A Theory of Selection for Gamut Mapping Colour Constancy

Graham Finlayson

University of Derby
Colour & Imaging Institute
Britannia Mill
Mackworth Road
Derby DE22 3BL

Steven Hordley

University of Derby
Colour & Imaging Institute
Britannia Mill
Mackworth Road
Derby DE22 3BL

Abstract

Gamut mapping colour constancy attempts to determine the set of diagonal matrices taking the gamut of image colours under an unknown illuminant into the gamut of colours observed under a standard illuminant. Forsyth [5] developed such an algorithm in rgb *sensor space which Finlayson [3] later modified to work in a 2-d chromaticity space. In this paper we prove that Forsyth's 3-d solution gamut is, when projected to 2-d, identical to the gamut recovered by the 2-d algorithm. Whilst this implies that there is no inherent disadvantage in working in chromaticity space, this algorithm has a number of problems; the 2-d solution set is distorted and contains practically non-feasible illuminants. These problems have been addressed separately in previous work [4, 3]; we address them together in this paper.*

Non-feasible illuminants are discarded by intersecting the solution gamut with a non-convex gamut of common illuminants. In 2-d this intersection is relatively simple, but to remove the distortion, both these sets should be represented as 3-d cones of mappings, and the intersection is more difficult. We present an algorithm which avoids performing this intersection explicitly and which is simple to implement. Tests of this algorithm on both real and synthetic images show that it performs significantly better than the best current algorithms.

1 Introduction

The colour of the light leaving a surface depends both on the reflective properties of the surface and on the colour of the incident illumination. Changing the illuminant changes the sensor *rgb* values measured at an imaging device. Computer vision tasks often require that these sensor values be independent of the illuminant. Colour constancy algorithms are designed to correct the *rgb* values by discounting the effect of the incident illumination. *Gamut mapping* colour constancy algorithms attempt to map image colours into the gamut of colours observed under standard (canonical) lighting. Forsyth [5] developed a two stage algorithm which maps image colours using 3-dimensional diagonal matrices. In the first stage a 3-d convex set of diagonal mappings is recovered, representing all the possible scene illuminants. Then a single mapping is selected from this set as an estimate of the unknown illumination. While the algorithm performs well on images of flat, matte, uniformly illuminated scenes its performance on images of more realistic scenes can be poor (it has a particular problem with specularities) and furthermore the algorithm is computationally intensive.

To overcome these problems Finlayson [3] modified the algorithm to work in a 2-d chromaticity space. This simplifies the computation and means that features that confound 3-d gamut mapping such as specularities, shape information, and varying illumination no longer affect algorithm operation. How the transformation to 2-d chromaticity space affects the algorithm's performance however, has not yet been fully addressed. We prove in this paper the important result that the set of 3-d diagonal matrices returned by Forsyth's algorithm, when projected to 2-d, is identical to the set of 2-d diagonal matrices recovered by the 2-d algorithm.

Having shown that the two algorithms return identical constraints on the possible scene illuminants we address the problem of selecting a single mapping from the feasible set. There are two factors to consider here. First, previous work [4] has shown that in chromaticity space the set of feasible mappings is distorted, significantly degrading the algorithm's performance. Second, Finlayson [3] has pointed out that though each of the mappings in the feasible set corresponds to a physically realisable illuminant, many of these illuminants do not occur in practice. The distortion problem can be overcome by converting the 2-d constraint set to a 3-d cone of possible mappings. To limit the mapping set to realistic illuminants Finlayson intersected a gamut of plausible illuminants with the feasible set. Unfortunately this illuminant gamut too, is distorted. We can reverse the distortion effects in the same way as we do for the feasible set. However, the illumination gamut is non-convex so intersecting it with the feasible set in 3 dimensions is difficult. In this paper we give details of an algorithm which avoids performing this intersection explicitly. Consequently the algorithm is simple to implement and has the advantage that practically all computation is performed in 2-d.

In the next section we review the gamut mapping theory of colour constancy and discuss its limitations. In the rest of the paper we show how these limitations can be overcome. First we prove equivalence of the 2-d and 3-d algorithms. We then show how the illumination constraint can be added to the gamut mapping algorithm and give some details of our algorithm. In the final sections of the paper we illustrate the performance of the algorithm using real and synthetic images and we draw some conclusions from this work.

