Illuminant Estimation from Projections on the Planckian Locus

Baptiste Mazin, Julie Delon, and Yann Gousseau

LTCI, Télécom-ParisTech, CNRS,
46 rue Barault, Paris 75013, France
{baptiste.mazin,julie.delon,yann.gousseau}@telecom-paristech.fr

Abstract. This paper deals with the automatic evaluation of the illuminant from a color photography. While many methods have been developed over the last years, this problem is still open since no method builds on hypotheses that are universal enough to deal with all possible situations. The proposed approach relies on a physical assumption about the possible set of illuminants and on the selection of grey pixels. Namely, a subset of pixels is automatically selected, which is then projected on the Planckian locus. Then, a simple voting procedure yields a robust estimation of the illuminant. As shown by experiments on two classical databases, the method offers state of the art performances among learning-free methods, at a reasonable computational cost.

1 Introduction

Human vision has the ability to perceive colors in a stable way under light sources having very different spectral contents. In contrast, cameras capture colors differently depending on the illuminant. Automatically canceling out this dependence is a difficult problem that is not yet fully solved. Specifically, the problem of estimating the illuminant of a scene has been the subject of a large body of works, as detailed for instance in the recent paper [1]. A classical application of such an estimation is the automatic correction of scenes acquired under non-standard illuminants. In a different context, color is an important feature for many computer vision tasks such as image comparison, object detection or motion estimation, tasks for which robustness to a change in illuminant may be achieved by a preliminary estimation.

The illumination estimation problem, where one tries to separate the intrinsic color of objects and the content of the illuminant, is basically an underconstrained problem. Indeed, a change in measured colors can be caused by a change in object colors or by a change of illuminant. As a consequence, automatically achieving color constancy necessitates some supplementary assumptions.

The first type of assumptions is based on scene composition. The most obvious one is to assume the presence of a white object in the scene [2], the so-called White-Patch assumption. This assumption is generally weakened to the presence of a perfectly reflective surface, yielding the Max-RGB algorithm. Another well known hypothesis is that the average world is grey, resulting in the Grey-World algorithm [3].

An alternative approach is to use physical models to constrain the problem. Forsyth [4] first proposed to use the fact that under a given illuminant, the range of observable colors is limited. This approach was followed by many improvements over the last years [5,6].

While most algorithms described above rely on a Lambertian assumption, other algorithms involve a dichromatic model assuming that measured light is composed of both a Lambertian and a specular components [7,8]. These methods are often complemented by the assumption that feasible illuminant colors are limited [5,7].

Despite offering correct estimations in a large number of situations, none of the previously mentioned hypotheses is universal and satisfied for all scene-illuminant combinations. Moreover, failure situations vary from one hypothesis to the other. For this reason, interesting results have been obtained by combining the joint use of several basic methods and of learning procedures [9].

In this paper, we present a new learning-free method for the automatic estimation of the illuminant. The method relies on a physical *a priori*, building on the Planckian locus, on the possible set of illuminants. It also involves a simple voting procedure from a selected subset of pixels. Experiments on two widely used databases and comparisons with results provided by Gijzenij and al. [1] show that this approach outperforms standards methods without requiring any learning step.

2 Background

In this section, we give some background on the formation of color images, on the distribution of illuminants under a black body assumption and on the use of chromaticity diagrams. We also summarize previous works on the problem of illuminant estimation.

Illuminant estimation is the first and the toughest nut to crack for color constancy. The second aspect of color constancy is image correction, for which various methods can be used [1]. This paper only deals with the first step of the process.

2.1 Color Image Acquisition

Most approaches used to solve the illuminant estimation problem rely on a dichromatic model [10] for image formation, simplified with a Lambertian assumption. In these conditions, for a pixel $\mathbf{x}$, values $\mathbf{p}(\mathbf{x}) = (p_R(\mathbf{x}), p_G(\mathbf{x}), p_B(\mathbf{x}))$ depend on the spectral distribution of the light source $L(\lambda)$, the spectral distribution of surface reflectance $S(\lambda, \mathbf{x})$, and the camera sensitivities functions $\rho(\lambda) = (\rho_R(\lambda), \rho_G(\lambda), \rho_B(\lambda))$:

$$p_c(\mathbf{x}) = \int_{\omega} L(\lambda) S(\lambda, \mathbf{x}) \rho_c(\lambda) d\lambda, \quad (1)$$

where $c = \{R, G, B\}$, and ω is the visible spectrum.

