# Robust Color Correction for Preserving Spatial Variations within Photographs

Daljit Singh J. Dhillon
Clemson University
Clemson, SC, USA
djdhillon@gmail.com

Parisha Joshi
Clemson University
Clemson, SC, USA
parishj@g.clemson.edu

Jessica Baron
Clemson University
North Charleston, SC, USA
jrbaron@clemson.edu

Eric K. Patterson
Clemson University
North Charleston, SC, USA
ekp@clemson.edu

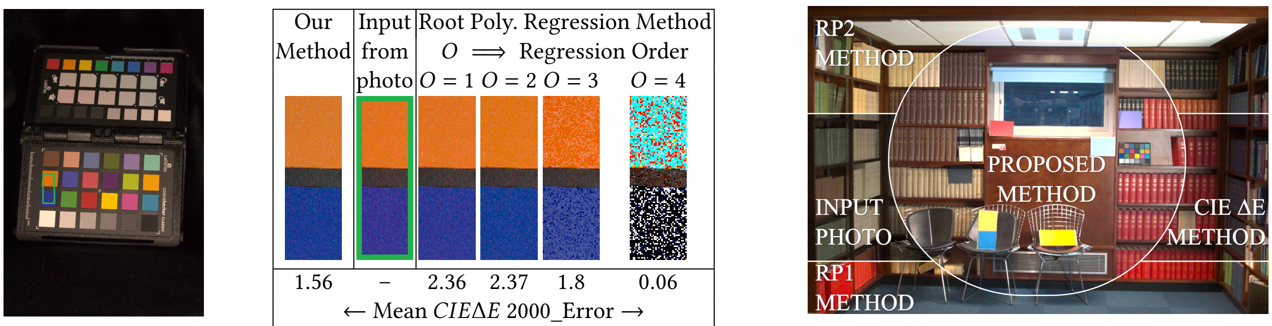


Colorchart Photo    Color correction results    Qualitative comparisons. Input from [Varghese et al. 2014]

**Figure 1:** *A reference colorchart (left image) is commonly used for color correction which is an ill-posed problem. State-of-the-art root-polynomial regression method reduces CIE $XYZ$ or linear-RGB color differences for the transformed reference blocks in the mean-sense. It improves significantly with the increasing regression order as CIEΔE is seen to drop with the increasing order. However, it does not account for spatial variations and produces serious artifacts as demonstrated here (center image). Proposed method improves color correction while preserving spatial variations, white-balancing appropriately and not over-damping the luminance as reported by Varghese et al. for their CIE_ΔE minimizing method.*

**ACM Reference Format:**
Daljit Singh J. Dhillon, Parisha Joshi, Jessica Baron, and Eric K. Patterson. 2023. Robust Color Correction for Preserving Spatial Variations within Photographs. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Posters (SIGGRAPH '23 Posters), August 06–10, 2023, Los Angeles, CA, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3588028.3603681

## 1 INTRODUCTION

We often expect colors captured by cameras to match perceptually well with their real-world counterparts. Color correction methods are devised to fulfill this goal. It is also an essential preprocessing step in colorimetry and image-based material acquisition pipelines for inverse-modeling and/or rendering where accuracy becomes paramount. However, color correction of individual images with mere three color channels is essentially an ill-posed problem (see Section. 2). Commonly, colorcharts with blocks of known reflectances are placed in a scene as references. Current state-of-the-art methods mainly focus on matching the source colors (mean or median values of imaged blocks) to their target (perceptually defined) counterparts. Some operate in *RGB* or $XYZ$ color spaces and use linear, polynomial or root-polynomial based least-mean square formulations to estimate correction transformations [Kucuk et al. 2022]. Others rely on minimizing *CIE_LAB* $\Delta E$ color-difference errors in transformed reference blocks through classical optimization or machine-learning methods. In this work, we show empirically that predicating the optimality of the desired color-correction transformation on reducing such reference color differences may only lead to an over-fitted solution even at its best. We then discuss the importance of spatial variations and propose a metric on local variations to take into consideration while evaluating the goodness of color-correction methods. We propose a novel method that is especially designed to preserve local spatial variations as well, in photos. We demonstrate that our method manages to do both: (a) stay within good tolerance limits for color matching differences, and (b) retain local spatial variations. In contrast, existing methods often overlook spatial variations while overfitting their solution to reduce color differences for reference blocks alone.

