https://doi.org/10.2352/CIC.2022.30.1.21

Color Reproduction in LED Wall Virtual Production Stages

Laurent Gudemann[1,2], Jan Fröhlich[1,2], and Harald Brendel[2]
[1]Stuttgart Media University, Germany [2]ARRI Munich, Germany

Abstract

Virtual production stages with LED walls utilize illumination, display, and camera equipment which was not designed with this use case in mind. Because the spectral sensitivity of a camera is different from a human observer, a device specific calibration is required. Furthermore, the illumination spectrum emitted by the display contains large gaps in the cyan and yellow wavelength ranges and is dissimilar to the light sources for which cameras are designed. This causes object colors to be reproduced by the camera in an unnatural manner, making cinematographers hesitant to use LED walls as their primary light source.

In this paper, a display calibration and camera color correction workflow for LED wall virtual production stages is proposed. A linear color correction matrix and the spectrum of a multi-channel LED fixture are jointly optimized to better reproduce object colors simultaneously illuminated by an LED display and the multi-channel fixture as they would appear under high CRI (Color Rendering Index) light sources. An alternative color correction method using root polynomials is found to further improve color reproduction. It is shown that the camera's response to the display can be characterized by a linear 3×3 *matrix and the display can be calibrated using the inverse of the color correction, allowing for a color accurate reproduction of a virtual environment.*

Introduction

LED video walls have seen widespread usage in outdoor applications since the late 90s [1] and offer a variety of benefits over other display technologies. Nowadays, they can reach high luminance levels, a wide color gamut, and a large dynamic range in dark viewing environments. Their novel use case in virtual production movie stages, however, differs from the technical requirements of a human viewer in some key aspects. The issues that this paper investigates are related to the reproduction of color when LED walls are photographed with a motion picture camera. These challenges can be divided into three categories:

Display Calibration

Within a virtual production workflow, the creators of the virtual environment need to make color critical decisions on a calibrated monitor and later get the result they expected after the scene is recorded with a camera. When a display is calibrated by the manufacturer, it is done so with a human observer in mind. Cameras, however, have spectral sensitivities different from humans and consequently require a separate calibration. We therefore introduce a display and camera specific color calibration and color management workflow which allows for an accurate reproduction of the virtual environment's colors.

Object Color Reproduction

The appearance of an object's color is dependant on the spectral sensitivity of the observer, the illumination, and the reflectance of the object itself. Digital cameras employ a color correction step in the image processing pipeline such that the color of common objects under common illumination spectra can be reproduced in an accurate and aesthetically pleasing manner. LED screens, on the other hand, are not designed to be used for object illumination. Their white spectrum is different from the common illuminants for which the camera's color correction is optimized and contains large gaps where cyan and yellow wavelengths have a low relative intensity. This paper investigates whether the color fidelity can be improved by altering the camera's image processing.

Color Consistency between Illuminants

In virtual production studios, the lighting provided by the LED walls is usually augmented using traditional film lighting equipment. Because these light fixtures do not exhibit the same spectral gaps as the LED wall, an object's color can appear very different depending on how much each light source contributes to the overall illumination.

Related Work

Reproducing background colors when directly photographing the LED display has been subject of multiple publications which share a similar process. James et al. [2], LeGendre et al. [3], and Weidlich et al.[4] showed that a linear transform can be obtained by measuring the camera's responses to the LED panel's primaries. LeGendre et al. [3] further investigated whether another *post-correction* matrix can be found that corrects the camera image such that objects illuminated by the LED display appear as they would under reference illuminants. By applying the inverse of the *post-correction* matrix to the display signal, they maintain the display calibration. Even though their experimental and theoretical results match closely, a non-negligible error remains in the lit color-checker's reproduction, possibly due to the limitations of the linear *post-correction* method.

In the field of image based lighting, efforts have been made to reproduce a reference illuminant's color rendition using a multispectral light source. LeGendre et al. [5] used a camera to capture color charts lit by a reference illuminant and by each of the multispectral light source's LEDs. They then found a linear combination of LED spectra by regression that most closely reproduced the reference color chart's appearance.

Our Approach

To avoid the limitations of the linear color correction (LCC) method, we chose the root-polynomial color correction method (RPCC) by Finlayson et al. [14]. Because the RPCC has more de-

grees of freedom, it is more prone to overfitting and can produce unstable results outside of the training data. To combat this, we used spectral measurements, camera spectral response functions, and a large dataset of object reflectances rather than photographing a color chart with a limited number of patches.

To improve the color consistency between the LED wall and additional lighting, we jointly optimize the spectrum of a 6-channel LED fixture and the camera's color correction so that the same correction can be applied regardless of each light source's contribution.

Similar to prior work [2, 3, 4], we characterize the camera's response to the LED panel and apply its inversion and the inverted color correction to pre-process the background image. Because an exact inverse of the RPCC cannot be found, we propose a method for approximating the inverse of the RPCC.

Background and Equations

Image Acquisition

Because our methods rely on radiometric simulations, we first explain how we predict the camera's response to the LED panel and to light reflected off an object. We later use this principle to optimize the color correction and the multi-channel illuminant.

The response of the camera sensor can be be computed from radiometric quantities if the light incident on the sensor is created by a light source with spectrum E reflecting off a lambertian object with a spectral reflectance R [6]:

$$c_k = \int_\omega E(\lambda) R(\lambda) Q_k(\lambda)\, d\lambda, \quad k \in R,G,B \tag{1}$$

with ω as the domain of visible light wavelengths and Q_k as the spectral sensitivities for each of the red, green, and blue the color channels k. The same relation can be written in discrete form:

$$\mathbf{c} = \mathbf{R}^\mathsf{T} \operatorname{diag}(\mathbf{E})\, Q, \quad \mathbf{c} \in \mathbb{R}^3 \tag{2}$$

where $\mathbf{R}$ is an n-element vector sampled from $R(\lambda)$, $\mathbf{E}$ is an n-element vector sampled from $E(\lambda)$, and Q is an $n \times 3$ matrix sampled from the color channels k in $Q_k(\lambda)$.