2 Background

Gamut mapping proceeds in two stages. First the set of all possible scene illuminants is recovered and then an estimate of the illuminant is found by selecting a single mapping from the set. Possible illuminants are represented as diagonal matrices which map the gamut of image *rgb*s to a set of canonical *rgb*s recorded under a canonical illuminant. The first stage of gamut mapping then is to find all diagonal matrices D such that

$$\forall \underline{p} \in \Gamma(I), D\underline{p} \in \Gamma(C) \tag{1}$$

where $\Gamma(I)$ and $\Gamma(C)$ are respectively, the image and canonical gamuts.

To begin we build the canonical gamut $\Gamma(C)$ of all sensor responses obtained by viewing all physically realisable surface reflectances under a canonical illuminant. If $\mathrm{C} = \{\underline{p}^{c,1}, \ldots, \underline{p}^{c,i}, \ldots, \underline{p}^{c,n}\}$ is a set of sensor responses to a range of n surfaces then any convex combination $\underline{p}^{c,x} = \sum_{i=1}^{n} \omega_i \underline{p}^{c,i}$ $(\forall i \omega_i > 0$ and $\sum_{i=1}^{n} \omega_i = 1)$, of vectors in C is also a physically realisable response. So we take as the canonical gamut, the set of all convex combinations of C which we denote $\Gamma(\mathrm{C})$. Since $\Gamma(\mathrm{C})$ is convex we need consider only its convex hull[1] C. Similarly we construct the image gamut from the set of image *rgb*s. If I is the set of sensor *rgb*s recorded under the unknown illumination, our image gamut is $\Gamma(I)$ the set of all convex combinations of I, the convex hull of I.

In Forsyth's algorithm we would next find all 3-d diagonal matrices satisfying equation (1) before selecting one of these as the illuminant estimate. Features such as specularities, shape information, and varying illumination intensity however, confound Forsyth's algorithm. But, since these factors only mitigate against the recovery of *rgb* intensity Finlayson modified the algorithm so that it recovers only the orientation of the scene illuminant, intensity information being discarded by transforming 3-d *rgb* colour vectors to 2-d chromaticity vectors:

$$r' = \frac{r}{b}, g' = \frac{g}{b}, 1 = \frac{b}{b} \tag{2}$$

Applying this transform to the 3-d gamuts, we obtain the corresponding 2-d gamuts $\Gamma(I')$ and $\Gamma(C')$ (the convexity of these gamuts is conserved under the transform [3]). Now we must find 2-d diagonal matrices D which map the 2-d image gamut to the 2-d canonical gamut. The uth point, $\underline{q}^{i,u}$ in the convex hull of the perspective image gamut I' can be mapped to the vth point $\underline{q}^{c,v}$ in the convex hull of the perspective canonical gamut C' by a diagonal matrix $D^{u,v}$ with kth diagonal element $D_{kk}^{u,v} = q_k^{c,v}/q_k^{i,u}$. In this way we can form the set of mappings which take $\underline{q}^{i,u}$ into C' which we denote $C'/\underline{q}^{i,u}$. We can represent these matrices as 2-vectors, hence:

$$C'/\underline{q}^{i,u} = D\underline{q} : \underline{q} \in C', D_{kk} = \frac{1}{q_k^{i,u}} \tag{3}$$

Since a point in $C'/\underline{q}^{i,u}$ maps $\underline{q}^{i,u}$ to C', $\underline{q}^{i,u}$ can be mapped anywhere in C' by a convex combination of points in $C'/\underline{q}^{i,u}$. We denote the set of all such convex combinations by $\Gamma(C'/\underline{q}^{i,u})$. We can find similar mapping sets for all points in I'. The intersection of all such sets, denoted $\Gamma(C'/I')$ is the set of mappings which take all points in the image gamut to the canonical gamut:

$$\Gamma(C'/I') = \bigcap_{j=1}^{n} \Gamma(C'/\underline{q}^{i,j}) \tag{4}$$

To solve for colour constancy we must now choose a single mapping from this set to represent the unknown scene illuminant. Finlayson (following Forsyth) adopted the heuristic approach of choosing the mapping which makes the image as colourful as possible. A more intuitive approach, first used by Barnard [1]