Assuming that camera sensitivity functions are Dirac delta functions associated to the wavelengths λ_R, λ_G and λ_B , Equation (1) can be rewritten as follows,

$$p_c(\mathbf{x}) = L(\lambda_c)S(\lambda_c, \mathbf{x}). \quad (2)$$

The illuminant estimation problem amounts to retrieve $L = (L_R, L_G, L_B) := (L(\lambda_R), L(\lambda_G), L(\lambda_B))$. This would for instance allow to render the scene under a given canonical illuminant L^u . Once the illuminant has been computed, the correction can be made using Von Kries model or more complex models. The main difficulty arises from the fact that measured pixel values $L(\lambda)$ depend on both the surface reflectance function and the spectrum of the light source.

2.2 Distribution of Feasible Illuminants

Grey-World and White-Patch hypotheses [2] (that will be recalled in the next paragraph) are certainly the most used assumptions to achieve illuminant estimation. There are, however, many situations in which these hypotheses are false and yield incorrect estimation. One of their limitation is that no *a priori* information on the chromaticity of the feasible illuminants is used. For instance, an image filled with a saturated green will lead to estimate a saturated green illuminant even though this is very unlikely.

Therefore, additional constraints are often added to limit the set of feasible illuminants. The blackbody radiator hypothesis is one of the most commonly used assumption. This hypothesis is based on the Planck model, which states that the spectrum of light emitted by an idealized physical body heated at a given temperature is given by

$$L(T, \lambda) = c_1 \lambda^{-5} \left[\exp \left(\frac{c_2}{\lambda T} \right) - 1 \right]^{-1}, \quad (3)$$

where T is the temperature in kelvins, λ is the wavelength and $c_1 = 3.74183 \times 10^{16} \text{ Wm}^2$ and $c_2 = 1.4388 \times 10^{-2} \text{ mK}$ are two constants.

Planck's formula is an accurate estimation of the spectrum of most illuminants including daylight [11]. It results that under the blackbody hypothesis, the possible colors of a grey pixel are limited. This can be conveniently visualized in a chromaticity diagram, the definition of which we now recall. A chromaticity diagram is a 2D projection of 3D color information discarding the intensity component. The CIE 1931 xy chromaticity diagram is the best known of such diagrams. It is derived from an RGB colorspace through an XYZ colorspace. This first step can be done by applying a device dependent matrix to RGB triplets. The xy chromaticities are then obtained by normalizing the vector XYZ such that $x + y + z = 1$. The relation between the Planckian locus and common illuminants (obtained from the database described in [12]) is illustrated on Figure 1.

2.3 Previous Work

In the context of illuminant estimation and color constancy, the most simple and widely used approaches are based on the White-Patch and on the Grey-World

assumptions. The White-Patch assumption assumes that a surface with a perfect reflectance is present in the scene. This hypothesis yields the well known Max-RGB algorithm which presumes that the brighter point corresponds to a perfect reflective surface. The Grey-World hypothesis is that the average reflectance of the scene is achromatic. Both assumptions have been combined by Finlayson [13], proposing the framework described by the following equation:

$$L_c = \left(\frac{\int p_c(\mathbf{x})^n d\mathbf{x}}{\int d\mathbf{x}} \right)^{1/n}. \quad (4)$$

When $n = \infty$, this equation yields the Max-RGB algorithm and when $n = 1$, it yields the Grey-World algorithm. A classical variant of this method, called general Grey World [14], consists in performing a preliminary convolution of the image with a Gaussian kernel.

This framework was extended by van de Weijer [15]. This extension includes the Grey-Edge hypothesis that the average of reflectance differences is achromatic, leading to the equation:

$$L_c = \left(\int \left| \frac{\partial^k p_c(\mathbf{x})^n}{\partial \mathbf{x}^k} \right| d\mathbf{x} \right)^{1/n}. \quad (5)$$

All algorithms given by Equation (5) use all pixels to identify the white point. As noticed in [16], a more intuitive approach is to first detect potentially grey surfaces and then to estimate the illuminant from the corresponding pixels. Several algorithms using this procedure have recently been proposed.