Similarly, the tristimulus response of a human observer to a radiance spectrum can be modeled:

$$\mathbf{c}^{\mathrm{XYZ}} = \mathbf{R}^\mathsf{T} \operatorname{diag}(\mathbf{E})\, Q^{\mathrm{XYZ}} \tag{3}$$

where Q^{XYZ} is obtained from the CIE 1931 2° XYZ color matching functions [7].

In the image processing pipeline of motion picture cameras, the image is then white balanced such that the response $\mathbf{c}^{wb}$ to the illuminant $\mathbf{E}$ has equal elements [8].

$$\mathbf{c}^{wb} = \operatorname{diag}(\mathbf{wb})\, \mathbf{c}; \quad \mathbf{wb} = \frac{1}{c_G(\mathbf{E})}\, \mathbf{c}(\mathbf{E}) \tag{4}$$

Color correction

A camera satisfies the *Maxwell-Ives criterion* (often also called *Luther condition*) if its spectral sensitivities $\mathbf{Q}_k^{\mathrm{cam}}$ are a linear combination of the CIE color matching functions $\mathbf{Q}_k^{\mathrm{XYZ}}$ [9]. Such a camera is called *colorimetric*. It is, however, not always desirable to design the sensor such that it satisfies this condition. Noise performance, for example, can be improved by deviating from the color matching functions [10]. Likewise, the sensors used in motion picture cameras also do not satisfy the Maxwell-Ives criterion and are not colorimetric [8]. This has two main consequences:

1. There is no 3×3 matrix that accurately converts all camera tristimulus values to the correct XYZ values.
2. Two spectra that produce a metameric match for a human observer may not create matching tristimulus values as observed by the camera.

Despite consequence 1, a 3×3 matrix $M_{rgb\to xyz}$ is commonly applied for the conversion to XYZ tristimulus values. While a nonlinear color correction transformation may improve the colorimetric accuracy of the camera, a well-optimized 3×3 matrix can still yield satisfactory results [8]. A matrix also has several advantages, including an easy to compute and numerically exact inverse transformation as well as exposure invariance.

The matrix optimization method recommended for motion picture cameras [8] is similar to the white point preserving color correction by Finlayson et al. [11]. A set of n measured or simulated camera responses is placed in an $n \times 3$ matrix L and the corresponding target values in XYZ are placed in another $n \times 3$ matrix N. The aim of the optimization is to find a matrix $M_{rgb\to xyz}$ such that:

$$\begin{aligned} N &\approx L M_{rgb\to xyz}^\mathsf{T} \\ \text{subject to} \quad M_{rgb\to xyz} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} &= \mathbf{w} \end{aligned} \tag{5}$$

Meaning that $M_{rgb\to xyz}$ is constrained to preserve white $\mathbf{w}$. Finlayson et al. [11] note that, for a white point $\mathbf{w} = (1,1,1)$, this constraint can be enforced by ensuring that the sum of each row in $M_{rgb\to xyz}$ is one. This observation can be generalized to any target white point $\mathbf{w}$ as shown in equation (6). The parameter given to the optimization function can be reduced to a 3×2 matrix U and then expanded into the 3×3 matrix $M_{rgb\to xyz}$ inside the error function as follows:

$$M_{rgb\to xyz}(U,w) = \begin{bmatrix} U_{1,1} & U_{1,2} & (w_1 - U_{1,1} - U_{1,2}) \\ U_{2,1} & U_{2,2} & (w_2 - U_{2,1} - U_{2,2}) \\ U_{3,1} & U_{3,2} & (w_3 - U_{3,1} - U_{3,2}) \end{bmatrix} \tag{6}$$

To find the approximation $\approx$ in equation (5), Finlayson et al. [11] propose the method of least squares:

$$\begin{aligned} U &:= \arg\min_U |N - L M_{rgb\to xyz}^\mathsf{T}(U,w)| \\ &= \arg\min_U \sqrt{\sum_{i=1}^{n} \sum_{j=1}^{3} (N_{i,j} - [L M_{rgb\to xyz}^\mathsf{T}(U,\mathbf{w})]_{i,j})^2} \end{aligned} \tag{7}$$

By slightly modifying the minimization problem, the optimization can be catered towards the color appearance of the human visual system. Instead of computing the squared differences in the target color space, the more perceptually uniform error space

CIELUV is used. The transformation from XYZ to CIELUV values is denoted $\mathscr{L}$.

$$U := \arg\min_{U} \sum_{i=1}^{n} \sqrt{\sum_{j=1}^{3} [\mathscr{L}(N)_{i,j} - \mathscr{L}(LM^{\mathsf{T}}_{rgb\to xyz}(U,w))_{i,j}]^2} \tag{8}$$

Overall, the color processing pipeline from the camera response to a colorimetric representation using a linear color correction method (LCC) can be expressed as follows:

$$\hat{\mathbf{c}}^{XYZ} = M_{rgb\to xyz}\ \mathrm{diag}(\mathbf{wb})\ \mathbf{c} \tag{9}$$

Image Presentation

In order to colorimetrically reproduce the color values encoded in the image on a display, a conversion must occur. First, a function must be known which maps the signal values received by the display to the corresponding spectral radiance. Assuming the signal values of the red, green, and blue color channels each map to the spectrum of the red, green, and blue display pixels respectively with a per-channel electro-optical transfer function EOTF, the spectrum emitted by the display is simple to compute:

$$\mathbf{E}_{disp} = S\ \mathrm{EOTF}(\mathbf{v}) \tag{10}$$

Where S is a $n \times 3$ matrix of sampled spectral radiances with each column representing the spectrum of the display's red, green, and blue color channels respectively.

