

Combining physical and statistical evidence for computational colour constancy

G. Schaefer¹, S.D. Hordley² and G.D. Finlayson²

¹*School of Computing and Informatics, Nottingham Trent University, Nottingham, U.K.*

²*School of Computing Sciences, University of East Anglia, Norwich, U.K.*

Corresponding author: G. Schaefer (Gerald.Schaefer@ntu.ac.uk)

ABSTRACT

Colour constancy algorithms are typically divided into physics-based techniques which exploit the physical image formation process and statistics-based method that use statistical knowledge on surfaces and lights. In this paper we introduce a combined physical and statistical algorithm for computational colour constancy. Estimates obtained from the statistical Colour by Correlation algorithm are integrated with those of a physics based method based on the dichromatic reflectance model. The algorithm not only provides an illuminant estimate but a complete set of likelihoods for a given set of reference lights. Recovery performance is shown to exceed that of purely statistical and purely physical algorithms.

1. INTRODUCTION

It is well known that the colour signals recorded by a digital camera do not depend solely on the surface characteristics of the captured objects. Rather the RGBs, the red, green and blue responses obtained by such a device, are a function of surface colour, scene illuminant and imaging geometry. It follows then that two identical scenes captured under two different light sources will result in two different digital images. In contrast, the human visual system seems capable of eliminating the effect of the scene illuminant – colours are perceived independent of the light source, an ability that is referred to as *colour constancy*.

Computational colour constancy, the equivalent of this ability for computer vision systems, is often defined as estimating the colour of the scene illuminant. Once this is known an image can be easily mapped to the same scene as captured under a different light [6]. Colour constancy algorithms can be divided into two kinds of approaches: physics-based and statistics-based algorithms [4]. Physical techniques are based on the underlying model of image formation and how certain physical processes such as highlights or interreflections manifest themselves in the resulting images. In contrast, statistical algorithms try to correlate the colours observed in an image with statistical knowledge on common surfaces and lights. Both approaches have their advantages and disadvantages. While statistics-based methods work well if there are many surfaces in a scene, they provide poor colour constancy for images with low colour diversity. On the other hand, physics-based methods are theoretically able to estimate the scene illuminant from two surfaces only [10] yet seem only to work well under controlled imaging condition [4].

In this paper we introduce a novel combined physical and statistical approach to colour constancy which aims to integrate the advantages of both approaches. An instance of the statistics-based Colour by Correlation framework [3] is combined with a robust physics-based algorithm that relies on the dichromatic reflection model [8]. Colour by Correlation correlates the colours recorded from a scene with statistical knowledge about the distribution of common surface and illuminant colours. This statistical information is expressed in a correlation matrix which encodes the likelihoods that a certain observed colour signal originated from a scene taken under a given reference light. Integrating the correlation matrix with the colour distribution of the actual images produces a set of likelihoods which inform us for a given set of lights how likely it is that the scene was captured under each of those illuminants. The physics-based algorithm employed exploits the dichromatic reflection model which states that the colour signals of an object fall on a two-dimensional plane in sensor space with the illuminant vector embedded in this plane. The intersection of two or more such dichromatic planes yields the colour of the light source. Following a coarse segmentation to extract dichromatic

planes from the image all pairwise plane intersections are calculated yet only those that are deemed useful to provide a good illuminant estimate are utilised to form a set of likelihoods similar to the one generated by the correlation framework. In a final step, the statistics-based and physics-based likelihoods are integrated to provide a combined physical/statistical likelihood vector for a given set of reference lights. The maximum likelihood estimate provides a good solution for the colour constancy problem yet clearly more information is available which could be passed on for further analysis.

The remainder of this paper is organised as follows. Section 2 describes the correlation matrix and dichromatic colour constancy algorithms employed as well as our novel combined physical and statistical algorithm. Section 3 provides experimental results on a large benchmark image database confirming the superiority of the approach presented in comparison to purely physical and purely statistical algorithms while Section 4 concludes the paper.

2. METHOD

2.1 Statistics-based Colour by Correlation

The Colour by Correlation approach [3] is based on the idea that the colours in an image hint at the possible scene illuminant. For instance, an image with a yellowish colour cast is more likely to have originated from a scene captured under yellow tungsten light than under bluish daylight. Colour by Correlation exploits this relationship by using the correlation between the colours recorded in an image and the range and probabilities of colours observable under each of a set of reference lights.

