

Edge-Based Color Constancy

Joost van de Weijer, Theo Gevers, and Arjan Gijsenij

Abstract—Color constancy is the ability to measure colors of objects independent of the color of the light source. A well-known color constancy method is based on the gray-world assumption which assumes that the average reflectance of surfaces in the world is achromatic. In this paper, we propose a new hypothesis for color constancy namely the gray-edge hypothesis, which assumes that the average edge difference in a scene is achromatic. Based on this hypothesis, we propose an algorithm for color constancy. Contrary to existing color constancy algorithms, which are computed from the zero-order structure of images, our method is based on the derivative structure of images. Furthermore, we propose a framework which unifies a variety of known (gray-world, max-RGB, Minkowski norm) and the newly proposed gray-edge and higher order gray-edge algorithms. The quality of the various instantiations of the framework is tested and compared to the state-of-the-art color constancy methods on two large data sets of images recording objects under a large number of different light sources. The experiments show that the proposed color constancy algorithms obtain comparable results as the state-of-the-art color constancy methods with the merit of being computationally more efficient.

Index Terms—Color constancy, object recognition, photometric invariance.

I. INTRODUCTION

COLOR constancy is the ability to recognize colors of objects independent of the color of the light source [1]. Obtaining color constancy is of importance for many computer vision applications, such as image retrieval, image classification, color object recognition, and object tracking [2]–[4].

Approaches to this problem can be divided into two groups. For the first group, the aim is to represent images by features which are invariant with respect to the light source, for example within the context of image retrieval. Such invariant representation have been proposed by Funt and Finlayson [5], Gevers and Smeulders [2], Geusebroek *et al.* [6], and Van de Weijer and Schmid [7]. For these methods, the actual estimation of the light source is not necessary. For the second group of approaches, the aim is to correct images for deviations from a canonical light source. Contrary to methods in the first group, solutions to this problem do estimate the color of the light source, be it explicitly or implicitly. Methods, either propose a light source estimation,

after which the image is corrected [8]–[11], or they directly estimate the color corrected image [1], [12], after which the light source can be derived. If desired, illuminant invariant features can subsequently be derived from the corrected image. In this paper, we look at color constancy approaches of the second group, i.e., methods from which a light source corrected image can be computed.

One of the most successful color constancy methods is gamut mapping proposed by Forsyth [1]. The method is based on the observation that only a limited set of RGB values can be observed under a given illuminant. The set of all possible RGB values for the canonical illuminant, typically a white illuminant, is called the canonical gamut. This canonical gamut is proven to be a convex hull in RGB space. The algorithm computes what transformations map an observed gamut into the canonical gamut. From these transformations, the illuminant color is derived. The gamut mapping algorithm provides among the best results in color constancy experiments [3]. Finlayson *et al.* [12] improve the gamut mapping algorithm by restricting the transformations to be plausible, meaning that only illuminants are allowed which correspond to existing illuminants. This adaptation of the gamut algorithm, called GCIE for gamut constrained illumination estimation, was shown to outperform the standard gamut algorithm. Further approaches to color constancy include probabilistic approaches [10] and learning-based methods [11]. A framework which unifies multiple color constancy algorithms together is presented by Finlayson and Hordley [9]. They propose to estimate the illuminant from the correlation of the image data, and the prior knowledge about which colors appear under a certain light.

Although the above-described algorithms arrive at reasonable color constancy accuracy, a drawback is that they are based on complex algorithms and all require an image data set with known light sources for calibration. In this paper, we will focus on color constancy based on less complex color constancy algorithms. To this end, fast algorithms are considered which are based on low-level image features, such as max-RGB and gray-world. Max-RGB is a simple and fast color constancy algorithm which estimates the light source color from the maximum response of the different color channels [13]. Another well-known simple color constancy method is based on the gray-world hypothesis [8], which assumes that the average reflectance in the scene is achromatic. If the images under evaluation are part of a coherent image data base, Gershon *et al.* [14] showed that assuming the average of a scene to be equal to the average reflectance of the database, improves the results over the standard gray-world method. As an example, they mention forest pictures full of green colors. In this case, most color constancy methods will predict light sources biased towards the green color. The database-compensated gray-world algorithm resolves this problem. These low-level methods are widely in

Manuscript received August 7, 2006; revised May 3, 2007. This work was supported by the Marie Curie European Fellowship Program of the Commission of the European Union. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Gabriel Marcu.

J. van de Weijer is with the LEAR team, INRIA Rhone-Alpes, 38330 Montbonnot, France, (e-mail: joost.van-de-weijer@inrialpes.fr).

T. Gevers and A. Gijsenij are with the Faculty of Science, University of Amsterdam, 1098 SJ Amsterdam, The Netherlands (e-mail: gevers@science.uva.nl; gijsenij@science.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.2007.901808

use, even in digital consumer cameras, due to their very low computational costs, i.e., taking the maximum (max-RGB) or average pixel values (gray-world).

Low-level approaches regained further interest recently after Finlayson and Trezzi [15] showed that only with minor adaptations results are obtained which are similar to those of complex color constancy algorithms. In fact, they showed that the max-RGB method and the gray-world method can be interpreted as the same algorithm applied with different instantiations of the error function. The max-RGB method is shown to be equal to applying the L^∞ Minkowski norm and gray-world is equal to using the L^1 norm. They further show that the best color constancy results are obtained with the L^6 norm. Although these simple color constancy algorithms are slightly outperformed by more elaborate methods, e.g., gamut mapping, they perform surprisingly well while they are conceptually simpler (for an extensive evaluation of multiple color constancy methods, see [3], [15], and [16]).

