Color Constancy with Specular and Non-Specular Surfaces

Kobus Barnard and Brian Funt
School of Computing Science,
Simon Fraser University
Burnaby, British Columbia
Canada V5A 1S6
{kobus, funt}@cs.sfu.ca

Abstract

There is a growing trend in machine color constancy research to use only image chromaticity information, ignoring the magnitude of the image pixels. This is natural because the main purpose is often to estimate only the chromaticity of the illuminant. However, the magnitudes of the image pixels also carry information about the chromaticity of the illuminant. One such source of information is through image specularities. As is well known in the computational color constancy field, specularities from inhomogeneous materials (such as plastics and painted surfaces) can be used for color constancy. This assumes that the image contains specularities, that they can be identified, and that they do not saturate the camera sensors. These provisos make it important that color constancy algorithms which make use of specularities also perform well when the they are absent. A further problem with using specularities is that the key assumption, namely that the specular component is the color of the illuminant, does not hold in the case of colored metals.

In this paper we investigate a number of color constancy algorithms in the context of specular and non-specular reflection. We then propose extensions to several variants of Forsyth's CRULE algorithm [1-4] which make use of specularities if they exist, but do not rely on their presence. In addition, our approach is easily extended to include colored metals, and is the first color constancy algorithm to deal with such surfaces. Finally, our method provides an estimate of the overall brightness, which chromaticity-based methods cannot do, and other RGB based algorithms do poorly when specularities are present.

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. This leads to a problem for many applications where the main interest is in the physical content of the scene. Consider, for example, a computer vision application which identifies objects by color. If the colors of the objects in a database are specified for tungsten illumination (reddish), then object recognition can fail when the system is used under the very blue illumination of a clear sky. This is because the change in the illumination affects object colors far beyond the tolerance required for reasonable object recognition. Thus the illumination must be controlled, determined, or otherwise taken into account.

The ability of a vision system to diminish, or in the ideal case, remove, the effect of the illumination, and therefore "see" the physical scene more precisely, is called color constancy. There is ample evidence that the human vision system exhibits some degree of color constancy. Interest in human vision, as well as robotics and image reproduction applications, has led to much research into computational methods to achieve color constancy. In this paper we build on this body of work, and propose and algorithm which combines the strengths of two different approaches to color constancy. We combine the information inherent in collections of matte surfaces and the information inherent in specularities. Interestingly, our method is easily extended to work with specularities from colored metals such as copper, and is the only method we know of which does so (but see [5-9] for related work).

The use of specularities for machine color constancy has its origin in the dichromatic model of reflectance [10, 11]. This model separates the light reflected from inhomogeneous materials such as plastics and paints into a diffuse (body) component, and a specular (interface) component. The body reflection blends the spectral reflectance properties of the object with that of the illumination, whereas the specular component has the same spectral makeup as the illuminant. Reflections from different parts of the same surface have varying amounts of the two reflection components due to changes in geometry, and various researchers have used this property to estimate the illuminant color [10, 12-16]. Alternatively, since the maximal specular reflection is typically much larger than the body reflection, a bright specularity can be a good estimate of the illuminant color as is, if it can be identified as a specularity. Either way, using specular reflection for color constancy typically requires an implicit physical segmentation of the image pixels, and the difficulties in doing this have, in part, inspired the present work.

In this paper we investigate a number of color constancy algorithms in the context of specular and non-specular reflection. We then propose extensions to several variants of Forsyth's CRULE algorithm [1-4] which make use of specularities if they exist, but do not rely on their presence. In addition, our approach is easily extended to include colored metals, and is the first color constancy algorithm to deal with such surfaces. Finally, our method provides an estimate of the overall brightness, which chromaticity-based methods cannot do, and other RGB based algorithms do poorly when specularities are present.

Approaches to Computational Color Constancy

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. This response will loosely be referred to as the color of the illuminant. It is most natural for that response to be the same dimension as the number of sensors in the vision system, and thus, for a standard color camera, the response would be the (R,G,B) of a white patch under that illuminant. However, it is often the case that we are most interested in the chromaticity of the illuminant, and an estimate of that chromaticity will suffice. This being the case, a number of color constancy algorithms have been developed which work entirely in some chromaticity spaces [2, 4, 17-20], and much progress has been made by taking advantage of the simplifications afforded by this strategy.

