

Colour Constancy for Scenes with Varying Illumination

Kobus Barnard¹, Graham Finlayson², and Brian Funt¹

¹School of Computing Science, Simon Fraser University, Burnaby BC, Canada,
V5A 1S6

kobus@cs.sfu.ca, funt@cs.sfu.ca

² Department of Computer Science, University of York, Heslington, York,
YO1 5DD, United Kingdom

graham@minster.cs.york.ac.uk

Abstract. We present an algorithm which uses information from both surface reflectance and illumination variation to solve for colour constancy. Most colour constancy algorithms assume that the illumination across a scene is constant, but this is very often not valid for real images. The method presented in this work identifies and removes the illumination variation, and in addition uses the variation to constrain the solution. The constraint is applied conjunctively to constraints found from surface reflectances. Thus the algorithm can provide good colour constancy when there is sufficient variation in surface reflectances, or sufficient illumination variation, or a combination of both. We present the results of running the algorithm on several real scenes, and the results are very encouraging.

1 Introduction

Many colour constancy algorithms have been developed, but all are subject to quite restrictive assumptions and few have been tested on real images. Of the existing algorithms we believe that the one by Finlayson [8] is currently the most general and most thoroughly tested. Nonetheless, it is restricted to scenes in which the illumination is constant or at least locally constant. This assumption is more often violated than one might at first suspect given that the incident illumination from any fixed direction does generally vary slowly as function of position. The problem is that the surface orientation even of smooth surfaces can vary quite rapidly with position so that light at nearby surface locations may be received from very different regions of the illumination field. Since these different regions of the illumination field often possess substantially different spectral power distributions—such as is the case in a room in which there is light from a light bulb mixed with daylight from a window—nearby points on the surface in fact can receive very different incident illumination.

This paper addresses the problem of colour constancy in scenes where the spectral power distribution of the incident illumination is allowed to vary with scene location. Finlayson et. al. [7], D’Zmura et. al. [18], and Tsukada et. al. [16] have shown that a difference in illumination, *once it has been identified*, provides additional constraints that can be exploited to obtain colour constancy, but they do not provide an automatic method of determining when such a difference exists. We

present a new algorithm that first uncovers the illumination variation in an image and then uses the additional constraint it provides to obtain better colour constancy. The algorithm presupposes a diagonal model for illumination changes, and that the illumination varies spatially slowly. However it is quite robust to moderate violations of these assumptions.

The colour constancy problem is the retrieval of an illumination-independent description of a scene's surface colours. This is essentially equivalent to modeling the illumination incident on the scene, since if the illumination is known the surface colours can be calculated. Following Forsyth [9] we interpret colour constancy as taking images of scenes under unknown illumination and determining the camera response to the same scene under a known, canonical light. In a general context this problem has proven difficult to solve, so to make progress, restrictive assumptions are made. In particular, it is common to assume that the scene is flat [9, 13, 14], that the illumination is constant throughout [2, 3, 9, 10, 15], and that all reflectances are matte. Finlayson [8] has shown that if we focus on solving only for surface chromaticity and forego estimating surface lightness then the restriction to flat matte surfaces can be relaxed. However, the assumption that the chromaticity of the illumination does not change is more troublesome.

The Retinex algorithm [1, 11, 13, 14] partially addresses the issue of varying illumination. At least in principle—it does not in fact work in practice—Retinex eliminates the variation in illumination and computes surface lightnesses for each of the three colour channels independently. Since eliminating the illumination and recovering the illumination are equivalent problems [4], if Retinex worked it could be used to recover the incident illumination. Retinex operates on the principle that within a single colour channel small changes in image intensity arise from changes in illumination while large changes indicate changes in surface colour. The small changes are thresholded away and the big changes are preserved so that the surface lightness can be reconstructed, essentially by integration. Unfortunately any error in classifying the intensity changes can lead to serious errors in the recovered result. In essence the Retinex algorithm uses a primitive, gradient-based-edge-detection strategy to identify the reflectance edges, so given the long history of edge-detection research, it should not be surprising that it does not perform well.

To overcome the weaknesses of Retinex's edge detection method, we incorporate knowledge about the set of plausible illuminants and from this set derive information about the kinds of chromaticity change that a change in illumination can induce within a region of uniform reflectance. This constraint is more global than local edge detection and using both types of constraint together yields good results. Once the illumination variation is uncovered it is combined with the other constraints arising from the set of colours found in the image as will be discussed below.

