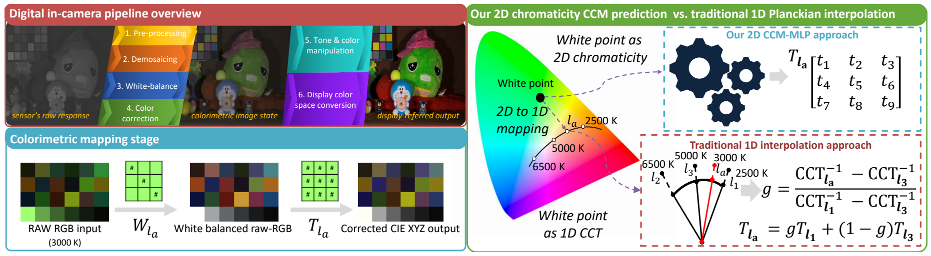

Fig. 2: Top Left: Camera pipeline overview. Early stages convert raw sensor responses to colorimetric values, while later stages transform colors to a photo-finished display-referred space [24]. Our work aims to optimize the colorimetric mapping stage. **Bottom Left: Colorimetric mapping stage.** When an image is captured under arbitrary illumination, the estimated illumination is used to remove color cast via a diagonal white balance 3×3 matrix (W_{l_a}), then mapped to CIE XYZ with an illumination-specific 3×3 CCM (T_{l_a}). **Right: 2D chromaticity CCM prediction vs. 1D CCT interpolation.** CCM interpolation approaches (e.g., $T_{l_a} = gT_{2500K} + (1-g)T_{6500K}$) assume CCM parameters can be derived from a CCT-based interpolation, given calibrated illuminants along the Planckian Locus. CCT-based interpolation works poorly for non-Planckian illuminants. We propose a 2D CCM-MLP that takes chromaticity coordinates as input and directly predicts the full 3×3 CCM.

sensor responses that are invariant to scene illumination in two steps: estimating scene illumination in raw sensor space, then removing the color cast.

White balance. Most color constancy works focus on illumination estimation using classical statistics [6, 17, 20, 50], gamut-based methods [16, 19], and deep-learning approaches [4, 33, 35, 46]. After illuminant estimation, correction typically applies a diagonal 3×3 matrix to the raw-RGB image. Simple WB correction only ensures proper correction of neutral colors ($R = G = B$ after correction).

Color balance. Additional correction is required for non-neutral colors. Thus, diagonal correction provides white balance (WB), while full correction is color balance (CB). Diagonal correction suffices for particular illuminations (e.g., broadband) to achieve full CB, but many illuminations require full color balancing [11]. Camera ISPs apply an illumination-specific CCM [25] after WB to correct non-neutral colors. Adobe DNG implements this by mapping illuminants to CCT, then interpolating between calibrated 3×3 CCM's that map to CIE XYZ. Initially, DNG supported two CCM's at high (∼6500 K) and low (∼2500 K) CCT. Karaimer and Brown [25] showed that an additional calibration point (∼4300 K) improves reproduction, resulting in DNG now supporting three CCMs. Other higher-dimensional mappings, such as polynomial [12] and root-polynomial [18], could be used in place of the 3×3 transformation when invertibility is not required.

Direct mapping. Recent methods such as CCCNN [36] and EXPINV [31] learn an illuminant-dependent colorimetric mapping by training MLPs that take raw pixel values and illuminant CCT as input to predict CIE XYZ values directly.

However, these networks require expensive per-pixel evaluation at test time and rely on access to camera spectral sensitivities during training. Their pixel-wise formulation introduces spatial inconsistencies and color casts, as we demonstrate in Sec. 4.2. Finally, their dependence on 1D CCT illumination coordinates limits these networks' ability to model non-Planckian illuminants. In contrast, our proposed CCM-MLP predicts a global CCM in a single forward pass, enforces spatial consistency, requires no spectral sensitivity data, and generalizes better across diverse light sources by operating in 2D chromaticity space.

3 Method

We aim to learn a mapping from illuminant chromaticity to a CCM using a two-stage method, visualized in Fig. 3, involving calibration data capture within a lightbox (Sec. 3.1) and training a lightweight CCM-MLP (Sec. 3.2). Building a drop-in replacement for the color correction stage allows us to maintain the two-stage approach [1] that currently dominates in-camera pipelines while keeping training requirements down.

3.1 Data Capture

Traditional color correction calibration methods calibrate CCMs by capturing color charts under standard illuminants such as CIE A or D65. However, we require our method to handle non-Planckian illuminants that a CCT cannot adequately describe. The first step in our process is capturing calibration data across diverse illuminants, providing broad coverage of the xy chromaticity space.

We use a modified GTI lightbox [21] to capture a ColorChecker [8] target under a wide range of lighting conditions using two smartphones and two DSLRs (Samsung S22, Google Pixel 9 Pro, Sony α57, and Canon EOS 5D Mark III). This modified lightbox replaces the original light sources with Telelumen multispectral LED lighting [48], which consists of seven narrow-band LEDs that can be adjusted to match a target spectral power distribution (SPD). Using the Telelumen lightbox, we capture the color chart under many SPDs, resulting in a large calibration dataset.

We capture each color chart under 790 different illuminant SPDs sourced from two methods. First, we utilize 390 illuminants from Punnappurath et al. [43] that estimate various light sources, including sunlight, fluorescent, incandescent, and LED. We use the original splits of 195 training, 78 validation, and 117 test illuminants. Then, we introduce a set of 400 illuminants generated by linearly combining the Telelumen LED SPDs with weights randomly sampled from a Dirichlet distribution (see Supplementary Sec. S2 for details). Using this distribution results in an even spread of illuminants across the chromaticity space and allows us to sample more illuminants farther from the Planckian Locus. We randomly split the dataset into 200 train, 80 validation, and 120 test illuminants. We combine both datasets for our calibration, resulting in 395 train, 158 validation, and 237 test illuminants.