If the display signal values are pre-processed by the inverse of the EOTF $\mathbf{v} = \mathrm{EOTF}^{-1}(\mathbf{v}_{lin})$, the reproduced stimulus $\mathbf{d}$ of an observer can be modeled using a linear $\mathbf{3} \times \mathbf{3}$ matrix $M_{disp\to obs}$.

$$\mathbf{d} = [S\ \mathrm{EOTF}(\mathrm{EOTF}^{-1}(\mathbf{v}_{lin}))]^{\mathsf{T}}\ Q = S^{\mathsf{T}} Q\ \mathbf{v}_{lin} \tag{11}$$

$$M_{disp\to obs} = S^{\mathsf{T}} Q \tag{12}$$

Categorizing Color Reproduction Errors

The choice of target values N in the equation (8), for which the color correction is optimized, should be discussed. Figure 1 shows a set of circumstances in which a color difference can occur.

Because the objective is to optimize a camera's color correction for specific non-reference light sources (in our case an LED display and a multi-channel LED fixture), only errors E3, E4, and E5 are considered.

E3: Camera error at light source

This error arises when comparing the colorimetric measurement of a given scene under a given illuminant with the reproduction of the same scene and illuminant when observed by a camera. If this error is subject to minimization for the color correction, any color differences caused by error E2 will remain in the final color reproduction.

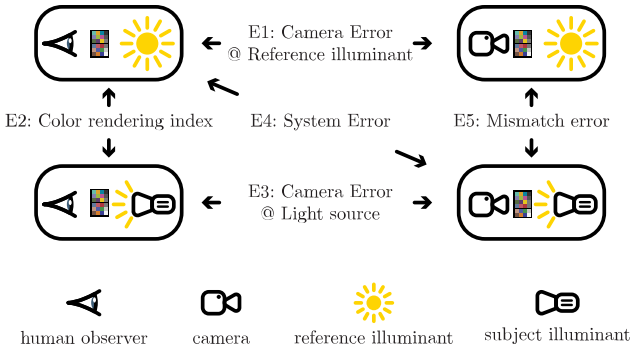

Figure 1: Categories of color errors between observers and illuminants. Based on a figure from [12].

E4: System error

The system error is produced by a camera under a given light source when compared with how a human observer would perceive the same scene under a reference illuminant. Minimizing this error can produce a more natural appearance of memory colors like skin tones [13].

E5: Mismatch error

If the spectrum of a given light source differs from a reference illuminant, the reflected object spectra differ as well and a camera will experience different tristimulus values. If two illuminants are present in a given scene, a color correction which affects the entire image cannot eliminate the mismatch error.

Hypotheses

Previously, linear models for the image acquisition and presentation were described. The calibration and color correction methods proposed in the following sections rely on these models and on a set of hypotheses. If all hypotheses are found to be true, a color management process can be stipulated with which accurate and aesthetically pleasing colors can be captured despite the challenging environment of a virtual production stage.

1. **Display characterization:** There is a linear $3x3$ matrix with which the tristimulus values of the camera can be predicted when photographing the display (equation (11)).
2. **Display calibration:** The inverse of the prediction matrix from hypothesis 1 and the inverse of the camera's color correction can be applied to a target color such that it is reproduced by the camera when photographing the display.
3. **Color correction optimization:** A color correction algorithm can be optimized for a known LED display spectrum such that the system error E4 is minimized. This reduces the unnatural color appearance caused by the low color rendering index E2 of the LED display illuminant.
4. **Multi-channel LED spectrum optimization:** The spectra of a multi-channel LED light fixture can be mixed such that the system error E4 is minimized when choosing the color correction from hypothesis 3. Thus, the same color correction works for both illuminants simultaneously.

Linear Color Correction and Multi-Channel Light Source Optimization

Optimizing a color correction algorithm for the LED display spectrum will negatively affect the color reproduction of commonly used light sources, for which a camera is traditionally optimized (see mismatch error E5 in Figure 1). However, if the secondary light source is similar to the display and produces a low mismatch error, the same color correction will work well under both spectra. To this end, a multi-channel LED fixture such as the ARRI Orbiter allows for precise control over its spectral power distribution (SPD). It emits light from red, green, blue, amber, cyan, and lime (RGBACL) LEDs which were confirmed to behave linearly when configured individually. The resulting spectra (Figure 2) offer a large selection of broad and narrow bandwidths.

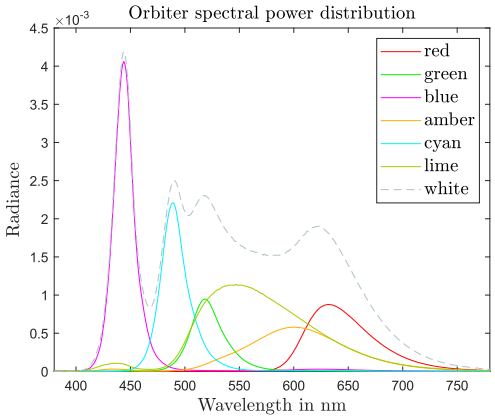

Figure 2: Spectral power distributions of the ARRI Orbiter LEDs reflected off the middle grey patch of an X-Rite ColorChecker Classic. The white spectrum is created by configuring all LEDs to the same intensity percentage.

Enforcing that the color of an illuminant with k channels is equal to the white color of the display as observed by the camera is done by reducing the degrees of freedom from k to $k-3$. In the following, a 6-channel illuminant is assumed. The 6 element vector $\mathbf{o}$ denotes the factors for each of the LEDs. First, an illuminant-to-camera prediction matrix A is obtained by placing the camera's response to the 6 LEDs into columns. Then, the camera's response $\mathbf{i}$ to 3 of the LEDs is computed from the last 3 elements of $\mathbf{o}$.

$$\mathbf{i} = \begin{bmatrix} A_{1,4} & A_{1,5} & A_{1,6} \\ A_{2,4} & A_{2,5} & A_{2,6} \\ A_{3,4} & A_{3,5} & A_{3,6} \end{bmatrix} \begin{bmatrix} o_4 \\ o_5 \\ o_6 \end{bmatrix} \tag{13}$$

The remaining camera response to the first 3 LEDs $\mathbf{j}$ is the difference between the target white response $\mathbf{c}(\mathbf{E})$ and the camera's response $\mathbf{i}$ to the last 3 LEDs:

$$\mathbf{j} = \mathbf{c}(\mathbf{E}) - \mathbf{i} = M_{disp \to obs} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} - \mathbf{i} \tag{14}$$