2 The Colour Constancy Algorithm

Our colour constancy algorithm has two main components: one to extract the illumination field and another to combine the constraints provided by the a priori knowledge of the surface and illumination gamuts with those obtained from the observed surfaces and the extracted illumination field. The constraint part of the algorithm will be described first.

2.1 Surface and Illumination Constraints

In order to represent the constraints efficiently, we make the approximation that the effect of the illumination can be modeled by a diagonal matrix [5, 6]. Specifically, if $[r, g, b]$ is the camera response of a surface under one illumination, then $[r, g, b]\mathbf{D} = [r\mathbf{D}_{11}, g\mathbf{D}_{22}, b\mathbf{D}_{33}]$, where \mathbf{D} is a diagonal matrix, is the camera response to the same surface under a second illumination. In other words, each camera channel is scaled independently. The accuracy of the diagonal approximation depends on the camera sensors, which for the camera used in the experiments is within 10% (magnitude of $[r, g, b]$ difference) of the general linear model. For sensors for which the diagonal model is too inaccurate, it is usually possible to improve it by spectrally sharpening the sensors [5].

Following Finlayson [8], we work in the chromaticity space $[r/b, g/b]$. This space preserves the diagonal model in the sense that if illumination was exactly modeled by a diagonal transform applied to $[r, g, b]$, then it will also be exactly modeled by a diagonal transform (now 2D) applied to $[r/b, g/b]$. If either the illumination is spatially constant or pre-processing has removed all illumination variation, then transforming the input image to what it would have looked like under the canonical illuminant requires simply transforming it by a single diagonal matrix. The goal of the colour constancy algorithm is to calculate this matrix.

The algorithm's basic approach is to constrain the set of possible diagonal maps by adding more and more information so that only a small set of possible maps remains. The first constraints are the those due to Forsyth [9]. He observed that the set of camera responses that could be obtained from all combinations of a large set of surfaces viewed under a fixed illuminant is a convex set which does not fill all of the $[r, g, b]$ colour space. This set is referred to as that illuminant's gamut, and in the case of the canonical illuminant is called the canonical gamut. For a typical scene under unknown illumination, the camera responses will lie in a subset of the unknown illuminant's full gamut. Since all possible surfaces are assumed to be represented within the canonical gamut, whatever the unknown illuminant is, it is constrained by the fact that it is a diagonal map projecting the scene's observed response set into the canonical gamut. There will be many possible diagonal maps satisfying this constraint because the scene's set is a subset of the full gamut and so it can 'fit' inside the larger gamut many different ways. Forsyth shows that the resulting constraint set of diagonal maps is convex. As shown in [8], all the required relationships hold in the $[r/b, g/b]$ chromaticity space.

The second source of constraint arises from considering the set of common illuminants as has been formulated by Finlayson [8]. After applying Forsyth's surface constraints, the resulting set of diagonal maps typically includes many that correspond to quite atypical illuminants. The illumination constraint excludes all the illuminants that are not contained in the set of typical illuminants. Finlayson restricted the illumination to the convex hull of the chromaticities of the 6 daylight phases provided by Judd et al [12], the CIE standard illuminants A, B, C [17], a 2000K Planckian radiator, and uniform white. We have improved upon this sampling of illuminants by using 100 measurements of illumination around the university campus, including both indoor and outdoor illumination. Some inter-reflected light was included such as that from concrete buildings and light filtering through trees, but illumination that was obviously unusual was excluded. The set of chromaticities of the measured illuminants is larger in area than the original set, but it does not

contain it entirely, as the 2000K Planckian radiator is more red than what is common.

It should be noted that the set of typical illuminants provides a constraint on mappings from the canonical to the unknown, which is the reverse of that for surfaces discussed above in which the restriction was on mappings from the unknown illuminant to the canonical illuminant. To make use of the constraint it must be inverted which means that the restriction on the set of illuminants becomes a non-convex set in the mapping space used for surface constraints. This potentially presents a problem since the sets must be intersected in order to combine constraints and in three-dimensions it is much faster to compute intersections of convex sets than non-convex ones. While in the two-dimensional case the set intersections can be directly computed, in practice the inverse of the measured illumination non-convex gamut was found to be close enough to its convex hull that for convenience the hull could be used anyway.

Varying illumination provides the third source of constraint. Our use of it here generalizes the algorithm presented in [7]. In that work the map taking the chromaticity of a single surface colour under an unknown illuminant to its chromaticity under the canonical illuminant is constrained to lie on a line. Effectively this amounts to assuming all the candidate illuminants are approximately Planckian radiators and their chromaticities lie roughly on the Planckian locus. The chromaticity of the same surface viewed under a second illuminant defines a second line. If the difference in the illuminations' chromaticities is non-trivial, the two lines will intersect, thereby constraining the surface's chromaticity to a unique value.