Li and al. [16] propose an algorithm which iteratively estimates a white point maximizing the number of grey pixels. This algorithm takes as a starting point the estimation given by Equation 5 and therefore aims at refining this estimation. However, this estimation can be inaccurate and as noticed by the authors this step is especially touchy.

Another way to detect potentially grey pixels is to select points that are close to the Planckian locus in a chromaticity space and then to average them [17]. While inconsistent estimations (such as a purple illuminant) can be avoided this way, the approach may fail in the case where there is two potentially grey surfaces in the scene. Indeed, in this case, the resulting estimation will be made of the midway point between these surfaces. In the spirit of this approach, we define a simple pixel based voting procedure to estimate the illuminant chromaticity.

Many voting procedures permitting to choose between concurrent solutions to the color constancy problem were proposed in the literature. Since the same pixels values can be obtained by different combinations of illuminant and reflectance, Sapiro proposes a voting procedure [18] where each pixel votes for a set of parameters defining the reflectance of the illuminant functions. A Hough transform is then used to select the best parameters set. Riess *et al.* [19] propose another voting procedure. Roughly, many patches are selected and for each of them an illuminant is estimated by using an inverse chromaticity space. All patches for which the estimation succeeds vote and the illuminant is selected from these votes.

Very recently Vasquez-Corral *et al.* [20] propose to refine such voting approaches by incorporating perceptual constraints. In their approach, illuminants are weighted by “their ability to map the corrected image onto specific colors”. These colors being chosen as universal color categories, this technique does not require any learning step.

3 Illuminant Selection via Grey Pixel Identification

Our illuminant selection algorithm draws on the assumption that most light sources are close to the Planckian locus. This hypothesis is well founded, as shown for instance by the experiment illustrated on Figure 1. Relying on this observation, our algorithm starts by finding pixels close to the Planckian locus in a well chosen color space. These pixels can be considered as the “most probably grey” in $\mathbf{p}$. We then find the color temperature of all these pixels by projecting them on the Planckian locus. We build the histogram of this set of temperatures and apply a voting procedure in order to determine the most probable temperature T^0 for our illuminant. Finally, the complete illuminant is estimated by computing the barycenter of all pixels p that voted for the temperature T^0 in the histogram (see Algorithm 1 for a complete description of the procedure).

The color space used in the whole algorithm is the CIE 1960 uv diagram, as recommended by the CIE. Indeed, in this space, isotherms are defined as the lines perpendicular to the Planckian locus [21]: the correlated color temperature (CCT) of a pixel is thus the one of its closest black body radiator. Also, the temperature scale used in the voting procedure must be chosen carefully. Indeed, the Kelvin temperature scale does not yield a regular sampling of the Planckian locus [22]. The correlated color temperature $T(p)$ of a each pixel p is thus converted into the MIREN (Micro Reciprocal Degree) scale $T_{Mired}(p) = 10^6/T(p)$ before computing the histogram. This conversion was found by Priest [22] as a better way to sample color temperature according to human perception. This MIREN scale is illustrated by Figure 1, right. Finally, in order to give a greater importance to brighter pixels in the voting procedure, the contribution of each pixel in the histogram is weighted by a power of its luminance in $\mathbf{p}$. This weighting limits the influence of dark pixels and can be seen as a trade off between the Grey-World and the White-Patch hypotheses [13]. An example of application of the method may be seen on Figure 1.

4 Algorithm Evaluation

In this section, we compare our algorithm to classical and recent state of the art approaches on two different databases, drawing from the results and codes from the recent evaluation paper [1]. These databases contain both indoor and outdoor scenes, taken under a large variety of illuminants.

The first database is the SFU grey-ball database [23], composed of 11,346 images and reprocessed in order to work on linear images, as explained in [1].