The remaining LED factors o_1, o_2 and o_3 which produce the remaining camera response $\mathbf{j}$ can now be computed.

$$\begin{bmatrix} o_1 \\ o_2 \\ o_3 \end{bmatrix} = \begin{bmatrix} A_{1,1} & A_{1,2} & A_{1,3} \\ A_{2,1} & A_{2,2} & A_{2,3} \\ A_{3,1} & A_{3,2} & A_{3,3} \end{bmatrix}^{-1} \mathbf{j} \tag{15}$$

The reduced color correction matrix U and the reduced LED configuration vector $\mathbf{o}_{456} = (o_4, o_5, o_6)$ are optimized by minimizing the cost function (16). Figure 3 shows a simplified overview. The weight ω determines the influence of the system error E4 over the mismatch error E5. The weight l determines whether the total system error E4 favors the display or the multi-channel fixture. D and O are $n \times 3$ matrices containing the reproduced colors after the color correction under the display and multi-channel illuminants respectively. They are computed from a spectral dataset of $n = 190$ real world object reflectances and compared with the the target values T under a reference illuminant as observed by the CIE 1931 standard observer.

$$\{U, \mathbf{o}_{456}\} = \underset{U, \mathbf{o}_{456}}{\arg\min} \sqrt{\omega\,(\Delta_{E4})^2 + (1-\omega)(\Delta_{E5})^2} \tag{16}$$

with

$$\begin{aligned} \Delta_{E4} &= l\,\Delta_{E4}^{disp} + (1-l)\Delta_{E4}^{orb} \\ \Delta_{E4}^{disp} &= \sum_{i=1}^{n} \sqrt{\sum_{j=1}^{3} (\mathscr{L}(D)_{i,j} - \mathscr{L}(T))_{i,j})^2} \\ \Delta_{E4}^{orb} &= \sum_{i=1}^{n} \sqrt{\sum_{j=1}^{3} (\mathscr{L}(O)_{i,j} - \mathscr{L}(T)_{i,j}))^2} \\ \Delta_{E5} &= \sum_{i=1}^{n} \sqrt{\sum_{j=1}^{3} (\mathscr{L}(D)_{i,j} - \mathscr{L}(O)_{i,j}))^2} \end{aligned} \tag{17}$$

The resulting matrix U and vector $\mathbf{o}_{456}$ are then expanded into color correction matrix $M_{rgb \to xyz}$ and multi-channel configuration vector $\mathbf{o}$ according to equations (6) and (15).

Root Polynomial Color Correction (RPCC)

To further reduce the system errors, advanced color correction methods need to be considered. For this purpose, the color correction using root-polynomial regression by Finlayson et al. [14] is chosen. This method expands the R, G, and B camera responses by additional terms, which are best explained by showing an example:

$$\bar{\mathbf{p}} = (R, G, B, \sqrt{RG}, \sqrt{GB}, \sqrt{RB}) \tag{18}$$

The 6 element vector $\bar{\mathbf{p}}$ contains the root polynomials of degree 2. In a similar fashion, the root polynomials of degree 3 can be defined, resulting in a 13 element vector. A root polynomial color correction matrix of size 3×6 and 3×13 respectively can be found by regression that, when right multiplied by $\bar{\mathbf{p}}$, minimizes the color errors. The increased degrees of freedom from 9 to 18 or 39 respectively, can produce a better fit than the linear color correction while maintaining the property of exposure invariance [14].

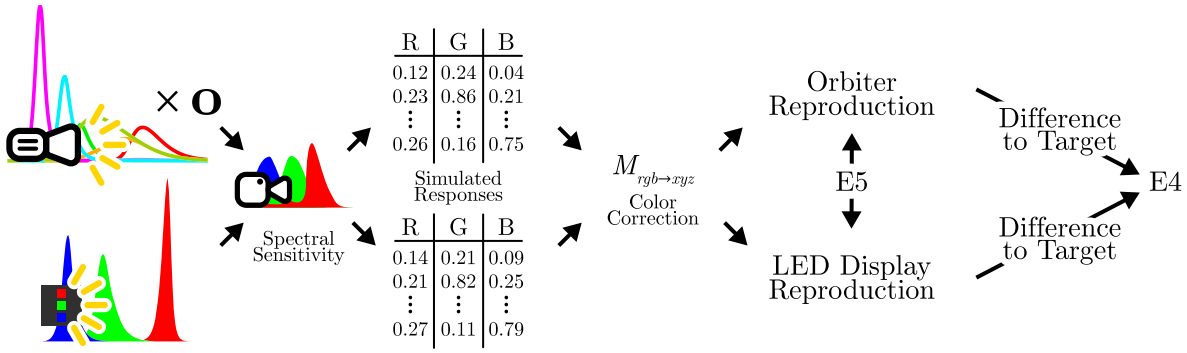

Figure 3: The architecture for the optimization of the multichannel configuration **o** and the color correction $M_{rgb\to xyz}$. A weighted sum of the system error $E4$ and the mismatch error $E5$ is minimized.

Similar to equation (6), a white point preservation can be enforced during the optimization:

$$\mathbf{m}_{last} = \frac{1}{p_d}(\mathbf{w} - U\begin{bmatrix}\bar{p}_1 & \bar{p}_2 & \cdots & \bar{p}_{d-1}\end{bmatrix}^{\mathsf{T}})$$
$$M_{rgb\to xyz} = \begin{bmatrix} U_{1,1} & \cdots & U_{1,n} & m_{last,1} \\ U_{2,1} & \cdots & U_{2,n} & m_{last,2} \\ U_{3,1} & \cdots & U_{3,n} & m_{last,3} \end{bmatrix} \tag{19}$$

where U is the reduced RPCC matrix subject to minimization, $\bar{\mathbf{p}}$ is the root polynomial vector of the camera's response to white after white balancing (usually $(1,1,1)$), and $\mathbf{w}$ is the target white point.

