

Adaptive Color Constancy Using Faces

Simone Bianco, *Member, IEEE* and Raimondo Schettini, *Member, IEEE*

Abstract—In this work we design an adaptive color constancy algorithm that, exploiting the skin regions found in faces, is able to estimate and correct the scene illumination. The algorithm automatically switches from global to spatially varying color correction on the basis of the illuminant estimations on the different faces detected in the image. An extensive comparison with both global and local color constancy algorithms is carried out to validate the effectiveness of the proposed algorithm in terms of both statistical and perceptual significance on a large heterogeneous data set of RAW images containing faces.

Index Terms—Color constancy, face detection, global illuminant estimation, local illuminant estimation

1 INTRODUCTION

COMPUTATIONAL color constancy aims to estimate the actual color in an acquired scene disregarding its illuminant rendering the color of the objects in the scene invariable in spite of changes in the illumination conditions. This can be done either by computing features that are invariant to the color of the scene illumination [54], or by estimating the scene illuminant from the recorded image and then by transforming it such that the effect of the color of the light source is removed [17]. This paper focuses on the illuminant estimation approach to color constancy, which is also called white balancing [36], [42]. Color constancy is useful for various computer vision applications [54], such as image segmentation, object recognition, scene classification, for digital photography [55] and image forensics [40]. Since color constancy is an ill-posed problem, as its solution lacks uniqueness and stability [12], many different solutions exist in the literature, each based on different assumptions [22]. It is known that when these assumptions are not met the algorithm's estimate of the actual illuminant can be very poor. It has been shown that the universal best and the universal worst algorithms do not exist [4]: the algorithm that performs best for a specific image depends on the image content. For this reason a recent research area which has shown promising results aims to improve illuminant estimation by using visual information automatically extracted from the images. The existing algorithms exploit both low-level [3], [15], intermediate-level [2] and high-level [30] visual information. Moreover, most state-of-the-art color constancy methods were designed assuming that in the scene a uniform illumination is present. In real-world images, this assumption is often violated as more than one light source, with different spectral distributions, are present.

Hansen et al. [20] showed that in the human visual system, memory colors could be used as hints to give an

estimate of the illuminant in the scene. Moreno et al. [25] obtained memory colors for three different objects (grass, snow and sky) using psychophysical experiments. They then used a supervised image segmentation method to detect memory color objects and exploit them to color correct the image using a weighted Von Kries formula. In [18] a color constancy approach is used to compensate for skin color variations to achieve accurate skin color segmentation, while the use of contextual information in the form of detected faces to improve skin pixel detection for tracking has been proposed by Soriano et al. [28]. Skin tones have been also used to automatically tune the parameters of different enhancement algorithms to correspond to human preferences regarding the appearance of people in an image [26], [41]. However, they assume that the image has already been white balanced and mapped in a standard color space. Montojo presents a post hoc chromatic adaptation method for semi-automatically removing color casts due to the incorrect application of in-camera white balance settings [37]. After choosing a reference image which exhibits the desired skin tones of a particular person, correction is done in a perceptually uniform opponent color space derived from the Munsell color system. Its evaluation is limited, since no performance measure is given and the method requires face recognition to get per-person skin color models. In [38] variations in the chromaticity gamut of varying types of pre-recognized human skin under varying illumination are characterized. Using a LED illuminator and a spectroradiometrically calibrated hyperspectral camera, they showed that human skin gamuts in cone-contrast space are characterized by a set of features that can be used to differentiate between similar illuminations, whose estimate can then be used to color correct an image. They tested their approach on hyperspectral images of hands taken from 8 different human subjects, under 39 distinct illuminations.

Preliminary findings reported in this paper appeared in [5], where we showed that skin colors provide enough and reliable information to estimate the scene illuminant. We showed that since skin colors tend to form a cluster in color spaces [13], [24], the diversity between the gamut of skin pixels of the detected faces and the skin canonical gamut can be affordably used to estimate the scene illuminant. The method was compared with state-of-the-art algorithms on a standard data set of RAW camera images having a known

- The authors are with the Department of Informatics, Systems and Communication, University of Milano-Bicocca, Viale Sarca 336, Edificio U14, Milano, Italy. E-mail: {bianco, schettini}@disco.unimib.it.

Manuscript received 26 July 2013; revised 3 Dec. 2013; accepted 27 Dec. 2013.
Date of publication 1 Jan. 2014; date of current version 10 July 2014.

Recommended for acceptance by D. Forsyth.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TPAMI.2013.2297710

color target [14]. This paper extends the work in [5] in several ways:

- Since one of the assumptions that is often violated in color constancy is the presence of a uniform illumination in the scene, we have extended the applicability of the proposed algorithm to the case of non-uniform illumination. The method is adaptive, being able to distinguish and process in different ways images of scenes taken under a uniform and those acquired under non-uniform illumination. To the best of our knowledge, this is the first algorithm which automatically modify its behavior from global to local color correction according to the analysis of the image content.
- We design a more efficient algorithm to estimate the scene illuminant from extracted skin regions using only their mean color value instead of the whole gamut. This algorithm is more suitable for resource-limited camera devices, such as consumer digital cameras and camera phones.
- Preliminary results reported in [5] included only 95 images having faces. We present a much more detailed experimental evaluation using a new portrait data set consisting of a total of 1,145 RAW images containing faces and at least one known color target for benchmarking. Statistical and perceptual significance tests are carried out. The results obtained confirm that the proposed methods are statistically and perceptually better than the other ones in the state of the art.

2 PROBLEM FORMULATION AND RELATED WORKS

The image values for a Lambertian surface located at the pixel with coordinates (x, y) can be seen as a function $\rho(x, y)$, mainly dependent on three physical factors: the illuminant spectral power distribution $I(x, y, \lambda)$, the surface spectral reflectance $S(x, y, \lambda)$ and the sensor spectral sensitivities $\mathbf{C}(\lambda)$. Using this notation $\rho(x, y)$ can be expressed as

$$\rho(x, y) = \int_{\omega} I(x, y, \lambda) S(x, y, \lambda) \mathbf{C}(\lambda) d\lambda, \quad (1)$$

where λ is the wavelength, ω is the wavelength range of the visible light spectrum, ρ and $\mathbf{C}(\lambda)$ are three-component vectors. Since the three sensor spectral sensitivities are usually respectively more sensitive to the low, medium and high wavelengths, the three-component vector of sensor responses $\rho = (\rho_1, \rho_2, \rho_3)$ is also referred to as the sensor or camera RGB = (R, G, B) triplet. The goal of color constancy is to estimate the color $\mathbf{I}(x, y)$ of the scene illuminant, i.e., the projection of $I(x, y, \lambda)$ on the sensor spectral sensitivities $\mathbf{C}(\lambda)$:

$$\mathbf{I}(x, y) = \int_{\omega} I(x, y, \lambda) \mathbf{C}(\lambda) d\lambda. \quad (2)$$

As it is more important to estimate the chromaticity of the scene illuminant than its overall intensity [21], usually the illuminant color is estimated up to a scale factor. Since the only information available are the sensor responses ρ across the image, color constancy is an

under-determined problem [12]; and thus further assumptions and/or knowledge are needed to solve it.

2.1 Color Constancy under Uniform Illumination

Several computational color constancy algorithms exist in the literature, each based on different assumptions. The most common assumption they make is that the color of the light source is uniform across the scene, i.e., $\mathbf{I}(x, y) = \mathbf{I}$. Van de Weijer et al. [29] have unified a variety of algorithms. These algorithms estimate the illuminant color \mathbf{I} by implementing instantiations of the following equation:

$$\mathbf{I}(n, p, \sigma) = \frac{1}{k} \left(\iint |\nabla^n \rho_{\sigma}(x, y)|^p dx dy \right)^{\frac{1}{p}}, \quad (3)$$

where n is the order of the derivative, p is the Minkowski norm, $\rho_{\sigma}(x, y) = \rho(x, y) \otimes G_{\sigma}(x, y)$ is the convolution of the image with a Gaussian filter $G_{\sigma}(x, y)$ with scale parameter σ , and k is a constant to be chosen such that the illuminant color \mathbf{I} has unit length (using the 2–norm). The integration is performed over all pixel coordinates. Different (n, p, σ) combinations correspond to different illuminant estimation algorithms, each based on a different assumption. For example, the Gray World (GW) algorithm [7]—generated setting $(n, p, \sigma) = (0, 1, 0)$ —is based on the assumption that the average color in the image is gray and that the illuminant color can be estimated as the shift from gray of the averages in the image color channels; the White Point (WP) algorithm [8]—generated setting $(n, p, \sigma) = (0, \infty, 0)$ —is based on the assumption that there is always a white patch in the scene and that the maximum values in each color channel are caused by the reflection of the illuminant on the white patch, and they can be thus used as the illuminant estimation; the Gray Edge algorithm [29]—generated setting for example $(n, p, \sigma) = (1, 0, 0)$ —is based on the assumption that the average color of the edges is gray and that the illuminant color can be estimated as the shift from gray of the averages of the edges in the image color channels.

The Gamut Mapping assumes that for a given illuminant, one observes only a limited gamut of colors [11]. It has a training phase in which a canonical illuminant is chosen and the canonical gamut is computed observing as many surfaces under the canonical illuminant as possible. Given an input image with an unknown illuminant, its gamut is computed and the illuminant is estimated as the mapping that can be applied to the gamut of the input image, resulting in a gamut that lies completely within the canonical gamut and produces the most colorful scene. If the spectral sensitivity functions of the camera are known, the color by correlation approach could be also used [10].

Bayesian approaches [14] model the variability of reflectance and of illuminant as random variables, and then estimate illuminant from the posterior distribution conditioned on image intensity data.

