

Color Constancy for Multiple Light Sources

Arjan Gijsenij, *Member, IEEE*, Rui Lu, and Theo Gevers, *Member, IEEE*

Abstract—Color constancy algorithms are generally based on the simplifying assumption that the spectral distribution of a light source is uniform across scenes. However, in reality, this assumption is often violated due to the presence of multiple light sources. In this paper, we will address more realistic scenarios where the uniform light-source assumption is too restrictive. First, a methodology is proposed to extend existing algorithms by applying color constancy locally to image patches, rather than globally to the entire image. After local (patch-based) illuminant estimation, these estimates are combined into more robust estimations, and a local correction is applied based on a modified diagonal model. Quantitative and qualitative experiments on spectral and real images show that the proposed methodology reduces the influence of two light sources simultaneously present in one scene. If the chromatic difference between these two illuminants is more than 1° , the proposed framework outperforms algorithms based on the uniform light-source assumption (with error-reduction up to approximately 30%). Otherwise, when the chromatic difference is less than 1° and the scene can be considered to contain one (approximately) uniform light source, the performance of the proposed method framework is similar to global color constancy methods.

Index Terms—Color constancy, computer vision, illuminant estimation.

I. INTRODUCTION

THE COLOR of a light source has a significant influence on object colors in the scene [1]. Therefore, the same object, taken by the same camera but under different illumination, may vary in its measured color values. This color variation may introduce undesirable effects in digital images. Moreover, it may negatively affect the performance of computer vision methods for different applications such as object recognition, tracking, and surveillance. The aim of color constancy is to correct for the effect of the illuminant color, either by computing invariant features, e.g., [2] and [3], or by transforming the input image such that the effects of the color of the light source are removed, e.g., [1] and [4]. This paper focuses on the latter definition, which is also called *white balancing*.

A considerable number of color constancy algorithms are proposed; see [1], [4], and [5] for reviews. In general, pixel



Fig. 1. Scenes with multiple different light sources, taken from the web.

values are exploited to estimate the color of the light source. Examples of such methods include approaches based on low-level features [6], [7], gamut-based algorithms [8], and other methods that use knowledge acquired in a learning phase [9]. Alternatively, methods that use derivatives (i.e., edges) and even higher order statistics are proposed [10]. These color constancy algorithms are based on the assumption that the light source across the scene is spectrally uniform. However, this assumption is often violated as there might be more than one light source illuminating the scene. For instance, indoor scenes could be affected by both indoor *and* outdoor illumination, each having distinct spectral power distributions. Moreover, interreflections can lead to the introduction of additional “virtual” light sources, which have similar effects as true light sources and could be therefore handled as such. For outdoor scenes, the shadowed regions are dominated by skylight, whereas nonshadowed regions are illuminated by a combination of skylight and sunlight. Examples of such scenarios are shown in Fig. 1.

Retinex [7] is one of the first color constancy methods developed and assumes 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. Numerous implementations have been proposed, e.g., using very large scale integration for real-time image processing [11], using center/surround for practical image processing applications [12], [13], or using MATLAB to standardize evaluation of the Retinex [14]. Moreover, various extensions have been proposed by adding additional knowledge about the scene into the method. For instance, Finlayson et al. [15] and Barnard et al. [16] propose a Retinex-based method that identifies and uses surface colors that are illuminated by two different light sources. Xiong and Funt [17] constructed an extension that uses stereo images to derive 3-D information on the surfaces that are present in the image. This information is used to more accurately distinguish material transitions from light color changes, but the stereo information is often not available and is not trivial to obtain. Ebner [18] also proposed a method that is based on the assumption that the illuminant transition is smooth. This method uses the local space average color (LSAC) for local estimation of the illuminant by convolving the image with a kernel function (e.g., a Gaussian or Exponential kernel). However, all these methods are based on the assumption that the illuminant transition is smooth, which

Manuscript received December 23, 2010; revised April 27, 2011 and July 31, 2011; accepted August 02, 2011. Date of publication August 18, 2011; date of current version January 18, 2012. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Jose M. Bioucas-Dias.

A. Gijsenij is with Alten PTS, 2909 LE Capelle a/d IJssel, The Netherlands (e-mail: arjan.gijsenij@gmail.com).

R. Lu is with the School of Computer and Information Technology, Beijing Jiaotong University, Beijing 100044, China.

T. Gevers is with the University of Amsterdam, 1098 XH Amsterdam, The Netherlands, and also with the Computer Vision Center, Universitat Autònoma de Barcelona, 08193 Bellaterra, Spain (e-mail: th.gevers@uva.nl).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2011.2165219

is often not the case (see for example the two leftmost images in Fig. 1). For such situations, Ebner proposes to use nonuniform averaging [19].

Other algorithms that consider multiple light sources include physics-based methods [20], biologically inspired models [21], and methods requiring manual intervention [22]. The first method [20] is specifically designed for outdoor images and distinguishes between shadow and nonshadow regions. Various other methods that distinguish between shadow and nonshadow regions have been proposed, e.g., [23] and [24], but such methods do not result in output images that have any visual similarity to the original input image. The second method [21] is based on retinal mechanisms and adaptation, simulating the properties of opponent and double-opponent cells. The latter method [22] requires spatial locations in the image that are illuminated by different light sources to be manually specified by a user.

In this paper, a new methodology is presented that enables color constancy under multiple light sources. The methodology is designed according to the following criteria: 1) it should be able to deal with scenes containing multiple light sources; 2) it should work on a single image; 3) no human intervention is required; and 4) no prior knowledge or restrictions on the spectral distributions of the light sources is required. Although the proposed framework is designed to handle multiple light sources, the focus in this paper is on scenes captured under *one or two distinct light sources* (including linear mixtures of two light sources), arguably the two most common scenarios in real-world images. Furthermore, not only images recorded under multiple light sources but also images that are recorded under only one light source should be properly processed. Hence, the improvement on multiple-light-source scenes should not be obtained at the expense of a decreased performance on single-light-source scenes.

To construct color constant images from scenes that are recorded under multiple sources, the proposed methodology makes use of local image patches, rather than the entire image. These image patches are assumed to have (local) uniform spectral illumination and can be selected by any sampling method. In this paper, grid-based sampling, key-point-based sampling, and segmentation-based sampling are evaluated. After sampling of the patches, illuminant estimation techniques are applied to obtain local illuminant estimates, and these estimates are combined into more robust estimations. This combination of local estimates is done with two different approaches, i.e., clustering and segmentation. The first approach is to cluster the illuminant estimates, taking the cluster centers as final illuminant estimate for each of the regions. The second approach is to take spatial relations between local estimates into account by applying segmentation on the back-projected local illuminant estimations. Finally, when the resulting illuminant is estimated, a modified diagonal transform is applied to obtain the color-corrected images.

This paper is organized as follows: First, in Section II, color constancy is discussed. Next, in Section III, the proposed methodology is explained in detail. Experiments are described in Section IV, and Section V presents a discussion of the

obtained results and some directions for future work in this line of research.

II. COLOR CONSTANCY

In general, the goal of computational color constancy is to estimate the chromaticity of the light source and then to correct the image to a canonical illumination using the diagonal model. Here, we will briefly outline this process.