We extend the idea of using the variation in illumination in two ways. First we use the entire illumination gamut instead of simply the Planckian radiators. Second we exploit the illumination variation across the entire image, as opposed to just that between two points on one surface patch. Thus the illumination over the entire image is both used, and solved for. The details follow.

For the moment assume that we already have the relative illumination field for the image. The relative illumination field for each pixel P is defined by the diagonal transform required to map the illumination at some chosen base pixel B to the illumination at P. The relative illumination field describes all the pixels only with respect to one another, so given it, the remaining problem is to solve for the illumination at B and hence establish the illumination everywhere in absolute terms.

The approach is motivated by the following argument. Suppose that the left side of the image is illuminated by a blue light. This means that the relative illumination field at a pixel on the left side transforms illuminants so that they are more blue. However, the illumination at the center of the image cannot be so blue that making it even more blue produces an illumination that falls outside the set of possible illuminants. Thus the illumination at the center is constrained by the jump towards blue. All entries in the field contribute this sort of constraint. This will now be made more formal.

First we verify the intuitive claim that the constraint provided by one of the values \mathbf{D} in the relative illumination field is the set of possible illuminants scaled by \mathbf{D}^{-1} . Consider the illumination gamut, I which is a convex set:

$$I = \left\{ \mathbf{X} \mid \mathbf{X} = \sum_i \lambda_i \mathbf{X}_i \text{ where } \sum_i \lambda_i = 1 \right\} \text{ for hull points } \{ \mathbf{X}_i \} \quad (1)$$

We have the constraint that we can map the illumination by the diagonal map \mathbf{D} and still be in this set:

$$\mathbf{XD} \in I \quad (2)$$

This means that:

$$\mathbf{XD} = \sum_i \lambda_i \mathbf{X}_i \quad \text{for some } \lambda_i \text{ with } \sum_i \lambda_i = 1, \lambda_i \geq 0 \quad (3)$$

And

$$\mathbf{X} = \sum_i \lambda_i (\mathbf{X}_i \mathbf{D}^{-1}) \quad \text{for some } \lambda_i \text{ with } \sum_i \lambda_i = 1, \lambda_i \geq 0 \quad (4)$$

So we define a new constraint set V as:

$$V = \left\{ \mathbf{X} \mid \mathbf{X} = \sum_i \lambda_i (\mathbf{X}_i \mathbf{D}^{-1}) \text{ where } \sum_i \lambda_i = 1, \lambda_i \geq 0 \right\} \quad (5)$$

It is clear that for all $\mathbf{X} \in V$, $\mathbf{XD} \in I$. Furthermore, the argument is reversible. That is, if $\mathbf{Y} = \mathbf{XD} \in I$, $\mathbf{X} \in V$. It should be noted that the above also shows that we can identify the convex constraint set with the mapped hull points $\mathbf{X}_i \mathbf{D}^{-1}$.

Next we note that the convex hull of these constraints is just as powerful as the entire set. The motivation for using the hull is that it saves a significant amount of processing time. We are free to use the hull regardless, but it is comforting to know that doing so does not weaken the algorithm.

Despite the details, the additional constraint is very simple in that it says that we have to be able to scale the illuminant by a certain amount and *still* satisfy the illumination constraint. This constraint is realized by simply scaling the set of illuminants by the inverse. As a simple example, consider the one-dimensional line segment $[0,1]$. If we have a condition on these points that when they are scaled by a factor of two the result must still be in that segment, then the set of points in our constrained set must be $[0, \frac{1}{2}]$. In other words, the set was scaled by the inverse of the scale factor.

2.2 Combining the Constraints

Given the above formulation of the various constraints they can be easily combined into a forceful colour constancy algorithm. First the relative illumination field is used to remove the illumination variation from the image leaving an image which is of the scene with chromaticities as they would have appear if it had been illuminated throughout by the illumination at the base point. Starting from this intermediate result a constraint on the possible illumination maps is derived for each of the surface chromaticities. The illumination constraint provided by the set of plausible illuminants is fixed by the initial measurement of the various illuminants around the campus. Each hull point of the set of the relative illumination field furnishes yet another constraint; namely, the illumination constraint multiplied by the appropriate diagonal transform. The illumination constraint and the transforms due to the relative illumination field are intersected, and the result is inverted. As mentioned above, this inverted set was approximated well by its convex hull. The inverted set is then intersected with the intersection of all the surface constraints.

