

## Chromagenic Colour Constancy

G. D. Finlayson, S. D. Hordley and P. M. Morovic

School of Computing Sciences,  
University of East Anglia,  
Norwich (United Kingdom)

Corresponding author: G.D. Finlayson ([graham@cmp.uea.ac.uk](mailto:graham@cmp.uea.ac.uk))

### ABSTRACT

A chromagenic camera takes two pictures of each scene. The first image is taken as normal and then the second is captured with a specially chosen chromagenic filter placed in front of the camera. In contradistinction to previous cameras that have more than 3 sensors, the aim of a chromagenic camera is not to measure more degrees of freedom in reflectance. Indeed, in a chromagenic camera the RGBs are, to a first approximation, linearly related to the filtered RGBs. However, the chromagenic filter is chosen so that this relationship depends on, and varies with illumination. The chromagenic camera is sensitive to the degrees of freedom in *illumination*. Chromagenic colour constancy proceeds in two stages. In pre-processing, for each light, the relation that takes filtered to unfiltered RGBs is computed. The input to colour constancy processing comprises the unfiltered and filtered RGBs captured for a given scene under unknown lighting conditions. To estimate the illuminant, the filtered responses are transformed by the pre-computed relations, and then these estimates are compared to the unfiltered counterparts. The transform that best predicts the data identifies the illuminant. Remarkably, this very simple approach works as well as, or better than, all other algorithms tested.

### 1. INTRODUCTION

From physics we know that when blue and red lights strike a white piece of paper, the reflected spectra are themselves blue and red respectively. Yet in both cases we will see the paper as being white. The mechanisms through which we might achieve this *colour constancy* are the subject of much research in the fields of human and artificial vision.

Simple approaches to solving colour constancy can deliver surprisingly good performance. Indeed, by simply dividing each image RGB by the average or maximum  $R$ ,  $G$  and  $B$  in the image we can often achieve good colour constancy (though of course it is easy to find images where such a simple approach cannot work). More advanced algorithms achieve better performance by incorporating additional information into the problem formulation. Specifically, it is reasonable to assume that some RGBs can occur under some lights but not others (the bluest blue response cannot occur under the reddest light) and also that some colours are themselves more or less likely (ultra saturated reds do not occur very often in nature). These constraints form the basis of Gamut Mapping<sup>1</sup> and Color by Correlation<sup>2</sup> respectively. With respect to standardised tests, these algorithms work much better than the simple approaches and, moreover, they also significantly outperform a large number of other published algorithms.

In this paper we consider how we might change the image capture in such a way as to make the colour constancy problem easier to solve. Specifically, we investigate whether colour constancy is a simpler problem if instead of measuring only  $R$ ,  $G$  and  $B$  we also measure  $R^F$ ,  $G^F$  and  $B^F$ : the same image captured through a chromagenic<sup>1</sup> coloured filter. This idea is in fact not new. Several authors have considered the multi-illuminant colour constancy problem: is it easier to recover true surface colour if we see the same surfaces under two or more lights?

Because, to a first approximation, an image formed by placing a filter in front of a camera is the same as the image formed when the same filter is placed in front of the light, chromagenic colour constancy is a special case of the 2-light colour constancy problem. However, since we know what

---

<sup>1</sup> The term “chromagenic” comes from “Chromagen”. Chromagen contact lenses (specially chosen colored filters) are prescribed to improve the vision of colour deficient and dyslexic observers.

our filter is, the second light cannot be truly independent of the first. This is the key insight in the chromagenic approach. Since we know the filter, and we know the typical lights that might impinge upon a scene, we should be able to predict the filtered RGBs (after all, we know that we can discount the effect of illumination by applying a suitable mapping<sup>1,2</sup>). So long as different lights imply different predictions, colour constancy is achieved by testing different predictions *in situ* and choosing the one that best fits the data.