Given a set computational color constancy algorithms, in [2] an image classifier is trained to classify the images as indoor and outdoor, and different experimental frameworks are proposed to exploit this information in order to select the best performing algorithm on each class. In [3] it has been shown how intrinsic, low level properties of the images can be used to drive the selection of the best

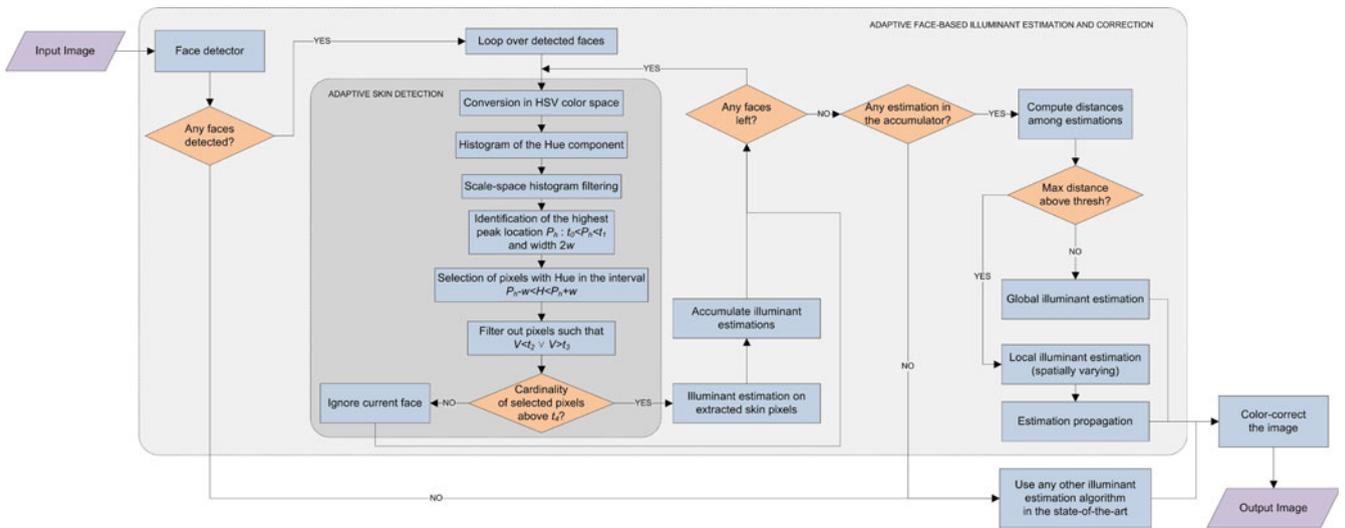


Fig. 1. The operation flowchart of the proposed adaptive method.

algorithm (or the best combination of algorithms) for a given image. The algorithm selection and combination is made by a decision forest composed of several trees on the basis of the values of a set of heterogeneous features.

In [15] the Weibull parameterization has been used to train a maximum likelihood classifier based on mixture of Gaussians to select the best performing color constancy method for a certain image.

In [9] a statistical model for the spatial distribution of colors in white balanced images is developed, and then used to infer illumination parameters as those being most likely under their model. High level visual information has been used to select the best illuminant out of a set of possible illuminants [30]. This is achieved by restating the problem in terms of semantic interpretability of the image. Several color constancy methods are applied to generate a set of illuminant hypotheses. For each illuminant hypothesis, they correct the image, evaluate the likelihood of the semantic content of the corrected image, and select the most likely illuminant color.

2.2 Color Constancy under Non-Uniform Illumination

The great majority of state-of-the-art color constancy methods assume that a uniform illumination is present in the scene. This assumption is often violated in real-world images. It is not trivial to extend the existing color constancy algorithms to work locally instead that globally, since the spatial support on which they accumulate the statistics is reduced, and the final local estimate could be biased by local image properties. One of the first color constancy methods developed is Retinex [35] which is able to deal with non-uniform illumination assuming that an abrupt change in chromaticity is caused by a change in reflectance properties. This implies that the illuminant smoothly varies across the image and does not change between adjacent or nearby locations. Ebner [34] proposed a method that assumes that the illuminant transition is smooth. The method uses the local space average color for local estimation of the illuminant by convolving the image with a Gaussian kernel function. Bleier et al. [39]

investigated whether existing color constancy methods, originally developed assuming uniform illumination, can be adapted to local illuminant color estimation using image sub-regions. Multiple independent estimations are then combined through regression to obtain a more robust final estimate. Gijssen et al. [36] proposed a method that makes use of local image patches, which can be selected by any sampling method. After sampling of the patches, illuminant estimation techniques are applied to obtain local illuminant estimates, and these estimates are combined into more robust estimations, since it is assumed that the number of different lights is less than the number of patches. This combination of local estimates is done with two different approaches: clustering if the number of lights is known, segmentation otherwise. A different class of algorithms is based on user guidance to deal with the case of two [47] and multiple lights [48].

3 THE PROPOSED APPROACH

In this work we propose a global/local adaptive color constancy method to estimate and correct the scene illuminant using faces. The operation flowchart of the proposed adaptive method is reported in Fig. 1. The face detector module is run on the input image to detect any faces. If no faces are detected, the input image may be processed with any other state-of-the-art illuminant estimation algorithm. If one or more faces are detected, a skin detection module is run on the detected faces to filter out any non-skin and unreliable pixels. A local illuminant estimation is made on the detected skin pixel of each face. If the maximum distance among the estimations on the different faces is lower than a fixed threshold, the local estimates are combined into a unique global illuminant estimate; otherwise the single face estimates are propagated to the rest of the image to give a local illuminant estimate for each pixel of the image.

Our method is based on three assumptions:

- skin colors form a sufficiently compact cluster in the color space in order to represent a valid clue for illuminant estimation [20];

- the illumination on each face is uniform;
- the illumination estimated on the faces properly sample the illumination distribution in the scene.

To have an idea of the applicability of the proposed method, we have generated different queries on Flickr (<http://www.flickr.com/>) using very generic tags such as cameras and mobile phones manufacturers. We found that among 30 and 60% of the returned images were portraits or included faces. More specific queries such as party, family, birthday, holiday, etc. usually contain a much higher number of faces. These results are consistent with a qualitative analysis of the camera phone photospace distribution reported by the International Imaging Industry Association (I3A) [19].

3.1 Adaptive Skin Detection

Looping on all the faces detected, the first step is a preliminary adaptive skin detection aimed to discard pixels that for sure are not skin pixels. Next the detected face pixels are converted in the HSV color space. Then a technique based on scale-space histogram filtering [33] is used to identify the highest peak location and width of the histogram of the hue component, within the hue interval corresponding to feasible skin colors. This involves generating a multi-scale description of the histogram $h(x)$ by convolving it with a series of Gaussians of increasing scale:

$$H(x, \sigma) = h(x) \otimes G_\sigma(x) = \int_{-\infty}^{+\infty} h(t) \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-t)^2}{2\sigma^2}} dt. \quad (4)$$

$H(x, \sigma)$ forms a two-dimensional surface called the scale-space image. The locations of valleys and peaks are easily located in terms of zero-crossing in the scale-space image, since at any value of σ , the extrema in the n th derivative of the smoothed histogram are given by the zero-crossing in the $(n+1)$ th derivative, computed using the relation

$$\frac{\partial^n H(x, \sigma)}{\partial x^n} = h(x) \otimes \frac{\partial^n G_\sigma(x)}{\partial x^n}, \quad (5)$$

where the derivatives of the Gaussian are readily obtained. Although this technique applies to zeros in any derivative we will focus our attention on those in the second derivative [33].

The highest peak location P_h and width $2w$ are identified within the hue interval corresponding to feasible skin colors, i.e., $P_h \in [t_0 \ t_1]$. Then only those pixels satisfying the condition $P_h - w < H < P_h + w$ are selected. This condition permits to implement a sort of adaptive skin detector, since the selected hue interval depends on the current peak location P_h . Any unreliable skin pixel as being too dark or too bright, and thus potentially clipped, is filtered out if it satisfies the condition $V < t_2 \vee V > t_3$. If the cardinality of the remaining skin pixels, normalized for the total number of pixels in the detected face, is above the threshold t_4 , they are converted into the YCbCr color space where they are luminance normalized such that the average luminance $\bar{Y} = 0.5$ and used to give an estimate of the local illuminant, which is then added in an accumulator; otherwise the current face is ignored. The thresholds

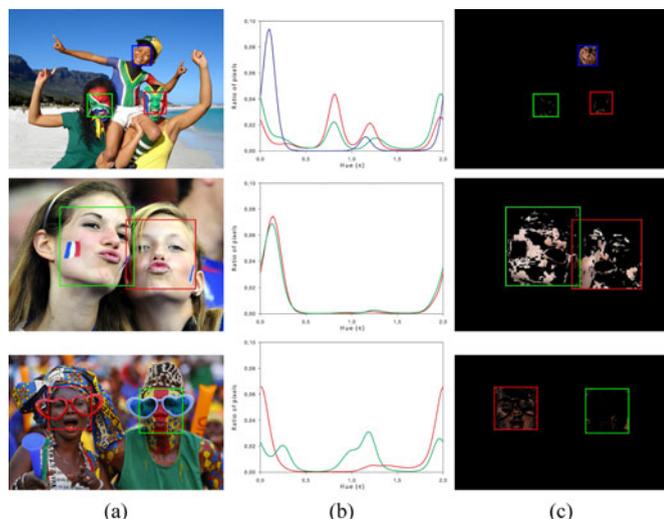


Fig. 2. An example of the combined effect of the implemented face and skin detector. The original image (a); the histograms of the hue component of the detected faces (b); the output of the adaptive skin detector (c).

t_0, \dots, t_4 could be either heuristically set or estimated from the synthetic data. In this work we have used an optimization procedure, described in the experimental section, exploiting a suitable data set.

It should be noted that the proposed method in the present form makes it possible to discard face regions with an unnatural or unreliable skin colors. An example of the combined effect of the implemented face and skin detector is reported in Fig. 2.

3.2 Skin-Based Illuminant Estimation

3.2.1 Skin-Based Gamut Mapping

Our approach applies gamut mapping to skin pixels only: the illuminant is estimated as the mapping that can be applied to the gamut of the skin colors in the input image, resulting in a gamut that lies completely within the skin canonical gamut.

The first step is the computation of the skin canonical gamut S_C : this is the convex hull of the skin colors of different people acquired under the chosen canonical illuminant.

Given an input image for which we want to estimate the illuminant, in which F faces are present, the masks $u_f(x, y)$, $f = 1, \dots, F$, are obtained:

$$u_f(x, y) = \begin{cases} 1 & \text{if } (x, y) \in \text{face no. } f \wedge \text{ is a skin pixel} \\ 0 & \text{otherwise,} \end{cases} \quad (6)$$

i.e., $u_f(x, y)$ assumes the value 1 only on the pixels inside the f -th detected face area, being classified as skin pixels. The extracted skin colors \mathbf{skin}_f are then computed as

$$\mathbf{skin}_f = \{\rho(x, y) : u_f(x, y) = 1\}. \quad (7)$$

Once we have extracted the skin colors for each detected face, these are converted into the YCbCr color space and luminance normalized, such that the average luminance $\bar{Y} = 0.5$. Given the face f , the skin gamut $S_{I,f}$ is then computed as the convex hull of the converted values of the skin pixels belonging to the current face, thus obtaining a

different gamut for each face in the image. The set of feasible mappings \mathcal{M}_f for each face is then determined: it consists of all mappings that can be applied to the skin gamut $S_{I,f}$ that result in a gamut that lies completely within the skin canonical gamut S_C :

$$\mathcal{M}_f = \{\mathcal{M}_{f,i} : \mathcal{M}_{f,i}S_{I,f} \in S_C\}. \quad (8)$$

In this work each transformation $\mathcal{M}_{f,i}$ is modeled by a diagonal mapping or von Kries Model [32]. The local illuminant color I_f is then estimated as the inverse of the average of the feasible set [1].