As a finite number of reference illuminants is chosen a priori, the algorithm is inherently discrete. Also, the choice of reference lights ensures that only physically possible lights are considered and is further being used to restrict the attention to common light sources only, which in turn increases the accuracy of the algorithm. At the heart of the correlation framework is a correlation matrix which encodes the probabilities that a certain colour was captured under a given light. In order to obtain this correlation matrix for a given device and illuminant set, a process that needs to be performed only once, one first has to choose an appropriate colour space and quantisation scheme thereof. For the experiments presented in Section 3 we have used a cube-root chromaticity space [3] which was uniformly quantised into 50x50 bins. Each colour space bin represents a row in the correlation matrix whereas each column describes one of the reference illuminants. Each entry of the matrix then holds the probability that a certain colour (encoded by its row number) was captured under a certain light (identified by its column number). Following a segmentation step, colours of an input image are also quantised according to the same scheme used for the correlation matrix. An image vector is then formed which simply encodes in binary fashion whether each colour appears in the image or not. Correlation is then performed by calculating the dot product between the illuminant columns of the correlation matrix and the image vector which results in a vector of likelihoods.

This vector contains one entry for each reference light which holds the likelihood that the given image was captured under this light. While the maximum likelihood estimate can be used to give a single answer to the colour constancy problem it is clear that the solution vector provides more information which could be exploited to obtain a confidence measure or which could be used by other algorithms for further processing.

2.2 Physics-based robust dichromatic colour constancy

The dichromatic reflectance model states that the colour signals of an object made of materials such as paints, plastics, papers etc. will fall on a two-dimensional plane in colour space [9]. Crucially, as the index of refraction is constant across the visible spectrum for these materials the vector of the scene illuminant is one of the vectors spanning this dichromatic plane. It follows then that colour constancy can be achieved by intersecting two or more dichromatic planes as the intersection will yield the illuminant vector [10]. Unfortunately, this approach has been shown to work well only under controlled conditions inside the lab but not on real images [4]. However, by adding some statistical knowledge on common light sources in form of an illuminant constraint much

improved recovery performance can be achieved [5], in fact, it has been shown that colour constancy performance similar to the best statistical methods is possible [7].

The dichromatic algorithm used in here follows the approach taken in [7]. A discrete set of reference lights is chosen a priori thus representing a statistical illumination constraint. The image is divided into small (30x30 pixel) non-overlapping sub-blocks. Each of these blocks is then tested for dichromatic properties where a block is deemed to be dichromatic if more than 98.5% of its variance is captured by the first 2 eigenvectors obtained from singular value decomposition. Next, all pairwise intersections of dichromatic planes are calculated resulting in a large set of illuminant estimates (each based on two image sub-blocks). However, due to noise and insufficient segmentation not all of these intersections will contain good estimates. Therefore all those vectors that fall outside the convex hull of a set of common lights [5] are discarded. Similarly, all those intersections are discarded where the corresponding dichromatic planes intersect each other at small angles (< 5 degrees) as these intersections are much more likely to have been affected by noise [4]. All the remaining plane intersections are combined to form a solution vector which contains one entry per reference light describing the likelihood that the image was taken under this light. These likelihoods are obtained by summing up the inverses of distances between each plane intersection and the reference light (over all intersections).

A single estimate can be derived as the maximum likelihood candidate. However, similar to the results provided by the Colour by Correlation algorithm a whole set of likelihoods is available rather than just a single solution.

2.3 Combined physical and statistical algorithm

Due to the inherent design of both the Colour by Correlation algorithm and the robust dichromatic colour constancy method combining the two is achieved in a rather simple way. Both algorithms provide a set of likelihoods for a given set of reference lights. Naturally, these reference illuminants are chosen to be identical for both algorithms. Each entry in the solution vector obtained from the physical dichromatic algorithm has then a corresponding counterpart in the Colour by Correlation solution vector.

Likelihoods are first normalised; then a non-linear (cube root) function is applied to attenuate any spikes in the distribution. After that a weighted sum of the two likelihood vectors is calculated (for the Experiments reported in Section 3 we weighted both algorithms equally) which yields a final, combined physical/statistical solution vector describing the likelihoods of the input image to have been captured under the reference lights. As above, the maximum likelihood estimate can be taken as a single solution to the colour constancy problem. However, again additional information is available. A few top ranking lights can be returned together with their corresponding likelihoods which can be used as a confidence measure. Alternatively, the complete solution vector can be used for further processing.