[1] The convex hull of a set C is defined to be [7] the minimal cardinality subset of C, denoted C such that $\Gamma(\mathrm{C}) = \Gamma(C)$.

is to assume that all illuminants in the mapping set are equally likely. The *mean* of this mapping set is then, in general, a good estimate of the unknown illuminant. Previous work [4] has shown however that taking the mean of the 2-d mapping as Barnard did is not the best approach. To see why this is the case consider a 3-d image vector $(r^i, g^i, b^i)^t$ mapped to a vector $(r^c, g^c, b^c)^t$ in the canonical gamut:

$$\begin{pmatrix} \alpha & 0 & 0 \\ 0 & \beta & 0 \\ 0 & 0 & \gamma \end{pmatrix} \begin{pmatrix} r^i \\ g^i \\ b^i \end{pmatrix} = \begin{pmatrix} r^c \\ g^c \\ b^c \end{pmatrix} \tag{5}$$

Under the perspective transform taking *rgb* vectors to 2-d chromaticity vectors this equation becomes:

$$\begin{pmatrix} \frac{\alpha}{\gamma} & 0 \\ 0 & \frac{\beta}{\gamma} \end{pmatrix} \begin{pmatrix} \frac{r^i}{b^i} \\ \frac{g^i}{b^i} \end{pmatrix} = \begin{pmatrix} \frac{r^c}{b^c} \\ \frac{g^c}{b^c} \end{pmatrix} \tag{6}$$

From this equation it is clear that a 2-d mapping $(\alpha', \beta')^t$ taking an image chromaticity to the canonical gamut can be considered as a 3-d mapping in the plane $\gamma = 1$. Taking the mean of the 2-d set ignores the fact that it is perspectively distorted and leads to an unwanted bias in our estimate of the unknown illuminant. To remove this bias we must reverse the effect of the perspective transform. To do this we form the 3-d cone of mappings whose vertex is at the origin and whose rays pass through the extreme points of the mapping set in the plane $\gamma = 1$. The mean of this 3-d set is not distorted and provides a good estimate of the unknown illuminant[4].

While perspective gamut mapping with cone-based selection can provide good colour constancy, it has been criticised in two respects. First, while it is easy to show that the illuminant constraint delivered by the perspective algorithm is no stronger than that provided by Forsyth's original 3-d theory it has proven harder to establish that the converse is also true. Yet proving the converse is crucial: the perspective algorithm should only be used in preference to 3-d gamut mapping if it delivers an equally strong constraint on the illuminant colour. A second criticism is that some of the mappings in the cone of feasible maps do not correspond to 'plausible' illuminants. For example, a purple light is physically realisable, but it does not occur in nature nor is it used for artificial lighting. So if a mapping corresponding to such an illuminant is recovered it should be rejected. Finlayson [3] addressed this problem in the 2-d perspective theory and showed how the illuminant plausibility constraint could be characterised as a 2-d non-convex set of maps. Enforcing the constraint involved intersecting this set with the set of feasible maps returned from gamut mapping. In the context of the cone-based algorithm this intersection must be performed in 3-d; a difficult problem but one that needs to be solved to improve the cone based selection procedure. We address these difficulties in the next section.

3 Simplifying Gamut Mapping

We first prove the following theorem:

Theorem: *The set of 3-d diagonal matrices recovered using Forsyth's algorithm is, when projected into 2-d chromaticity space, identical to the set of 2-d diagonal matrices recovered by Finlayson's algorithm.*

Proof of Equivalence

Consider first a single 3-d mapping set $\Gamma(C/\underline{p}^{i,u})$. A mapping $(\alpha, \beta, \gamma)^t$ in this set maps the uth image point $\underline{p}^{i,u} = (r^i, g^i, b^i)^t$ to a point $(r^c, g^c, b^c)^t$ somewhere in the canonical gamut. Under the perspective transform we have:

$$\begin{pmatrix} \frac{\alpha}{\gamma} & 0 & 0 \\ 0 & \frac{\beta}{\gamma} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \frac{r^i}{b^i} \\ \frac{g^i}{b^i} \\ 1 \end{pmatrix} = \begin{pmatrix} \frac{r^c}{b^c} \\ \frac{g^c}{b^c} \\ 1 \end{pmatrix} \tag{7}$$

