https://doi.org/10.2352/CIC.2024.32.1.2
This work is licensed under the Creative Commons
Attribution 4.0 International License.

Colour Meets Geometry in Colorimetric Filter Design

Graham D. Finlayson*[1,2] *and Ivar Farup*[2], [1] *School of Computing Science, University of East Anglia, UK;* [2] *Colourlab, Department of Computer Science, NTNU, Norway

Abstract

Panoramas are formed by stitching together two or more images of a scene viewed from different positions. Part of the solution to this stitching problem is 'solving for the homography': where the detail in one image is geometrically warped so it appears in the coordinate frame of another. In this paper, we view the spectral loci for a given camera and the human visual system (i.e. their respective chromaticity diagrams) as two pictures of the same 'scene' and warp one to the other by finding the best homography. When this geometric distortion renders the two loci to be identical then there exists a unique colour filter (that falls gracefully from the derivation without further calculation) which makes the camera colorimetric (the camera+filter measures RGBs that are exactly linearly related to XYZs). When the best homography is not exact the filter derived by this method still makes cameras approximately colorimetric. Experiments validate our method.

Introduction

One of the central tenets of colour science is that if the tristimulus response – the three numbers found by integrating a spectrum with the three XYZ colour-matching functions [1] (here and throughout this paper we use the 1931 standard observer curves, sometimes denoted CMFs) – for two different spectra are the same then they will also look the same to the average human observer. When two spectra match, they are called metamers. For various reasons – not least the limited number of dyes that can be deposited on imaging sensors – the spectral sensitivities of the cameras are not the same as or linearly transformable to the human vision colour-matching functions. The import of this is that there exist spectral pairs that look identical to the camera and different to the eye and vice versa. Consequently, no mathematical means can exist to discount this metamer mismatching exactly [2] though some non-linear functions can mitigate the problem e.g. [3, 4] (i.e. they deliver lower colour error than a linear transform).

Of course, many of the colours we see on our devices are perfectly good reproductions of our colour percepts. Yet, there are always colours that are incorrectly reproduced. Moreover, in the domains of industrial inspection and medical applications the differences between the desired XYZs and that which can be approximated measured by a camera are significant, e.g. [5, 6]. One solution to the metamerism problem is to make more than 3 measurements (e.g. use a 6 sensor camera [7]) or, indeed, to capture the spectral radiance at each point in the scene [8]. In the latter scenario, the correct tristimulus response is easily found by integrating with the XYZ CMFs. Unfortunately, spectral imaging is often not a viable solution to the metamerism problem. Spectral imagers are expensive, are sensitive measurement devices, require long integration times and are, generally, unsuitable for use in the field.

Recently, it has been shown that is possible to design a special coloured filter – with a particular transmittance profile – that, when placed in front of a typical RGB colour camera makes that camera almost colorimetric, i.e. the filtered RGB camera measurements are almost within a linear 3×3 matrix correction from corresponding XYZ tristimulus values [9, 10]. Figure 1 illustrates how the method works. In, respectively, panels (a) and (b) a Canon EOS 50D camera and its spectral sensitivities [11] are shown. In panel (d) we plot the XYZ colour-matching functions [1] and in panel (e) the simple least-squares fit of the camera's sensitivities to best match these CMFs. If the Canon 50D camera was colorimetric, then the fitted and actual XYZ curves would be the same. The optimal derived filter (solved for according to [9]) for this camera is shown in panel (c)). This filter is most transmissive in the long and short wavelengths. As a visualization, a colored filter is shown to be in front of the 50D camera in panel (a)). Finally, in panel (f) we show the least-squares fit of the camera sensitivities multiplied by this filter to the XYZ curves. These new fitted filter curves are almost the same as the XYZ colour-matching curves, indicating the filter camera is almost colorimetric.

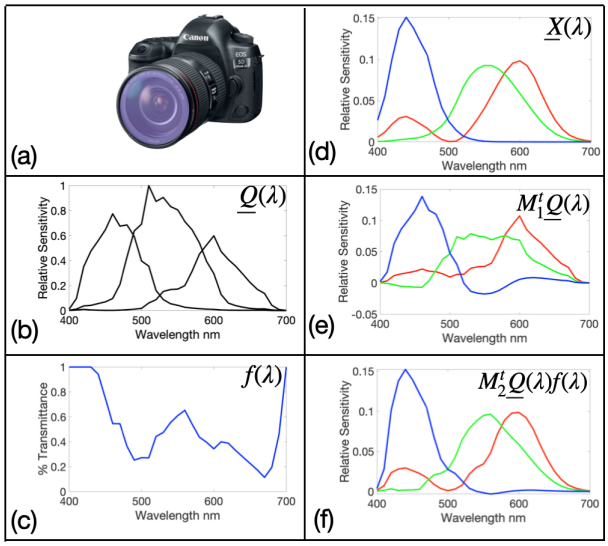

Figure 1: Respectively (a) through (d) show a Canon 50D camera, its spectral sensitivities, a transmissive colour filter and the XYZ CMFs. Panel (e) shows the best fit of the Canon sensitivities (M_1^t) to the XYZ CMFs. In panel (f) the filtered Canon sensitivities are fitted to the XYZs (M_2^t).