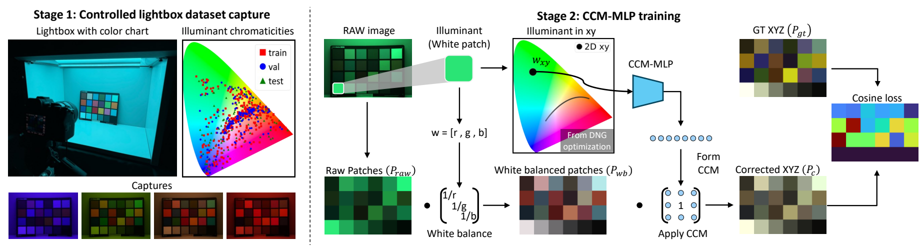

Fig. 3: Our 2D chromaticity CCM prediction framework. *(Stage 1) Controlled lightbox dataset capture:* A ColorChecker [8] target is captured under diverse illumination, including traditional both Planckian an LED illuminants covering on-locus and off-locus illuminants. (Stage 2) CCM-MLP training: Our MLP takes 2D chromaticity coordinates as input and predicts the color correction matrix. To train CCM-MLP, we begin with raw color patches P_{raw} and apply white balancing using the estimated illuminant w, resulting in white balanced patches P_{wb}. We then convert the illuminant w into chromaticity coordinates w_{xy} via the DNG optimization procedure (see Supplementary Sec. S5). The chromaticity w_{xy} serves as input to our CCM-MLP, which predicts the CCM applied to P_{wb}. Finally, we train with a cosine loss between the corrected patches P_c and the ground-truth patches P_{gt} in XYZ space.

3.2 Model

We propose a model, CCM-MLP, that maps a white point chromaticity w_{xy} to a CCM T. Following Karaimer and Brown [25], we use the setting of mapping white-balanced raw images into the XYZ canonical space, and assume the raw image I_{raw} and ground-truth white point chromaticity $w = [r/g, b/g]$ are known. Following previous work [25], we use the ground-truth white point from the color chart white patch and compute the white-balanced raw image I_{wb} by applying a diagonal correction [4]. For our model, we convert the raw-space chromaticity w to its xy-space counterpart w_{xy} using the DNG [1] optimization procedure described in Supplementary Sec. S5. This algorithm is fast, commonly implemented on modern phones [25], and executes in an average of 250 μs on an iPhone 16.

CCM-MLP is a ReLU [40] MLP with one hidden layer of 32 neurons and a linear output layer of eight neurons, trained with the Adam optimizer [30]. Since the angular-error loss defines the CCM only up to scale, we fix its central element to 1 to remove this ambiguity, leaving eight free entries for the network to predict. These are reshaped into the 3×3 CCM. If invertibility is not a requirement, we also discuss predicting polynomial and root-polynomial mappings in Sec. 4.

To train our method, we extract the mean color of the 24 patches in the raw image $P_{raw} \in \mathbb{R}^{24 \times 3}$ by averaging across an 11×11 pixel region per patch. We then apply a diagonal white balance to each patch using the white patch w, resulting in white-balanced patches P_{wb}. For a given illumination, the model aims to predict a 3×3 CCM that, when applied to P_{wb} minimizes the angular error between the corrected XYZ patches $P_c \in \mathbb{R}^{24 \times 3}$ and the ground-truth XYZ $P_{gt} \in \mathbb{R}^{24 \times 3}$.

We utilize our calibration dataset to train the model and use the ground-truth XYZ values provided by X-Rite [8] for the color chart. Following Karaimer and Brown [25], we aim to minimize the angular error of the corrected XYZ colors and the ground-truth XYZ colors. We found direct angular error loss was unstable during optimization; thus, we follow other works [44, 49] and use a cosine loss, which still minimizes angular errors [49]. We define the loss function as the average cosine loss over all color patches:

$$L(P_c, P_{gt}) = \frac{1}{N} \sum_{i=1}^{N} \left(1 - \cos\left(\hat{P_c^i}, \hat{P_{gt}^i}\right)\right), \tag{1}$$

where $\hat{P_c^i}$ and $\hat{P_{gt}^i}$ are the normalized colors at index i.

4 Evaluation

First, we evaluate 1D CCT vs. 2D chromaticity-based approaches using the lightbox dataset described in Sec. 3.1. Next, we test our lightbox calibration across a set of in-the-wild scenes. Finally, we evaluate on LSMI [28] under both single- and multi-illuminant settings, and the NUS [10] dataset.

4.1 Evaluation Details

Baselines. For comparison, we implement a variety of state-of-the-art baselines. First, to contextualize the upper bound of CCM prediction, we compute the Oracle, the best-fit 3×3 matrix that minimizes the angular error given the full color chart. Then, we compare with the standard 2-CCM interpolation described by the DNG [1], which uses matrices recalibrated by capturing the color chart under Standard Illuminant A (2850 K) and D65 (6500 K) lighting within the lightbox. We also compare to the 3-CCM [25] interpolation with an additional calibrated CCM at 4250 K. Next, we implement a nearest neighbors (NN) [42] baseline that maps each training white point with its best-fit 3×3 CCM (minimizing angular error on the color chart). At test time, given a white point coordinate, we retrieve the nearest illuminant by L2 distance and apply its corresponding matrix. Finally, we compare direct MLP methods, CCCNN [36] and EXPINV [31], that perform per-pixel correction given an input color and white point.
1D CCT vs 2D Chromaticity. To demonstrate the importance of using 2D chromaticity coordinates, we also compare our CCM-MLP with a 1D CCT-based variant. Additionally, we adapt the 1D CCT-based baseline methods (NN, EXPINV, and CCCNN) into 2D chromaticity-based versions to show the benefits of our approach.
Polynomial mappings. To evaluate the effect of higher-dimensional transforms, we extend our CCM-MLP to output polynomial [12] and root-polynomial [18] mappings. We compare standard linear 3×3 transforms against polynomial and root-polynomial mappings of varying sizes and find that root-polynomial

Table 1: Quantitative evaluation on lightbox dataset. We compare angular error and ΔE_{2000} of traditional 2-CCM/3-CCM interpolation [1, 25], NN [42], CCCNN [36], EXPINV [31], and our CCM-MLP. For NN, CCCNN, EXPINV, and CCM-MLP, we evaluate both 1D CCT implementations and our proposed 2D chromaticity adaptations (highlighted in blue). Methods Marked (RP) predict a root-polynomial [18] mapping.

