The bright-chromagenic algorithm for illuminant estimation

Clément Fredembach[1,2] ***and Graham Finlayson***[1]
1) School of Computing Sciences, University of East Anglia, Norwich, UK
2)Audio-Visual Communications Laboratory, Ecole Polytechnique Federale de Lausanne (EPFL), Lausanne, Switzerland
clement.fredembach@epfl.ch, graham@cmp.uea.ac.uk

Abstract

In this paper, we propose a new algorithm for illuminant estimation. We begin by reviewing the concept of chromagenic colour constancy, where two pictures are taken from each scene: a normal one and one where a coloured filter is placed in front of the camera, and look at parameters known to affect its performance such as filters and sensor choice.

We show that the basic formulation of the chromagenic algorithm has inherent weaknesses: a need for perfectly registered images and occasional large errors in illuminant estimation. Our first contribution is to analyse the algorithm performance with respect to the reflectances present in a scene and demonstrate that fairly bright and desaturated reflectances (e.g., achromatic and pastel colours) provide significantly better chromagenic illuminant estimation.

This analysis leads to the bright-chromagenic algorithm. We show that it not only remedies the large error problem but also allows us to relax the image registration constraint. Experiments performed on a variety of synthetic and real data show that the newly designed bright-chromagenic algorithm significantly -in a strict statistical sense- outperforms current illuminant estimation methods, including those having a substantially higher complexity.

Introduction

The human visual system is, to a certain extent, colour constant [1, 4, 5]; that is, it discounts the colour of the illumination. This is why, for example, snow always appears white, no matter which illuminant it is observed under.

However, it has proven difficult to emulate this colour constancy in manufactured devices. This is not only a problem in image reproduction but also for a variety of computer vision tasks, such as tracking [19], indexing [24] and scene analysis [20] where stable measures of reflectance are sought or assumed for objects in a scene.

Colour constancy is generally broken down into two parts: first, the colour of the prevailing illuminant is estimated. At a second stage, the colour bias due to illumination is removed. This second part is, in fact, relatively easy [18] and so most colour constancy algorithms focus on the illuminant estimation problem.

Numerous algorithms for illuminant estimation have been proposed and can broadly be split in two groups. Algorithms in the first group make simple assumptions about the scene being observed, such as MaxRGB -a maximally reflective patch exist in the image-, Gray World -the average reflectance in a scene is gray [6] or some sort of gray [12, 26]. Another group of algorithms comprises more sophisticated approaches such as neural networks [7], colour by correlation -a bayesian method that correlates the RGBs in the image with plausible RGBs under various illuminants to find the best illuminant- [11] and Gamut Mapping methods [13, 14]. Generally, the most complex algorithms perform better, but at the expense of a -much- greater computational complexity.

It has been proposed that a chromagenic camera, which takes two pictures of every scene (with and without a coloured filter), makes illuminant estimation easier. This idea of using two pictures has also been explored in flash/non-flash image pairs that can be used to estimate the original scene illuminant [8, 21, 23].

The standard chromagenic colour constancy algorithm [10] works in two stages: the *training* stage is a preprocessing step where the relationship between filtered and unfiltered RGBs is calculated using a given filter, camera sensitivities and a number of candidate lights. Then, those relations are *tested* on other images in order to estimate the actual scene illuminant. While the general outcome of the algorithm shows good performance, two problems usually remain: for certain combinations of reflectances and illuminants, the error between the estimated and actual light can be large. Also, in order to achieve good performance, one has to compare RGBs transitions that occur between *identical* reflectances. It is, in essence, a pixel-level method and so the algorithm can fail when the pair of images is not *exactly* registered.

