

A New Color Constancy Algorithm Based on the Histogram of Feasible Mappings*

Jaume Vergés–Llahí and Alberto Sanfeliu

Institut de Robòtica i Informàtica Industrial,
Technological Park of Barcelona, U Building
{jverges, asanfeliu}@iri.upc.edu

Abstract. Color is an important cue both in machine vision and image processing applications, despite its dependence upon illumination changes. We propose a color constancy algorithm that estimates both the set and the likelihood of feasible color mappings in respect to their frequency and effectiveness. The best among this set is selected to rendered back image colors as seen under a canonical light. Experiments were done to evaluate its performance compared to Finlayson’s 2D gamut–mapping algorithm, outperforming it. Our approach is a helpful alternative whenever illumination is poorly known since it employs only image data.

Keywords: Color, color mappings, color constancy, color histograms.

1 Introduction

In a number of applications from machine vision to multimedia or even mobile robotics, it is important that the recorded colors remain constant under changes in the scene illumination. Hence, a preliminary step when using color must be to remove the distracting effects of such a variation. This problem is usually referred to as *color constancy*, i.e., the stability of surface color appearance under varying illumination. Part of the difficulty is its entanglement with other confounding phenomena such as object shape, viewing and illumination geometry, besides changes in the nature of light and the reflectance properties of surfaces.

Generally, color constancy (CC) is understood as the recovery of *descriptors* for the surfaces in a scene as they would be seen under a canonical illuminant. This is similar to recover an estimate of the illumination color from an image under an unknown light, since it is relatively straightforward to map image colors back to illuminant independent descriptors [1]. Therefore, *finding a mapping between colors or the color of the scene illuminant are equivalent problems*.

Some algorithms followed such a path, specially those related to the *gamut–mapping* approach [2,3,4,5]. Recently, the trend has slightly changed to make a guess on the illumination, as in *color–by–correlation* [1] or *color–voting* [6], rather than attempting to recover only one estimate. A measure of the *likelihood* that each of a set of feasible candidates was the scene illuminant is set out instead, which is afterwards used to render the image back into the canonical illuminant.

* Sup. by the Spanish Min. of Edu. & Sci. under proj. TIC2003-09291 & DPI2004-5414.

Some difficulties arise among the approaches above. The main one for us is the dependence on the fact that *all* possible colors under a canonic illuminant be available *a priori*. Besides, while gamut–mapping employs a selection strategy blind to any information about the likelihood of a particular mapping, color–by–correlation can only chose an illumination out of a discrete set of lights known *a priori*. Real–world applications usually do not fit such requirements.

Hence, in this paper we propose a new color constancy algorithm that can be used in tasks with little knowledge on the scene illumination. Our algorithm computes the histogram of the feasible maps and recovers one to changes image colors as if seen under a canonic light only using a canonic image picturing a similar scene. The performance of this algorithm is compared to that of the Finlayson’s 2D gamut–mapping, outperforming it and obtaining similar figures to those of the color–by–correlation approach.

2 The Color Constancy Algorithm

The goal of this algorithm is to recover an image as seen under some canonic illumination. This algorithm gives to each feasible color mapping a particular likelihood related to its *frequency* and *performance*. The frequency is estimated from the histogram of feasible mappings, while the performance evaluates how close colors are rendered to those in the canonic set. The combined measure is used afterwards to select the best color mapping. Hereafter, we describe the basis of our algorithm along with its elements, namely, likelihood function, color coordinates, color change model, and mapping estimation and selection.

Let I_c and I_a be the *canonic* and the *actual* images, respectively, picturing similar scenes under two different lights. The algorithm only employs image raw data and no segmentation is needed. Our aim is to find the likeliest color transformation T mapping colors of image I_a as close to those of I_c as possible. We note as $\mathcal{I} \subset \mathbb{R}^d$ a set of colors, where d is the color space dimension. Its origin can be either a specified color gamut or an image. The color histogram of \mathcal{I} is noted as $\mathcal{H}_{\mathcal{I}}$. If a mapping $T \in \langle T \rangle$ is applied to each color in \mathcal{I} , a transformed set $T(\mathcal{I})$ is obtained. $\langle T \rangle$ is the set of *feasible color mappings*. Analogously, $T(\mathcal{H}_{\mathcal{I}})$ represents the transformation of the histogram $\mathcal{H}_{\mathcal{I}}$ by the map T .

