Relaxed Grey-World: Computational Colour Constancy by Surface Matching

Francesc Tous, María Vanrell, and Ramón Baldrich

Computer Vision Center, Dept. Informàtica
Universitat Autònoma de Barcelona
08193 Bellaterra (Barcelona), Spain
{ftous,maria,ramon}@cvc.uab.es

Abstract. In this paper we present a new approach to computational colour constancy problem based on the process of surface matching. Classical colour constancy methods do not usually rely on this important source of information and they often use only partial information in the images. Our proposal is to introduce the use of a set of canonical surfaces and its matching versus the content of the image using a 'relaxed' grey-world assumption to perform colour constancy. Therefore, our approach takes into account information not considered in previous methods, which normally rely on statistical information in the image like highest luminance or image gamuts. Nevertheless the selection of the canonical surfaces is not a trivial process and should be studied deeply.

1 Introduction

The human visual system has the capability to perceive the same colour for a given surface regardless the colour of the illuminating light. This is a fundamental property to colour vision and pursues the perception of a stable coloured world, even though the stimulus reaching the retina differs for the same surface under different conditions of illumination. The perceived colour of a white patch under a blue sky compared to the same patch in a room with a light bulb is perceived as the same colour. Actually in the first situation the reflected light reaching the eye has a bluish spectrum compared to the reddish reflected light of the second. This ability is known as colour constancy, the constant appearance of surface colours despite changes in the colour of the illumination. The mechanisms of human colour constancy have not yet been completely understood, and there are different approaches trying to explain them [1–4].

2 Background

RGB images are formed by the light reflected from different surfaces reaching three sensors that integrate the incident light at different wavelengths. The color of a surface depends on the surface reflectance and the colour of the incident light. The aim of computational colour constancy is to find an illuminant invariant description of a scene from an image taken under unknown lighting conditions.

J.S. Marques et al. (Eds.): IbPRIA 2005, LNCS 3522, pp. 192–199, 2005.

This process is often performed in two steps: (1) estimate the illuminant parameters and (2) use those parameters to build illuminant independent description of the scene. For these methods a canonical illuminant must be defined, i.e. an illuminant for which the camera is balanced and the colours appear in a trustworthy form. Under this illuminant, the RGB values of an image of a scene can be taken as descriptors of the surfaces. There is a wide literature on computational colour constancy methods [5–10]. None of them performs perfectly on all kind of images under weak assumptions.

Many of these methods directly estimate the illumination change from the unknown illuminant to the canonical illuminant. Considering the von Kries adaptation model [11], the transform of an illuminant change can be modelled by a linear diagonal model, as proven in [12]. For example, the RGB response of a camera to a white patch under an unknown illuminant is (R_w^U, G_w^U, B_w^U) and the response under the canonical illuminant is (R_w^C, G_w^C, B_w^C), the illuminant change from the unknown to the canonical illuminant can be obtained by scaling the three channels by $R_w^C/R_w^U, G_w^C/G_w^U, B_w^C/B_w^U$ respectively. Thus, the colour of the illuminant of an RGB image can be modified by a diagonal change (1),

$$(R^C, G^C, B^C) = \begin{pmatrix} \alpha & 0 & 0 \\ 0 & \beta & 0 \\ 0 & 0 & \gamma \end{pmatrix} \begin{pmatrix} R^U \\ G^U \\ B^U \end{pmatrix} \tag{1}$$

where $\alpha = R_w^C/R_w^U, \beta = G_w^C/G_w^U, \gamma = B_w^C/B_w^U$. In a typical colour constancy problem, we have acquired the image under an unknown illuminant, (R^U, G^U, B^U), and try to obtain the surface descriptors, (R^C, G^C, B^C). The triplet (α, β, γ) is called a map, and knowing the actual map implies a guessing of the unknown illuminant.

The different methods proposed in the literature can be sorted in different classes regarding the assumptions they are based on. The first family of algorithms are established upon the Retinex theory of human vision [13], which goes beyond simple illuminant estimation. The theory assumes that slight spatial changes in the response are due to changes in illumination or noise, and large changes correspond to surface changes. The idea is to run random paths from every surface and compute the ratio of the responses in each channel. The descriptor of a pixel is given by the average of the ratios from different paths beginning at the same pixel.