Our approach here starts by asking the question: "Everything else (lights, camera sensitivities and the chromagenic filter) being equal, what is the influence of changing reflectances on the transforms and the estimation error?". To answer this question, we first select 287 typical lights and 1995 reflectances from the Simon Fraser database [3] and a filter from possible Wratten photographic filters. Using these elements, we create synthetic images composed of 1 randomly selected illuminant and of 1-8 distinct reflectances. Testing the algorithm on these images allow us to evaluate which RGBs exhibit a good -very low errors- or bad -very high errors- estimate of the illuminant. The result show that achromatic reflectances yield lower errors than strongly chromatic ones. Restricting the testing on the least chromatic reflectances thus enhances the overall performance of the algorithm and greatly decreases the maximum observed error. This is, however not enough to allow the algorithm to work on real images. Indeed, im-

age registration is still a problem and noise levels in dark -chromatic or achromatic- regions can impact performance.

We address this registration problem by first noting that, independently of the image capture conditions, a white reflectance (which is achromatic) will still be the brightest value after filtering. We therefore restrict the algorithm to test correspondences between the "whitest" RGBs of both filtered and unfiltered images as this will allow us to use the information that the most achromatic reflectances are more reliable for the chromagenic algorithm while restricting ourselves to parts of the image that have a high signal to noise ratio. Practically, we average the RGB values of the brightest 3% of the original image and compare them to the brightest 3% values of the filtered image.

We validate our experiments by testing this modified algorithm on three different image databases and compare the results we obtain with those achieved by both simple and complex methods. We find that our algorithm performs as well as any other available method using the Wilcoxon sign test to discriminate between algorithm performance, as recommended in [17].

The Chromagenic Algorithm

The standard chromagenic illuminant estimation algorithm proceeds as follows: Let $S(\lambda)$ be the descriptor of surface reflectances, $E(\lambda)$ the scene illuminant SPD, $Q_k(\lambda)$ the camera sensitivities (we consider here trichromatic cameras, so $k = \{R, G, B\}$) and $F(\lambda)$ be the transmittance of the colour filter placed in front of the camera.

The sensor responses of the unfiltered, $\underline{\rho}$, and filtered, $\underline{\rho}^F$, image can be written as:

$$\rho_k = \int_\omega E(\lambda)S(\lambda)Q_k(\lambda)d\lambda \tag{1}$$

$$\rho_k^F = \int_\omega E(\lambda)S(\lambda)F(\lambda)Q_k(\lambda)d\lambda \tag{2}$$

thus, for each scene we recover six responses per pixel that form the input to the illuminant estimation problem. For the purposes of this work, we set out to recover $\underline{\rho}_E$: the RGB of a white surface under the scene illuminant E.

Let us first consider the equations of filtered and unfiltered image formation (1) and (2). We can approximate the filtered image by posing a second illuminant, $E^F(\lambda)$ so that it is equivalent to putting the filter $F(\lambda)$ in front of the light source $E(\lambda)$, i.e., $E^F(\lambda) = F(\lambda)E(\lambda)$. We can therefore think of $\underline{\rho}$ and $\underline{\rho}^F$ as sensor responses of a single surface under two different illuminants. Thus, we can then write:

$$\underline{\rho}^F = T_E^F \underline{\rho} \tag{3}$$

where T_E^F is a 3×3 linear transform that depends on both the chromagenic filter and the scene illuminant.

Barring the cases where camera sensors behave like Dirac functions and where the filter used is a neutral density one (where chromagenic colour constancy is not possible [10]), the transforms can be pre-computed by choosing a set of n typical scene illuminants: $E_i(\lambda),\ i = 1, \ldots, n$ and a set of m surface reflectances: $S_j(\lambda),\ j = 1, \ldots, m$ representative of the real world. For each illuminant i, we create a $3 \times m$ matrix $\mathcal{Q}_i$ whose j^{th} column contains the sensor response of the j^{th} surface illuminated by the i^{th} illuminant. We also create $\mathcal{Q}_i^F$, which contains the equivalent filtered responses. For each illuminant, we can then define the transform matrix as:

$$\mathcal{T}_i = \mathcal{Q}_i^F \mathcal{Q}_i^+ \tag{4}$$

where $+$ denotes the Moore-Penrose pseudo-inverse operator: $\mathcal{Q}^+ = \mathcal{Q}^{\mathrm{T}}(\mathcal{Q}\mathcal{Q}^{\mathrm{T}})^{-1}$. $\mathcal{T}_i$ can then be described as the transform that best maps, in a least square sense, unfiltered to filtered responses under illuminant i.