In general, given two color sets, \mathcal{I}_a and \mathcal{I}_c , a *model of color change* consists in a mapping $T \in \langle T \rangle$ so that

$$\begin{aligned} T: \mathcal{I}_a &\longrightarrow \mathcal{I}_c \\ s &\longmapsto T(s) = q \end{aligned} \tag{1}$$

where $s \in \mathcal{S}$ and $q \in \mathcal{Q}$ are *corresponding* colors. The set of feasible maps is $\langle T \rangle = \{T = \Delta(\mathcal{S}, \mathcal{Q}) \mid \forall \mathcal{S} \subset \mathcal{I}_a \text{ and } \forall \mathcal{Q} \subset \mathcal{I}_c\}$ and Δ is a *mapping estimation scheme* computing a map T out of two corresponding sets \mathcal{S} and \mathcal{Q} .

2.1 Likelihood Function

The color constancy algorithm must select the likeliest transformation \hat{T} from the set $\langle T \rangle$ accordingly to a likelihood function $\mathcal{L}_{\Delta}(T \mid \mathcal{I}_a, \mathcal{I}_c)$ as

$$\hat{T} = \underset{T \in \langle T \rangle}{\operatorname{argmax}} \{ \mathcal{L}_\Delta(T | \mathcal{I}_a, \mathcal{I}_c) \} \quad (2)$$

Any likelihood \mathcal{L}_Δ is related to a probability function \Pr as $\mathcal{L}_\Delta(T | \mathcal{I}_a, \mathcal{I}_c) = \log(\Pr(T | \mathcal{I}_a, \mathcal{I}_c))$, noting that the map maximizing $\mathcal{L}_\Delta(T | \mathcal{I}_a, \mathcal{I}_c)$ also maximizes $\Pr(T | \mathcal{I}_a, \mathcal{I}_c)$, and *vice versa*. Hence, we must first get a value for the probability of any mapping T . As an estimate of $\Pr(T | \mathcal{I}_a, \mathcal{I}_c)$ we use the *histogram* of the set of feasible mappings, noted as $\mathcal{H}_{\langle T \rangle}$. In order to compute frequencies we employ colors in the sets \mathcal{I}_a and \mathcal{I}_c , as well as the mapping estimator Δ . *The point is that the likelier a mapping is, the more frequent it should be in the histogram $\mathcal{H}_{\langle T \rangle}$, and the other way round.*

Nonetheless, a particular mapping T can be produced from different groups of colors. Thus, it is useful to define the set of all pairs $(\mathcal{S}, \mathcal{Q})$ giving rise to a certain map T from the estimator Δ , i.e., $\Delta^{-1}(T) = \{(\mathcal{S}, \mathcal{Q}) \in 2^{\mathcal{I}_a} \times 2^{\mathcal{I}_c} \mid \Delta(\mathcal{S}, \mathcal{Q}) = T\}$, where $2^{\mathcal{I}}$ stands for the set of all subsets of \mathcal{I} . Therefore, the set $\Delta^{-1}(T)$ is equivalent to T and can be taken instead of it since $\Delta^{-1}(T) = \Delta^{-1}(T') \Leftrightarrow T = T'$. Hence, it is true that $\Pr(T | \mathcal{I}_a, \mathcal{I}_c) = \Pr(\Delta^{-1}(T) | \mathcal{I}_a, \mathcal{I}_c)$.

