

Color Homography: Theory and Applications

Graham Finlayson, *Member, IEEE*, Han Gong^{1b}, and Robert B. Fisher^{1b}

Abstract—Images of co-planar points in 3-dimensional space taken from different camera positions are a homography apart. Homographies are at the heart of geometric methods in computer vision and are used in geometric camera calibration, 3D reconstruction, stereo vision and image mosaicking among other tasks. In this paper we show the surprising result that homographies are the apposite tool for relating image colors of the same scene when the capture conditions—illumination color, shading and device—change. Three applications of color homographies are investigated. First, we show that color calibration is correctly formulated as a homography problem. Second, we compare the chromaticity distributions of an image of colorful objects to a database of object chromaticity distributions using homography matching. In the color transfer problem, the colors in one image are mapped so that the resulting image color style matches that of a target image. We show that natural image color transfer can be re-interpreted as a color homography mapping. Experiments demonstrate that solving the color homography problem leads to more accurate calibration, improved color-based object recognition, and we present a new direction for developing natural color transfer algorithms.

Index Terms—Color homography, illumination estimation, color correction, color indexing, color transfer

1 INTRODUCTION

IN image formation there are two important parts, the geometry of how points in space map to image locations and the photometry of how illumination, surface reflectances and camera sensors combine to form the colors in an image. Broadly, the mathematical tools underlying our understanding of image geometry are non-linear reflecting the non-linear perspective nature of image formation. Important non-linear concepts include “solving for the homography” (e.g., relating subsequent frames in panorama stitching [1]) and epipolar geometry in stereo vision [2]). In contrast, the majority of methods in color/photometric computer vision are linear which, at least for simplified scenes such as the eponymous Mondrian world [3], [4] (the world consists of a patchwork of flat co-planar reflectances), reflects the physics of how images are formed. Linear color problems include, color correction [5], [6], [7] (e.g., mapping RAW colors from camera to display RGB) and modeling illuminant color change [8] e.g., for color object recognition [9].

In Fig. 1a, we illustrate a homography as the term applies in geometric computer vision. Here π_1 might denote the image—a perspective projection—of a plane (in 3-dimensions) and π_2 denotes the same plane viewed in a second image. The homography H relates the two planes.

In Fig. 1b, Ball_1 is the image of the side-view of a 4-color ball where the ball is lit from behind the camera with a

- G. Finlayson and H. Gong are with the School of Computing Sciences, University of East Anglia, Norwich, East Anglia NR4 7TJ, United Kingdom. H. Gong is a co-first author with equal major contributions. E-mail: g.finlayson@uea.ac.uk, gong@fedoraproject.org.
- R.B. Fisher is with the School of Informatics, University of Edinburgh, Edinburgh, Scotland EH8 9YL, United Kingdom. E-mail: rbf@inf.ed.ac.uk.

Manuscript received 10 Aug. 2016; revised 26 July 2017; accepted 19 Sept. 2017. Date of publication 3 Dec. 2017; date of current version 12 Dec. 2018. (Corresponding author: Han Gong.)

Recommended for acceptance by Y. Matsushita.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TPAMI.2017.2760833

white light. The same ball is lit from above with a bluish light, image Ball_2. The images are carefully registered so they are in pixel-wise correspondence. In an attempt to color correct panel Ball_2 to match Ball_1, we, naively, carry out a linear regression—for locations where both images have non-zero response—and find the best 3×3 matrix M mapping the corresponding pixels. The result of the regression results in image Ball_3. Notice that while the color cast due to the bluish light appears lessened, viewed closely, the colors are incorrectly mapped. In particular, notice that the red color segment looks wrong. Now, we now transform Ball_2 image to the image Ball_4 using the *correct* linear transform H .

In this paper, we propose that to map one *photometric* view to another we must map the colors correctly independent of shading. Since shading only affects the brightness, or magnitude, of the RGB vectors we wish, in effect, to find the 3×3 map which maps the color *rays* (the RGBs with arbitrary scalings) in one photometric view to corresponding rays in another. We note that this “ray matching” is precisely the circumstance in geometric computer vision when co-planar points in an image are mapped, via a homography, to corresponding points in a second image [2]. In analogy to the geometric case, at least 4 non-coplanar rays are required to solve for a color homography.

An RGB measurement without shading can be encoded as the (r, g) chromaticity coordinate: $r = R/(R + G + B)$ and $g = G/(R + G + B)$ (since the vector $[R \ G \ B]^T$ has the same orientation as $[r \ g \ 1 - r - g]^T$). In Fig. 1c, the 4 reflectances from the ball correspond to 4 points in an rg-chromaticity diagram and these define the quadrilaterals shown in the left and right of the panel (for respectively for the images Ball_2 and Ball_1). Assuming the illumination color change is linear, the mapping between the two chromaticity diagrams is precisely a homography (a fact we formally prove later).

We apply color homographies to help solve problems in three applications. In Fig. 2a, we show the picture of an

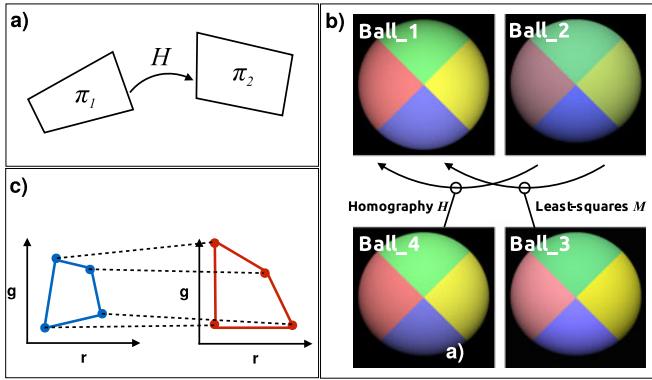


Fig. 1. Top left, panel (a), images of two planes are related by a homography. Right, panel (b), 4 images of a colored ball are shown. Ball_1 is the reference image where the illumination color is white and placed behind the camera. Ball_2 is the object illuminated with a blue light from above. Respectively, Ball_3 and Ball_4 are the least-squares mapping and the homography match (in both cases the aim is to *correctly* undo the illumination color) from Ball_2 to Ball_1, Bottom right, panel (c), the chromaticities from Ball_2 matched to corresponding chromaticities in Ball_1.

image in the RAW RGB space of a camera and the corresponding reproduction when the colors are corrected for display (where both images are also tone mapped for printing in this paper). In computer vision, the idea that pervades color correction is that all we need to do is find the best least-squares transform mapping the color checker shown in the RAW image to pre-measured correct display RGBs with some of the problem focus directed towards automatically finding the checker in the image [11], [12]. However, consistent with other recent work [13], [14], we

found that the illumination intensity could vary significantly over the image of a checker. When shading varies, the color correction transform is found by solving for the homography relating the colors from the RAW to reference display RGBs.

Consider that we have a database of colorful objects where the color content of each image is represented by its chromaticity distribution. By matching color distributions we can obtain surprisingly accurate color object recognition [15], [16]. However, object recognition performance degrades when the light color changes [17]. The role of homographies in color object recognition is summarized in Fig. 2b. Clearly, image I_1 matches image I_2 (it is of the same object) but image I_3 is an image of a different object. The distribution of rg chromaticities for image I_1 is shown in purple in both chromaticity diagrams in the middle of Fig. 2b. The rg distribution of images I_2 and I_3 are overlaid in green. Notice, however, that the distributions do not match in either case. Indeed, even though images I_1 and I_2 are of the same object their chromaticity distributions don't match because images I_2 and I_3 are taken with respect to a warmer illumination color.

Our hypothesis is that the chromaticity distributions for the same scene lit by two different lights but where the image shading might change will be related by a homography. In the last row of Fig. 2b, we show the output of homography matching for chromaticity distributions for I_1 matched to I_2 and I_1 matched to I_3 . We see the chromaticity distribution of $H(I_1, I_2)$ (the colors in I_1 homographically transformed to match those in I_2) now overlap and we can conclude the two images plausibly have the same color content. In contrast $H(I_1, I_3)$ fails to match the chromaticity