Once the n transforms have been pre-computed, the illuminant estimation proceeds as follows: let $\mathcal{Q}$ and $\mathcal{Q}^F$ denote the $3 \times m$ matrices of unfiltered and filtered RGBs of arbitrary reflectances under an unknown light. For each plausible illuminant we calculate the fitting error, e_i, as:

$$e_i = \|\mathcal{T}_i\mathcal{Q} - \mathcal{Q}^F\|,\ \ i = 1, \ldots, n \tag{5}$$

under the assumption that $E_i(\lambda)$ is the actual scene illuminant. We then choose the transform that minimizes the error and surmise that it corresponds to the scene illuminant. Our estimated illuminant is $E_{\text{est}}(\lambda)$ where

$$\text{est} = \arg\min_i(e_i)\ \ i = 1, \cdots, n \tag{6}$$

While in general the chromagenic algorithm can deliver good colour constancy, it has two major weaknesses: the first one is that, though good on *average*, the performance can, on occasion, be (very) poor. The second problem comes from equation (5) where we see that the fitting error for each candidate illuminant is evaluated on a per-pixel basis. This implies that, for the algorithm to deliver an optimal performance, the two images have to be perfectly registered, a demanding requirement when images are taken outside of the lab. Registration methods can be of some help, but since we are looking for an exact registration at pixel-level, they may not be sufficient.

The Bright-Chromagenic Algorithm

We propose here a modification of the chromagenic algorithm that has three outcomes: it improves the average illuminant estimation performance, it reduces the maximal errors observed when the estimation is erroneous and, more importantly, it allows the algorithm to be used on unregistered images.

Let us first look back at equations (1) and (2) where we see that the sensor responses depend on the scene illuminant, the chromagenic filter, the sensor sensitivities and the surface reflectances. It follows that the linear transforms $\mathcal{T}_i$ also depend -to some degree- on those factors. Among them, the illuminant is what we aim to recover so it has to remain a variable; achievable improvements due to choosing both the filter and the sensors sensitivities have been explored in [9]. We therefore set out to explore the only unknown that remain in the equation: the scene reflectances.

Building a model based on reflectances can be difficult for one has, in general, no control as to which reflectances are present in a scene. This uncertainty is the reason why

simple estimation methods such as gray-world and Max-RGB are unreliable (if every scene contained a "white-like" patch, Max-RGB would be very accurate). Another issue is linked to the actual input of most illuminant estimation algorithms: RGBs. Indeed, both the training and testing steps of the chromagenic algorithm are RGBs, which are composed of all of the image formation process parameters.

Reflectances Analysis

For the chromagenic algorithm to work well, the transforms $\mathcal{T}_i$ that map RGBs to their filtered counterparts should depend as much as possible on the illuminant. Here, we want to quantify the variance of the transforms when the illuminant changes and compare it to the variance observed when the illuminant is fixed but the reflectances vary.

To perform this assessment, we follow the methodology of [9] and [10] in our choice of parameters. The illuminants belong to a set of 87 measured illuminants Spectral Power Distributions (SPD) that include most common light sources. These SPDs are sampled every 10nm, from 380nm to 780nm. More details about the set can be found in [3] while the set itself is available online [2]. For surface reflectances, we use a synthesized set of 1995 Munsell surface reflectances [25]. The reflectances are also sampled every 10nm from 380nm to 780nm, more details about that set can be found in [22].

Concerning the choice of camera sensitivities and filter, we use the sensors of a Sony DXC-930 camera (as in [3] and [17]) and a Wratten 81B filter (a yellowish filter). Both the filter and the sensor sensitivities are shown in Fig. 1.