3. RESULTS

We used the Mondrian subset of the SFU image database [1] which provides a benchmark dataset for the evaluation of colour constancy algorithms. It consists of 22 objects captured under up to 11 different lights; in total there are 223 images. For each of the images the colour of the actual scene illuminant is known and serves as the ground truth for the dataset.

We estimate the scene illuminant using the Colour by Correlation algorithm as outlined in Section 2.1 and the robust dichromatic algorithm explained in Section 2.2 (all algorithms use the 87 reference lights used in [1]). The Colour by Correlation technique is known as one of the best performing colour constancy to date. In [7] it was shown (using the same dataset) to outperform gamut mapping colour constancy [6] which in turn was demonstrated to perform better than many other algorithms [1].

Finally we calculated the illuminant estimate based on the maximum likelihood solution of the combined physical/statistical algorithm introduced in this paper. The results, expressed in terms of median, mean and maximum error in angular degrees in RGB space between the estimate and the actual scene illuminant are shown in Table 1.

From there we can see that while the correlation matrix and the dichromatic algorithms perform fairly similar, the combined physical/statistical method clearly outperforms them both. The median error is only 2.2 degrees compared to 2.72 and 2.52 for the correlation framework and dichromatic colour constancy respectively. The differences in terms of mean angular error are even higher, here the 4.5 degrees obtained from the combined algorithm compare favourably with the 5.59 and 5.68 achieved by the purely statistical and purely physical approaches.

Table 1: Performance of combined physical/statistical colour constancy algorithm compared to purely statistics- and physics-based approaches based on an image database of 223 images. Results are given in terms of median, mean and maximum angular error over the whole dataset.

	median error	mean error	max. error
Colour by Correlation	2.72	5.59	28.78
Robust dichromatic colour constancy	2.52	5.68	36.98
Combined physical/statistical colour constancy	2.20	4.50	29.88

4. CONCLUSIONS

A combined physical and statistical approach to computational colour constancy has been presented. A variant of the statistics-based Colour by Correlation framework is integrated with a physics-based algorithm based on the dichromatic reflectance model. Experimental results on a large benchmark dataset confirm the superiority of this approach to purely statistical and physical algorithms.

References

- [1] K. Barnard, "Practical colour constancy," *PhD Thesis*, Simon Fraser University, School of Computing, 1999.
- [2] V.C. Cardei and B.F. Funt, "Committee based colour constancy," In *IS&T and SID's 7th Colour Imaging Conference*, pp. 311-313, 1999.
- [3] G. Finlayson, P. Hubel, and S. Hordley, "Colour by correlation: A simply, unifying approach to colour constancy," In *7th IEEE Int. Conference on Computer Vision*, pp. 835-842, 1999.
- [4] G. Finlayson and G. Schaefer, "Convex and non-convex illuminant constraints for dichromatic colour constancy," In *IEEE Int. Conference on Computer Vision and Pattern Recognition*, Vol. 1, pp. 598-604, 2001.
- [5] G. Finlayson and G. Schaefer, "Solving for colour constancy using a constrained dichromatic reflectance model," *Int. Journal of Computer Vision*, Vol. 42, No. 3, pp. 127-144, 2001.
- [6] D. Forsyth, "A novel approach for color constancy," *Int. Journal of Computer Vision*, Vol. 5, pp. 5-35, 1990.
- [7] S. Hordley and G. Finlayson, "Re-evaluating colour constancy algorithms," In *17th IEEE Int. Conference on Pattern Recognition*, 2004.
- [8] G. Schaefer, "Robust dichromatic colour constancy," In *Int. Conference on Image Analysis and Pattern Recognition*, Springer Lecture Notes in Computer Science, LNCS Vol. 3212, pp. 257-264, 2004.
- [9] S. Shafer, "Using color to separate reflection components," *Color Research Applications*, Vol. 10, No. 4, pp. 210-218, 1985.
- [10] H.C. Lee, "Method for computing the scene illuminant from specular highlights," *Journal of the Optical Society of America A*, Vol. 3, No. 10, pp. 1694-1699, 1986.