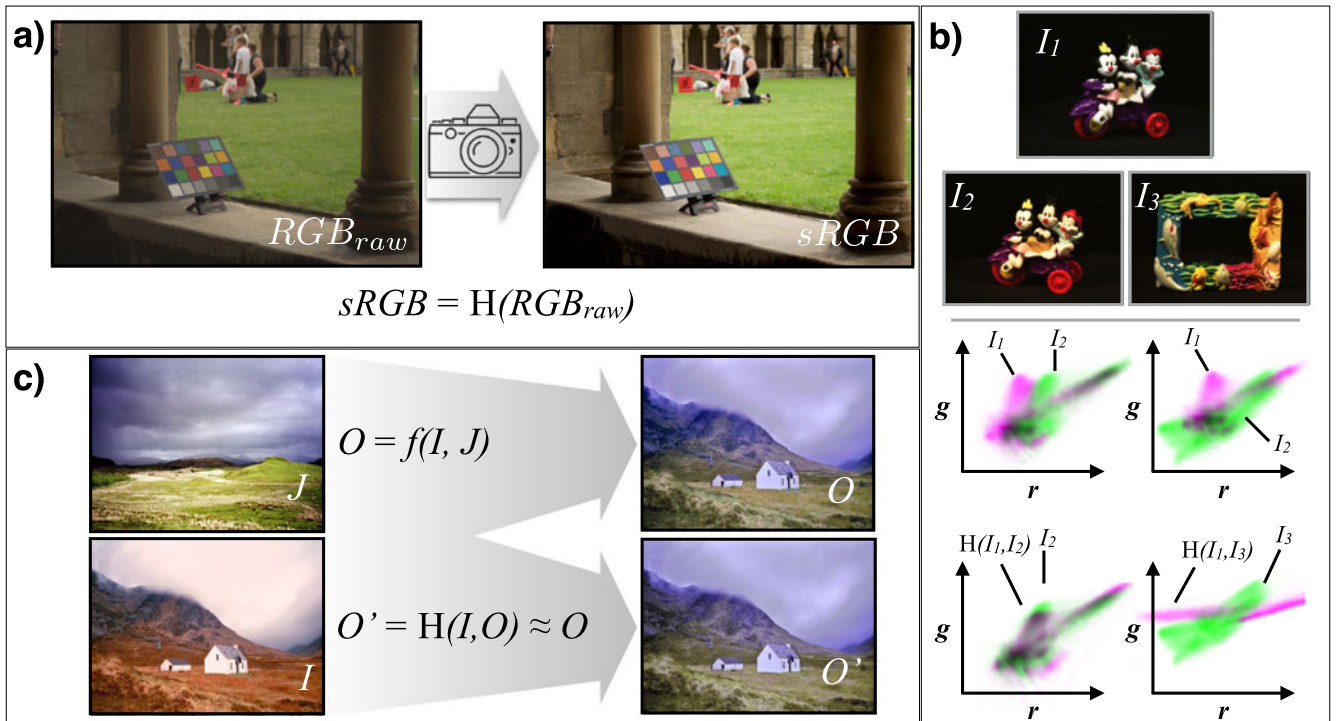


Fig. 2. a) Color correction (mapping RAW to display sRGB [10]) is a homography problem. b) The top contains 3 images of colorful objects. Histogram matching of chromaticity distributions of images I_1 with I_2 and I_1 with I_3 are shown in the middle (chromaticity distribution for I_1 is in purple and for I_2 and I_3 is in green). Solving for the color homographies best mapping I_1 to I_2 and I_1 to I_3 (respectively, $H(I_1, I_2)$ and $H(I_1, I_3)$) results in the histograms shown in the bottom row. The distributions for I_1 and I_2 now match and the object is correctly identified. c) Color transfer for matching the colors of an input image I to the colors of a target image J can be reinterpreted as a simple color homography mapping $H(I, O)$.

distributions. Note this latter match is poor (even poorer than the input) due to the mechanics of how the match is made which we discuss in Section 3.2. The matching method in this case actually informs us that no good match is possible, but for purposes of illustration, we show the *weak* match that was found.

Finally, we consider color homographies in the context of image color transfer. Automated color transfer is often required in professional photo editing. Artists can choose a desired target picture and manipulate another picture to match its color styles to the target. In Fig. 2c, Image I has its colors transferred to map those shown in Image J and the result of the method [18] is denoted $O = f(I, J)$. Using our color homography theorem, we can approximate the color transfer result O by a color homography model which produces a close approximation image $O' = H(I, O)$ (I mapped by a homography to approximate O). By enforcing color space homography the resulting image can be physically interpreted as being an image of the same scene as the input image but under a different illumination. We can re-interpret most color transfer effects using a color homography model. This result is useful because, unlike many of the color transfer algorithms, color homographies can be computed quickly. Indeed we have found that we can run a computationally expensive algorithm on a thumbnail image, compute the homography and then apply the homography to the full-resolution image. Another benefit of our homography-based model is that we often remove the artifacts introduced by some color transfer algorithms. That is, we run a color transfer algorithm (whose output has artifacts) and approximate it as a homography. Because of the simpler transfer that is enforced by the homography many of the spatial artifacts either disappear or are mitigated.

Experiments demonstrate the power of our color space homography idea. Regarding color correction we report a significant improvement in color accuracy compared to the commonly used color correction methods. For object recognition on the Amsterdam Object Image Database [19] (> 100000 images encompassing a range of capture conditions), homographically matching chromaticity distributions supports state-of-the-art color-based object recognition. Finally, we found that existing color transfer algorithms can be re-interpreted as a color homography mapping.

In Section 2, we review the homography problem from geometric computer vision and relate this to linear image formation in color/photometric vision. Section 3 presents the color homography theorem together with a discussion on how to, in practice, solve for a color homography. Experiments in color correction, color object recognition and color transfer are reported in Section 4. We discuss application of homographies to non-RAW images in Section 5. The paper concludes in Section 6.

2 BACKGROUND

2.1 Geometry

For the geometric planar homography problem, we write

$$\begin{bmatrix} \alpha x \\ \alpha y \\ \alpha \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} \triangleq \underline{x} = H(\underline{x}'). \quad (1)$$

In Equation (1), (x, y) and (x', y') denote corresponding image points—the same physical feature—in two images. In homogeneous coordinates the vector $[a \ b \ c]^\top$ maps to the coordinates $[a/c \ b/c]^\top$ and so, in Equation (1), the scalar α cancels to form the image coordinate (x, y) . For all pairs of corresponding points (x, y) and (x', y') that lie on the same plane in 3 dimensional space, Equation (1) exactly characterises the relationship between their images [2]. To solve for a homography (e.g., for image mosaicking), we need to find distinctive feature points in pair of images, match them to find candidate corresponding points, then solve for the *best* homography and finally warp the image to bring one image into the coordinate frame of the other [20], [21]. Homographies are at the heart of geometric methods in computer vision and are used in geometric camera calibration [22], 3D reconstruction [23], stereo vision [24] and image mosaicking [25] amongst other tasks.

2.2 Color

A physically accurate model of Lambertian image formation where the illumination impinging on a scene is a single color is written as

$$\underline{\rho}^x = \alpha^x \int_{\omega} E(\lambda) S^x(\lambda) \underline{Q}(\lambda) d\lambda, \quad (2)$$

where respectively $E(\lambda)$, $S^x(\lambda)$ and $\underline{Q}(\lambda)$ denote the spectral power distribution of the light, the spectral reflectance of a surface, and the vector of R-, G- and B-spectral sensitivities of the camera. And, the integral is taken over the visible spectrum ω . The superscript x denotes spatial dependency and α^x is a scaling factor encoding brightness changes due to the relative interplay between where the lights are positioned and the orientation of the surface in the scene (e.g., Lambert's law) and the quantity of light (e.g., to model extended light sources). It is well established [26], [27], [28] that surface reflectance can often be written to a tolerable approximation as the weighted sum of three basis functions

$$S(\lambda) \approx \sum_{j=1}^3 S_j(\lambda) \sigma_j. \quad (3)$$

When $\alpha^x = 1$ (so we can ignore spatial dependency) we have the so-called Mondrian-world [3]. There we have no shading in images and this simplifies consideration of how colors change with illumination and/or imaging device. Under the Mondrian-world assumption, we write

$$\begin{aligned} \underline{\rho} &= \Lambda^{E(\lambda), \underline{Q}(\lambda)} \underline{\sigma} \\ \Lambda_{ij}^{E(\lambda), \underline{Q}(\lambda)} &= \int_{\omega} E(\lambda) S_j(\lambda) Q_i(\lambda) d\lambda. \end{aligned} \quad (4)$$

In Equation (4), image formation is a 3×3 linear matrix Λ multiplying a 3×1 weight vector $\underline{\sigma}$. Two important results follow from Equation (4). First, that the 3×3 matrix

$$\Lambda^{E_2(\lambda), \underline{Q}(\lambda)} [\Lambda^{E_1(\lambda), \underline{Q}(\lambda)}]^{-1}, \quad (5)$$

maps colors viewed under illuminant $E_1(\lambda)$ to those recorded under $E_2(\lambda)$ assuming the same camera sensitivities $\underline{Q}(\lambda)$.

Second that the 3×3 transform relating colors recorded by cameras with the respective sensitivities $\underline{Q}_1(\lambda)$ and $\underline{Q}_2(\lambda)$ is written as

$$\Lambda^{E(\lambda), \underline{Q}_2(\lambda)} [\Lambda^{E(\lambda), \underline{Q}_1(\lambda)}]^{-1}. \quad (6)$$

assuming the same illuminant $E(\lambda)$.

Of course it is well known that Equation (3) is only approximate. Indeed, illuminant metamerism [29]—the phenomenon that two surfaces look the same under one light can look different under another—cannot be described by a 3×3 matrix of illuminant change. Yet, metamerism (illuminant or sensor) is rare. Moreover Marimont and Wandell [30] extended the linear model formalism to incorporate image formation into the derivation of the optimal linear basis and found that 3×3 matrices could well account for illuminant and sensor change. A similar result was reported by Funt et al. [31], [32]. While most of the literature assumes a single global illuminant, other recent work are also proposed to solve multi-illuminant color constancy [33], [34], [35], which is an emerging and challenging area of research.

That illumination (or device) colors map across images using a linear transform is a common assumption and is widely reported in the literature including, in color correction [5], [6], [7] color object recognition [9], [36] and illuminant estimation [4]. However, the Mondrian-world does not generally hold. Indeed, color intensity can and does vary on a per-pixel basis due to the relative position of light and surface and also due to the quantity of light varying across a scene. When illumination color is held fixed then a chromaticity representation of color can be used to factor out per pixel shading variation e.g., [37].

To deal with a changing light color, the so-called “diagonal” model of image formation is often employed with respect to which color change across images is modeled as a diagonal matrix [38], [39] multiplying the image colors. The diagonal model is at the heart of the “comprehensive color image normalization” [40] and the $m1m2m3$ coordinate system [41] both of which are light color plus shading invariant image features. Further, specially chosen spectral band ratios (for example, R/G and B/G) have an analogous diagonal property (the 2-d band ratios are mapped across illumination by a 2×2 diagonal matrix). This property is exploited in illuminant color and shading-independent histogram matching for object recognition [42] and illuminant estimation [43], [44].

Of course moving to a chromaticity representation means one of the 3 degrees of freedom measured in an RGB image has been lost. In calibrated color correction—mapping the RAW RGBs recorded for a known color chart to a standard color space—it is possible to find the full 3×3 matrix mapping the orientation of input color vectors to the orientation of output colors by a searching strategy [13]. Alternately, in [14], solving for the best shading while simultaneously finding the color correction matrix was formulated as an Alternating Least-squares approach. Both these methods deliver significantly lower correction error compared with a shading blind linear least-squares optimization. In display calibration [45] it was shown that 4 chromaticities sufficed for color calibration.