In this paper, we pursue this line of color constancy based on low-level image features. First, we propose the gray-edge hypothesis, which assumes that the average edge difference in the scene is achromatic. The method is based on the observation that the distribution of color derivatives exhibits the largest variation in the light source direction. The Minkowski norm of these derivatives is used to approximate this direction. The method is further extended to also include higher order derivatives. To our knowledge, this is the first work in which color constancy based on image derivatives is pursued. Finally, we propose a new framework of color constancy based on low-level image features which includes the known algorithms (gray-world, max-RGB, shades of gray) and the newly proposed gray-edge and higher order gray-edge algorithms.

This paper is organized as follows. In Section II, color constancy based on the gray-world and the max-RGB hypothesis is discussed. In Section III, we propose the gray-edge hypothesis, which estimates the illuminant color based on the distribution of the color derivatives. We further extend the color constancy framework of Finlayson and Trezzi [15] to also include color constancy methods derived from the gray-edge hypothesis. We further extend it with a parameter to regulate the amount of local averaging. Section IV, contains experiments on two large databases of images. Section V contains a discussion of results and indications for future research. Section VI finishes with concluding remarks.

II. GRAY-WORLD HYPOTHESIS

The image values, $\mathbf{f} = (R, G, B)^T$, for a Lambertian surface are dependent on the light source $e(\lambda)$, where λ is the wavelength, the surface reflectance $s(\lambda)$ and the camera sensitivity functions $\mathbf{c}(\lambda) = (R(\lambda), G(\lambda), B(\lambda))$

$$\mathbf{f} = \int_{\omega} e(\lambda) s(\lambda) \mathbf{c}(\lambda) d\lambda \quad (1)$$

where ω is the visible spectrum and bold fonts are applied for vectors. We assume that the scene is illuminated by a single light source. The goal of color constancy is to estimate the light

source color $e(\lambda)$, or its projection on the *RGB*-kernels

$$\mathbf{e} = \begin{pmatrix} R_e \\ G_e \\ B_e \end{pmatrix} = \int_{\omega} e(\lambda) \mathbf{c}(\lambda) d\lambda \quad (2)$$

given the image values $\mathbf{f}(\mathbf{x})$, where \mathbf{x} is the spatial coordinate in the image. The task of color constancy is not attainable without further assumptions.

Buchsbaum [8] proposes the gray-world hypothesis which assumes that the average reflectance in a scene is achromatic. In the original work, the hypothesis is used to derive that the average reflectance for the short-wave, middle-wave and long-wave regions is equal. Here, we employ a stronger definition of the achromatic reflectance of a scene (as also used in [15])

$$\frac{\int s(\lambda, \mathbf{x}) d\mathbf{x}}{\int d\mathbf{x}} = g(\lambda) = k \quad (3)$$

which avoids to make further assumptions. Buchsbaum [8], for example, needed to make further assumptions on the basis functions for the camera sensitivities, the surface reflectances, and the light source spectra. The constant k is between 0 for no reflectance (black) and 1 for total reflectance (white) of the incident light, and the integral is over the domain of the scene. For such a scene with achromatic reflectance, it holds that the reflected color is equal to the color of the light source, since

$$\frac{\int \mathbf{f}(\mathbf{x}) d\mathbf{x}}{\int d\mathbf{x}} = \frac{1}{\int d\mathbf{x}} \int \int_{\omega} e(\lambda) s(\lambda, \mathbf{x}) \mathbf{c}(\lambda) d\lambda d\mathbf{x} \quad (4)$$

$$= \int_{\omega} e(\lambda) \mathbf{c}(\lambda) \left(\frac{\int s(\lambda, \mathbf{x}) d\mathbf{x}}{\int d\mathbf{x}} \right) d\lambda \quad (5)$$

$$= k \int_{\omega} e(\lambda) \mathbf{c}(\lambda) d\lambda = k\mathbf{e} \quad (6)$$

where we applied the theorem of Fubini to exchange the order of integration. The normalized light source color is computed with $\hat{\mathbf{e}} = k\mathbf{e}/|k\mathbf{e}|$.

Another popular color constancy method is called max-RGB [17]. It is based on the assumption that the reflectance which is achieved for each of the three channels is equal

$$\max_{\mathbf{x}} \mathbf{f}(\mathbf{x}) = k\mathbf{e} \quad (7)$$

where the max operation is executed on the separate channels

$$\max_{\mathbf{x}} \mathbf{f}(\mathbf{x}) = \left(\max_{\mathbf{x}} R(\mathbf{x}), \max_{\mathbf{x}} G(\mathbf{x}), \max_{\mathbf{x}} B(\mathbf{x}) \right). \quad (8)$$

This method is sometimes explained as being derived from the white-patch hypothesis. Since a white patch reflects all the incident light, its position in the image can be found by searching for the maximum RGB values. It should be noted, however, that the max-RGB methods does not require the maxima of the separate channels to be on the same location; hence, it also obtains correct illuminant estimation results when the maximum reflectance is equal for the three channels.

Recently, Finlayson and Trezzi [15] showed that the gray-world and the max-RGB algorithm are two different instantiations of a more general color constancy algorithm based on the

Minkowski norm. Their method is called shades of gray and is computed by

$$\left(\frac{\int (\mathbf{f}(\mathbf{x}))^p d\mathbf{x}}{\int d\mathbf{x}} \right)^{1/p} = k\mathbf{e}. \quad (9)$$

For $p = 1$, the equation is equal to the gray-world assumption. For $p = \infty$, it is equal to color constancy by max-RGB. They investigated the performance of the illuminant estimation as a function of the Minkowski norm and found that the best results are obtained with a Minkowski norm with $p = 6$.