Nonetheless, if we now consider the case where specularities are present, we observe that certain RGB based algorithms, such as the original CRULE algorithm, estimate the illuminant chromaticity surprisingly well—even though they were not designed to optimize chromaticity estimation [3, 21]. The success of these algorithms when specularities are present is limited by the dynamic range of the vision system. We expect more dynamic range to become available to machine vision systems (see [22] for information about one high dynamic range camera), but currently, specularities tend to be clipped, and such pixels must be excluded as unreliable. As clipping becomes severe, these methods degrade, especially Retinex [21]. We also note that using these algorithms for illumination brightness estimation fails when strong specularities are present.

Chromaticity-based approaches, on the other hand, cannot use specular information on a pixel by pixel basis, and cannot provide illuminant brightness estimation. However, as noted above, we are often most interested in illuminant chromaticity estimation, and these approaches tend to be robust with respect to specularities. This is because specularities in chromaticity space simply desaturate colors, leading to colors which are perhaps less useful to the algorithm, but are nonetheless plausible [2], and thus the degradation is graceful. The essence of this observation also applies in the case of colored metals.

In contrast to the above algorithms (and others), which we analyze post hoc with respect to their abilities to ignore or take advantage of specularities, several researchers have developed computational color constancy methods which explicitly use and rely on specularities [10, 12-16]. In favorable situations, these methods can work well, but strong specularities are not always present, and as noted above, are often clipped. Furthermore, specularities from colored metals are not the same color as the illuminant, and these methods do not address this. These considerations lead us to propose extensions to several of the variants of Forsyth's CRULE method which take advantage of specularities if they exist, but continue to be strong algorithms if there are no specularities present.

Extending CRULE for specularities

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. The set is convex and is represented by its convex hull. We will refer to this set as the canonical gamut. The set of all possible RGB under the unknown illuminant is similarly represented by its convex hull. Now, under the diagonal assumption of illumination change, these two hulls are a unique diagonal mapping (a simple 3D stretch) of each other. To understand this assumption further, suppose that the RGB of white under the unknown illuminant is (Wr, Wg, Wb), and the RGB of white under the canonical illuminant is (Wr', Wg', Wb'). Then the RGB of white in the unknown gamut is mapped to the corresponding RGB in the canonical gamut by multiplication by the matrix DIAG(Wr'/Wr, Wg'/Wg, Wb'/Wb). To the extent that this same mapping applies to other, non-white surfaces, we say that we have a diagonal model of illumination change. The efficacy of this model is partly a function of the vision system sensors, and is a good approximation for our camera.

The gamut mapping strategy is to constrain the set of possible diagonal maps, with each map corresponding to an illuminant estimate. Figure 1 illustrates the situation using triangles for the gamuts. 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, 23] for more details on gamut mapping algorithms). Another important contribution was the observation that the set of maps could further be constrained by restricting them to ones corresponding to common or expected 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.

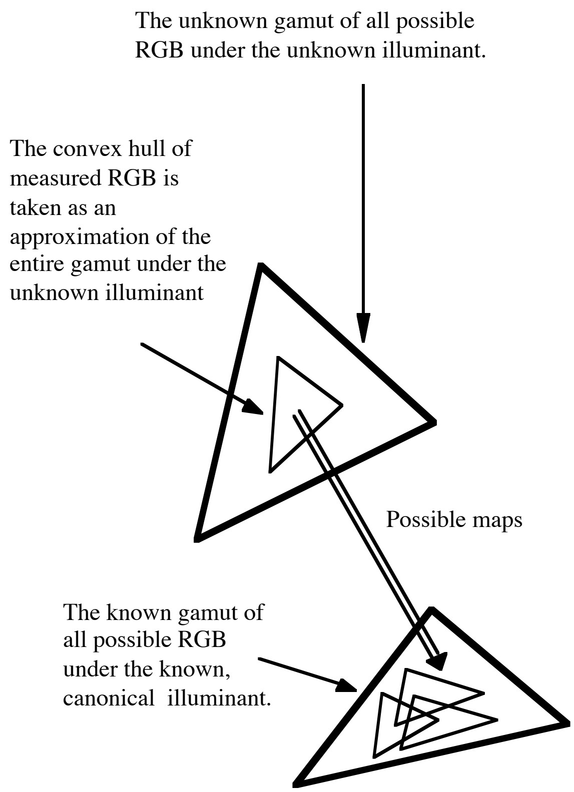

Figure 1: Illustration of the basic idea of gamut mapping color constancy.