In practice, $\Delta^{-1}(T)$ can be thought as a finite disjoint union of singletons $\{(\mathcal{S}, \mathcal{Q})\}$, being each set a combination of colors from sets \mathcal{I}_a and \mathcal{I}_c , respectively. Besides, to compute $\Pr((\mathcal{S}, \mathcal{Q}) | \mathcal{I}_a, \mathcal{I}_c)$ we must only remind that each $\{(\mathcal{S}, \mathcal{Q})\}$ is formed by two independent sets, namely, $\mathcal{S} \subset \mathcal{I}_a$ and $\mathcal{Q} \subset \mathcal{I}_c$. Therefore,

$$\Pr((\mathcal{S}, \mathcal{Q}) | \mathcal{I}_a, \mathcal{I}_c) = \Pr(\mathcal{S} | \mathcal{I}_a) \cdot \Pr(\mathcal{Q} | \mathcal{I}_c) \quad (3)$$

Then, we get that

$$\Pr(T | \mathcal{I}_a, \mathcal{I}_c) = \sum_{\forall (\mathcal{S}, \mathcal{Q}) \in \Delta^{-1}(T)} \Pr(\mathcal{S} | \mathcal{I}_a) \cdot \Pr(\mathcal{Q} | \mathcal{I}_c) \quad (4)$$

that is, frequency of a mapping T in the histogram of feasible mappings $\mathcal{H}_{\langle T \rangle}$ is computed adding the product of frequencies corresponding to all the pairs \mathcal{S} and \mathcal{Q} giving rise to the mapping T by means of the mapping estimator Δ .

$\Pr(\mathcal{S} | \mathcal{I}_a)$ and $\Pr(\mathcal{Q} | \mathcal{I}_c)$ are approximated by the frequency of the corresponding bins in histograms $\mathcal{H}_{\mathcal{I}_a}$ and $\mathcal{H}_{\mathcal{I}_c}$. In case $\mathcal{S} = \{s\}$, it is straightforward that $\Pr(\mathcal{S} | \mathcal{I}) \approx \mathcal{H}_{\mathcal{I}}(s)$. Hence, in general, if $\mathcal{S} = \{s_i\}_{i=1, \dots, n}$, the probability of the set under the hypothesis of independence of colors is

$$\Pr(\mathcal{S} | \mathcal{I}) \approx \prod_{i=1}^n \mathcal{H}_{\mathcal{I}}(s_i) \quad (5)$$

On the other hand, usually some spurious peaks appear in $\mathcal{H}_{\langle T \rangle}$, which might mislead the algorithm. To improve the robustness of Eq. (4), a measure of similarity between the transformed set $T(\mathcal{I}_a)$ and the canonic set \mathcal{I}_c is taken into account, evaluating the efficiency of a particular mapping T . We use the Swain & Ballard intersection-measure between histograms defined in [7] as

$$\cap(\mathcal{H}_1, \mathcal{H}_2) = \sum_k \min\{\mathcal{H}_1(k), \mathcal{H}_2(k)\} \in [0, 1] \quad (6)$$

The advantages of such a measure are that it is very fast to compute if compared to other matching functions. Besides, if the histograms are sparse and colors equally probable, it is a robust way of comparing images [7]. This step helps in practice to eliminate outlier mappings among candidates.

As a consequence, we finally define our likelihood function by joining the probability of a mapping and its performance in a single expression as follows

$$\mathcal{L}_\Delta(\mathbf{T} \mid \mathcal{I}_a, \mathcal{I}_c) = \log(\cap(\mathbf{T}(\mathcal{H}_{\mathcal{I}_a}), \mathcal{H}_{\mathcal{I}_c})) \cdot \Pr(\mathbf{T} \mid \mathcal{I}_a, \mathcal{I}_c), \quad \forall \mathbf{T} \in \langle \mathbf{T} \rangle \quad (7)$$

where $\mathbf{T}(\mathcal{H}_{\mathcal{I}_a})$ is the transformation of $\mathcal{H}_{\mathcal{I}_a}$ by a mapping \mathbf{T} and $\Pr(\mathbf{T} \mid \mathcal{I}_a, \mathcal{I}_c)$ is the frequency of \mathbf{T} in the histogram $\mathcal{H}_{\langle \mathbf{T} \rangle}$ computed with Eq. (4).