Fig. 1. Left: CIE 1931 chromaticity diagram. Right: CIE 1960 uv chromaticity diagram. On both diagrams, the green line represents the Planckian locus, and blue points correspond to the groundtruth illuminants of the Gehler database [12]. On the right, some isothermatures curves are plotted in the CIE 1960 uv chromaticity diagram, with a maximum distance to the Planckian locus equal to 0.02.

input : Image $\mathbf{p}$ in the space sRGB, threshold δ , bin number N , standard deviation σ , power n , canonical illuminant $(u^{\text{ref}}, v^{\text{ref}})$.

output: Illuminant (u^0, v^0) in the CIE uv diagram.

Initialize histogram $H = 0$ on N bins;

for each pixel p **do**

$p_{uv} = \text{Convert}_{sRGB \rightarrow uv}(p)$;

$[p_{\text{Planck}}, T(p)] = \text{Proj}_{\text{Planck}}(p_{uv})$;

// p_{Planck} is the projection of p_{uv} on the Planckian locus in the CIE 1960 uv diagram and $T(p)$ is the temperature (in Kelvins) of the projection;

$d = \text{distance}(p_{uv}, p_{\text{Planck}})$;

if $d < \delta$ **and** $2000K \leq T(p) \leq 20000K$ **then**

$T_{\text{MIREL}}(p) = 10^6 / T(p)$;

$\text{weight}(p) = \text{luminance}(p)^n$;

Add $\text{weight}(p)$ to bin $T_{\text{MIREL}}(p)$ in H ;

end

end

if $H == \emptyset$ **then**

return $(u^0, v^0) = (u^{\text{ref}}, v^{\text{ref}})$;

end

Filter histogram H with a normalized Gaussian of standard deviation σ ;

Find $T^0 = \arg \max(H)$;

Find the barycenter of $H[T^0 - \sigma, T^0 + \sigma]$;

$\mathcal{L} = \text{list of pixels } p \text{ such that } T_{\text{MIREL}}(p) \text{ is in } [T^0 - \sigma, T^0 + \sigma]$;

return $(u^0, v^0) = \text{Barycenter}(\mathcal{L}, \text{weight}(\mathcal{L}))$;

Algorithm 1. Illuminant selection algorithm. In practice, the parameters used in the experimental section are $\delta = 0.01$, $N = 1000$, $\sigma = 2$ and $n = 3$.

Fig. 2. Illustration of method results (for (a) and (b) errors are computed in the sRGB colorspace). (c) shows points close to the Planckian locus and (d) shows the histogram built from their projection on the Planckian locus, weighted by the luminance. On (c) and (d), red points correspond to a CCT < 4000 and blue points to a CCT $\in [4000, 6000]$.

The second is the Gehler database, composed of 568 images. Gehler *et al.* [24] provide both RAW and TIFF versions of each image. TIFF images are generated automatically by the camera and hence contain many non-linear corrections (demosaicing, clipping, gamma correction). In order to avoid the influence of these corrections, we choose, as in [24] to rely only on RAW versions. In our experiments, we provide two different evaluations, each one relying on a different reprocessing of RAW images. In the first one, images are reprocessed in the color space of the camera, as in [1]. In the second one, images are reprocessed in the sRGB colorspace, which is a less arbitrary choice, as explained below. This linear transformation is enabled by the fact that we know the cameras used to create the database¹.

For each database, the experimental protocol is then exactly as described in [1] and consists in computing for each image of the database some similarity measure between the estimated illuminant $\mathbf{I}^e$ and the groundtruth illuminant $\mathbf{I}^g$. Both illuminants define a “grey” axis in the color space of the image. The quality of the estimation is then measured by using the angle between these axes, *i.e.*

¹ The conversion from the camera colorspace to sRGB is specific to the camera and involves the use of the intermediary space XYZ.

$$E_a(\mathbf{I}^e, \mathbf{I}^g) = \cos^{-1} \left(\frac{\mathbf{I}^e \cdot \mathbf{I}^g}{\|\mathbf{I}^e\| \cdot \|\mathbf{I}^g\|} \right). \quad (6)$$

For the images of the SFU grey-ball database, this error is computed directly in the sRGB colorspace. For the Gehler database, this error can also be computed in the camera RGB colorspace, as chosen in [1]. This choice presents two problems. First, the three color components of a camera have very different dynamic ranges, which makes the error computation somehow biased. Most importantly, two different cameras are used to create the database, and averaging errors obtained in two different colorspace seems inconsistent. This further motivates the use of the sRGB space, both for the developing and for the error computation.