When an optimal correction filter is deployed, the colour errors between actual and estimated XYZ tristimulus values can in principle be reduced by 2/3 or more [10] compared to the colour correction without a filter. Placing a coloured filter in front of the

light source (rather than the camera) will, for many capture scenarios, result in the same filtered RGB measurements. In [12], a multispectral illuminator is used to simulate the optimal filter idea and significant reductions in colour-measurement errors are obtained, validating the filtering idea.

One of the disadvantages of the prior art filter design work is that the associated optimization is not closed form and the *best*, iteratively, derived filter can vary depending on an initialisation condition for the optimization. Consequently, in [10], thousands of initial conditions – starting points for the filter optimization – were required to find the best filter overall. Another potential weakness of the prior-art is that colour error in measured in tristimulus space rather than in more perceptually uniform colour spaces (e.g. CIELAB and CIELUV [1]).

In this paper, we present a surprising result. We show that to find the best filter to make a camera colorimetric we need to *warp* – i.e. non-linearly spatially transform – the chromaticity diagram of the camera to best match the corresponding xy or $u'v'$ chromaticity diagram. The image warp is exactly a homography, which is a non-linear transform more commonly found in geometric computer vision where images are geometrically mapped and then stitched together (e.g. in panorama formation [13] or epipolar geometry in stereo vision [14]). Once the globally optimal warp is found (here there is no need for careful initialisation), the best colorimetric filter – which we call the *geometric warp filter* – falls out of the optimisation without further optimization.

In experiments, we, indirectly, find the filter that makes a wide variety of cameras approximately colorimetric by geometrically warping their respective chromaticity diagrams. We then evaluate how well these geometric warp filters make the camera colorimetric, in general (i.e. for a range of typical lights and surfaces and not only the colours on the spectral locus). A hybrid method is also considered, where the geometric warp filter is refined by the prior art optimisation [9].

Background

Geometric image warping: In geometric image processing, one image is often warped in order to match another. Figure 1, illustrates this idea with two pictures taken – by one of the authors – from different positions of the same scene. Panel (a) shows a picture of a lake with vegetation in the background. From a different position, we see the same scene in panel (b). Note there is content in common and content that differs, and as such, if we can merge the two images we might recover a larger format image with all the available detail. Four matching pairs of points are shown, linked by the dotted lines (though many more matching pairs can be found). Given these point correspondences, we can warp the image (b) to be in the same coordinate frame as (a). Placing the warped image beside image (a) allows us to make a panorama images (where we can see all the detail from both images). When we warp image (a) then the result is spatially distorted, which is why the combined image (c) appears to have a non-square shape. In black, we frame the part of the combined image that has the usual rectangular shape and combines information between both the images.

So, how do we map one image onto another? Let the corresponding pairs of 4 points in the images (a) and (b) be respectively denoted by the coordinate pairs (x_i, y_i) and (u_i, v_i) $(i = \{1,2,3,4\})$. In geometric warping, it is often useful to invert

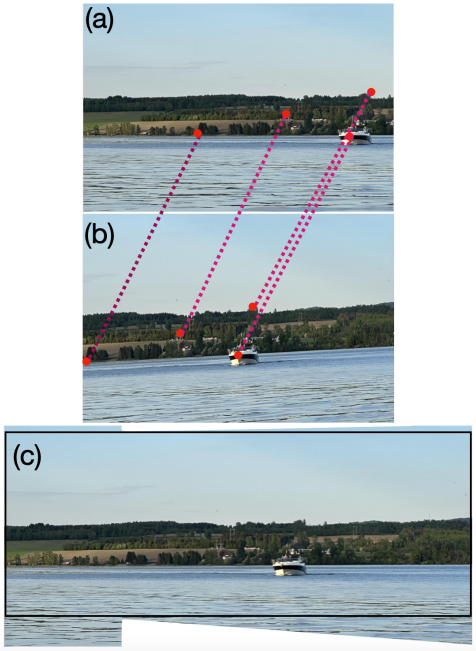

Figure 2: Image (b) is warped into the coordinate frame of image (a) [by matching the corresponding points linked by the dotted lines] and then combined with (a) to form the panorama image in (c)

the image formation (that maps points in 3D to 2D pixel location counterparts). To achieve this, points are mapped to 3-D corresponding homogeneous vectors: $\underline{p}_i = [x_i\ y_i\ 1]^t$ and $\underline{q}_i = [u_i\ v_i\ 1]^t$ (here and throughout this paper ^{t} denotes the transpose operator). The intuition here is that we represent the points in the image as having a *depth* of 1.