That is, the corresponding chromaticity co-ordinates are mapped by a diagonal mapping $(\frac{\alpha}{\gamma}, \frac{\beta}{\gamma}, 1)^t$; a vector in the same direction as $(\alpha, \beta, \gamma)^t$ but with a different magnitude. Hence, the direction of vectors in the perspective mapping set is equivalent to that of vectors in the 3-d set. Here we have ignored maps of the form $(\alpha, \beta, 0)^t$ which have no corresponding perspective counterparts. However, such mappings never occur and if they did then all blue responses would be mapped to 0 making colour constancy impossible. Of course because it is always possible to switch the light source off we add the map $(0, 0, 0)^t$ (representing black light) to our perspective set.

Next we consider the effect of intersecting two mapping sets. In 3-d each mapping set is a bounded, convex polyhedron, which we can represent as the intersection of a set of halfspaces [7] [2] H. Halfspaces can be written as constraints in one of two forms:

$$\begin{aligned} a_1\alpha + a_2\beta + a_3\gamma &>= 0 \\ a_1\alpha + a_2\beta + a_3\gamma &>= b, b \neq 0 \end{aligned} \tag{8}$$

If the dividing plane passes through the origin the constraint is of the first kind, if not, it is of the second. Importantly only constraints of the first kind exclude points according to their direction. To see that this is the case, consider a sphere with radius r centred at the origin. When $b \neq 0$ the dividing plane does not

[2] A plane divides 3-dimensional space into two halves. A half space is the set of points lying on the plane together with the points in one of these two spaces[7].

pass through the origin and we can choose r such that the sphere lies entirely within the halfspace. Since a sphere contains points in all directions, the halfspace does not exclude points on the basis of their direction. If $b = 0$ however, the plane cuts through the sphere so whatever value we choose for r some points will lie outside the halfspace; points of certain orientations are excluded. So we can represent the mapping set H as the union of two sets of halfspaces: $H = A \cup B$ where A contains constraints such that $b = 0$ and B constraints with $b \neq 0$. All mapping sets must contain the point $(0,0,0)^t$ corresponding to black light. This observation coupled with the fact that all sensor responses are positive, suffices to ensure that A is non-empty. Also, it is immediate that under a perspective transform the mapping set is equal to the unbounded 3-d cone represented by the constraint set A (the range of directions in the perspective set is identical to that in the 3-d set). Consider two 3-d mapping sets whose halfspace constraints are represented by $H_1 = A_1 \cup B_1$ and $H_2 = A_2 \cup B_2$. For a point $\underline{d}$ to lie in the intersection of these two sets it must satisfy the constraints in H_1 and H_2, in which case it lies in the set of directions bounded by the convex cones A_1 and A_2 and its length is bounded by the intersection of B_1 and B_2. That is, the direction of $\underline{d}$ is bounded only by the answer to perspective mapping colour constancy: $\underline{d} \in A_1 \cap A_2$. This completes the proof.

3.1 The Illumination Constraint

We now wish to add an illuminant plausibility constraint to further limit the size of the solution set $A_1 \cap A_2$. We do this in two stages. First we characterize the set of plausible illuminant maps. Second, we show how this constraint might actually be applied. The method of application turns out to be crucial for effective selection for gamut mapping colour constancy.

Each mapping in the feasible set $\Gamma(C/I)$, recovered at the first stage of the gamut mapping algorithm, represents a physically realisable illuminant. In practice however the range of illuminants which occur in the world is quite restricted. So it is likely that the feasible set contains illuminants which occur very rarely or not at all. To reflect this Finlayson [3] introduced an extra constraint on the solution set in the form of an illuminant gamut. Let E be the set of perspective chromaticities for a surface s imaged under n different illuminants. We denote the convex hull of this set by E and the set of all convex combinations of chromaticities in E, which we call the illuminant gamut, by $\Gamma(E)$. Now consider $\underline{q}^{c,s}$, the chromaticity of s under the canonical illuminant. We can find a convex set of diagonal mappings which take this chromaticity to any other chromaticity in $\Gamma(E)$. This set we write as $\Gamma(E/\underline{q}^{c,s})$. A diagonal mapping D taking the image gamut into the canonical gamut is valid only if its inverse mapping D^{-1} takes the canonical illuminant to another illuminant in the illuminant gamut. That is $D^{-1} \in \Gamma(E/\underline{q}^{c,s})$. Now if $\Gamma^{-1}(E/\underline{q}^{c,s})$ is the non-convex set of diagonal mappings taking illuminants in the illuminant gamut to the canonical illuminant then