## 2 COLOR CORRECTION PROBLEM

*Ill-posedness.* To understand the ill-posedness, we first note that integrating continuous spectral radiances incident on camera pixels into three-channel color vector is a linear operation. Mathematically, it implies that: $\mathbf{T}[(a\cdot s_1(\lambda)+b\cdot s_2(\lambda)] = a\cdot\mathbf{T}[s_1(\lambda)]+b\cdot\mathbf{T}[s_2(\lambda)]$, where $\mathbf{T}$ is the transform from the spectral space to, say, linear *RGB* color space, $s_1$ and $s_2$ are two arbitrary spectral radiance profiles within the visible bandwidth, and $a, b$ are arbitrary scalars. However, **R**, **G** and **B** dimensions for a given camera's colorspace are represented as three $(x, y)$ points in CIE's $x-y$ chromaticity space and


*SIGGRAPH '23 Posters, August 06–10, 2023, Los Angeles, CA, USA*

ACM ISBN 979-8-4007-0152-8/23/08.
https://doi.org/10.1145/3588028.3603681

**Table 1: Error Statistics**

| Method (Order) | CIE ΔE_2000 Error | | | 100×Mean \|ΔCoV\| | | |
|---|---|---|---|---|---|---|
| | Mean | Median | Max | $x$ | $y$ | $Y$ (Luma) |
| Root-Poly (1) | 2.36 | 2.12 | 6.70 | 0.10 | 0.06 | 0.01 |
| Root-Poly (2) | 2.37 | 1.86 | 5.55 | 0.20 | 0.10 | 0.12 |
| Root-Poly (3) | 1.80 | 1.09 | 5.25 | 2.42 | 1.40 | 3.56 |
| Root-Poly (4) | 0.06 | 6E-4 | 0.68 | 466 | 411 | 169 |
| Our (L-2,C-2) | 1.56 | 1.46 | 3.81 | 0.27 | 0.11 | 0.13 |

these $(x, y)$ points vary depending on the camera's color functions. Thus, an accurate mapping between two different *RGB* colorspaces defined by two distinct sets of color functions is essentially a many-to-many association as each camera in this case exhibits a different metamerism (See Figure 1 in the supplemental). Now, if one of these cameras emulates CIE's color functions, then solving for the said mapping is essentially the color correction problem. Even when the reference blocks map to distinct points in the chromaticity space, perceptually metameric spectral radiances can map to different points in the x-y chromaticity space when imaged.

*Camera Noise and Colorchart Inaccuracies.* The mapping problem is further complicated by camera noise and material irregularities (see Fig. 2 in supplemental); especially, by color aberrations introduced by Bayer demosaicing. We found that even pre-filtering for noise does not remedy these issues. As inverse rendering and colorimetric problems often use the raw camera recordings, these aberrations may lead to serious visual artifacts (see Fig. 1).

*Non-uniform Illumination.* Furthermore, the scene including the colorchart blocks may not be illuminated uniformly. Thus there is a non-linearity or even ambiguity added to the problem as a single given level of luminance could result in different luminance values in different spatial regions in the photo due to unequal illumination.

