DOI: 10.1111/cgf.14959
Pacific Graphics 2023
R. Chaine, Z. Deng, and M. H. Kim
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 42 (2023), Number 7

Robust Distribution-Aware Color Correction for Single-Shot Images

Daljit Singh J. Dhillon[1] Parisha Joshi[1] Jessica Baron[2] Eric K. Patterson[2]

[1]Clemson University, Clemson, SC, USA [2]Clemson University, North Charleston, SC, USA

Figure 1: *A reference color chart (left image, bottom-left corner) is commonly used for color correction which is an ill-posed problem. State-of-the-art root-polynomial regression method reduces CIE ΔE_2000 color differences for the corrected reference blocks in the mean-sense. It improves significantly with increasing regression order as CIE Δ E_2000 drops with the increasing order. However, it completely ignores the spatial color variations or aberrations that produce serious artifacts as demonstrated here (right image). The proposed method improves color correction while preserving spatial variations and white-balancing accurately. It produces better tonal match with the reference ground truth (left image) as well. It is robust against strong illumination effects in the input photos and consistently outperforms reference state-of-the-art methods [FMH15, VWM14a].*

Abstract

Color correction for photographed images is an ill-posed problem. It is also a crucial initial step towards material acquisition for inverse rendering methods or pipelines. Several state-of-the-art methods rely on reducing color differences for imaged reference color chart blocks of known color values to devise or optimize their solution. In this paper, we first establish through simulations the limitation of this minimality criteria which in principle results in overfitting. Next, we study and propose a few spatial distribution measures to augment the evaluation criteria. Thereafter, we propose a novel patch-based, white-point centric approach that processes luminance and chrominance information separately to improve on the color matching task. We compare our method qualitatively with several state-of-the art methods using our augmented evaluation criteria along with quantitative examinations. Finally, we perform rigorous experiments and demonstrate results to clearly establish the benefits of our proposed method.

Keywords: *Color correction, material acquisition, inverse appearance modeling, inverse rendering*

CCS Concepts

• ***Computing methodologies*** → ***Image processing;*** *Image-based rendering; Reflectance modeling;*

1 Introduction

We often expect photographs of real-world objects to look 'real', i.e., to immerse us with the same visual perception that one may experience upon physically replacing the camera with ourselves. For instance, consumers buying goods online would expect the real cars, shoes, dresses, paints, artefacts, ornaments and other materials to look the same as their photographs relayed over sales portals. Amongst other factors, authentic reconstruction of the color information for real-world objects through camera captures remains an open and active area of research. Color correction is traditionally the pre-processing step that transforms the colors that a camera captures somewhat differently in comparison to the human visual system to recover their natural fidelity. Furthermore, it is also an essential pre-processing step in colorimetry and image-based material acquisition pipelines for inverse modeling and/or rendering. Such modeling and simulations often find applications in medical diagnoses, scientific

investigations and military operations where the accuracy of this task often becomes paramount.

Color correction of individual images with just three input color channels is essentially an ill-posed problem. We discuss this ill-posedness at length in Section 3 while examining the basic mathematical formulation for representing the color information as deduced from the incident spectral power distributions and the related practical limitation. There exists a substantial body of literature approaching this problem from many different angles (see Section 2 for details). Methods aiming at measurable levels of fidelity in color reconstruction mostly rely on color charts. A color chart has variedly colored blocks of known reflectances, and it is placed in a scene as a reference (left image in Figure 1). A color correction method then works to transform colors in the photograph such that the color chart blocks match their defined color values under some known illumination, as closely as possible.

Current state-of-the-art methods focus solely 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 [KFMA22]. Others rely on minimizing CIELAB ΔE errors in transforming reference blocks through classical optimization methods or machine learning methods. In this work, we first show empirically that predicating the optimality of the desired color-correction transformation on reducing such reference color differences may only lead to an overfitted solution even at its best.

We study the factors influencing the accuracy of the color correction task (Section 3) and devise practical strategies to counter each strong error contributing factor (Section 4). In particular, we examine the roots and role 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 simultaneously while reducing reference color differences in photos (Section 4). Based on our studies and experiments, we devise it to: (i) treat luminance and chrominance information separately, (ii) use whole reference block patches instead of just mean/median statistics, (iii) estimate a white-point centric initial transformation matrix, and (iv) use local spatial-variation statistics within reference patches to guide further optimization of the chrominance transformation matrix. We demonstrate that our method manages to do both: (a) stay within good tolerance limits for color differences for transformed reference blocks, and (b) retain local spatial variations (Section 5). In contrast, existing methods often overlook spatial variations while overfitting their solution to reduce color differences for reference blocks alone. Furthermore, unlike previous methods, we show our method to be robust against variations in the illumination conditions. We quantitatively as well as qualitatively evaluate our method including a few example images for a publicly available image dataset to establish comparative benefits. We also show our method to be resilient to different levels of noises in the imaging pipeline as well as robust to illumination variations and tonal biases.

2 Related Work

Substantial literature exists that deals with color correction and other transformational methods aimed at accurate material appearance capturing [Pal99]. This includes tone mapping, white balancing, recoloring, spectral reconstruction as well as computational hyperspectral imaging [Vrh93, KC14, TB05, GGVDW11, ABB22, SSFJ22, WSF*19, ABS16]. Given the scope of this work, we focus only on discussing state-of-the-art methods that can color correct individual photographs with just three color channels, using color charts specifically. Such methods can be classified broadly as *transformation matrix*, *nonlinear minimization*, or *machine learning* methods.

Transformation Matrix Approach. Simplest of the transformation matrix methods involve estimating a 3×3 matrix $\mathbf{L}$ using a least mean square formulation to match the source (imaged) colors to the target (reference) colors for the color chart blocks [KC14]. Often, a single mean or median color value per color chart block is used in such methods. Such methods are fundamentally targeted at using metrically linear and additive color spaces such as linear–RGB or XYZ color spaces. Conceptually, any such method is shifting, stretching and shearing the input source color space to best match a given target color space by minimizing the average Euclidean distance between the corrected color values and the reference color values for the blocks in the color charts. A somewhat nuanced extension of this approach involves piecewise linear transformation [MAF16]. In this adaptation, the source color space is partitioned into convex hulls using the chart block's mean colors as vertices in that space or some other criteria such as few fixed hue planes. This way, linear transformations can operate on nonlinear, non-additive color spaces such as HSV and CIELAB as well. However, space partitioning approaches exact a heavy computational cost by imposing a per-pixel search in determining which linear matrix must be used to transform it.

Instead of piecewise linear approximation of nonlinear color spaces, Hong et al. [HLR01] examine polynomial regression schemes for estimating the transformation matrix $\mathbf{P}$. This approach works better than the linear or piecewise linear regression approaches. However, polynomial regression is brightness dependent and it only works well for pixels with their luminance in the range expressed by the imaged set of color chart blocks. Finlayson et al. [FMH15] devised a root-polynomial regression method to make their color correction step independent of the pixel brightness. Root-polynomial method works well and it is the best state-of-the-art method operating with color charts that have an analytic form for directly estimating the transformation matrix [ABB22, APCB19b]. However, for orders higher than 2, it starts exhibiting visual artifacts and is recommended to be used with order 2 regression in a later study involving one of the original authors [KFMA22]. Furthermore, we found above methods not to be robust against varying illumination conditions (see Section 5 for details).

Nonlinear minimization. Varghese et al. [VWM14b] have developed one of the simplest and effective approaches beyond analytic formulations. They propose estimating the linear transformation matrix $\mathbf{L}$ through nonlinear minimization of CIELAB ΔE errors in reconstructing the reference color chart blocks instead of using a least mean square formulation in the source or target's (linear) color space. This method works well and improves color corrections subjectively. However, we found it to be sub-optimal in white-balancing images taken under illuminations with extreme spectral variations. Furthermore, estimating measurable quantities (relative radiance) through subjective color space (CIELAB) optimizations that do not

express one-to-one inverse relations seems counter-intuitive for applications where those measurements are put to some direct physical interpretations; for example, in medical diagnosis or other scientific applications of colorimetry. [CRA*21] provides greater insight into this issue.

Machine Learning Methods. For a long time, ML-based methods have been proposed to address color correction [Vrh03]. While the initial attempts had reasonable success under known illumination [ANU96], this illumination dependent re-training requirement is also the limitation of some of the recent basic Neural-Network based approaches, as noted by Kucuk et al. [KFMA22] while examining the method proposed by MacDonald and Mayer [MM21]. Recently, Afifi et al. [ABB22] propose a deep neural-network based method for auto white-balancing photos even without the need for color charts. Their approach is shown to produce excellent white-balancing for scenes under non-uniform illuminations. However, like many DNN based methods, this approach has notable memory footprint requirements.

Other Approaches. Few other methods involve spectral reconstruction [GGDG22, SSFJ22, WSF*19, TV91]. These methods are either specialized towards computational photography or require multiple images while changing illumination. Our goal is to enable reliable radiance estimation for inverse rendering pipelines or more serious colorimetric applications with practical or economical imaging setups. We thus focus only on methods that work for the generic imaging *modus operandi* that involves capturing individual images with or without any control over environmental lighting while still aiming for accurate color correction.

