

Color Constancy with Fluorescent Surfaces

Kobus Barnard

School of Computing Science,

Simon Fraser University

Burnaby, British Columbia, Canada

kobus@cs.sfu.ca

Abstract

Fluorescent surfaces are common in the modern world, but they present problems for machine color constancy because fluorescent reflection typically violates the assumptions needed by most algorithms. The complexity of fluorescent reflection is likely one of the reasons why fluorescent surfaces have escaped the attention of computational color constancy researchers. In this paper we take some initial steps to rectify this omission. We begin by introducing a simple method for characterizing fluorescent surfaces. It is based on direct measurements, and thus has low error and avoids the need to develop a comprehensive and accurate physical model. We then modify and extend several modern color constancy algorithms to address fluorescence. The algorithms considered are CRULE and derivatives [1-4], Color by Correlation [5], and neural net methods [6-8]. Adding fluorescence to Color by Correlation and neural net methods is relatively straight forward, but CRULE requires modification so that its complete reliance on diagonal models can be relaxed. We present results for both synthetic and real image data for fluorescent capable versions of CRULE and Color by Correlation, and we compare the results with the standard versions of these and other algorithms.

Introduction

The image recorded by a camera depends on three factors: The physical content of the scene, the illumination incident on the scene, and the characteristics of the camera. It is the goal of computational color constancy to identify, separate, or mitigate the effects of these factors. Doing so has applications in computer vision and image reproduction. Here we address computational color constancy in the case where fluorescent surfaces may be present in the scene. Such surfaces are common because fluorescent inks are often used to provide strong color. However, machine color constancy research has not yet addressed the problem of fluorescent reflection, likely due to the difficulties presented. Nonetheless, we feel that it is necessary to investigate this problem because some of the most effective color constancy

algorithms are sensitive to fluorescent surfaces, and can have poor results when they are present.

Characterizing Fluorescent Surfaces

We begin by introducing a simple method for characterizing fluorescent surfaces. It is based on direct measurements, and thus has low error and avoids the need to develop, fit, and test physical models. Such models are necessarily quite complex and limited to the kinds of surfaces exhibiting the processes being modeled (an elegant model for one case is developed in [9]). We remind the reader that the key characteristic of fluorescent surfaces is that some of the light energy they absorb is re-emitted at longer wavelengths (lower energy). If we represent the incident light spectra as a vector of samples over wavelength, then reflectance can be described by the multiplication of that input vector by a triangular matrix. This is much more complex than the non-fluorescent case where a diagonal matrix is sufficient. Although it is possible to measure this matrix, doing this effectively requires equipment which is not readily available. Thus we introduce a more direct method for obtaining the data required.

Given a fluorescent surface candidate, we measure the spectra of the reflected light under a number of illuminants using a Photoresearch PR-650 spectroradiometer. We also measure the spectra of the illuminants providing the input energy to the fluorescent surface. Then, to simulate the surface under a new illuminant spectra, we first compute the positive linear combination of the test illuminants which is closest to the new illuminant spectra using constrained least squares optimization. The reflected energy of the fluorescent spectra under the new illuminant is then approximately that same linear combination applied to the measured test response spectra set. A simple example should make this clear. Assume that when the fluorescent surface is illuminated by a spectra A, the result is spectra A', and similarly, let B' be the response to stimulus B. Then if a illuminant C is roughly $A+2B$, then the response, C', is roughly $A'+2B'$. This procedure is used to simulate fluorescent reflection to obtain the data sets required by color constancy algorithms.

Color Constancy with Fluorescent Surfaces

We now turn to the algorithms themselves. We feel that the most effective computational color constancy methods currently available are CRULE and derivatives [1-4], Color by Correlation [5], and neural net methods [6-8]. Adding fluorescence to Color by Correlation and neural net methods is relatively straight forward. In either case, our characterization of fluorescent surfaces is used to augment the world used for training (neural nets) or building correlation matrices (Color by Correlation).