## 2. CHROMAGENIC THEORY

**Theory:** Marimont and Wandell<sup>3</sup> showed that the dimensionality of light and surface, in terms of how they combine and project to form RGBs, could be modelled by 3-dimensional linear models. Relative to this model we can define an RGB response as:

$$\underline{p} = \Lambda(\underline{\varepsilon})\underline{\sigma}, \quad \Lambda(\underline{\varepsilon}) = \sum_{i=1}^3 \Lambda_i \varepsilon_i \quad (1)$$

where  $\underline{p}$  and  $\underline{\sigma}$  denote an RGB vector and a 3-dimensional weight vector describing the colour of the surface. The  $3 \times 3$  Lighting matrix  $\Lambda$  is itself a linear sum of 3 basis matrices.

In a chromagenic camera we propose that at each pixel we have an RGB plus the RGB measured after placing a coloured filter in front of the camera. That is, we have 6 measurements at each point in the image. We can extend the lighting matrix formalism and describe the 6-response for a chromagenic camera as:

$$\begin{bmatrix} \underline{p} \\ \underline{p}^F \end{bmatrix} = \begin{bmatrix} \Lambda(\underline{\varepsilon}) \\ \Lambda^F(\underline{\varepsilon}) \end{bmatrix} \underline{\sigma} \quad (2)$$

Equation (2) is interesting as it demonstrates that, although we are making six measurements, colour formation is, for a fixed light, a 3-dimensional process. Another way to demonstrate this inherent 3 dimensionality is to write the measured RGBs in terms of their filtered counterparts:

$$\underline{p} = \Lambda(\underline{\varepsilon})[\Lambda^F(\underline{\varepsilon})]^{-1} \underline{p}^F, \quad \underline{p} = T(\underline{\varepsilon})\underline{p}^F \quad (3)$$

Clearly, if for two different lights, (defined by the epsilon vector) the corresponding transforms  $T(\underline{\varepsilon}_1) = T(\underline{\varepsilon}_2)$  then we will not be able to solve for colour constancy. Thus, a prerequisite for solving for colour constancy is that for different lights the corresponding filtered to unfiltered transform should be unique. Assuming 3-dimensional models of light and surface we show elsewhere that for typical cameras uniqueness is guaranteed, so that it is plausible that we can solve for colour constancy<sup>4</sup>. However, it is worth considering when constancy is not possible. Trivially, if we place a neutral density filter in front of the camera then all transforms will be proportional (in the same proportion) to the identity matrix. While this represents the worst case, empirically it can be shown that some filters support better constancy computation than others. We propose the following figure of merit to gauge the effectiveness of a filter  $F$ :

$$\text{Goodness}(F) = \frac{\sum_j \sum_i |p_{i,j} - T_j^F p_{i,j}^F|}{\sum_j |T_j^F|} \quad (4)$$

The transforms  $T_j^F$  depend on the filter  $F$  and the light  $j$ , and are found by regression on a training set of RGBs and filtered RGBs. The goodness measure is calculated over a set of training surfaces indexed by  $i$ . The denominator term of (4) is a measure of the spread of the transforms: we wish different lights to require different transforms. The numerator term measures the error in

mapping filtered RGBs to unfiltered counterparts (this will be non zero as a 3-dimensional model of light and surface is approximate and there will be noise in realistic imaging conditions).

### 3. CHROMAGENIC ALGORITHM

Let us assume that the model in (1) approximately holds and that unique transforms are associated with known lights. Then we can solve for colour constancy by applying the following pre-processing and operation steps:

**Pre-processing:** Let  $P_j$  and  $P_j^F$  denote  $N \times 3$  matrices of RGBs, filtered and unfiltered, for  $N$  surfaces measured under the  $j$ th, of  $M$ , measured lights. We calculate the corresponding transform matrices:

$$T_j = [P_j^F]^+ P_j \quad (5)$$

where  $+$  denotes pseudo inverse.