Since we only have a 2D image we cannot recover the 3D geometry, but the actual 3D point should lie on the *ray* (all points that are scalings from their homogeneous coordinate vectors). In many circumstances, given an image of a 3D scene, we can predict the scene captured by a camera in a different position without having to understand the depth of points in a scene. Rather, it suffices to multiply the rays by a 3×3 *homography* matrix to generate new rays. Then we map rays to points (we simulate 3D to 2D image formation). To a good approximation, if we take two pictures from a camera at the same position (though different orientations) or view scenes from sufficiently far away (and can assume that the observed scene points are at a similar depth) then this (i) point-to-rays, (ii) apply 3×3 matrix and then (iii) reproject methodology is the correct way to warp one image into the coordinate frame of another. A useful property of the homography mapping – that has to be true given our visual example – is that the resulting homography between the 2D coordinates will map straight lines onto straight lines. For more details, see [14].

Mapping Colour: We recapitulate colour image formation at a point as

$$\underline{\rho} = \int_{\omega} \underline{Q}(\lambda) C(\lambda) d\lambda \tag{1}$$

where respectively, $\underline{\rho}$, ω, $\underline{Q}(\lambda)$ and $C(\lambda)$ represent the RGB sensor response, the visible spectrum (here, assumed to be 400 to

700 nanometres), the 3-spectral sensitivities of the camera and the spectrum of light entering the camera. Often in colour measurement, we are interested in mapping $\underline{\rho}$ to a corresponding XYZ tristimulus $\underline{\chi}$ (since the latter is referenced to how *we* see colours in the world). Denoting $\underline{X}(\lambda)$ as the CIE colour-matching curves [1] (we use the 2 degree standard observer throughout), $\underline{\chi}$ is calculated as

$$\underline{\chi} = \int_{\omega} \underline{X}(\lambda) C(\lambda) d\lambda \tag{2}$$

Often, the mapping from RGB to XYZ is a linear 3×3 matrix M such that

$$\underline{\rho}^t M \approx \underline{\chi}^t \tag{3}$$

here ^{t} denotes the transpose operator. For the purposes of derivations set forth in the next section , it is useful for M to be a post-multiplying correction (hence we transpose $\underline{\rho}$ and $\underline{\chi}$). In Figure 1 – where we use vector functions – we, instead, transpose the matrix). The *meaning* of M is the same throughout this paper.

Now let us relate (r,g) chromaticities with corresponding (x,y) counterparts where we calculate *rg* and *xy* chromaticities as:

$$\begin{aligned} (r,g) &= \left(\frac{\rho_R}{\rho_R+\rho_G+\rho_B}, \frac{\rho_G}{\rho_R+\rho_G+\rho_B} \right) \\ (x,y) &= \left(\frac{\chi_1}{\chi_1+\chi_2+\chi_3}, \frac{\chi_2}{\chi_1+\chi_2+\chi_3} \right) \end{aligned} \tag{4}$$

Clearly, the vector $[r\ g\ 1]^t$ (point to ray) is proportional to the vector $[\rho_R\ \rho_G\ \rho_R+\rho_G+\rho_B]^t$. If $\underline{\rho}^t$ is mapped to $\underline{\chi}^t$ by M (Equation 3), then there must exist a matrix M' which maps $\underline{\rho}' = [r\ g\ 1]^t$ to $\underline{\chi}' = k[x\ y\ 1]^t$ (where k is a scalar). Let's define

$$C = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \tag{5}$$

It follows that

$$M' = C^{-1} M C \tag{6}$$

and

$$[\underline{\chi}']^t = [\underline{\rho}']^t M' \tag{7}$$

Mapping rays to points to calculate chromaticities $(\chi'_1/\chi'_3, \chi'_2/\chi'_3) = (x', y')$. The mapping from *rg* to *xy* chromaticities is the same as the homography image warping discussed previously. We map *rg* chromaticities to homogeneous coordinates (points to rays), apply a 3×3 colour homography matrix (M') and reproject to recover corresponding *xy* chromaticities. There is a small difference with the geometric case, however. Specifically, we only expect the colour mapping to be approximate. In the geometric case, solving for the homography can precisely model the actual geometric distortion.

Another important difference to the geometric mapping case is that in solving for the *colour homography*, we generally have access to, and are interested in, the 3rd dimension (depth in the geometric analogue). One consequence of this is that the solution strategy for colour homographies is, hitherto, posed quite differently from its geometric analogue.

Solving for a colour homography:. Suppose we have a pair of $n \times 3$ matrices R and X containing, respectively, RGBs and XYZs induced by n radiance spectra. In colour homography, we seek to jointly find a 3×3 matrix M and an $n \times n$ diagonal matrix D that minimizes the Frobenius norm ($||.||_F$):

$$\min_{D,M} ||DRM - X||_F \tag{8}$$

The closed-form solution: in the case where $n = 4$, Equation 8 can be solved exactly, e.g. see [10]. The intuition on why this result is true is that for $n = 4$ the matrix R has 12 measurements. The diagonal matrix D and matrix M have respectively 4 and 9 components, so there are 13 unknowns. But, because there is a scaling indeterminacy between D and M (we can multiply one by k and the other by $1/k$ and get the same result) there are only 12 unknowns to solve for. When $n > 4$ there are more knowns than unknowns and an iterative algorithm has to be used to find the best fit of the data.