Afifi et al. [APCB19b] proposed a large image corpus search based method for white-balancing. Their approach is very robust in white-balancing, and we use some of their dataset images in our studies. We found our method to consistently color correct images generated under different tone mappings where all correspond to the same given white-balanced ground-truth image. Our method produces no noticeable color distortion in the white blocks for the color chart in the image (see Fig. 1).

Irrespective of the approach used, all the methods discussed above rely solely on reducing color differences in the mean-sense to estimate their transformation matrix or numerically tuning in a transformer model for color correction.This optimality criteria has serious, inherent limitations and we elaborate on them in the following section while also proposing a meaningful augmentation to it. This is followed by our proposed method which improves the current state-of-the-art approaches for accurate and robust color corrections.

3 Color Correction Problem

Ill-posedness. To understand the underlying ambiguity and ill-posedness of the color correction problem, let us begin by considering the mathematics of pixel coloration. Continuous spectral radiance $s(\lambda)$ is integrated under three different spectral response (color) functions of the camera, say $\mathbf{f}_c(\lambda) = [r(\lambda),\ g(\lambda),\ b(\lambda)]^T$. Registered pixel color is then defined as:

$$\mathbf{c} = [r_c, g_c, b_c]^T = \mathcal{T}_c[s(\lambda)] = \int s(\lambda)\mathbf{f}_c(\lambda)d\lambda. \qquad (1)$$

Here $\mathcal{T}_c$ represents the transformation from the spectral space to a linear color space, say linear–RGB. This transformation $\mathcal{T}_c$ is linear in the sense that $\mathcal{T}_c[a\cdot s_1(\lambda) + b\cdot s_2(\lambda)] = a\mathcal{T}_c[s_1(\lambda)] + b\mathcal{T}_c[s_2(\lambda)]$

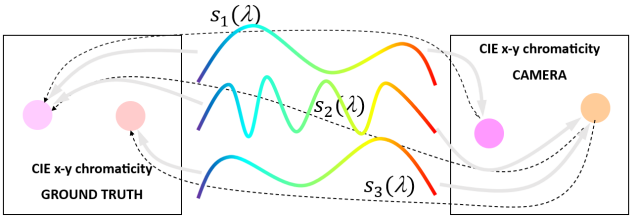

Figure 2: *Ambiguity of the many-to-many problem is depicted using chromaticity plots. Color correction requires solving this ill-posed problem along with luminance corrections.*

for any two arbitrary scalars a, b and any two arbitrary spectral radiances $s_1(\lambda)$ and $s_2(\lambda)$. Here we note that the **R**, **G** and **B** dimensions for a given camera's color space are represented as three (x,y) points in CIE's $x-y$ chromaticity space and these three points vary depending on the camera's spectral response (color) functions. Figure 2 depicts two such cameras with two different $\mathbf{f}_c$ functions. $s_1(\lambda)$, $s_2(\lambda)$ and $s_3(\lambda)$ are three spectral radiance profiles incident on three different camera pixels. Both s_1 and s_2 map to the same RGB color value and thus the same (x,y) chromaticity position (on the left in Figure 2) in CIE color space, due to Camera_1 specific metamerism. Similarly, s_2 and s_3 map to the same chromaticity position in CIE color space due to Camera_2's metamerism. This example illustrates two one-to-many mapping instances that together result in a many-to-many mapping. Thus, an accurate mapping between two different RGB color spaces 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. Now if one of those cameras emulates CIE's color functions, then the mapping represents the color correction problem. Even when the reference blocks map to distinct points in the chromaticity space, it does not preclude perceptually metameric spectral radiances to map to different points in the x-y chromaticity space. Thus color correction is inherently ambiguous and ill-posed, in general.

To understand camera specific metamerism, we can quickly visit the discrete formulation which is often used for principal component analysis (*PCA*) based methods for color correction. For simplicity, often times the integration expressed under transformation $\mathcal{T}_c$ is approximated with discretized counterparts $\mathbf{s}_{n\times 1} \equiv s(\lambda)$ and $\mathbf{F}_{3\times n} = [\mathbf{r},\mathbf{g},\mathbf{b}]^T \equiv \mathbf{f}_c(\lambda)$ and expressed using linear algebra as $\mathbf{c} = \mathcal{T}[\mathbf{s}] = \mathbf{F}\mathbf{s}$. Here, n is the number of discrete spectral bands used to approximate $s(\lambda)$. Clearly, $\mathbf{F}$ has $n-3$ dimensional null-space which results in projection losses, or in other words metamerism.

Now, let $\mathbf{F}_{\text{cie}}$ represent discretized CIE color functions. So, $\mathbf{c}_{\text{cie}} = (\mathbf{F}_{\text{cie}}\mathbf{s})$ is a perceptually correct mapping for $\mathbf{s}$ under chosen discretization. With $\mathbf{F}_{\text{cie}}$ having a different null-space than $\mathbf{F}$, a naive linear transformation $\mathbf{T}_{3\times 3} = \mathbf{F}_{\text{cie}}(\mathbf{F}^T\mathbf{F})^{-1}_{n\times n}\mathbf{F}^T$, to map $\mathbf{c}_{\text{cie}} = \mathbf{T}\mathbf{c}$ does not exist when $n > 3$ as $(\mathbf{F}^T\mathbf{F})$ becomes rank deficient and cannot be inverted. In such scenarios, we solve for $\mathbf{T}$ in a least mean square sense by using an adequate ($\geq n$) number of $< \mathbf{c}_{\text{cie}}, \mathbf{c} >$ tuples where each pair is known to result from same spectral profile $\mathbf{s}$ under two projections $<\mathbf{F}_{\text{cie}}, \mathbf{F}>$. In practice, we use color charts to give us those tuple sets. Since the set of such tuples is not unique, $\mathbf{T}$ does not have a unique solution. In other-words, we can have two different color charts where their blocks differ in their spectral reflectances and they will result in different 'best' solutions for $\mathbf{T}$. This is the mathematical way of saying that color correction with

just three color channel information per pixel is inherently an ambiguous and ill-posed a problem. Also, it mathematically implies that by just minimizing color differences with the references, we cannot resolve the color correction problem. At best, we can find the best solution under given constraints. In this paper, we propose to extend the constraint set to include additional spatial statistics with the justifications as discussed in the following.

Camera Noises and Color Chart Inaccuracies: The mapping problem is further complicated by camera noise and material irregularities, even of brand new color charts whose irregularities only increase with degradation over time (see Figure 3). Especially, color aberrations introduced by Bayer demosaicing pose a serious challenge. We can expect that even pre-filtering these aberrations do not remedy inaccuracies due to the color correction problem's inherent ambiguity. As inverse rendering and colorimetric problems often use the raw camera recordings, these aberrations lead to serious visual artifacts if not handled well (see Fig. 1 (right image)).

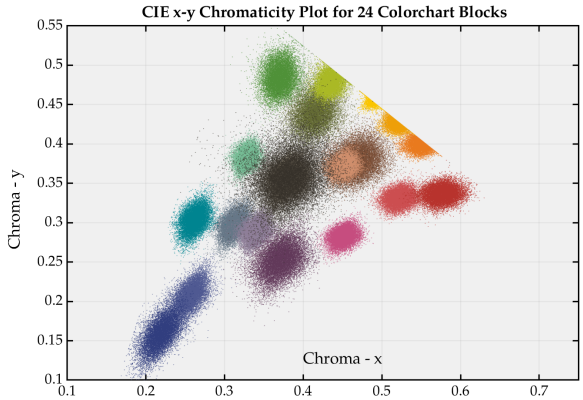

Figure 3: *Xrite color chart block patches from an example photo taken in our dark room laboratory settings are plotted pixelwise in CIE x-y chromaticity space. Raw pixel colors are assumed to be in CIE RGB space.*

Now, let us begin by assuming that each pixel is an independent variable. This assumption is immediately challenged by the existence of Bayer filter that imposes local biases. Since such biases are spatially imposed, spectral radiance incident on an individual pixel expresses an imprint over its neighborhood. This justifies the use of spatial statistics to devise additional constraints on solving any transformation matrix **T**. Similarly, other degrading factors discussed above impose such spatial biases that must remain intact under a reliable and robust color correction process.

Based upon above understanding of the color correction problem and further empirical studies, we propose a novel evaluation criteria that incorporates spatial statistics in estimating a transformation matrix **T**. We also propose a new method that is devised to allow for optimizing on this new evaluation criteria systematically instead of just a brute-force minimization. Lastly, we state that all the discussions in this section about the linear 3×3 transformation matrix **T** also hold true for estimating any higher-order regression based transformation matrix. Next, we present our proposed method.

4 Proposed Method

Based on our studies, we propose a novel method that: (step i) treats luminance and chrominance information separately, (step ii) uses whole reference block patches instead of just mean/median statistics, (step iii) estimates a white-point centric initial transformation matrix, and (step iv) uses local spatial variation statistics within reference patches to guide further optimization of the chrominance transformation matrix. To understand the relevance and justifications driving these design choices, let us begin by considering the limitations of the commonly relied upon error metric, i.e. CIE ΔE in CIELAB color space. While we always employ CIE ΔE_2000 throughout our research, where the context warrants generality, we refer to it as CIE ΔE or just ΔE for the sake of brevity.