				Angular Error (°)															
				Sony				**Canon**				**Pixel**				**Samsung**			
Method	**Dim**	**Size (KB)**	**MACs (M)**	Mean	25%	50%	90%	Mean	25%	50%	90%	Mean	25%	50%	90%	Mean	25%	50%	90%
Oracle	1D	0.00	0.01	1.71	0.54	1.15	3.68	2.06	0.64	1.31	4.65	2.79	0.95	2.00	5.94	2.58	0.84	1.83	5.47
2-CCM [1]	1D	0.07	0.01	3.31	0.80	1.98	7.91	3.86	0.96	2.31	9.59	4.56	1.48	3.25	9.94	4.51	1.53	3.19	9.95
3-CCM [25]	1D	0.11	0.01	3.31	0.80	2.01	7.91	3.86	0.94	2.32	9.68	4.53	1.48	3.22	9.95	4.51	1.51	3.18	9.94
NN [42]	1D	15.43	0.01	3.93	0.94	2.33	9.59	4.68	1.05	2.74	11.62	5.24	1.54	3.56	11.73	5.48	1.82	3.59	12.27
CCCNN [36]	1D	68.51	13.39	3.28	1.09	2.13	7.38	3.58	1.19	2.42	8.02	4.39	1.47	2.96	9.67	4.71	1.76	3.35	10.00
EXPINV [31]	1D	68.53	13.39	3.13	0.91	1.86	7.42	3.61	1.08	2.33	8.46	4.31	1.26	2.85	9.88	4.68	1.43	3.31	10.40
CCM-MLP	1D	1.41	0.01	3.09	0.91	1.85	7.35	3.52	0.98	2.06	8.60	4.11	1.37	2.62	9.55	4.05	1.29	2.60	9.42
CCM-MLP(RP)	1D	2.57	0.01	2.82	0.85	1.73	6.65	3.26	0.96	2.05	7.68	3.77	1.34	2.45	8.57	3.70	1.21	2.43	8.52
NN [42]	2D	16.97	0.01	3.24	0.82	1.79	7.73	3.56	0.90	2.03	8.51	4.30	1.37	2.84	9.81	4.15	1.33	2.72	9.48
CCCNN [36]	2D	69.01	13.48	2.89	0.97	1.99	6.43	3.00	1.13	2.20	6.19	4.00	1.50	2.93	8.76	4.17	1.78	3.32	8.52
EXPINV [31]	2D	69.03	13.48	2.66	0.87	1.78	5.97	3.11	0.94	2.26	6.89	3.84	1.32	2.71	8.41	4.14	1.49	3.05	8.86
CCM-MLP	2D	1.54	0.01	2.66	0.82	1.70	6.14	2.89	0.87	1.85	6.67	3.60	1.22	2.38	8.18	3.60	1.20	2.43	8.23
CCM-MLP(RP)	2D	2.70	0.01	2.36	0.71	1.44	5.53	2.60	0.71	1.58	6.17	3.27	1.13	2.14	7.47	3.28	1.17	2.20	7.25
				$\boldsymbol{\Delta E_{2000}}$															
Oracle	1D	0.00	0.01	5.37	2.91	4.70	9.49	5.83	3.22	5.20	10.17	7.07	4.55	6.41	12.14	6.51	4.09	5.94	11.06
2-CCM [1]	1D	0.07	0.01	7.21	3.78	6.06	13.18	8.15	4.41	6.63	15.08	9.18	5.68	7.91	15.74	8.61	5.28	7.64	14.76
3-CCM [25]	1D	0.11	0.01	7.20	3.77	6.02	13.19	8.13	4.40	6.65	15.10	9.16	5.68	7.93	15.65	8.62	5.28	7.65	14.76
NN [42]	1D	15.43	0.01	8.02	4.00	6.63	15.03	9.21	4.57	7.20	17.91	9.53	5.57	8.07	17.03	9.49	5.36	8.18	16.81
CCCNN [36]	1D	68.51	13.39	7.00	3.53	5.90	12.55	7.93	4.16	6.53	14.42	8.36	4.91	7.40	14.39	8.45	5.05	7.96	13.93
EXPINV [31]	1D	68.53	13.39	7.03	3.55	5.98	12.42	7.76	3.99	6.75	14.25	8.09	4.83	7.14	13.97	8.17	5.07	7.29	13.98
CCM-MLP	1D	1.41	0.01	7.09	3.86	6.20	12.56	7.81	4.25	6.87	14.35	8.48	4.99	7.45	14.58	8.14	4.65	7.51	14.10
CCM-MLP(RP)	1D	2.57	0.01	7.60	4.97	6.83	12.60	8.28	5.36	7.25	14.01	8.19	4.60	6.89	14.82	8.39	4.84	7.09	14.65
NN [42]	2D	16.97	0.01	7.29	3.60	5.99	13.17	7.79	3.92	6.37	14.21	8.74	5.26	7.35	15.30	8.37	4.80	7.16	14.71
CCCNN [36]	2D	69.01	13.48	6.63	3.69	5.62	11.40	6.78	3.58	6.07	11.61	8.07	4.88	7.42	12.91	7.76	4.86	7.66	11.96
EXPINV [31]	2D	69.03	13.48	6.48	3.45	5.84	11.03	7.07	3.98	6.52	11.93	7.77	4.74	7.10	12.63	7.68	4.37	7.51	12.20
CCM-MLP	2D	1.54	0.01	6.73	3.66	6.07	11.58	7.17	3.96	6.41	12.55	7.95	4.78	7.21	13.26	7.74	4.38	7.29	13.11
CCM-MLP(RP)	2D	2.70	0.01	7.16	4.57	6.48	11.56	7.47	4.28	6.83	12.68	7.74	4.50	6.64	13.84	8.08	4.56	7.10	13.97

3×6 mappings yield the lowest angular error (see Tab. 2). We denote the 3×6 root-polynomial version as CCM-MLP (RP) in all subsequent tables.

Training details. We use the dataset splits in Sec. 3.1 to train and evaluate all methods except CCCNN [36] and EXPINV [31], which learn a per-pixel mapping; following their original papers, we train them on a larger simulated dataset. Additionally, we adapt EXPINV with the CCT conditioning utilized in CCCNN [36]. For simulation, we measure each camera's spectral sensitivities using CamspecV2 [23], collect illumination SPDs from our lightbox capture setup (see Sec. 3.1), and utilize 1269 reflectance spectra from the Munsell Book of Color [37] measured by Chromaxion [39]. Finally, we synthesize data using the image formation model described in CCCNN [36] and EXPINV [31].

All MLP methods are trained for 100,000 iterations using the Adam [30] optimizer with a learning rate of 0.001. Our CCM-MLP is trained using the cosine loss described in Sec. 3.2, while CCCNN [36] and EXPINV [31] are trained with the recommended ΔE loss. Spectral sensitivities are unavailable for the LSMI [28] and NUS [10] datasets, so CCCNN and EXPINV are excluded from those evaluations.