3 COLOR HOMOGRAPHY

Let us map an RGB $\underline{\rho}^\top$ to a corresponding RGI (red-green-intensity) \underline{c}^\top using a 3×3 matrix C

$$\underline{\rho}^\top C = \underline{c}^\top$$

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} R \\ G \\ R+G+B \end{bmatrix}^\top, \quad (7)$$

interpreting the right-hand-side of Equation (7) as a homogeneous coordinate (e.g., here, as chromaticities) we see that

$$\underline{c} \propto [r \quad g \quad 1]^\top, \quad r = \frac{R}{R+G+B}, \quad g = \frac{G}{R+G+B}. \quad (8)$$

In the following proof it is useful to represent 2-d chromaticities by their corresponding 3-d homogeneous coordinates.

Theorem 1 (Color Homography). *Chromaticities across a linear change in capture condition (light color, shading and imaging device) are a homography apart.*

Proof. First we assume that across a change in illumination or a change in device where the shading is the same (e.g., for the Mondrian-world) the corresponding RGBs are related by a linear transform M (i.e., Equations (5) and (6) hold). Clearly, $H = C^{-1}MC$ maps colors in RGI form between illuminants. Due to different shading, the RGI triple under a second light is represented as $\underline{c}'^\top = \alpha \underline{c}^\top H$, where α denotes the unknown scaling and $^\top$ denotes transpose. Without loss of generality let us interpret \underline{c} as a homogeneous coordinate i.e., assume its third component is 1. Then, chromaticity coordinates \underline{c}^\top and \underline{c}'^\top are a homography apart. \square

In geometry, homographies are applied for mapping spatial coordinates in one image to correspondences in another. In color homography we are also interested in this 2D-2D match problem (chromaticity to chromaticity mapping). However, for some applications—e.g., color transfer—the homography that maps 3D colors to 3D color matches (in a shading independent way) is the apposite tool.

3.1 Color Homography Estimation by Alternating Least Squares

Suppose A and B denote respectively $n \times 3$ matrices of n corresponding pixel RGBs with respect to two images of the same scene where the illumination changes (and, also possibly the camera properties). The color change is modeled as a linear transform (Equations (5) and (6)) but because of the relative positions of light and surfaces there might also be per-pixel shading perturbations. Assuming that the Lambertian image formation is an accurate physical model, the relationship between A and B is written as

$$DAH \approx B, \quad (9)$$

where D is an $n \times n$ diagonal matrix of scalar shading factors and H is a 3×3 homography color correction matrix. If both D and H are applied to A , then we call this a *shading homography*. We can solve Equation (9) by using Alternating Least-Squares (ALS) [14], [46], [47] described in Algorithm 1. There

$\|\cdot\|_F$ denotes the Frobenius norm. And, at iteration i , H^i and D^i are found using the closed form Moore-Penrose inverse. The least squares fit of A to B is equal to $(A^\top A)^{-1} A^\top B$. Each scalar component at the j th diagonal element of D^i is the result of a least squares fit of the j th row of A^{i-1} to the j th row of B (where again we use Moore-Penrose inverse). The effect of the individual H^i and D^i can be combined into a single matrix $D = \prod_i D^i$ and $H = \prod_i H^i$ (where the product is taken by post-multiplying matrices). That is, $DAH \approx B$. ALS converges to a local minimum [48].

Algorithm 1. Homography from Alternating Least-Squares

```

1  $i = 0, \min_{D^0} \|D^0 A - B\|_F, A^0 = D^0 A;$ 
2 repeat
3    $i = i + 1;$ 
4    $\min_{H^i} \|A^{i-1} H^i - B\|_F;$ 
5    $\min_{D^i} \|D^i A^{i-1} H^i - B\|_F;$ 
6    $A^i = D^i A^{i-1} H^i;$ 
7 unit  $\|A^i - A^{i-1}\|_F < \epsilon;$ 

```

Finally, note how we initialize the matrix A^0 to be the closest least-squares fit of the rows of A to the rows of B . This initialization is performed for two reasons. First, we find slightly better visual results if we applied D before H . Second, it simplifies the proof of Theorem 2.

To motivate the theorem below we know from geometric computer vision that given the images of the two planar regions each enclosed by 4 corner points in two views (assuming certain assumptions) there is a unique homography and that this can be found directly [2]. In Theorem 2, we capture the conditions where the ALS method can find the unique homography given 4 pairs of colors captured across viewing conditions.

Theorem 2 (ALS Uniqueness). *The ALS method finds the homography—unique, up to a scalar multiplier—when there are 4 pairs of corresponding non-collinear colors A and B each of which is a 4×3 full-rank matrix (i.e., $\text{rank}(A) = \text{rank}(B) = 3$) and the null space—which is represented by the 4-vector \underline{v} —of $(A^i - B)$ at convergence is not sparse. A vector \underline{v}^i is sparse if 1, 2 or 3 (but not all 4) of its entries are 0.*

Proof. In this case, A and B are 4×3 matrices. Because the elements of A and B are strictly positive and full rank, the 3×3 cross product matrices $A^\top A$ and $A^\top B$ are also full rank. We point this out because this in turn implies that at every time we calculate D^i and H^i in the ALS method that these matrices themselves are full rank. That is, so long as A and B begin as full rank matrices the matrices A^i , B , D^i and H^i are all themselves full rank i.e., we never encounter the circumstance where an individual least-squares fit introduces rank deficiency.

Now let us suppose that, the stopping condition of the ALS procedure is met in the i th iteration. To prove Theorem 2, we need to first prove the following lemma. \square

Lemma 1. *When the ALS algorithm converges at step i then $D^i = \mathcal{I}_{4 \times 4}$ and $H^i = \mathcal{I}_{3 \times 3}$ where \mathcal{I} denotes the identity matrix.*

Proof. Let us assume Lemma 1 is false then when the algorithm converges it must follow we can write $DA^{i-1}H = A^i$

(where here D and H are respectively a 4×4 diagonal matrix and an arbitrary 3×3 linear transform). This can only be true if the rows of A^{i-1} are row Eigenvectors of H . This cannot be the case because then two of the rows of A^{i-1} would be the same (up to a scalar) i.e., collinear and in the statement of the theorem this was assumed not to be the case (i.e., to get to A^i involves multiplying A by a series of full rank matrices. If the rows of A are not collinear then neither are the rows of A^i). \square

Let us now prove Theorem 2 by contradiction. We will show that if the alternating least-squares procedure has converged and the unique homography has not been found then the assumption of convergence cannot hold (and, yet, we know the algorithm converges [48]).

Let us assume that $A^i \neq B$. It follows that

$$B = A^i + \underline{v} \underline{w}^\top \quad \text{s.t.} \quad \underline{v}^\top A^i = \mathbf{0},$$

where respectively \underline{v} and \underline{w} are 4×1 and 3×1 vectors, and \underline{v} spans the column null-space of A^i . From Lemma 1, on convergence, both D^i and H^i are identity matrices. Let us consider the final step in the algorithm (step 5 just before the algorithm stops). By Assumption the algorithm converges and $A^i \neq B$ and $H^i = \mathcal{I}_{3 \times 3}$ (from Lemma 1):

$$A^i + \underline{v} \underline{w}^\top = B = D^i A^i H^i = D^i A^i \mathcal{I}.$$

The least-squares per-row scalar returned in Line 5 of Algorithm 1 can be written in closed form (we are simply using the Moore-Penrose inverse)

$$D_{jj}^i = \frac{\underline{a}_j^\top (\underline{a}_j + v_j \underline{w})}{\underline{a}_j^\top \underline{a}_j} = \frac{\|\underline{a}_j\|^2 + v_j \underline{a}_j^\top \underline{w}}{\|\underline{a}_j\|^2},$$

where the j th row of A^i , denoted \underline{a}_j^\top , has a magnitude $\|\underline{a}_j\| > 0$, because A^i has full rank and no two points are collinear. It can be shown that the scaling of the j th row of A^i that best matches the j th row of B in a least-squares sense can be written as

$$D_{jj}^i = 1 + \frac{v_j}{\|\underline{a}_j\|^2} \underline{a}_j^\top \underline{w},$$

which implies $D_{jj}^i = 1 \Leftrightarrow \underline{a}_j \underline{w} = \mathbf{0}$. Because A^i is assumed to be full rank, it is possible that three of its four rows can satisfy $\underline{a}_j^\top \underline{w} = 0$. However, this can not hold for the fourth row otherwise A is not full rank. And by our assumption of *non-sparsity* all elements of \underline{v} are either non-zero or zero. If they are all non-zero i.e., the algorithm has not converged (when by assumption it has). We have a contradiction that $D^i \neq \mathcal{I}$.

Of course the reader will have noticed a sleight of hand on our part. We made a *non-sparsity* assumption on the null-space of $A^i - B$ at convergence. Albeit rarely, we found the alternating least-squares procedure might terminate with a sparse null space vector and in this case with $A^i \neq B$ (the algorithm stops but we do not find the unique homography).