Another group are the Grey World methods. They are based on the assumption that the scene is colorimetrically unbiased (no particular colour predominates). In other words, supposes that a complex scene contains a wide range of reflectances, whose mean is a grey reflectance (for instance, a uniform reflectance with half of the maximum energy). Therefore, to correct the illumination of an image the map that takes the average of the image to the average of the canonical gamut is used as an estimation of the illuminant change.

One of the most important groups to date are the Gamut Mapping methods. All of them are based on the idea of canonical gamut firstly introduced by Forsyth in [5]. If we consider all the possible reflectances under a canonical illuminant we obtain a convex set of RGB values, which are the whole set of values that can be perceived under the canonical illuminant for a given camera. This introduces a device/illuminant restriction, and it can be used to build a set of illuminant changes that are feasible, i.e. which map the image gamut within the canonical gamut. To build the feasible set of illuminant changes, the image gamut is computed first. All the maps from a single colour in the image gamut to each colour in the canonical gamut form a convex set. The intersection of the convex sets obtained for each vertex in the image gamut results in a convex set of feasible maps. This feasible set, which is given in the map space, $\alpha\beta\gamma$-space, normally contains a wide range of assorted maps unless the gamut of the image is large enough to reduce the possible bindings of the image gamut inside the canonical gamut. A selection step is needed to choose the optimal map inside the feasible set, i.e. the best approximation to the unknown illuminant. Different heuristics have been used to obtain a single answer. The most successful heuristic [14] is the selection of the map that maximises the volume of the mapped image gamut, i.e. the map that makes the image gamut as colorful as possible within the bounds of the canonical gamut, also known as CRULE. Other heuristics like the average map of the feasible set have also been studied. Several methods have derived from Forsyth first approach, [9, 15].

Another kind of methods are those based on Colour by Correlation which propose to study the chromaticities of an image to decide among a set of proposed illuminants the one that is more compatible with the chromaticities found [16]. A correlation matrix is pre-computed and describes for each of the selected illuminants the occurrence of image chromaticities. Each row in the matrix corresponds to a different training illuminant and matrix columns to possible chromaticity ranges.

An interesting study comparing the preformance of these different methods described can be found in [14]. There are more contributions which are important in colour constancy but they do not adapt to the context we work in, as they deal with the recovery of surface spectral reflectances using reduced sets of linear bases [6].

3 Surface Matching

The method we propose in this paper tries to introduce the surface matching phenomenon, previously studied as one of the cues of how the human visual system performs colour constancy [4, 17], to reduce the number of possible map solutions. Nevertheless the idea has not yet been explored when performing computational colour constancy. In the process of guessing the illuminant of an image, it is likely to match the colours that we find in the image with colours that we have previously learned, which are a set of colours we already know for its significance. It can be easily assumed that when looking at an image a part of the colour constancy process is the matching of the colours that we see in the image with colours that we 'expect' to find in the image. This refers to a previously learned knowledge of common colours as seen under an ideal, canonical, illuminant.

Considering this idea, we can pair the colours that are present in our image with 'reference' colours. The values of these colours as they would be seen under the canonical illuminant can be computed and they can be named as canonical colours or 'canonical surfaces'. Therefore, we can match every surface in our image with a 'canonical surface'. This is the surface matching process, also known as 'asymmetric colour matching' and depicted in [4]. To perform the 'surface matching' process, we need the set of surfaces to match with. In our surface matching approach, we propose to use a reduced set of 'canonical surfaces', carefully selected to represent the most important and frequent colours. The selection of these canonical surfaces is a hard goal that should be addressed.

4 Relaxed Grey-World

Surface matching implies to match every image surface with every canonical surface, that is to generate all the possible combinations of matchings. Even using a reduced and significant set of image surfaces and a small set of canonical surfaces the set of pairs of matches that can be derived is too large and introduces lots of non-consistent pairs of matchings (if a reddish image surface is matched with a bluish canonical surface, it is not coherent to match another bluish image surface with a reddish canonical surface). This leads us to introduce an assumption to constrict the set of matchings, in order to build a consistent set losing minimum performance.