2.2 Color Coordinates

Colors can be represented as vectors in \mathbb{R}^d . In our case, to alleviate problems with specularities or shape and to reduce at the same time the computational burden, we use the *perspective color coordinates* $(r, g) = (R/B, G/B)$ proposed by Finlayson in [3]. Finlayson and Hordley also proved in [8] for these coordinates that the set of feasible mappings computed in a 3D space and projected into a 2D space afterward is the same as the one directly computed in a 2D space.

2.3 Color Change Model

It is been stated that any color change could be described using a homogeneous linear relation [2]. Thus, the mapping \mathbf{T} is specified in coordinates as follows

$$\begin{aligned} \mathbf{T}: \mathcal{I}_a &\longrightarrow \mathcal{I}_c \\ \mathbf{s} &\longmapsto \mathbf{T}(\mathbf{s}) = \mathbf{T} \mathbf{s}^t = \mathbf{q}^t \end{aligned} \quad (8)$$

A reasonable tradeoff between simplicity and performance is attained employing a *diagonal model* [2,1,3]. This model assumes color sensors are completely uncorrelated and any change in the light falling into them equates to independently scaling each channel, that is, $\mathbf{T} = \text{diag}(t_1, \dots, t_d)$. Equivalently, \mathbf{T} can be expressed as a vector $\mathbf{t} = (t_1, \dots, t_d) \in \mathbb{R}^d$, so that, $\mathbf{T} = \mathbf{t} \mathbf{I}_d$. This is the model used in this paper, despite the color constancy algorithm also works with any other more complete model, involving more computational time as a result.

2.4 Mapping Estimation

Formally, the mapping estimator Δ is a function computing a mapping $\mathbf{T} \in \langle \mathbf{T} \rangle$ from two sets $\mathcal{S} = \{s_i\}_{i=1, \dots, n} \subset \mathcal{I}_a$ and $\mathcal{Q} = \{q_i\}_{i=1, \dots, n} \subset \mathcal{I}_c$, i.e.,

$$\begin{aligned} \Delta: 2^{\mathcal{I}_a} \times 2^{\mathcal{I}_c} &\longrightarrow \langle \mathbf{T} \rangle \\ (\mathcal{S}, \mathcal{Q}) &\longmapsto \Delta(\mathcal{S}, \mathcal{Q}) = \mathbf{T} \end{aligned} \quad (9)$$

A mapping \mathbf{T} can be expressed as a matrix \mathbf{T} so that $q = \mathbf{T}(s)$ is equivalent to $\mathbf{T} \mathbf{s}^t = \mathbf{q}^t$. In general, it is not possible to find a matrix \mathbf{T} out of just a

pair of corresponding vectors \mathbf{s} and \mathbf{q} , since we need, at least, as many linearly independent pairs as the space dimension d . Taking advantage of the nature of \mathbf{T} and two sets of $n \geq d$ one-to-one corresponding vectors, namely, $\{\mathbf{s}_i\}_{i=1,\dots,n}$ and $\{\mathbf{q}_i\}_{i=1,\dots,n}$, it is true that $\mathbf{T}(\mathbf{s}_1^t | \dots | \mathbf{s}_n^t) = (\mathbf{q}_1^t | \dots | \mathbf{q}_n^t)$.

If $\mathbf{S} = (\mathbf{s}_1^t | \dots | \mathbf{s}_n^t) \in \mathcal{M}_{d \times n}(\mathbb{R})$ and $\mathbf{Q} = (\mathbf{q}_1^t | \dots | \mathbf{q}_n^t) \in \mathcal{M}_{d \times n}(\mathbb{R})$ are introduced by joining vectors columnwise, finding a matrix \mathbf{T} consists in solving the linear system $\mathbf{Q} = \mathbf{T}\mathbf{S}$, where \mathbf{T} is the unknown. The usual method to solve this system is by a SVD to compute pseudo-inverse \mathbf{S}^+ . Therefore, the version of Eq. (9) in coordinates is

$$\begin{aligned} \Delta: \mathcal{M}_{d \times n}(\mathbb{R}) \times \mathcal{M}_{d \times n}(\mathbb{R}) &\longrightarrow \mathcal{M}_d(\mathbb{R}) \\ (\mathbf{S}, \mathbf{Q}) &\longmapsto \Delta(\mathbf{S}, \mathbf{Q}) = \mathbf{Q}\mathbf{S}^+ = \mathbf{T} \end{aligned} \quad (10)$$