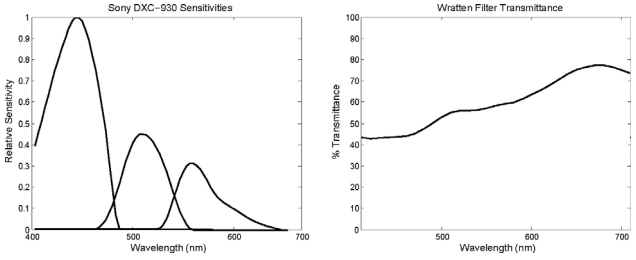

Figure 1. *Left: The Sony DXC-930 sensitivities. Right: The transmittance of the 81B Wratten filter used in the experiments.*

Modelling In and Outliers

We know that the estimation accuracy will partly depend on the set of tested reflectances and we would like to model that dependency. Because there is a multitude of combinations of illuminants and reflectances, we will analyse the performance of the algorithm on scenes with a single reflectance (where good estimates can still be obtained [10]). For each scene we estimate we estimate the RGB of the light and calculate the angle to the actual RGB of the illuminant.

The transforms are calculated as before, thus creating 87 of them. The test set for this experiment consists of all possible single reflectance scenes under 287 illuminants, i.e., ∼570,000 pairs of filtered and unfiltered RGBs. This larger illuminant set used in testing covers the same gamut as the 87 training lights; the chromagenic algorithm will select one of the 87 lights as the scene illuminant.

The angular errors for all the 570,000 scenes range from 0 to 42 degrees, with a mean of 9.3 and a median of 5. For this particular dataset, our experiments indicate that an angular error of 3 degrees or less is necessary for acceptable colour cast removal.

To understand what is happening, we look at the RGBs that comprise the top and bottom 20% of the cumulative histogram. We plot the brightness and saturation (i.e., S and V of the HSV space) of these RGBs in Fig. 2a for the highest errors and Fig. 2b for the lowest ones. It is clear that low errors correlate with fairly desaturated RGBs (pastel tones and achromatic) whereas high errors correlate with dark and saturated RGBs. More interesting perhaps is the fact that bright achromatic RGBs are not at all present amongst the RGBs linked to high errors. We also observed the behaviour of the algorithm with respect to hue but found no significant hue dependency.

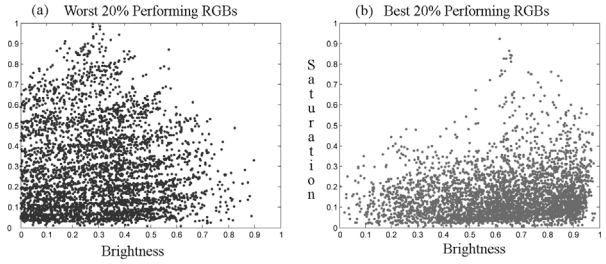

Figure 2. (a) Brightness-Saturation scatter plot of the 20% worst performing RGBs. (b) Brightness-Saturation scatter plot of the 20% best performing

Assuming a uniform distribution of colours in an image, we propose that it is easy to find RGBs and their filtered counterparts that belong to this preferred set. We simply look for a small percentage of the brightest image regions. We therefore modify the chromagenic algorithm formulation so that only bright image pixels are considered.

The **Bright-Chromagenic** algorithm is defined as:

Preprocessing: For a database of m lights $E_i(\lambda)$ and n surfaces $S_j(\lambda)$ calculate $\mathcal{T}_i \approx \mathcal{Q}_i^F \mathcal{Q}_i^+$ where $\mathcal{Q}_i$ and $\mathcal{Q}_i^F$ represent the matrices of unfiltered and filtered sensor responses to the n surfaces under the ith light and $+$ denotes a *pseudo-type-inverse*

Operation: Given P surfaces in an image we have $3 \times P$ matrices Q and Q^F. From these matrices we choose the r% brightest pixels giving the matrices $\mathcal{Q}$ and $\mathcal{Q}^F$, where the brightest pixels are the ones with the largest $R^2+G^2+B^2$ value. Then the estimate of the scene illuminant is $\underline{\rho}_{est}$ where

$$est = \arg\min_i(err_i) \;\; (i = 1, 2, \cdots, m)$$

and

$$err_i = ||\mathcal{T}_i \mathcal{Q} - \mathcal{Q}^F||$$

This formulation is robust since it does not make assumptions about which reflectances might or might not be present in the scene, i.e., if there are no bright reflectances in the image, the bright-chromagenic algorithm will still have an equivalent performance to the original chromagenic algorithm.