Table 1. Performances on the colorchecker dataset in sRGB space. Results against learning-free methods ([2,3,14,15]) and methods involving a training phase ([4,6,8]).

Method	Mean	Median	Trimean	Best-25%(μ)	Worst-25%(μ)
White-Patch [2]	9.6°	7.7°	8.6°	2.2°	20.3°
Grey-World [3]	9.1°	8.6°	8.8°	2.9°	16.2°
general Grey-World [14]	6.7°	4.9°	5.4°	1.4°	14.9°
1st – order Grey-Edge [15]	7.0°	6.3°	6.4°	2.8°	12.5°
Pixel-based Gamut Mapping [4]	5.8°	3.4°	4.1°	0.7°	14.4°
Edge-based Gamut Mapping [6]	10.1°	7.4°	8.2°	2.8°	21.7°
ML (category-wise prior) [8]	4.3°	3.6°	3.7°	1.4°	8.4°
Our method	4.1°	2.9°	3.4°	0.7°	12.8°

Table 2. Performances on the colorchecker dataset in camera RGB space

Method	Mean	Median	Trimean	Best-25%(μ)	Worst-25%(μ)
White-Patch [2]	7.5°	5.7°	6.4°	1.5°	16.2°
Grey-World [3]	6.4°	6.3°	6.3°	2.3°	10.6°
general Grey-World [14]	4.7°	3.5°	3.8°	1.0°	10.2°
1st – order Grey-Edge [15]	5.4°	4.5°	4.8°	1.9°	10.0°
Pixel-based Gamut Mapping [4]	4.2°	2.3°	2.9°	0.5°	10.8°
Edge-based Gamut Mapping [6]	6.7°	5.0°	5.5°	2.0°	14.0°
ML (category-wise prior) [8]	3.1°	2.3°	2.5°	0.9°	6.5°
Our method	4.1°	3.1°	3.3°	0.9°	9.0°

We compare our algorithm against standards learning-free illuminant estimation schemes and against some of the best methods requiring a training phase. Results are summarized in Table 1,2 and 3 . The first Table shows the results obtained in the sRGB colorspace ² for the Gehler database. The same comparisons performed in the camera RGB space are given in Table 2, enabling a direct comparison with the results from [1]. Table 3 provides the results on the SFU

² In order to compute errors in sRGB for other approaches, we project the illuminant estimations provided by [1] in sRGB. This may differ from a direct estimation of the illuminant in sRGB.

Table 3. Performances on the linear SFU grey-ball database in sRGB space

Method	Mean	Median	Trimean	Best-25%(μ)	Worst-25%(μ)
White-Patch [2]	12.7°	10.5°	11.3°	2.5°	26.2°
Grey-World [3]	13.0°	11.0°	11.5°	3.1°	26.0°
general Grey-World [14]	12.6°	11.1°	11.6°	3.8°	23.9°
1st – order Grey-Edge [15]	11.1°	9.5°	9.8°	3.2°	21.7°
Pixel-based Gamut Mapping [4]	11.8°	8.9°	10.0°	2.8°	24.9°
Edge-based Gamut Mapping [6]	13.7°	11.9°	12.3°	3.7°	26.9°
ML [8]	10.3°	8.9°	9.2°	2.8°	20.3°
Our method	10.9°	8.5°	9.1°	2.2°	23.5°

grey-ball database. Our approach shows better results than standard, learning-free methods, both in sRGB and in the camera colorspace, and both in term of mean and median of errors. The method is quite stable, providing good results on two databases with the same set of parameters. Our results are also not far from those obtained with more complex learning-based methods, at a very reasonable computational cost. Observe that errors in the sRGB colorspace are higher than those obtained in the camera colorspace. This can be explained by the reduced dynamic range of the camera color component.

5 Conclusion

In this paper, we have presented an automatic method for the estimation of the illuminant from a digital image. The method relies on projections on the Planck locus and on a voting scheme. Its efficiency has been demonstrated through various comparisons on two classical databases. One of the main interests of the method is a detailed analysis of the informative points in the image, obtained through a histogram of weighted projections on the Planckian locus. An interesting perspective of the work is a finer analysis of this histogram, an application of which could be the estimation of multiple illuminants of a scene.