Inversion of the Root Polynomial Method

In order to maintain the display calibration, the inverse of the root polynomial method needs to be applied to the display signal so that the recorded image remains accurate after the color correction is performed. While this is simple for the linear color correction, the RPCC is not as trivial to invert. Because the root polynomial matrix is not square and reduces a 6 or 13 dimensional root polynomial to 3 dimensional XYZ coordinates, no exact inverse matrix can be found.

Instead, another RPCC matrix M^{-1}_{rpcc} can be found which approximates the inverse of the root polynomial method by minimizing:

$$\underset{M^{-1}_{rpcc}}{\arg\min} \sum_{i=1}^{n} \sqrt{\sum_{j=1}^{3} (\mathscr{L}(T)_{i,j} - \mathscr{L}(M_{rpcc}\mathscr{P}(M^{-1}_{rpcc}\mathscr{P}(T)))_{i,j})^2} \tag{20}$$

where $\mathscr{P}(M)$ is the row-wise root polynomial of a matrix and M^{-1}_{rpcc} is constrained to map the target white point $\mathbf{w}$ from equation (19) to the camera's response to white after white balancing. In other words, the minimization optimizes a pseudo-inverse RPCC matrix such that it can be applied to a target image followed by the root polynomial color correction matrix. The resulting color should look identical to the initial target color. This method is exposure invariant and produces low errors as illustrated in Figure 4. In Table 1, the min, median, and max ΔE of the pseudo-inverse for a set of real-world objects as observed by the CIE 1931 standard observer and for colors produced by a ROE Ruby LED wall as observed by an ARRI ALEXA Mini LF is listed.

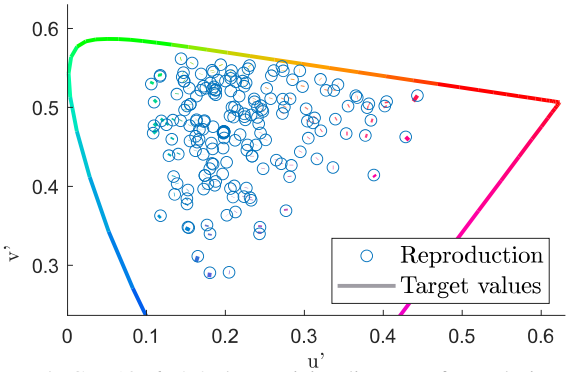

Figure 4: CIE 1976 u'v' chromaticity diagram of pseudo-inverse RPCC errors

	ΔE		
Dataset	mean	median	max
S190	0.59	0.40	2.60
Display gamut	0.70	0.49	2.99

Table 1: CIELAB color errors of the pseudo-inverse RPCC targeting common colors from a set of 190 object colors (top) and scattered synthetic targets within the ROE Ruby display gamut (bottom).

Results

Simulated Results

The minimization of the linear color correction is performed once and the color reproduction errors are simulated separately for the LED display illuminant and an optimized spectrum of the ARRI Orbiter fixture. Figures 6a and 6b show the resulting color errors. To generate the Figures 5a and 5b, a multispectral image from the *CAVE* database is used [15]. The target image is computed using the CIE D65 illuminant and the CIE 1931 standard observer. The reproduced images inside the squares are illuminated by the LED display spectrum and observed by the ARRI ALEXA Mini LF mean spectral sensitivity.

The simulations show that the optimized multi-channel illuminant spectrum can produce a low mismatch error E5. Overall, reproduced colors are a significant improvement over the default matrix but saturated red and green objects are rendered with a greater error. Skin tones are very close to their target chromaticity and only appear slightly lighter.

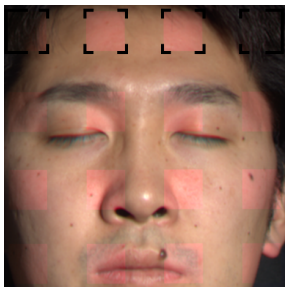

(a) Unoptimized linear color correction (inside squares) vs. target values (outside of squares).

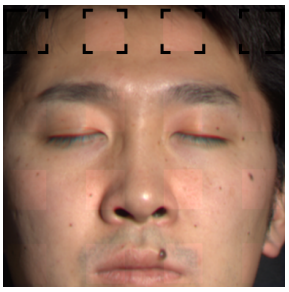

(b) Optimized linear color correction (inside squares) vs. target values (outside of squares).

Figure 5: Simulated colors before (left) and after (right) optimizing the linear color correction.

While the optimized linear color correction is an improvement over the default matrix, it does not produce a low error for all object colors. The inherent limitation of the linear color correction method is that it can only perform linear transformations like rotation, scale, and shear. Improvements to the saturated reds would negatively affect the accuracy of saturated orange or less saturated red objects, causing the optimization to find a compromise with the lowest overall error according to the cost function.

The optimization of the root polynomial method is performed on the same reflectance dataset and the results can be seen in Figures 6c and 6d. A value of 0.8 is once again chosen for the error and illuminant weights. There is a significant improvement in saturated red and green objects over the linear color correction.

Experimental Results and Verification

So far, the color correction and calibration results presented are all simulations based on radiometric measurements. To confirm the efficacy of the methods proposed in this paper, we must apply them in a real world environment and compare the reproduction with the simulated data.

Color correction

The LED wall is configured to display a white image and the ARRI Orbiter's RGBACL color channels are set up according to the optimized vector **o** from equation (16). A test scene containing a person with Caucasian skin, an X-Rite ColorChecker Classic, and various color-critical objects is illuminated from the front by the display and multi-channel illuminant separately. These objects are not included in the training data of the color correction.

A visual comparison highlighting areas of interest is shown in Figure 7. The overall appearance of the colors corrected by the linear and root polynomial methods is similar, especially when looking at skin tones. The default LCC matrix (left) exhibits a strong skin color shift towards magenta-red, which both the LCC and RPCC correct for. The ColorChecker illuminated by the optimized multichannel spectrum (Fig. 7c bottom) does not match the targets as well as the display illuminant (top). This is expected, because the weights chosen during for the optimization favored the display illuminant.