A. Reflection Model

Image color $\mathbf{I} = (I_R, I_G, I_B)^T$ for a Lambertian surface at location \mathbf{x} can be modeled as

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

where $c \in \{R, G, B\}$ and $E(\lambda, \mathbf{x})$, $S(\lambda)$ and $\rho_c(\lambda, \mathbf{x})$ are the illuminant spectrum distribution, surface reflectance, and camera sensitivity, respectively. Furthermore, ω is the visible spectrum. Then, for a given location \mathbf{x} , the color of the light source can be computed as follows:

$$\mathbf{L}(\mathbf{x}) = \begin{pmatrix} L_R(\mathbf{x}) \\ L_G(\mathbf{x}) \\ L_B(\mathbf{x}) \end{pmatrix} = \int_{\omega} E(\lambda, \mathbf{x}) \boldsymbol{\rho}(\lambda) d\lambda \quad (2)$$

where it should be noted that, typically, color constancy is involved with estimating the chromaticity of the light source (i.e., intensity information is not recovered). Estimating this chromaticity from a single image is an underconstrained problem as both $E(\lambda, \mathbf{x})$ and $\boldsymbol{\rho}(\lambda) = (\rho_R, \rho_G, \rho_B)^T$ are unknown. Therefore, assumptions are imposed on the imaging conditions. Typically, assumptions are made about statistical properties of the illuminants or surface reflectance properties. Moreover, most color constancy algorithms are based on the assumption that the illumination is uniform across the scene (i.e., $E(\lambda, \mathbf{x}) = E(\lambda)$). However, for real-world scenes, this assumption is very restrictive and often violated.

B. Illuminant Estimation: One Light Source

Most color constancy algorithms proposed are based on the assumption that the color of the light source is uniform across the scene. For instance, the white-patch algorithm [7] is based on the assumption that the maximum response in a scene is white, and gray-world algorithm [6] is based on the assumption that the average color in a scene is achromatic. These assumptions are then used to make a global estimate of the light source and correspondingly correct the images.

The framework proposed in [10] allows for systematically generating color constancy as follows:

$$\left(\int \left| \frac{\partial^n I_{c,\sigma}(\mathbf{x})}{\partial \mathbf{x}} \right|^p d\mathbf{x} \right)^{\frac{1}{p}} = k L_c^{n,p,\sigma} \quad (3)$$

where $L_c^{n,p,\sigma}$ is used to denote different instantiations of the framework. Furthermore, $\|\cdot\|$ is the Frobenius norm, $c = \{R, G, B\}$, n is the order of the derivative, p is the

Minkowski norm, and $I_{c,\sigma} = I_c \otimes G_\sigma$ is the convolution of the image with a Gaussian filter with scale parameter σ . According to the characteristics of the Gaussian filter, the derivative can be further described by

$$\frac{\partial^{a+b} I_{c,\sigma}}{\partial x^a \partial y^b} = I_c * \frac{\partial^{a+b} G_\sigma}{\partial x^a \partial y^b} \quad (4)$$

where $*$ denotes the convolution and $a + b = n$.

Using (3), many different color constancy algorithms can be derived by varying one or more parameters (i.e., n , p , and/or σ). In this paper, we focus on the following methods [10]:

- Pixel-based color constancy algorithms (i.e., $n = 0$). Minkowski norm p and smoothing parameter σ are dependent on the images in the data set; thus, in this paper we will show results for three variations, i.e., the gray-world algorithm (with $n = 0$, the Minkowski-norm $p = 1$ and the smoothing filter size $\sigma = 0$, i.e., $\mathbf{L}^{0,1,0}$), the white-patch algorithm (with Minkowski-norm $p = \infty$, i.e., $\mathbf{L}^{0,\infty,0}$), and one specific instantiation of the general gray-world algorithm, i.e., $\mathbf{L}^{0,8,1}$;
- Higher order color constancy methods (i.e., $n = 1$ and $n = 2$), resulting in the first-order gray-edge ($\mathbf{L}^{1,1,1}$) and the 2nd-order gray-edge ($\mathbf{L}^{2,1,1}$).

C. Correction: Diagonal Model

After the color of the light source is estimated, the aim is to transform the input images, taken under an unknown light source, into colors as if they appear under a canonical light source. Usually, this is done using the diagonal model or *von Kries model*, i.e.,

$$\mathbf{I}^c = \Lambda^{u,c} \mathbf{I}^u \quad (5)$$

where \mathbf{I}^u is the image taken under an unknown light source while \mathbf{I}^c is the image transformed, which appears as if it is taken under the canonical illuminant. $\Lambda^{u,c}$ is the mapping diagonal matrix, i.e.,

$$\Lambda^{u,c} = \begin{pmatrix} \alpha & 0 & 0 \\ 0 & \beta & 0 \\ 0 & 0 & \gamma \end{pmatrix} = \begin{pmatrix} \frac{L_R^c}{L_R^u} & 0 & 0 \\ 0 & \frac{L_G^c}{L_G^u} & 0 \\ 0 & 0 & \frac{L_B^c}{L_B^u} \end{pmatrix} \quad (6)$$

where \mathbf{L}^u is the unknown light source and \mathbf{L}^c is the canonical light source.

III. COLOR CONSTANCY FOR MULTIPLE LIGHT SOURCES

The majority of color constancy methods are based on the assumption of spectrally uniform lighting (e.g., one light source or multiple identical light sources). However, in real-world scenes, this assumption is often violated as more than one light source, with different spectral distribution, is present. Here, a methodology is proposed that extends traditional methods to estimate N different light sources, where N is the number of light sources that affect an image. The outline of this process is shown in Fig. 2. The proposed framework consists of the following five steps.

- Step 1) *Sampling of image patches.* The first step is to sample P patches from the image. For each patch, estimation is computed of the light source valid for that patch. It is assumed that the color of the light source is uniform over each patch, which is a reasonable assumption in practice. Therefore, the patch size should be limited but of sufficient size to extract enough image properties to accurately estimate the light source. Different sampling strategies can be used, e.g., dense sampling, interest points, and segmentation. Dense or grid-based sampling and segmentation-based sampling ensure that the union of all patches covers the full image. Furthermore, segmentation-based sampling can result in boundaries between segments that naturally follows the boundary between light sources (as most segmentation algorithms are sensitive to changes in the illuminant). Grid-based sampling has the advantage that the patches contain varied amount of information, whereas patches that are selected using segmentation will generally contain similar colors (and hence less variation). Finally, key-point sampling is specifically suited for edge-based color constancy methods as the key points are located around edges and junctions. In this paper, key points are located using the Harris detector at multiple scales [25] (scales 1, 1.5, and 2), using a similar patch size as the grid-based sampling, whereas segmentation is performed using the c [26].
- Step 2) *Patch-based illuminant estimation.* As the illuminant for each patch is assumed to be spectrally uniform, traditional color constancy methods are applied on every patch to estimate the local illuminant. In this paper, several color constancy algorithms are used, as explained in Section II-B. Color constancy is applied on each patch resulting in one estimate of the illuminant chromaticity per patch.
- Step 3) *Combination of estimates.* Since there is only a limited amount of information available when using a relatively small patch for the estimation of the light source, this may introduce estimation errors. To overcome this lack of information, patches that are taken from parts of the image that are illuminated by the same light source are combined to form a larger patch (and consequently result in a more accurate estimate). Patches that are illuminated by the same light source are likely to vote for the same illuminant, as illustrated in Fig. 2. Assuming that the number of clusters is known (we assume $N = 2$ here), the chromaticities can be grouped together using any clustering algorithm.
- Step 4) *Back-projection of clusters.* After the different estimates are grouped together into N groups, the result can be back-projected onto the original image to identify the locations in the image that are illuminated by each of the estimated light sources. This results in an illuminant classification, where