Noise augmentation. When queried with out-of-domain chromaticities not seen in our training set, our 2D CCM-MLP can sometimes fail to generalize. As illustrated in Fig. 4, the training distribution does not contain all in-the-wild illuminants; these chromaticities are outside of the range of our lightbox LEDs. To let our models generalize beyond this range, we add Gaussian noise to the input chromaticity coordinates during training [32]. However, a lightbox with a broader range of narrow-band LEDs, or synthetic data augmentation discussed in Supplementary Sec. S1.1, could provide a similar coverage. For CCM-MLP, we add noise with standard deviation $\sigma_0 = 0.05$ to the normalized input coordinates for all training iterations. We show the positive effects of this augmentation in our ablation study (Tab. 2). Finally, we exclude this augmentation for direct-prediction MLP baselines (EXPINV [31], CCCNN [36]) as it can degrade their performance (see Supplementary Sec. S1.3).

Fig. 4: We visualize the chromaticities of our train (lightbox) and in-the-wild illuminants, which include a few out-of-distribution illuminants (circled purple).

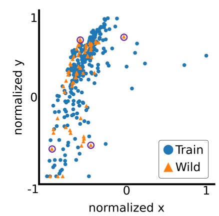

Metrics. For evaluation, we utilize angular error following previous work [25]. We also compare ΔE_{2000} after applying a uniform scaling to all colors to match the Y channel of the third gray patch (bottom row of color chart). Finally, we report model parameter sizes (in KB) and inference costs (in MACs) when applied to a 1080×720 image.

4.2 Lightbox Dataset Results

We compare all methods quantitatively on our dataset in Tab. 1 and qualitatively in Fig. 5 (additional visualizations in Supplementary Sec. S2). We find that *all methods benefit from using a 2D chromaticity coordinate* instead of a 1D CCT, which confirms our motivation for this work. Additionally, our 2D CCM-MLP surpasses traditional interpolation and nearest-neighbor methods, providing a strong solution that integrates seamlessly into the existing DNG [1] pipeline, which relies on linear transformations.

In-the-wild results. After calibrating the CCM prediction methods with lightbox images, we evaluate on a smaller set of 150 in-the-wild photos captured with both smartphones (Pixel and Samsung). As shown in Fig. 6, methods that rely on 2D chromaticity deliver qualitatively superior results, enabling more accurate color correction under challenging LED illumination across diverse regions containing skin tones, posters, and craft materials. We also observe that EXPINV [31], a per-pixel prediction method, exhibits color casts when applied to full-size images. Additional comparisons are provided in Supplementary Sec. S2.

Comparing CCM-MLP with per-pixel MLPs. We find our 2D CCM-MLP outperforms 1D MLP-based direct prediction baselines, CCCNN [36] and EXPINV [31]. Additionally, highlighting the strength of utilizing 2D chromaticity input, our 2D chromaticity-based variants of CCCNN [36] and EXPINV [31] achieve lower ΔE_{2000} in many cases and surpass CCM-MLP under challenging

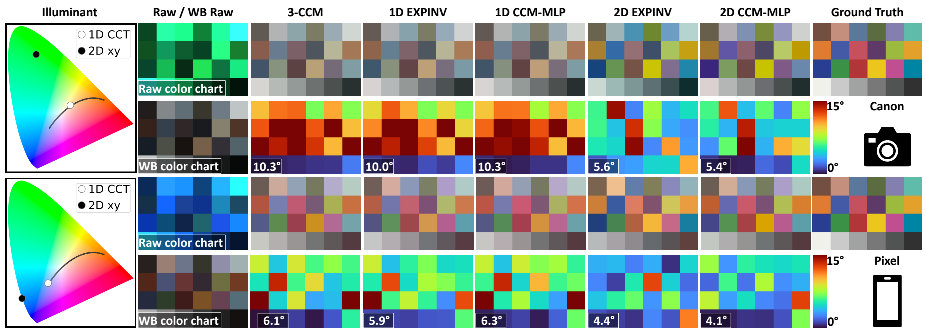

Fig. 5: Colorimetric mapping results on laboratory dataset. Results for two non-Planckian LED sources: green LED (top half) and blue LED (bottom half) illuminants, captured using Canon and Pixel cameras. Each half starts with the illuminant's CCT and xy chromaticity and contains two rows: the top row shows the raw-RGB color chart and display-referred renderings; the bottom row shows the white-balanced color chart and error maps. Methods: 3-CCM [25], 1D EXPINV [31], 1D CCM-MLP, 2D EXPINV [31], 2D CCM-MLP, and ground truth.

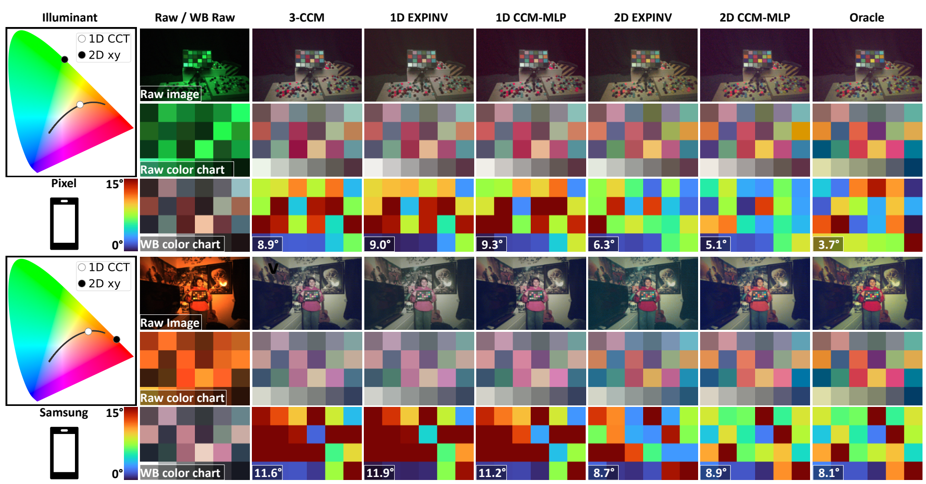

Fig. 6: Colorimetric mapping results on in-the-wild captures. Results for two non-Planckian LED sources: green LED (top half) and orange LED (bottom half) illuminants, captured using Pixel and Samsung cameras. Each half starts with the illuminant's CCT and xy chromaticity and contains three rows: the top row shows raw-RGB image, and display-referred renderings; the middle row shows raw-RGB color chart and display-referred renderings; the bottom row shows white-balanced color chart and angular error maps. Methods: 3-CCM [25], 1D EXPINV [31], 1D CCM-MLP, 2D EXPINV [31], 2D CCM-MLP, and Oracle.