A similar approach has been proposed by Barnard [18] to select the appropriate transformation from the feasible set of transformations computed with Gamut mapping method. In the original work, Forsyth [1] proposed to take the transformation belonging to the gamut with the maximum volume. Instead, Barnard [18] considered various exponentials of the geometric mean of the transformation vector to select the best transformation from the feasible set. He showed that by varying the exponential, the selection criterion changes from taking the average over all transformation to the maximum volume heuristic. Like in the case of the shades of gray method, intermediate exponentials were shown to obtain better results.

As a final extension of the gray-world algorithm, we consider local averaging. The norm computation, as given by (9), is a global averaging operation, which ignores the important local correlation between pixels. This local correlation can be used to reduce the influence of noise. Local smoothing as a preprocessing step was proven to be beneficial for color constancy algorithms, as discussed in Barnard's study [3]. To exploit this local correlation, we introduce a local smoothing with a Gaussian filter, \mathbf{G}^σ , with standard deviation σ

$$\left(\frac{\int (\mathbf{f}^\sigma(\mathbf{x}))^p d\mathbf{x}}{\int d\mathbf{x}} \right)^{1/p} = k\mathbf{e} \quad (10)$$

where $\mathbf{f}^\sigma = \mathbf{f} \otimes \mathbf{G}^\sigma$.

III. GRAY-EDGE HYPOTHESIS

As an alternative to the gray-world hypothesis, we propose the *gray-edge hypothesis*: *the average of the reflectance differences in a scene is achromatic*

$$\frac{\int |s_{\mathbf{x}}^\sigma(\lambda, \mathbf{x})| d\mathbf{x}}{\int d\mathbf{x}} = g(\lambda) = k. \quad (11)$$

The subscript \mathbf{x} indicates the spatial derivative at scale σ . With the gray-edge assumption, the light source color can be computed from the average color derivative in the image given by

$$\frac{\int |\mathbf{f}_{\mathbf{x}}(\mathbf{x})| d\mathbf{x}}{\int d\mathbf{x}} = \frac{1}{\int d\mathbf{x}} \int \int e(\lambda) |s_{\mathbf{x}}(\lambda, \mathbf{x})| c(\lambda) d\lambda d\mathbf{x} \quad (12)$$

$$= \int_{\omega} e(\lambda) \left(\frac{\int |s_{\mathbf{x}}(\lambda, \mathbf{x})| d\mathbf{x}}{\int d\mathbf{x}} \right) c(\lambda) d\lambda \quad (13)$$

$$= k \int_{\omega} e(\lambda) c(\lambda) d\lambda = k\mathbf{e} \quad (14)$$

where $|\mathbf{f}_{\mathbf{x}}(\mathbf{x})| = (|R_{\mathbf{x}}(\mathbf{x})|, |G_{\mathbf{x}}(\mathbf{x})|, |B_{\mathbf{x}}(\mathbf{x})|)^T$. The gray-edge hypothesis originates from the observation that the color

derivative distribution of images forms a relatively regular, ellipsoid-like shape, of which the long axis coincides with the light source color [20]. In Fig. 1, the color derivative distribution is depicted for three images. The color derivatives are rotated to the opponent color space as follows:

$$\begin{aligned} O1_{\mathbf{x}} &= \frac{R_{\mathbf{x}} - G_{\mathbf{x}}}{\sqrt{2}} \\ O2_{\mathbf{x}} &= \frac{R_{\mathbf{x}} + G_{\mathbf{x}} - 2B_{\mathbf{x}}}{\sqrt{6}} \\ O3_{\mathbf{x}} &= \frac{R_{\mathbf{x}} + G_{\mathbf{x}} + B_{\mathbf{x}}}{\sqrt{3}}. \end{aligned} \quad (15)$$

In the opponent color space, $O3$ coincides with the white light direction. For the scene under white light (the leftmost picture), the distribution of the derivatives are centered along the $O3$ i.e., the white-light axis. Once we change the color of the light source, as in the second and third picture, the distribution of the color derivatives no longer align with the white-light axis. In other words, color constancy based on the gray-edge assumption can be interpreted as skewing the color derivative distribution such that the average derivative is in the $O3$ orientation.

Similar to the gray-world based color constancy, the gray-edge hypothesis can also be adapted to incorporate the Minkowski norm

$$\left(\frac{\int |\mathbf{f}_{\mathbf{x}}^\sigma(\mathbf{x})|^p d\mathbf{x}}{\int d\mathbf{x}} \right)^{1/p} = k\mathbf{e}. \quad (16)$$

Color constancy based on this equation assumes that the p th Minkowski norm of the derivative of the reflectance in a scene is achromatic. We distinguish two special cases. For $p = 1$, the illuminant is derived by a normal averaging operation over the derivatives of the channels. For $p = \infty$, the illuminant is computed from the maximum derivative in the scene. The resemblance between the color constancy derivation from the gray-world and gray-edge hypothesis is apparent. Both methods can be combined in a single framework of color constancy methods based on low-level image features derived from the following general hypothesis:

$$\left(\int \left| \frac{\partial^n \mathbf{f}^\sigma(\mathbf{x})}{\partial \mathbf{x}^n} \right|^p d\mathbf{x} \right)^{1/p} = k\mathbf{e}^{n,p,\sigma}. \quad (17)$$

The division by $\int d\mathbf{x}$ has been incorporated into the constant k . Next to the already discussed hypotheses (gray-world, max-RGB, Minkowski norm, and the newly proposed gray-edge), it is obvious that this framework also includes higher order-based color constancy. High-order derivatives have correspondences with the center-surround mechanism of the human eyes for color constancy such as exploited in the well-known center-surround retinex algorithm [21]. The influence of the color intensities could be weighted according to their distance to the center of the receptive field generally calculated by a difference of Gaussian functions.