4.1 CIE ΔE Limitation and Proposed Improvement

Simulations examining ΔE reliability in color matching. Several state-of-the-art color correction methods solely rely on minimizing color differences between the transformed and the expected color values for a reference set. We use one such highly effective approach, namely, root polynomial regression [FMH15, APCB19b, KFMA22] in simulated conditions to empirically establish the shortcomings of this metric. We simulate perfect imaging conditions for transforming a million spectral radiance $s_i(\lambda)$ profiles using: (a) Canon 600D's *RGB* spectral response functions (giving source data) [Sob19], and (b) CIE *XYZ* color matching functions under D50 illumination (giving ground truth). We do the same for 24 blocks of Xrite's color chart [Bab23]. Source code for these simulations are included in the supplemental material. Next, using the least-squares root-polynomial regression fitting, we estimate the color transformation matrix and evaluate errors arising from its use. Table 1(first column-set) shows error statistics for different orders of root-polynomial terms where both regression and test sets are limited to color chart blocks alone. In this scenario, CIE ΔE_2000 color difference statistics fall significantly below the aspiring Just-Noticeable Difference *JND* level ($\Delta E \leq 1$) with the increasing regression order. However, for the million spectral radiance samples not used in the regression process, differences between corrected color values and corresponding ground-truth values are comparatively high. Lastly, we regress over all of the million samples to estimate the color transformation matrix. The last column in Table 1 shows that for all the orders of root-polynomial terms ΔE statistics remain notably higher than *JND*. This clearly points towards data overfitting by the transformation matrix **T** (described in Section 3). This overfitting can lead to noticeable visual artifacts such an increase in spatial chromatic aberrations.

Proposed spatial metric. To find an accurate solution, we need to estimate the spectral power distribution of the radiance incident on a given pixel accurately. This is not doable with just three color channels. At the same time, we noted that the major fallout of methods relying on minimizing color differences is that spatial artifacts or anomalies are introduced. Also, in Section 3 we mathematically examined why using color differences alone cannot resolve the ambiguity of the color correction problem. We also discussed the justification to use some measure of spatial variations to help preserve them. Zhang and Wandell [ZW97] with their S–CIELAB metric suggest multi-level, lowpass pre-filtering of images to incorporate spatial variations in ΔE measures [JF03]. Johnson and Fairchild [JF01] further improve the filters used in S–CIELAB metric while Wang and Hardeberg adapt S–CIELAB to employ six channels of information (lightness, chroma, hue, compression, noise and

CIE ΔE_2000 Error Measures

Regression Set →	24 Color chart blocks				24 Color chart blocks				Million random $\{s_i(\lambda) : i = 1 \cdots 10^6\}$			
Testing Set →	24 Color chart blocks				Million random $\{s_i(\lambda) : i = 1 \cdots 10^6\}$				Million random $\{s_i(\lambda) : i = 1 \cdots 10^6\}$			
Root poly. order ↓	mean	median	95%	max	mean	median	95%	max	mean	median	95%	max
one	0.9030	0.7114	1.7194	2.7205	2.3761	2.1606	4.7560	12.4777	2.2546	2.0203	4.6439	13.2740
two	0.5385	0.4101	1.1666	1.6317	2.3372	2.1170	4.7104	12.8163	2.2533	2.0182	4.6442	13.3730
three	0.3672	0.2482	1.0046	1.0946	2.3868	2.1645	4.8065	12.9689	2.2532	2.0181	4.6443	13.3673
four	0.0121	3.1823e-05	0.0604	0.1405	5.5213	4.2903	14.4383	77.3033	2.2532	2.0180	4.6448	13.3722

Table 1: *Examining the effectiveness of reducing CIE ΔE errors in optimizing the color correction matrix. Regression results against millions of random spectral radiance profiles indicates that using only few color chart blocks as references results in overfitting. These experiments also show that using a very large gamut results in mean ΔE around 2.25 units.*

sharpness) subject to bilateral filtering [WH09, WH12]. Similarly, Simone et al. [SOF09a] incorporate the use of log-compressed OSA-UCS space [SOF09b] to adapt S–CIELAB into a new metric called S–DEE. Also, Pedersen and Hardeberg [PH09] incorporate spatial properties of the human visual system in extending a hue-angle algorithm [HL02] to devise a new S–CIELAB like metric called SHAME. All these metrics are particularly focused at comparing local image features such as edges, sharpness and color flow. We, on the other hand, want to study (non-local) mass distribution properties over reasonably sized patches with rather flatness in expected variations. More importantly, we do not want to penalize differences in spatial variations on an absolute scale (say standard deviation σ for a patch). Instead, we want to preserve spatial variations under some scale normalization. We thus propose to incorporate a spatial variation metric which is the ratio of the standard deviation to the mean for a patch to the optimization criteria. It is called coefficient of variation or CoV. We found that augmenting color differences with CoV differences allows for improvements. In Section 5 we establish how to relatively weight ΔCoV against ΔE errors to estimate net error measures that we aim to minimize.

4.2 Key Contributive Ideas

Before detailing the algorithm for our method, we first present the details and justifications for our four key contributive steps.

Step (i): Processing chrominance and luminance data separately. Generally, we are motivated to impose higher-order regressions to better match given reference color data by accounting for the nonlinearities warranted in corrections. However, based on our experiments, we found that notable spatial artifacts are introduced with higher-order root-polynomial regression, which is mainly due to the amplification of local variations in the luminance channel for the input photo. We thus firstly partition the regression step to separately process the luminance and chrominance data. This enables us to better control and limit the visual artifacts by applying a low-order root-polynomial regression method to correct the luminance channel information when appropriate. For instance, non-uniform illumination conditions over the physical color chart can lead to greater artifacts due to overfitting on the luminance data by higher-order regression. To process luminance separately, we first estimate robust mean color vectors as explained further in Step (ii). All 24 robust mean color vectors are then used in a regression fitting step to correct the luminance channel as detailed in Algorithm 1 in supplemental.

Also, color correction is essentially a problem about fixing the $x-y$ chromaticities. We thus directly work in CIE $x-y$ chromaticity space to find the chroma transformation matrix with higher-order regression. By processing chromaticity channels separately, we also avoid the potential compromise on chromatic corrections that result from overfitting the out of balance luminance data that may have been degraded due to noise or illumination non-uniformities. We note that a linear transformation in CIE XYZ space results in a homographic transformation in $x-y$ space. We thus work with homogeneous chroma coordinates. By extension, we use homogeneous coordinates for nonlinear regression as well.

Step (ii): Using patch-based robust statistics. For processing luminance data, we compute robust mean color vectors in CIE XYZ color space for 24 color chart block patches in the input photo. For this we use only the inliers in each patch to compute its respective mean color vector. Also, due to camera noises, demosaicing errors or color chart anomalies in the imaged $x-y$ chroma data exhibit a high amount of color aberration. This shows up as a serious overlap between samples from different color blocks of the color chart (see Fig. 3). We found that instead of using mean/median chroma for each block, using all the additional samples for each reference color block to match a single expected target $x-y$ pair allows for influencing the chroma transformation matrix by individual patch distributions. Our approach in effect works to retain the local spatial variational distributions. Each block patch is additionally used as a reference to quantify relative changes in spatial variations due to color correction. We compute patchwise ΔCoV for explicit minimization as explained in Step (iv) in the following.

Step (iii): White-point centric correction. To ensure greater efficacy in white-balancing the results from color correction, it is recommended to employ constrained regression [FD97]. Finlayson et al. [FMH15] also conclude that using this constrained regression approach can help to preserve the white-point while correcting colors in a luminance independent manner. However, we observe a subtle luminance-dependent shift in the *corrected* white point in the $x-y$ space even after using one of the reference color chart white blocks to constrain root-polynomial regressions. Section 5.2 elaborates on this anomaly with an illustration in Fig. 7. These subtle white point drifts result from spatial chromatic aberrations that occur due to noisy imaging. The aberrations cause marginally different impact on chroma estimates at different brightness levels owing to differences in relative signal-to-noise ratios between these levels. To draw an analogy, the influence of the scale of the noise on estimating a camera's spectral sensitivities is studied and tabulated by Jiang et al. in their research work (Section 5.2) [JLGS13].

In effect, differences in the impact of noise at different brightness levels lead to minor yet systematic shears in chromaticity estimates. Such shearing effects are not amenable to correction through constrained regression for white point preservation, as demonstrated in Section 5.2. Furthermore, we aim to make our color correction algorithm to be robust against shadows on the color chart and similar non-uniformities in the illumination intensity. We thus explicitly apply a white-point centric transformation approach where imaged white blocks from the color chart act as representative 'white points' at their recorded luminance levels respectively. This allows for approximating a corrective shear for the luminance Y axis in $x-y-Y$ space, prior to estimating chroma transformation in the $x-y$ space alone. Using the white point $x-y$ as the origin in the homogeneous chroma coordinate space allows for an easier, deterministic control for white-point matching. It also results in chroma-correction errors to distribute evenly around the white point. Sub-steps for our white-point centric corrections are presented in Algorithm 1 in the supplemental.