illumination conditions (higher percentiles). This performance gain arises because per-pixel MLPs can model more complex mappings that depend on both

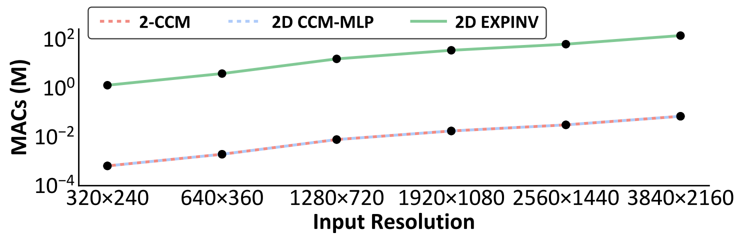

Fig. 7: Computational cost comparison. We plot the multiply–accumulate operations (MACs, log scale) at different resolutions for 2-CCM [1] interpolation, EXPINV [31], and our 2D CCM-MLP. Our CCM-MLP has nearly the same computational cost as traditional methods, since the CCM is computed once and applied to the entire image. In contrast, approaches such as EXPINV must evaluate an MLP at every pixel.

input color and white point, whereas our method predicts a global transformation that depends only on the white point. However, per-pixel methods remain unsuitable for two reasons. First, CCM-MLP has a footprint of 1.5 KB and 0.01 MMACs, compared to 69 KB and 13 MMACs for per-pixel methods. CCM-MLP is efficient because it predicts a single global transform. In Fig. 7, we show that the computational cost of a per-pixel MLP far exceeds our CCM-MLP at various input resolutions. Second, as shown in Fig. 6, applying per-pixel MLP methods to full-size images can introduce color casts, likely because each pixel is transformed independently, leading to spatial inconsistencies. In contrast, CCM-MLP applies a single transformation across the image, ensuring spatial consistency.

Ablation of CCM-MLP. In Tab. 2, we ablate three design choices. First, we compare specifying the whitepoint in 2D xy chromaticity, which yields better performance than using 1D CCT space or 2D raw space. Second, we find

Table 2: Ablation of CCM-MLP. We evaluate CCM-MLP on Pixel camera images from both the lightbox and in-the-wild scenes. We test 1D and 2D input coordinates with and without augmentation. Our results show that using xy chromaticity yields the lowest angular error among all white point parameterizations. Finally, we extend our analysis to higher-dimensional mappings using the xy chromaticity white point.

Input coord.	Mapping	Noise Aug.	Lightbox		In-the-wild	
			Ang. (°)	ΔE	Ang. (°)	ΔE
1D (CCT)	Linear (3×3)	✗	4.11	8.48	3.68	7.05
2D (raw)	Linear (3×3)	✗	3.58	7.99	4.02	7.81
	Linear (3×3)	✓	3.75	8.05	4.14	8.33
2D (xy)	Linear (3×3)	✗	3.53	7.92	3.64	7.33
	Linear (3×3)	✓	3.60	7.95	3.32	6.71
	Polynomial (3×9)	✓	3.67	8.78	3.67	8.78
	Root Polynomial (3×6)	✗	3.13	7.45	3.40	7.15
	Root Polynomial (3×6)	✓	3.27	7.73	3.07	6.77

that our noise augmentation strategy is crucial for our method to generalize to out-of-domain illuminant chromaticities. Finally, we compare the performance of predicting different transforms and find that linear and root-polynomial mappings yield the best results. We further ablate CCM-MLP with varying layer and node counts in Supplementary Sec. S1.3.

Sensitivity to imperfect input white point. Our experiments assume that the ground-truth illuminant is known to isolate the performance of various CCM prediction methods. In practice, this assumption does not hold, as white-balance estimation modules [3,27] are imperfect. To better assess real-world performance, we evaluate our method (trained with ground truth white points) alongside several baselines on our lightbox dataset with controlled white point offsets, and on the in-the-wild dataset using white points derived from an off-the-shelf cross-camera white point estimator [27].

Lightbox dataset. For each illuminant, an offset of $0°$, $1°$, $2°$, $3°$, or $10°$ was applied in random directions. Our results show CCM-MLP achieves the lowest mean angular error when the white point is offset by $0° - 2°$, which falls within the typical range of white point estimators such as CCMNet [27] and Time-Aware White Balance [3], which achieve $1.01° - 2.23°$ mean and $0.60° - 1.71°$ median angular error on datasets that include LED illuminants. When the white point offset is large ($3°$,$10°$), our 2D CCM-MLP performs similarly to 1D CCM-MLP, but still outperforms 2-CCM interpolation.

In-the-wild dataset. We apply CCMNet [27], a cross-camera illumination estimation method, to our in-the-wild data as a worst-case illumination estimate, while masking the color chart in each image. On the in-the-wild dataset, CCMNet recovers the illuminant with a mean angular error of $3.21°$. Again, we find that all methods benefit from using 2D chromaticity input.

Table 3: White-point error ablation. Angular and ΔE error for 2-CCM, 3-CCM, CCCNN, EXPINV, NN, and CCM-MLP under varying white point offsets on the lightbox and in-the-wild datasets (Pixel). Each offset is applied in a random direction; for in-the-wild, we report results with ground truth and CCMNet [27] estimated white-points ($3.21°$ estimation error). 2D-input methods are highlighted in blue. CCM-MLP performs best at $0°$–$2°$; 1D and 2D methods are comparable under large offsets ($\geq 3°$).