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. The original method was to choose the mapping which maximized the volume of the mapped set. Although it was not designed to do so, we have observed that this method is a good choice for estimating the illuminant chromaticity, especially in the presence of specularities. (Note that the selected diagonal map implicitly specifies an estimate of the illuminant color as defined above). A second method for choosing a solution is to average the possible maps. This method potentially gives a better estimate of the illuminant RGB [3, 21], and can be more robust under clipping.

In order to use Forsyth's method in the case of specularities, we model specular reflection and extend the canonical gamuts appropriately. The canonical gamuts are polytopes in RGB space, having roughly the shape of two multi-faceted pyramids which are joined together at their identical bases. We normally include the origin as one of the vertices (and thus it is the apex of one of the pyramids), because, a priori, the observed RGB could all be due to surfaces which are arbitrarily dark as a result of being obliquely illuminated. At the other extreme (the apex of the other pyramid) there is a vertex corresponding to the whitest reflectance. To include specularities we take that vertex, and move it away from the origin, along the line connecting to the origin. Thus the hull facets adjacent to the origin remain the same, but the ones adjacent to the RGB of white are stretched away from the origin. In other words, we add a single reflectance to our world which is a multiple of a uniform reflectance. The multiple should be large enough to accommodate a bright specularity taking the dynamic range of the vision system into account, but the exact specification of the value is not very important. (We have experimented with factors of 2, 4, and 8). The concept is illustrated using two dimensions in Figure 2.

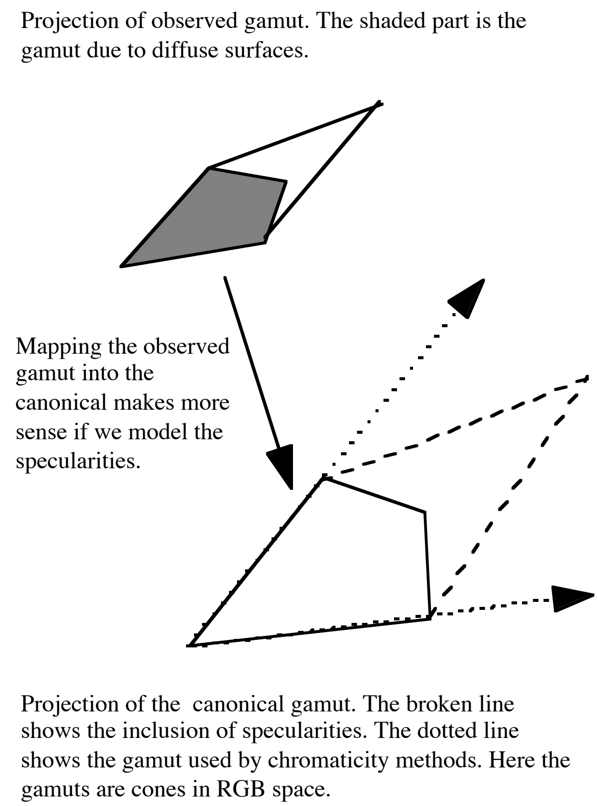

Figure 2. Illustration of gamut extension used for specularities. The gamuts are actually polytopes in 3 dimensional RGB space.

While very simple, the method naturally models real specularities which are always a combination of the specular reflection and the underlying body reflection. Both the specular reflection and the body reflection are part of the convex hull, and thus any convex combination of them is also in the hull. Finally, to include the specular reflection of colored metals (brass, copper, gold), we add multiples of the reflections for these substances into the canonical gamut. The color of specularities is still quite restricted, being somewhere between white and the color of copper, but the existence of metallic specularities will now work with, instead of against, the information provided by the other colors.

The new canonical gamut is then used as part of standard RGB based gamut-mapping algorithms. As in [3] and [21] we used Finlayson's illumination constraint. We investigate the two methods for choosing a solution from the constraint set that were mentioned above, those being the original maximum volume method and the average over the feasible set. When the illumination constraint is used, this set is not precisely convex, and we numerically integrate to obtain the average.