The Grey-World assumption, as depicted before, supposes the average of an image is grey. Even though this is a strong assumption it can help us to find the consistent constriction that maintains the colour structure of the image gamut. In order to relax this assumption we propose another one:

Relaxed Grey-World Assumption. *The image gamut under the canonical illuminant contains grey or its average is close to grey.*

Considering this assumption the set of canonical surfaces that can be paired with each image surface can be reduced to the canonical surfaces which are close to the image surfaces when the grey-world map is applied to the image, figure 1. That is, the grey-world assumption is relaxed in order to find the solutions near the grey-world, enabling some sort of flexibility near this solution.

The relaxed grey world asumption combined with surface matching lead us to the new approach we propose in this paper. The method matches the image surfaces with canonical surfaces that we have previously selected, but only with the surfaces that are consistent with the relaxed grey-world assumption, i.e. the canonical colours near a neighbourhood in the grey world transform.

First of all we need to select a representative set of surfaces and compute their RGB values for the canonical illuminant, which is selected to be well balanced with our sensor. Hence we have a set of k canonical surfaces, denoted as $S^C = \{S_1^C, S_2^C, \ldots, S_k^C\}$.

Thus, for a given image, I, acquired under an unknown illuminant U, the matching algorithm is carried out with the following steps:

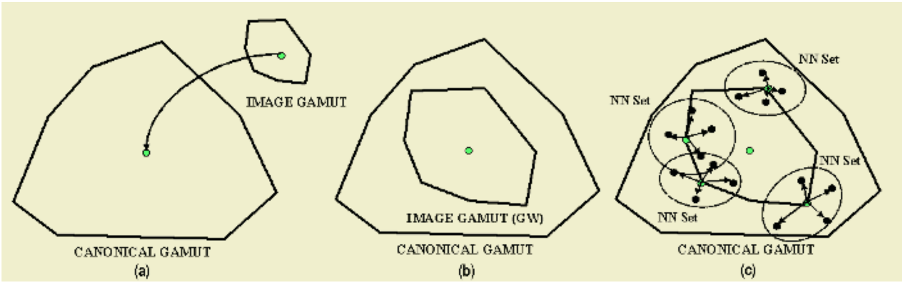

Fig. 1. The relaxed grey-world assumption leads us to find a set of nearest-neighbour canonical surfaces for each image surface. The image is maped to the center of the canonical gamut (a),(b) and there the nearest-neighbour canonical surfaces for each image surface are selected (c).

1. Getting RGB values of surfaces from the image I, denoted as $S^U(I) = \{S_1^U, S_2^U, \ldots, S_n^U\}$, where n is the number of surfaces.
2. Applying the grey world transform to $S^U(I)$, which places the center of the image gamut in the center of the canonical gamut (fig. 1 a,b). It is denoted as $S^{GW}(I)$.
3. For each surface, $i = 1 \ldots n$, of $S^{GW}(I)$ we select the m nearest neighbours surfaces from the canonical surfaces (fig.1 c), S^C, we denote this set as S_i^{NN}.
4. Computing the set of all possible correspondences between each S_i^U with all the surfaces in S_i^{NN}, we name this set $RCorr = \{S_1^U = S_{1,p_1}^{NN}, S_2^U = S_{2,p_2}^{NN}, \ldots, S_n^U = S_{n,p_n}^{NN}; \forall p_i = 1, \ldots, m\}$, where $\#RCorr = m^n$.
5. For each element of $RCorr$, the corresponding $\alpha\beta\gamma$ map is computed, and we obtain a set of maps, $MAP_{\alpha\beta\gamma}^{RCorr}$.
6. All the maps in $MAP_{\alpha\beta\gamma}^{RCorr}$ out of the feasible set are removed, as we do not want to deal with impossible maps.

Once we have generated the set of maps, $MAP_{\alpha\beta\gamma}^{RCorr}$, we propose to use one of the existing heuristics to select one map within this set. In the following section we show the results using the heuristics of maximum gamut volume and average of the set. A simplification of the process can be seen in figure 2.

5 Experiments and Results

To evaluate our method in this first approach we have looked at its performance using only synthetic data. This is a first way to evaluate methods because performance is not affected by image noise and we are able to evaluate performance over hundreds of synthetic images and thus obtain a reliable performance statistic. Otherwise, with real data these problems arise, and also the available datasets are not large enough to extensively test the method.