The final step of the algorithm is to choose a solution from the set of possible solutions. In [8, 9] the solution chosen maximizes the volume of the mapped set, which is equivalent to maximizing the product of the components of the mapping. In this work, however, we use the centroid of the solution set, which is more natural. This choice can be shown to minimize the expected error if all solutions are equally likely and error is measured by the distance from the choice to the actual solution. Furthermore, in both synthesized and real images, the centroid was found to give better results.

The colour constancy algorithm that incorporates all the different constraints was tested first on generated data. One thousand sets containing 1, 2, 4, 8, and 16 surfaces were randomly generated and used in conjunction with 1 of 5 illuminants, with 0 through 4 of the remaining lights playing the role of additional illuminants arising as a result of varying illumination. Table 1 gives the results which are exactly as wished. As either the number of surfaces or the number of extra lights increases, the answer consistently improves. Thus it was verified that varying illumination is a powerful constraint, and furthermore, it can be effectively integrated with the other constraints.

3 Finding the Relative Illumination Field

We now detail an algorithm for finding the relative illumination field describing the variation in the incident illumination. This algorithm can be divided into two parts. The first is a new technique for image segmentation appropriate for scenes with varying illumination. The second part uses the segmentation to determine the illumination map robustly.

Unless the illumination is known to be constant, it is essential that a segmentation method be able to accommodate varying illumination. In general, the segmentation problem is quite difficult, especially with varying illumination, as in this case an area of all one reflectance can exhibit a wide range of colour. Fortunately for our purposes, it is not critical if an occasional region is mistakenly divided into two pieces, nor if two regions which have almost the same colour are incorrectly merged. This is because the goal at this point is simply to identify the illumination, not the surfaces. Nonetheless, the better the segmentation, the more reliable and accurate the possible colour constancy.

One approach to segmentation is that used by Retinex theory [13, 14]. In Retinex small changes in pixel values at neighboring locations are assumed to be due to changes in the illumination and large changes to changes in surface reflectance. This idea can be used to segment an image into regions of constant surface reflectance properties by growing regions by including pixels only if they are less than some small threshold different from their neighbours. The threshold must be large enough to allow for both noise and the illumination changes and yet not admit small changes in surface reflectance—a balance which is of course impossible to establish.

We use this method as part of our algorithm, but alone, it is not sufficient. The problem is that two dissimilar regions will eventually mistakenly merge if there exists a sequence of small jumps connecting them. This can occur if the edge is gradual or because of noise. In essence, a threshold large enough to allow for noise (and varying illumination) allows for enough drift in the pixel values to include an entirely dissimilar region. Local information alone is insufficient, so we resolve the problem by adding a more global condition involving illumination constraints.

Number of Surfaces

	1	2	4	8	16
BF	0.073	0.073	0.073	0.073	0.073
BDT	0.116	0.116	0.116	0.116	0.116
GW	1.62	1.01	0.69	0.513	0.428
RET	1.62	1.10	0.72	0.478	0.354
S	12.4	4.4	1.65	0.585	0.285
SI	2.275	1.68	0.99	0.480	0.271
SIV1	1.65	1.26	0.79	0.420	0.256
SIV2	1.154	0.896	0.620	0.351	0.242
SIV3	0.800	0.656	0.488	0.311	0.231
SIV4	0.384	0.36	0.317	0.274	0.228

Solution Method

Solution Method Key

BF	Error of best possible solution using full linear map
BDT	Error of best possible solution using a diagonal map
GW	Naive Grey World Algorithm (scale each channel by average)
RET	Naive Retinex Algorithm (scale each channel by maximum)
S	Surface constraints alone
SI	Surface and illumination constraints
SIV	Surface and illumination constraints with view under one extra illuminant
SIV2	Surface and illumination constraints with view under 2 extra illuminants
SIV3	Surface and illumination constraints with view under 3 extra illuminants
SIV4	Surface and illumination constraints with view under 4 extra illuminants

Table 1. Results of color constancy experiments for 1000 sets of 1, 2, 4, 8, and 16 surfaces under all combinations of test lights and extra lights for varying illumination. The canonical illuminant was a Philips CW fluorescent light. The values shown are the average magnitude of the chromaticity vector difference between the estimate and the desired answer, averaged over all results.