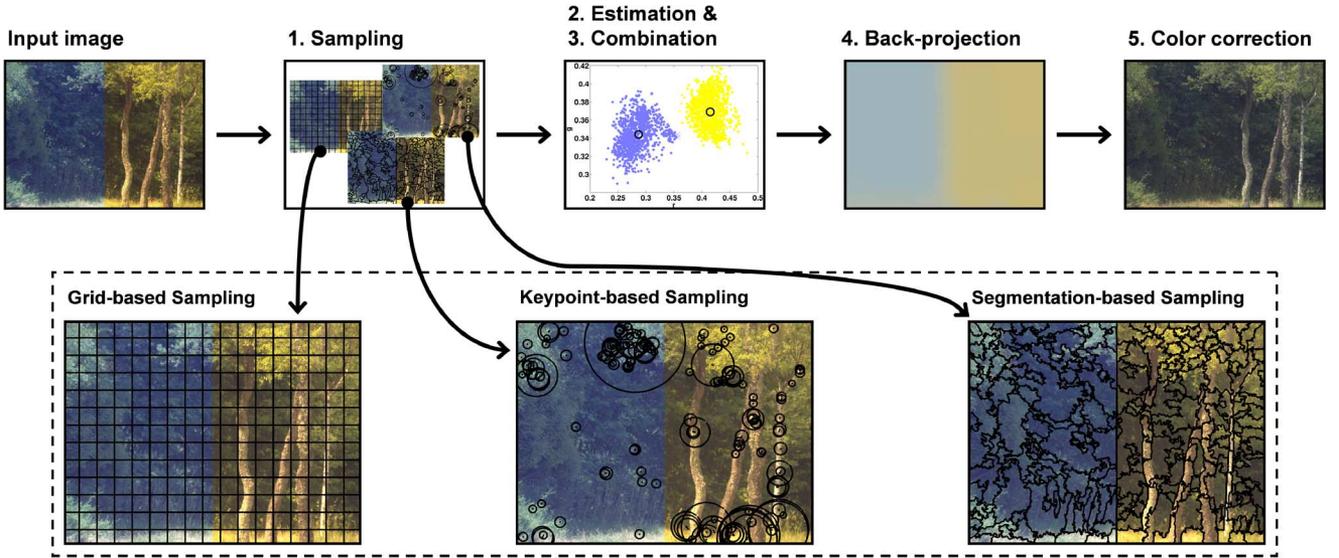


Fig. 2. Illustration of the proposed methodology. The input image, recorded under two different illuminants, is first sampled to obtain image patches. Different approach for sampling can be employed, i.e., grid-based sampling, key-point-based sampling (at multiple scales, i.e., 1, 1.5, and 2), and segmentation-based sampling (using graph-cut segmentation). Then, patchwise illuminant estimation is applied, and the obtained estimates are combined. Back-projecting the illuminant estimates onto the image results in a pixel-based illuminant estimate, which can be finally used to color correct the input image.

every pixel is assigned to one of the estimated light sources. After back-projection, a pixelwise illuminant estimate is obtained.

Step 5) *Color correction*. Finally, using the pixelwise estimates, the output image is constructed. Transforming the input image so that it appears to be taken under a white light source is an instantiation of chromatic adaptation, e.g., [27]. Many methods to do this exist, e.g., [28] and [29], but all assume that the color of the light source in the input image is known. Since the focus in this paper is to estimate the illuminant, the diagonal model, or von Kries model [30], is used, as described in Section II.

Note that steps 3 and 4 could be merged into one segmentation step, depending on the sampling strategy in step 1: If a pixelwise illuminant estimate is obtained, which covers the entire image, then the combination of local estimates can be done using any segmentation algorithm that naturally incorporates spatial information. This alleviates the requirement of presetting the number of light sources, at the expense of introducing more parameters into the methodology. In this paper, we will adopt both approaches.

Further note that the pipeline of traditional color constancy methods assuming a uniform illuminant is a special instance of the proposed algorithm by assuming that there are $N = 1$ light sources and sampling $P = 1$ patches, where the patch size is the same dimensions as the original image. In this situation, steps 3 and 4 are trivial and can be omitted, and step 5 reduces to the regular diagonal model.

A. Light-Source Estimators

Images in the proposed method are divided into patches, which are assumed to be small enough such that it is consistent with the uniform spectral assumption. For each patch, illuminant estimation is obtained by using a standard color constancy

algorithm (based on a uniform light source). For simplicity, although other color constancy methods can be used, we focus on the five instantiations described in Section II-B, which include pixel and derivative-based methods. Multiple light-source estimates can be simultaneously taken into account, but in this paper, the focus is on single estimates per patch. Since the used algorithms merely estimate the chromaticity of the light source, every estimate is normalized for intensity, i.e.,

$$\begin{cases} r = R/(R + G + B) \\ g = G/(R + G + B). \end{cases} \quad (7)$$

The illuminant over each patch is represented by a 1×2 vector. The theoretical situation could occur where the illuminant estimate results in a black light source, i.e., $R = G = B = 0$. In such situations, albeit unlikely, we propose to use a white-light source, i.e., $r = g = 1/3$.

B. Overlapping Light Sources

The underlying assumption of the proposed framework is that the different light sources are locally constant. For instance, in outdoor images, the two light sources are light from the blue sky and direct sunlight, which can be clearly distinguished by the shadow boundary. However, for some images, the boundary between the different light sources might not be so obvious. In these cases, the true light source at locations in between two light sources is a linear mixture of those two light sources. Enforcing a sharp distinction will result not only in quantitative errors but also in severe qualitative errors. To overcome this issue, step 4 of the proposed method can be augmented with a filtering of the back-projected clusters to smooth the transition from one light source to another.

More formally, let $d_j(\mathbf{x})$ denote the chromatic distance (Euclidean distance is used throughout this paper) of the estimated illuminant of the patch located at spatial coordinate \mathbf{x} in the image to the j th illuminant (where $j \in \{1, \dots, N\}$). First, from

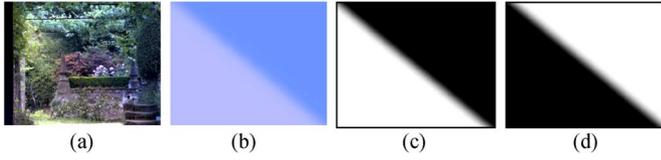


Fig. 3. Example of light mask maps. (a) There are two light sources in the synthetic image. The illumination for each pixel is present in (b). (c)–(d) Illuminant mask maps: The contribution each light source makes to the light mixture in which white means a big ratio, whereas dark, a small ratio.

distance $d_j(\mathbf{x})$, we compute $d'_j(\mathbf{x})$, which denotes the similarity between the estimated illuminant of the patch and the j^{th} illuminant, i.e.,

$$d'_j(\mathbf{x}) = \frac{\sum_{\mathbf{x}} d_j(\mathbf{x})}{d_j(\mathbf{x})}. \quad (8)$$

Then, mask map $m_j(\mathbf{x})$, indicating the estimated probability of the j^{th} illuminant, is defined as the ratio of $d'_j(\mathbf{x})$ to the sum of distances to all illuminants as

$$m_j(\mathbf{x}) = \frac{d'_j(\mathbf{x})}{\sum_{k=1}^N d'_k(\mathbf{x})}. \quad (9)$$

Filters are applied on this mask map to get the smooth illuminant distributions. In this paper, both linear and nonlinear filters are considered. Specifically, Gaussian and median filters are used. A Gaussian filter takes spatial information into consideration so that a pixelwise probability can be computed of the range of the estimated light sources. The advantage of the median filter is that it preserves the edges, which could be suited as some scenes can have sharp transitions from one light source to another. Such sharp transitions cannot be handled by Gaussian smoothing filters (in fact, using Gaussian smoothing filters implies an assumption of discontinuous illuminants).

C. Image Correction

Before image correction, the illumination for each pixel is estimated as follows:

$$L_e(\mathbf{x}) = \sum_{i=1}^N L_{e,i} m_i(\mathbf{x}) \quad (10)$$

where L_e is the illuminant estimation over the scene, $L_{e,i}$ is the estimation for the i^{th} illumination, and $m_i(\mathbf{x})$ is the contribution of the i^{th} light-source estimation to pixel \mathbf{x} . A linear mixture of light sources is controlled by this variable: A higher value of m_i indicates a larger influence of the i^{th} light source on this pixel ($m_i(\mathbf{x}) = 1$ indicates that pixel \mathbf{x} is solely illuminated by light source i). As shown in Fig. 3, the light mask map is of the same size as the input image.