To build the RGB of the canonical surfaces, we have chosen a synthetic planckian illuminant with CCT=6500K (fig. 3 (a)). A gausian narrow-band sensor has been built, with centers in 450, 540 and 610 nm (fig. 3 (b)). Hence, the

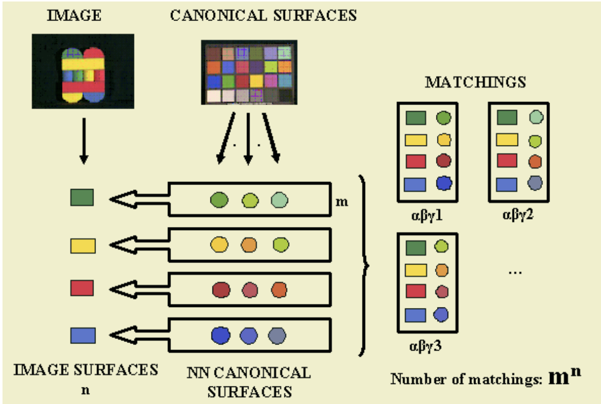

Fig. 2. An illustration of how the relaxed grey world algorithm proceeds.

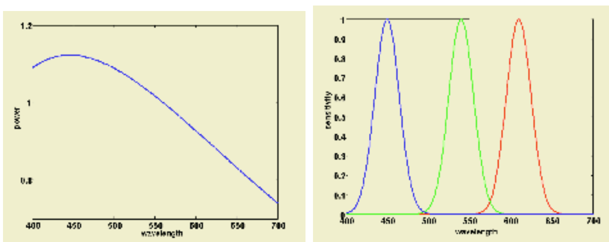

(a) (b)

Fig. 3. The synthetic illuminant (a) and sensor (b) used in the experiments.

1995 reflectances of the Munsell chips have been used to synthesise the RGB values of our canonical set of surfaces.

Once we have selected the canonical surfaces we generate synthetic images to test the algorithm. 400 images consisting of 10 reflectances per image (from Munsell chips randomly selected) under a random illuminant, chosen from a frequently used selection of 11 different illuminants [14]. To test the method, we have selected 6 surfaces from each image and found their 5 nearest neighbours surfaces from the canonical surfaces, that is $n = 6$ and $m = 5$.

We have used as recovery error the angular error between the RGB of the estimated illuminant, $\widehat{RGB}_w^C$, and the RGB of the canonical illuminant used, RGB_w^C (as it is done in [14]). These RGB values of the illuminants are normally unknown in real images, but they can be computed easily working with synthetic data.

$$\text{recovery error} = angle(\widehat{RGB}_w^C, RGB_w^C)$$

In table 1 we can see the performance of the proposed method versus one of the most significant colour constancy algorithms that normally achieves best results [14], CRULE (introduced by Forsyth in [5]). The performance varies

Table 1. Comparison of the performance of the two methods. The value shown is the root mean square of the angular errors computed for the 400 synthetic images.

Heuristic	CRULE	Relaxed Grey-World
Maximum Volume map	7.09°	7.55°
Average map	9.35°	6.62°

depending on the heuristic used to select the optimal map within the computed maps. As it can be seen, the best performance is obtained taking the average map of the proposed Relaxed Grey World. This improvement reinforces the use of the relaxed grey-world assumption. Also, in figure 4 the different sets of maps generated with the two algorithms can be compared. With our method, we avoid to generate a large set of maps that includes the worse maps. We look for a reduced set of maps which includes the best solutions. In this sense we have computed the average value of the best angular error for each of the 400 images and it has resulted to be 1.9°, which means that an optimal map is included in our set of maps in the most of the cases. This result combined with the performance of our method using the average as heuristic justifies the use of the reduced set of maps.