3.2.2 The Skin Patch (SP) Illuminant Estimation

We also experiment here the use of the mean values of the extracted skin colors estimating the scene illuminant as the ratio among the computed average skin color and the reference average skin color. Using the same formalism defined in the previous section, the illuminant estimation on the face f can be written as:

$$I_f = \frac{\overline{skin}_f}{\rho_{skin}} = \frac{\sum_{(x,y) \in f} \rho(x,y)u_f(x,y)}{\rho_{skin} \sum_{(x,y) \in f} u_f(x,y)}, \quad (9)$$

where \overline{skin}_f is the computed average skin color over the face f and ρ_{skin} is the reference average skin color. Given the analogy with the White Patch algorithm, we call this method skin patch. The main advantage of the SP algorithm with respect to the one described in Section 3.2.1 is its efficiency, since the latter requires a convex optimization to be solved and it is not guaranteed that such a solution exists.

We want now to understand if skin reflectances are statistically equivalent to a neutral patch for the estimation of the scene illuminant. To this end, the data set of skin reflectances has been extracted from the ISO [23], for a total of 697 heterogeneous samples. The same analysis carried out in [6] is performed, where we have shown that there is a subset of colors in the Macbeth Color-Checker (MCC) DC and Munsell Atlas which are statistically equivalent to a neutral patch for the estimation of the scene illuminant. The statistical equivalence is tested under 321 different illuminants with a Correlated Color Temperature (CCT) ranging from about 2,000 to 13,000 K. Given a canonical illuminant $I_c(\lambda)$, for each illuminant and for each skin reflectance, the cumulative histogram of the normalized distance metric (ND) [43] is computed with respect to the elements of the Munsell Atlas using two different illuminant corrections. The first one is computed as the element-wise ratio between the camera values of a neutral patch with reflectance $n(\lambda)$ under the actual illuminant $I_i(\lambda)$ and those of the same neutral patch under the canonical illuminant $I_c(\lambda)$, i.e.:

$$\frac{\rho_{n,c}}{\rho_{n,i}} = \frac{\int_{\omega} I_c(\lambda)n(\lambda)\mathbf{C}(\lambda)d\lambda}{\int_{\omega} I_i(\lambda)n(\lambda)\mathbf{C}(\lambda)d\lambda}, \quad (10)$$

while the second one is computed as the element-wise ratio between the camera values of the skin reflectance $s_j(\lambda)$

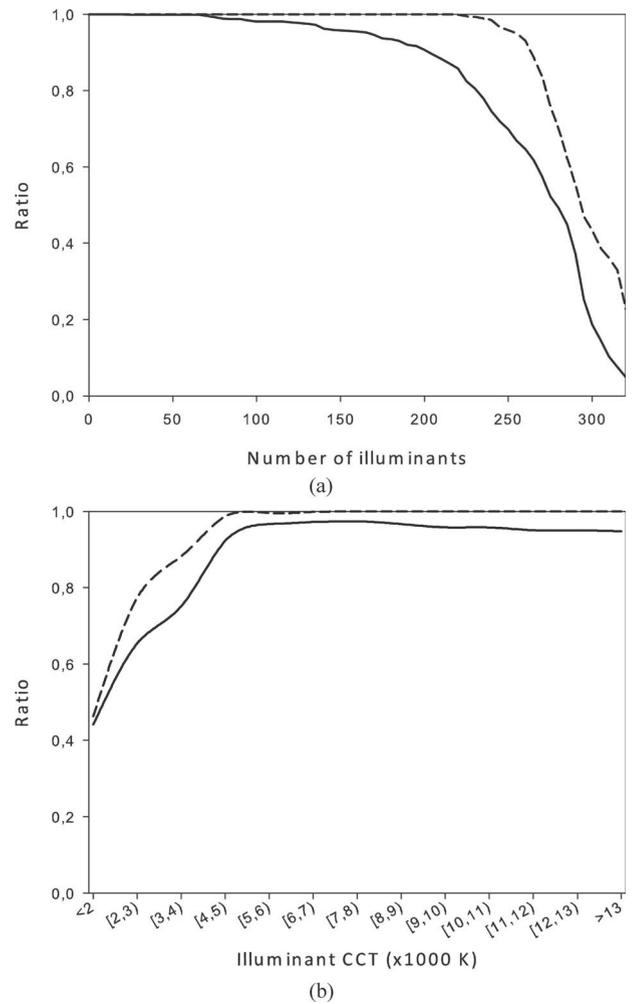


Fig. 3. Results of the statistical equivalence tests using Eq. (11) (dashed line) and using the average skin reflectance in the denominator of Eq. (11) (solid line). Ratio of skin reflectances judged statistically equivalent to a neutral patch under a given number of illuminants (a); ratios for different illuminant CCT ranges (b).

under the actual illuminant and those of the same skin reflectance under the canonical illuminant, i.e.:

$$\frac{\rho_{j,c}}{\rho_{j,i}} = \frac{\int_{\omega} I_c(\lambda)s_j(\lambda)\mathbf{C}(\lambda)d\lambda}{\int_{\omega} I_i(\lambda)s_j(\lambda)\mathbf{C}(\lambda)d\lambda}. \quad (11)$$

The skin reflectance $s_j(\lambda)$ is considered statistically equivalent to a white surface under the illuminant $I_i(\lambda)$ if the respective ND cumulative histograms are statistically equivalent using the Wilcoxon Signed-Rank Test (WST) [44]. A second statistical equivalence test is run substituting the actual skin reflectance $s_j(\lambda)$ in the denominator of Eq. (11) with the average skin reflectance $\bar{s}(\lambda)$ of the whole data set. The results of the two statistical equivalence tests are reported in Fig. 3a. The lines represent the ratio of skin reflectances judged statistically equivalent to a white surface under a given number of illuminants. The ratios of skin reflectances judged statistically equivalent to a neutral patch for different illuminant CCTs are reported in Fig. 3b.

The plots reported in Fig. 3b reveal that almost all the skin reflectances are equivalent to a neutral patch under illuminants with a CCT above about 4,000 K, while this ratio decreases for lower CCTs.

Proved that using the ratio among the extracted skin color and the average skin color for estimating the illuminant is not equivalent to a neutral patch under all the possible illuminants, we could still use it as long as its illuminant estimate is good enough. To this end, the angular error [21] among the actual and estimated illuminant is computed for all the illuminants and skin reflectances used. Using the estimator in Eq. (11) gives a median angular error of 1.06° . The median error is of 1.70° when the ratio of the skin reflectance under the actual illuminant and that of the average skin reflectance under the canonical illuminant is used. Since Hordley suggested that an angular error of less than 2° represents good enough color constancy performance on complex images [22], we have decided to include the skin patch algorithm in the experiments.

In order to use the proposed algorithms, we have to first calculate the thresholds for the adaptive skin detector, the convex hull of the skin colors under the canonical illuminant for what concerns the skin-based gamut mapping, and the reference skin ρ_{skin} for the skin patch algorithm. The thresholds, S_C and ρ_{skin} are obtained through an optimization procedure aimed at minimizing the angular error on a set of training images. The objective function adopted is given as pseudocode in Algorithm 1, where g_i and o_i , $i = 1, 2$, are gain and offset terms.

Algorithm 1 Input: $((t_0, \dots, t_4), g_1, o_1, g_2, o_2)$

- 1: Run the adaptive skin detection on training images
 - 2: Color correct extracted skin pixels using ground truth illuminants
 - 3: Intensity normalize the color corrected skin pixels
 - 4: Compute Canonical Gamut S_C
 - 5: $S_C \leftarrow S_C g_1 + o_1$
 - 6: Compute skin average color ρ_{skin}
 - 7: $\rho_{\text{skin}} \leftarrow \rho_{\text{skin}} g_2 + o_2$
 - 8: **for all** Training images **do**
 - 9: Run the adaptive skin detection
 - 10: Intensity normalize the extracted skin pixels
 - 11: Compute the actual gamut
 - 12: Estimate the illuminant with skin-based Gamut Mapping
 - 13: Compute $e_{\text{ang,GM}}$ among estimated and ground truth illuminants
 - 14: Compute average skin color
 - 15: Estimate the illuminant with Skin Patch algorithm
 - 16: Compute $e_{\text{ang,SP}}$ among estimated and ground truth illuminants
 - 17: **end for**
 - 18: **return** $e_{\text{ang,GM}}, e_{\text{ang,SP}}, S_C, \rho_{\text{skin}}$
-

3.3 Adaptive Illuminant Estimation and Image Correction

Independently on which skin-based illuminant estimation is used, when the looping on all the detected faces is finished, the accumulator containing the local illuminant estimates is

analyzed. If it is empty, any other algorithm in the state-of-the-art can be used to estimate the illuminant in the scene. Otherwise, the angular distances among all the local estimates in the accumulator are computed. If the maximum distance is below a predefined threshold, the algorithm considers the image as having a uniform illumination and the estimates are combined to give a unique global illuminant estimate. The final global estimate is obtained by averaging the L_1 -normalized values of the local estimates. Once we have estimated the global illuminant color $\mathbf{I} = [I_R, I_G, I_B]$, given the choice of diagonal mappings, each pixel in the image is color corrected using the von Kries model [32], i.e.: $\rho_{\text{out}}(x, y) = \text{diag}(1/\mathbf{I})\rho_{\text{in}}(x, y)$.

If the maximum distance among the local estimates in the accumulator is above the predefined threshold, the algorithm considers the image as having non-uniform illumination and thus a spatially varying illuminant estimation and correction is made. Since our local illuminant estimates are localized on the detected faces, in order to obtain a per-pixel illuminant estimate, the spatially varying illuminant estimate $\mathbf{I}(x, y)$ is obtained by taking the local estimates as seeds and propagating them to the rest of the image with a nearest neighbor diffusion scheme: given n local estimates \mathbf{I}_f , $f = 1, \dots, n$ located at the corresponding centroids $c_f = (x_f, y_f)$, $f = 1, \dots, n$ of the corresponding detected faces, the illuminant estimate at a generic pixel (x, y) is computed as:

$$\mathbf{I}(x, y) = \frac{1}{k} \sum_{i=1}^k \mathbf{I}_{j_i} \quad \text{s.t.} \quad \|c_j - (x, y)\|_2 = \min_f \|c_f - (x, y)\|_2. \quad (12)$$

Finally, a Gaussian filter is applied to the spatially varying illuminant estimate to get a smooth illuminant estimate. The spatially varying corrected image is then obtained applying a different von Kries model to each pixel, i.e.: $\rho_{\text{out}}(x, y) = \text{diag}(1/\mathbf{I}(x, y))\rho_{\text{in}}(x, y)$.