In case the diagonal model is used, the above function is greatly simplified

$$\begin{aligned} \Delta: \mathbb{R}^d \times \mathbb{R}^d &\longrightarrow \mathbb{R}^d \\ (\mathbf{s}, \mathbf{q}) &\longmapsto \Delta(\mathbf{s}, \mathbf{q}) = \left(\frac{q_1}{s_1}, \dots, \frac{q_d}{s_d}\right)^t = \mathbf{t} \end{aligned} \quad (11)$$

The method for carrying out the mapping computations between two color sets consists in first computing the histograms of both sets and applying the estimator Δ to every possible pair of color sets. Colors in \mathcal{S} need not to be really in a correspondence with those in \mathcal{Q} , since *all* possible combinations are checked. Nevertheless, in a real case some heuristic are necessary to speed up the process. In our case, only the colors with nonzero frequency are considered, which reduces the amount of computations to less than $O(p^2)$, where $p \ll n$ is the number of histogram bins and n the number of colors in a set.

2.5 Mapping Selection

In addition to Eq. (2), other heuristics can be useful to improve the results

Max: As said in Eq. (2), the mapping with the highest likelihood is selected.

Mean: Mean among the candidates, after weighting them by their likelihood.

CMax: The same as in **Max**, after constraining the set of feasible illuminants, as proposed by Finlayson.

CMean: The same as in **Mean**, after constraining the illuminant set.

CM: The center of mass of the convex hull from the set of feasible mappings.

3 Experiments and Results

We show now the performance of our color constancy algorithm compared to that of Finlayson's 2D gamut mapping. To that goal, the database¹ considered consisting in a set of 220 images from 20 objects viewed under 11 illuminants.

¹ http://www.cs.sfu.ca/~color/data/objects_under_different_lights/index.html

Basically, these are images (Fig. 1) of single colorful objects in a black background to avoid interreflections. The pose of the objects varied every time lights were changed. Additionally, authors took care of analyzing the lights used to produce the images, quite helpful in building the set of feasible illuminants.

A canonic image of the scene is employed as the reference for colors. Experiments carried out basically consists in computing a distance between colors from each scene and those of the canonic one before and after the CC stage to measure the amount of reduction in color dissimilarity. This is done for the two algorithms and one canonic image per set. The more the processed images resemble to the canonic, the greater the color dissimilarity reduction. Each algorithm has various selection heuristics. Our algorithm used *CM*, *Mean*, *Max*, *CMean*, and *CMax*, as explained in Section 2.5. In Finlayson’s 2D-GM we employed the center of mass (*FinCM*), the mean (*FinMean*), and the 3D mean proposed by Hordley and Finlayson in [8] (*FinHord*), all constrained by the set of lights.

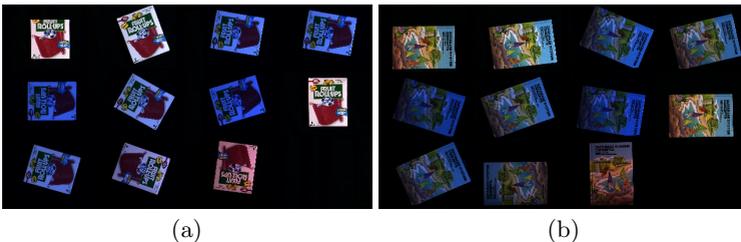


Fig. 1. Sets of some objects under all light variation: (a) rollups, (b) book

Two measures of color dissimilarity were applied to estimate the performance of the two CC methods, namely, the Swain & Ballard difference [7] and a RMS error. The first accounts for the difference between the canonic and the processed color histograms. Here, an extra result (*True*) is obtained using the maps from the true lights in the database. Instead, RMSE compares two images obtained from each actual image after transforming it with the *true* mapping and the one from the particular CC algorithm. The canonic image is useless here because RMSE is pixelwise and objects were moved from one image to another.