Extending Forsyth's CRULE method is more involved. This method explicitly assumes that an illumination change can be modeled by a diagonal transform. For example, if a the camera response to a white surface changes from (W_r, W_g, W_b) to (W_r', W_g', W_b') due to an illumination change, then we assume that the illumination change for the other (non-white) surfaces are modeled by multiplication by the diagonal matrix formed from the vector $(W_r'/W_r, W_g'/W_g, W_b'/W_b)$. The accuracy of this assumption in the non-fluorescent case has been well studied [10-13], and is known to be strongly dependent on the camera sensors. The problem is that the very nature of fluorescent reflection is contrary to this assumption. To deal with this we first propose a modification to Forsyth's method which makes it more resilient to diagonal model failure. Once we make this modification, incorporating fluorescent surfaces is straightforward.

We will now provide some additional details of the extension beginning with a brief review of Forsyth's method [1]. First we form the set of all possible RGB due to surfaces in the world under a known, "canonical" illuminant. This set is convex and is represented by its convex hull. The set of all possible RGB under the unknown illuminant is similarly represented by its convex hull. Under the diagonal assumption of illumination change, these two hulls are a unique diagonal mapping (a simple 3D stretch) of each other.

Figure 1 illustrates the situation using triangles to represent the gamuts. In the full RGB version of the algorithm, the gamuts are actually three dimensional polytopes. The upper thicker triangle represents the unknown gamut of the possible sensor responses under the unknown illuminant, and the lower thicker triangle represents the known gamut of sensor responses under the canonical illuminant. We seek the mapping between the sets, but since the one set is not known, we estimate it by the observed sensor responses, which form a subset, illustrated by the thinner triangle. Because the observed set is normally a proper subset, the mapping to the canonical is not unique, and Forsyth provides a method for effectively computing the set of possible diagonal maps. (See [1-4, 12] for more details on gamut mapping algorithms).

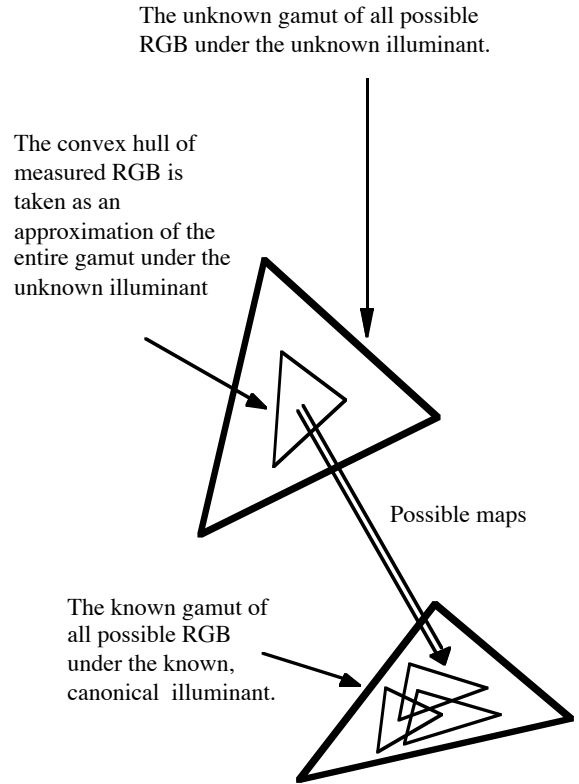


Figure 1: Illustration of the basic idea of gamut mapping color constancy.

Once the set of possible maps has been computed, an important second stage of the algorithm is to choose a solution from the feasible set. Several different methods for doing this lead to different variants of the algorithm [1, 3, 4, 14]. Another group of variants work in an appropriate chromaticity space rather than RGB [2]. Finally, Finlayson showed that it is possible to further constrain the solution by restricting the solutions to those corresponding to common or likely illuminants [2]. We will make use of this extra constraint in this study, and we will denote algorithms using them as "extended" CRULE, or E-CRULE for short.

We now consider the case where the diagonal model is less appropriate. Here it may be possible that an observed set of illuminants does not map into the canonical set with a single diagonal transform. This corresponds to an empty solution set. In earlier work we forced a solution by assuming that such null intersections were due to measurement error, and we thus increased various error estimates until a solution was found. However, this method does not give very good results in the case of extreme diagonal failures, such as those due to fluorescent surfaces.