Step (iv): Explicit optimization for preserving spatial variations. Lastly, to fine tune our chroma transformation matrix, we directly employ patch-based coefficient of variance measures, i.e. ΔCoV. It augments and regularizes our objective function that also includes CIE ΔE_2000 color differences for reference blocks. However, for computing color differences (say ΔE_{chroma}) in this step, we first equalize the luminance Y channels in the CIE $x-y-Y$ color space for the color vectors being compared. In other words, we assign the luminance value of the reference color vector $\mathbf{c_{cie}}$ to the luminance value for the corresponding color corrected color vector $\mathbf{c_r}$ to give us a new color vector $\bar{\mathbf{c}}_\mathbf{r}$. Note that all three color vectors $\mathbf{c_{cie}}$, $\mathbf{c_r}$ and $\bar{\mathbf{c}}_\mathbf{r}$ are in CIE $x-y-Y$ color space. We then compute ΔE_{chroma} as the CIE ΔE_2000 between $\mathbf{c_{cie}}$ and $\bar{\mathbf{c}}_\mathbf{r}$. The minimization problem is then defined as:

$$\arg\min_{\mathbf{T_C}} \sum_i \|\Delta E_{\text{chroma}}(\mathbf{c}^i_{\mathbf{cie}}, \mathbf{c}^i_{\mathbf{r}})\|_2^2 + \sum_j \|\alpha\lambda\Delta\text{CoV}_j\|_2^2, \quad \textit{where} \quad (2)$$

$$\mathbf{c}^i_{\mathbf{r}} = < \mathbf{T_C} \cdot \mathbf{c}^i_{\mathbf{rp}} >, \quad \textit{and} \quad (3)$$

$$\Delta E_{\text{chroma}}(\mathbf{c}^i_{\mathbf{cie}}, \mathbf{c}^i_{\mathbf{r}}) = \Delta E(\mathbf{c}^i_{\mathbf{cie}}, \bar{\mathbf{c}}^i_{\mathbf{r}}) \quad (4)$$

Here, $\mathbf{c}^i_{\mathbf{rp}}$ is the root-polynomial vector in homogeneous $x-y-1$ coordinates, $\mathbf{T_C}$ is the chroma transformation matrix, $<\cdot>$ indicates transformation into regular $x-y$ coordinates, α is the number of inliers to be used in each patch and λ is set as discussed in Sec. 5.1. 'i' identifies inlier pixels and 'j' identifies individual patches. ΔCoV only considers $x-y$ channels.

4.3 Our Algorithm

For clarity, we provide a step-by-step complete pseudo-code for our method in the supplemental material (See Algorithm 1 in it). Here, we include pseudo-code for two core parts of our method: (a) estimating the luminance transformation matrix $\mathbf{T_L}$ (see Algorithm 1) and (b) estimating the chrominance transformation matrix $\mathbf{T_C}$ (see Algorithm 2). Next, we emphasize some of the important technical details that warrant further clarification.

- To compute root-polynomial vectors we use the following formula and logic. Consider a pixel i, with its photographed linear color vector $\mathbf{p}_i = [r,g,b]^T$. Its root-polynomial vector $\mathbf{r}_i$ of order n enumerates all elements of the set:

$$S_n = \{(r^\alpha g^\beta b^\gamma)^{1/t} : t = (\alpha+\beta+\gamma) \leq n\}. \quad (5)$$

When there are less than 3 elements in the color vector, above equation is adapted accordingly. When the color vector for a pixel is in homogeneous coordinates, we do not compute its root-polynomial color vector directly. Instead, we first compute the root polynomial vector for the equivalent non-homogeneous color vector and augment it with the homogeneity coordinate of value 1. This way, we ensure that the scale independent operations devised by Finlayson et al. [FMH15] are applied correctly.
- When correcting luminance, ideally we would want to regress over Y channel alone. However, we assume that the captured photo is in linear RGB color space while converting it into CIE XYZ space. Now, the camera's spectral sensitivity functions differ from CIE prescribed RGB color functions. However, both X and Y channels that we compute from the input photo are expected to have high correlation with the true image luminance. We thus use both of these channels in regressive corrections.
- When correcting luminance with $r_{\text{luma}} = 1$ in Algorithm 1, the inlier patch statistics produce the same result as with using the robust mean color vectors alone. For higher-orders, all inlier statistics as used in Alg. 1 to provide even better resilience against non-uniform illumination conditions.
- White-point centric corrections result in negative chromatic coordinates that involve complex number operations in root-polynomial based regression. We estimate final corrections based on the magnitude of these complex number calculations.
- For the patch-based operations we used inliers for robust statistics for ΔE computations. However, we estimate ΔCoV errors using entire patches.

Algorithm 1 Estimate Luminance Transformation Matrix $\mathbf{T_L}$

Input:
- $\mathbf{S}_{n\times m\times 3}$: An array of photo pixel colors for 'n' chart blocks, each with 'm' inliers and expressed in 3 color channels, namely XYZ
- $\mathbf{T}_{n\times 3}$: An array of target reference pixel colors for 'n' chart blocks that are expressed in 3 color channels, namely XYZ
- r_{luma}: Regression order for processing luminance data

Output:
$\mathbf{T_L}$: Luminance Transformation Matrix

function COMPUTELUMATRANSFORMATION($\mathbf{S}, \mathbf{T}, r_{\text{luma}}$)
  $p \leftarrow 0$
  for $i = 1$ to n **do**
    for $j = 1$ to m **do**
      $p \leftarrow p+1$
      $\mathbf{R}_{\text{matrix}}(p,:) \leftarrow$ GiveRootPolyVector($\mathbf{S}(i,j,1:2), r_{\text{luma}}$) ▷ See Eqn. 5, {using only first two color channels here}
      $\mathbf{t}_{\text{vector}}(p) \leftarrow \mathbf{T}(i,2)$ ▷ Reference Luminance
    end for
  end for
  $\mathbf{T_L} \leftarrow \arg\min_{\mathbf{T}^*_\mathbf{L}} \|\mathbf{t}_{\text{vector}} - \mathbf{R}_{\text{matrix}}\mathbf{T}^*_\mathbf{L}(:)\|_2^2$ ▷ linsolve() in MATLAB
  return $\mathbf{T_L}$
end function

5 Experiments and Results

We performed both quantitative as well as qualitative evaluations

Algorithm 2 Estimate Chrominance Transformation Matrix $\mathbf{T_C}$

Input:

- $\mathbf{S}_{n\times m\times 2}$: An array of photo pixel colors for 'n' chart blocks, each with 'm' inliers and expressed in 2 color channels for CIE chromaticity x, y. Note that pixel colors could have undergone a shift in their reference white points as explained in Step (iii) for Section 4.2
- $\mathbf{T}_{n\times 2}$: An array of target reference pixel colors for 'n' chart blocks that are expressed in 2 color channels representing CIE chromaticity x, y. Note that target reference colors could have undergone a shift in their reference white points as explained in Step (iii) for Section 4.2
- r_c: Regression order for processing chrominance data
- o_{flag}: Flag indicating that Step (iv) in Sec. 4.2 must be applied.

Output:

$\mathbf{T_C}$: Chrominance Transformation Matrix

function COMPUTECHROMATRANSFORMATION($\mathbf{S}, \mathbf{T}, r_c$)
  $p \leftarrow 0$
  for $i = 1$ to n **do**
    for $j = 1$ to m **do**
      $p \leftarrow p + 1$
      $\mathbf{R}(p,:) \leftarrow$ GiveRootPolyVector($\mathbf{S}(i,j,:), r_c$) ▷ See Eqn. 5
      $\mathbf{t1}(p) \leftarrow \mathbf{T}(i,1)$; $\mathbf{t2}(p) \leftarrow \mathbf{T}(i,2)$
    end for
  end for
▷ Make adjustments to deal with the 'homogeneous' chroma coordinates
  for $p1 = 1$ to p **do**
    $\mathbf{D1}(p1,p1) \leftarrow \mathbf{t1}(p1)$; $\mathbf{D2}(p1,p1) \leftarrow \mathbf{t2}(p1)$
  end for
  $\mathbf{A}_{\text{matrix}} \leftarrow \begin{bmatrix} \mathbf{R} & 0\times\mathbf{R} & -\mathbf{D1}\times\mathbf{R} \\ 0\times\mathbf{R} & \mathbf{R} & -\mathbf{D2}\times\mathbf{R} \end{bmatrix}$; $\mathbf{t}_{\text{vector}} \leftarrow \begin{bmatrix} \mathbf{t1} \\ \mathbf{t2} \end{bmatrix}$
  $\mathbf{T_C} \leftarrow \arg\min_{\mathbf{T_C^*}} \|\mathbf{t}_{\text{vector}} - \mathbf{A}_{\text{matrix}}\mathbf{T_C^*}(:)\|_2^2$ ▷ linsolve() in MATLAB
  $\mathbf{T_C} \leftarrow \begin{bmatrix} \mathbf{T_C} & \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T \end{bmatrix}$ ▷ For Homogeneous Coordinates
  if o_{flag} is true **then**
    $\mathbf{T_C} \leftarrow$ Eqn. 2▷ Perform nonlinear minimization starting with the initial value of $\mathbf{T_C}$ from the previous step. See lsqnonlin() in MATLAB
  end if
  return $\mathbf{T_C}$
end function

for our method. We implemented our method in *MATLAB*. When opening raw photograph files, we use *DCRAW* with default options. For photographs in other color spaces, we require and use additional metadata to convert them to linear–RGB first. Unless otherwise specified, we set the regression order for luminance correction to 2. Increasing it further produces spatial artifacts. Using linear regression does not help minor nonlinearities that might creep in due to camera's properties or variations in the environmental lighting.

5.1 Quantitative Evaluations

We perform several quantitative examinations of our method. We use *CanonDarkroom_*01.*CR*2 from Figure 4 as the raw input image for these evaluations. It has a resolution of 4022×6024 pixels.