The illuminant estimation of (17) describes a framework for low-level based illuminant estimation. This framework produces different estimations for the illuminant color based on three variables.

- 1) The order n of the image structure is the parameter determining if the method is a gray-world or a gray-edge algorithm. The gray-world methods are based on the *RGB*

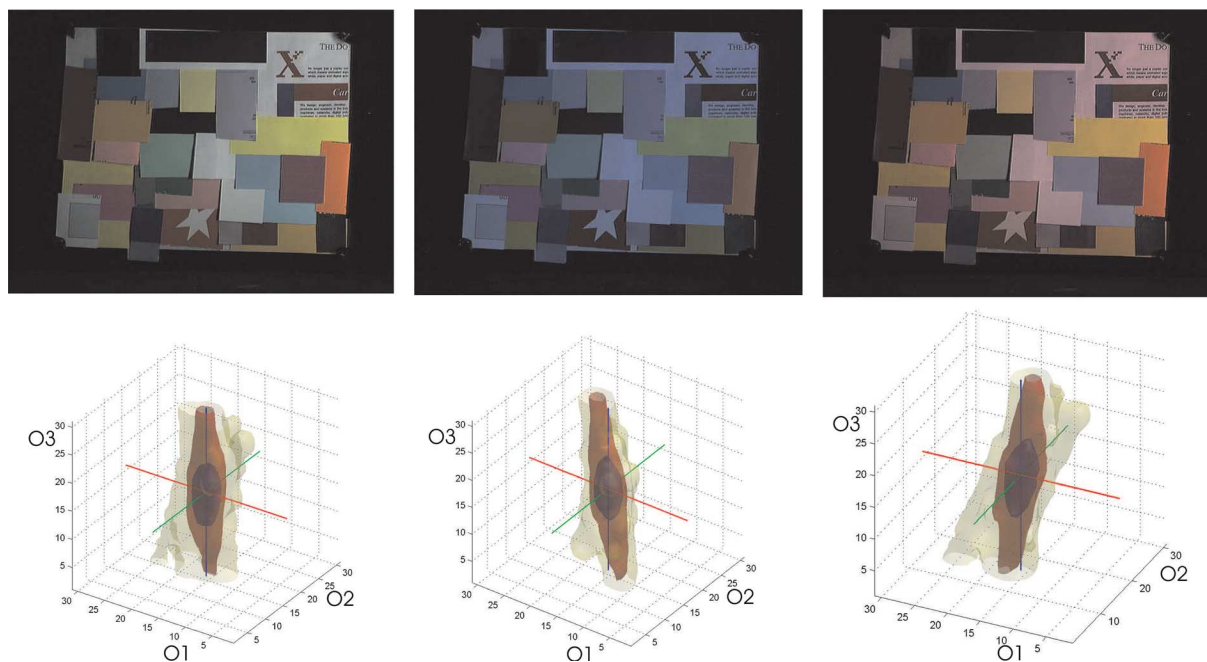


Fig. 1. Three acquisitions of the same scene under different light sources [19]. In the bottom row, the color derivative distributions are shown, where the axes are the opponent color derivatives and the surfaces indicate derivative values with equal occurrence and darker surfaces indicating a more dense distribution. Note the shift of the orientation of the distribution of the derivatives with the changing of the light source.

TABLE I
OVERVIEW OF THE DIFFERENT ILLUMINANT ESTIMATIONS METHODS TOGETHER WITH THEIR
HYPOTHESES. THESE ILLUMINANT ESTIMATIONS ARE ALL INSTANTIATIONS OF (17)

name	symbol	equation	hypothesis
Grey-World	$e^{0,1,0}$	$\left(\int \mathbf{f}(\mathbf{x}) d\mathbf{x}\right) = k\mathbf{e}$	the average reflectance in a scene is achromatic
max-RGB	$e^{0,\infty,0}$	$\left(\int \mathbf{f}(\mathbf{x}) ^\infty d\mathbf{x}\right)^{\frac{1}{p}} = k\mathbf{e}$	the maximum reflectance in a scene is achromatic
Shades of Grey	$e^{0,p,0}$	$\left(\int \mathbf{f}(\mathbf{x}) ^p d\mathbf{x}\right)^{\frac{1}{p}} = k\mathbf{e}$	the p th-Minkowsky norm of a scene is achromatic
General Grey-World	$e^{0,p,\sigma}$	$\left(\int \mathbf{f}^\sigma(\mathbf{x}) ^p d\mathbf{x}\right)^{\frac{1}{p}} = k\mathbf{e}$	the p th-Minkowsky norm of a scene is achromatic after local smoothing
Grey-Edge	$e^{1,p,\sigma}$	$\left(\int \mathbf{f}_x^\sigma(\mathbf{x}) ^p d\mathbf{x}\right)^{\frac{1}{p}} = k\mathbf{e}$	the p th-Minkowsky norm of the image derivative in a scene is achromatic
Max-Edge	$e^{1,\infty,\sigma}$	$\left(\int \mathbf{f}_x^\sigma(\mathbf{x}) ^\infty d\mathbf{x}\right)^{\frac{1}{p}} = k\mathbf{e}$	the maximum reflectance difference in a scene is achromatic
2nd order Grey-Edge	$e^{2,p,\sigma}$	$\left(\int \mathbf{f}_{xx}^\sigma(\mathbf{x}) ^p d\mathbf{x}\right)^{\frac{1}{p}} = k\mathbf{e}$	the p th-Minkowsky norm of the second order derivative in a scene is achromatic