$$D \in \Gamma^{-1}(E/\underline{q}^{c,s}) \Rightarrow D^{-1} \in \Gamma(E/\underline{q}^{c,s}) \qquad (9)$$

This symmetry between the inverted illuminant gamut $\Gamma^{-1}(E)$ and the inverted diagonal matrix D shown above is crucial in implementing the cone-based selection that exploits the illuminant constraint. It shows that we do not actually have to calculate the non-convex illuminant gamut. Rather we can invert a feasible map and check for plausibility in terms of the convex set $\Gamma(E)$. More details are described as part of the whole gamut mapping process which is summarised below.

3.2 The Algorithm

1. Form the canonical gamut $\Gamma(C)$:
   a. *Image n surfaces under a canonical illuminant*
   b. *For each* rgb *calculate a chromaticity*
   c. *Find the convex hull of all chromaticities*
2. Form the illuminant constraint: $\Gamma(E)$:
   a. *Calculate the set of maps from the canonical to m plausible lights*
   b. *For each map calculate the corresponding perspective map*
   c. *Find the convex hull of the perspective maps*
3. Form the image gamut: $\Gamma(I)$:
   a. *For each image* rgb *calculate the corresponding chromaticity* $(\frac{r}{b}, \frac{g}{b})$
   b. *Find the convex hull of the image chromaticities*
4. Run Gamut mapping, calculate $\Gamma(I/C)$:
   a. *For each point in I calculate the corresponding set of mappings*
   b. *Intersect all the mapping sets*
5. Select a single mapping:
   a. *Set* $\underline{\mu} = \underline{0}$ *and* $c = 0$
   b. *Choose a vector* $\underline{v}$ *at random from within the unit sphere*
   c. *If* $(\frac{\alpha}{\gamma}, \frac{\beta}{\gamma}) \notin F$ *goto b.*
   d. *If* $(\frac{\alpha}{\gamma}, \frac{\beta}{\gamma}) \notin E$ *goto b.*
   e. $\underline{\mu} = \underline{\mu} + \underline{v}, c = c + 1$
   f. *If* $c < U$ *goto b.*
   g. $\underline{\mu} = \underline{\mu}/U$

In preprocessing steps (1) and (2) the canonical and illuminant gamuts are built. For a given image the image gamut is constructed in step (3) and from this the set of feasible maps in (4). It remains to select a mapping belonging both to the feasible set and satisfying the illuminant constraint. Step (5) details an algorithm that returns the mean mapping, calculated in 3-d (without the perspective distortion), that satisfies both constraints. It is apparent that this mean selection proceeds by Monte Carlo simulation. That is 3-d mappings are generated at random and are checked against the feasibility and illuminant constraints (this is done in 2-d). This process is repeated U times where U is sufficiently large to provide a stable mean statistic.