IV. EXPERIMENTS

This section presents the experimental validation of the proposed method. First, the influence of several parameters on the performance of the proposed method is studied using hyperspectral data. Then, the proposed method is applied on two data sets of real images.



Fig. 4. Example images: The first image is the original hyperspectral image, whereas the others are generated using two light sources on the original image. Note that all experiments are performed on linear images; gamma correction is applied only for improved visualization.

A. Performance Measurement

Two performance measures are used in this paper: The clustering performance is measured using misclassification rate, whereas the angular error is used to evaluate the performance of color constancy algorithms.

Misclassification Rate: Pixels are considered to be misclassified if they are illuminated by L_i , but the method assigns them to another illuminant. Given that there are N light sources, the clustering performance is measured by

$$\eta = \frac{1}{T} \sum_{i=1}^N S_i \quad (11)$$

where S_i is the number of misclassified pixels illuminated by the i^{th} light source, whereas T is the total number of pixels in the image. Note that, in order to compute η , all pixels are assigned to either of the light sources.

Angular Error: Given a pixel in an image, $\mathbf{L}_t(\mathbf{x})$ is the ground truth of the light source illuminating it, whereas $\mathbf{L}_e(\mathbf{x})$ is the corresponding estimation; then, the angular error ε is

$$\varepsilon(\mathbf{x}) = \cos^{-1} \left(\hat{\mathbf{L}}_t(\mathbf{x}) \cdot \hat{\mathbf{L}}_e(\mathbf{x}) \right) \quad (12)$$

where the hat indicates the normalized vector. As scenes are illuminated by varying illuminations, the angular error is computed pixel by pixel throughout the image. Then, the average angular error across the scene is considered as the measurement.

Statistical Significance: Evaluating color constancy methods on a range of images requires the need for a summarizing statistic. Hordley and Finlayson [31] showed that the median is an appropriate statistic in the field of color constancy since the angular error is angular errors over a large set of images is generally not normally distributed. Moreover, to test the significance of the difference between two distributions, we adopt the Wilcoxon sign test [31]. This test is based on the assumption that if two methods A and B have the same median error on a set of N images (the null hypothesis), then the number of images on which method A performs better than method B (denoted W) is as high as the number of images on which method B performs better than method A . In other words, if the null hypothesis is true, then W is binomially distributed. The Wilcoxon sign test is based on this statistic.

B. Hyperspectral Data

First experiments are performed on a data set that is generated using hyperspectral data taken by Foster *et al.* [32]. This data set consists of a mixture of rural and urban scenes. Images (the sizes of the images slightly vary but approximately

TABLE I

PERFORMANCE OF COLOR CONSTANCY ALGORITHMS COMPUTED FOR THE HYPERSPECTRAL DATA SET. THE COLOR CONSTANCY ALGORITHMS BASED ON A SINGLE LIGHT SOURCE ASSUMPTION ARE USED AS BASELINE (I.E., THE RELATIVE IMPROVEMENT DENOTED BETWEEN PARENTHESES ARE COMPUTED WITH RESPECT TO THESE BASELINES); THE LSAC METHOD IS THE METHOD PROPOSED BY [18]. DIFFERENT SAMPLING STRATEGIES FOR THE PROPOSED METHOD ARE EVALUATED, AND SEGMENTATION-BASED COMBINATION (WITH NO ASSUMPTIONS ON THE NUMBER OF LIGHT SOURCES) IS APPLIED TO GRID-BASED SAMPLING

Method		Median	Method		Median
Do Nothing (DN)		26.3°	general Grey-World (gGW)		11.5°
Grey-World (GW)		14.0°	1 st -order Grey-Edge (GE-1)		16.8°
White-Patch (WP)		12.5°	2 nd -order Grey-Edge (GE-2)		16.7°
LSAC	Exponential filter (Impl. from [14])	10.5°	LSAC	Gaussian filter (Impl. from [34])	10.2°
Retinex		11.7°	Retinex		9.9°
Proposed: grid-based sampling, then clustering (k-means, k = 2)	GW	10.0° (-29%)	Proposed: keypoint-based sampling, then clustering (k-means, k = 2)	GW	10.9° (-22%)
	WP	9.5° (-24%)		WP	12.8° (+2%)
	GGW	9.9° (-14%)		GGW	12.1° (+5%)
	GE-1	15.4° (-8%)		GE-1	15.9° (-5%)
	GE-2	15.0° (-10%)		GE-2	15.4° (-8%)
Proposed: grid-based sampling, then segmentation (mean-shift)	GW	10.0° (-29%)	Proposed: segmentation-based sampling then clustering (k-means, k = 2)	GW	12.3° (-12%)
	WP	10.2° (-18%)		WP	10.1° (-19%)
	GGW	10.2° (-11%)		GGW	13.2° (+15%)
	GE-1	16.8° (-0%)		GE-1	13.0° (-23%)
	GE-2	16.1° (-4%)		GE-2	15.4° (-8%)

correspond to 335×255 pixels) are illuminated by two different light sources, randomly selected from a set of 81 illuminant spectra [33]. Together with the camera sensitivity functions specified by [33] and (1), the hyperspectral data are converted into (R, G, B) -values. Gaussian filtering is used to generate smooth transitions from one light source to the other, and different filter sizes are used. Using this approach, a data set with 1437 images is generated, containing a wide variety of combinations of light sources, with chromatic differences between the two light sources ranging from 0° to roughly 40° . Scenes with small chromatic differences between the two light sources are considered to be illuminated by one (approximate) uniform light source. Several examples are shown in Fig. 4. Note that the experiments on this data set are mainly used to systematically demonstrate the characteristics and behavior of the proposed algorithm.

Single Light Source: In order to validate how the proposed method improves the color constancy performance, we first perform experiments based on the assumption that there is only one light source illuminating the scene, while in fact there are two distinct light sources. The performance of five color constancy algorithms is measured using median angular errors shown in Table I. These results are considered as the baseline as they represent the performance of standard color constancy methods assuming a single light source.

Multiple Light Sources: The experiments on the hyperspectral data set are used to evaluate the effects of different parameter settings of the proposed method. In Table I, the results of the different sampling strategies are shown. It can be observed that grid-based sampling outperforms all other sampling strategies (except segmentation-based sampling combined with first-order gray-edge, which performs better than grid-based sampling). The reason for this is the combination of total coverage of the image by patches and the variety of image colors in every patch. The grid-based sampling has the advantage over keypoint-based sampling that the full image is covered by patches. Segmentation-based sampling obtains this same effect, but segmentation will group “similar” pixels into segments. The similarity of the pixels in one segment will result in local estimations

that are less accurate than the local estimations of the grid-based sampling. All results are obtained using k -means clustering. Differences with clustering based on Gaussian mixture modeling are small (not shown here). Furthermore, it is important to note that all differences between baseline methods and proposed extensions are statistically significant (at 95% confidence level), according to the Wilcoxon sign test [31].

An important parameter for the grid-based sampling is the patch size that is considered. Fig. 5 shows the influence of the patch size on the performance of the different color constancy algorithms. It can be derived that the performance roughly decreases as the patch size increases. Since a larger patch size logically implies more information, one might expect that an increased patch size would result in a better performance. However, a larger patch size includes more pixels; therefore, a misclassified patch at a large scale results in more misclassified pixels than a misclassification at a small scale. Even though the *probability* of misclassification decreases, the *penalty* of misclassification increases as a function of the patch size. Moreover, a large patch size implies that a patch covers a large illumination variation, hence decreasing the classification accuracy. Misclassification in this context means that pixels that are actually recorded under illuminant i are used to estimate illuminant j . Obviously, there is a tradeoff between the amount of information required for the local estimation of the illuminant on one side and the misclassification rate on the other. Generally, the trend for the different algorithms is the same.