To quantify the results, a color difference is computed between the reproduced ColorChecker patches after color correction

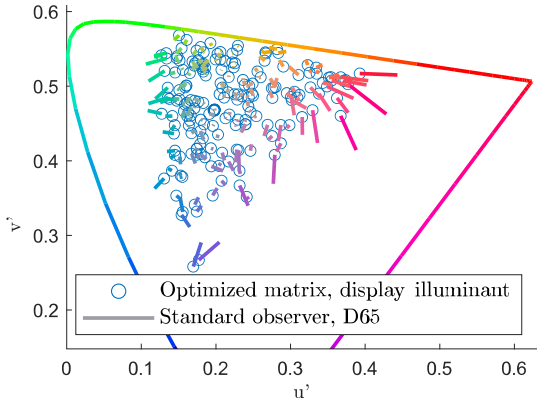

(a) Linear color correction illuminated by the LED display

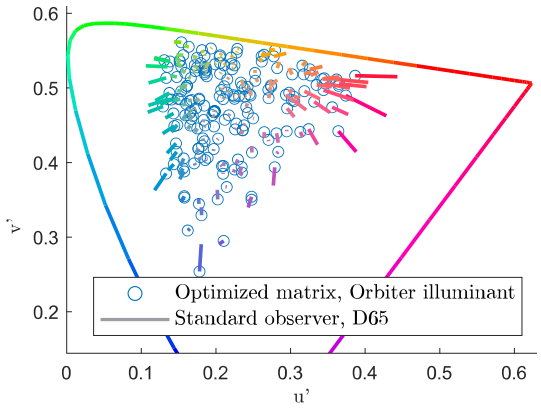

(b) Linear color correction illuminated by the optimized ARRI Orbiter

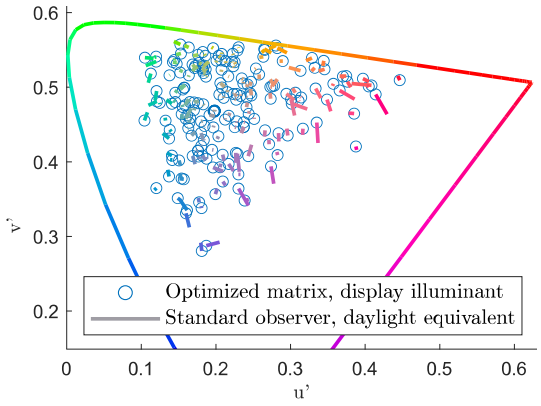

(c) RPCC of degree 3 illuminated by the LED display

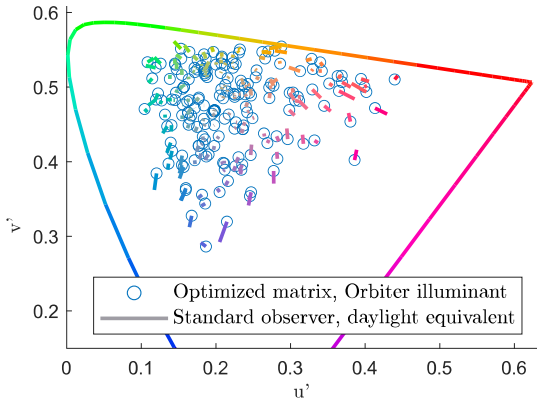

(d) RPCC of degree 3 illuminated by the optimized ARRI Orbiter

Figure 6: Simulated color errors using weights $\omega = 0.8$, $l = 0.8$

and the target color under a reference illuminant. The color space CIELAB is chosen, as it is commonly used for color appearance prediction and has sufficient perceptual uniformity when computing euclidean distances between color coordinates (ΔE) [13]. The results in Table 2 confirm our previous findings. While most errors are greater than the just-noticeable-distance (JND) threshold of about 2.3 [16], both the linear and root polynomial method produce a significantly lower error than the default matrix in every metric. The RPCC is able to reduce the mean, median, and maximum errors better than the LCC and reproduced most ColorChecker patches more accurately. Especially outliers (max in Table 2) were much improved by both methods, with the RPCC achieving an improvement of $\Delta E = 15.5$ and the LCC reducing the error by $\Delta E = 13.5$. Figure 8 shows the system error from the same ColorCheckers patches as u'v' chromaticity diagrams. The RPCC is able to reproduce the chromaticities better for most ColorChecker patches.

Next, the recorded colors are compared with the radiometric simulations (Table 3). If the simulated conditions perfectly matched those of the real world test, we would expect an error of $\Delta E \approx 0$. The remaining differences are most likely caused by inaccuracies in the ColorChecker spectral reflectance and the camera spectral sensitivity data.

Display Calibration

The display calibration is performed by the inverted conversion matrices used in the color processing of the camera image and the inverse of the display-to-camera characterization from equations (11) & (12).
For the linear color correction, the calibration is performed as follows:

$$\mathbf{v} = \mathrm{EOTF}^{-1}(M^{-1}_{disp\rightarrow obs}\,\mathrm{diag}(\mathbf{wb})^{-1} M^{-1}_{rgb\rightarrow xyz}\,\mathbf{v}^{XYZ}_{lin}) \quad (21)$$

The root polynomial method is calibrated in a similar manner:

$$\mathbf{v} = \mathrm{EOTF}^{-1}(M^{-1}_{disp\rightarrow obs}\,\mathrm{diag}(\mathbf{wb})^{-1} M^{-1}_{rpcc}\,\bar{\mathbf{p}}^{XYZ}_{lin}) \quad (22)$$

Where $\bar{\mathbf{p}}^{XYZ}_{lin}$ is the root polynomial of the target color's XYZ coordinates and M^{-1}_{rpcc} is the pseudo-inverse RPCC matrix.

Rather than computing the matrix $M_{disp\rightarrow obs}$ from radiometric functions, we can directly measure it with the camera to improve the accuracy of the matrix. The key insight is that each column in the matrix represents the camera's response to the red, green, and blue spectra respectively.

The results in Figure 9 show that the display calibration is highly accurate for both the linear and root polynomial method at any white balance.