The iterative solution: The alternating least-squares algorithm is summarised in Algorithm 1.

Algorithm 1 Alternating Least-Squares

$D_0 = I_{n\times n}$, $M_0 = I_{3\times 3}$, $stop =$ **false**, $i = 0$
while $\sim stop$ **do**
  $i = i + 1$
  $\min_{M_i} ||D_{i-1} R M_i - X||_F$
  $\min_{D_i} ||D_i R M_i - X||_F$
  $stop = ||D_{i-1} R M_{i-1} - D_i R M_i||_F < \varepsilon$
end while

where ε is a small non-negative real number (set by the user). $I_{3\times 3}$ and $I_{31\times 31}$ respectively denote the 3×3 and 31×31 identity matrices.

There are 4 types of colour homography reported in the literature. When R and X are, respectively, the spectral sensitivities of the camera and the XYZ colour-matching functions (we discretely sample across the wavelength domain), then, when we solve for the optimisation, the diagonal matrix D can be interpreted as a physical filter, $f(\lambda)$ which, when placed in front of the camera, results in sensitivities well-related by a linear transform [9]. Returning to Figure 1, in panel (e) we see the best least-squares fit of a Canon 50D camera to the XYZ colour-matching functions by a linear transform. The fit is rather poor, and this indicates that this camera will see differently – at least to some stimuli – to how we see these colours ourselves. In panel (f) we see the filtered sensitivities of the camera fitted to match the XYZ CMFs. The fit is excellent, indicating that the camera+filter is quite colorimetric. In this figure – as makes sense in the continuous domain – we show the two colour transforms as M_1^t and M_2^t (as here they premultiply).

In the second type of homography, we simply have RGB response and XYZ tristimuli data. Here, the matrix D performs a shading correction [10, 15]. In the 3rd type of colour homography, as well as having measured camera sensitivities and CMFs we also have measured spectral data [10]. One can then formulate a more complex optimization statement than Equation 8 and solve for the optimal filter given the measured spectral data. Finally, it is possible to find the best homography mapping *rg* image chromaticities to corresponding *xy* chromaticities [16]. All four of these applications of colour homography are solved using alternating least-squares. The underpinning theory of colour homography has been developed and extended in several interesting ways. This includes considering the role of noise in filter design [17], the minimum transmissivity of the filter [18, 19] and evaluating commercially available filters [20]. Colour homography theory has been applied to colour correction [16], to colour transfer [21] and image indexing [10] among other problems.

Finding the Geometric Warp Filter

In this section, we will consider the colour homography mapping between the *rg* spectral loci of a camera system to $u'v'$ chromaticity diagram. Figure 3 illustrates our method.

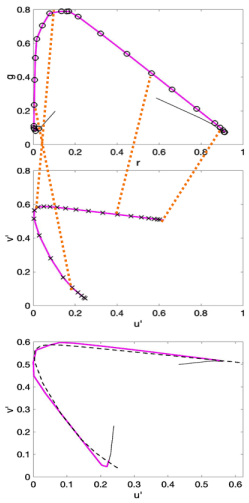

Figure 3: Top: the *rg* locus for a Canon 50D Camera. Middle: the $u'v'$ spectral locus. We match 4 *rg* and $u'v'$ chromaticities for the same wavelengths (they are joined by the orange dotted lines). Bottom: the *rg* chromaticity diagram warped into the $u'v'$ coordinate frame is shown (solid magenta) and contrasted with the actual $u'v'$ locus (dashed black line).

In the top panel, we show the *native rg* chromaticity diagram for the Canon 50D camera (also the subject of Figure 1). From 420 through 670 nm, the locus is plotted in thick, solid magenta. At 10 nm sampling, the corresponding locus chromaticities are marked with an 'O'. Because the camera sensitivities from 400 to 420 and 670 and 700 nm are below a criterion amount (the camera is not very sensitive to these wavelengths) the locus here is shown in a thin black line. In the middle panel, we plot the $u'v'$ diagram. Here, per 10 nm wavelengths are marked with an 'X'. The four orange dotted straight lines linking the *rg* and $u'v'$ diagrams join corresponding wavelengths in the two loci. Given 4 pairs of points, we can solve for the homography that maps the top chromaticity diagram into the middle's $(u'v')$ coordinate frame. The homography-based geometric warp, applied to all the locus points, is shown in the bottom panel (solid magenta). Where, the actual $u'v'$ locus is, for reference, shown as a dashed black line. The fit is excellent.