Because we "exclude" -select them only if no other are available- the worst performing RGBs, we expect the bright chromagenic algorithm to significantly reduce the worst errors.

We also make the following observation: if we assume that scenes admit a diversity of reflectances, then it follows that -if the filter does not vary too drastically across the spectrum- the brightest unfiltered RGBs will most likely be mapped onto the brightest filtered RGBs. If we are relatively conservative with the number of bright pixels we use to estimate the illuminant (we typically use the top 1-3% of the brightest pixels[1]), the bright-chromagenic algorithm will then be able to estimate illuminants even when the images are not registered. Both these properties are verified in our experiments.

Experiments

In this section, we analyse the performance of the bright-chromagenic algorithm, and compare it to various other illuminant estimation methods on three datasets of increasing difficulty, ranging from perfect synthetic data to real images taken with a "prosumer" digital camera whose sensitivities are unknown.

We evaluate algorithms according to the framework of Hordley and Finlayson [17] where it was shown that, if one wants to summarize the performance of an illuminant estimation algorithm over a dataset, one should prefer the median angular over the mean or Root Mean Square error.

Angular error is an intensity independent error measure that is widely used in the literature [3, 15, 17]. It is the measure between the sensor responses of a white reflectance under both the estimated and actual scene illuminant. If we denote these responses by $\underline{\rho}_{\text{est}}$ and $\underline{\rho}_E$ respectively, the angular error e_{Ang} is calculated as:

$$e_{\text{Ang}} = \text{acos}\left(\frac{\underline{\rho}_E^{\text{T}}\underline{\rho}_{\text{est}}}{\|\underline{\rho}_E\|\,\|\underline{\rho}_{\text{est}}\|}\right) \quad (7)$$

The use of a median statistic permits to assess if the difference of performance between to algorithms is statistically significant at chosen confidence level. That significance is given by using the Wilcoxon unranked sign test [16] at a 95% confidence level.

To simplify the writing, we will use the following notations: S_M is the set of 1995 synthetised Munsell reflectances of [25], E_{87} and E_{287} are the sets of respectively 87 and 287 illuminants from [2].

Synthetic Reflectances and Lights

The test on synthetic images is run according to the testing protocol proposed by Barnard et al. in [3]:

[1]To avoid saturated pixels, we use the brightest 3% pixels of *non-maximal* value

♯ surfaces	1	2	4	8	16	32	rank
Chromagenic	6	5.2	4.5	3.5	3	2.2	2
Max RGB	9.7	7.9	6.1	4	2.9	2.6	6
Grey World	9.1	7.3	5.8	4.9	4.8	4.8	8
Database GW	9.5	6.7	4.8	3.4	2.8	2.5	4
LP GM	9.6	6.7	4.8	3.3	2.7	2.4	4
Neural Network	8.8	7.1	5	4	2.9	2.6	6
Colour by Corr.	6.9	5	3.5	3.1	2.4	2.3	2
Bright-Chroma.	6	5.2	4.1	2.8	2.1	0.9	**1**

Average median angular error for 1000 tests at each complexity level. The last column is the rank, based on the 32 surfaces test, according to the Wilcoxon sign test with a confidence level of 95%

Training: The linear transforms are created by imaging the whole of S_M under E_{87}, thus generating 87 transforms. We use the Sony-DXC camera sensitivities and the 81B Wratten filter, whose transmittance is shown in Fig. 1.

Testing: We generate 1000 images containing n different reflectances,
$n = \{1, 2, 4, 8, 16, 32\}$, randomly taken from S_M. We then illuminate these images with one light taken at random from E_{287}.

We estimate the illuminant of each image using both the bright-chromagenic and the original chromagenic algorithms. For images where $n > 4$, the bright-chromagenic version estimates the illuminant based on the four brightest reflectances only.

The results are displayed in Table 1 where the last column indicates the ranking of the considered algorithms, taking in to account the results of Wilcoxon's sign test. An algorithm is ranked better than another if its median is lower *and* if the difference is statistically significant at the 95% level. If the sign test is inconclusive, the algorithms will be ranked equally.