The threshold used to decide if the current scene is illuminated by a single light or by multiple light sources could be heuristically set. Following [22] we set this threshold equal to 3° , since it has been judged to be a noticeable but acceptable difference.

4 EXPERIMENTAL SETUP

The aim of this section is to investigate if the proposed algorithm can outperform state-of-the-art algorithms in the illuminant estimation on images containing at least one face.

4.1 Image Data Sets and Evaluation Procedure

To test the performance of the proposed algorithm, two data sets of RAW camera images having a known color target are used. Both data sets are captured using high-quality digital SLR cameras in RAW format, and are therefore free of any color correction. The first data set [14] was originally available in sRGB-format, but Shi and Funt [27] reprocessed the raw data to obtain linear images with a higher dynamic range (14 bits as opposed to standard 8 bits). The data set has been acquired using a Canon 5D and a Canon 1D DSLR cameras and consists of a total of 568 images, 95 of which are portraits or include faces and constitute what we have



Fig. 4. Examples of images within the Cambridge portrait data set.



Fig. 5. Examples of images within the Milan portrait data set.

called the Cambridge portrait data set. This portrait data set contains people of different nationalities. The Macbeth ColorChecker chart is included in every scene acquired, and this allows to accurately estimate the actual illuminant of each acquired image. Examples of images within the Cambridge portrait data set are reported in Fig. 4.

The Milan portrait data set has been acquired in order to further evaluate the algorithms proposed on a larger number of images. It has been acquired in RAW format using four different DSLR cameras: Canon 40D, Canon 350D, Canon 400D, and Nikon D700. The data set is the union of different subsets that have been acquired in three different world locations: Italy, Taiwan, and Japan. The data set includes portraits of a single person with a single MCC up to multiple persons with multiple MCCs. The only difference among the subsets is that the Taiwan subset contains a neutral-color reference card instead of the MCC. The Milan data set composition is summarized in Table 1, while some examples of images within the data set are reported in Fig. 5.

For both the portrait data sets, during experiments the MCC has been masked to avoid biasing the algorithms.

4.2 Benchmark Algorithms

Preliminary results on the Cambridge portrait data set were reported in [5]. Different benchmarking algorithms for color constancy were considered. Since each image of the data set contains only one MCC, only global color constancy algorithms based on the assumption of uniform illumination can be compared. Six of them were generated varying the three variables (n, p, σ) in Equation (3), and correspond to well known and widely used color constancy algorithms. The values chosen for (n, p, σ) are reported in Table 2 and set as in [17]. The algorithms are used in the original authors' implementation which is

TABLE 1
Milan Portrait Data Set Composition

	Single MCC	Multiple MCCs
Single person	920	0
Multiple persons	28	197

freely available online (<http://lear.inrialpes.fr/people/vandeweyer/code/ColorConstancy.zip>). The seventh algorithm is the pixel-based Gamut Mapping [16]. The value chosen for σ is also reported in Table 2. The other algorithms considered were the Bayesian (BAY [14]); the Natural Image Statistics (NIS [15]); the High Level Visual Information [30]: bottom-up (HLVI BU), top-down (HLVI TD), and their combination (HLVI BU&TD); the Spatio-Spectral statistics [9]: with Maximum Likelihood estimation (SS ML), and with General Priors (SS GP); the Automatic color constancy Algorithm Selection (AAS) [3] and the Automatic Algorithm Combination (AAC) [3].

The last algorithm considered was the Do Nothing (DN) algorithm which gives the same estimation for the color of the illuminant ($\mathbf{I} = [1 \ 1 \ 1]$) for every image, i.e., it assumes that the image is already correctly balanced.

For what concerns the color constancy algorithms able to deal with non-uniform illumination, the algorithms considered are: the Retinex [35]; the Local Space Average Color with Gaussian kernel function (LSAC) [34]; two image sub-regions fusion schemes [39]; gradient tree boosting (GTB) and random forest regression (RFR); multiple light sources (MLS) [36] using White Point (WP) and Gray World (GW) algorithms, grid based sampling, in the clustering version, i.e., setting the number of lights equal to the number of targets in the scene.

5 RESULTS AND DISCUSSION

The state-of-the-art algorithms considered are run on both the portrait data sets. The error metric considered,

TABLE 2
Values Chosen for (n, p, σ) for the State-of-the-Art Algorithms Which Are Instantiations of Eq (3)

Algorithm	n	p	σ
Gray World (GW)	0	1	0
White Point (WP)	0	∞	0
Shades of Gray (SoG)	0	4	0
general Gray World (gGW)	0	9	9
1st-order Gray Edge (GE1)	1	1	6
2nd-order Gray Edge (GE2)	2	1	1
Gamut Mapping (GM)	0	0	4

TABLE 3
Angular Error Statistics Obtained by the State-of-the-Art Algorithms Considered on the Cambridge Portrait Data Set

Algorithm	Min	10 th prc	Med	Avg	90 th prc	Max
DN	6.43	9.93	13.20	13.10	15.77	25.36
GW	0.25	1.68	5.38	5.75	10.02	14.84
WP	0.19	1.42	6.22	7.62	15.63	40.54
SoG	0.42	1.28	3.80	4.61	9.12	13.90
gGW	0.19	1.04	3.44	4.51	9.29	14.87
GE1	0.49	1.74	4.44	4.90	8.34	11.50
GE2	0.68	1.53	3.90	4.34	7.27	12.87
GM	0.05	0.40	2.61	4.24	11.33	20.00
BAY	0.10	1.22	4.05	4.79	9.02	17.33
NIS	0.61	1.40	3.34	4.05	7.33	12.88
HLVI BU	0.33	0.89	2.95	3.45	6.71	14.96
HLVI TD	0.37	0.84	2.88	3.67	6.87	15.88
HLVI BU&TD	0.37	0.80	2.59	3.43	6.44	15.88
SS ML	0.34	0.92	3.45	3.78	7.55	12.30
SS GP	0.43	0.85	3.28	3.67	7.14	12.09
AAS	0.19	1.17	3.57	4.70	10.18	18.08
AAC	0.12	0.81	2.87	3.79	7.83	13.31
FACE _{ideal} GM	0.13	0.68	1.81	2.34	4.87	8.53
FACE _{real} GM	0.13	0.77	1.93	2.53	5.05	8.53
FACE _{ideal} SP	0.36	0.76	1.91	2.21	4.19	7.94
FACE _{real} SP	0.36	0.76	2.04	2.39	4.24	7.94

as suggested by Hordley and Finlayson [21], is the angle between the RGB triplet of estimated illuminant (ρ_w) and the RGB triplet of the measured ground truth illuminant ($\hat{\rho}_w$):

$$e_{\text{ANG}} = \arccos\left(\frac{\rho_w^T \hat{\rho}_w}{\|\rho_w\| \|\hat{\rho}_w\|}\right). \quad (13)$$

In order to apply the proposed algorithms, we have to first calculate the thresholds for the adaptive skin detector, the convex hull of the skin colors under the canonical illuminant for what concerns the skin-based gamut mapping, and the reference skin ρ_{skin} for the skin patch algorithm. Since we are considering RAW images, free of color correction, it is impossible to use information about skin clusters available in the state of the art [13]. The thresholds, S_C and ρ_{skin} are obtained through an optimization procedure aimed at minimizing the angular error on a set of training images. The objective function adopted is given as pseudocode in Algorithm 1, where g_i and o_i , $i = 1, 2$, are gain and offset terms. The results on the Cambridge portrait data set are obtained using a leave-one-out cross validation. This means that the skin canonical gamut S_C and reference skin color ρ_{skin} are recomputed for each image. When calculating S_C and ρ_{skin} for each image, all the images containing the same person were left out too. The results on the Milan portrait data set use the Cambridge portrait data set as training set.

In Table 3 the minimum, the 10th-percentile, the median, the average, the 90th-percentile, and the maximum of the angular errors obtained by the considered state-of-the-art algorithms and the global versions of proposed algorithms on the Cambridge portrait data set are reported. For each of the proposed methods two different results are reported: the first one refers to the ideal results that would be obtained using a perfect face detector (i.e., the faces are manually detected), the second one refers to the results

obtained using a real face detector which is an implementation of the widely used Viola-Jones [31], trained on an independent data set to detect frontal faces. The statistics using the face detector are referred to only those images in which it was able to detect at least one face (i.e., 82.11% of the images of the Cambridge portrait data set).

It can be noticed from the comparison of the results in Table 3 that the proposed skin-based Gamut Mapping—in the instantiation with the manual face detector—is able to estimate the illuminant in the scene with the highest accuracy, reducing the median angular error by 30.12% with respect to the HLVI BU&TD (which had the lowest median error among the algorithms in the state of the art). Furthermore the maximum angular error is reduced by 25.83% with respect to the GE1 (which had the lowest maximum error). The second best median error is obtained by the proposed skin patch algorithm, again in the instantiation with the manual face detector, which improves over the HLVI BU&TD by 26.25%. The maximum error is reduced by 30.96%. Coupled with the real face detector, the improvements in the median error are of 25.48 and 21.24% respectively. In Fig. 6, three images of the Milan portrait data set on which the proposed SP algorithm makes an angular errors above the 95th-percentile are reported.

In Table 4 the results obtained by the proposed algorithms on the Milan portrait data set are reported. They are relative to the images containing a single reference target. The proposed skin-based methods, when no faces are detected in the image by the face detector, estimate the illuminant in the scene using the GW algorithm. Experiments have been run using also the other algorithms considered, but the results obtained are substantially identical since there are only few images on which the face detector was unable to detect any face.

From the analysis of the results reported in Table 4 it is possible to notice that the best median and maximum errors are obtained by the proposed skin-based gamut mapping (using a manual face detector), which improves by 17.81 and 27.11% respectively with respect to SS ML. The lowest median and maximum errors obtained using a real face detector improve by 16.19 and 7.36% respectively with respect to SS ML.