Since the proposed method makes use of a clustering algorithm, the patch-based illuminant estimates should be easily separable. However, when an image is illuminated by two similar light sources, the patch-based illuminant estimates will be similar as well. Consequently, the clustering performance is influenced by the nature of the light sources. Fig. 6 (top row) shows the clustering performance in terms of misclassification rate as a function of the chromatic distance between the two light sources (note that the misclassification rate for all baseline algorithms is the same for all scenes. Since both light sources cover approximately half of the image, the misclassification rate when assuming that there is only one light source is 50%). As ex-

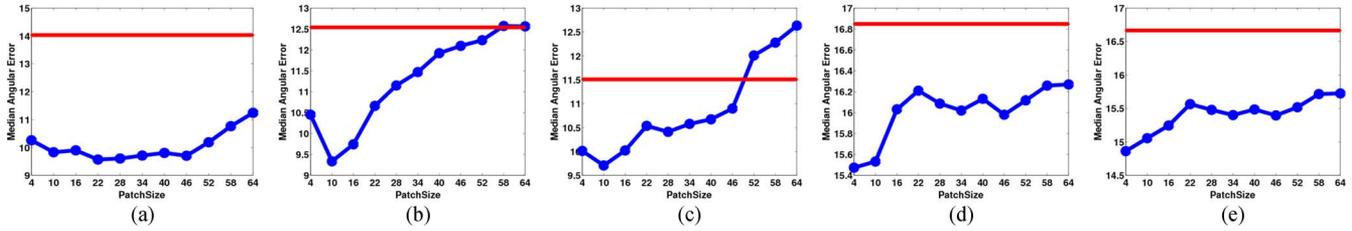


Fig. 5. Relationship between the patch size and the algorithm performance. (From left to right) The performance of gray world (GW), white patch (WP), general gray world (GGW), the first gray edge (GE-1), and the second gray edge (GE-2). Note that, in this experiment, no postprocessing is done. In each graph, the horizontal axis is the patch size, whereas the vertical axis is the median angular error. (Red line) The performance assuming only one light source is present (while in fact there are two).

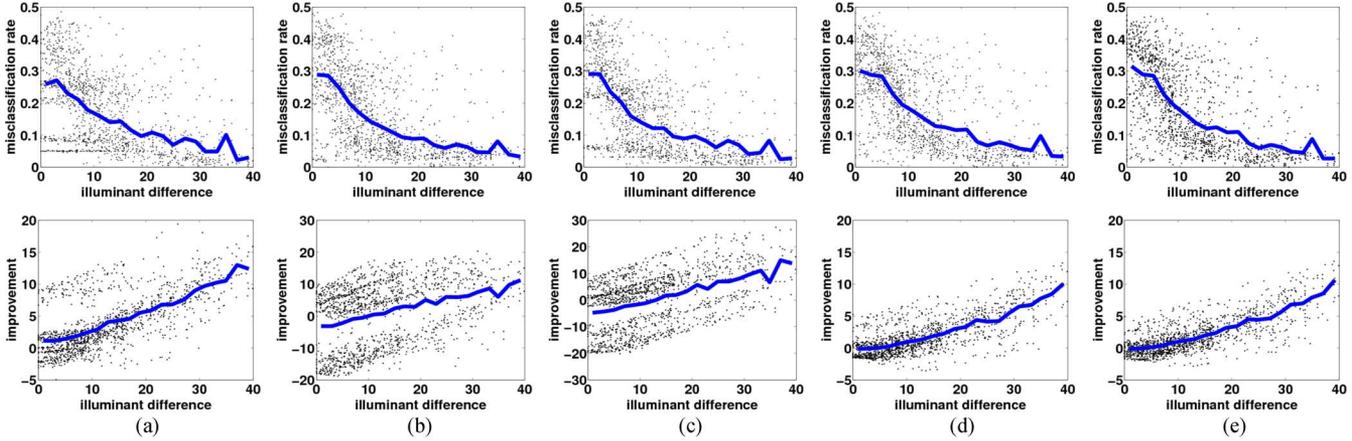


Fig. 6. Influence of illuminant differentiation. (Top row) The influence of the chromatic difference between the light sources for clustering. (Bottom row) The influence on the color constancy performance. No postprocessing is done in the experiment. Note that the blue lines are obtained by averaging over points with similar illumination difference and is superimposed to illustrate the trend.

TABLE II

PERFORMANCE OF COLOR CONSTANCY ALGORITHMS COMPUTED FOR A SUBSET OF THE HYPERSPECTRAL DATA SET CONTAINING IMAGES WITH ROUGHLY GLOBAL ILLUMINATION (I.E., THE ILLUMINATION DIFFERENCE WITHIN THE IMAGE IS LESS THAN 1°). THE COLOR CONSTANCY ALGORITHMS BASED ON A SINGLE LIGHT-SOURCE ASSUMPTION ARE USED AS BASELINE (I.E., THE RELATIVE IMPROVEMENT DENOTED BETWEEN PARENTHESES ARE COMPUTED WITH RESPECT TO THESE BASELINES). THE PROPOSED METHODOLOGY IS APPLIED USING GRID-BASED SAMPLING

Method	Median	Method	Median
Grey-World (GW)	8.9°	Proposed methodology using GW	$8.1^\circ (-9\%)$
White-Patch (WP)	7.9°	Proposed methodology using WP	$8.4^\circ (+6\%)$
general Grey-World (gGW)	5.5°	Proposed methodology using gGW	$8.1^\circ (+47\%)$
1 st -order Grey-Edge (GE-1)	15.1°	Proposed methodology using GE-1	$15.0^\circ (-1\%)$
2 nd -order Grey-Edge (GE-2)	14.9°	Proposed methodology using GE-2	$14.8^\circ (-1\%)$

pected, it can be observed that the misclassification decreases as the chromatic difference between light sources increases. Obviously, a higher misclassification rate is of influence on the color constancy performance: If more pixels are biased toward the wrong illuminant, then the overall estimation accuracy will diminish. The effect of the misclassification rate is shown in Fig. 6 (bottom row): Better illuminant estimates are achieved for scenes with a large chromatic difference between the two light sources. However, the proposed method performs (at least) as good on scenes with a small chromatic difference between the light sources (i.e., scenes with approximate uniform illumination) as the baseline algorithms. To illustrate this, we extracted a subset of images from the hyperspectral data set, consisting of images with roughly uniform illumination (i.e., images with an illuminant difference of less than 1° , a total of 74 images meet this requirement). In Table II, the performance of this subset is shown for both the algorithms assuming global illumination and the proposed methodology assuming two illuminants. It is

shown that the performance of the proposed methodology is similar to the algorithms assuming global illumination, except when the general gray world is used. In the other situations, the performance of the proposed methodology is similar to the baseline method (assuming only one illuminant is present). This important conclusion shows that the proposed method can, without loss of accuracy, be applied to images with one and two light sources present.

The proposed methodology can be also applied to images without the assumption that there are at most two distinct light sources in a scene. The local illuminant estimates can be obtained using mean-shift segmentation [35] on the patch-based illuminant estimates: Instead of clustering the estimates, the spatial relations between the different patches can be taken into account by applying segmentation, rather than clustering. The advantage of this approach with respect to clustering is that no additional information on the number of light sources is required. Results in Table I demonstrate that an increase in performance

with respect to the baseline is obtained even without the assumption of two different light sources. However, note that the performance of this approach is not as good as the clustering approach reported in Table I. This is explained by the additional knowledge that is available to the clustering-based approach: When applying clustering to the local estimates, the number of light sources is assumed to be fixed (and known). When using segmentation to combine the local illuminant estimates, this assumption is relaxed, and therefore, the performance slightly decreases with respect to clustering-based combination. Hence, the advantage of segmentation-based combination is that there are no assumptions on the number of light sources. However, this comes at the cost of introducing additional parameters to take into account and a slightly decreasing performance, although this approach still improves with respect to algorithms assuming a single uniform light source.