Compare with state-of-the-art methods. Firstly we compared our method with both reference methods, namely Root-Polynomial (**RP**) and Varghese et al. (**VEM**). We applied **RP** method with linear to fourth-order of regression. For our method, we used 1% patch data for robust mean estimates, second-order regression for both luma and chroma channels, kept white point shearing on and optimized transformation matrix using nonlinear minimization as the last step. Table 2 lists both color differences as well

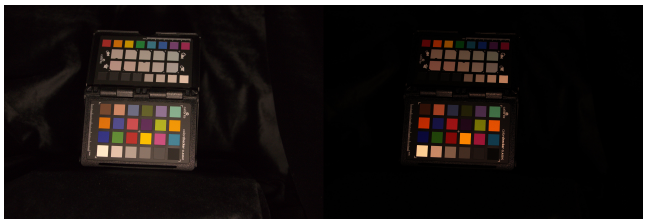

Figure 4: *CanonDarkroom_*01.*CR*2 *in sRGB and linear-RGB (right).*

as patch-based coefficient of variation errors for all the methods. Also, Figure 5 shows distribution of these errors across different chart blocks for few method configurations. Looking at the table and these images, it seems that **RP** method with regression order 4 does the best in reducing the color differences. Unfortunately, this is a misleading observation as the reduction in ΔE errors for mean color values is achieved through overfitting and at the cost of introducing visual artifacts. Subjective evaluation of these artifacts is done in the following subsection. In a quantitative sense, those artifacts reflect as high numbers in mean absolute Δ*CoV* (last three columns in Table 2). This table also shows that at comparable levels of *CoV* errors (order 2 for **RP** method), our method does better in reducing mean ΔE error. Similar to this case, we found our method did better than other methods in reducing the net error in general as well.

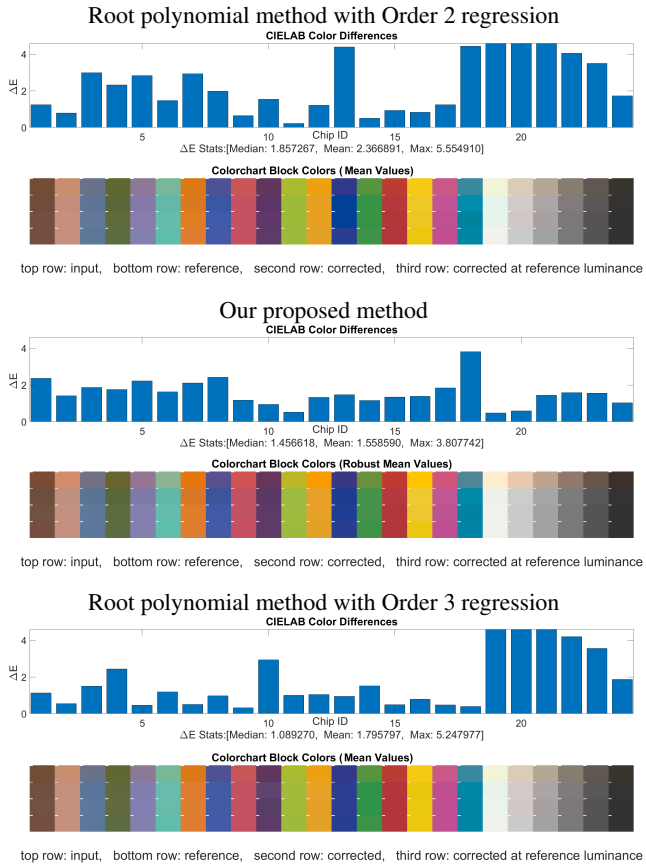

Figure 5: *Barcharts showing CIE* Δ*E_*2000 *errors for different color blocks for all the examined methods in different configurations.*

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
Varghese et al.	1.95	1.52	8.53	0.42	0.34	0.11
Our (L-2,C-2)	1.56	1.46	3.81	0.27	0.11	0.13

Table 2: *Quantitative comparison of different methods. Color differences as well as differences in the patch-level coefficient of variance are reported here. Proposed method shows lower mean Δ than other two methods at comparable level of variance preservation.*

Chroma Order [**R**]	Whitepoint sheared [**S**]	Patch Used [**P**]	Step (iv) Used [**O**]	CIE ΔE_2000 Error			100×Mean		|ΔCoV|
				Mean	Median	Max	x	y	Y
2	×	×	×	2.028	1.756	8.602	0.542	0.440	0.188
2	×	×	✓	1.969	1.777	8.254	0.409	0.467	0.188
2	×	✓	×	2.032	1.824	8.603	0.525	0.432	0.187
2	×	✓	✓	2.886	2.758	6.944	0.026	0.011	0.187
2	✓	×	×	1.425	1.435	3.303	0.503	0.427	0.188
2	✓	×	✓	1.429	1.403	3.298	0.481	0.365	0.188
2	✓	✓	×	1.418	1.463	3.430	0.473	0.407	0.187
2	✓	✓	✓	2.475	2.515	5.049	0.050	0.039	0.187
3	×	×	×	2.013	1.767	8.681	0.555	0.617	0.188
3	×	×	✓	2.063	1.837	8.279	0.496	0.563	0.188
3	×	✓	×	2.018	1.766	8.653	0.526	0.585	0.187
3	×	✓	✓	3.363	3.295	6.842	0.159	0.300	0.187
3	✓	×	×	1.412	1.369	3.246	0.510	0.520	0.188
3	✓	×	✓	1.373	1.361	3.138	0.474	0.375	0.188
3	✓	✓	×	1.410	1.385	3.385	0.475	0.510	0.187
3	✓	✓	✓	2.480	2.622	4.597	0.057	0.045	0.187

Table 3: *Error statistics for various configurations of the proposed method. Using second order regression for chromatic correction along with all other options activated gives best results in general. Patch statistics and white-point centric must almost always be applied. Final optimization step reduces spatial color aberration errors considerably only when patch statistics are used.*

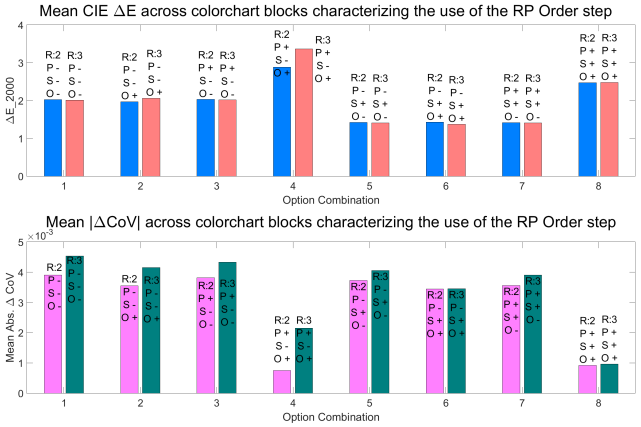

Figure 6: *Examining the effects of varying the root polynomial order for chromatic corrections under different configurations*

Examining constrained root polynomial regressions. We incorporated constrained least-square regressions [FD97] in the **RP** method by Finlayson et al. [FMH15]. Figure 7 illustrates our findings. Case C1 uses the brightest white block as a constraint. It results in brownish shade for the white blocks with lower luminance. C2 and C3 use white blocks with lower reflectances but they result

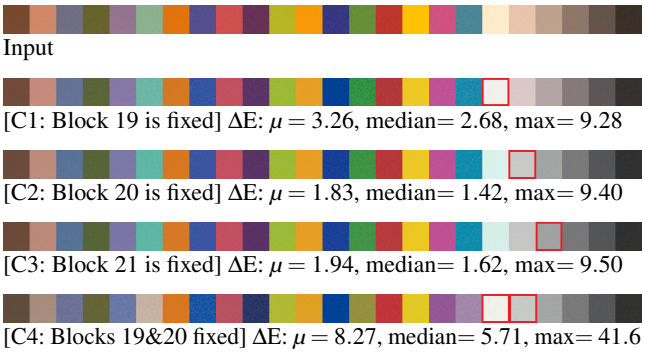

Figure 7: *Examining the impact of constrained regressions on RP method by Finlayson et al. [FMH15]. Patches in red boxes above were used to enforce color equality constrains on their robust mean color vectors. Second-order regressions were applied to all cases.*

in bluish shift in the white point at higher luminance. Using more than one constrain (C4) produces serious artifacts in other color blocks. ΔE statistics included in the figure reinforce these findings quantitatively. This case study highlights the challenges faced by the **RP** method in white point preservation using constrains alone.

Examining the impact of different configurable modules. We experimented with all combinations of all four configurable parameters for our method. Table 3 shows error statistics for it. We first look at the impact of the regression order for chromatic correction. Using 3^{rd} order regression does not improve much over 2^{nd} order results in reducing ΔE errors. Also, it mostly resulted in slightly higher *CoV* errors in comparison to 2^{nd} order regression (see Fig. 6). This reduction in higher-order effectiveness is mainly due to separate processing of chromatic and luminance information. Higher-order regression appear to achieve reductions in ΔE errors by overfitting to the luminance data. In summary, we find our method to work better with 2^{nd} order regression. In comparison to **RP** method, we found our method to have less severe *CoV* degradations at higher-order regression, but we leave those results out for the sake of brevity.