Let us now examine the computation illustrated in Figure 3 including showing how this relates to colour filter design. Assuming n spectral samples, we can place the spectral sensitivities of the camera and colour-matching functions in the columns of the $n \times 3$ matrices R and X. In forming spectral locus of the camera, we calculate $R' = RC$ (C is defined in Equation 5). The first two columns of R' (which we denote as $R'_{(\bullet,1:2)}$) are unchanged and the last column, denoted $R'_{(\bullet,3)}$, is the sum of R, G and B.The matrix R' is in a useful form for calculating chromaticities. The ith chromaticity is equal to $(r_i, g_i) = (R'_{i,1}/R'_{i,3}, R'_{i,2}/R'_{i,3})$, However, for our purposes it is useful to think of chromaticities as the 3-component homogeneous vectors: $[r_i \; g_i \; 1]^t$. We calculate the $n \times 3$ matrix, $\mathscr{R}$, of homogeneous chromaticity coordinates as

$$\mathscr{R} = p(R') = \text{diag}(\underline{d}^{R'})^{-1} R' \;,\; \underline{d}^{R'} = R'_{(\bullet,3)} \tag{9}$$

where the diag() maps a vector to a diagonal matrix. To ease notation, $p(R')$ denotes the projection of the vectors in R' (dividing by the 3rd coordinate is the perspective projection that for an *image*). An immediate, and useful, property of $p()$ is that it is independent of row scaling. In reviewing colour homography in the last section, an $n \times n$ diagonal matrix was denoted D. By construction:

$$p(DR') = p(R') \tag{10}$$

We could consider mapping *rg* chromaticities to any perceptually relevant chromaticity space. We propose to use $u'v'$ chromaticities here, as the colour space is approximately perceptually uniform[1]. Thus, to facilitate the computation of $u'v'$ chromaticities, we introduce a new transform matrix C_1 (that follows from the standard u', v' definitions, e.g. found in CIELUV [1].

$$C_1 = \begin{bmatrix} 4 & 0 & 1 \\ 0 & 9 & 15 \\ 0 & 0 & 3 \end{bmatrix} \tag{11}$$

Letting, $X' = XC_1$ then the $n \times 3$ matrix of $u'v'$ homogeneous coordinates is calculated as

$$\mathscr{X} = p(X') = \text{diag}(\underline{d}^{X'})^{-1} X' \;,\; \underline{d}^{X'} = X'_{(\bullet,3)} \tag{12}$$

The first two columns of $\mathscr{X}$ contain the $u'v'$ chromaticity coordinates. By construction, the 3rd column is a vector of 1s.

Suppose we wish to map $\mathscr{R}$ to $\mathscr{X}$ in the usual colour homography way we could solve

$$D'\mathscr{R}M' \approx \mathscr{X} \tag{13}$$

(where we use $'$ to signify we are mapping homogeneous coordinates).

From discussion around Equation 8, we know that if $n = 4$ then we can find a unique D' and M' (up to a scalar) such that $D'\mathscr{R}M' = \mathscr{X}$. Further from Equation 10, we know that (for these four points) and the optimal M'

$$\mathscr{P} = p(\mathscr{R}M) = p(D'\mathscr{R}M') = \mathscr{X} \tag{14}$$

Under the assumption that we wish to map 4 points exactly, we can find the *best* homography relating $\mathscr{R}$ to $\mathscr{X}$ using Geometric Warping, Algorithm 2 below. There, to index the i, j, k and lth row of a matrix we use the subscript $_{([i,j,k,l],\bullet)}$. In this algorithm, we, by brute force, find the exact – in closed-form – homography relating the chromaticity coordinates at 4 wavelengths[10]. These 4 camera chromaticities are mapped exactly to 4 corresponding $u'v'$ counterparts. According to each of these mappings, we calculate how well all the points are mapped to one another. The best overall homography mapping is chosen.

Algorithm 2 Geometric Warping

```
ERR=∞
for ∀([i,j,k,l] ∈ {K ⊂ {1,...,n} | |K| = 4}) do
    Find M'' s.t. p(R_([i,j,k,l],•) M'') = X_([i,j,k,l],•)
    err = ||p(R M'') − X||_F
    if err < ERR then
        M' = M'', ERR = err
    end if
end for
```

Let us now consider how this geometric fit relates to filter design. As for the rg and $u'v'$ coordinates, let us explicitly write out how the warped homogeneous coordinates are mapped to chromaticities. Let the final approximate $u'v'$ coordinates, $\hat{\mathscr{R}}$, be defined as

$$\hat{\mathscr{R}} = \text{diag}(\underline{d}^{\mathscr{R}M'})^{-1}\mathscr{R}M' \;,\; \underline{d}^{\mathscr{R}M'} = [\mathscr{R}M']_{(\bullet,3)} \tag{15}$$