The results show that the bright-chromagenic algorithm performs significantly better than all other methods. We also see that the more complex methods form a group that is, in turn, significantly better than the simpler scene assumptions algorithms.

An additional result, shown in Fig. 3, is the reduction in maximal error achieved by the bright chromagenic algorithm. This experiment validates our selection of the bright RGBs to reduce the high max errors observed with the original chromagenic algorithm.

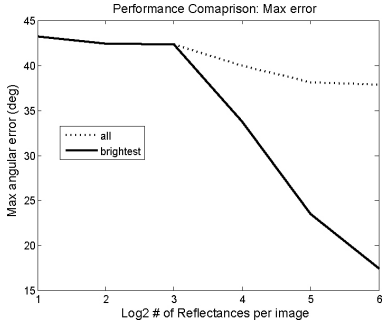

Figure 3. Comparison of the max angular error between the original and the bright chromagenic algorithm, one can see the significant reduction achieved by selecting only the brightest RGBs.

SFU dataset

The next set we evaluate our algorithm on is the non-specular Simon Fraser University (SFU) dataset, which is described in detail in [3].

The data set consists of 31 colourful objects captured under 11 illuminants. Importantly, the images are not registered (in fact, the objects were rotated in between two pictures when creating the dataset).

This experiment differs from the previous ones in the sense that here we are directly provided with the RGBs of the images instead of reflectances. This, plus the non-registration of the image will provide a difficult test for the bright-chromagenic algorithm. The SFU dataset has been used in several illuminant estimation comparisons because ground truth is provided in addition to the images themselves. That is, both the SPD of the 11 illuminants (they are actually a subset of E_{87}) and the camera sensitivities are given (the camera used to take the images is the Sony-DXC 930 whose sensitivities are shown in Fig. 1 and that we used in the previous tests).

To perform chromagenic illuminant estimation, we require pairs of images taken with and without a coloured filter. However, if we only have image RGBs at our disposal we cannot retrospectively model the filtered responses. As it turns out, 8 out of the 11 illuminants present in the set come in pairs: the original lamp lights and those lights filtered with a blue filter. Since the actual illuminant SPDs are known, we can derive the filter that was used -we do not actually know what it is- by dividing the spectra of the lights by the filtered ones. The eight -two pairs of four- lights that are considered and the derived filter are shown in Fig. 4.

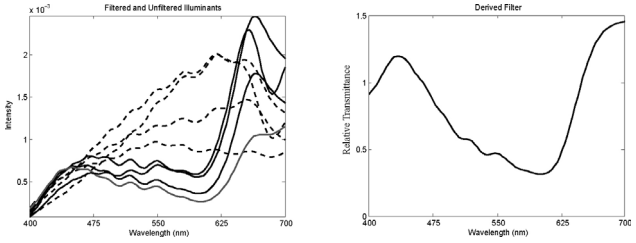

Figure 4. Left: The 8 light sources considered in this experiment. The dashed lines are spectra of the light sources, while the continuous ones are from the filtered sources. Right: The Filter derived from the light source data.

Training: The transforms $\mathcal{T}_i$ are obtained by imaging the *synthetic* reflectances S_M under the illuminants of E_{87}. As filter, we use the one derived from the eight illuminants shown in Fig. 4; the camera sensitivities are the same as in previous experiments.

Testing: To test the algorithm, we estimate the illuminant of all the possible pairs of images (124 pairs in total) using the top 3% of the brightest pixels in both filtered and unfiltered images independently. These pixels typically belong to one or two of the surfaces in the scene (we do not need a white reflectance *per se*, we simply use the brightest ones available). Since the images are not registered, we are able to test our hypothesis that, in general, the brightest pixels in both images come from the same surfaces, and that the bright-chromagenic algorithm does away with the need for registration.

Algorithm	Mean	Median	Rank
Bright-Chromagenic	4.8	3.4	**1**
Max RGB	6.4	4.1	4
Grey World	11.9	9.3	7
Database GW	10	7	5
Neural Network	8.9	7.8	5
LP Gamut Mapping	5.5	3.8	2
Colour by Corr.	6	3.6	2

Mean and median angular errors over the SFU dataset. The ranks are significant at the 95% level.