To investigate convergence empirically, we uniformly randomly generated 4×3 matrices A (matrix elements in the interval $[0, 1]$ to 2 decimal places) and then randomly generated matrices D (elements chosen uniformly and

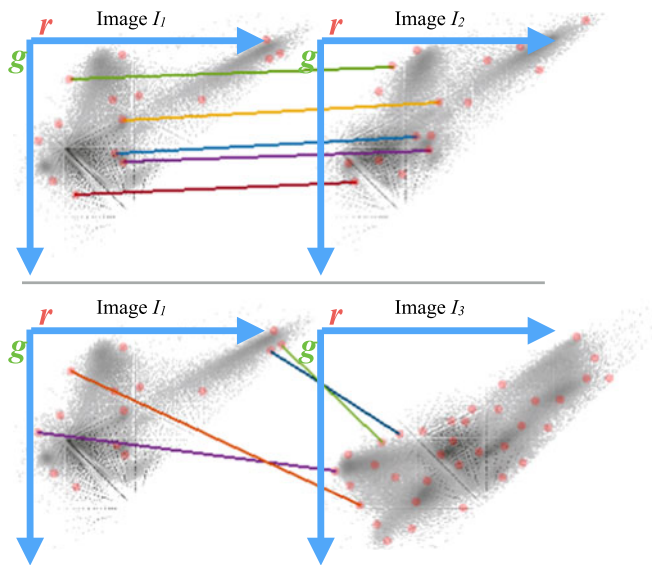


Fig. 3. Top: ASIFT [21] can match the chromaticity distributions of images I_1 and I_2 . Bottom: ASIFT cannot match the chromaticity distributions of image I_1 to image I_3 (images I_1 , I_2 and I_3 shown in Fig. 2b).

randomly from $[0, 1)$ and H (elements chosen from the standard normal distribution $N(0, 1)$). Then, we compute $B = DAH$. The matrix B created in this way could plausibly be matching colors under different lights or different devices. We then ran the ALS procedure to solve for the homography, to discover D and H .

Over 10,000 runs the ALS procedure converged to the correct answer (the non-sparsity condition was satisfied) over 96.5 percent of the time. For the remaining 3.5 percent of the cases that did not converge we mapped A and B to new matrices A' and B' according to

$$A' = D^A A H^A, \quad B' = D^B A H^B,$$

where D^A and D^B are randomly chosen positive matrices (entries in $[0, 1)$) and H^A and H^B are 3×3 matrices with elements drawn from $N(0, 1)$. We then run our procedure a second time. Assuming convergence to the correct answer i.e., we find matrices D' and H' such that $D' A' H' = B'$. It follows that $D = [D^B]^{-1} D' D^A$ and $H = H^A H' [H^B]^{-1}$. In all cases, this kind of random perturbation sufficed to make ALS converge to the correct answer. In all cases (even without the perturbation), the % difference $\|DAH - B\|/\|B\|$ between the actual and fitted homography, between DAH and B was less than 0.5 percent (and usually orders of magnitude smaller). Indeed, we have not found a compelling visual example where the output of the ALS procedure—when it converges to the wrong answer—appears different from the overall best answer (zero error for the 4 point homography fitting). In Appendix A we present a numerical example where the 4 point ALS minimization fails to solve for the homography.

Finally, in thinking about solving for a homography by mapping the original matrices A and B to counterparts (which, for ALS then converges correctly) we were, as a side-effect, able to write the homography calculation directly as a simple matrix computation (see Appendix B). While wholly equivalent to the “direct method” [2], the form of the equation appears novel.

3.2 Homography-Based Chromaticity Matching

In geometric computer vision we find the homography relating two images in three steps. First, we find distinctive features, second we find candidate matching locations by pairing distinctive features which also have similar underlying image structure and third, we find the best homography match accounting for as many of our paired image points as possible. See [20] for a general discussion for finding corresponding feature points between two images. Here, we treat the chromaticity distributions *as images* and seek to find and match these distributions analogously to the geometric case (i.e., we find and match interest points found for chromaticity distributions). An interesting technical issue is that we would, ideally, like to be able to find the same feature points robust against the image transformation we are trying to discover. The ASIFT [21] algorithm is a methodology for matching images that is fully invariant to affine image modification. In the top of Fig. 3, we show the ASIFT match for the color distributions of images I_1 and I_2 from Fig. 2. Visually, we see that structurally similar (corresponding to the shape of the chromaticity distribution) points are found in both images. Pairs of correspondences found by RANSAC [2] matching are also shown (as the lines joining the ASIFT points in the two distributions). Among other popular feature point matching algorithms, such as SIFT [20], MSER [49], SURF [50], and Harris corner feature [51], we found ASIFT delivers the best chromaticity matching result.

The result of solving for the homography relating matched points is shown in the chromaticity diagram shown bottom left of Fig. 2b. There the green distribution shows the chromaticities of image I_2 and in purple the homographically transformed chromaticities from image I_1 . The distributions, and so the images, match. The result of applying the best homography relating the distributions for images I_1 and I_3 (based on the matched points shown in the bottom of Fig. 3), does not match the chromaticity distributions (see bottom right, Fig. 2b). We can conclude that the object in image I_1 does not match the object in image I_3 .

While there are many ways of matching color distributions we found that a simple structural match score—a measure of how well a homography matches ASIFT points—provided a powerful way to determine whether one chromaticity distribution matched another. Suppose ASIFT returned N and M points of interests from a pair of chromaticity distributions and that (via RANSAC) we found that we could match m pairs of points by finding the best homography. Then, our structural match is defined as

$$\text{structural_match} = \frac{m}{\sqrt{MN}}. \quad (10)$$

A structural match score of 1 means all ASIFT points found in both chromaticity distributions are, placed in correspondence, a homography apart. Section 4.2 shows an application of object recognition which adopts this structural match score as a measure of image similarity. Later we will see that the number of ASIFT points extracted from a chromaticity distribution—which we call its *structural complexity*—can be used to predict whether a homography-based distribution match will, in fact, be possible.

TABLE 1
Errors for Color Correction (X-Rite Classic Color Checker)

Method	Mean	Median	95%	Max
ΔE (Lab)				
Least-Squares	6.16	5.67	12.27	13.83
Root-Polynomial	5.67	4.67	14.60	16.97
ALS	3.71	3.27	8.24	9.02
Homography	3.40	2.59	9.20	10.28
ΔE (Luv)				
Least-Squares	7.02	6.63	14.23	15.55
Root-Polynomial	6.69	5.46	16.88	19.19
ALS	4.17	3.70	8.93	9.89
Homography	3.88	2.97	9.97	10.76
ΔE RGB ($\times 10^{-2}$)				
Least-Squares	7.58	6.34	19.07	30.11
Root-Polynomial	9.07	7.21	24.72	30.60
ALS	4.36	3.16	14.01	26.05
Homography	4.01	2.82	14.62	24.44

It is known that ASIFT is about 13.5 times the cost of SIFT [21]. Importantly, efficient implementations of SIFT exist including in real-time on mobile devices. Also, in general SIFT operates on large (> 1 megapixel) images. In contrast, though we use the more expensive ASIFT to match thumbnail images: our “images” are small: chromaticity distribution that have just 320×320 bins. Thus, despite the $13.5 \times$ cost multiplier compared to SIFT, the cost of ASIFT in matching chromaticity distributions is similar to running SIFT on full size images.

4 EXPERIMENTS

4.1 Color Correction

4.1.1 Color Correction Using a Color Chart

In color correction—mapping RAW RGBs to a display color space—the target RGBs are known to vary in intensity. Indeed, serious photographers will take a picture of a color target (such as the Macbeth color checker [52] shown in Fig. 2a) and a second picture of a uniform gray target with same size in the same location. By dividing the RGB image of the checker by the image of the gray-target the shading is removed and then the shading corrected RGBs can be mapped to known reference display color coordinates using simple least-squares. However, in computer vision, and even for the vast majority of photographers, this two-step approach to calibration is rarely taken (e.g., see [11], [12]). If nothing else it is invasive and in some cases cannot be done at all (e.g., in an on-going surveillance situation). However, the photographic “best practice” allows us to measure useful data to evaluate homography versus linear least-squares color correction. We can find the best least-squares 3×3 matrix mapping the non-shading corrected checker to a reference target and then apply this matrix to the shading corrected target. Or, we transform the target using a homography. This experimental methodology is described in detail in [13].

Because the display RGBs are measured in coordinates relevant to human vision, e.g., sRGB [10], the color error can be converted to the CIE Lab, CIE Luv, RGB error metrics. The CIE formulas [29] return a ΔE error of 1 if two patches



Fig. 4. Two example images with non-uniform shading used for color correction test.

are just discernible from each other. The ΔE RGB error between two RGB vectors \underline{p} and \underline{q} is calculated as $\|\underline{p} - \underline{q}\|$. Similarly, ΔE in Table 1 is also for the Luv and Lab representations. In Fig. 2a, we show the RAW and rendered (JPEG) images of one of our images. Two examples containing outdoor and indoor lighting conditions for our color correction evaluation are shown in Fig. 4. In total, we captured 13 images at Norwich cathedral and around the University of East Anglia campus. Our image set contains both indoor and outdoor illumination. To calculate the best homography we randomly chose 4 matching colors (according to the known correspondences) and using RANSAC, chose the homography that minimized correction error (ΔE Luv). The mean, median, 95 percent quantile and max ΔE errors calculated over the 13 images are reported in Table 1 where we compare performance to simple least-squares, root-polynomial [53], and ALS [14] (which is shown to be improved on [13]). Compared with the ALS color correction, all mean errors are improved by about 9 percent, and all median errors are improved by about 21 percent, at the cost of getting a slightly higher 95 percent quantile error and maximum error (except for RGB metric). Compared to the least squares and root-polynomial methods—which are shading dependent—the advantage of using the homography based formalism is quite significant (e.g., the median errors—for all 3 error metrics—are halved).

4.1.2 Color Correction for Single Color Objects

The solution for a color homography requires at least 4 non-collinear chromaticities. For a perfect convex-shaped monochrome Lambertian surface viewed under a single light source, the resulting chromaticity distribution—in the idealized case—will comprise a single point. However, most objects are not perfectly Lambertian—there is a highlight component—and are not convex and so there could be inter-reflections. Both specular highlights [54], [55] and inter-reflections [56], [57], [58] lead to the chromaticity distribution of a single object being more than a single point. Indeed, in the presence of specular highlights and inter-reflections we expect a spread of points in 3D RGB space and a spread (not all collinear) points in chromaticity space.