The results are summarized in Fig. 2 through Fig. 5. Plots in Fig. 2 and Fig. 3 show the results corresponding to the S&B distance, while Fig. 4 and Fig. 5 illustrate those using the RMSE distance. These graphics consist in three plots each: mean, median, and a distribution boxplot², where colors represents versions of the algorithms. Additionally, the numerical values for the mean distances and the percent reduction of the color dissimilarity are accounted in Table 1 (Finlayson’s 2D-GM) and Table 2 (our CC algorithm), where *No CC* correspond to the distance between images before any CC is carried out.

First of all, when considering the performance of our algorithm, we state the results from the Finlayson’s 2D-GM – 0.27 (RMSE) – are comparable to those in

² A boxplot is a statistic descriptive tool showing, at once, a distribution of values.

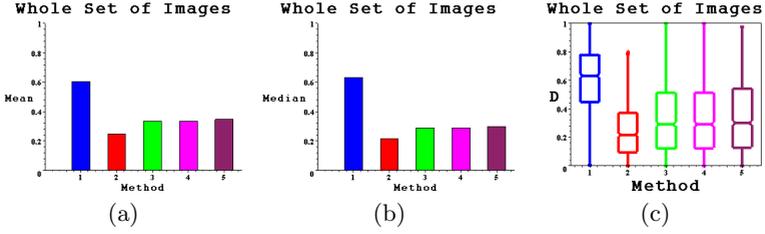


Fig. 2. Results using S&B and Finlayson's 2D-GM. Blue: *No CC*. Red: *True*. Green: *FinCM*. Magenta: *FinMean*. Violet: *FinHord*. (a) Mean, (b) median, and (c) boxplot.

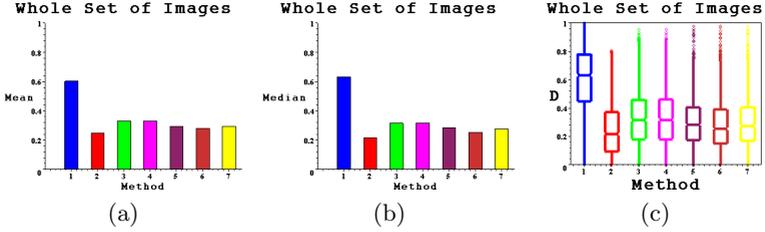


Fig. 3. Results using S&B and our algorithm. Blue: *No CC*. Red: *True*. Green: *CM*. Magenta: *Mean*. Violet: *Max*. Orange: *CMean*. Yellow: *CMax*. (a) Mean, (b) median, and (c) boxplot.

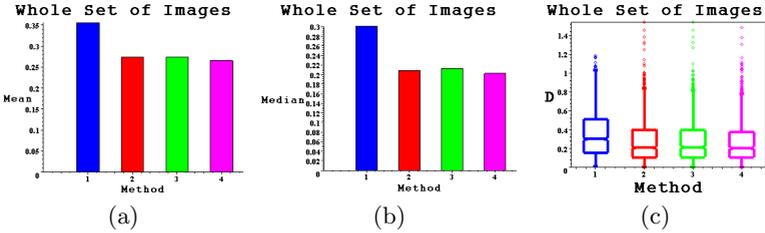


Fig. 4. Results using RMSE and Finlayson's 2D-GM. Blue: *No CC*. Red: *FinCM*. Green: *FinMean*. Magenta: *FinHord*. (a) Mean, (b) median, and (c) boxplot.