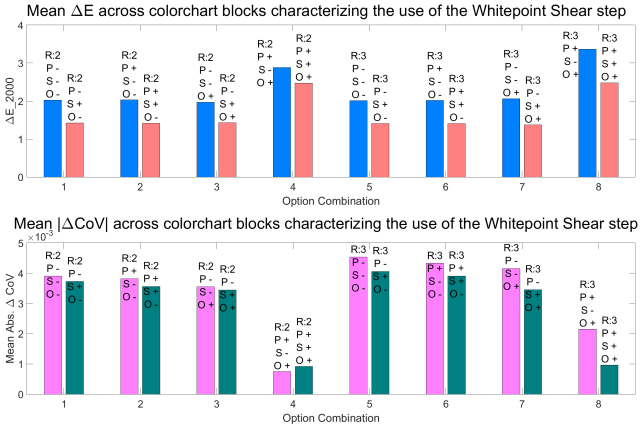

Figure 8: *Examining the effects of applying luminance dependent white shifts for chromatic correction under different configurations*

Our next important step of white-point centric corrections shows reductions in mean ΔE error in every possible configuration. Figure 8 visually depicts this trend. Furthermore, it almost always reduces

CoV errors except when patch statistics and nonlinear optimization, both, are disabled. Even in that case, the degradation is marginal.

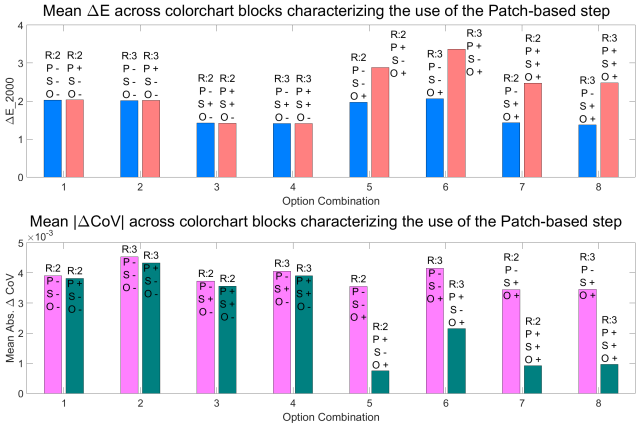

Figure 9: *Examining the effects of using patch statistics for color correction under different configurations*

We examined the impact of using 5% of patch inliers for robust mean estimation on our method's performance. Figure 9 illustrates its differentiable effects. Using patch statistics does not impact ΔE errors unless nonlinear optimizations are used. In that case, ΔE errors increase notably but not without significant benefits in reducing *CoV* errors. On overall basis, using patch statistics result in net benefits. Finally, using last step minimizations produce the obvious effect of reducing *CoV* errors and increasing ΔE. But as per its objective function's design, these optimizations always result in net error reductions.

In summary, using second-order regression for chromatic correction along with all other options activated gives best results in general. Patch statistics and white-point centric must almost always be applied. Final optimization step reduces spatial color aberration errors considerably only when patch statistics are used.

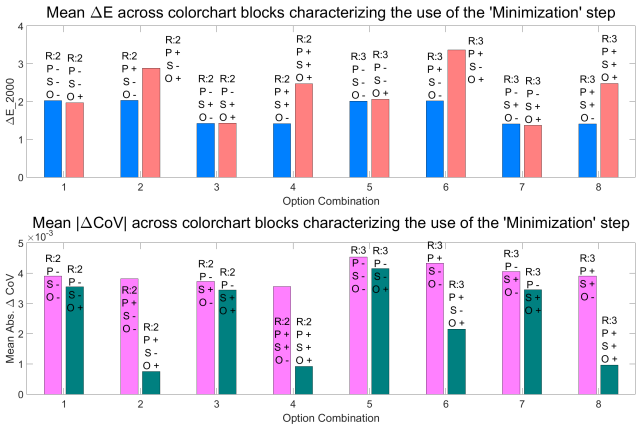

Figure 10: *Examining the effects of applying nonlinear minimization on weighted cumulative errors under different configurations*

Impact of relative patch size. Using patch statistics plays an important role in making our method robust. Instead of patch median colors, we chose to use robust mean color values estimated over a limited set of pixels closest to the patch median. While the median color values help ward off the outlier's influences, they do not take into account the influence of non-uniform illumination or systematic degradations in the color chart blocks. We found that using robust means makes our method robust against such degrading adversarial scenarios. However, establishing the threshold for inclusion of patch pixels contributing towards mean estimation is a non-trivial task. We took an empirical approach to resolve this criteria. We tried different inlier percentages ($\mathbf{C}_{in}$ in Algorithm 1 in the supplemental) and studied their impact on net color matching errors. For these studies, we first enabled white-point centric operations and post nonlinear minimizations while setting the chroma regression order to 2. To begin with, we set the weight factor $\lambda = 50$ for combining per pixel ΔE and *CoV* errors as $E_{\text{total}} = \Delta\text{E} + \lambda\Delta CoV$ for the minimization step. With this we color corrected our input photo for different patch sizes and computed mean error for all pixels as ΔE_{μ} and CoV_{μ} as mean errors over all patches. Next, we traced plots for $E_{\text{ref}} = \Delta\text{E}_{\mu} + \lambda_{\text{raw}} \times \Delta CoV_{\mu}$ against $\mathbf{C}_{in}$ varying between 1% and 90% , for different values of λ_{raw}. Since the proportion of true outliers dictate the patch size that should be used for robust estimates, we tried different levels of outlier noise in spatial chromatic aberrations. To do this we used different levels of wavelet filtering to smoothen raw images while reading with *DCRAW* utility [Cof18].

Figure 11 shows all the plots for different combinations of λ_{raw} and wavelet filtering level. We found that for lower levels of λ_{raw}, ΔE dominates the minimization energy. For very high values of λ_{raw}, Δ*CoV* dominates this energy term. For a range of $\lambda_{\text{raw}} \in (700, 2000)$ we found synergetic combination of these two error types. We can use λraw to make better informed a choice about λ. To put things in perspective, for 24 blocks with two chroma channels and three channels for computing net mean *CoV* over all patches, $\lambda \times 24 \times 2/3 = \lambda_{\text{raw}}$, i.e. $16 \times \lambda = \lambda_{\text{raw}}$. Thus a range of $\lambda \in (40, 125)$ works well for different levels of outliers. Also, when $\Delta\text{E} = \Delta\text{E}_{\text{chroma}}$ is used, $\lambda \in (29, 84)$. Higher values of λ are warranted for photos with less noise. Our initial choice of $\lambda = 50$ is well within this range and we can rely on these plots to pick a good range of patch sizes. All three plots in Figure 11 indicate that the net synergetic energy has the best chance of minimization when the patch size is between 1% and 10%, with most plots indicating a minima near 5%. We thus empirically establish that around 5% of inliers must be used to compute robust patch mean color values for error minimization with our method.

General note on ΔCoV errors. Multiplying ΔCoV by 100 in Tables 2 and 3 above indicates changes in the scale-normalized variations as a percentage-wise reading. We found that even 1% change in the 'normalized' σ (i.e., CoV) even with a random spread across the patch, manifests itself into a disturbing visual artifact. Less than 0.3% error (i.e., $\Delta\text{CoV} = 0.003$) is not noticeable.

5.2 Qualitative Evaluations

We now look into qualitative evaluations of our method. For all qualitative evaluations, unless otherwise specified, we configured our method to use patches with 5% as the inlier threshold for robust mean estimation, along with nonlinear minimization at the end and white-point centric operations. The chromatic regression order is 2.

Aberration, Spatial Noise and White-balancing. One of the primary motivations was to introduce statistical measures into our color correction method to reduce spatial chromatic aberrations

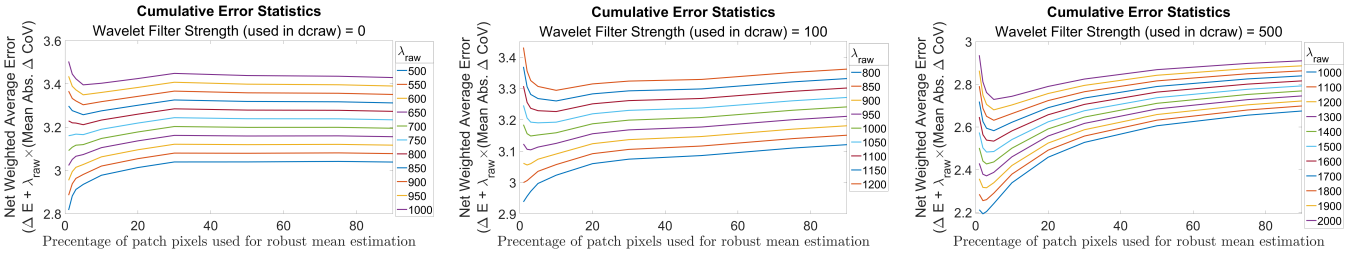

Figure 11: *Examining the impact of the patch fraction used for estimating robust statistics.*

from getting introduced in the process. For noisy photos, one can state the goal here is to retain the level of spatial variations in the output of the color correction step. We thus first make a comparative evaluation of noise levels between the input and the output for different color correction methods. For these tests, we use the same input photo, namely *CanonDarkroom*_01.*CR*2. Figure 12 and Figure 1(right image) show notable or even disturbing noisy chromatic aberrations for **RP** method with regression order greater than two. While at lower orders, **RP** method results in noise levels that are similar in subjective impressions about the spatial luminance and chrominance variations in the output blocks, as in the input blocks. Varghese et al. use linear regression and thus have similar pleasant correlation between the input and output noise levels. However, both of these methods are slightly off on resolving the white-balance accurately. **RP** method shows a slightly yellowish white shading and **VEM** method show a bluish tinge in the white blocks.