**Operation:** Let  $Q_j$  and  $Q_j^F$  denote  $n \times 3$  matrices of RGBs, filtered and unfiltered, for  $n$  surfaces measured under an unknown illuminant. The estimated illuminant is found by minimising:

$$\min_j |Q_j - Q_j^F T_j| \quad (6)$$


Notice how simple the chromagenic algorithm is: it can be written in just a few lines of code (for example, in a language like Matlab). This contrasts favourably with other recent theory including gamut mapping<sup>1</sup> and Color by Correlation<sup>2</sup> which require more sophisticated programs and more involved pre-processing.

### 4. RESULTS

The Simon Fraser testing protocol<sup>1</sup> is now widely used to test the performance of colour constancy algorithms. The test involves a training and operation phase. In the training phase, algorithm parameters are estimated given knowledge of 1995 measured reflectances, the spectral sensitivities of a camera and 87 test lights. In the operation phase 287 test lights are used (which include the 87 training lights) but the reflectances and camera are the same as for pre-processing. Synthetic images are created for 1000 images with respectively, 2, 4, 8, 16, 32 or 64 distinct surfaces (randomly drawn from the 1995 reflectances). For each of the 6000 images, 6 algorithms are used to estimate the illuminant colour.

Each algorithm returns an estimate of the RGB of the illuminant for a given synthetic image. This is then compared to the true RGB calculated through direct measurement. Because we cannot recover the intensity of the estimate, the error in estimation is defined to be the angle between the true and estimated light. Median angular error measures, as a function of the number of surfaces in the scene, are shown for a variety of algorithms in Figure 1.

It is clear that the simple algorithms, Grey-World, Modified Grey-World and Max RGB perform less well than the other more advanced algorithms (see Barnard *et al*<sup>1</sup> for a review of these algorithms). It is also clear that the chromagenic algorithm performs better for small numbers of surfaces than all other algorithms. We carried out statistical testing to see which algorithms performed better, or worse, in a statistical sense. We applied the Wilcoxon Sign Test and found that the chromagenic algorithm delivers statistically better performance, at the 99% significance level, than all other algorithms except the probabilistic Color by Correlation approach where performance was found to be similar.



**Figure 1:** The Performance of 6 colour constancy algorithms.

In other work<sup>4</sup> we have also extended the basic theory in a Gamut Mapping version of the chromagenic algorithm (we incorporate knowledge of which RGBs can physically occur under particular lights). Here we found chromagenic performance is significantly better than all other algorithms. The chromagenic algorithm has also been tested on a set of real images with good results.

#### 4. CONCLUSIONS

A chromagenic camera takes two pictures of each scene: one as normal and another through a specially chosen filter. The filter is chosen so that the relationship between filtered and unfiltered RGBs depends on, and varies with, illumination. By testing pre-computed relations *in situ*, it was shown how it is possible to estimate the illumination. Experiments demonstrate that chromagenic colour constancy outperforms other algorithms

#### Acknowledgements

This research is supported by EPSRC grant GR/R65978. G. D. Finlayson also gratefully acknowledges the support of the Leverhulme trust.

#### References

1. K. Barnard, V. Cardei and B. Funt, "A Comparison of Computational Color Constancy Algorithms (I and II)," IEEE Trans. Image Processing, 972-996 (2002).
2. G. Finlayson, S. Hordley and P. Hubel, "Color by Correlation: A simple, Unifying Framework for Color Constancy," IEEE Trans. Pattern Analysis and Machine Intelligence, 1209-1221 (2001).
3. D. Marimont and B. Wandell, "Linear Models of Surface and Illuminant Spectra," Journal of the Optical Society of America A, 1905-1913 (1992).
4. G. Finlayson, S. Hordley and P. Morovic, "Solving for Colour constancy using the chromagenic constraint," IEEE Conference on Computer Vision and Pattern Recognition, (June 2005).