To deal with this problem, we propose the following modification. Consider the gamut of possible RGB under a single test illuminant. Call this the test illuminant gamut. Now consider the diagonal map which takes the RGB for white under the test illuminant to the RGB for white under the canonical illuminant. If we apply that diagonal map to our test illuminant gamut, then we

will get a convex set similar to the canonical gamut, the degree of difference reflecting the failure of the diagonal model. If we extend the canonical gamut to include this mapping of the test set, then there will always be a diagonal mapping from the observed RGB of scenes under the test illuminant to the canonical gamut. We repeat this procedure over a representative set of illuminants to produce a canonical gamut which is applicable to those illuminants as well as any convex combination of them. The basic idea is illustrated in Figure 2. Note that the test illuminant gamuts can include fluorescent surfaces modeled by the method described in the preceding section.

The gamuts of all possible RGB under three training illuminants.

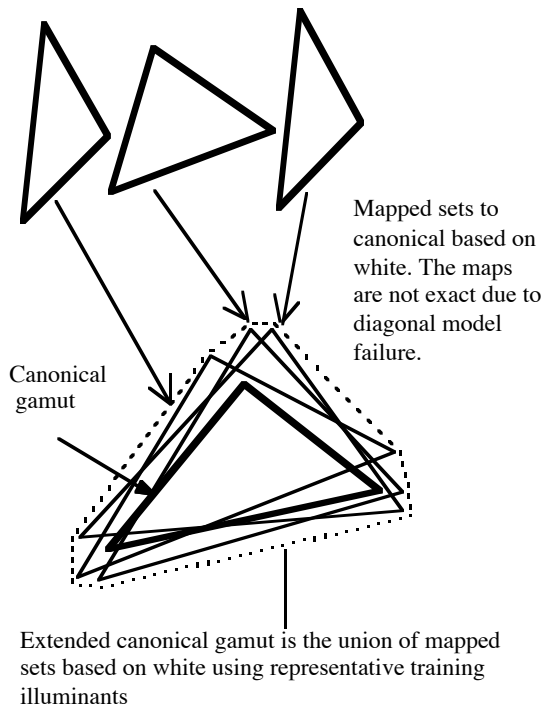


Figure 2: Illustration of the modification to the gamut mapping method to enable the handling of fluorescent surfaces.

If the diagonal model holds fairly well, as is the case with non-fluorescent surfaces and our Sony DXC-930 video camera [3, 13], then the canonical hull is extended slightly, and under some conditions several variants give better results than the same algorithms without the modification, especially as the number of available colors increases. If we model fluorescent surfaces, then the canonical gamut will be extended a fair amount. In this case, the performance on data devoid of fluorescent surfaces is slightly degraded, as the constraints on the solutions are less strict, but when there are fluorescent surfaces the performance can be substantially better.

Experimental Results

As an initial step in our investigation of fluorescence we measured a number of candidate surfaces, and trimmed this set down to 9 strongly fluorescent ones. These included 3 printed surfaces from a laundry detergent box, 2 surfaces from a multi-colored child's cloth ball, 2 different colors of flagging tape, and 2 different vividly colored pieces of paper. As described above, our method of characterizing the fluorescent surfaces required measuring their reflectance spectra under a number of representative illuminants. For non-fluorescent spectra we used a set of roughly 2000 spectra collected from several sources.

The illuminant set for algorithm calibration (training) was carefully chosen in the following manner. We started with 11 illuminants selected to cover, with relatively uniformity, the region of common natural and man made illuminants in the $(R/(R+G+B), B/(R+G+B))$ chromaticity space. This set of 11 illuminants is also the one used for the image data. We then added additional measured illuminants, and linear combinations thereof, to complete the uniform coverage with higher density. This second set of illuminants included both additional common sources and the illumination found in 90 random indoor and outdoor locations. The resultant set was used for the construction of gamuts and correlation matrices, and thus played the role of a training set. The illuminant set used for testing in the case of generated data was constructed in the same manner, but the chromaticity space was filled 4 times more densely.

Our camera was calibrated as described in [15]. We used the resultant camera model for synthetic scene generation, as well as the computation of the various "training" data sets mentioned above. Thus the same algorithm calibrations were used for both synthetic and image data experiments.