4 Results

To assess our algorithm we tested it using both synthetic and real images. A canonical gamut was first constructed. The surface reflectances of 462 Munsell [8] chips together with the most spectrally uniform of the Judd [6] daylight phases (D55) was used to form the set of sensor responses of a camera with spectral sensitivities $\underline{R}(\lambda)$. Given a surface reflectance $S(\lambda)$, an illuminant spectral power distribution $E(\lambda)$, and a 3-vector of camera spectral sensitivities $\underline{R}(\lambda)$, the camera's sensor responses can be calculated as: $\underline{p} = \int_\omega \underline{R}(\lambda)E(\lambda)S(\lambda)d\lambda$. The convex hull of the corresponding chromaticities $\underline{q}$ is used as the canonical gamut. For our illuminant gamut 37 common illuminants were selected including the Judd daylight phases, CIE standard illuminants [8] A, F11, and Planckian black body radiators. The sensor responses to a perfectly reflecting white surface under each of these illuminants was used to create the mapping set taking the canonical illuminant to all other illuminants. A subset of n random surface reflectances from the set of Munsell chips forms a synthetic image. Sensor responses to these surfaces under an illuminant from the illuminant gamut are calculated and the convex hull of the corresponding n image chromaticities form the input to the algorithm. The algorithm returns a 3-d diagonal mapping $\underline{d}$ as an estimate of the unknown illuminant. We use two metrics to measure the accuracy of the algorithm. First we calculate the angle between our estimate $\underline{d}$ and the correct diagonal mapping. Given n surfaces in an image the angle will varying according to the range of chromaticities in the image, therefore to get an accurate statistic we averaged the angular error over many different synthetic images. Figure 1 plots this statistic against the number of surfaces in the image.

The second error measure we call the worst case error. Our algorithm returns a set of diagonal mappings representing all possible scene illuminants. From this set we select a single mapping but in fact any mapping in this set could be the correct answer. So to measure how wrong our estimate can be, we calculate the maximum angular error between our estimate and any other mapping in the feasible set. The results for this measure are plotted in figure 2, again the error is averaged over many different scenes for each value of n. We compared our new algorithm to the previous best algorithm (2-d gamut mapping with the 3-d cone selection) and also to the case of performing no colour constancy. Importantly, both algorithms always do significantly better than doing no colour constancy at all (red line). It is clear from the graphs that our algorithm (green line) does significantly better than the previous best algorithm (blue line). The results are particularly impressive for the worst case error measure. In this case, for small numbers of surfaces, our new algorithm does 50%-65% better than the previous best algorithm.

A more important test of the algorithm, is its performance on real images. To test this we took a number of images of the same scene under different lights. Figure 3 shows the results of these experiments for

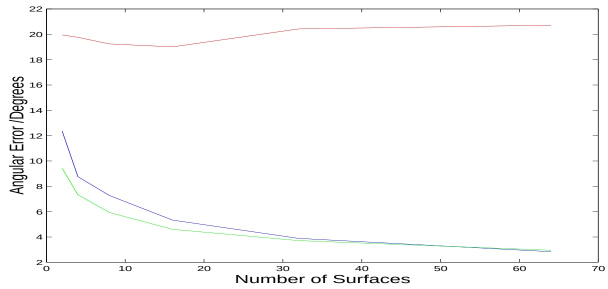

Figure 1: Average actual angular error for synthetic images.

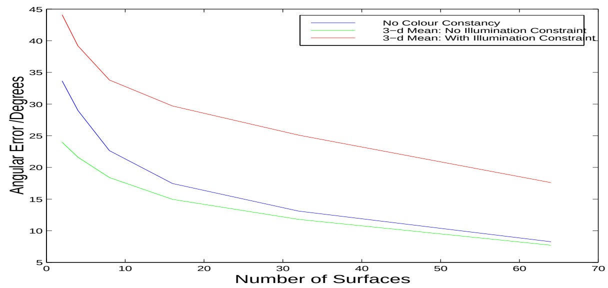

Figure 2: Average worst case angular error for synthetic images.

two scenes [3] In the top two images we show the same scene imaged under a reddish tungsten (left) and D65 (right); the colour shift is clear. The next two images show the result of correcting the tungsten image using the grey world algorithm [2] (left) and our new algorithm (right). The second set of four images are similar except that this time the image to be corrected is taken under a fluorescent light. It is clear from these images that our algorithm produces excellent colour constancy. It is noticeable also that the grey world assumption (the colour of the averaged light from a scene is grey) breaks down here and grey world colour constancy performance is poor.

5 Conclusions

In this paper we proved that 2D and 3D gamut mapping algorithms return identical constraints on the colour of the illuminant. This result is very important since the 2D algorithm has much lower computational complexity. However, because, the 3D problem is mapped to 2D by a perspective transform the 2D constraint (while equivalent) is somewhat distorted. Unfortunately, this distortion causes problems when we come to choose a single answer from the constraint set. Methods developed in this paper show how this distortion can be dealt with by inverting the perspective transform. In an important extension to previous work the constraint on illuminant colour plausibility is incorporated into this process. Colour constancy experiments were carried out on synthetic and real images. In all cases the new algorithm provided significantly better colour constancy performance than existing methods.