		White-point offset:										**In-the-wild**			
		0°		**1°**		**2°**		**3°**		**10°**		**0°**		**CCMNet [27]**	
Method	**Dim**	**Ang.**	**ΔE**	**Ang.**	**ΔE**	**Ang.**	**ΔE**	**Ang.**	**ΔE**	**Ang.**	**ΔE**	**Ang.**	**ΔE**	**Ang.**	**ΔE**
2-CCM	1D	4.56	9.18	5.08	9.80	6.28	11.30	7.87	13.16	21.05	26.87	3.86	7.04	6.23	10.51
3-CCM	1D	4.53	9.16	5.05	9.78	6.25	11.27	7.84	13.13	20.97	26.80	3.81	7.00	6.14	10.42
NN [42]	1D	5.24	9.53	6.07	10.75	6.65	11.72	7.90	13.44	20.27	26.99	5.43	9.53	7.47	12.69
CCCNN [36]	1D	4.39	8.36	4.93	9.10	5.91	10.62	7.13	12.31	17.21	23.84	3.58	7.06	7.96	14.67
EXPINV [31]	1D	4.31	8.09	4.88	9.05	5.98	10.67	7.32	12.42	17.29	23.80	3.35	6.82	5.54	11.03
CCM-MLP	1D	4.11	8.48	4.61	9.17	5.86	10.86	7.46	12.82	19.88	25.93	3.68	7.05	6.09	11.01
CCM-MLP(RP)	1D	3.77	8.19	4.42	9.00	5.80	10.72	7.52	12.71	18.50	24.91	3.45	7.63	6.04	11.37
NN [42]	2D	4.30	8.74	5.27	9.94	7.03	12.03	9.15	14.33	22.68	29.08	5.06	9.13	7.45	12.61
CCCNN [36]	2D	4.00	8.07	4.68	8.86	5.86	10.41	7.27	12.20	17.44	23.94	3.47	6.79	10.77	15.73
EXPINV [31]	2D	3.84	7.77	4.62	8.70	5.95	10.40	7.40	12.23	17.31	23.67	3.07	6.42	5.84	10.64
CCM-MLP	2D	3.60	7.95	4.30	8.82	5.76	10.61	7.54	12.64	20.48	25.96	3.32	6.71	6.01	10.88
CCM-MLP(RP)	2D	3.27	7.74	4.03	8.59	5.59	10.39	7.42	12.42	18.31	24.36	3.07	6.77	5.88	10.79

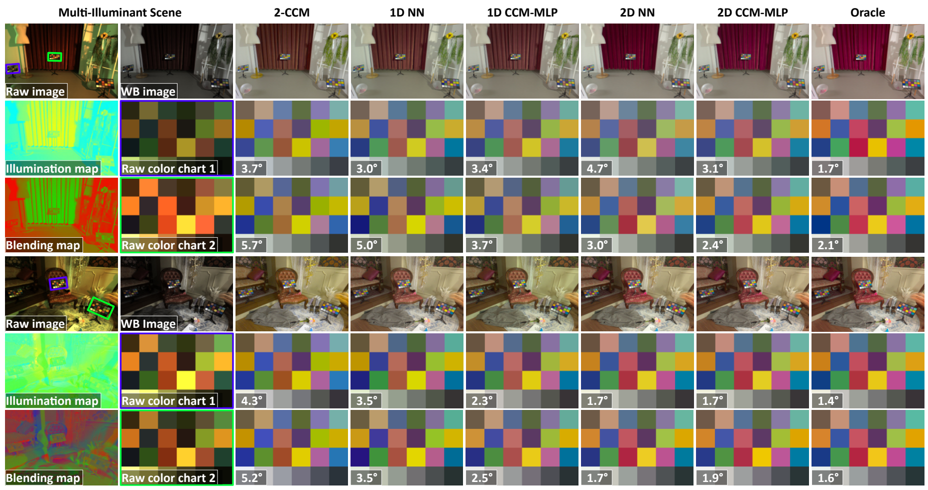

Fig. 8: Colorimetric mapping results on multi-illuminant data from LSMI [28]. We visualize two multi-illuminant scenes from the Galaxy camera. **Row structure for each half:** Top row displays raw-RGB image, white-balanced raw-RGB image, and display-referred renderings; middle row shows contrast-stretched illumination map, first raw-RGB color chart, and display-referred renderings; bottom row presents blending map, second raw-RGB color chart, and display-referred renderings. **Methods from left to right:** 2-CCM [1], 1D NN [42], 1D CCM-MLP, 2D NN [42], 2D CCM-MLP, and Oracle. Angular errors are displayed at the bottom left of each color chart image.

4.3 LSMI Dataset Results

We evaluate CCM-MLP on both single- and multi-illuminant scenes from the Large-Scale Multi-Illuminant (LSMI) dataset [28]. LSMI was originally proposed for white-balance; we repurpose it for colorimetric evaluation by utilizing the color charts in each scene. LSMI consists of 2,700 scenes captured using 2 smartphones and 2 DSLR cameras. Each scene contains multiple color charts and two to three distinct illuminants, providing both single- and multi-illuminant images.

Single-illuminant results. We train our CCM-MLP and all baseline models on single-illuminant images from the LSMI dataset, since we do not have access to lightbox captures. We filter the LSMI dataset to include scenes where all patches on all color charts are not clipped, resulting in 1504 images for Galaxy and 1414 for Sony. We then split the data into train, validation, and test sets using 50%, 20%, and 30% of the images, respectively. For this method, we are unable to compare with CCCNN [36] and EXPINV [31] because these methods require access to the cameras' spectral sensitivities, which are not provided with the LSMI dataset. Summarized results comparing our 2D CCM-MLP with 2-CCM interpolation (using DNG CCMs), nearest-neighbor (NN), and 1D CCM-MLP baselines are presented in Tab. 4 (left). We find that our 2D CCM-MLP

Table 4: LSMI results under single and mixed-illuminant settings. Comparison of 2D chromaticity-based CCM-MLP and NN [42] (highlighted blue) with baselines on Galaxy and Sony data.

		Single-illuminant				Multi-illuminant			
		Galaxy		Sony		Galaxy		Sony	
Method	**Dim**	**Ang.**	ΔE	**Ang.**	ΔE	**Ang.**	ΔE	**Ang.**	ΔE
Oracle	1D	2.38	4.90	1.59	3.52	2.06	3.92	1.44	3.01
2-CCM [1]	1D	5.00	8.43	8.40	9.80	3.88	6.87	8.14	9.08
NN [42]	1D	4.46	7.40	3.07	5.19	3.23	5.42	2.43	4.26
CCM-MLP	1D	3.77	6.30	2.70	4.55	2.52	4.28	1.88	3.40
CCM-MLP (RP)	1D	3.60	6.54	2.37	4.45	2.35	4.62	1.74	3.55
NN [42]	2D	3.98	6.56	2.48	4.39	2.83	4.76	1.93	3.61
CCM-MLP	2D	3.49	5.82	2.49	4.22	2.26	3.88	1.82	3.28
CCM-MLP (RP)	2D	3.37	6.12	2.31	4.50	2.14	4.20	1.69	3.63

achieves the highest performance on this dataset. Additional details and results are provided in Supplementary Sec. S3.