For the purposes of this study, we will assume that the goal of the algorithms is to estimate the response of the vision system to a perfect white patch. However, it is often the case that we are most interested in the chromaticity of the illuminant, and several of the algorithms of interest only compute the illuminant chromaticity. Hence, we only report chromaticity results. The specific error metric used considers the illuminant RGB and the corresponding estimate thereof as vectors in RGB space, and computes the angle between these two vectors in degrees.

We first present some results using generated data. The use of generated data eliminates calibration problems, and simplifies analysis of the effects of statistical assumptions. We present the results of the various algorithms when subjected both to completely non-fluorescent data, as well as a mix of fluorescent and non-fluorescent data. In each case, 8 randomly chosen surfaces was used. For the second case we arranged for the fluorescent surfaces to be represented roughly 30% of the time.

The results for generated data are shown in Table 2. The first conclusion is that the presence of fluorescent surfaces does, as predicted, degrade every algorithm not

designed to deal with them. On the other hand, the extensions to gamut mapping and Color by Correlation work well to reduce this performance degradation. This reduction is most extreme with the extended Color by Correlation method, but this is likely due in part to an unnatural advantage that it does not enjoy in the case of real image data. Specifically, the second Color by Correlation algorithm was trained on data statistically similar to the test data. What is more promising is that the Color by Correlation method trained with fluorescent surfaces works well on data devoid of such surfaces, fairs better than the gamut mapping algorithms in this regard. Finally, it should also be noted that our gray world algorithm has access to the actual average of our data base of non-fluorescent surfaces, and thus has a unnatural advantage in the non-fluorescent case.

We have also tested algorithms on real image data. We constructed 7 scenes which all included known or suspected fluorescent surfaces, and took images of these under 11 different illuminants, resulting in 77 images. Seven images were culled due to problems with the experiment, leaving a total of 70 input images. Again, we feel it necessary to look at the performance of the algorithms when fluorescent surfaces are absent. Thus we also present the results for 321 input images from 33 scenes which are relatively free of fluorescent surfaces. The dynamic range of all images was extended using reduced illumination levels and averaging multiple frames. This gives us the opportunity to explore color constancy in the context of a high dynamic range vision system, as well as more standard vision systems which can be simulated by truncating the higher range data. The effect on the results is to give the Retinex based algorithm, and the maximum volume algorithms, an advantage. This is especially true when there are specularities.

In general, results from this real image data demonstrate that modeling fluorescence is again beneficial, although the large improvement in the case of Color by Correlation has been reduced to quite a modest increase. This is likely due in part to the mismatch between the statistics used for training and the somewhat arbitrary statistics in the image data. We also note that Color by Correlation has many possible implementations, and we are still working on finding a robust set of parameters for that algorithm. In the case of the gamut mapping algorithms, we see that the performance on the real image data is excellent. As noted above, the extended dynamic range of the data, enables the maximum volume algorithms to use specularities, which are often present in real image data, for illuminant chromaticity estimation. However, it is interesting to note that the Retinex algorithm, which does very well on the non-fluorescent data for the same reason, is badly degraded when used on the fluorescent data. The FL-ECRULE-MV algorithm on the other hand, handles both cases well, and is the overall top performer on our image database.

Table 1: Key to algorithms

ECRULE	CRULE with illumination constraint
MV	Solutions are chosen by max volume heuristic
AVE	Solutions are the average over the feasible set
FL	Algorithm is extended for fluorescence.
Retinex	Estimate illuminant by the max RGB in each channel.
Gray World	Estimate illuminant color by image average
C-by-C	Color by Correlation [5], with a Gaussian mask to smooth the correlation matrix and maximum likelihood estimate. For the FL-C-by-C variant, an abundance of fluorescent surfaces are included in the construction of the correlation matrix
Neural Net	Neural net trained to estimate illuminant chromaticity based on the observed image colors [6-8].