## 3 PROPOSED METRIC AND METHOD

*Simulations examining color differences based regression.* Several state-of-the-art color correction methods solely rely on minimizing color matching differences for a reference set. We use one such highly effective approach, namely, root polynomial regression (RPR) [Afifi et al. 2019; Finlayson et al. 2015; Kucuk et al. 2022] in simulated conditions to empirically establish the shortcomings of this metric. We simulate perfect imaging conditions for transforming a million spectral radiance profiles $\{s_i(\lambda)\}$ using: (a) Canon 650D's *RGB* spectral response functions (source data), and (b) CIE $XYZ$ color matching functions under D50 illumination (ground truth). We do the same for 24 blocks of Xrite's colorchart. Details for these simulations are presented in the supplemental material. Next, using the least-squares RPR fitting, we estimate the color transformation matrix and examine mapping errors. Table.1(first column-set) in the *supplemental pdf* shows error statistics for different orders of RPR where both regression and test sets are limited to colorchart blocks alone. In this case, $CIE\_\Delta E\_2000$ statistics fall significantly below the aspiring Just-Noticeable Difference $JND$ level ($\Delta E \leq 1$) as the regression order increases. However, for the million $\{s_i(\lambda)\}$ samples that are not used for regression, differences between corrected color values and corresponding ground truth values are comparatively high. Lastly, we regress over all of $\{s_i(\lambda)\}$ samples to estimate the color transformation matrix. Last column in the table from the supplemental shows that for all the orders of RPR, color differences statistics remain notably higher than $JND$.

*Proposed spatial metric.* To find an accurate solution, we need to estimate spectral radiances incident on given pixels accurately. We found that the major fallout of methods relying on minimizing three-channel color differences is that spatial anomalies are introduced. Fig. 1(center) illustrates these visual artifacts. We thus propose to incorporate a spatial variation metric, namely coefficient of variation (CoV) in the optimization criteria. We found that augmenting color differences with CoV differences allows for improvements.

*Proposed method.* We propose a novel method that: (i) treats luma and chroma information separately, (ii) uses whole reference block patches for robust statistics, (iii) estimates a whitepoint centric initial transformation matrix, and (iv) uses CoV statistics within reference patches to guide further optimization of the chrominance transformation matrix. See the supplemental Appx. B for details.

## 4 EVALUATION AND FUTURE-WORK

We first captured several photos of an X-rite colorchart under various illumination profiles in a laboratory setup. We then performed color correction using state-of-the-art regression method as well as our method. Table 1 presents our findings in error statistics for the photo in Fig. 1. As the order of regression increases $CIE\Delta E$ errors reduce but mean difference in CoV increases drastically. Visual artifacts due to these increases are dramatic and shown in Fig. 1 (center). We found our method to be comparable the second-order root-polynomial regression method in ΔCoV errors while it significantly improves on $\Delta E$ errors. We also took a photo from [Varghese et al. 2014] and color corrected it using various methods. Fig. 1 (right) shows similarly white-balanced results for qualitative inspection. Varghese et al. have reported that their $CIE\Delta E$ minimization method somehow results in unexplained darker shades. First-order regression results in overly saturated colors. Second-order regressions improves on color tones but has its white-point shifted towards a yellowish shade as seen for the ceiling. Our method avoids over-saturation, estimates overall brightness correctly and white-balances without any noticeable artifacts.

We found our method to produce promising results. In future, we plan to examine our method under large varieties of lighting conditions, color filtering and against public datasets. We trust our method to significantly benefit future inverse-rendering pipelines.

## ACKNOWLEDGMENTS

This work is fully supported by NSF Grant #2007974.

## REFERENCES

M. Afifi, B. Price, S. Cohen, and M.S. Brown. 2019. When color constancy goes wrong: Correcting improperly white-balanced images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 1535–1544.

G.D. Finlayson, M. Mackiewicz, and A. Hurlbert. 2015. Color correction using root-polynomial regression. *IEEE Trans. on Image Processing* 24, 5 (2015), 1460–1470.

A. Kucuk, G.D. Finlayson, R. Mantiuk, and M. Ashraf. 2022. Comparison of regression methods and neural networks for colour correction. (2022).

D. Varghese, R. Wanat, and R.K. Mantiuk. 2014. Colorimetric calibration of high dynamic range images with a ColorChecker chart. *Proceedings of the HDRi* (2014).