For comparison, two versions of the Retinex are evaluated, i.e., [13] and [14], as well as the color constancy algorithm based on the LSAC [18]. Note that both Retinex methods actually correct for more than just the chromaticity of the light source and do not explicitly use the diagonal model. In order to be still able to quantify the angular error, (5) is used on the input image and the output image of the Retinex-method to retrieve the diagonal model and hence extract the chromaticity of the light source at each pixel. Furthermore, several parameter settings for all algorithms are evaluated. The results reported on the LSAC-methods are obtained with the parameter setting as implied in [18], i.e., 0.18 s for the Gaussian kernel and 0.17 s for the exponential kernel (where $s = \max(n_x, n_y)/2$ and n_x and n_y are the width and height of the image, respectively). The results of the Retinex are optimal for a single scale (in the case of [13]) and for a single iteration (in the case of [14]). It is shown in Table I that both the Retinex and the LSAC outperforms the baseline algorithms. However, using grid-based sampling combined with the white-patch algorithm results in a median angular error of 9.5°, whereas the Retinex results in a median angular error of 9.9°. This difference is statistically significant at 95% confidence level, using the Wilcoxon sign test [31]. Hence, it can be concluded that when the number of light sources is *known*, the proposed method in combination with the gray-world algorithm can significantly outperform the Retinex method.

When assuming *no a priori* knowledge on the number of light sources, i.e., replacing the clustering step with a segmentation step, the proposed methodology in combination with the gray-world algorithm obtains a median angular error of 10.0°. The median angular error of the Retinex method is slightly lower (9.9°), but the Wilcoxon sign test still favors the proposed methodology. In other words, the proposed method (in combination with the gray world, the white patch, or the general gray world) statistically significantly outperforms the Retinex method. This can be explained by the relatively stable performance of the Retinex method, compared with the proposed method, i.e., the Retinex method results in fewer outliers than the proposed methodology as the proposed method shows a poor performance on a small number of images.

To conclude, the proposed methodology increases the performance of all algorithms based on the assumption of a single light source, either when assuming the number of distinct light

sources is at most two or without assumptions on the number of light sources. For instance, the median angular error of the gray-world algorithm is improved with 29% with respect to the corresponding baseline method (i.e., the color constancy algorithm based on a single light source assumption), whereas the median angular error of the gray edge is improved with 8%.

C. Real Images

Here, the proposed method is validated on a new data set of real *RGB* images. This data set consists of images captured under laboratory settings and of natural real scenes. All images differ in size but roughly correspond to 0.1 megapixels (e.g., 384×256). The set is captured using a Sigma SD10 camera with Foveon X3 sensor. The white balancing of the camera is kept fixed at the preset *overcast*, and the lens is Sigma DC 18-200, 3.5-6.3. The images are captured in raw format and transferred using Sigma Photo Pro (v2.1). All images are stored in *sRGB* color space but are converted to linear *RGB* space for the experiments.

In the images under laboratory settings, two halogen lights with the same specification are used. Four color filters are used to obtain different colors of the light source. To ensure diffuse radiation of the light sources, a white tent is installed between the light sources and the objects. For the ease of the ground truth of each illuminant, gray boards are fixed in the scene. In all the settings, to get the ground truth of each pixel, the area of every light source is manually annotated. Seven different objects are illuminated by a combination of two light sources. Images of the same scene are aligned at the pixel level for comparison, and after removing the images that are misaligned, 59 images remain.

The natural scenes are obtained outdoors around the campus. To provide the ground truth of the light sources, several gray balls are placed in the scene. To reduce the amount of inter-reflection, the balls are posed on tripods. The ground truth of the illumination is computed by applying the gray-world algorithm to the brightest region gray balls (of which the exact locations are manually annotated). In total, nine outdoor images are acquired, with illuminant chromaticity differences ranging from 1.5° to 13°. Some examples are shown in Fig. 7.

Results: The data sets are processed with the following parameters: grid-based dense sampling with patch-size 10×10 (approximately 5% of the image-size). The results on the scenes that are illuminated with two distinct light sources are shown in Table III(a). It can be derived from this data set that the obtained improvement is similar to the obtained improvement on the hyperspectral images. For instance, the obtained improvement of the proposed method using the white-patch algorithm on this set of images is 11%, whereas the improvement on the hyperspectral data set is 9%. The results on the natural scenes are shown in Table III(b). The performance of the proposed method on these images considerably increases, even though this data set consists of images with similar *and* dissimilar light sources. Moreover, the proposed methodology outperforms the LSAC and Retinex methods on the natural scenes [see Table III(b)]. On the images under laboratory settings, the proposed methodology using the gray world and the second-order gray edge outperforms the LSAC and the Retinex. The pairwise comparisons



Fig. 7. Examples of the real-world data set. (Upper row) The natural scenes and (bottom row) images captured in laboratory settings. Note that all experiments are performed on linear images; gamma correction is applied only for improved visualization.

TABLE III
PERFORMANCE OF COLOR CONSTANCY ALGORITHMS FOR THE REAL-WORLD IMAGES. THE PERFORMANCE ON THE IMAGES UNDER LABORATORY SETTING ARE SHOWN IN TABLE (A), AND THE PERFORMANCE ON THE NATURAL SCENES ARE SHOWN IN TABLE (B)

Method		Median	Method		Median
Do Nothing (DN)		18.7°	Do Nothing (DN)		3.6°
Grey-World (GW)		12.8°	Grey-World (GW)		8.9°
White-Patch (WP)		14.8°	White-Patch (WP)		7.8°
general GW (GGW)		14.9°	general GW (GGW)		8.9°
1 st -order GE (GE-1)		14.4°	1 st -order GE (GE-1)		6.4°
2 nd -order GE (GE-2)		14.6°	2 nd -order GE (GE-2)		5.0°
LSAC	Exp. filter	13.2°	LSAC	Exp. filter	7.4°
	Gauss. filter	12.9°		Gauss. filter	7.4°
Retinex	(Impl. from [14])	13.0°	Retinex	(Impl. from [14])	7.7°
	(Impl. from [34])	14.1°		(Impl. from [34])	7.6°
Proposed: grid based sampling	GW	11.7° (-9%)	Proposed: grid based sampling	GW	6.4° (-28%)
	WP	13.2° (-11%)		WP	6.7° (-14%)
	GGW	13.1° (-12%)		GGW	7.0° (-21%)
	GE-1	13.4° (-7%)		GE-1	5.6° (-13%)
	GE-2	12.3° (-16%)		GE-2	5.1° (+2%)

(a)

(b)

between these methods are shown to be statistically significant using the Wilcoxon sign test [31]. The methodology using the remaining three methods (i.e., the white patch, the general gray world, and the first-order gray edge) on the images under laboratory settings resulted in performance that is on par with the LSAC and the Retinex methods [see Table III(a)], again computed using the Wilcoxon sign test on the pairwise comparisons.

Results on the natural scenes are interesting. Although the proposed methodology using grid-based sampling outperforms both global algorithms, as well as Retinex and LSAC, the best performance is obtained by doing nothing. This remarkable fact is explained by the small variation in the data set (consisting of only nine images, recorded in small time frame). Hence, for some images, the colors of the two light sources deviate only marginally from white. However, for those images that do strongly deviate from white (which is the case for four images), the proposed methodology outperforms the “do nothing” approach (with average improvements of 4% to 9%).