This method works well even if there are no specularities. The work of Finlayson and Hordley [4] suggests that the most important facets in the non-specular case are the ones adjacent to the origin; specifically the ones not modified by our method. The arguments in that work also imply that our method should be at least as strong as any chromaticity-based gamut-mapping algorithm, regardless of the presence of specularities. Of course, when specularities are present, our algorithm should excel. Finally, when there are strong un-clipped specularities, our algorithm estimates the overall illuminant brightness better than all other algorithms.

Experiments and Results

We have tested the above methods both on synthetic image data, and on real image data. For the former, we generated data without specularities, with non-metallic specularities, and with a mixture of metallic and non-metallic specularities. To model the metallic specularities we measured the specular reflectance of a number of metallic objects using a Photoresearch PR-650 spectraradiometer. The metallic samples included several brass and copper surfaces, as well as gray metallic surfaces such as aluminum and stainless steel. We modeled non-specular reflectance using a database of roughly 2000 reflectance spectra obtained from a number of sources including some of our own measurements. For each simulated "world" we ran all the algorithms on 200 randomly selected groups of 4 surfaces under randomly selected illuminants. Color constancy on such a small number of surfaces is difficult on average, and thus doing well requires specularities. For each set of generated data we also simulated pixel clipping.

We provide the results of the algorithms using two different error metrics. The first measures the ability of the algorithms to estimate the chromaticity of the illuminant. Here we consider the illuminant RGB and the corresponding estimate thereof as vectors in RGB space, and compute the angle between these two vectors in degrees. The second error metric is simply the Euclidean distance between these two vectors.

In general, the results are very encouraging. We see that the original CRULE algorithm (with Finlayson's illumination constraint) and the Retinex method work well when there are good specularities, but that these algorithms are more sensitive than the others to clipping, and give poor illuminant brightness results (Table 3). On the other hand, the new algorithms, as exemplified by "SP-ECRULE-AVE", do not have these problems. We also found, not surprisingly, that if metallic specularities are present, then modeling metallic specularities yields better results. On the negative side, the error in the illuminant chromaticity estimates increases somewhat when such surfaces are absent. Further work is needed to characterize this tradeoff.

We also provide some results on real images from two data sets (Table 4). The first data set consisted of scenes with and without significant specularities, but with few metallic specularities. In this data set there are 33 scenes taken under 11 different illuminants. Several images were culled due to problems, leaving 321 test images. For the second set we used seven scenes with metallic specularities under the same 11 illuminants. Again, some images were culled, leaving 71. The images were taken at low enough light to minimize clipping due to specularities, and the dynamic range was extended by averaging multiple frames. This allows us to investigate strong specularities, and the possibilities afforded by higher dynamic range cameras.

The image data results generally confirm the results found with synthetic data. Overall the gamut mapping algorithms do well compared to the other algorithms on real data. However, part of their success may be due to the use of images with extended dynamic range, and therefore future work will look at the effect of artificially clipping such data.

Conclusions

We have considered computational color constancy in the context of scenes with both specular and non-specular surfaces. We have also proposed an algorithm which is explicitly designed to make use of both types of information. This is in contrast to most current color constancy algorithms which generally focus on using only the matte surfaces or only the specularities. Unlike other algorithms using specular information, our method does not need to identify groups of pixels as corresponding to the same surface under different geometry. Instead, the method implicitly uses the information inherent in the brightness of the image pixels. However, since the method extends the already capable gamut mapping approach, the method can give good results even when specularities are not present. In addition, our method is easily extended to deal with specular reflection from colored metals, and is the first color constancy algorithm to do so. The new algorithms yield good results on both synthetic and real image data.

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.

Table 1: Key to algorithms

ECRULE	CRULE with illumination constraint
MV	Solution maximizes volume of mapped set
AVE	Solution is the average of feasible set
Retinex	Estimate illuminant color by the maximum of each channel.
Gray World	Estimate illuminant color by image average

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

	No specularities	Simulated specularities	Simulated specularities and simulated clipping	Simulated metallic and non-metallic specularities	Simulated metallic and non-metallic specularities and clipping
ECRULE-MV	8.4	3.63	8.93	7.44	9.29
ECRULE-AVE	8.0	6.55	8.43	7.00	8.57
SP-ECRULE-MV	8.6	4.22	9.37	7.69	9.56
SP-ECRULE-AVE	7.6	3.84	8.03	6.46	8.59
MET-ECRULE-MV	9.4	5.88	10.49	7.41	10.17
MET-ECRULE-AVE	8.7	7.01	9.55	6.25	9.52
Retinex	13.2	5.12	12.88	10.30	12.80
Gray World	11.9	6.66	13.38	9.95	12.41
Color by correlation	8.0	5.12	9.24	6.81	9.08
Neural Net	7.3	6.12	8.13	6.59	8.52