Discussion

When it comes to colorimetric accuracy, the root polynomial method is shown to be superior to the linear method for certain object colors. Especially saturated red and green objects can be reproduced more faithfully. However, when judging the overall appearance of the image, both methods offer a similar improvement over the default matrix. Most objects, including skin, are not saturated enough to produce a visible improvement with the RPCC.

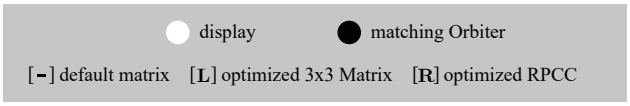

(a) Legend for test configuration indicators

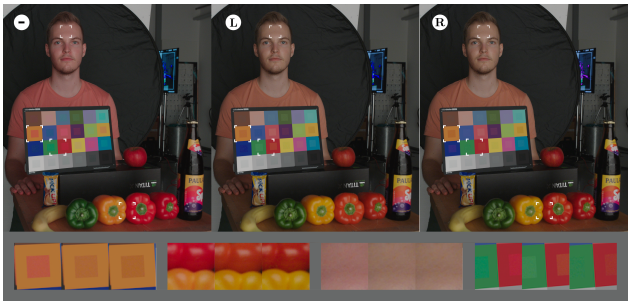

(b) The same image illuminated by the LED display and color corrected with the default correction **-** (left), with the linear method **L** (center), and with the root polynomial method **R** (right)

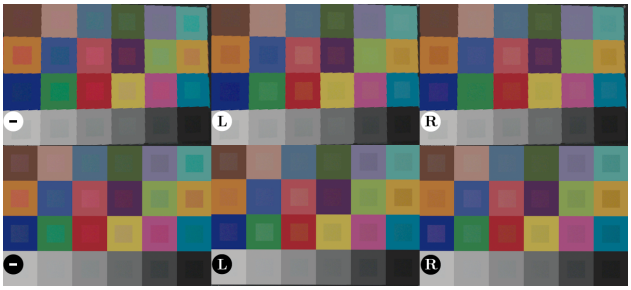

(c) Color checkers illuminated by the display (top) and the optimized ARRI Orbiter (bottom) using the default color correction (left), the linear matrix (center), and the RPCC (right)

Figure 7: Isolated results showing ColorCheckers illuminated by the LED wall and ARRI Orbiter separately

Description	ΔE mean	ΔE median	ΔE max
Default LCC 3200 K display	11.4	10.1	24.5
RPCC 3200 K display	4.7	4.2	9.0
LCC 3200 K display	5.5	4.3	11.0
RPCC 3200 K Orbiter	6.0	6.1	9.8
LCC 3200 K Orbiter	6.2	6.1	12.4

Table 2: Color reproduction errors comparing the ColorChecker patches after color correction with simulated target colors illuminated by a 3200 K black-body reference illuminant.

Description	ΔE mean	ΔE median	ΔE max
RPCC 6025 K display	4.8	4.6	7.8
RPCC 6025 K Orbiter	4.8	4.7	8.2
RPCC 3200 K display	4.3	4.2	8.2
RPCC 3200 K Orbiter	5.2	5.1	10.1

Table 3: Simulation errors comparing the ColorChecker patches after color correction with the simulated color reproduction based on radiometric functions.

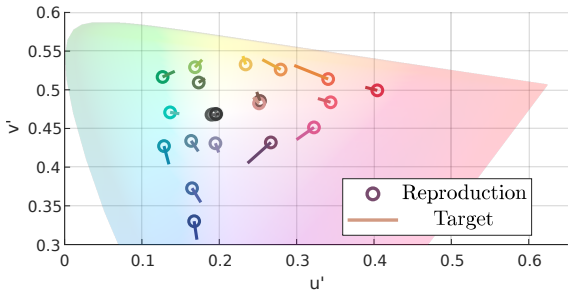

(a) Default LCC

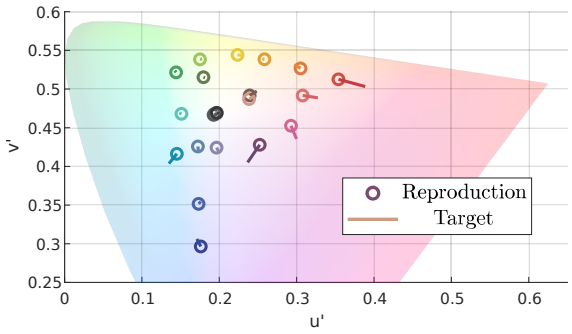

(b) Optimized LCC

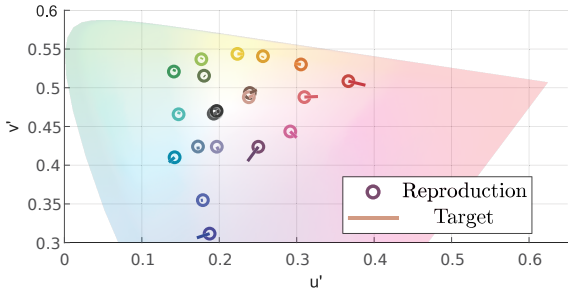

(c) Optimized RPCC

Figure 8: u'v' chromaticity diagram comparing: **Lines:** the simulated ColorChecker target colors illuminated by a 3200 K blackbody reference illuminant as observed by the standard observer. **Markers:** the real-world reproduction using the default LCC (top), the optimized LCC (middle), and the optimized RPCC (bottom) matrix under display illumination.

Comparing the experimental results in Figure 8 with the simulations in Figure 6c and 6d, it is clear that even the most saturated red and green ColorChecker patches are not as critical as the saturated objects used in the training data. It can be argued that the linear method is accurate enough and produces satisfactory results for most object colors.

Conclusion

In this paper, a display calibration and color correction method specific to an LED wall virtual production studio is proposed. It is shown that a highly accurate display calibration matrix can be obtained by measuring the camera's response to each of the display's LEDs.