values, whereas the gray-edge methods are based on the spatial derivatives of order n . In this paper, we will investigate higher order-based color constancy up to order $n = 2$.

- 2) The Minkowski norm p which determines the relative weights of the multiple measurements from which the final illuminant color is estimated. A high Minkowski norm emphasizes larger measurements whereas a low Minkowski norm equally distributes weights among the measurements.
- 3) The scale of the local measurements as denoted by σ . For first- or higher order estimation, this local scale is combined with the differentiation operation computed with the Gaussian derivative. For zero-order gray-world methods, this local scale is imposed by a Gaussian smoothing operation.

An overview of the instantiations of the illuminant estimation given by the framework of (17), which are considered in this paper, is given in Table I.

An advantage of the color constancy methods based on (17) is that they are all based on low computational demanding operations. In fact, the p th Minkowski norm of (smoothed) RGB values or derivatives can be computed extremely fast (even real-time on dedicated hardware). Furthermore, the method does not require an image database taken under a known light source for calibration as is necessary for more complex color constancy methods such as color gamut mapping, and color by correlation [1], [15].

IV. EXPERIMENTS

In the previous section, a general formulation for color illuminant estimation has been proposed based on low-level image features. In this section, the performance is tested for various parameter settings on a set of colorful objects in a controlled indoor setting, and on a real-world data set containing mainly outdoor scenes. For both data sets, the illuminant color of the scene is provided as additional information (i.e., ground-truth).

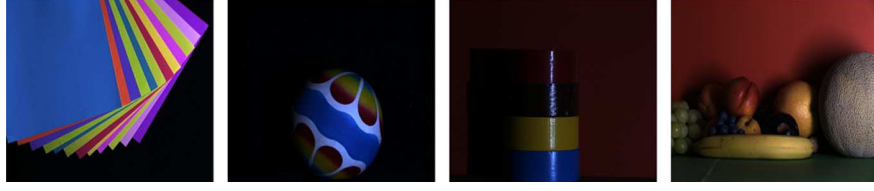


Fig. 2. Examples of the images in data set [19].

For evaluation, the angular error between the estimated light source \mathbf{e}_e and the actual light source \mathbf{e}_l is used as an error measure

$$\text{angular error} = \cos^{-1}(\hat{\mathbf{e}}_l \cdot \hat{\mathbf{e}}_e) \quad (18)$$

where the (\cdot) indicates a normalized vector. For the three data sets tested here, we report the median angular error, which is considered appropriate to assess the performance of color constancy algorithms [22].

A. Controlled Indoor Image Set

First, the proposed algorithms are tested on a large data set of colorful objects under different light sources [19]. The set consists of images taken under 11 varying light sources of 30 different scenes containing both matte and specular objects (see examples in Fig. 2). Several images were found to be unusable, resulting in a data set of 321 images. Results of other color constancy algorithms on this standard data set are available in [3], [12], [15], and [22].

In Table II, the results of multiple methods are summarized. Let us first consider the results obtained by the gray-world and its performance as a function of the Minkowski norm and local smoothing. Both parameters significantly improve the performance. For the zero-order instantiation, only varying the Minkowski norm reduces the error from 7.0° for gray-world, to 3.7° for a Minkowski norm equal to 7. A further increase in performance of 15%, to an error of 3.2° , is obtained by combining a simple local smoothing with $\sigma = 1$ and a Minkowski norm of 11 (indicated by $e^{0,11,1}$). The illuminant estimations based on the gray-edge hypothesis, both the first and second-order, exhibit similar behavior. Again, a significant drop in the error is obtained by an appropriate choice of the Minkowski norm and local smoothing. The best results, an error of 2.7° , is obtained with the second-order gray-edge method.

In Fig. 3, the median error is given as a function of the Minkowski norm and local smoothing for the indoor image set. These figures show that a small local scale or a low Minkowski norm perform significantly less. The performance increases by augmenting the Minkowski norm and the local scale. The performance quickly levels off to a plateau where the performance changes little. In Table III, we have indicated for which regions of parameter settings a comparable performance, i.e., within 10% of the optimal, is obtained.