Acknowledgement. This work is partially founded by the ANR projects Callisto and FUI project CEDCA.

References

1. Gijsenij, A., Gevers, T., van de Weijer, J.: Computational color constancy: Survey and experiments. *IEEE Trans. Image Process.* 20, 2475–2489 (2011)
2. Land, E., McCann, J.J.: Lightness and retinex theory. *J. Opt. Soc. Am.* 61, 1–11 (1971)
3. Buchsbaum, G.: A spatial processor model for object colour perception. *J. Franklin Inst.* 310, 1–26 (1980)
4. Forsyth, D.A.: A novel algorithm for color constancy. *Int. J. Comput. Vision* 5, 5–36 (1990)

5. Finlayson, G.D., Hordley, S.D.: Gamut constrained illuminant estimation. *Int. J. Comput. Vision* 67 (2006)
6. Gijsenij, A., Gevers, T., van de Weijer, J.: Generalized gamut mapping using image derivative structures for color constancy. *Int. J. Comput. Vision* 86, 127–139 (2010)
7. Finlayson, G.D., Schaefer, G.: Solving for colour constancy using a constrained dichromatic reflection model. *Int. J. Comput. Vision* 42, 127–144 (2001)
8. Chakrabarti, A., Hirakawa, K., Zickler, T.: Color constancy with spatio-spectral statistics. *IEEE Trans. Pattern Anal. Mach. Intell.* (2012)
9. Gijsenij, A., Gevers, T.: Color constancy using natural image statistics and scene semantics. *IEEE Trans. Pattern Anal. Mach. Intell.* 33, 687–698 (2011)
10. Shafer, S.A.: Using color to separate reflection components, pp. 43–51 (1992)
11. Judd, D.B., Macadam, D.L., Wyszecki, G., Budde, H.W., Condit, H.R., Henderson, S.T., Simonds, J.L.: Spectral distribution of typical daylight as a function of correlated color temperature. *J. Opt. Soc. Am.* 54, 1031–1036 (1964)
12. Gehler, P.V., Rother, C., Blake, A., Minka, T., Sharp, T.: Bayesian color constancy revisited. In: *Conference on Computer Vision and Pattern Recognition*, pp. 1–8 (2008)
13. Finlayson, G.D., Trezzi, E.: Shades of gray and colour constancy. In: *Color Imaging Conference*, pp. 37–41 (2004)
14. Barnard, K., Martin, L., Coath, A., Funt, B.: A comparison of computational color constancy algorithms. *IEEE Trans. Image Process.* 11, 2002 (2002)
15. Van De Weijer, J., Gevers, T., Gijsenij, A.: Edge-based color constancy. *IEEE Trans. Image Process.* 16, 2207–2214 (2007)
16. Li, B., Xu, D., Xiong, W., Feng, S.: Color constancy using achromatic surface. *Color Res. Appl.* 35, 304–312 (2010)
17. Kwon, H.J., Lee, S.H., Bae, T.W., Sohng, K.I.: Compensation of de-saturation effect in hdr imaging using a real scene adaptation model. *J. Visual Commun. Image Represent.* (2012)
18. Sapiro, G.: Color and illuminant voting. *IEEE Trans. Pattern Anal. Mach. Intell.* 21, 1210–1215 (1999)
19. Riess, C., Eibenberger, E., Angelopoulou, E.: Illuminant estimation by voting (2009)
20. Vazquez-Corral, J., Vanrell, M., Baldrich, R., Tous, F.: Color constancy by category correlation. *IEEE Transactions on Image Processing* 21, 1997–2007 (2012)
21. Wyszecki, G., Stiles, W.S.: *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd edn. Wiley-Interscience (2000)
22. Priest, I.G.: A proposed scale for use in specifying the chromaticity of incandescent illuminants and various phases of daylight. *J. Opt. Soc. Am.* 23, 41–45 (1933)
23. Ciurea, F., Funt, B.: A large image database for color constancy research (2003)
24. Shi, L., Brian: Re-processed version of the gehler color constancy dataset of 568 images (2010)