Interestingly, because we know how $\mathscr{X}$ was derived from the colour-matching functions X, we can map these approximate chromaticities to approximated XYZ CMFs, $\hat{X}$:

$$\hat{X} = \text{diag}(d^{X'})\hat{\mathscr{R}}C_1^{-1} \tag{16}$$

Here the diagonal matrix $\text{diag}(\hat{d}^{X'})$ *undoes* the perspective projection (Equation 12) and post-multiplying by C_1^{-1} them maps us back to the XYZ CMFs. In fact, because we know how we mapped RGBs to homogeneous coordinates, the homography M' and the projection of the transformed coordinates to the $u'v'$ chromaticity plane, we can explicitly write these operations too: R (the original camera sensitivities) relate to X (the XYZ CMFs) according to:

$$X \approx \text{diag}(\underline{d}^{X'})\,\text{diag}(\underline{d}^{\mathscr{R}M'})^{-1}\,\text{diag}(d^{R'})^{-1}RCM'C_1^{-1} \tag{17}$$

Since R represents the spectral sensitivities of the camera, the meaning of Equation 17 is that a diagonal matrix premultiplies the sensitivities and then is post multiplied by a 3×3 matrix. Physically, the diagonal matrix can be thought of as being equivalent to the coloured filter illustrated in Figure 1. We call this filter the geometric warp filter, and it is defined in Equation 18.

$$f(\lambda) \equiv \text{diag}(\underline{d}^{X'})\,\text{diag}(\underline{d}^{\mathscr{R}M'})^{-1}\,\text{diag}(d^{R'})^{-1} \tag{18}$$

Note here there is no explicit computation of the filter. Simply, if falls out of the known projective geometry at hand. Finally, note if we substitute $M = CM'C_1^{-1}$ In Equation 17 we are back in the colour homography form, Equation 8.

Experiments

We return to the example set forth in Figure 1. Now, using Algorithm 2 we warp the rg chromaticities of a Canon 50D camera to $u'v'$ counterparts using Algorithm 2. And once we find the best geometric warp, the corresponding optimal filter is returned by Equation 18. The result of filter design by geometric warping (red line) versus using alternating least-squares (blue line) is shown in Figure 4. It is clear the filters are similar – peaks and toughs are located at the same wavelengths – but different. This is not too surprising as the geometric warp delivers the best filter for matching chromaticities and the colour homography matches full 3-dimensional sensitivities. Predictably, if you change the error metric, the best solution to a problem can also change.

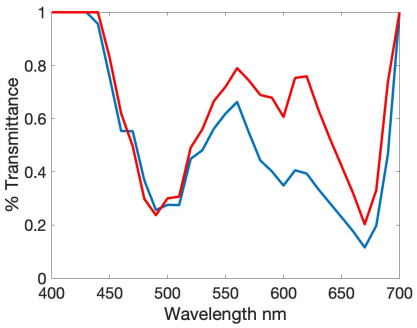

Figure 4: Blue and Red lines respectively filters designed by Alternating Least-Squares and Geometric Warping

In Figure 5, we summarise the results of an experiment for mapping rg loci for 28 cameras[11] to the $u'v'$ locus. Here each camera is represented by a 31×3 matrix numbers corresponding to the R, G, B colour channels and where the wavelength domain is from 400 to 700 and sampled every 10 nm . The XYZ CMFs are similarly represented as a 31×3 matrix. For each camera, we solve for 3 mappings M_i, $i = \{1,2,3\}$. The mapping M_1 is simply the least squares fit from the camera sensitivities to the XYZ CMFs. We then map the fitted values to 31 $u'v'$ coordinates. In the second fit, we run the prior art filter design method[9] to find a diagonal matrix D and a 3×3 matrix M_2 (we use Algorithm 1 and minimize Equation 8). Because D only scales the rows of the camera sensitivity matrix X it plays no role in the formation of $u'v'$ coordinates (we need to calculate D to find M_2 but once found we can ignore it here). Lastly, we find the fit, M_3 according to the geometric warp described in the last section (Algorithm 2).

In Figure 5, we plot for each of the 28 cameras the mean $u'v'$ errors (the Euclidean distance between the corresponding per wavelength actual and approximate $u'v'$ coordinates) multiplied

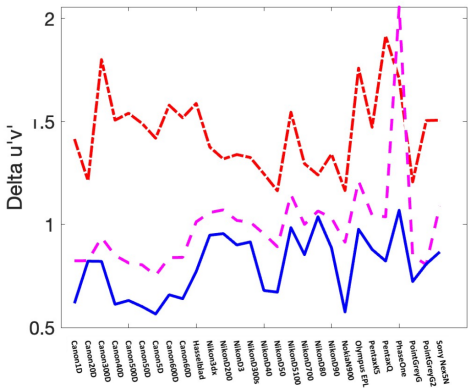

Figure 5: Solid Red, $u'v'$ error for simple least-squares. Dashed purple shows same error for colour homography (Algorithm 1). Solid blue, $u'v'$ error for the homography found by geometric mapping (Algorithm 2).

by 13 (since when the *Luv* colour space is used to gauge colour differences for reflectance samples the multiplication by 13 yields distances corresponding roughly to a JND of 1 [22]). The, least-squares mean error (M_1) is shown in red, colour homography error (M_2 found using Algorithm 1) in dashed magenta and the new geometric warp method error (M_3) is shown in solid blue. Compared to a least-squares fit, the geometric warp returns about half the mean Delta $u'v'$ error.