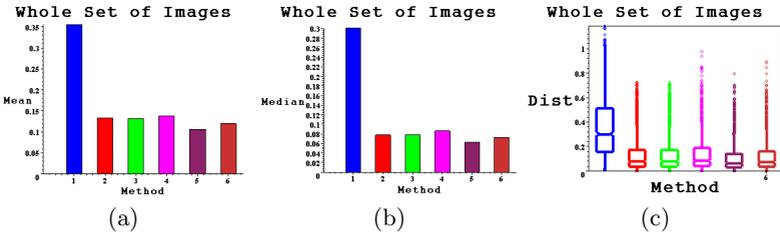


Fig. 5. Results using RMSE and our algorithm. Blue: *No CC*. Red: *CM*. Green: *Mean*. Magenta: *Max*. Violet: *CMean*. Orange: *CMax*. (a) Mean, (b) median, and (c) boxplot.

Table 1. Results using Finlayson’s 2D–GM algorithm

Finlayson’s 2D–GM		FinCM	FinMean	FinHord	No CC	True
S&B	Mean Dist.	0.3367	0.3357	0.3464	0.6036	0.2488
	% Red.	44.22	44.38	42.61	~	58.78
RMSE	Mean Dist.	0.2732	0.2738	0.2655	0.3559	
	% Red.	23.24	23.07	25.40	~	

Table 2. Results using our color constancy algorithm

Our CC alg.		CM	Mean	Max	CMean	CMax	No CC	True
S&B	Mean Dist.	0.3301	0.3301	0.2956	0.2805	0.2945	0.6036	0.2488
	% Red.	45.31	45.31	51.03	53.53	51.21	~	58.78
RMSE	Mean Dist.	0.1330	0.1313	0.1382	0.1060	0.1200	0.3559	
	% Red.	62.63	63.11	61.17	70.32	66.28	~	

[1] – 0.21 (RMSE) –. In Table 1 and 2 it is shown that with the S&B distance the Finlayson’s 2D–GM best result is achieved with *FinMean* (44.38%), whereas for our algorithm the best result is obtained with *CMean* (53.53%), which is pretty close to the figure from the true lights mappings (58.78%). In respect to the RMSE results, figures obtained for our algorithm (70.32%) are definitely better than those of Finlayson’s 2D–GM (25.40%) and totally comparable to the best CC algorithm so far, namely, color–by–correlation [1] – 0.11 (RMSE) –, which is fairly a good result since our approach works with less information.

4 Conclusions

This paper describes a new color constancy algorithm. This is based on the computation of the histogram of feasible maps between two light conditions and selecting the best mapping on the base of a likelihood measure for each mapping encompassing both its frequency and effectiveness. Experiments were done to evaluate the performance of our algorithm in comparison to that of Finlayson’s 2D–GM, outperforming it and attaining similar results to those of the color–by–correlation scheme. Our algorithm works with only a canonic image as a reference and resulting maps are not restricted to a discrete set of feasible ones. Besides using less information than other methods, it is a useful alternative in tasks where little knowledge on the illumination is at hand.

References

1. Finlayson, G., Hordley, S., Hubel, P.: Colour by correlation: A simple, unifying framework for colour constancy. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **23** (2001) 1209–1221
2. Forsyth, D.: A novel algorithm for color constancy. *Int. Journal of Computer Vision* **5** (1990) 5–36

3. Finlayson, G.: Color in perspective. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **18** (1996) 1034–1038
4. Barnard, K., Cardei, V., Funt, B.: A comparison of computational colour constancy algorithms: Part one: Methodology and experiments with synthesized data. *IEEE Trans. on Image Processing* **11** (2002) 972–983
5. Barnard, K., Martin, L., Coath, A., Funt, B.: A comparison of computational colour constancy algorithms: Part two: Experiments with image data. *IEEE Trans. on Image Processing* **11** (2002) 985–996
6. Sapiro, G.: Color and illuminant voting. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **21** (1999) 1210–1215
7. Swain, M., Ballard, D.: Indexing via color histograms. In: *Proc. Int. Conf. on Computer Vision*. (1990) 390–393
8. Finlayson, G., Hordley, S.: Improving gamut mapping color constancy. *IEEE Trans. on Image Processing* **9** (2000) 1774–1783