Similarly, we tested our method along with reference methods on few of the images from an openly available dataset [APCB19a]. This dataset contains poorly white-balanced input photos and their white-balanced counterparts. Figure 1 shows color correction results for different methods as applied to one tiger image. Figure 1 also shows that the white blocks and the tiger's eyebrows as well as upper lips are slightly yellowish for **VEM** method. **RP**(2) retains an overall bluish tinge as noticeable near the inscription on the rock in Figure 1 (middle image). Similarly, Figure 14 show that **RP**(3) overshoots orange color estimates for the tiger's body. We have found these slight offsetting of white blocks to often show up in the results for both reference state-of-the-art methods.

These subjective evaluations fairly balance the interpretation of color difference reduction statistics presented in the subsection above. In comparison, our method retains spatial variational characteristics while accurately white-balancing and color correcting a given photo.

Method Robustness. We tested our method against different levels of spatial aberrations, illumination conditions and few challenging reference cases as inputs. We first examined our method for different levels of spatial variations within the color chart blocks. With *CanonDarkroom*_01.*CR*2 as the input image subject to different levels of wavelet filtering by *DCRAW*, we created different levels of spatial smoothness for the color chart blocks. Figure 13 confirms that our method is consistent in retaining relative spatial variations in the output across different levels of local smoothness.

Next, we applied our method and both reference methods to the tiger image set where the input photos depict *Cloudy*, *Daylight*, *Fluorescent*, *Shady* and *Tungsten* lighting conditions. We color corrected these differently illuminated (re-rendered) photos with our

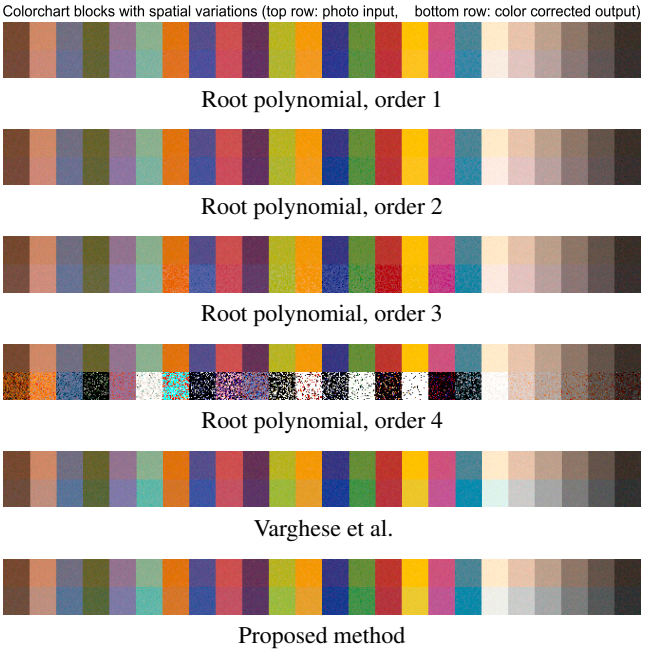

Figure 12: *Color chart block patches cutout from the input and output images.* ***RP*** *method shows yellowish shade for the white blocks while Varghese et al. results in a slight bluish tinge. Also, as the regression order increases, spatial chromatic aberrations get added to the results.*

method, **RP** method and **VEM** method. Figure 14 shows the results for various methods with different input illumination. Our method produces consistently white-balanced results across all input conditions. In comparison, other methods depict notable overall tonal differences in their results for different input conditions. To further compare the robustness in chromatic corrections, we took the color chart cut-outs from each row in Figure 14. Next, for each set of images corresponding to only one row in Figure 14, we picked individual pixels randomly from any one of the images within that set. Since the comparison is for chromatic accuracy, we fixed the output luminance to that for respective row's output corresponding to the cloudy tone mapping. Figure 15 shows the results of this random merging along with the cutout from the dataset's ground-truth photo. Upon closer inspection, this figure shows that overall our method has least local variations in those largely smooth chart blocks.

Finally, we tested our method against one of the difficult input case from the reference dataset [APCB19a]. Figure 16 shows an input image of a painting that has heavy red bias. Our method is able to reproduce the finer shades of red, yellow and purple on

the front woman's dress along with orange and green color shades around her. Our results are close to the given ground truth as well as Afifi et al.'s results. In contrast, **RP** and **VEM** retain stronger overall orangish tinge from the input image. This is in spite of having removed those effects reasonably well from the color chart blocks, given the near extreme input conditions for this case. In summary, we found our method to be fairly robust against several challenging input conditions, as demonstrated in this subsection.

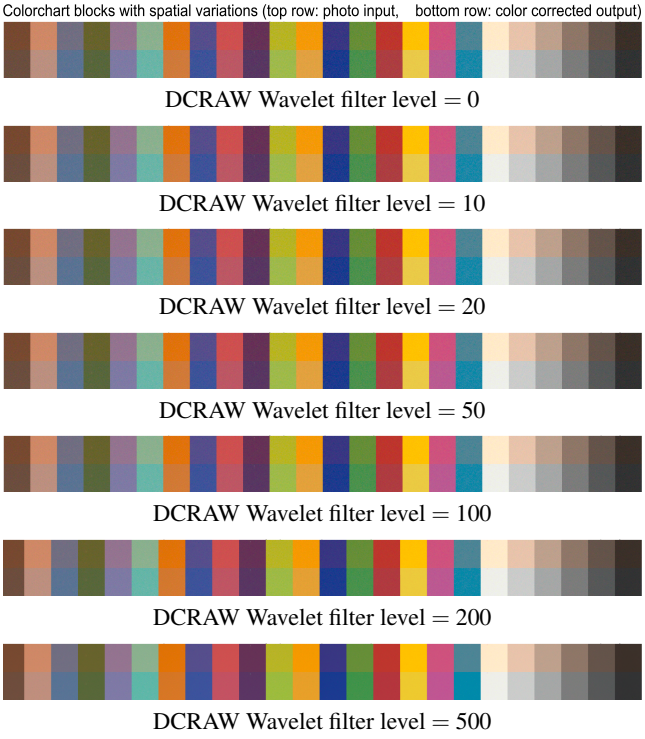

Figure 13: *Examining proposed method for robustness against different levels of input chromatic noise in the input photo. CanonDarkroom_01.CR2 is subject to different levels of wavelet filtering by DCRAW to smooth out spatial noises. In each case, our method produces subjectively similar impressions about noise intensity as well colorful variations in each patch.*

6 Summary and Future Work

In this work, we revisited the mathematical basis for transformation matrix based color correction approaches. We examined the inherent ill-posedness of this problem through mathematical constructs. We empirically examined limitations of color difference reduction as the sole criteria for estimating such transformation matrices. Based on our studies, we proposed augmenting the optimality criteria to incorporate spatial statistical measures. In essence, our strategy is to retain spatial variational relations during the color correction process. We achieve this by incorporating patch-based coefficient of variance differences to the objective function for non-linear minimization. We also devise a new method that processes luminance and chrominance data in a photo separately to reduce the impact of overfitting with ΔE errors. Furthermore, our method uses a good proportion of each color chart block patch to matching chroma. This allows spatial variations in the input photo to guide the estimation of the chroma transformation matrix. Our white-point centric corrections result in reliable white-balancing. We empirically established good range of parameter values to use for configuring proposed method. Our patch based minimization was found to converge in few hundred iterations with the step size reducing to the order of 10^{-7}. In contrast, other single mean (or median) color value based method requires thousands of iterations to reduce the step size to the order of 10^{-5}.

We evaluated our method both quantitatively as well as qualitatively to demonstrate that it performs better than current state-of-the-art methods in its league. Our method has a low memory foot print as it does not use data-driven or machine learning techniques. As one might expect, our method cannot deal with input conditions that express gamut loss due to heavy narrow band lighting conditions. Other than that, we showed that our method is consistent and robust against broad level illumination changes, spatial variations, heavy tone biases as well as different levels of spatial chromatic aberrations. In the future, we will research the use of this method in a material acquisition pipeline. Our method also holds good promise for the traditional digital photography related applications.

Acknowledgment

This work is fully supported by NSF grant #2007974. We are grateful to the anonymous reviewers for their constructive feedback.