Table 2: Average angular error in RGB space of illuminant estimate (generated data)

	No Fluorescence	With 30% Fluorescence
ECRULE-MV	6.2	9.9
ECRULE-AVE	6.7	9.1
FL-ECRULE-MV	8.2	7.1
FL-ECRULE-HA	9.6	7.7
Retinex	9.7	13.1
Gray World	6.7	12.0
C-by-C	5.9	10.6
FL-C-by-C	6.3	4.1
Neural Net	5.5	6.9

Table 3: Average angular error in RGB space of illuminant estimate (image data)

	Scenes without fluorescent surfaces	Scenes with fluorescent surfaces
ECRULE-MV	5.3	13.2
ECRULE-HA	6.2	10.7
FL-ECRULE-MV	6.4	10.7
FL-ECRULE-AVE	9.0	10.3
Retinex	7.8	18.0
Gray World	12.0	17.5
C-by-C	10.2	12.0
FL-C-by-C	10.0	11.7
Neural Net	9.8	11.4

Conclusions

We have shown how to modify the three leading machine color constancy methods to deal with fluorescent surfaces. Dealing with such surfaces has been ignored until now, but we argue that doing so is important, as such surfaces are common in the modern world, and yet they dramatically degrade the performance of existing algorithms. Although further work is needed to estimate the frequency of occurrence of such surfaces, we pass on to the reader the following anecdotal datum. Our interest in exploring fluorescent surfaces arose because such surfaces were present in 20% of the randomly constructed scenes used to provide preliminary data for research into color constancy performance. Clearly we had to deal with fluorescence before we could proceed towards our goal of having colour constancy algorithms for real world applications.

Acknowledgments

We are grateful for the support of Hewlett-Packard Corporation and the Natural Sciences and Engineering Council of Canada. In addition we acknowledge the efforts of Lindsay Martin and Adam Coath who helped greatly with the data collection.

References

1. D. Forsyth, A novel algorithm for color constancy, *International Journal of Computer Vision*, **5**, pp. 5-36 (1990).
2. G. D. Finlayson, Color in perspective, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **18**, pp. 1034-1038 (1996).
3. K. Barnard, Computational colour constancy: taking theory into practice, : Simon Fraser University, School of Computing (1995).
4. G. Finlayson and S. Hordley, A theory of selection for gamut mapping colour constancy, *Proc. IEEE Conference on Computer Vision and Pattern Recognition* (1998).
5. G. D. Finlayson, P. H. Hubel, and S. Hordley, Color by Correlation, *Proc. IS&T/SID Fifth Color Imaging Conference: Color Science, Systems and Applications*, pp. 6-11 (1997).
6. B. Funt, V. Cardei, and K. Barnard, Learning Color Constancy, *Proc. IS&T/SID Fourth Color Imaging Conference: Color Science, Systems and Applications*, pp. 58-60 (1996).
7. V. Cardei, B. Funt, and K. Barnard, Modeling color constancy with neural networks, *Proc. International Conference on Vision Recognition, Action: Neural Models of Mind and Machine* (1997).
8. V. Cardei, B. Funt, and K. Barnard, Adaptive Illuminant Estimation Using Neural Networks, *Proc. International Conference on Artificial Neural Networks* (1998).
9. P. Emmel and R. D. Hersch, Spectral colour prediction model for a transparent fluorescent ink on paper, *Proc. IS&T/SID Sixth Color Imaging Conference: Color Science, Systems and Applications*, pp. 116-122 (1998).
10. G. D. Finlayson, M. S. Drew, and B. V. Funt, Spectral Sharpening: Sensor Transformations for Improved Color Constancy, *Journal of the Optical Society of America A*, **11**, pp. 1553-1563 (1994).
11. G. D. Finlayson, M. S. Drew, and B. V. Funt, Color Constancy: Generalized Diagonal Transforms Suffice, *Journal of the Optical Society of America A*, **11**, pp. 3011-3020 (1994).
12. G. D. Finlayson, Coefficient Color Constancy, : Simon Fraser University, School of Computing (1995).
13. K. Barnard and B. Funt, Experiments in Sensor Sharpening for Color Constancy, *Proc. IS&T/SID Sixth Color Imaging Conference: Color Science, Systems and Applications*, pp. 43-46 (1998).
14. B. Funt, K. Barnard, and L. Martin, Is Colour Constancy Good Enough?, *Proc. 5th European Conference on Computer Vision*, pp. I:445-459 (1998).
15. K. Barnard and B. Funt, Camera calibration for color research, *Proc. Human Vision and Electronic Imaging IV*, pp. 576-585 (1999).