The angular errors reported in the first two columns of Table 3 show that, despite its simplicity, the bright chromagenic algorithm outperforms in terms of both mean and median angular error all other algorithms at the 95% confidence level. The original chromagenic algorithm is not shown here since its registration requirement is not fulfilled.

Perhaps the most remarkable aspect of the bright-chromagenic algorithm is that, despite modelling the transforms on synthetic data with a filter derived from measurements, it is still able to estimate accurately the illuminants of real, significantly non-registered, images.

Real Images

The last, experiment is designed to evaluate the performance of the bright-chromagenic algorithm *in situ*. Whereas the previous datasets were obtained in "controlled conditions" (purely synthetic data and controlled lighting environment), we use here a set of real-world images taken with a digital camera whose specifications are unknown.

Chromagenic Photography:

For the illuminant estimation to be meaningful, we must take a couple of precautions when capturing the images. The camera we used is a Nikon D70: a prosumer Single Lens Reflex camera. The camera was setup to capture linear Raw, unprocessed, images[2]. To prevent the camera from using a different white balance between filtered and unfiltered images, all images were captured with the white balance set to "daylight". We also used a tripod and and a remote shutter release to minimize registration errors.

For the filter, we used an actual 81B-type Wratten filter. The captured images were then exported, using Nikon capture, as 16bits/channel linear tiff images.

The dataset consists of 86 pair of images taken under a variety of indoor and outdoor illuminants. In every scene, we placed a Macbeth colour checker that is used to accurately determine the colour of the prevailing light, thereby providing a ground truth to assess the accuracy of illuminant estimation algorithms.

From the dataset, we then created separate training and testing sets. The training set consists of the 24 Macbeth patches present in all the images. The testing set is created by blacking the colour checker from the images. Images from the original (with the colour chart) and the testing set are shown in Fig. 5.

[2]In fact, even with Raw settings, the camera and associated software will process the image somewhat; it is however as unprocessed as one can have with a general purpose digital camera.

Algorithm	Mean	Median	Rank
Bright-Chromagenic	7.09	4.15	**1**
Max RGB	7.87	7	3
Grey World	11	10.8	5
Chromagenic	7.96	5.1	2
ℓ^4 Gray	10.3	9.6	4

Mean and median angular errors over the 86 real images. The ranks are significant at the 95% level.

We note that, despite the precautions taken, the registration between images is not perfect and some image regions can be over-exposed. Additionally, multiple illumination is sometimes present in images, which can lead to errors when estimating the prevailing illuminant.

Figure 5. Top row: unfiltered/filtered image pairs from our real image set. Bottom row: the testing images, where the charts have been cropped out.

Training: We create 86 linear transforms, using the 24 RGBs of the colour checker present in each image.
Testing: We estimate the illuminant for both the original and bright-chromagenic (using the top 3% brightest pixels) algorithms on the 86 pairs of images that have the colour chart clipped out.

The results are shown in Table 3 and illustrate that the most accurate illuminant estimation is given by the bright-chromagenic algorithm. The results otherwise exhibit the same behaviour than previous experiments.

Conclusion

In this paper, a detailed error analysis demonstrated that bright pixels in images lead to smaller chromagenic estimation errors. This led to the bright-chromagenic algorithm, which bases its estimation only on a fixed percentage of the brightest pixels in the filtered and unfiltered images. Importantly, these pixels are chosen independently in each image so there is no need for image registration. Experiments on various sets of synthetic and real data demonstrate that the bright-chromagenic algorithm delivers a statistically significant better illuminant estimation than all other tested algorithms.

References

[1] L. Arend and A. Reeves. Simultaneous color constancy. *JOSA part A*, 3:1743–1751, 1986.

[2] K. Barnard. Data for computer vision and computational colour vision. http://www.cs.sfu.ca/∼color/data, 2002.