The new global segmentation condition is based on the idea that in order for two pixels—no matter how far apart they are spatially—to be considered part of the same region, a plausible illumination change between them must exist. The set of plausible illuminant changes can be derived in advance from the initial set plausible illuminants. This condition binding pixels within a single region based on plausible illumination changes is called patch coherence. The patch coherence condition differs from the Retinex condition in two important ways. First, the cumulative drift in pixel values along a path is limited, as opposed to growing linearly with the pixel distance. Second, the allowable drift is constrained more in certain directions due to the nature of the set of common illuminants. For example, green illuminants are rare, which means that the set of common illuminants is narrow in the green direction, and thus overall, the drift towards or away from green is more restricted than that towards or away from blue.

It was found to be useful to retain the Retinex condition as well as the patch coherence method described above for two reasons. First, the Retinex condition is

faster to compute, and thus can be used to reject pixels that do not need to be tested further for inclusion. Second, if a comprehensive set of possible illuminants is used, then an occasional surface boundary change will also be a possible illumination change. Since the Retinex method by itself works much of the time, these exceptional cases in which a surface change mimics an illumination change generally will be covered by the Retinex condition.

In detail our segmentation algorithm begins with an arbitrary initial starting point in a region and assumes that the illumination at that point is constrained to be in the set of plausible illuminants. It is important to update the constraints on the illumination at the starting point each time a new point is added to the region. Each newly included point further constrains the possible illuminations at the starting point. Updating the constraints is similar to using the relative illumination field to solve for colour constancy as described above. The element-wise ratio of the chromaticities of the new point to that of the initial point induces a constraint set V defined by (5). Specifically, the illumination gamut is transformed by the inverse of the ratio interpreted as a diagonal transform. This convex set is intersected with the current constraint set. If the intersection is null, then the new point is excluded and the constraint set is left unchanged. If it is not null, then the intersection becomes the updated constraint set and the new point is added to the region.

Similar to the situation when solving for colour constancy, it is sufficient to perform the intersection only when the new transform to be applied to the illumination gamut is outside the convex hull of the preceding transforms. Although calculations are relative to the initial point, this procedure ensures that all points in the region can be assigned illuminants from the set of plausible illuminants which are consistent with the illumination jumps between them. Furthermore, the inclusion of any of the rejected points would violate this condition.

Given our segmentation we reduce the problem of finding the relative illumination field to that of finding the illumination at the center of each region relative to that at the center of the base region. Since the center of a region, as defined by the center of mass, need not be inside the region, the implementation uses a point in the region close to the center of mass, preferably a few pixels from the boundary. The illumination at a point relative to that of the region center is simply the ratio of its response to the response of the center point. This follows directly from the assumption that the pixels are from the same surface, given that we accept a diagonal model for illumination change. Thus the map at an arbitrary point is simply the map at the center, adjusted by this relative jump.

To determine the maps at the center points we make the assumption that illumination does not change significantly at the region boundaries. Thus every jump across a boundary gives a condition on the relative maps of the centers of the two adjacent regions. More specifically, consider two regions A and B, with centers C_A and C_B , and boundary points B_A and B_B close to each other. Denote responses by R subscripted by the point label and denote the diagonal map relative to the grand central point as D , also subscripted by the point label. Each channel or chromaticity component is dealt with independently, so the quantities in the equations are scalars. The assumption that the illumination does not change significantly at the boundary is simply:

$$D_{B_A} = D_{B_B} \quad (6)$$

Since we are assuming a diagonal model of illumination change, and C_A is on the same surface as B_A , and similarly for the center and boundary of surface B, we have:

$$D_{B_A} = D_{C_A} \left(\frac{R_{B_A}}{R_{C_A}} \right) \quad \text{and} \quad D_{B_B} = D_{C_B} \left(\frac{R_{B_B}}{R_{C_B}} \right) \quad (7)$$

Combining (10) and (11) yields:

$$D_{C_A} \left(\frac{R_{B_A}}{R_{C_A}} \right) = D_{C_B} \left(\frac{R_{B_B}}{R_{C_B}} \right) \quad (8)$$

Taking logarithms of both sides of (12), and rearranging terms gives:

$$\ln(D_{C_A}) - \ln(D_{C_B}) = \ln(R_{B_B}) - \ln(R_{B_A}) + \ln(R_{C_A}) - \ln(R_{C_B}) \quad (9)$$

This final equation is at the heart of the method. Here we have a condition on the component of the map for two of the regions. Other boundary point pairs produce additional equations. In order to have a robust method, one would like long boundaries to have more weight in the process than short ones, since the latter may be due to a small region consisting entirely of noise. But this is exactly what we will get if we enter one equation for each boundary pair and solve the resulting system of equations in the least squares sense. Furthermore, some boundary pairs can be identified as being more reliable and these are weighted even more by scaling the equation by a number greater than one (typically five). In addition, some boundary pairs should contribute less, and their equations are scaled by a number less than unity.