Finally, in Table 5 the results obtained by proposed adaptive algorithms on the Milan portrait data set are reported and compared with some spatially varying algorithms in the state of the art. They are evaluated on the images of the data set containing multiple targets. Also global color constancy algorithms are reported in the comparison, together with the global versions of the proposed algorithms. As for the previous experiment, the proposed skin-based methods, when no faces are detected in the image by the face detector, estimate the illuminant in the scene using the GW algorithm. Experiments have been run using also the other algorithms considered, but the results obtained are substantially identical since there are only few images on which the face detector was unable to detect any face.

From the analysis of the results reported in Table 5 it is possible to notice that the best median error is obtained by the proposed adaptive skin patch algorithm using a real face detector, which improves by 28.23% with respect to SS

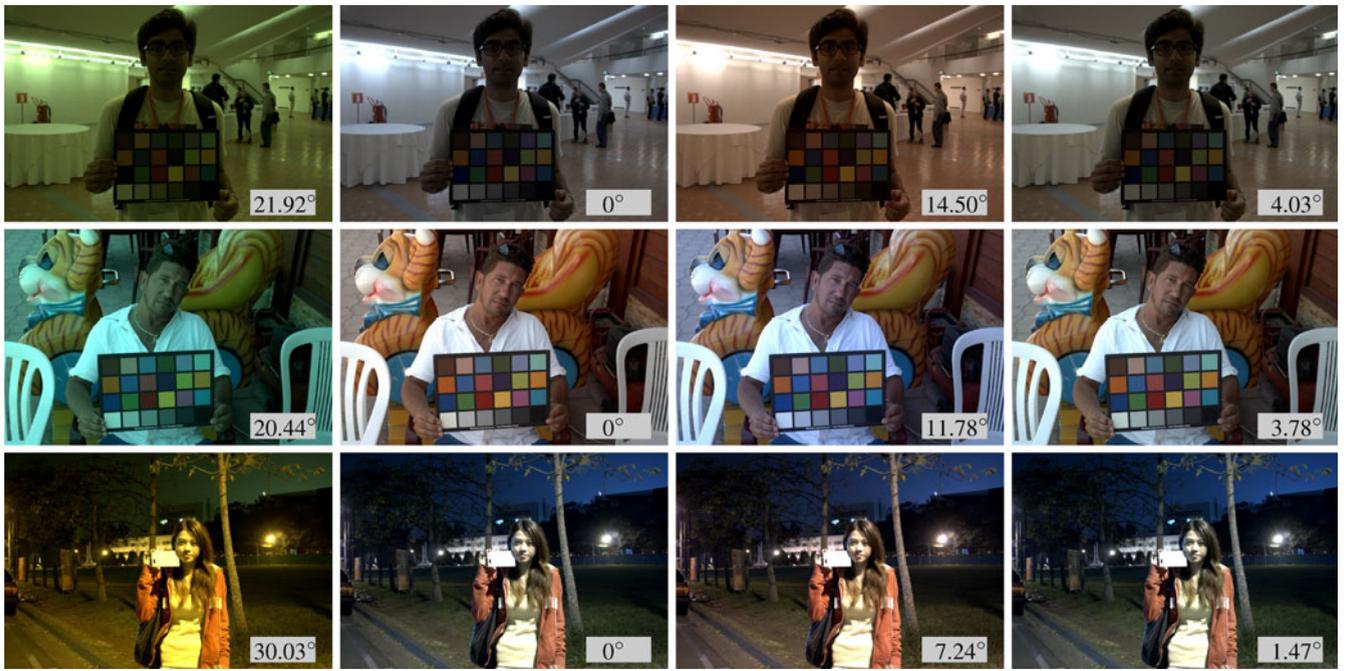


Fig. 6. Example of images in the Milan portrait data set on which the proposed SP algorithm makes an angular errors above the 95th-percentile. Left to right: original image; ideal correction based on the MCC; correction with the proposed algorithm; correction with the best algorithm in the state-of-the-art among the ones considered, i.e., SS ML.

ML. The improvement is even higher (i.e., 34.27%) with respect to MLS+WP which is the best spatially varying algorithm tested. The lowest maximum error is obtained by the proposed skin patch algorithm (in the spatially varying version) using the manual face detector, which is 32.29% lower than that obtained by GE2.

In Fig. 7 an example image that the proposed adaptive SP method considers having a non-uniform illumination is reported. The maximum distance among the illuminant estimations on the seven detected faces is 23.4° and thus the local version of the algorithm is used. In Fig. 8 an example image that the proposed adaptive SP method considers having a uniform illumination is reported. The illuminant estimations on the two detected faces differ by 1.1° and thus the global version of the algorithm is used.

TABLE 4
Angular Error Statistics Obtained on the Milan Portrait Data Set (Single Target)

Algorithm	Min	10 th prc	Med	Avg	90 th prc	Max
DN	8.98	16.80	18.89	19.30	21.89	30.03
WP	0.11	1.98	14.61	12.73	21.42	30.03
GW	0.05	1.00	3.54	4.55	9.01	27.16
SoG	0.11	0.74	2.74	4.50	11.19	29.17
gGW	0.11	0.84	3.61	5.62	13.27	31.13
GE1	0.07	0.95	3.33	4.42	9.48	25.82
GE2	0.04	2.19	4.60	5.50	10.02	24.00
GM	0.07	0.84	3.10	3.45	7.73	20.97
NIS	0.04	1.01	3.66	5.05	11.11	24.44
HLVI BU	0.05	0.92	2.91	4.12	8.84	21.31
HLVI TD	0.04	1.03	3.10	4.25	9.02	23.46
HLVI BU&TD	0.04	1.03	3.07	4.23	9.02	23.46
SS ML	0.10	0.67	2.47	3.47	7.69	19.44
FACE _{ideal} GM	0.02	0.80	2.03	2.63	4.99	14.17
FACE _{ideal} SP	0.05	0.75	2.16	2.62	4.76	14.50
FACE _{real} GM+GW	0.04	0.83	2.07	2.69	5.17	18.01
FACE _{real} SP+GW	0.04	0.84	2.19	2.74	4.85	18.01

5.1 Statistical Significance

In order to assess if the difference in performance among the different algorithms reported in Table 5 are statistically significant, we have used the Wilcoxon Signed-Rank Test (WST) [44] as suggested in [21]. Given two algorithms for which we want to test the statistical significance, the WST is run on the corresponding angular error distributions on the whole data set. The results of the statistical significance test for all the couples of algorithms considered are reported in Table 6a for the single target case, and in Table 6b for the multiple targets case. A

TABLE 5
Angular Error Statistics Obtained on the Milan Portrait Data Set (Multiple Targets)

Algorithm	Min	10 th prc	Med	Avg	90 th prc	Max
DN	10.84	15.71	17.30	17.53	19.74	28.60
WP	0.31	1.74	12.16	11.39	19.08	28.60
GW	0.20	1.18	4.26	4.86	9.26	20.04
SoG	0.16	0.88	4.39	5.93	14.03	20.02
gGW	0.11	0.83	5.25	6.42	15.07	20.80
GE1	0.19	1.52	4.59	5.08	9.50	18.22
GE2	0.36	2.45	4.93	5.39	9.69	15.36
SS ML	0.09	0.67	2.94	3.72	8.03	16.28
LSAC	0.20	1.59	4.23	4.79	8.66	18.99
RETINEX	0.12	1.56	4.28	4.83	8.39	20.54
MLS + WP	0.15	0.99	3.21	4.04	7.55	17.19
MLS + GW	0.11	0.92	3.33	4.18	8.82	17.97
Fusion Grad. Tree Boost.	0.03	1.51	4.48	5.29	9.95	31.26
Fusion Rand. Forest Regr.	0.03	1.07	3.23	3.96	7.61	27.76
FACE _{ideal} GM glo	0.20	1.21	2.53	3.09	5.78	13.05
FACE _{ideal} GM SV	0.18	1.12	2.62	3.01	5.35	12.20
FACE _{ideal} SP glo	0.11	1.05	2.12	2.74	5.43	11.57
FACE _{ideal} SP SV	0.11	0.87	2.14	2.61	4.88	10.40
FACE _{real} GM glo+ GW	0.03	1.16	2.59	3.13	5.88	11.70
FACE _{real} GM SV + GW	0.03	1.09	2.69	3.10	5.48	11.23
ADA FACE _{real} GM + GW	0.03	1.15	2.48	3.02	5.31	11.13
FACE _{real} SP glo+ GW	0.11	1.00	2.18	2.81	5.60	11.40
FACE _{real} SP SV + GW	0.11	0.86	2.21	2.73	5.04	11.43
ADA FACE _{real} SP + GW	0.11	0.87	2.11	2.66	5.15	11.43

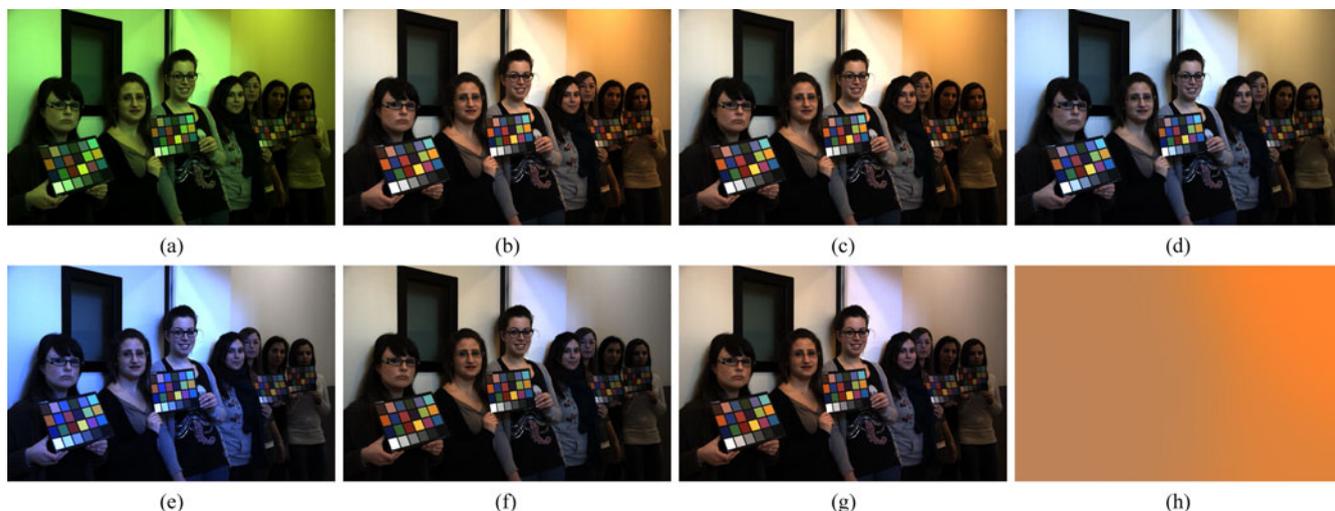


Fig. 7. Example image identified as having non-uniform illumination by the adaptive SP method. Original image (a); ideal corrections using from left to right: the first target (b), the second one (c), the third one (d), and the fourth one (e); ideal spatially varying correction using all the four targets (f); local correction using the adaptive SP method (g); estimated illuminant color map (h) corrected with the camera daylight multipliers for better visualization.

positive sign (+) at location (i, j) indicates that the median of the method i is significantly lower than that of method j at the 95% confidence level. A negative sign (-) indicates the opposite, and an equal sign (=) indicates that there is no significant difference between the two methods. In the last column of Table 6 the score indicating the number of methods with respect to which the corresponding method has been considered significantly better is reported.