Acknowledgements

This work was supported by the EPSRC and by Hewlett-Packard Incorporated.

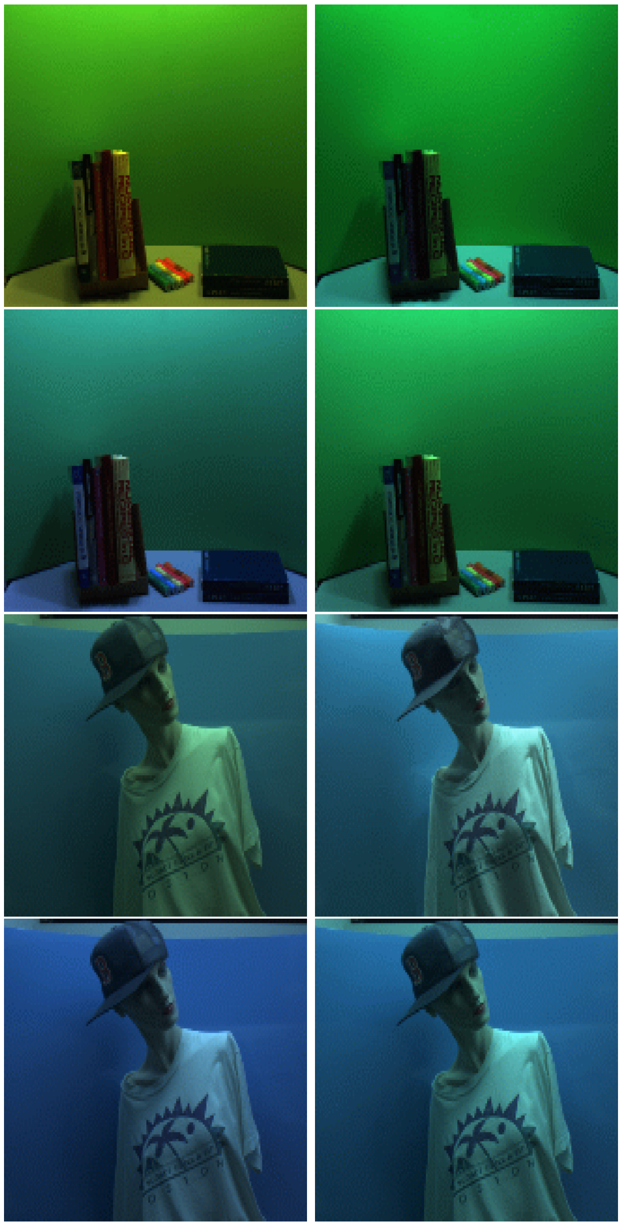

Figure 3: Algorithm performance for real images.

References

[1] K. Barnard. Computational Color Constancy: Taking Theory Into Practice. Master's thesis, Simon Fraser Univ., School of Computing Science, 1995.

[2] G. Buchsbaum. A spatial processor model for object colour perception. *Journal of the Franklin Institute*, 310:1–26, 1980.

[3] G. D. Finlayson. Color in Perspective. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(10):1034–1038, 1996.

[4] Graham Finlayson and Steven Hordley. Selection for Gamut Mapping Colour Constancy. In *Proceedings of the 8th British Machine Vision Conference*, volume 2, pages 630–639. BMVA, September 1997.

[5] D. A. Forsyth. A Novel Algorithm for Colour Constancy. *International Journal of Computer Vision*, 5(1):5–36, 1990.

[6] D. B. Judd, D. L. MacAdam, and G. W. Wyszecki. Spectral distribution of typical daylight as a function of correlated colour temperature. *Journal of the Optical Society of America, A*, 54:1031–1040, 1964.

[7] F.P. Preparata and M.I. Shamos. *Computational Geometry: An Introduction*. New York:Springer Verlag, 1985.

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

[3] A colour version of this paper is available at `http://colour.derby.ac.uk/colour/abstracts/98/hor98a.html`.