Multi-illuminant results. Assuming access to per-pixel illuminant information and blending maps that quantify each illuminant's contribution, we can apply our MLP trained on single-illuminant images to multi-illuminant scenes. Specifically, for each illuminant, we predict and apply the CCM to the white-balanced image to generate multiple corrected versions. We then merge these corrected images using the blending map as weights. A detailed description of this procedure is provided in Supplementary Sec. S3. In Tab. 4 (right), we present quantitative comparisons of our method. Additionally, in Fig. 8, we present qualitative results on 2- and 3-illuminant scenes from the LSMI dataset. We observe that our 2D CCM-MLP outperforms other methods quantitatively and qualitatively, demonstrating the benefits of using a 2D chromaticity representation in multi-illuminant scenarios. Additional details and results are provided in Supplementary Sec. S3.

4.4 NUS Dataset Results

We compare on four cameras from the NUS [10] dataset. We report quantitative results for two cameras in Table 5. Similarly to our lightbox dataset, the 2D CCM-MLP outperforms the 2-CCM baseline but performs equally well as the 1D CCM-MLP. This is due to the illuminant distribution of the NUS dataset, which is dominated by illuminants that can be effectively modeled with CCT. These results indicate that our method matters most in the presence of LED illuminants.

Table 5: Quantitative evaluation on the NUS dataset. We compare the 2-CCM [1] and NN [42] methods against our proposed CCM-MLP using angular error and ΔE_{2000}. We find that the difference between 1D and 2D CCM-MLPs is small, likely because the dataset's light sources do not strongly deviate from the Planckian Locus. Methods using our proposed 2D chromaticity input are highlighted in blue.

Method	Dim	Size (KB)	MACs (M)	Nikon		Olympus		Panasonic		Samsung	
				Ang.	ΔE	Ang.	ΔE	Ang.	ΔE	Ang.	ΔE
Oracle	1D	0.00	0.01	1.00	2.29	1.16	2.60	1.21	2.67	0.92	2.08
2-CCM [1]	1D	0.07	0.01	1.65	2.88	2.10	3.50	2.11	3.38	1.69	2.83
NN [42]	1D	3.67	0.01	1.38	2.57	1.67	3.07	1.53	2.91	1.31	2.42
CCM-MLP	1D	1.41	0.01	1.29	2.47	1.47	2.89	1.45	2.80	1.13	2.22
CCM-MLP(RP)	1D	2.57	0.01	1.12	2.84	1.21	3.20	1.18	3.28	1.02	2.05
NN [42]	2D	4.04	0.01	1.41	2.62	1.64	3.06	1.52	2.91	1.37	2.47
CCM-MLP	2D	1.54	0.01	1.28	2.46	1.46	2.87	1.44	2.80	1.15	2.23
CCM-MLP(RP)	2D	2.70	0.01	1.11	3.17	1.20	2.86	1.16	3.25	1.01	2.03

5 Conclusion

This paper demonstrates that shifting from conventional 1D CCT-based interpolation to 2D chromaticity-based neural prediction can improve in-camera colorimetric accuracy. Additionally, our investigation finds that *using the 2D chromaticity space improves accuracy for all methods* by capturing illuminant characteristics that are lost in CCT parameterization. Moreover, our CCM-MLP leverages this representation effectively while remaining lightweight and efficient. We demonstrate that our 2D CCM-MLP outperforms existing methods on a newly introduced lightbox dataset, on in-the-wild scenes, and on both single- and mixed-illuminant data from the LSMI [28] dataset. Our approach has a computational cost nearly identical to current CCT-based interpolation methods, making it suitable for industry adoption and laying a foundation for next-generation camera color processing robust to illumination diversity.

Limitations. Our method shares constraints with other color correction methods. First, it depends on imperfect white-balance estimators, whose errors propagate into all downstream tasks. Second, our work is limited by metamers [34]—distinct spectra that produce identical camera responses. In our setting with LED illuminants, we observe an increase in metamers due to their narrow-band characteristics [47]. Finally, our advantage over baseline methods diminishes near the Planckian locus (see Section 4.4), indicating that our method should be interpreted as an improvement focused on modern artificial lighting rather than for all illuminants. Please see Supplementary Sec. S1.4 for more discussion.

Acknowledgements

We thank Trevor Canham, Jason J. Yu, and Konstantinos Derpanis for their valuable discussions and feedback throughout this project.

References

1. Adobe Inc.: Adobe Digital Negative (DNG) Specification. https://helpx.adobe.com/camera-raw/digital-negative.html (2004), accessed: 2026-03-01
2. Adobe Inc.: Adobe Lightroom. https://lightroom.adobe.com/ (2007), accessed: 2026-03-01
3. Afifi, M., Zhao, L., Punnappurath, A., Abdelsalam, M.A., Zhang, R., Brown, M.S.: Time-aware auto white balance in mobile photography. In: ICCV (2025)
4. Barron, J.T.: Convolutional Color Constancy. In: ICCV (2015)
5. Boyce, P.R.: Human Factors in Lighting. CRC Press (2014)
6. Buchsbaum, G.: A spatial processor model for object colour perception. Journal of the Franklin Institute (1980)
7. Buzzelli, M.: Visualizing Perceptual Differences in White Color Constancy. Journal of Imaging Science and Technology (2023)
8. Calibrite LLC: ColorChecker Classic. https://calibrite.com/us/product/colorchecker-classic/ (1976), accessed: 2026-03-01
9. Chang, C.W., Fan, C.D., Chang, C.C., Lo, Y.C., Tseng, Y.C., Huang, J.L., Liu, Y.L.: GCC: Generative Color Constancy via Diffusing a Color Checker. In: CVPR (2025)
10. 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 (2014)
11. Cheng, D., Price, B., Cohen, S., Brown, M.S.: Beyond White: Ground Truth Colors for Color Constancy Correction. In: ICCV (2015)
12. Cheung, V., Westland, S., Connah, D., Ripamonti, C.: A Comparative Study of the Characterisation of Colour Cameras by Means of Neural Networks and Polynomial Transforms. Coloration Technology **120**(1), 19–25 (2004)
13. Chiu, Y.C., Chen, G.R., Chen, Z., Peng, Y.T.: ABC-Former: Auxiliary Bimodal Cross-domain Transformer with Interactive Channel Attention for White Balance. In: CVPR (2025)
14. de l'Éclairage (CIE), C.I.: Commission Internationale de l'Éclairage Proceedings. Cambridge University Press, Cambridge, England (1931)
15. Delbracio, M., Kelly, D., Brown, M.S., Milanfar, P.: Mobile computational photography: A tour. Annual Review of Vision Science (2021)
16. Finlayson, G.D.: Corrected-moment illuminant estimation. In: ICCV (2013)
17. Finlayson, G.D., Hordley, S.D., Hubel, P.M.: Color by correlation: A simple, unifying framework for color constancy. IEEE TPAMI (2001)
18. Finlayson, G.D., Mackiewicz, M., Hurlbert, A.: Color Correction Using Root-Polynomial Regression. IEEE Transactions on Image Processing **24**(5), 1460–1470 (2015)
19. Forsyth, D.A.: A novel algorithm for color constancy. International Journal of Computer Vision (1990)
20. Gijsenij, A., Gevers, T., van de Weijer, J.: Improving color constancy by photometric edge weighting. IEEE TPAMI (2012)
21. GTI Graphic Technology Inc.: GTI ColorMatcher Series. https://www.gtilite.com/products/color-matching-systems/gti-colormatcher-series/ (2004), accessed: 2026-03-01
22. Hernandez-Juarez, D., Parisot, S., Busam, B., Leonardis, A., Slabaugh, G., McDonagh, S.: A multi-hypothesis approach to color constancy. In: CVPR (2020)