Table 3: Average RMS error of illuminant RGB estimate (generated data)

	No specularities	Simulated specularities	Simulated specularities and simulated clipping	Simulated metallic and non-metallic specularities	Simulated metallic and non-metallic specularities and clipping
ECRULE-MV	147	1103	159	4574	163
ECRULE-AVE	125	2020	149	7653	152
SP-ECRULE-MV	206	166	241	1944	237
SP-ECRULE-AVE	138	236	169	3327	176
MET-ECRULE-MV	227	204	261	478	257
MET-ECRULE-AVE	161	171	205	905	204
Retinex	201	884	198	3487	204
Gray World	141	1065	173	3046	162

Table 4: Average angular error in RGB space of illuminant estimate (image data).

	Images without metallic surfaces	Images with metallic surfaces
ECRULE-MV	5.3	12.0
ECRULE-AVE	6.6	11.2
SP-ECRULE-MV	6.3	10.9
SP-ECRULE-AVE	6.4	10.9
MET-ECRULE-MV	7.2	10.4
MET-ECRULE-AVE	7.5	10.6
Retinex	7.8	13.9
Gray World	12.0	16.4
Color by correlation	11.0	13.6
Neural Net	9.8	12.2

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. J. Klinker, S. A. Shafer, and T. Kanade, A physical approach to color image understanding, *International Journal of Computer Vision*, **4**, pp. 7-38 (1990).
6. B. Maxwell, Segmentation and interpretation using multiple physical hypotheses of image formation, CMU, (1996).
7. B. A. Maxwell and S. A. Shafer, Physics-based segmentation: moving beyond color, *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 742-749 (1996).
8. B. A. Maxwell and S. A. Shafer, Physics-based segmentation of complex objects using multiple hypotheses of image formation, *Computer Vision and Image Understanding*, **65**, pp. 269-295 (1997).
9. G. Healey, Using color for geometry-insensitive segmentation, *Journal of the Optical Society of America A*, **6**, pp. 920-937 (1989).
10. S. A. Shafer, Using color to separate reflection components, *COLOR Research and Application*, **10**, pp. 210-218 (1985).
11. S. K. Nayar, K. Ikeuchi, and T. Kanade, Surface reflection: physical and geometric perspectives, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **13**, pp. 611-634 (1991).
12. M. D'Zmura and P. Lennie, Mechanisms of color constancy, *Journal of the Optical Society of America A*, **3**, pp. 1662-1672 (1986).
13. H.-C. Lee, Method for computing the scene-illuminant chromaticity from specular highlights, *Journal of the Optical Society of America A*, **3** pp. 1694-1699 (1986).
14. S. Tominaga and B. A. Wandell, Standard surface-reflectance model and illuminant estimation, *Journal of the Optical Society of America A*, **6**, pp. 576-584 (1989).
15. S. Tominaga, Realization of color constancy using the dichromatic reflection model, *Proc. IS&T/SID 2nd. Color Imaging Conference: Color Science, Systems and Applications*, pp. 37-40 (1994).
16. W. M. Richard, Automated detection of effective scene illuminant chromaticity from specular highlights in digital images, : Center for Imaging Science, Rochester Institute of Technology (1995).
17. 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).
18. 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).
19. 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).
20. V. Cardei, B. Funt, and K. Barnard, Adaptive Illuminant Estimation Using Neural Networks, *Proc. International Conference on Artificial Neural Networks* (1998).
21. B. Funt, K. Barnard, and L. Martin, Is Colour Constancy Good Enough?, *Proc. 5th European Conference on Computer Vision*, pp. I:445-459 (1998).
22. U. Seger, G. H.-G. Graf, and M. E. Landraf, Vision assistance in scenes with extreme contrast, *IEEE Micro*, **13**, pp. 50-56 (1993).
23. G. D. Finlayson, Coefficient Color Constancy, Simon Fraser University, School of Computing (1995).