It is demonstrated that a multi-channel light fixture like the ARRI Orbiter can be configured in conjunction with the camera's color correction to emit a spectrum which produces a low mis-

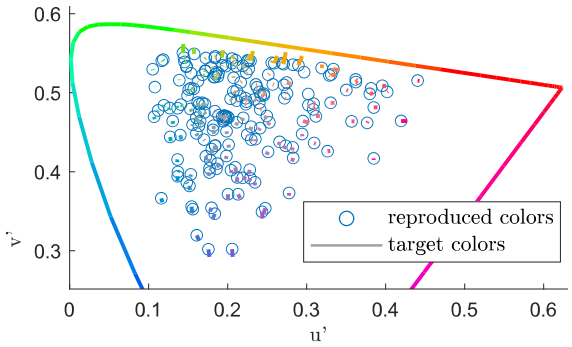

(a) Measured calibration error using the inverse LCC matrix and calibrating the display for a white balance of 4000 K

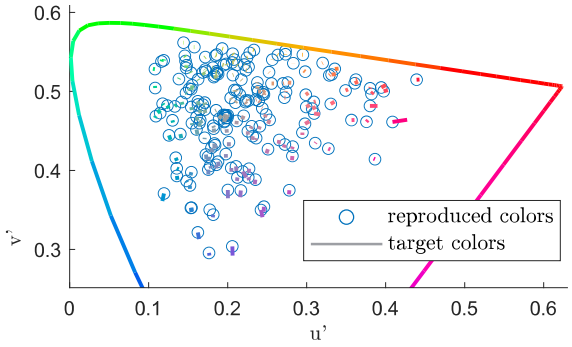

(b) Measured calibration error using the inverse RPCC at the native display white CCT of 6025 K

Figure 9: Real-world calibration performance

match error to the LED display. The same color correction can be applied to objects illuminated by either illuminant, resulting in a consistent and natural presentation.

If the illuminant is carefully chosen, the color correction method presented in this paper offers a robust and inexpensive solution before LED displays with more than three primaries become available. Our method can be performed with any multi-channel light source, as long as its spectral power distribution can be precisely configured and accurately predicted.

References

[1] F. Nguyen, "Challenges in the design of a RGB LED display for indoor applications", In: *Synthetic Metals 122.1. Proceedings of the E-MRS 2000 Spring Meeting, Symposium I*, pp. 215–219, (2001).

[2] O. James, R. Achard, J. Bird, and S. Cooper., "Colour-Managed LED Walls for Virtual Production", In: *ACM SIGGRAPH 2021 Talks*, (2021).

[3] C. LeGendre, L. Lepicovsky, and P. Debevec, "Jointly Optimizing Color Rendition and In-Camera Backgrounds in an RGB Virtual Production Stage", In: *The Digital Production Symposium, ACM*, (2022).

[4] A. Weidlich, A. Forsythe, S. Dyer, T. Mansencal, J. Hanika, A. Wilkie, L. Emrose, and A. Langlands, "Spectral imaging in production: course notes Siggraph 2021", In: *ACM SIGGRAPH 2021 Courses (SIGGRAPH '21)*, (2022).

[5] C. LeGendre, Xueming Yu, Dai Liu, J. Busch, A. Jones, S. Pattanaik, and P. Debevec, "Practical multispectral lighting reproduction", In: *ACM Trans. Graph. 35, 4, Article 32 (July 2016)*, (2016).

[6] Y. Zhu, G. D. Finlayson, "Designing a Color Filter with High Over-

all Transmittance for Improving the Color Accuracy of Digital Cameras", In: *Proc. IS&T 29th Color and Imaging Conf.*, pp 1 - 6, (2001).
[7] CIE, "Commission Internationale de l'Éclairage Proceedings, 1931", Cambridge: Cambridge University Press, (1932).
[8] J. Goldstone et al., "P-2013-001 Recommended Procedures for the Creation and Use of Digital Camera System Input Device Transforms", Academy of Motion Picture Arts and Sciences, (2009).
[9] R. Luther, "Aus dem Gebiet der Farbreizmetrik", In: *Zeitschrift für technische Physik 8*, pp. 540–558. (1927).
[10] H. Kuniba and R. S. Berns, "Spectral sensitivity optimization of color image sensors considering photon shot noise", In: *Journal of Electronic Imaging 18.2*, p. 023002, (2009).
[11] G. Finlayson and M. Drew, "White-point preserving color correction", In: *Color Imaging Conference.*, (1997).
[12] M. Leonhardt, "Optimierung von Mehrkanal-LED-Scheinwerfern für eine digitale Kamera, und vice versa." unpublished, MA thesis, Hochschule Furtwangen University, (2015).
[13] M. D. Fairchild, "Color Appearance Models", (2005).
[14] G. D. Finlayson, Mackiewicz M., and A. Hurlbert, "Color Correction Using Root-Polynomial Regression", In: *IEEE Transactions on Image Processing 24.5*, pp. 1460–1470, (2015).
[15] F. Yasuma, T. Mitsunaga, D. Iso, and S. Nayar, "Generalized Assorted Pixel Camera: Post-Capture Control of Resolution, Dynamic Range and Spectrum", In: *IEEE Transactions on Image Processing 99*, (2010).
[16] W. Mokrzycki, M. Tatol, "Color difference Delta E - A survey. Machine Graphics and Vision", In: *MG&V 20, 4*, 383-411, (2011).

Author Biography

Laurent Gudemann is an image engineer and computer science master's student with a bachelor's degree in audiovisual media. His prior work as part of his bachelor's thesis in the Image Science department of ARRI Munich dealt with the optimization of illuminant spectra, color correction methods, and display calibration in the context of LED wall virtual production.

Jan Fröhlich is Professor for Motion Picture Engineering at Stuttgart Media University (HdM). His research focus is high dynamic range (HDR) and wide color gamut (WCG) image encoding, color rendering and camera metrology. Before becoming a professor he worked as Senior Image Scientist at ARRI Munich and for Dolby Laboratories contributing to the development of the ITU Rec.2100 ICtCp color space and the Dolby Vision HDR and WCG content distribution platform.

Harald Brendel is the head of the Image Science team at ARRI Munich, where he is responsible for the color processing in the AMIRA, ALEXA, and ALEXA35 cameras. He contributed to the development of the ARRI LogC logarithmic color spaces and the AMPAS Academy Color Encoding System (ACES).