Perhaps unsurprisingly, the geometric warp optimization always minimizes the mean Delta $u'v'$ error works best for all cameras. Also, in general, the mapping found using the colour homography method (Algorithm 1) works well. Though in one case (for the PointGrayG2 camera) a simple least-squares fit results in lower $u'v'$ fitting error. This camera – of all 28 – is the least colorimetric after filter compensation. Moreover, Algorithm 1 does not attempt to minimize the $u'v'$ error. These facts taken together account for this unusual (1 out of 28 cameras) circumstance where solving for a filter - using the prior art method [9] - results in less accurate mapped chromaticities.

While a filter designed based on the spectral sensitivities of a camera or its corresponding *rg* locus, we are, in effect, attempting to find the best *worst-case* filter. Indeed, we are assuming that any spectra might be measured by a camera, including monochromatic stimuli. However, in everyday imaging scenarios, the scenes we capture contain surfaces with smooth spectral reflectances, and they are illuminated by orange, yellow, white, and blue lights (but rarely by purple or green lights). Thus, we wish to look at filter performance for a set of real lights and surfaces.

In this experiment, we use the same 28 cameras and the SFU set of 1995 reflectances and 102 lights [23]. Per camera and light, we make 3 sets of raw RGBs: the native camera RGBs and two sets where we use filtered camera spectral sensitivities. The first filter is the one returned by our geometric warping idea (Algorithm 2). Now, we use the geometric warp filter to initialise Algorithm 1 (find the colour homography) and recover the *Refined Geometric Warp* Filter. The rationale here is that, in this experiment, we seek a good filter not for projective coordinates but full 3-dimensional camera RGBs and XYZ tristimulus values. For each filter condition – no filter, geometric warp filter, and refined geometric warp filter, we calculate the CIELAB Delta E between the measured per camera and per light RGBs and the XYZs for the same light. Per filter condition, this returns $1995 * 102 * 28$ ΔE errors. The corresponding box plot for the 3 filter condition errors is shown on the left of Figure 6. We also show the No-Filter and Refined Geometric Warp filter box plot just for the Canon 50D camera, on the right.

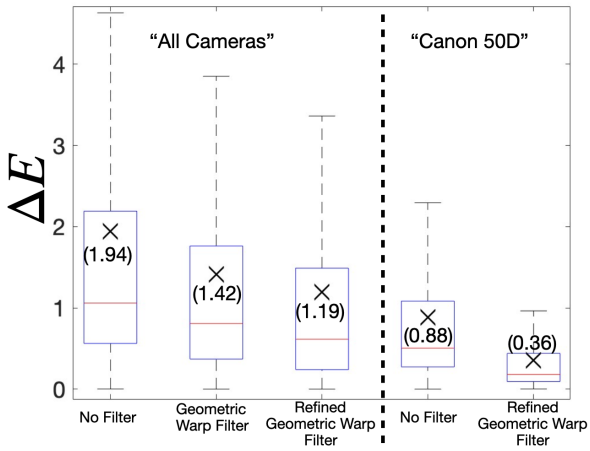

Figure 6: Colour measurement performance - with and without a filter correction - for all cameras and, separately, the Canon D50.

The line in the centre of each box records the median error. The top and the bottom of the box demarcates 25 and 75% quantile errors, and the top and bottom of the dotted lines the 99 and 1% quantiles. The 'X' marks the mean errors. In broad sweep, this experiment returns statistics similar to the prior art [9]. The mean error without the refined filter is over 50% larger than with it in place. The performance increment is even larger for typical photographic cameras. For the Canon D50 camera, the mean, median and 95% quantile errors are less than half for the Refined Geometric Warp Filter than the non filter case.

Conclusion

In colorimetric filter design, a special transmissive filter is designed which, when it is placed in front of a camera, makes the camera more colorimetric. Hitherto the optimisation for solving for a good filter has, arguably, had two weaknesses. First, the efficacy of the optimisation depends on choosing a good initialisation condition. Second, the prior art optimisation as formulated does not minimise a perceptually relevant metric.

In this paper, we presented a surprising result. We demonstrated that the problem of geometrically warping an *rg* camera chromaticity diagram to one derived from human vision colour-matching functions – we use $u'v'$ here – is the same as conventional homography mapping in the spatial domain (e.g. for making panoramas). For all cameras tested, we can geometrically warp the *rg* camera loci onto $u'v'$ to a good approximation. Moreover, so long as we keep a record of how chromaticities are formed, the best colorimetric filter falls from the warp operation *for free* (without any explicit optimisation). Our method both finds a globally optimal colorimetric filter and – as we are using $u'v'$ target chromaticity coordinates – does so in a perceptually relevant way.