5.2 Perceptual Significance

In addition to the statistical significance test, the perceptual significance is also investigated here. In fact, the fact that the difference between two algorithms is statistically significant might not always justify the conclusion that one algorithm is better than another [45]. The statistical significance test reported in Table 6 permits in fact to conclude that the error of one algorithm is often lower than that of another algorithm, but it does not tell if this difference is noticeable to a human observer. In [45] the just noticeable difference (JND) formula for the angular error between methods i and j with corresponding angular errors e_{ANG_i} and e_{ANG_j} is derived as $JND_{angular} = 0.06 \cdot \max\{e_{ANG_i}, e_{ANG_j}\}$. Following [45] and using this JND definition, a perceptual significance test on the results reported in Tables 4 and 5 is run by computing the quantity $p_{i,j} = (e_{ANG_i} - e_{ANG_j}) / \max\{e_{ANG_i}, e_{ANG_j}\}$. The results are reported in Table 7a for the single target case, and in Table 7b for the multiple targets case. A positive sign (+) at location (i, j) indicates that the method i is perceptually better than the method j (i.e., $p_{i,j} \leq -0.06$). A negative sign (-) indicates the opposite (i.e., $p_{i,j} \geq 0.06$), and an equal sign (=) indicates that there is no perceptual difference between the two methods (i.e., $-0.06 < p_{i,j} < 0.06$). The score in the last column indicates the number of methods with respect to which the corresponding method has been considered perceptually better.

The results of the statistical and perceptual significance tests on the Milan portrait data set containing a single target (Tables 6a and 7a) show that the proposed skin-based

gamut mapping using a manual face detector is both statistically and perceptually the best algorithm. The second statistically best algorithm is the instantiation of the skin-based gamut mapping using a real face detector, followed by the skin patch algorithm. The perceptual significance test shows that this difference is not significant. From the results of the statistical and perceptual significance tests

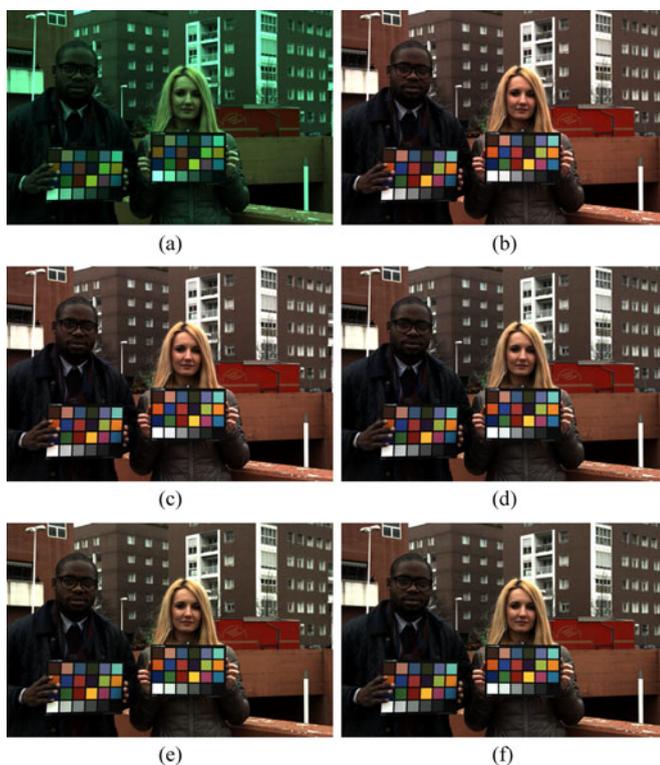


Fig. 8. Example image identified as having uniform illumination by the adaptive SP method. Original image (a); ideal corrections using: the target on the left (b), and the one on the right (c); correction using the illuminant estimated on the face on the left (d), and the face on the right (e); global correction using the adaptive SP method (f).

TABLE 6

Results of the Statistical Significance Test on the Milan Portrait Data Set: Single Target (a), Multiple Targets (b)

	1) DN	2) WP	3) GW	4) SoG	5) gGW	6) GE1	7) GE2	8) GM	9) NIS	10) HLVI BU	11) HLVI TD	12) HLVI BU & TD	13) SS ML	14) FACE _{ideal} GM	15) FACE _{ideal} SP	16) FACE _{real} GM + GW	17) FACE _{real} SP + GW	Score
1)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	0
2)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	1
3)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	4
4)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	11
5)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	3
6)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	6
7)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	2
8)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	3
9)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	7
10)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	10
11)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	7
12)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	7
13)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	12
14)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	15
15)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	13
16)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	14
17)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	13

(a)

	1) DN	2) WP	3) GW	4) SoG	5) gGW	6) GE1	7) GE2	8) SS ML	9) LSAC	10) RETINEX	11) MLS + WP	12) MLS + GW	13) Fusion Grad. Tree Boost.	14) Fusion Rand. Forest Regr.	15) FACE _{ideal} GM glo	16) FACE _{ideal} GM SV	17) FACE _{ideal} SP glo	18) FACE _{ideal} SP SV	19) FACE _{real} GM gbo+ GW	20) FACE _{real} GM SV + GW	21) ADA FACE _{real} GM + GW	22) FACE _{real} SP glo+ GW	23) FACE _{real} SP SV + GW	24) ADA FACE _{real} SP + GW	Score
1)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	0
2)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	1
3)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	7
4)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	5
5)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	2
6)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	4
7)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	3
8)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	13
9)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	7
10)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	7
11)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	11
12)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	10
13)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	5
14)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	11
15)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	15
16)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	14
17)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	19
18)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	19
19)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	15
20)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	14
21)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	17
22)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	19
23)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	19
24)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	20

(b)

A positive sign (+) at location (i, j) indicates that the median of the method i is significantly lower than the median of method j at the 95% confidence level. A negative sign (-) indicates the opposite, and an equal sign (=) indicates that there is no significant difference between the two methods. The score in the last column indicates the number of methods with respect to which the corresponding method has been considered significantly better.

on the Milan portrait data set containing multiple targets respectively reported in (Tables 6b and 7b) it is possible to see that the highest scores are clustered in the lower part of the tables. The proposed adaptive skin patch using a real face detector is statistically the best algorithm followed by the proposed skin patch independently from the face detector used (in both global and spatially varying versions). The differences among all these algorithms are not perceptually significant and all rank at the first place on a tie.

TABLE 7

Results of the Perceptual Significance Test on the Milan Portrait Data Set: Single Target (a), Multiple Targets (b)

	1) DN	2) WP	3) GW	4) SoG	5) gGW	6) GE1	7) GE2	8) GM	9) NIS	10) HLVI BU	11) HLVI TD	12) HLVI BU & TD	13) SS ML	14) FACE _{ideal} GM	15) FACE _{ideal} SP	16) FACE _{real} GM + GW	17) FACE _{real} SP + GW	Score
1)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	0
2)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	1
3)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	3
4)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	10
5)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	3
6)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	5
7)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	2
8)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	7
9)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	3
10)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	9
11)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	7
12)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	7
13)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	12
14)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	15
15)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	13
16)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	13
17)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	13

(a)

	1) DN	2) WP	3) GW	4) SoG	5) gGW	6) GE1	7) GE2	8) SS ML	9) LSAC	10) RETINEX	11) MLS + WP	12) MLS + GW	13) Fusion Grad. Tree Boost.	14) Fusion Rand. Forest Regr.	15) FACE _{ideal} GM glo	16) FACE _{ideal} GM SV	17) FACE _{ideal} SP glo	18) FACE _{ideal} SP SV	19) FACE _{real} GM gbo+ GW	20) FACE _{real} GM SV + GW	21) ADA FACE _{real} GM + GW	22) FACE _{real} SP glo+ GW	23) FACE _{real} SP SV + GW	24) ADA FACE _{real} SP + GW	Score
1)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	0
2)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	1
3)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	5
4)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	4
5)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	2
6)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	4
7)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	3
8)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	13
9)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	5
10)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	5
11)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	10
12)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	10
13)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	4
14)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	10
15)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	14
16)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	14
17)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	19
18)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	19
19)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	14
20)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	14
21)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	15
22)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	19
23)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	19
24)	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	19

(b)

A positive sign (+) at location (i, j) indicates that the method i is perceptually better than the method j. A negative sign (-) indicates the opposite, and an equal sign (=) indicates that there is no perceptual difference between the two methods. The score in the last column indicates the number of methods with respect to which the corresponding method has been considered perceptually better.