In order to have a solution to the set of equations, it must be insured that all segments connect to each other through the boundary pairs. This might be accomplished simply by assigning a region to every point, and using each break in either the horizontal or vertical directions to produce a boundary pair. This is usually not an option because often some parts of the image should not be used; for example, when an area is too dark. Therefore the likelihood of connectedness between regions was increased in the following manner. Boundary pairs were assigned at each horizontal and vertical change of region. If one of the regions was to be ignored, a good region was sought in the same direction, taking as many pixels as required. The resulting equation was weighted inversely to the distance taken to find a good region. Thus such a boundary would contribute little to the solution, but connectivity was not a problem for reasonable images (it is still possible to construct an image which will lack connectivity).

Several additional steps were taken to improve robustness. First very small regions were excluded from the computations. Second, it was found to be better to use pixels one unit towards the insides of the respective regions, if these were available. This way the pixels would tend to have contributions that were solely due to a single surface, as opposed to the possibility that they straddled more than one surface. These boundary pairs were weighted by a factor of five compared to ones where it was necessary to use pixels exactly on the boundary.

The final step in determining the relative illumination field is to interpolate over any excluded areas.

4 Results

The algorithm has been tested on a set of images of real scenes. In all cases the 'unknown' illumination consists of light from an incandescent bulb coming from

one direction mixed with light from a Philip’s CW fluorescent tube covered by a pale blue filter coming from another. The latter is similar in colour temperature to sky light. Thus the scenes mimic a common real world situation—an office with a window.

Unfortunately, qualitative evaluation of the results requires access to colour reproductions which are not available here. However, the grey scale counterparts reproduced in Figure 1 should give the reader some idea of the nature of the input used. The first image shown is a three-dimensional “Mondrian” made by affixing coloured construction paper to a conical waste paper bin. The bin is lying on its side with the incandescent light shining from above and the blue light from below. The top has blue and green papers on it, in the middle is a grey patch, and near the bottom are red and yellow papers. This illumination causes a distinct reddish tinge at the top and a clear bluish tinge at the bottom.

The second image is a simple two-dimensional “Mondrian” made by attaching eight sheets of coloured construction paper to the lab wall such that substantial parts of the wall remained visible. The third is a multi-coloured cloth ball. The cloth ball is interesting because the cloth has more texture than the construction paper. The fourth image shown in Figure 1 is of a person in front of a grey background with the sky-like light on the left and the incandescent light on the right. Under this illumination the left side of the grey background appears quite blue, and the flesh tones on the left are noticeably incorrect. It is not possible to obtain a canonical image for comparison because people move too much in the time it takes to set up the new illumination conditions. However, the qualitative results are quite promising. An additional image used for quantitative results is of a single piece of green poster board (not shown).

The numerical results in Table 2 reflect the RMS difference (over the entire image) between the $[r/b, g/b]$ chromaticities at corresponding pixels in the recovered and canonical images. There are few colour constancy algorithms designed to deal with scenes of the generality addressed here so it is difficult to make comparisons with existing algorithms without violating their assumptions. As a first measure, we compare the solution obtained by a straight least squares fit of the input image to the canonical using a full linear model and then with a diagonal model. Without accounting for the illumination variation no algorithm working on the 3D Mondrian image can do better than the 0.78 error of the full linear case. In contrast, by

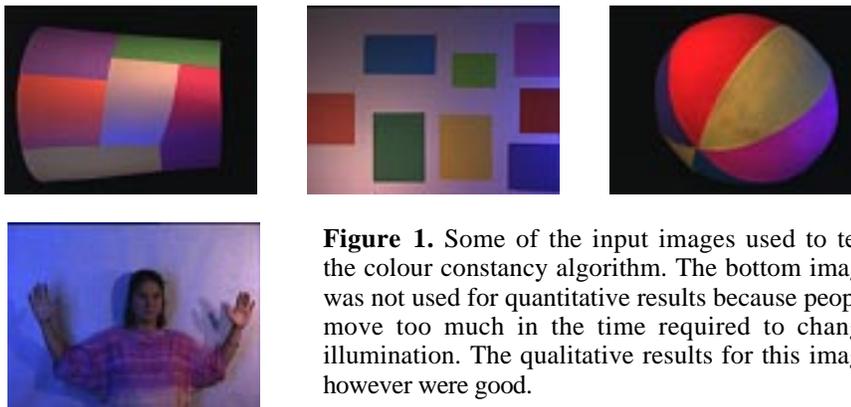


Figure 1. Some of the input images used to test the colour constancy algorithm. The bottom image was not used for quantitative results because people move too much in the time required to change illumination. The qualitative results for this image however were good.