[3] K. Barnard, V. Cardei, and B. Funt. A comparison of computational color constancy algorithms- part i: methodology and experiments with synthetized data. *IEEE Trans. on Image Processing*, 11:972–984, 2002.

[4] D.H. Brainard. Color constancy in the nearly natural image: 2 achromatic loci. *Journal of The Optical Scociety of America part A*, 15:307–325, 1998.

[5] D.H. Brainard and W.T. Freeman. Bayesian color constancy. *JOSA part A*, 14:1393–1411, 1997.

[6] G. Buchsbaum. A spatial processor model for object colour perception. *Franklin Institute*, 310:1–26, 1980.

[7] V.C. Cardei, B. Funt, and K. Barnard. Estimating the scene illuminant chromaticity using a neural network. *Journal of The Optical Scociety of America part A*, 19:2374–2386, 2002.

[8] J.M. DiCarlo, F. Xiao, and B.A. Wandell. Illuminating illumination. In *Proc. of the ninth Color Imaging Conference*, pages 27–34, 2001.

[9] G. Finlayson, S. Hordley, and P. Morovic. Chromagenic filter design. In *Proc. of the 10th AIC*, pages 1079–1083, 2005.

[10] G. Finlayson, S. Hordley, and P. Morovic. Colour constancy using the chromagenic constraint. In *Computer Vision and Pattern Recognition (CVPR) 2005*, pages 1079–1086, 2005.

[11] G.D. Finlayson, S.D. Hordley, and P.M. Hubel. Color by correlation: A simple, unifying framework for color constancy. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 23:1209–1221, 2001.

[12] G.D. Finlayson and E. Trezzi. Shades of gray and colour constancy. In *Proc. of the Twelfth Color Imaging Conference*, pages 37–41, 2004.

[13] G.D. Finlayson and R. Xu. Convex progamming colour constancy. In *IEEE Workshop on Color and Photometric Methods in Computer Vision*, pages –, 2003.

[14] D.A. Forsyth. A novel algorithm for colour constancy. *Intl Journal of Computer Vision*, 5:5–36, 1990.

[15] B. Funt, K. Barnard, and L. Martin. Is machine color constancy good enough? In *Proc. of the Fifth Color Imaging Conference*, pages 455–459, 1998.

[16] R.V. Hogg and E.A. Tanis. *probability and Statistical Inference*. Prentice Hall, 2001.

[17] S.D. Hordley and G.D. Finlayson. Reevaluation of color constancy algorithm performance. *Journal of the Optical Society of America A*, 24:1008–1020, 2006.

[18] P.M: Hubel, J. Holm, G.D. Finlayson, and M.S. Drew. Matrix calculations for digital photography. In *Proc. the Fith color Imaging Conference*, pages 105–111, 1997.

[19] H. Jiang and M. Drew. Tracking objects with shadows. In *CME03: International Conference on Multimedia and Expo,*, pages 100–105, 2003.

[20] G. J. Klinker, S. A. Shafer, and T. Kanade. A physical approach to color image understanding,. *International Journal of Computer Vision*, 4:7–38, 1990.

[21] C. Lu and M.S. Drew. Practical scene illuminant estimation via flash/no-flash pairs. In *Proc. of the fourteenth Color Imaging Conference*, pages 1–1, 2006.

[22] J. Parkkinen and T. Jaaskelainen. Characteristic spectra of munsell colors. *Journal of The Optical Scociety of America part A*, 6:318–322, 1989.

[23] G. Petschnigg, R. Szeliski, M. Agrawala, M.F. Cohen, H. Hoppe, and K. Toyama. Digital photography with flash and no-flash image pairs. *ACM Trans. on Graphics*, 23:664–672, 2004.

[24] M. J. Swain and D. H. Ballard. Color indexing. *International Journal of Computer Vision*, 7:11–32, 1991.

[25] Joensuu University. Musell colors matt [online]. spectral.joensuu.fi/databases/download/munsell_atof.htm

[26] J. van de Weijer and T. Gevers. Color constancy based on the grey-edge hypothesis. In *Proc. of the International Conference on Image Processing, Genoa*, pages 722–725, 2005.