Acknowledgments

This research was funded by the Research Council of Norway for the project 'Individualised Colour Vision-based Image Optimisation', grant number 287209. Graham Finlayson is also grateful to support from grant EPSRC EP/S028730/1.

References

[1] G. Wyszecki and W.S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulas*. Wiley, New York, 2nd edition, 1982.

[2] B.K.P. Horn. Exact reproduction of colored images. *Comp. Vision, Graphics and Image Proc.*, 26:135–167, 1984.

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

[4] G. Hong, M. R. Luo, and P. A. Rhodes. A study of digital camera colorimetric characterization based on polynomial modeling. *Color Research & Application*, 26(1):76–84, 2001.

[5] A. Badano et al. Consistency and standardization of color in medical imaging: a consensus report. *J. Digital Imaging*, 28(1):41–52, 2015.

[6] C. Boukouvalas, J. Kittler, R. Marik, and M. Petrou. Color grading of randomly textured ceramic tiles using color histograms. *IEEE Transactions on Industrial Electronics*, 46(1):219–226, 1999.

[7] M. Tsuchida et al. Stereo one-shot six-band camera system for accurate color reproduction. *Journal of Electronic Imaging*, 22(3):033025, 2013.

[8] J. Behmann et al. Specim IQ: evaluation of a new, miniaturized handheld hyperspectral camera and its application for plant phenotyping and disease detection. *Sensors*, 18(22):441, 2018.

[9] G. D. Finlayson, Y. Zhu, and H. Gong. Using a simple colour prefilter to make cameras more colorimetric. In *The 26th Color and Imaging Conference*, pages 182–187. 2018.

[10] G. D. Finlayson and Y. Zhu. Designing color filters that make cameras more colorimetric. *IEEE Transactions on Image Processing*, 30:853–867, 2021.

[11] J. Jiang, D. Liu, J. Gu, and S. Süsstrunk. What is the space of spectral sensitivity functions for digital color cameras? In *IEEE Workshop on Applications of Computer Vision*, pages 168–179, 2013.

[12] Y. Zhu and G. D. Finlayson. Matched illumination: using light modulation as a proxy for a color filter that makes a camera more colorimetric. *Opt. Express*, 30(12):22006–22024, 2022.

[13] M. Brown and D. G. Lowe. Automatic panoramic image stitching using invariant features. *International Journal of Computer Vision*, 74(1):59–73, 2007.

[14] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press,, second edition, 2004.

[15] G. D. Finlayson, M. Mohammadzadeh Darrodi, and M. Mackiewicz. The alternating least squares technique for nonuniform intensity color correction. *Color Research & Application*, 40(3):232–242, 2015.

[16] G. D. Finlayson, H. Gong, and R.B. Fisher. Color homography color correction. In *The 24th Color and Imaging Conference*, pages 310–315. 2016.

[17] M. J. Vrhel. Improved camera color accuracy in the presence of noise with a color prefilter. In *The 28th Color and Imaging Conference*, pages 187–192. 2020.

[18] Y. Zhu and G. D. Finlayson. Designing a color filter with high overall transmittance for improving the color accuracy of digital cameras. In *The 29th Color and Imaging Conference*, pages 1–6. 2021.

[19] H. J. Rivertz. On filters making an imaging sensor more colorimetric. In *The 28th Color and Imaging Conference*, pages 169–174. 2020.

[20] M. J. Vrhel and H. J. Trussell. Selection of optimal external filter for colorimetric camera. *The 29th Color and Imaging Conference*, pages 141–146, 2021.

[21] H. Gong, G. D. Finlayson, R. B. Fisher, and F. Fang. 3D color homography model for photo-realistic color transfer re-coding. *The Visual Computer*, 35:323–333, 2019.

[22] Günter Wyszecki. Proposal for a new color-difference formula. *J. Opt. Soc. Am.*, 53(11):1318–1319, 1963.

[23] K. Barnard, L. Martin, B. Funt, and A. Coath. A data set for color research. *Color Research & Application*, 27(3):147–151, 2002.

Author Biography

Graham Finlayson received his BSc in Computer Science from the University of Strathclyde in 1989 and, respectively, in 1992 and 1995 his MSc and PhD in Computing Sciences from Simon Fraser University. He is a professor in Computer Science at the University of East Anglia where he is the Director of the Colour & Imaging Lab. Graham's research spans colour image processing, physics-based computer vision and visual perception. He is interested in taking the creative spark of an idea, developing the underlying theory and algorithms and then implementing and commercialising the technology.

Ivar Farup received his M.Sc. in physics from the Norwegian Institute of Technology (NTH) in Trondheim, Norway in 1994 and a PhD in applied mathematics from the University of Oslo, Norway in 2000. He is currently a professor of computer science at Norwegian University of Science and Technology, Gjøvik, Norway, mainly focusing on computational methods in colour science and image processing.