Examples: Finally, some results are shown on several images captured from the web. Since these images come without ground truth, the comparison between the algorithms can be only done qualitatively. Fig. 8 shows results on three images. Each of the original images, shown in Fig. 8(a), is influenced by two different light sources. Fig. 8(b) and (c) are created by

applying *global* correction using either of the two (manually annotated) light sources. It can be observed that global correction for these images is inappropriate: The mountain image, for instance, corrected for the reddish sunlight results in an image that contains shadow regions that are too blue. On the other hand, when the shadow regions are correctly adjusted, then the sunlit regions appear too red. A tradeoff is shown in Fig. 8(d), obtained by applying the gray-world algorithm to the image. Fig. 8(e) shows the result of applying local correction using the proposed methodology. Although the *color appearance* of these images could be question of debate, it can be observed that the effects of the two different light sources are less visible in the images in Fig. 8(e). For instance, the snow in the top (e) is corrected to white, without overcorrecting the shadow region, as in (b). Furthermore, the blue cast that appears in the middle row of (c) and (d) is not visible in (e), whereas this does not come at the cost of a yellow cast, as in (b). Aesthetically, one might prefer other corrections, but this discussion is beyond the scope of this article.

V. DISCUSSION

In this paper, a new methodology that can be used to apply color constancy to images that are recorded in the presence of multiple different light sources has been proposed. Most

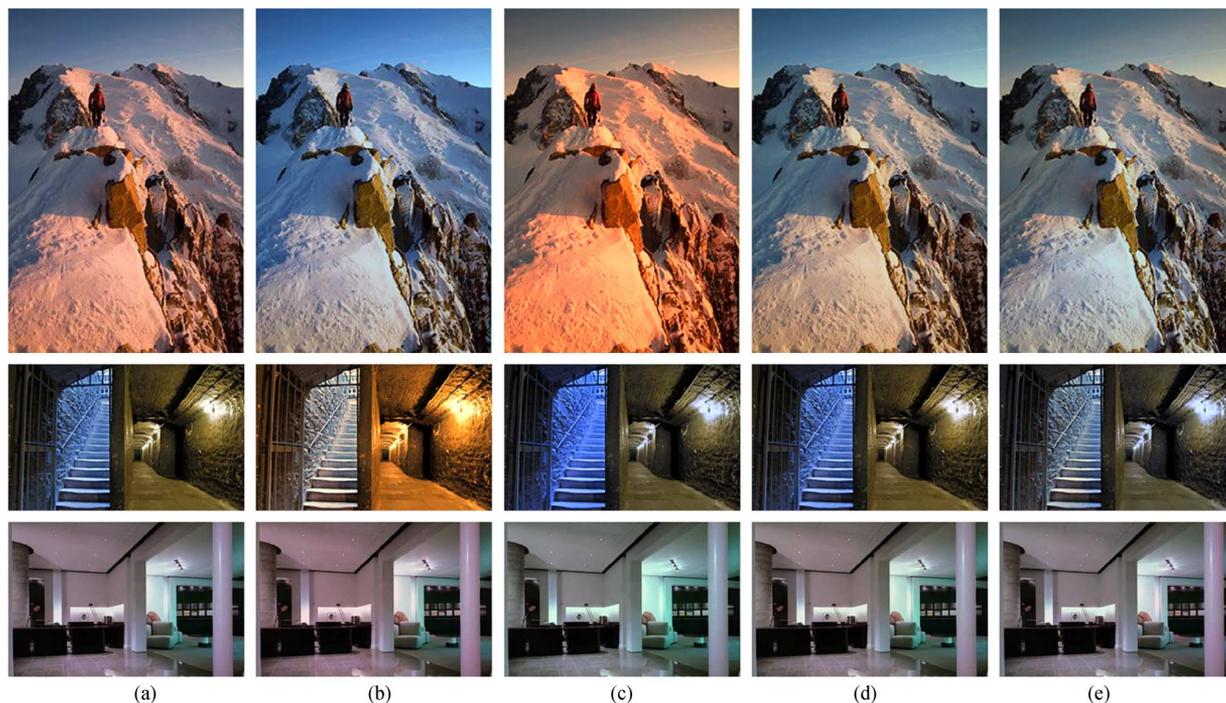


Fig. 8. Results of color constancy on some real-world images, taken from the web. (a) The original image. (b)–(c) The result of *global correction* with one of the two illuminants that are present (manually approximated). (d) The result of global correction with the gray-world algorithm and (e) the result of *local correction* with the proposed method (using gray world). Note that all experiments are performed on linear images; gamma correction is applied only for improved visualization.

existing algorithms are based on the assumption that the spectral distribution of the light source is uniform across the scene, but this assumption is often violated in reality. We have shown that many existing methods can be locally applied to image patches, rather than globally to the entire image. Interestingly, grid-based sampling outperformed segmentation-based and key-point-based sampling. This observation can be explained by the nature of the information that is sampled. Segmentation-based sampling will generally result in patches with similar colors (and, hence, less variation), whereas key-point-based sampling increases the likelihood of misclassifications (interest-point detectors usually trigger on boundaries that could be also a light-source transition). As a future work, texture-preserving segmentation methods can be used to overcome the disadvantage of segmentation-based sampling. Alternatively, other color constancy methods could be used in combination with segmentation-based sampling that are less sensitive to the limited variety of colors (for instance, physics-based methods using the dichromatic reflection model like [36]). One important parameter that needs to be taken into account when implementing the grid-based sampling procedure is the patch size. This parameter is closely related to the used color constancy algorithm, as well as to the contents of the image. Ideally, implementation should be combined with a calibration procedure to find the appropriate patch size for the used type of images and color constancy algorithm. Without such calibration, it is best to estimate the patch size conservatively (e.g., a patch size of approximately 2%–5% of the image dimensions).

Experiments on spectral and real images show that the proposed method properly reduces the influence of two light

sources present in one scene simultaneously. If the chromatic difference between these two illuminants is more than 1° , the proposed method outperforms algorithms based on the uniform light-source assumption (with an average error reduction of roughly 10%–15%). Otherwise, when the chromatic difference is less than 1° and the scene can be considered to contain one (approximately) uniform light source, the performance of the proposed method is similar to existing methods. The proposed methodology is able to improve existing algorithms even without assumptions on the number of light sources by applying mean-shift segmentation [35] on the back-projected illuminant estimates. However, the performance of this segmentation approach is less than the clustering-based approach as clustering-based approaches have the advantage of additional information on the number of light sources. An alternative to the segmentation-based approach would be the automatic detection of the number of light sources (and, hence, the number of clusters). To this end, a dynamic clustering algorithm that automatically determines the correct number of clusters could be adapted. Alternatively, spatial relations between the image patches could be taken into account by applying any segmentation algorithm to the patchwise illuminant estimates. Both approaches, however, come at the cost of additional parameters to learn.

To conclude, the methodology in this paper has been shown to be able to extend existing methods to more realistic scenarios where the uniform light-source assumption is too restrictive. We have shown that patch-based illuminant estimation can be as accurate as global illuminant estimation when the light source is (approximately) uniform. Furthermore, when there are two distinct light sources present in an image, the proposed method-

ology is able to increase the performance of existing algorithms considerably.