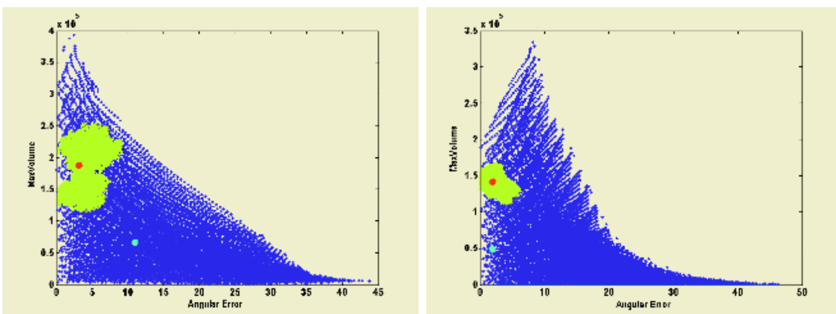

Fig. 4. Comparison of the sets of maps generated with CRULE (dark dots) versus the set of maps generated with our method (bright dots) for 2 different images. In the x-axis is represented the angular error and in the y-axis the maximum volume heuristic.

6 Discussion

As it has been proven, the introduction of the surface matching approach to solve computational colour constancy opens a new line of research in this problem that can help in reducing the error of current methods, that ignore image information that can be introduced by surface matching. The method proposed performs good in the synthetic world and this encourages us to go on with its improvement. The selection of canonical surfaces is an important step to pay more attention and to be focus of a deep study. Indeed, the number of canonical surfaces used in our experiments may seem too large to depict representative colours, but it has been used as a first approach to the surface matching method, to test how good it could perform. Further work needs to be done in the selection of the set of canonical surfaces, as they should represent more trustworthily our knowledge of colours. When done, this part of the process of colour constancy in the human visual system will be enabled to take part in computational approaches.

Acknowledgments

This work has been partially supported by project TIN 2004-02970 (Ministerio de Educación y Ciencia).

References

1. Land, E.H., McCann, J.J.: Lightness and retinex theory. Journal of the Optical Society of America **61** (1971) 1–11
2. Spitzer, H., Semo, S.: Color constancy: a biological model and its application for still and video images. Pattern Recognition **35** (2002) 1645–1659
3. Wandell, B.A.: Foundations of Vision. Sinauer Associates, Inc. (1995)
4. Foster, D.H.: Does colour constancy exist? TRENDS in Cognitive Sciences **7** (2003) 439–443
5. Forsyth, D.A.: A novel algorithm for color constancy. Int. J. Comput. Vision **5** (1990) 5–36
6. Maloney, L.T., Wandell, B.A.: Color constancy: a method for recovering surface spectral reflectance. Journal of the Optical Society of America **3** (1986) 29–33
7. Sapiro, G.: Color and illuminant voting. Transactions on Pattern Analysis and Machine Intelligence **21** (1999) 1210–1215
8. Funt, B., Ciurea, F., McCann, J.: Retinex in matlab. In: Proceedings of the IST/SID Eighth Color Imaging Conference: Color Science, Systems and Applications. (2000) 112–121
9. Finlayson, G.: Color in perspective. Transactions on Pattern Analysis and Machine Intelligence **18** (1996) 1034–1038
10. Funt, B., Barnard, K., Martin, L.: Is colour constancy good enough? In: 5th European Conference on Computer Vision. (1998) 445–459
11. von Kries, J.: Beitrag zur physiologie der gesichtempfindung. Arch. Anat. Physiol. **2** (1878) 505–524
12. Finlayson, G., Drew, M., Funt, B.: Diagonal transforms suffice for color constancy. In: ICCV93. (1993) 164–171
13. Land, E.H.: The retinex theory of color vision. Scientific American (1977) 108–128
14. Barnard, K., Cardei, V., Funt, B.: A comparison of computational colour constancy algorithms. part one. methodology and experiments with synthesized data. IEEE Transactions on Image Processing **11** (2002) 972–984
15. Finlayson, G.D.: Color constancy in diagonal chromaticity space. In: Proceedings of the Fifth International Conference on Computer Vision. (1995) 218–223
16. Finlayson, G.D., Hordley, S., Hubel, P.M.: Colour by correlation: A simple, unifying framework for colour constancy. Transactions on Pattern Analysis and Machine Intelligence **23** (2001) 1209–1221
17. Troost, J.M., Weert, C.M.D.: Naming versus matching in color constancy. Perception and Psychophysics **50** (1991) 591–602