Object

	3D Mondrian	Ball	Green Card	Paper on Wall
BT	0.78	0.536	0.195	0.265
BDT	0.80	0.603	0.240	0.273
N	0.96	0.93	0.279	0.405
GW	0.88	1.37	0.312	0.349
RET	0.87	0.75	0.246	0.290
S	0.86	0.92	1.030	0.476
SI	0.87	0.73	0.540	0.472
VIR-BT	0.148	0.224	0.042	0.093
VIR-BDT	0.176	0.227	0.044	0.133
VIR-N	0.44	0.595	0.262	0.530
VIR-GW	0.68	1.36	0.275	0.306
VIR-RET	0.61	0.95	0.275	0.228
VIR-S	0.33	1.11	1.929	0.211
VIR-SI	0.39	0.38	0.269	0.163
VIR-SIV	0.26	0.45	0.073	0.151

Error between the result and the canonical using the various solution methods

BT	Best possible linear map solution
BDT	Best possible diagonal map
N	No processing. Simply use the input image as the result.
GW	Naive Grey World Algorithm (scale each channel by average)
RET	Naive Retinex Algorithm (scale each channel by maximum)
S	Using Surface Constraints
SI	Using Surface and illumination constraints
VIR-BT	Best linear map applied after varying illumination is removed
VIR-BDT	Best diagonal map applied after varying illumination is removed
VIR-N	No processing applied after varying illumination is removed
VIR-GW	Naive Grey World Algorithm with varying illumination removed.
VIR-RET	Naive Retinex Algorithm with varying illumination removed.
VIR-S	Using surface constraints with varying illumination removed.
VIR-SI	Using surface and illumination constraints with varying illumination removed.
VIR-SIV	Complete new algorithm using surface, illumination, and varying illumination constraints with varying illumination removed.

Table 2. Results of color constancy algorithms applied to four images. The canonical illuminant was a Philips CW fluorescent light. The values shown are the RMS (over all pixels) magnitude of the chromaticity vector difference between the estimate and the desired answer, which is a view of the scene under the canonical light.

accounting for and utilizing the illumination variation our new algorithm reduces the error to 0.26.

Table 2 also shows the chromaticity error for the case of doing nothing at all. The grey world algorithm, which uses the average image chromaticity as an illumination estimate, and the Retinex normalization strategy of taking the maximum response from each colour band as an illumination estimate are tried even though the comparison is somewhat unfair because they both assume the illumination to be constant. Similar tests are run using surface constraints alone and surface constraints with the additional constraints on the set of plausible illuminants.

To make the comparison fairer, we also include similar tests with these algorithms but applied after the illumination variation has been discounted. In other words, we combined the first part of our algorithm (removal of the illumination variation) with each of the other algorithms. In this case, the other algorithms are applied to data which does not violate the constant illumination assumption, but they still do not exploit the information contained in the illumination variation. The Retinex normalization applied in this way gives an algorithm which is close, in theory, to the original Retinex idea.

The results show first that if the varying illumination is not accounted for, then all the colour constancy algorithms perform poorly. In all cases, the complete new VIR-SIV algorithm did better than any algorithm which assumed the chromaticity of the illumination to be constant. In fact, the performance is better than that of the best diagonal and best linear fits. The complete algorithm also performed better than the others applied to the data with the varying illumination removed, except when compared to applying the combination of surface and illumination constraints to the ball image. Most importantly, the algorithm performs better than the Retinex scaling applied to the data with the variation removed. As mentioned above, this procedure is close to the spirit of the original Retinex algorithm, which is unique as an alternative to our algorithm, even though its testing has been limited to scenes with more controlled illumination.

The results for the green card are included to illustrate that the varying illumination constraint can be very useful in the case when there is a paucity of other information. Most colour constancy algorithms require a good selection of surfaces for reliable performance. The ability of this algorithm to go beyond that in the case of varying illumination is encouraging.