23. Image Engineering GmbH & Co. KG: camSPECS. https://www.image-engineering.de/products/equipment/measurement-devices/588-camspecs-express (2025), accessed: 2026-03-01
24. Karaimer, H.C., Brown, M.S.: A Software Platform for Manipulating the Camera Imaging Pipeline. In: ECCV (2016)
25. Karaimer, H.C., Brown, M.S.: Improving Color Reproduction Accuracy on Cameras. In: CVPR (2018)
26. Khanna, V.K.: Fundamentals of Solid-State Lighting. CRC Press (2014)
27. Kim, D., Afifi, M., Kim, D., Brown, M.S., Kim, S.J.: Ccmnet: Leveraging calibrated color correction matrices for cross-camera color constancy. In: ICCV (2025)
28. Kim, D., Kim, J., Nam, S., Lee, D., Lee, Y., Kang, N., Lee, H.E., Yoo, B., Han, J.J., Kim, S.J.: Large Scale Multi-Illuminant (LSMI) Dataset for Developing White Balance Algorithm Under Mixed Illumination. In: ICCV (2021)
29. Kim, D., Kim, J., Yu, J., Kim, S.J.: Attentive Illumination Decomposition Model for Multi-Illuminant White Balancing. In: CVPR (2024)
30. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. In: ICLR (2015)
31. Kucuk, A., Finlayson, G.D., Mantiuk, R., Ashraf, M.: An Exposure Invariant Neural Network for Colour Correction. In: Color and Imaging Conference (2022)
32. Ling, S., Nimier-David, M., Jacobson, A., Sharp, N.: Stochastic preconditioning for neural field optimization. In: SIGGRAPH (2025)
33. Lo, Y.C., Chang, C., Chiu, H., Huang, Y.H., Chen, C.P., Chang, Y.L., Jou, K.: Clcc: Contrastive learning for color constancy. In: CVPR (2021)
34. Logvinenko, A.D., Funt, B., Mirzaei, H., Tokunaga, R.: Rethinking Colour Constancy. PLOS One (2015)
35. Lou, Z., Gevers, T., Hu, N., Lucassen, M.P.: Color constancy by deep learning. In: BMVC (2015)
36. MacDonald, L., Mayer, K.: Camera Colour Correction using Neural Networks. In: London Imaging Meeting (2021)
37. Munsell, A.H.: A Color Notation (12th Edition). Munsell Color Company (1971)
38. Murdoch, M.J.: Characterization and Control of a Multi-primary LED Light Lab. Optics Express (2017)
39. Myers, R.D.: Chromaxion: Information for working with color. https://www.chromaxion.com/ (2025), accessed: 2026-06-23
40. Nair, V., Hinton, G.E.: Rectified linear units improve restricted boltzmann machines. In: ICML (2010)
41. Ohno, Y.: Practical Use and Calculation of CCT and Duv. LEUKOS, The Journal of the Illuminating Engineering Society (2013)
42. Peterson, L.E.: K-nearest neighbor. Scholarpedia (2009)
43. Punnappurath, A., Zhao, L., Le, H., Abdelhamed, A., Tedla, S.K., Brown, M.S.: Improved mapping between illuminations and sensors for RAW images. Journal of the Optical Society of America A **42**(8), 1182–1190 (2025)
44. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning Transferable Visual Models From Natural Language Supervision. In: ICML (2021)
45. Schubert, E.F.: Light-Emitting Diodes (2023)
46. Serrano-Lozano, D., Arora, A., Herranz, L., Derpanis, K.G., Brown, M.S., Vazquez-Corral, J.: Revisiting image fusion for multi-illuminant white-balance correction. In: ICCV (2025)
47. Tedla, S., Wang, Y., Patel, M., Brown, M.S.: Analyzing color imaging failure on consumer-grade cameras. Journal of the Optical Society of America (2022)

48. Telelumen LLC: Octa Light Player. https://telelumen.com/wp-content/uploads/2023/11/telelumen-booth-option-brochure-230217.pdf (2004), accessed: 2026-03-01
49. Wang, H., Wang, Y., Zhou, Z., Ji, X., Gong, D., Zhou, J., Li, Z., Liu, W.: Cosface: Large margin cosine loss for deep face recognition. In: CVPR (2018)
50. van de Weijer, J., Gevers, T., Gijsenij, A.: Edge-based color constancy. IEEE TIP (2007)
51. Wyszecki, G., Stiles, W.S.: Color Science (2nd Edition). Wiley (2000)
52. Zhan, F., Yu, Y., Zhang, C., Wu, R., Hu, W., Lu, S., Ma, F., Xie, X., Shao, L.: Gmlight: Lighting estimation via geometric distribution approximation. IEEE Transactions on Image Processing **31**, 2268–2278 (2022)