Results of more complex color constancy methods, such as gamut mapping, neural network algorithm, and color-by-correlation, have been reported in literature for the images in group A [3], [12], [22]. These results have been also included in Table II. The more complex algorithms obtained comparable results to the ones we reported. Only the *GCIE* algorithms obtain better

TABLE II
MEDIAN ANGULAR ERROR (DEGREES) ON INDOOR IMAGE DATA SET FOR VARIOUS COLOR CONSTANCY METHODS

indoor set	symbol	median
Grey-World	$e^{0,1,0}$	7.0
Max-RGB	$e^{0,\infty,0}$	6.5
Shades of Grey	$e^{0,7,0}$	3.7
general Grey-World	$e^{0,11,1}$	3.2
Grey-Edge	$e^{1,7,4}$	3.2
2nd order Grey-Edge	$e^{2,7,5}$	2.7
Color by Correlation	-	3.2
Gamut Mapping	-	2.9
Neural Networks	-	7.8
GCIE Version 3, 11 lights	-	1.3
GCIE Version 3, 87 lights	-	2.6

TABLE III
PARAMETER SETTINGS FOR WHICH THE PERFORMANCE REMAINS WITHIN 10% OF OPTIMAL PERFORMANCE AS GIVEN IN TABLE II

method	local scale	Minkowski norm
general Grey-World	$1 \leq \sigma \leq 5$	$8 \leq p \leq 18$
Grey-Edge	$3 \leq \sigma \leq 5$	$6 \leq p \leq 14$
2nd order Grey-Edge	$4 \leq \sigma \leq 7$	$5 \leq p \leq 11$

results with an error of 1.3° . However, it should be noted that this algorithm uses the 11 illuminants which were used during the image acquisition as prior knowledge. If this prior knowledge is put aside, the performance drops to an error of 2.6° (GCIE Version 3, 87 lights).

B. Real-World Image Set

Next, the color constancy algorithms are tested on a database presented by Ciurea and Funt [23]. The database contains 11 000 images extracted from 2 h of digital video. Both indoor and outdoor scenes from a wide variety of locations are represented (see Fig. 5). A small gray sphere was mounted onto the video camera, appearing for all the images in the right bottom corner of the images. The sphere is exploited to estimate the illuminant color in the scene. This color illuminant estimation is available with the database and is used as a ground truth in this experiment. The original images were extracted from 15 different film clips taken at different locations. Because of the high correlation between the images in the database, the experiments are performed on a subset of 150 randomly chosen images containing ten images from each of the 15 video clips. The pixels in the right bottom

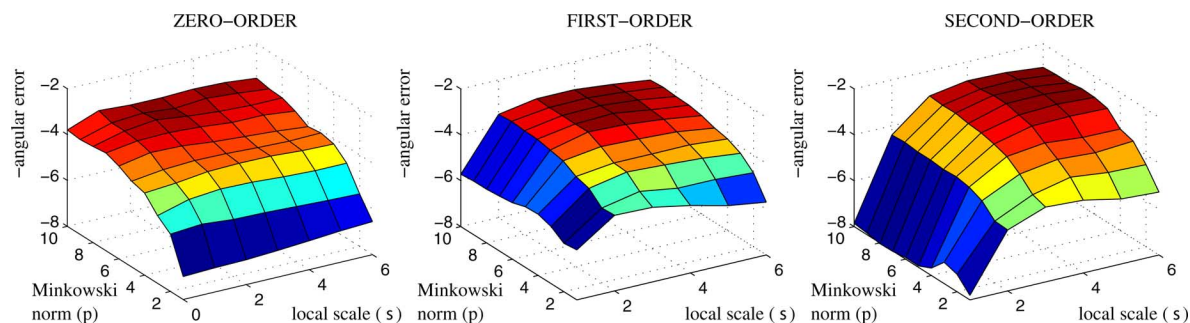


Fig. 3. Median angular error of the general gray-world, first-order, and second-order gray-edge method as a function of the Minkowski norm and local smoothing. The angular error axis is inverted for visualization purposes.



Fig. 4. Examples of the images from the real-world data set [23].