5 Conclusion

We have presented a new algorithm for colour constancy which builds upon the recent gamut-based algorithms of Forsyth [9] and Finlayson [7, 8]. The new algorithm models the illumination via a diagonal transform or equivalently a coefficient rule model. Within the diagonal model framework, the algorithm combines the constraints provided by the observed gamut of image colours, by a priori knowledge of the set of likely illuminants, and by the variation in illumination across a scene. Existing algorithms that use the information inherent in illumination variation assume that some unspecified pre-processing stage has already identified the variation, and thus are not fully automated.

Identifying illumination variation is in itself a difficult problem. The Retinex algorithm is the only alternative colour constancy algorithm designed for scene conditions similar to those investigated in this paper. Nonetheless, it was restricted to flat Mondrian scenes and is known not to work very well for a variety of reasons.

The new algorithm is more powerful than Retinex both because it incorporates a new technique for identifying the relative illumination field and because the it actually uses the illumination variation when solving for colour constancy. While many improvements are still possible, tests using both synthetic and real image data for three-dimensional scenes verify that the algorithm works well.

6 References

1. A. Blake, "Boundary conditions for lightness computation in Mondrian world", *Computer Vision, Graphics, and Image Processing*, **32**, pp. 314-327, (1985)
2. G. Buchsbaum, "A spatial processor model for object colour perception", *Journal of the Franklin Institute*, **310**, pp. 1-26, (1980)
3. W. Freeman and David Brainard, "Bayesian Decision Theory, the Maximum Local Mass Estimate, and Color Constancy", in *Proceedings: Fifth International Conference on Computer Vision*, pp 210-217, (IEEE Computer Society Press, 1995)
4. B. V. Funt, M. S. Drew, M. Brockington, "Recovering Shading from Color Images", in *Proceedings: Second European Conference on Computer Vision*, G. Sandini, ed., 1992.
5. G.D. Finlayson and M.S. Drew and B.V. Funt, "Spectral Sharpening: Sensor Transformations for Improved Color Constancy", *J. Opt. Soc. Am. A*, **11**, 5, pp. 1553-1563, (1994)
6. G.D. Finlayson and M.S. Drew and B.V. Funt, "Color Constancy: Generalized Diagonal Transforms Suffice", *J. Opt. Soc. Am. A*, **11**, 11, pp. 3011-3020, (1994)
7. G. D. Finlayson, B. V. Funt, and K. Barnard, "Color Constancy Under Varying Illumination", in *Proceedings: Fifth International Conference on Computer Vision*, pp 720-725, 1995.
8. G. D. Finlayson, "Color Constancy in Diagonal Chromaticity Space", in *Proceedings: Fifth International Conference on Computer Vision*, pp 218-223, (IEEE Computer Society Press, 1995).
9. D. Forsyth, "A novel algorithm for color constancy", *Int. J. Computer. Vision*, **5**, pp. 5-36, (1990)
10. R. Gershon and A.D. Jepson and J.K. Tsotsos, "Ambient illumination and the determination of material changes", *J. Opt. Soc. Am. A*, **3**, 10, pp. 1700-1707, (1986)
11. B.K.P. Horn, "Determining lightness from an image", *Computer Vision, Graphics, and Image Processing*, **3**, pp. 277-299, (1974)
12. D.B. Judd and D.L. MacAdam and G. Wyszecki, "Spectral Distribution of Typical Daylight as a Function of Correlated Color Temperature", *J. Opt. Soc. Am.* , **54**, 8, pp. 1031-1040, (August 1964)
13. E.H. Land, "The Retinex theory of Color Vision", *Scientific American*, 108-129, (1977)
14. John J. McCann, Suzanne P. McKee, and Thomas H. Taylor, "Quantitative Studies in Retinex Theory", *Vision Research*, **16**, pp. 445-458, (1976)
15. L.T. Maloney and B.A. Wandell, "Color constancy: a method for recovering surface spectral reflectance", *J. Opt. Soc. Am. A*, **3**, 1, pp. 29-33, (1986)
16. M. Tsukada and Y. Ohta, "An Approach to Color Constancy Using Multiple Images", in *Proceedings Third International Conference on Computer Vision*, (IEEE Computer Society, 1990)
17. G. Wyszecki and W.S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulas*, 2nd edition, (Wiley, New York, 1982)
18. M. D'Zmura and G. Iverson, "Color constancy. I. Basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces", *J. Opt. Soc. Am. A*, **10**, 10, pp. 2148-2165, (1993)