In columns A and B of Fig. 5 we show, respectively, 4 monochrome objects viewed in 2 viewing conditions (different colored light in different positions). In column C we show, in green, the chromaticity distributions of images in column A and in pink the chromaticities of the images in B. Clearly, there is a spread in chromaticities in both viewing conditions. If we are to match chromaticity distributions by solving for a homography then it is a necessary condition that the images (and their chromaticities) in columns A and B are related by a homography. Using the ALS method (the




Fig. 5. Given pixel-wise correspondence between Images A and B, a color homography is solved for to convert Image A to Images D (with shading correction) and E (without shading). Image F is the worse conversion result when linear least-squares is adopted. Image C shows the rg chromaticity spread (Image A in green, Image B in pink).

pixels are in correspondence) we solve for the best 3×3 matrix and shading correction that relates the images in column A to counterparts in column B. Mapping the images in A with the solved shading and color corrections results in the fitted images, shown in column D. Visually, the fit is excellent. In column E we apply just the 3×3 homography matrix (no shading correction) to the inputs from A. Now we have the object mapped to viewing condition B but the original shading preserved. Finally, in column F we find a pixel-wise least-squares fit. Here the wrong 3×3 matrix is found (since it is attempting to best compensate shading and colour correction in a single matrix).

4.2 Color Object Recognition

The Amsterdam Object Image Database [19] is a large 1,000 object database widely used to benchmark color-based object recognition. Importantly, each object in the database has its image captured with respect to 72 rotational views (range from 0° to 355° , at 5° resolution) and 24 illumination angles, and 12 illuminant color temperatures (2175K to 3075K). In total the ALOI dataset comprises in excess of 100,000 images. We match objects by matching their underlying distributions. Simply, in color-distribution-based object recognition, if two images have similar underlying color distributions then this is taken as evidence of a possible object match.

We wished to evaluate how well homography-based chromaticity matching would support color-based object recognition. In our approach we use the ASIFT procedure to find distinctive points in the chromaticity distributions of all the images in the database. Then, for each testing condition (viewing angle, illumination angle or illumination color), we use a standard reference dataset and use the remaining images as query images. The standard reference capture condition is: frontal view (i.e., 0 degree of rotation), frontally lit and the light has a 3075K yellow color.

We compare the homography-based color object recognition (H) with Swain’s color indexing algorithm which uses an rg-chromaticity space [15] (rg), comprehensive color image normalization [40] (CN) and Gevers and Smeulders’ [41] $m1m2m3$ approach. The Comprehensive Normalization and $m1m2m3$ techniques are chosen both because they explicitly incorporate shading invariance into their formalisms and

TABLE 2
Average Match Percentile Results for All 1,000 Objects

Test	rg	m1m2m3	CN	H	Hybrid H
Viewing Angle	98.8	95.3	96.8	97.3	97.0
Illumination Angle	92.2	88.3	94.8	93.8	95.1
Illumination Color	87.3	98.2	98.0	97.2	98.2
Average	92.8	93.9	96.5	96.2	96.8

also they supersede early work (e.g., Comprehensive Normalization delivers better results than aligning chromaticity distributions by matching their means with a diagonal matrix [38], [39]).

The recommended bin sizes for rg and CN are 16×16 . The recommended bin size for $m1m2m3$ is 32×32 (note the distribution of these bins is non-linear with the partition found by a calibration procedure [41]). For our ASIFT-based chromaticity distribution match, a 320×320 histogram is used. For all methods we assess recognition performance using the simple Average Match Percentile (AMP). An average match percentile of 99 percent informs us that the correct matching answer is in the top 1 percent of matches.

We run the recognition experiment given all 1,000 objects where, relative to the standard viewing condition, either the viewing angle, illumination angle or illumination color are individually varied. As shown in Table 2, the pure homography-based chromaticity matching supports the second best performance for viewing angle and illumination angle tests. Of course, as alluded to in the color correction section, to match chromaticity histograms using homographies the underlying distributions need to have sufficient 2-dimensional structure. Visually, we found that $\sim 80\%$ objects in the ALOI dataset are comprised of objects with 3 or fewer reflectance colors and we know we need 4 points to solve for a homography. This said, the overall performance of the homography-based method is good i.e., it performs well for many objects with fewer than 3 colors present. Potentially, a good match remains possible because the effect of specular highlights and inter-reflection increase the dimensionality of the RGB signal recorded for a single object (see discussion at the end Section 4.1.2).

Intrinsic to our ASIFT based matching is the notion of color *structural complexity* which is simply defined as the number ASIFT points found for a given chromaticity distribution. We hypothesize that, empirically, we can find a threshold of the structural complexity to determine whether color homography indexing is likely to work (e.g., there is sufficient structure in the chromaticity distribution to find matches using homographies in a database).

To test this hypothesis, for the “Illumination Color” test, in Fig. 6, we plot the percentiles of the structural complexity (i.e., we rank images according to the number of ASIFT points found in their chromaticity distributions) and plot against the corresponding match percentiles. Clearly, there is a good correlation between the two measures and this indicates that as structural complexity increases so the homography-based measure delivers better object recognition. We find a similar positive correlation for the $m1m2m3$ and comprehensive normalization approaches.

In this example, assuming we wish to find the correct match in the top 1 percent of a database we need to use a




Fig. 6. Plot of color structure complexity percentile and matching percentile. The result is filtered by a moving Gaussian convolution kernel (width = 10 percent of total color structure complexity span, $\sigma = 1/6$ kernel width). The dashed line indicates the threshold, over which homography-based method works better.

threshold structural complexity as defined by the 75th percentile image. For our test the 75th percentile image and above have a structural complexity of at least 520 points. Importantly, notice, that in this top quantile range, the homography match is better than either the CN and m1m2m3 methods. Note that a threshold of 75 percent for the color structural complexity means that only 25 percent of the query images, can be homographically matched.

Suppose, we seek to match only the chromaticity distributions with 520 or more ASIFT points. We report the results in Table 3. The percentage shown are the number of chromaticity distributions in the database that have more than 520 ASIFT points. Note this is less than 25 percent (compared to our test in Fig. 6). Because under some illumination and viewing angles, there are fewer object colors so fewer ASIFT points. For images whose chromaticity distributions have sufficient structural complexity, the homography method delivers better results compared with competing techniques. In the last column of Table 2, we show the performance of a Hybrid method (Hybrid H): if the chromaticity distribution has more than 520 ASIFT points then homography-based matching is used. Otherwise, we adopt the CN method. The hybrid method delivers the best results for the illuminant colour and illumination angle test conditions and the best result overall.

4.3 Color Transfer

Color transfer is an image editing process that adjusts the colors of an input image I to match the palette of a target image J . Instead of adopting often computationally costly non-linear color mappings (e.g., [59]), we investigate whether color transfer can also be interpreted as a simple linear color homography mapping which re-visualizes an image with respect to real physical scene changes (e.g., from summer to autumn) and/or illumination. Our experiment is a continuation of our previous work [60] which is built on recent research [61]. The recent work [61] has demonstrated that it is possible to find a simple 3D similarity transform to linearly approximate some of the effects of color transfer.

There are three important reasons for visualizing color transfer as possibly being a color homography. First, if true, it would indicate a surprising result. Specifically, color transfer, though defined quite generally, tends to generate images which have a real-world physical interpretation. Second, the homography formulation is simpler compared

TABLE 3
Average Match Percentile Results for Query Image with Color Structure Complexity

Test	rg	m1m2m3	CN	H
Viewing Angle (12%)	98.6	94.8	96.3	97.4
Illumination Angle (19.4%)	91.5	90.1	94.8	96.3
Illumination Color (24.4%)	92.1	99.1	99.3	99.5
Average	93.4	94.3	96.4	97.7

Numbers in bracket indicate percentage of color-homography-compatible queries.

with some of the transfer algorithms. This simplicity means that spatial artifacts are less likely to be introduced (a fact borne out by our experiments). Lastly, the homography calculation is fast. As we will show below, we can calculate an expensive color transfer on a thumbnail, calculate the homography and then apply the result to a full resolution image.

We start with the color transfer output Image O and try to approximate its color transfer result by color homography. An example is shown in Fig. 2c where Image O' is the color transfer approximation result. In the discussion that follows it is useful to think of the image I (or O) being a simple $n \times 3$ matrix of RGB pixels, which can be reconstituted into an image grid for display purposes. Mathematically, we write

$$O' = H(I, O) \approx O. \quad (11)$$

To solve Equation (11) using our ALS algorithm, we respectively convert the images I and O images to the corresponding $n \times 3$ matrices A and B . Further, the ALS computed output DAH mapped back to an image is denoted O' . Here, the homography matrix H can be understood as a global chromaticity shift and distortion, the matrix D can be interpreted as shading change factors which simulate the change of surface reflectance or position of illuminant.

The visual results of color transfer approximations of four color transfer methods [18], [62], [63], [64] are shown in Fig. 7. In our experiments, the number of ALS iterations is set as 10. As can be seen, the global 3D Similarity mapping [61] does not perfectly reproduce the shading adjustments of color transfer. Our homography-based method offers a closer color transfer approximation.

We can also quantify this visual closeness by calculating the error between the color transfer approximation result and the original color transfer result. We adopt three error metrics:

PSNR (Peak Signal-to-Noise Ratio). PSNR is the ratio between the maximum possible value (power) of a signal and the power of distorting noise that affects the quality of its representation. Acceptable values for wireless image transmission quality loss are considered to be over 20 dB [65].

SSIM (Structural SIMilarity) [66]. SSIM is a model that considers image degradation as perceived change in structural information. SSIM can be used to assess the artifacts of color transfer. A SSIM value "1" indicates a perfect match.