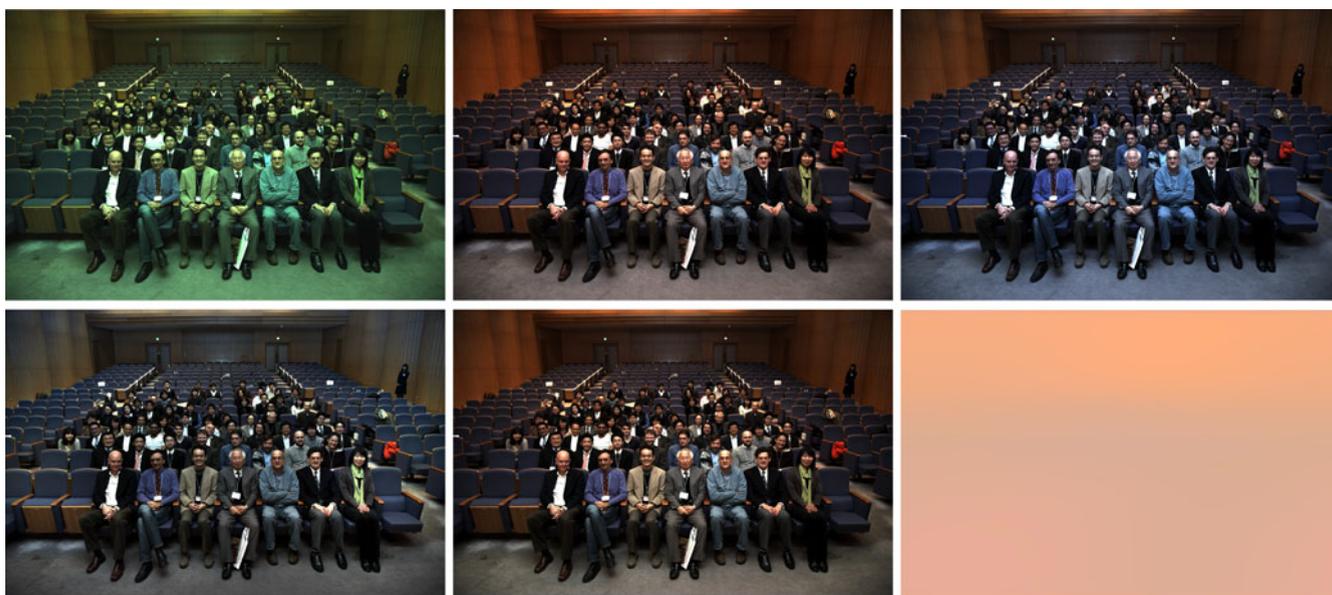


Fig. 9. Example on real-world RAW image. First row: original RAW image, correction balanced for the first rows, correction balanced for the furthest rows. Second row: correction with LSAC, local correction with the adaptive skin patch algorithm and its estimated illuminant color map (balanced with the camera daylight multipliers for better visualization).

flash was able to light only the first rows of seats. The first row of images reported represents the original RAW image, the ideal correction for the first rows, and the ideal correction for the furthest rows respectively. The second row of images represents the correction with the LSAC algorithm, the correction with the adaptive skin patch algorithm, and the estimated illuminant color map (which has been corrected with the camera daylight multipliers for better visualization).

In order to be able to correct the images for which we do not have the RAW camera file of the images, as for example for the images that we could find on the web, we could proceed in two different ways. On one side we could estimate

the RAW image from its sRGB version by knowing the camera model and using the method described in [42]. On the other side we could re-train our method to work directly in sRGB. The latter approach has been used here. The results on a couple of images taken from the web are reported in Fig. 10. In the first column the original image is reported, followed by the image corrected with the adaptive skin patch algorithm and by the estimated illuminant color map. Both the images are considered having a non-uniform illumination by our adaptive method (the maximum distance among the local illuminant estimates are about 7.3° and 30.4° within each image respectively) and thus the local illuminant estimation and correction are applied. As a comparison, the

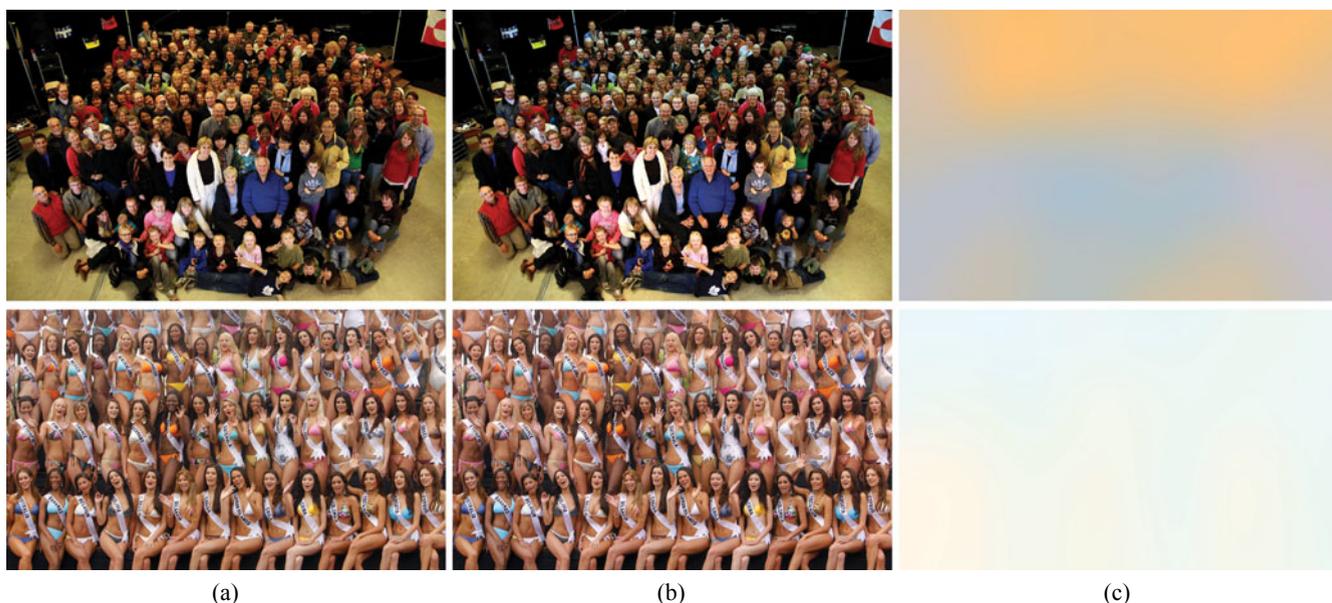


Fig. 10. Examples on real-world sRGB images taken from the web. Original image (left), local correction with adaptive skin patch algorithm (middle), and estimated illuminant color maps (right).

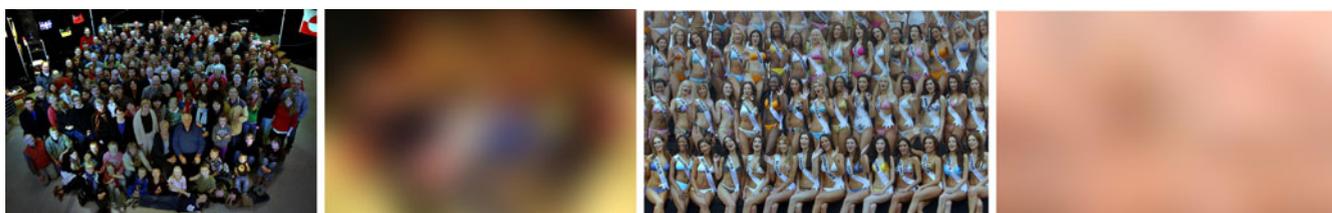


Fig. 11. Example of correction with LSAC for the same images reported in Fig. 10 and the corresponding estimated illuminant color maps.

images corrected with the LSAC and the corresponding estimated illuminant color maps are reported in Fig. 11. Although the color appearance of the examples images reported in Figs. 9 and 10 could be question of debate, it can be observed that the effects of the different light sources are less visible in the images corrected with our method.

6 CONCLUSION

In this paper we have designed an adaptive color constancy algorithm that exploiting the skin regions found in faces is able to estimate and correct the scene illumination. The algorithm automatically switches from global to spatially varying color correction on the basis of the illuminant estimations on the different faces detected in the image. When no faces are detected any other algorithm in the state of the art can be applied. An extensive comparison with both global and local color constancy algorithms has proved the effectiveness of the proposed algorithm in terms of both statistical and perceptual significance on a large heterogeneous data set of RAW images. To the best of our knowledge this is the first algorithm that automatically modifies its behavior from global to local on the basis of the image content. The proposed method can also be applied to sRGB images.

The integration of our algorithm in actual digital still camera processing pipelines is straightforward since many cameras have already embedded a face detector working on gray level images. With respect to our preliminary work that exploited skin-based gamut mapping, we have also designed a new illuminant estimation algorithm, also based on skin, that is more efficient and therefore more suitable for resource-limited camera devices, such as consumer digital cameras and camera phones.

Our method could be easily extended to also use reference objects that have intrinsic colors. In fact, the computational procedure adopted here to investigate if skin color is a valid cue for color constancy could be applied to any single color or colors combination.

Our spatially varying color correction assumes that the illumination on each face is constant and that the illuminants estimated on the faces properly sample the illumination distribution in the acquired scene. In fact, the illuminant estimated on each face is propagated to the rest of the image. There are therefore cases in which our algorithm can fail. For example, if some areas are not properly sampled or there is a light source behind a subject, the resulting image could be miscorrected. To deal with this types of situations, we plan to investigate the integration of our algorithm with other spatially varying algorithms or with algorithms that estimate the scene geometry, e.g., [46]. To relax the assumption of uniform illumination on each face it could be investigated if a generic or specific [49] 3D face shape prior can be used in the

shape-from-shading framework proposed by Huang and Smith [50]. A further hint for having information about lighting conditions could be obtained by eye reflections [51], [52], [53] as suggested by a reviewer.

ACKNOWLEDGMENTS

The authors would like to thank Prof. Hwann-Tzong Chen for making available the RAW files of the Taiwan subset. The source code and the Milan portrait data set are available at <http://www.ivl.disco.unimib.it/research/FaceColorConstancy/>.

REFERENCES

- [1] K. Barnard, "Improvements to Gamut Mapping Colour Constancy Algorithms," *Proc. European Conf. Computer Vision*, pp. 390-403, 2000.
- [2] S. Bianco, G. Ciocca, C. Cusano, and R. Schettini, "Improving Color Constancy Using Indoor-Outdoor Image Classification," *IEEE Trans. Image Processing*, vol. 17, no. 12, pp. 2381-2392, Dec. 2008.
- [3] S. Bianco, G. Ciocca, C. Cusano, and R. Schettini, "Automatic Color Constancy Algorithm Selection and Combination," *Pattern Recognition*, vol. 43, no. 3, pp. 695-705, 2010.
- [4] S. Bianco, F. Gasparini, and R. Schettini, "Consensus Based Framework for Illuminant Chromaticity Estimation," *J. Electronic Imaging*, vol. 17, no. 2, 2007.
- [5] S. Bianco and R. Schettini, "Color Constancy Using Faces," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 65-72, 2012.
- [6] S. Bianco, "Color Constancy Using Single Colors," *Proc. IEEE European Conf. Computer Vision Workshops and Demonstrations*, pp. 390-400, 2012.
- [7] G. Buchsbaum, "A Spatial Processor Model for Object Color Perception," *J. Franklin Inst.*, vol. 310, pp. 1-26, 1980.
- [8] V. Cardei, B. Funt, and K. Barnard, "White Point Estimation for Uncalibrated Images," *Proc. Seventh IS&T/SID Color Imaging Conf.*, pp. 97-100, 1999.
- [9] A. Chakrabarti, K. Hirakawa, and T. Zickler, "Color Constancy with Spatio-Spectral Statistics," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 34, no. 8, pp. 1509-1519, Aug. 2012, DOI:10.1109/TPAMI.2011.252.
- [10] G. Finlayson, S. Hordley, and P. Hubel, "Color by Correlation: A Simple, Unifying Framework for Color Constancy," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 23, no. 11, pp. 1209-1221, Nov. 2001.
- [11] D.A. Forsyth, "A Novel Algorithm for Color Constancy," *Int'l J. Computer Vision*, vol. 5, no. 1, pp. 5-36, 1990.
- [12] B. Funt, K. Barnard, and L. Martin, "Is Machine Colour Constancy Good Enough?" *Proc. European Conf. Computer Vision*, pp. 445-459, 1998.
- [13] F. Gasparini, S. Corchs, and R. Schettini, "Recall or Precision Oriented Strategies for Binary Classification of Skin Pixels," *J. Electronic Imaging*, vol. 17, p. 023017, 2008.
- [14] P. Gehler, C. Rother, A. Blake, T. Minka, and T. Sharp, "Bayesian Color Constancy Revisited," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 1-8, 2008.
- [15] A. Gijsenij and T. Gevers, "Color Constancy using Natural Image Statistics and Scene Semantics," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 4, no. 33, pp. 687-698, Apr. 2011.
- [16] A. Gijsenij, T. Gevers, and J. van de Weijer, "Generalized Gamut Mapping Using Image Derivative Structures for Color Constancy," *Int'l J. Computer Vision*, vol. 86, no. 2-3, pp. 127-139, 2010.