REFERENCES

- [1] M. Ebner, *Color Constancy*. Hoboken, NJ: Wiley, 2007.
- [2] B. Funt and G. Finlayson, "Color constant color indexing," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 17, no. 5, pp. 522–529, May 1995.
- [3] T. Gevers and A. Smeulders, "Color-based object recognition," *Pattern Recognit.*, vol. 32, no. 3, pp. 453–464, 1999.
- [4] G. Hordley, "Scene illuminant estimation: Past, present, and future," *Color Res. Appl.*, vol. 31, no. 4, pp. 303–314, Aug. 2006.
- [5] A. Gijzenij, T. Gevers, and J. van de Weijer, "Computational color constancy: Survey and experiments," *IEEE Trans. Image Process.*, vol. 20, no. 9, pp. 2475–2489, Sep. 2011.
- [6] G. Buchsbaum, "A spatial processor model for object colour perception," *J. Franklin Inst.*, vol. 310, no. 1, pp. 1–26, Jul. 1980.
- [7] E. H. Land, "The retinex theory of color vision," *Sci. Amer.*, vol. 237, no. 6, pp. 108–128, Dec. 1977.
- [8] D. Forsyth, "A novel algorithm for color constancy," *Int. J. Comput. Vis.*, vol. 5, no. 1, pp. 5–36, Aug. 1990.
- [9] A. Gijzenij and T. Gevers, "Color constancy using natural image statistics," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2007, pp. 1–8.
- [10] J. van de Weijer, T. Gevers, and A. Gijzenij, "Edge-based color constancy," *IEEE Trans. Image Process.*, vol. 16, no. 9, pp. 2207–2214, Sep. 2007.
- [11] A. Moore, J. Allman, and R. Goodman, "A real-time neural system for color constancy," *IEEE Trans. Neural Netw.*, vol. 2, no. 2, pp. 237–247, Mar. 1991.
- [12] D. Jobson, Z. Rahman, and G. Woodell, "Properties and performance of a center/surround retinex," *IEEE Trans. Image Process.*, vol. 6, no. 3, pp. 451–462, Mar. 1997.
- [13] D. Jobson, Z. Rahman, and G. Woodell, "A multiscale retinex for bridging the gap between color images and the human observation of scenes," *IEEE Trans. Image Process.*, vol. 6, no. 7, pp. 965–976, Jul. 1997.
- [14] B. Funt, F. Ciurea, and J. McCann, "Retinex in MATLAB," *J. Electron. Imag.*, vol. 13, no. 1, pp. 48–57, Jan. 2004.
- [15] G. Finlayson, B. Funt, and K. Barnard, "Color constancy under varying illumination," in *Proc. IEEE Int. Conf. Comput. Vis.*, 1995, pp. 720–725.
- [16] K. Barnard, G. Finlayson, and B. Funt, "Colour constancy for scenes with varying illumination," *Comput. Vis. Image Understand.*, vol. 65, no. 2, pp. 311–321, Feb. 1997.
- [17] W. Xiong and B. Funt, "Stereo retinex," *Image Vis. Comput.*, vol. 27, no. 1/2, pp. 178–188, Jan. 2009.
- [18] M. Ebner, "Color constancy based on local space average color," *Mach. Vis. Appl.*, vol. 20, no. 5, pp. 283–301, Jun. 2009.
- [19] M. Ebner, "Estimating the color of the illuminant using anisotropic diffusion," in *Proc. Comput. Anal. Images Patterns*, 2007, vol. 4673, pp. 441–449.
- [20] R. Kawakami, K. Ikeuchi, and R. T. Tan, "Consistent surface color for texturing large objects in outdoor scenes," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2005, pp. 1200–1207.
- [21] H. Spitzer and S. Semo, "Color constancy: A biological model and its application for still and video images," *Pattern Recognit.*, vol. 35, no. 8, pp. 1645–1659, 2002.
- [22] E. Hsu, T. Mertens, S. Paris, S. Avidan, and F. Durand, "Light mixture estimation for spatially varying white balance," *ACM Trans. Graph.*, vol. 27, no. 3, pp. 70:1–70:7, Aug. 2008.
- [23] G. Finlayson and S. Hordley, "Color constancy at a pixel," *Opt. Soc. Amer.*, vol. 18, no. 2, pp. 253–264, Feb. 2001.
- [24] C. Lu and M. Drew, "Shadow segmentation and shadow-free chromaticity via Markov random fields," in *Proc. IS T/SID's Color Imag. Conf.*, 2005, pp. 125–129.
- [25] K. Mikołajczyk and C. Schmid, "Scale and affine invariant interest point detectors," in *Proc. Int. J. Comput. Vis.*, Oct. 2004, vol. 60, no. 1, pp. 63–86.
- [26] P. F. Felzenszwalb and D. Huttenlocher, "Efficient graph-based image segmentation," *Int. J. Comput. Vis.*, vol. 59, no. 2, pp. 1–26, Sep. 2004.
- [27] M. D. Fairchild, *Color Appearance Models*. Hoboken, NJ: Wiley, 2005.

- [28] G. Finlayson, M. Drew, and B. Funt, "Spectral sharpening: Sensor transformations for improved color constancy," *Opt. Soc. Amer. A*, vol. 11, no. 5, pp. 1553–1563, May 1994.
- [29] H. Y. Chong, S. J. Gortler, and T. Zickler, "The von Kries hypothesis and a basis for color constancy," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2007, pp. 1–8.
- [30] J. von Kries, *Influence of Adaptation on the Effects Produced by Luminous Stimuli*. Cambridge, MA: MIT Press, 1970.
- [31] S. Hordley and G. Finlayson, "Reevaluation of color constancy algorithm performance," *J. Opt. Soc. Amer. A*, vol. 23, no. 5, pp. 1008–1020, May 2006.
- [32] D. Foster, S. Nascimento, and K. Amano, "Information limits on neural identification of coloured surfaces in natural scenes," *Vis. Neurosci.*, vol. 21, no. 3, pp. 331–336, Sep. 2004.
- [33] K. Barnard, L. Martin, B. Funt, and A. Coath, "A data set for colour research," *Color Res. Appl.*, vol. 27, no. 3, pp. 147–151, Jun. 2002.
- [34] V. Štruc and N. Pavešić, "Photometric normalization techniques for illumination invariance," in *Advances in Face Image Analysis: Techniques and Technologies*, Y. Zhang, Ed. Hershey, PA: IGI Global, 2010, ch. 15.
- [35] D. Comaniciu and P. Meer, "Mean shift: A robust approach toward feature space analysis," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 5, pp. 603–619, May 2002.
- [36] R. Tan, K. Nishino, and K. Ikeuchi, "Color constancy through inverse-intensity chromaticity space," *J. Opt. Soc. Amer. A*, vol. 21, no. 3, pp. 321–334, Mar. 2004.



Arjan Gijzenij (M'09) received the M.Sc. degree in computer science from the University of Groningen, Groningen, The Netherlands, in 2005 and the Ph.D. degree from the University of Amsterdam, Amsterdam, The Netherlands, in 2010.

Formerly, he was a Postdoctoral Researcher at the University of Amsterdam. He is currently a Scientific Software Consultant for Alten PTS, Eindhoven, The Netherlands. He has published papers in high-impact journals and conferences and served on the program committee of several conferences and workshops. His main research interests are in the field of color in computer vision and psychophysics.



Rui Lu is currently working toward the Ph.D. degree in the School of Computer and Information Technology, Beijing Jiaotong University, Beijing, China. She has been a continuous Master–Doctor student since 2005.

Her research interest include computer vision, image processing, and statistical pattern recognition. Currently, she focuses on computational color constancy.



Theo Gevers (M'11) is an Associate Professor of computer science with the University of Amsterdam, Amsterdam, The Netherlands, where he is also the Teaching Director of the M.Sc. degree of artificial intelligence. He is also currently with Computer Vision Center, Universitat Autònoma de Barcelona, Barcelona, Spain. He is a Lecturer of postdoctoral courses given at various major conferences (CVPR, ICPR, SPIE, CGIV). His main research interests are in the fundamentals of content-based image retrieval, color image processing, and computer

vision, specifically in the theoretical foundation of geometric and photometric invariants.

Dr. Gevers is the Chair of various conferences, and he is an Associate Editor of the IEEE TRANSACTIONS ON IMAGE PROCESSING. Furthermore, he is a program committee member of a number of conferences and an invited speaker at major conferences. He was the recipient of the VICI Award (for excellent researchers) from the Dutch Organization for Scientific Research.