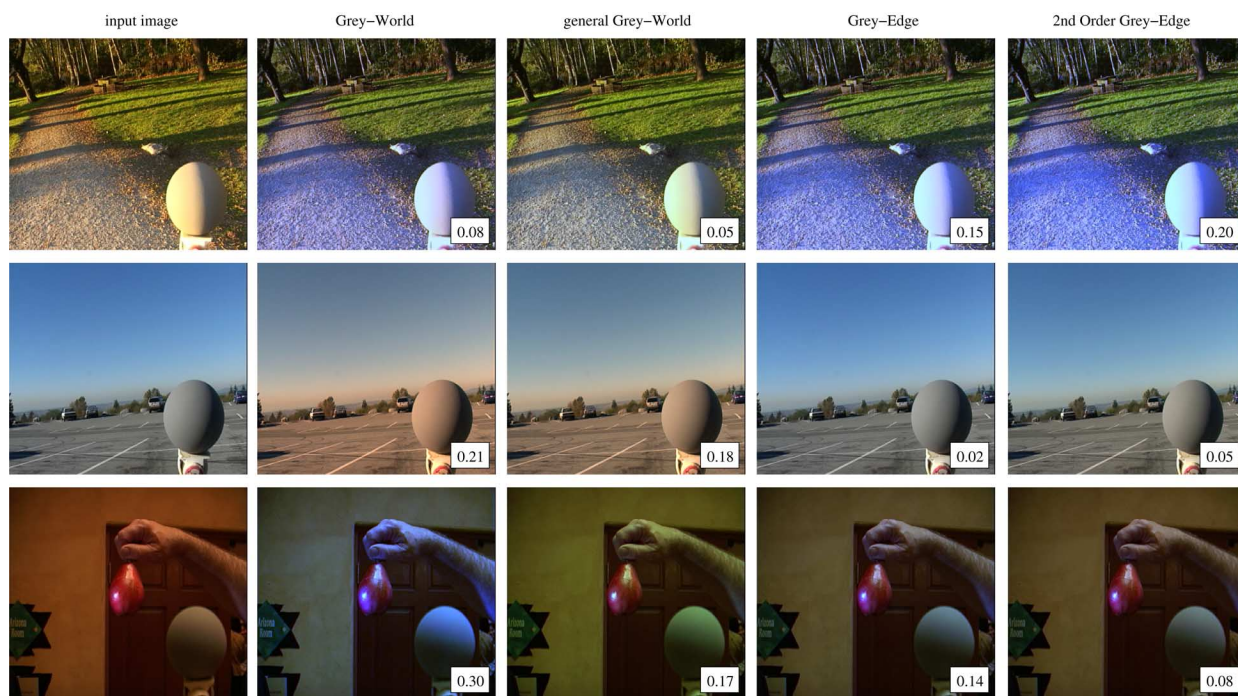


Fig. 5. Color constancy results of gray-world, general gray-world, gray-edge, and second-order gray-edge on real-world data set. The angular error is indicated in the right bottom corner. The first row depicts a failure of the edge-based approaches, whereas the gray-world methods give acceptable results. The second and third rows show examples where the gray-world methods fail and the gray-edge methods obtain superior results.

corner, which contains the gray sphere, are excluded from the color constancy computation.

The results on the real-world data set are summarized in Table IV. Again, an appropriate choice of the parameters significantly improves the results. Interestingly, the gray-edge performs best on this set of real-world images. It improves the color constancy results with about 40% compared to the Max-RGB.

In Fig. 5, corrected images based on gray-world and gray-edge methods are given for the real-world set. We applied the optimal parameters as given in Table IV. We often found a high correlation between the two edge-based methods. However, the estimations based on the gray-world and the two edge-based methods differ often. For example, in row one of Fig. 5, the edge-based methods fail and their estimation of the light source is too blue, whereas the general gray-world obtains acceptable

TABLE IV
MEDIAN ANGULAR ERROR (DEGREES) FOR VARIOUS COLOR
CONSTANCY METHODS ON REAL-WORLD IMAGE SET

Real-World set	symbol	median
Grey-World	$e^{0,1,0}$	7.3
Max-RGB	$e^{0,\infty,0}$	6.7
general Grey-World	$e^{0,13,2}$	4.7
Grey-Edge	$e^{1,1,6}$	4.1
2nd order Grey-Edge	$e^{2,1,5}$	4.3

results. The second row shows an example where the large blue sky results in an light source estimation which is much too blue for the gray-world methods. The edge-based methods obtain better results for this image. The last row shows an example where the second-order gray-edge method outperforms the other methods.

V. DISCUSSION

In the previous experiments, we have demonstrated that the proposed color constancy algorithms obtain comparable results to more complex color constancy algorithms. However, the optimal parameter setting vary for the different data sets. Important to note is the difference in the Minkowski norm for the edge-based methods. On the real-world set, a Minkowski norm of one is optimal whereas for the indoor sets a higher Minkowski norm between 6 and 15 is optimal. This is probably caused by the black background present in the indoor scenes. For the edge-based methods, the illumination estimate is computed by averaging the edge differences in a scene. Depending on the Minkowski norm, more weight is given to prominent edges. For the edge-based methods, the fact that the background in the images of sets A and B is black, significantly increases the chance that the highest edge in the image is black-white. This explains why a high Minkowski norm for these data sets is preferred. In case of the real-world data, the chance that the highest edge-response is caused by a black-white edge is smaller and consequently it is advantageous to take multiple measurements in the image into account. This is reflected in a lower Minkowski norm.

The proposed gray-edge algorithm can be obtained from the gray-world algorithm by simply exchanging the RGB values for the spatial image derivatives. The zero-order image structure, which provides the building stones for the gray-world method, is replaced by the higher order image structure. In the case of the gray world, the change to higher orders proved beneficial and the gray-edge methods outperform the gray-world methods. To our knowledge this paper is the first to propose color constancy derived from image derivatives. Methods such as Gamut mapping, neural network based color constancy and color by correlation are all based on the zero-order structure of images. However, there is no restriction which prevents them from using the higher order structure of images, and it would be interesting to see how these methods performed once based on the derivatives of images, or based on both zero-order and higher order structure of images.

Further work also includes searching for more elaborate ways to combine the low-level building blocks proposed in this paper. Can we find ways to decide what method is expected to perform best on a particular image? Furthermore, we believe that automatic estimation of the parameters separately per image, instead of for the whole data set, will improve the color constancy results, possibly within a learning context, such as proposed by Cardei *et al.* [11]. Another interesting research direction would be to constrain the possible illuminants to be physically feasible as is done in [12].

VI. CONCLUSION

In this paper, we have investigated edge-based color constancy. The method is derived from the gray-edge hypothesis which assumes that the average edge difference in a scene is achromatic. In contrast to existing methods, which are based on zero-order structure of the image, our method is based on the higher order structure of images. Furthermore, we introduce a framework of color constancy based on low-level image features which includes the known algorithms (gray-world, max-RGB, Minkowski norm) as well as the newly proposed gray-edge and higher order gray-edge algorithms. The quality of the various instantiations of the framework is tested on two large data sets of images recording objects under a large number of different light sources. The experimental results show that the newly proposed simple color constancy algorithms obtain similar results as more complex state-of-the-art color constancy methods. Furthermore, the results show that color constancy based on the gray-edge hypothesis obtains better results than those obtained with the gray-world method for real-world images.

ACKNOWLEDGMENT

The authors would like to thank F. Ciurea for providing the real-world database.