References

[ABB22] AFIFI M., BRUBAKER M. A., BROWN M. S.: Auto white-balance correction for mixed-illuminant scenes. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (2022), pp. 1210–1219. 2, 3

[ABS16] ARAD B., BEN-SHAHAR O.: Sparse recovery of hyperspectral signal from natural rgb images. In *Computer Vision–ECCV 2016* (2016), Springer, pp. 19–34. 2

[ANU96] ARAI Y., NAKAUCHI S., USUI S.: Color correction method based on the spectral reflectance estimation using a neural network. In *Color and Imaging Conference* (1996), vol. 1996, Society for Imaging Science and Technology, pp. 5–9. 3

[APCB19a] AFIFI M., PRICE B., COHEN S., BROWN M.: Image recoloring based on object color distributions. 10

[APCB19b] AFIFI M., PRICE B., COHEN S., BROWN M. S.: When color constancy goes wrong: Correcting improperly white-balanced images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (2019), pp. 1535–1544. 2, 3, 4, 12, 13

[Bab23] BABELCOLOR: Xrite color chart spectral reflectances. https://babelcolor.com/, 2023. Last accessed: 26th August, 2023. 4

[Cof18] COFFIN D.: Decoding raw digital photos in linux. https://www.dechifro.org/dcraw/, 2018. Last accessed: 26th August, 2023. 9

[CRA*21] CARVALHO P. H., ROCHA I., AZEVEDO F., PEIXOTO P. S., SEGUNDO M. A., OLIVEIRA H. P.: Cost-efficient color correction approach on uncontrolled lighting conditions. In *Computer Analysis of Images and Patterns: 19th International Conference, CAIP 2021, Virtual Event, September 28–30, 2021, Proceedings, Part I 19* (2021), Springer, pp. 90–99. 3

[FD97] FINLAYSON G. D., DREW M. S.: Constrained least-squares regression in color spaces. *J. Electronic Imaging 6* (1997), 484–493. URL: https://api.semanticscholar.org/CorpusID:15397623. 5, 8

[FMH15] FINLAYSON G. D., MACKIEWICZ M., HURLBERT A.: Color correction using root-polynomial regression. *IEEE Transactions on Image Processing 24*, 5 (2015), 1460–1470. 1, 2, 4, 5, 6, 8

[GGDG22] GUARNERA G. C., GITLINA Y., DESCHAINTRE V., GHOSH

Figure 14: *Examining robustness against different illumination conditions. Input images were obtained from an open source dataset [APCB19b]. Our method produces consistent results across all input lighting conditions shown here.*

A.: Spectral Upsampling Approaches for RGB Illumination. In *Eurographics Symposium on Rendering* (2022), Ghosh A., Wei L.-Y., (Eds.), The Eurographics Association. doi:10.2312/sr.20221150. 3

[GGVDW11] GIJSENIJ A., GEVERS T., VAN DE WEIJER J.: Computational color constancy: Survey and experiments. *IEEE transactions on image processing 20*, 9 (2011), 2475–2489. 2

[HL02] HONG G., LUO M. R.: Perceptually-based color difference for complex images. In *9th Congress of the International Colour Association* (2002), vol. 4421, SPIE, pp. 618–621. 5

[HLR01] HONG G., LUO M. R., RHODES P. A.: A study of digital camera colorimetric characterization based on polynomial modeling. *Color Research & Application 26*, 1 (2001), 76–84. 2

[JF01] JOHNSON G. M., FAIRCHILD M. D.: Darwinism of color image difference models. In *Color and imaging conference* (2001), vol. 2001, Society for Imaging Science and Technology, pp. 108–112. 4

[JF03] JOHNSON G. M., FAIRCHILD M. D.: A top down description of s-cielab and ciede2000. *Color Research & Application 28*, 6 (2003), 425–435. 4

[JLGS13] JIANG J., LIU D., GU J., SÜSSTRUNK S.: What is the space of spectral sensitivity functions for digital color cameras? In *2013 IEEE Workshop on Applications of Computer Vision (WACV)* (2013), IEEE, pp. 168–179. 5

[KC14] KUMAR A., CHOUDHURY R.: Principles of colour appearance and measurement volume 1: Object appearance, colour perception and instrumental measurement. *Woodhead Publishing Limited, Cambridge 16* (2014), 2020. 2

[KFMA22] KUCUK A., FINLAYSON G., MANTIUK R., ASHRAF M.: Comparison of regression methods and neural networks for colour correction. 2, 3, 4

[MAF16] MACKIEWICZ M., ANDERSEN C. F., FINLAYSON G.: Method for hue plane preserving color correction. *JOSA A 33*, 11 (2016), 2166–2177. 2

[MM21] MACDONALD L., MAYER K.: Camera colour correction using

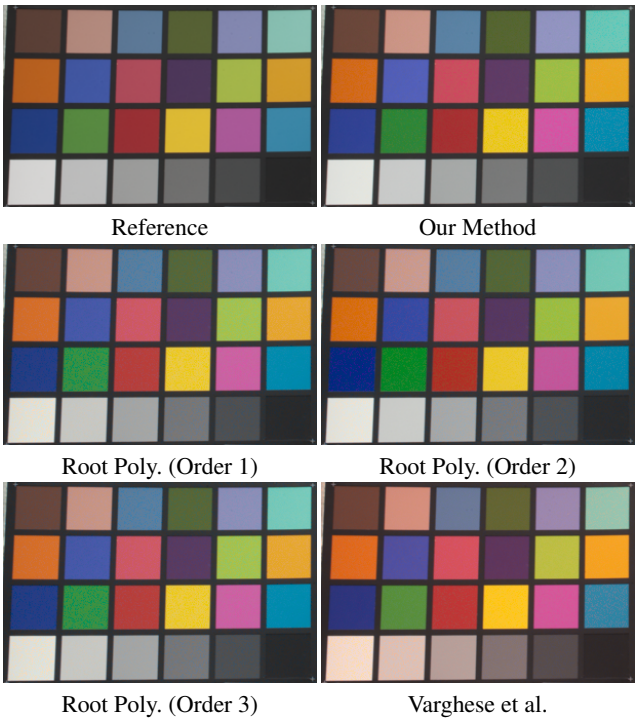

Figure 15: *Testing method robustness through random merging. Each row in Figure 14 is collapsed into a singe image by picking random pixels from random columns in the same row. Only cutout of the color chart is shown here. Input photo has relatively smooth chart blocks. Our method shows very low spatial distortions in this random merge test.*

Figure 16: *Testing method robustness against strong tonal biases in the input. Input photo sourced from Afifi et al. [APCB19b]. Our method reconstructs vivid pinkish, reddish and greenish tones that are closer to the ground truth than the **RP** and **VEM** methods.*

neural networks. In *London Imaging Meeting* (2021), vol. 2021, Society for Imaging Science and Technology, pp. 54–57. 3

[Pal99] PALMER S. E.: *Vision science: Photons to phenomenology*. MIT press, 1999. 2

[PH09] PEDERSEN M., HARDEBERG J. Y.: Shame: A new spatial hue angle metric for perceptual image difference. *Journal of Vision 9*, 8 (2009), 343–343. 5

[Sob19] SOBOTKA T. J.: Camera sensitivity source code. https://github.com/sobotka/camera-spectral-sensitivity, 2019. Last accessed: 26th August, 2023. 4

[SOF09a] SIMONE G., OLEARI C., FARUP I.: An alternative color difference formula for computing image difference. In *Proceedings from Gjøvik Color Imaging Symposium* (2009), no. 4, pp. 8–11. 5

[SOF09b] SIMONE G., OLEARI C., FARUP I.: Performance of the euclidean color-difference formula in log-compressed osa-ucs space applied to modified-image-difference metrics. In *11th Congress of the International Colour Association (AIC)* (2009), Sydney Australia, pp. 81–81. 5

[SSFJ22] SALESIN K., SEYB D., FRIDAY S., JAROSZ W.: Diy hyperspectral imaging via polarization-induced spectral filters. In *2022 IEEE International Conference on Computational Photography (ICCP)* (2022), IEEE, pp. 1–12. 2, 3

[TB05] TZENG D.-Y., BERNS R. S.: A review of principal component analysis and its applications to color technology. *Color Research & Application 30*, 2 (2005), 84–98. 2

[TV91] TRUSSELL H., VRHEL M.: Estimation of illumination for color correction. In *[Proceedings] ICASSP 91: 1991 International Conference on Acoustics, Speech, and Signal Processing* (1991), pp. 2513–2516 vol.4. doi:10.1109/ICASSP.1991.150912. 3

[Vrh93] VRHEL M. J.: *Mathematical methods of color correction*. North Carolina State University, 1993. 2

[Vrh03] VRHEL M.: Approximation of color characterization mluts with artificial neural networks. In *Proceedings 2003 International Conference on Image Processing (Cat. No.03CH37429)* (2003), vol. 1, pp. I–465. doi:10.1109/ICIP.2003.1246999. 3

[VWM14a] VARGHESE D., WANAT R., MANTIUK R.: Colorimetric calibration of high dynamic range images with a colorchecker chart. *Proceedings of the HDRi* (2014). 1

[VWM14b] VARGHESE D., WANAT R., MANTIUK R.: Colorimetric calibration of high dynamic range images with a colorchecker chart. 2

[WH09] WANG Z., HARDEBERG J. Y.: An adaptive bilateral filter for predicting color image difference. In *Color and Imaging Conference* (2009), vol. 2009, Society for Imaging Science and Technology, pp. 27–31. 5

[WH12] WANG Z., HARDEBERG J. Y.: Development of an adaptive bilateral filter for evaluating color image difference. *Journal of Electronic Imaging 21*, 2 (2012), 023021–023021. 5

[WSF*19] WANG L., SUN C., FU Y., KIM M. H., HUANG H.: Hyperspectral image reconstruction using a deep spatial-spectral prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019). 2, 3

[ZW97] ZHANG X., WANDELL B. A.: A spatial extension of cielab for digital color-image reproduction. *Journal of the society for information display 5*, 1 (1997), 61–63. 4