HI (Histogram Intersection) of rgb Chromaticities. This is a score for measuring the similarity between two rgb

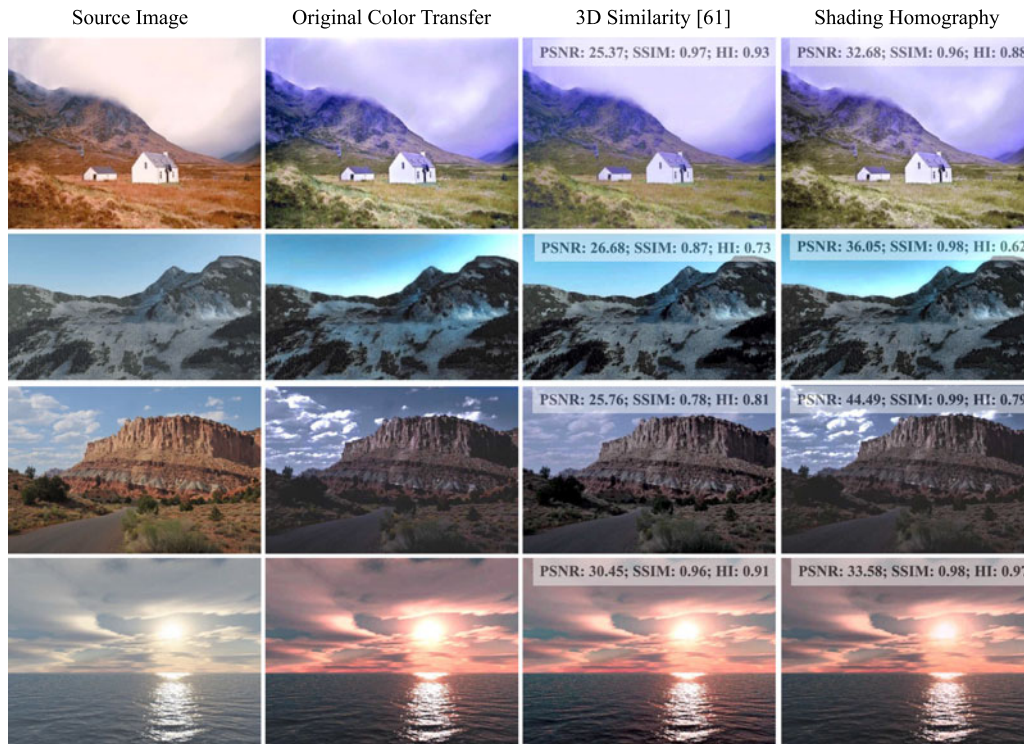


Fig. 7. Visual result of color transfer approximations (in the order of [18], [62], [63], [64]). The images in Column 4 (Homography) are generally more similar to those in Column 2 than those shown in Column 3 (3D Similarity).

chromaticity distributions, which is also shading-independent. A similarity score “1” indicates a perfect match.

We show the average color transfer approximation result in Table 4 (see the supplementary material, which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TPAMI.2017.2760833> for the complete table and their visual results). The quantitative test is based on 7 classic color transfer image pairs and 4 color transfer methods [18], [62], [63], [64]. Our color homography transfer produces the best results overall for the PSNR and SSIM tests. For the chromaticity mapping test (HI), our color homography approximation performs better for Pouli and Reinhard [63], [64].

We note that both the PSNR and SSIM metrics operate on the images output from our method. But, the histogram intersections are on the chromaticities of the original and matched images. The chromaticity mapping takes all RGBs and scales them so they sum to 1. This has a dramatic effect

especially for dark pixel values. Yet, these values are precisely those that are de-weighted in the ALS method (which operates by least-squares). They are also the pixels in the image that we cannot see.

To reduce the computational cost, it is also possible to estimate the shading homography (i.e., H and D) from the down-sampled images. In addition to the original ALS, we also upsample the smaller shading matrix returned by ALS (from the downsampled input images) by using Joint Bilateral Upsampling [67] (guided by the chromaticity-transferred result AH). We find that image down-sampling barely affects the color homography mapping quality. In Fig. 8, we show an example corresponding to the result of Row 1, Fig. 7. This can be useful because we can run a computationally costly color transfer algorithm on thumbnail images, extract the color transfer effect from the down-sampled results, and apply the extracted effect to a full-resolution source image. For some computationally costly color transfer algorithms (e.g., [18], [62], [63]), this offers a similar

TABLE 4
Errors of Color Transfer Approximation

	Ngu. [18]	Pitie [62]	Pouli [63]	Rein. [64]
PSNR (Peak Signal-to-Noise Ratio)				
3D Similarity [61]	26.85	26.04	26.92	28.49
Homography	34.05	31.91	37.28	36.31
SSIM (Structural Similarity)				
3D Similarity [61]	0.91	0.85	0.84	0.88
Homography	0.94	0.91	0.97	0.98
HI (Histogram Intersection)				
3D Similarity [61]	0.87	0.75	0.75	0.81
Homography	0.78	0.71	0.76	0.87




Fig. 8. Color transfer approximation for [18] from downsampled images. The sizes of source images and target images (i.e., I and J) are reduced by the corresponding factors. The original color transfer (MATLAB) takes about 3.630s. The 3 evaluation measurements and the total estimation time with down-sampling are shown over the images. As it is shown, image down-sampling barely affects the color transfer approximation quality. And, it takes less time to color transfer an image by using our down-sampling trick.



Fig. 9. Color transfer enhancement. A source image is color transferred by [63] with some noticeable imperfections (on the clouds). This issue is fixed by approximating the original color transfer effect with a *shading-smoothed* shading homography. Visually, the enhanced result also preserves more texture details compared with the result obtained by directly smoothing the original approximation RGB.

quality color transfer result but requires a reduced amount of processing time.

Color homography is also useful for color transfer enhancement. Fig. 9 shows an example where the original color transfer result contains some obvious artifacts. These artifacts are usually caused by sharp image gradient changes. Since a color homography transform barely modifies the original image gradient, the decomposed shading component absorbs most of the artifacts (e.g., original shading in Fig. 9). By spatially smoothing the original noisy shading component with a bilateral filter [68], we can remove the gradient artifacts in the shading component (e.g., smoothed shading in Fig. 9). The original color transfer result is improved by applying the modified shading homography (i.e., original color homography + smoothed shading).

When we proved the *Uniqueness* theorem—that the ALS procedure could be used to find the homography for the 4 point case—we introduced the idea of perturbing the data to find the unique homography. The data was perturbed when, albeit very rarely, the ALS procedure did not converge to the unique homography (see discussion after the proof of Theorem 2 in Section 3). For the color transfer problem we explored perturbing the input image by a random homography to investigate whether, overall, a better color transfer (a better fit) is found. We did not find that this was the case. This result is perhaps not surprising since for the 4 point case when the ALS procedure converged to the wrong answer the fitting error was, empirically, very good (typically less than 1 percent).

5 RAW VERSUS RENDERED IMAGES

That a color homography is a tool for describing color change across viewing conditions appears, at first glance, to hold only for RAW linear images. Empirically, our work on color transfer also indicates we can apply the homography method to non-linear (rendered) images. An example of RAW-to-rendered mapping approximation by our shading homography is shown in Fig. 10 where the approximation result is visually close to the rendered—actual camera output—image where the RAW and rendered pair is drawn from the Middlebury dataset [69]. But, why should a shading homography relate a RAW input to a camera output?

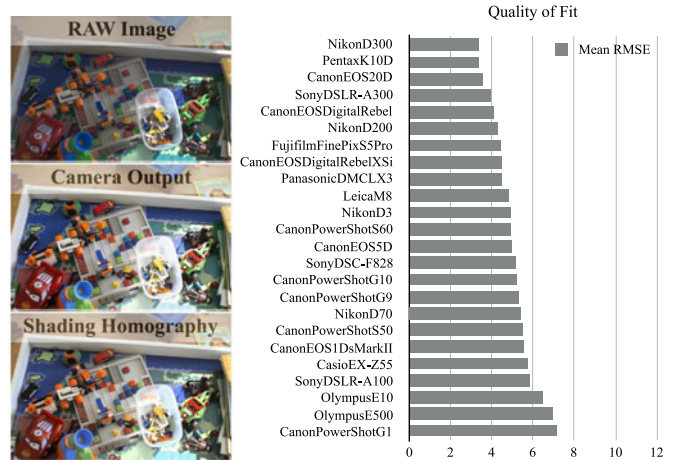


Fig. 10. RAW-to-JPEG Approximation. Left: The non-linear mapping from a RAW image to its rendered camera output image can be well approximated by a shading homography. Right: RMSE between estimation and ground truth using our shading homography model of the 24 RAW-capable cameras in the Middlebury dataset [69].

To a first approximation that an $n \times 3$ matrix of pixels, A , are captured under RAW (i.e., linear conditions) and that a corresponding sRGB image B , to a tolerable approximation [70]—for real cameras—can be represented as

$$B \approx f(AH), \quad (12)$$

where H is a 3×3 matrix and $f()$ is monotonically increasing function (a camera curve). This formalism is similar to the homography formalism. Indeed the form of f is far from arbitrary. If $\hat{b}_j^T = \underline{a}_j^T H$ (the color corrected j th pixel value), then one of the requirements for $f()$ is that

$$f(\hat{b}_j) \approx \gamma_j \hat{b}_j, \quad (13)$$

where γ_j is a scale factor. That is, the function $f()$ changes the magnitude of the vector but not its direction. If this was not the case then, as the same object is seen at different brightness levels in the same image, its color will change significantly. If we map RAW to rendered using a homography then

$$f(\hat{b}_j) = \gamma_j \hat{b}_j. \quad (14)$$

In [69], for a series of 24 cameras, several RAW and sRGB images (with intensities in $[0, 255]$, for different illuminants), of a 140 patch Macbeth SG color checker are captured. Chakrabarti et al. analyzed the RMSE error between the predicted sRGB images and those actually recorded where $f()$ in Equation (12) is a monotonically increasing 5th order polynomial. We repeat the same experiment for our homography-based approach and RMSE errors are reported in Fig. 10 where for convenience we sort the cameras from lowest to highest error. The reader is referred to Fig. 4 in [69] for comparison.