REFERENCES

- [1] D. Forsyth, "A novel algorithm for color constancy," *Int. J. Comput. Vis.*, vol. 5, no. 1, pp. 5–36, 1990.
- [2] T. Gevers and A. Smeulders, "Color based object recognition," *Pattern Recognit.*, vol. 32, pp. 453–464.
- [3] K. Barnard, L. Martin, A. Coath, and B. Funt, "A comparison of computational color constancy algorithms—Part II: Experiments with image data," *IEEE Trans. Image Process.*, vol. 11, no. 9, pp. 985–996, Sep. 2002.
- [4] J.-P. Renno, D. Makris, T. Ellis, and G. Jones, "Application and evaluation of colour constancy in visual surveillance," presented at the Joint IEEE Int. Workshop Visual Surveillance and Performance Evaluation, Beijing, China, 2005.
- [5] B. Funt and G. Finlayson, "Color constant color indexing," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 17, no. 5, pp. 522–529, 1995.
- [6] J. M. Geusebroek, R. van den Boomgaard, A. W. M. Smeulders, and T. Gevers, "A physical basis for color constancy," *Pattern Recognit. Lett.*, vol. 24, pp. 1653–1662, 2003.
- [7] J. van de Weijer and C. Schmid, "Blur robust and color constant image description," presented at the Int. Conf. Image Processing, 2006.
- [8] G. Buchsbaum, "A spatial processor model for object colour perception," *J. Frank. Inst.*, vol. 310, 1980.
- [9] G. Finlayson, S. Hordley, and P. Hubel, "Color by correlation: A simple, unifying framework for color constancy," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 11, pp. 1209–1221, Nov. 2001.
- [10] D. Brainard and W. Freeman, "Bayesian color constancy," *J. Opt. Soc. Amer. A*, vol. 14, no. 7, p. 1393.

- [11] V. Cardei, B. Funt, and K. Barnard, "Estimating the scene illumination chromaticity using a neural network," *J. Opt. Soc. Amer. A*, vol. 19, no. 12, 2002.
- [12] G. Finlayson and S. Hordley, "Gamut constrained illumination estimation," *Int. J. Comput. Vis.*, vol. 67, no. 1, pp. 93–109, 2006.
- [13] E. Land and J. McCann, "Lightness and retinex theory," *J. Opt. Soc. Amer. A*, vol. 61, no. 1, pp. 1–11, Jan. 1971.
- [14] R. Gershon, A. D. Jepson, and J. K. Tsotsos, "From [r,g,b] to surface reflectance: Computing color constant descriptors in images," in *Proc. 10th IJCAI*, Milan, Italy, 1987, pp. 755–758.
- [15] G. Finlayson and E. Trezzi, "Shades of gray and colour constancy," in *Proc. IS&T/SID 12th Color Imaging Conf.*, 2004, pp. 37–41.
- [16] K. Barnard, V. Cardei, and B. Funt, "A comparison of computational color constancy algorithms—Part I: Methodology and experiments with synthesized data," *IEEE Trans. Image Process.*, vol. 11, no. 9, pp. 972–984, Sep. 2002.
- [17] E. H. Land, "The retinex theory of color vision," *Sci. Amer.*, vol. 237, no. 6, pp. 108–120, 122, 123, 126, 128, 1977.
- [18] K. Barnard, "Improvements to gamut mapping colour constancy algorithms," in *Proc. Eur. Conf. Computer Vision*, 2000, pp. 390–403.
- [19] 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, 2002.
- [20] J. van de Weijer, T. Gevers, and A. Bagdanov, "Boosting color saliency in image feature detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 1, pp. 150–156, Jan. 2006.
- [21] E. Land, "An alternative technique for the computation of the designator in the retinex theory of color vision," in *Proc. Nat. Acad. Sci.*, 1986, vol. 83, pp. 3078–3080.
- [22] 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.
- [23] F. Ciurea and B. Funt, "A large image database for color constancy research," in *Proc. IS&T/SID's Color Imaging Conf.*, Scottsdale, AZ, 2004, pp. 160–64.



Joost van de Weijer received the M.Sc. degree in applied physics from the Delft University of Technology, Delft, The Netherlands, in 1998, and the Ph.D. degree in computer science from the University of Amsterdam, Amsterdam, The Netherlands, in 2005.

He is currently a Marie Curie Intra-European Fellow with the LEAR team, INRIA Rhone-Alpes, France. His research interests include color constancy, object recognition, color image processing, color naming, and machine learning for computer vision.



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

Dr. Gevers is the Chair of the Internet Imaging Conference (SPIE 2005 and 2006) and Multimedia Content Access: Algorithms and Systems (SPIE 2008), and coorganizer of the First International Workshop on Image Databases and Multi Media Search (1996), the International Conference on Visual Information Systems (1999, 2005), and the Conference on Multimedia and Expo (ICME 2005). He is the Guest Editor of the *International Journal of Computer Vision* special issue on content-based image retrieval and the *Computer Vision and Image Understanding* special issue on color for image indexing and retrieval. Further, he is a Lecturer for postdoctoral courses given at various major conferences (CVPR, ICCV, ICPR, ICIP, SPIE, and CGIV).



Arjan Gijsenij received the M.Sc. degree in computer science from the University of Groningen, Groningen, The Netherlands, in 2005. Currently, he is pursuing the Ph.D. degree at the University of Amsterdam, Amsterdam, The Netherlands.

His research interests include computer vision, (color) image processing, and statistical pattern recognition. His current research focuses on color constancy algorithms.