- [17] A. Gijsenij, T. Gevers, and J. van de Weijer, "Computational Color Constancy: Survey and Experiments," *IEEE Trans. Image Processing*, vol. 20, no. 9, pp. 2475-2489, Sept. 2011.
- [18] R. Gottumukkal and V. Asari, "Skin Color Constancy for Illumination Invariant Skin Segmentation," *Proc. SPIE-IS&T Electronic Imaging*, pp. 969-976, 2005.
- [19] C.I. Group, "Camera Phone Image Quality Phase 1: Fundamentals and Review of Considered Test Methods," *IBA CPIQ Phase 1 Document Repository*, pp. 1-82, 2007.
- [20] T. Hansen, M. Olkkonen, S. Walter, and K.R. Gegenfurtner, "Memory Modulates Color Appearance," *Nature Neuroscience*, vol. 9, pp. 1367-1368, 2006.
- [21] S. Hordley and G. Finlayson, "Re-Evaluating Color Constancy Algorithms," *Proc. Int'l Conf. Pattern Recognition*, pp. 76-79, 2004.
- [22] S.D. Hordley, "Scene Illuminant Estimation: Past, Present, and Future," *Color Research & Application*, vol. 31, no. 4, pp. 303-314, 2006.
- [23] ISO, "Graphic Technology-Standard Object Colour Spectra Database for Colour Reproduction Evaluation (Socs)," Technical Report ISO/TR 16066:2003(E), Int'l Organization for Standardization, 2003.
- [24] P. Kakumanu, S. Makrogiannis, and N. Bourbakis, "A Survey of Skin-Color Modeling and Detection Methods," *Pattern Recognition*, vol. 40, pp. 1106-1122, 2007.
- [25] A. Moreno, B. Fernando, B. Kani, S. Saha, and S. Karaoglu, "Color Correction: A Novel Weighted von Kries Model Based on Memory Colors," *Proc. Computational Color Imaging Workshop*, pp. 165-175, 2011.
- [26] H. Nachlieli, R. Bergman, D. Greig, C. Staelin, B. Oicherman, G. Ruckenstein, and D. Shaked, "Skin-Sensitive Automatic Color Correction," HP Technical Report: HPL-2009-13 pp. 1-13, HP Laboratories, 2009.
- [27] L. Shi and B.V. Funt, "Re-Processed Version of the Gehler Color Constancy Database of 568 Images," <http://www.cs.sfu.ca/colour/data/>, July, 2013.
- [28] M. Soriano, B. Martinkauppi, S. Huovinen, and M. Laaksonen, "Using the Skin Locus to Cope with Changing Illumination Conditions in Color-Based Face Tracking," *Proc. IEEE Nordic Signal Processing Symp.*, pp. 383-386, 2000.
- [29] J. van de Weijer, T. Gevers, and A. Gijsenij, "Edge-Based Color Constancy," *IEEE Trans. Image Processing*, vol. 16, no. 9, pp. 2207-2214, Sept. 2007.
- [30] J. van de Weijer, C. Schmid, and J. Verbeek, "Using High-Level Visual Information for Color Constancy," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 1-8, 2007.
- [31] P. Viola and M. Jones, "Rapid Object Detection Using a Boosted Cascade of Simple Features," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 511-518, 2001.
- [32] J. von Kries, "Chromatic Adaptation," *Festschrift der Albrecht-Ludwigs-Universität, MIT Press*, 1902 [Translation: D.L. MacAdam, Colorimetry-Fundamentals, SPIE Milestone Series, vol. MS 77, 1993].
- [33] A. Witkin, "Scale-Space Filtering: A New Approach to Multi-Scale Description," *Proc. IEEE Int'l Conf. Acoustics, Speech, and Signal Processing*, vol. 9, pp. 150-153, 1984.
- [34] M. Ebner, "Color Constancy Based on Local Space Average Color," *Machine Vision and Applications*, vol. 20, no. 5, pp. 283-301, 2009.
- [35] E.H. Land, "The Retinex Theory of Color Vision," *Scientific Am.*, vol. 237, no. 6, pp. 108-128, 1977.
- [36] A. Gijsenij, R. Lu, and T. Gevers, "Color Constancy for Multiple Light Sources," *IEEE Trans. Image Processing*, vol. 21, no. 2, pp. 697-707, Feb 2012.
- [37] J. Montojo, "Face-Based Chromatic Adaptation for Tagged Photo Collections," 2009.
- [38] S. Crichton, J. Pichat, M. Mackiewicz, G.-Y. Tian, and A. Hurlbert, "Skin Chromaticity Gamuts for Illumination Recovery," *Proc. Sixth European Conf. Colour in Graphics, Imaging, and Vision*, pp. 266-271, 2012.
- [39] M. Bleier, C. Riess, S. Beigpour, E. Eibenberger, E. Angelopoulou, T. Troger, and A. Kaup, "Color Constancy and Non-Uniform Illumination: Can Existing Algorithms Work?" *Proc. IEEE Int'l Conf. Computer Vision Workshops*, pp. 774-81, 2011.
- [40] T.J. de Carvalho, C. Riess, E. Angelopoulou, H. Pedrini, and A. de Rezende Rocha, "Exposing Digital Image Forgeries by Illumination Color Classification," *IEEE Trans. Information Forensics and Security*, vol. 8, no. 7, pp. 1182-1194, July 2013.
- [41] J. Marguier, N. Bhatti, H. Baker, M. Harville, and S. Süsstrunk, "Assessing Human Skin Color from Uncalibrated Images," *Int'l J. Imaging, Systems and Technology*, vol. 17, no. 3, pp. 143-151, 2007.
- [42] S.J. Kim, H.T. Lin, Z. Lu, S. Süsstrunk, S. Lin, and M.S. Brown, "A New In-Camera Model for Color Computer Vision and Its Application," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 34, no. 12, pp. 2289-2302, Dec 2012.
- [43] G.D. Finlayson, M.S. Drew, and B.V. Funt, "Color Constancy: Generalized Diagonal Transforms Suffice," *J. Optical Soc. of Am. A*, vol. 11, no. 11, pp. 3011-3019, 2004.
- [44] F. Wilcoxon, "Individual Comparisons by Ranking Methods," *Biometrics*, vol. 1, pp. 80-83, 1945.
- [45] A. Gijsenij, T. Gevers, and M.P. Lucassen, "Perceptual Analysis of Distance Measures for Color Constancy Algorithms," *J. Optical Soc. of Am. A*, vol. 26, no. 10, pp. 2243-2256, 2009.
- [46] R. Lu, A. Gijsenij, T. Gevers, D. Xu, V. Nedovic, and J.M. Geusebroek, "Color Constancy Using 3D Scene Geometry," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 1749-1456, 2009.
- [47] E. Hsu, T. Mertens, S. Paris, S. Avidan, and F. Durand, "Light Mixture Estimation for Spatially Varying White Balance," *ACM Trans. Graphics*, vol. 27, no. 3, article 70, 2008.
- [48] I. Boyadzhiev, K. Bala, S. Paris, and F. Durand, "User-Guided White Balance for Mixed Lighting Conditions," *ACM Trans. Graphics*, vol. 31, no. 6, article 200, 2012.
- [49] I. Kemelmacher-Shlizerman and S.M. Seitz, "Face Reconstruction in the Wild," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 1746-1753, 2011.
- [50] R. Huang and W.A.P. Smith, "Shape-from-Shading under Complex Natural Illumination," *Proc. IEEE Int'l Conf. Image Processing*, pp. 13-16, 2011.
- [51] K. Nishino and S. Nayar, "Eyes for Relighting," *ACM Trans. Graphics*, vol. 23, no. 3, pp. 704-711, 2004.
- [52] K. Nishino and S. Nayar, "The World in an Eye [Eye Image Interpretation]," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 444-451, 2004.
- [53] K. Nishino, P. Belhumeur, and S. Nayar, "Using Eye Reflections for Face Recognition under Varying Illumination," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 519-526, 2005.
- [54] T. Gevers, A. Gijsenij, J. van de Weijer, and J.-M. Geusebroek, *Color in Computer Vision: Fundamentals and Applications*. Wiley, 2012.
- [55] R. Ramanath, W.E. Snyder, Y. Yoo, and M.S. Drew, "Color Image Processing Pipeline," *Signal Processing Magazine*, vol. 22, no. 1, pp. 34-43, 2005.



Simone Bianco received the BSc and the MSc degrees in mathematics from the University of Milano-Bicocca, Italy, in 2003 and 2006, respectively. He received the PhD degree in computer science from the Department of Informatics, Systems and Communication of the University of Milano-Bicocca, Italy, in 2010, where he is currently a postdoctoral researcher. His research interests include computer vision, optimization algorithms, and color imaging. He is a member of the IEEE.



Raimondo Schettini is a full professor at the University of Milano-Bicocca (Italy), and head of Imaging and Vision Lab (www.ivl.disco.unimib.it). He has published more than 250 refereed papers and six patents about color reproduction, and image processing, analysis and classification. He is an IAPR fellow for his contributions to pattern recognition research and color image analysis. He is a member of the IEEE.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.