Broadly, the errors we found by homography fitting are on the same order as those found in the antecedent work [69]. However, some camera data is fit with lower error (e.g., the Canon Powershot P1 has a mean RMSE of 7.1 for the homography-based method but 12 using [69]). In contrast for the Olympus E500, the homography error is

larger (6.9 compared with 2). However, overall the range of fitting error of the homography based method and [69] is about the same (both have an overall RMSE of ~ 5).

6 CONCLUSION

In this paper, we demonstrated the surprising result that colors across a change in viewing condition (changing light color, shading and camera) are well related by a homography. We apply color homographies to color correction (mapping RAW RGBs to display counterparts), color object recognition and color transfer. Our homography-based color correction algorithm delivers improved color fidelity compared with the state-of-the-art. Matching chromaticity distributions using homographies delivers leading color-based object recognition. Re-interpreting color transfer as color homography mapping inspires a new direction for natural color transfer algorithm development.

APPENDIX A

NUMERICAL ALS CONVERGENCE TO THE WRONG ANSWER

The matrices A and B defined below are—when rows are interpreted as “rays” as a consequence of the Planar Homography theorem [2]—a homography apart. Equivalently, there exists a 4×4 diagonal matrix D and a 3×3 linear transform H such that $B = DAH$. Suppose,

$$A = \begin{bmatrix} 5 & 5 & 8 \\ 6 & 9 & 2 \\ 1 & 7 & 3 \\ 4 & 7 & 10 \end{bmatrix} \quad B = \begin{bmatrix} 78 & 76 & 96 \\ 107 & 82 & 71 \\ 531 & 270 & 423 \\ 87 & 74 & 111 \end{bmatrix}. \quad (15)$$

Let us find a matrix $K = DAH \approx B$ using ALS to find D and H . On convergence, to 2 decimal places, we find K

$$K = \begin{bmatrix} 78.31 & 76.23 & 96.41 \\ 109.33 & 56.28 & 85.89 \\ 532.63 & 276.44 & 421.57 \\ 87.37 & 73.94 & 111.62 \end{bmatrix}. \quad (16)$$

The % error ($\|K - B\|/\|B\|$) between K and B is almost 4 percent. According to our Theorem 2, when we converge to the wrong answer the null-space vector of $B - K$ should be sparse (1, 2, or 3 of the elements should be non-zero but not all 4). For this example, the unit length vector \underline{v} which is orthogonal to $A - B$, again to 2 decimal places is: $\underline{v} = [-0.74 \ 0 \ 0 \ 0.68]^T$ (as a numerical check, $\underline{v}^T(A - B) = [0 \ 0 \ 0]$, to 2 decimal places).

APPENDIX B

CLOSED-FORM HOMOGRAPHY

Here we present a new closed form solution for solving for the Homography matrix H that maps 4 points in one view to corresponding points in a second view. Let us denote the paired matching points in the 4×2 matrices A and B (the x and y coordinates are in the first two columns). Moving to homogeneous coordinates—we add a vector of 1’s to each matrix—to make 4×3 matrices A and B . Let the operator $\text{diag}(\underline{v})$ return a diagonal matrix with components of \underline{v}

along its diagonal. Let $M_{1:3}$ denote the first 3 rows of a matrix and that M_k is the k th row vector.

Theorem 3 (Closed-Form Homography). *The Homography matrix H relating the 4×3 matrices A and B (for the 4×2 matched A and B , where no three points in either matrix are collinear) can be written in closed-form as: $H = [A_{1:3}]^{-1} \text{diag}(A_4[A_{1:3}]^{-1})^{-1} \text{diag}(B_4[B_{1:3}]^{-1})B_{1:3}$.*

Proof. From the non-collinearity assumption $A_{1:3}$ and $B_{1:3}$ are full rank 3×3 invertible matrices. We define X :

$$X = A[A_{1:3}]^{-1} = \begin{bmatrix} \mathcal{I} \\ A_4[A_{1:3}]^{-1} \end{bmatrix}$$

$$\begin{aligned} & X \text{diag}(A_4[A_{1:3}]^{-1})^{-1} \text{diag}(B_4[B_{1:3}]^{-1}) \\ &= \begin{bmatrix} \text{diag}(A_4[A_{1:3}]^{-1})^{-1} \text{diag}(B_4[B_{1:3}]^{-1}) \\ B_4[B_{1:3}]^{-1} \end{bmatrix}. \end{aligned}$$

Let D denote a 4×4 diagonal matrix such that the upper left 3×3 sub-matrix (first 3 rows and columns) is

$$\text{diag}(B_4[B_{1:3}]^{-1})^{-1} \text{diag}(A_4[A_{1:3}]^{-1}),$$

and the 4th component along the diagonal of D is 1. We define Y

$$\begin{aligned} Y &= DX \text{diag}(A_4[A_{1:3}]^{-1})^{-1} \text{diag}(B_4[B_{1:3}]^{-1}) \\ &= \begin{bmatrix} \mathcal{I} \\ B_4[B_{1:3}]^{-1} \end{bmatrix}, \end{aligned}$$

which follows that $YB_{1:3} = B$ i.e., the matrix B . Substitute for Y and X and post-multiply by $B_{1:3}$, we can write

$$A[A_{1:3}]^{-1} \text{diag}(A_4[A_{1:3}]^{-1})^{-1} \text{diag}(B_4[B_{1:3}]^{-1})B_{1:3} = DB.$$

Dividing by the first two columns of DB by the third we end up with the original point set B . \square

ACKNOWLEDGMENTS

This work was supported by EPSRC Grant EP/M001768/1. We also acknowledge Prof. Roberto Cipolla (University of Cambridge) and Prof. Mark Drew (Simon Fraser University) for their useful discussion on this work. We would like to thank EPSRC for funding this project.

REFERENCES

- [1] M. Brown and D. G. Lowe, “Automatic panoramic image stitching using invariant features,” *Int. J. Comput. Vis.*, vol. 74, no. 1, pp. 59–73, 2007.
- [2] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge, U.K.: Cambridge University Press, ISBN: 0521540518, 2004.
- [3] E. Land, “The retinex theory of color vision,” *Sci. Amer.*, vol. 237, no. 6, pp. 108–129, 1977.
- [4] D. Forsyth, “A novel algorithm for color constancy,” *Int. J. Comput. Vis.*, vol. 5, pp. 5–36, 1990.
- [5] P. Hubel, J. Holm, and G. Finlayson, “Illuminant estimation and colour correction,” in *Proc. Colour Multimedia Conf.*, 1998, pp. 97–105.
- [6] M. J. Vrhel and H. J. Trussell, “The mathematics of color calibration,” in *Proc. Int. Conf. Image Process.*, 1998, pp. 181–185.
- [7] M. Drew and B. V. Funt, “Natural metamers,” *Graph. Models Image Process.*, vol. 56, pp. 139–151, 1992.

- [8] B. Wandell, "The synthesis and analysis of color images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 9, no. 1, pp. 2–13, Jan. 1987.
- [9] R. Lenz, L. V. Tran, and P. Meer, "Moment based normalization of color images," in *Proc. IEEE 3rd Workshop Multimedia Signal Process.*, 1999, pp. 103–108.
- [10] M. Anderson, R. Motta, S. Chandrasekar, and M. Stokes, "Proposal for a standard default color space for the internet: sRGB," in *Proc. Color Imag. Conf.*, 1995, pp. 238–245.
- [11] S. Bianco and C. Cusano, "Color target localization under varying illumination conditions," in *Computational Color Imaging*. Berlin Heidelberg, Germany: Springer, 2011, vol. 6626, pp. 245–255.
- [12] A. Minagawa, Y. Katsuyama, H. Takebe, and Y. Hotta, "A color chart detection method for automatic color correction," in *Proc. Int. Conf. Pattern Recog.*, Nov 2012, pp. 1912–1915.
- [13] B. Funt and P. Bastani, "Irradiance-independent camera color calibration," *Color Res. Appl.*, vol. 39, no. 6, pp. 540–548, 2014.
- [14] G. D. Finlayson, M. Mohammadzadeh Darrodi, and M. Mackiewicz, "The alternating least squares technique for non-uniform intensity color correction," *Color Res. Appl.*, vol. 40, no. 3, pp. 232–242, 2015.
- [15] M. Swain and D. Ballard, "Color indexing," *Int. J. Comput. Vis.*, vol. 7, no. 11, pp. 11–32, 1991.
- [16] W. Niblack and R. Barber, "The QBIC project: Querying images by content using color, texture and shape," in *Storage and Retrieval for Image and Video Databases I SPIE Proceedings Series*. New York, NY, USA: IBM Research Division, 1993.
- [17] B. Funt and G. Finlayson, "Color constant color indexing," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 17, no. 5, pp. 522–529, May 1995.
- [18] R. M. H. Nguyen, S. J. Kim, and M. S. Brown, "Illuminant aware gamut-based color transfer," *Comput. Graph. Forum*, vol. 33, no. 7, pp. 319–328, Oct. 2014.
- [19] J. Geusebroek, G. J. Burghouts, and A. W. M. Smeulders, "The amsterdam library of object images," *Int. J. Comput. Vis.*, vol. 61, no. 1, pp. 103–112, 2005.
- [20] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. J. Comput. Vis.*, vol. 60, no. 2, pp. 91–110, 2004.
- [21] G. Yu and J.-M. Morel, "ASIFT: An algorithm for fully affine invariant comparison," *Image Process. Line*, vol. 1, pp. 11–38, 2011.
- [22] Z. Zhang, "A flexible new technique for camera calibration," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 11, pp. 1330–1334, Nov. 2000.
- [23] S. Gefen, Y. Fan, L. Bertrand, and J. Nissarov, "Symmetry-based 3D brain reconstruction," in *Proc. Int. Symp. Biomed. Imag.*, 2004, pp. 744–747.
- [24] C. T. Loop and Z. Zhang, "Computing rectifying homographies for stereo vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 1999, pp. 1125–1131.
- [25] P. H. S. Torr and A. Zisserman, "Feature based methods for structure and motion estimation," in *Proc. Int. Workshop Vis. Algorithms: Theory Practice*, 2000, pp. 278–294.
- [26] L. Maloney, "Evaluation of linear models of surface spectral reflectance with small numbers of parameters," *J. Opt. Soc. America A*, vol. 3, pp. 1673–1683, 1986.
- [27] J. Parkkinen, J. Hallikainen, and T. Jaaskelainen, "Characteristic spectra of munsell colors," *J. Opt. Soc. America A*, vol. 6, pp. 318–322, 1989.
- [28] M. Vrhel and H. Trussel, "Color correction using principal components," *Color Res. Appl.*, vol. 17, pp. 328–338, 1992.
- [29] G. Wyszecki and W. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulas*. New York, NY, USA: Wiley, 1982.
- [30] D. Marimont and B. Wandell, "Linear models of surface and illuminant spectra," *J. Opt. Soc. America A*, vol. 9, no. 11, pp. 1905–1913, 1992.
- [31] B. Funt and H. Jiang, "Nondiagonal color correction," in *Proc. Int. Conf. Image Process.*, Sep. 2003, vol. 1, pp. 1–481–4.
- [32] G. D. Finlayson, B. V. Funt, and H. Jiang, "Predicting cone quantum catches under illumination change," in *Proc. Color Imag. Conf.*, 2003, pp. 170–174.
- [33] A. Gijsenij, R. Lu, and T. Gevers, "Color constancy for multiple light sources," *IEEE Trans. Image Process.*, vol. 21, no. 2, pp. 697–707, 2012.
- [34] I. Boyadzhiev, K. Bala, S. Paris, and F. Durand, "User-guided white balance for mixed lighting conditions," *ACM Trans. Graph.*, vol. 31, no. 6, 2012, Art. no. 200.
- [35] S. Beigpour, C. Riess, J. Van de Weijer, and E. Angelopoulou, "Multi-illuminant estimation with conditional random fields," *IEEE Trans. Image Process.*, vol. 23, no. 1, pp. 83–96, 2014.
- [36] R. Kondepudy and G. Healey, "Use of invariants for recognition of 3-dimensional color textures," *J. Opt. Soc. America A*, vol. 11, no. 11, pp. 3037–3049, Nov. 1994.
- [37] J. Berens and G. Finlayson, "Log-opponent chromaticity coding of colour space," in *Proc. Int. Conf. Pattern Recog.*, 2000, vol. 1, pp. 206–211.
- [38] G. Finlayson, M. Drew, and B. Funt, "Spectral sharpening: Sensor transformations for improved color constancy," *J. Opt. Soc. America A*, vol. 11, no. 5, pp. 1553–1563, May 1994.
- [39] H. Y. Chong, S. J. Gortler, and T. Zickler, "The von Kries hypothesis and a basis for color constancy," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2007, pp. 1–8.
- [40] G. Finlayson, B. Schiele, and J. Crowley, "Comprehensive colour image normalization," in *Proc. Eur. Conf. Comput. Vis.*, 1998, pp. 475–489.
- [41] T. Gevers and A. W. M. Smeulders, "Color-based object recognition," *Int. J. Pattern Recognit. Artif. Intell.*, vol. 32, no. 3, pp. 453–464, 1999.
- [42] D. Berwick and S. Lee, "A chromaticity space for specularly-, illumination color- and illumination pose-invariant 3-D object recognition," in *Proc. IEEE Int. Conf. Comput. Vis.*, 1998, pp. 165–170.
- [43] P. Hubel, G. Finlayson, and S. Hordley, "White point estimation using color by convolution," U. S. Patent No. 20030194125, Oct. 16, 2003.
- [44] J. T. Barron, "Convolutional color constancy," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 379–387.
- [45] Y. Ohno and J. Hardis, "Four-color matrix method for correction of tristimulus colorimeters," in *Proc. Color Imag. Conf.*, 1997, pp. 301–305.
- [46] G. Finlayson, H. Gong, and R. Fisher, "Color homography," in *Progress Colour Stud.*, 2016.
- [47] G. D. Finlayson, H. Gong, and R. B. Fisher, "Color homography color correction," in *Proc. Color Imag. Conf.*, 2016, pp. 310–314.
- [48] T. Zhang and G. H. Golub, "Rank-one approximation to high order tensors," *SIAM J. Matrix Anal. Appl.*, vol. 23, no. 2, pp. 534–550, 2001.
- [49] J. Matas, O. Chum, M. Urban, and T. Pajdla, "Robust wide baseline stereo from maximally stable extremal regions," in *Proc. Brit. Mach. Vis. Conf.*, 2002, pp. 36.1–36.10, doi:10.5244/C.16.36.
- [50] H. Bay, T. Tuytelaars, and L. Van Gool, "SURF: Speeded up robust features," in *Proc. Eur. Conf. Comput. Vis.* Berlin, Germany: Springer, 2006, pp. 404–417.
- [51] C. Harris and M. Stephens, "A combined corner and edge detector," in *Alvey Vision Conference*, New York, NY, USA: Cite-seer, 1988, vol. 15, Art. no. 50.
- [52] C. S. McCamy, H. Marcus, and J. G. Davidson, et al., "A color-rendition chart," *J. Appl. Photographic Eng.*, vol. 2, no. 3, pp. 95–99, 1976.
- [53] G. D. Finlayson, M. Mackiewicz, and A. Hurlbert, "Color correction using root-polynomial regression," *IEEE Trans. Image Process.*, vol. 24, no. 5, pp. 1460–1470, 2015.
- [54] C. L. Novak and S. A. Shafer, "Method for estimating scene parameters from color histograms," *J. Opt. Soc. America*, vol. 11, no. 11, pp. 3020–3036, 1994.
- [55] G. Healey, "Estimating spectral reflectance using highlights," *Image Vis. Comput.*, vol. 9, no. 5, pp. 333–337, 1991.
- [56] B. V. Funt and M. S. Drew, "Color space analysis of mutual illumination," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 15, no. 12, pp. 1319–1326, Dec. 1993.
- [57] R. Bajcsy, S. W. Lee, and A. Leonardis, "Color image segmentation with detection of highlights and local illumination induced by inter-reflections," in *Proc. Int. Conf. Pattern Recog.*, 1990, vol. 1, pp. 785–790.
- [58] S. Tominaga and T. Okayama, "A method for analyzing color images containing highlight and interreflection effects," *Syst. Comput. Jpn.*, vol. 28, no. 13, pp. 8–18, 1997.
- [59] Y. Hwang, J.-Y. Lee, I. S. Kweon, and S. J. Kim, "Color transfer using probabilistic moving least squares," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2014, pp. 3342–3349.
- [60] H. Gong, G. D. Finlayson, and R. B. Fisher, "Recoding color transfer as a color homography," in *Proc. Brit. Mach. Vision Conf.*, 2016.
- [61] F. Pitié and A. Kokaram, "The linear Monge-Kantorovitch linear colour mapping for example-based colour transfer," in *Proc. Eur. Conf. Vis. Media Prod.*, 2007, pp. 1–9.
- [62] F. Pitié, A. C. Kokaram, and R. Dahiya, "Automated colour grading using colour distribution transfer," *Comput. Vis. Image Understanding*, vol. 107, no. 1–2, pp. 123–137, Jul. 2007.

- [63] T. Pouli and E. Reinhard, "Progressive histogram reshaping for creative color transfer and tone reproduction," in *Proc. Int. Symp. Non-Photorealistic Animation and Rendering*. New York, NY, USA: ACM, 2010, pp. 81–90.
- [64] E. Reinhard, M. Ashikhmin, B. Gooch, and P. Shirley, "Color transfer between images," *IEEE Comput. Graph. & Appl.*, vol. 21, no. 5, pp. 34–41, Sep. 2001.
- [65] N. Thomos, N. V. Boulgouris, and M. G. Strintzis, "Optimized transmission of jpeg2000 streams over wireless channels," *IEEE Trans. Image Process.*, vol. 15, no. 1, pp. 54–67, 2006.
- [66] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, 2004.
- [67] J. Kopf, M. F. Cohen, D. Lischinski, and M. Uyttendaele, "Joint bilateral upsampling," *ACM Trans. Graph.*, vol. 26, no. 3, 2007, Art. no. 96.
- [68] S. Paris and F. Durand, "A fast approximation of the bilateral filter using a signal processing approach," in *Proc. Eur. Conf. Comput. Vis.* Berlin, Germany: Springer, 2006, pp. 568–580.
- [69] A. Chakrabarti, D. Scharstein, and T. Zickler, "An empirical camera model for internet color vision," in *Proc. Brit. Mach. Vis. Conf.*, 2009, vol. 1, no. 2, Art. no. 4.
- [70] S. J. Kim, H. T. Lin, Z. Lu, S. Susstrunk, S. Lin, and M. S. BrownHai, "A new in-camera imaging model for color computer vision and its application," in *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 12, pp. 2289–2302, Dec. 2012.



Graham D. Finlayson is currently a professor of computing sciences with the University of East Anglia. His research interests span color, physics-based computer vision, image processing, and the engineering required to embed technology in devices. He is a fellow of the Royal Photographic Society, the Society for Imaging Science and Technology, and the Institution for Engineering Technology. He is a member of the IEEE.



Han Gong received the PhD degree from the University of Bath. He is currently a senior research associate with the University of East Anglia. His research interests include computer vision and computational photography.



Robert B. Fisher received the BS (Hons.) degree in mathematics from the California Institute of Technology, the MS degree in computer science from Stanford University, and the PhD degree from the University of Edinburgh, in 1974, 1978, and 1987. He has been an academic in the School of Informatics, University of Edinburgh since 1984 and a full professor since 2003.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.