

Color Constancy Using Natural Image Statistics and Scene Semantics

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

Abstract—Existing color constancy methods are all based on specific assumptions such as the spatial and spectral characteristics of images. As a consequence, no algorithm can be considered as universal. However, with the large variety of available methods, the question is how to select the method that performs best for a specific image. To achieve selection and combining of color constancy algorithms, in this paper natural image statistics are used to identify the most important characteristics of color images. Then, based on these image characteristics, the proper color constancy algorithm (or best combination of algorithms) is selected for a specific image. To capture the image characteristics, the Weibull parameterization (e.g., grain size and contrast) is used. It is shown that the Weibull parameterization is related to the image attributes to which the used color constancy methods are sensitive. An MoG-classifier is used to learn the correlation and weighting between the Weibull-parameters and the image attributes (number of edges, amount of texture, and SNR). The output of the classifier is the selection of the best performing color constancy method for a certain image. Experimental results show a large improvement over state-of-the-art single algorithms. On a data set consisting of more than 11,000 images, an increase in color constancy performance up to 20 percent (median angular error) can be obtained compared to the best-performing single algorithm. Further, it is shown that for certain scene categories, one specific color constancy algorithm can be used instead of the classifier considering several algorithms.

Index Terms—Color constancy, illuminant estimation, natural image statistics, scene semantics, computer vision.

1 INTRODUCTION

DIFFERENCES in illumination cause measurements of object colors to be biased toward the color of the light source. Fortunately, humans have, to some extent, the ability of color constancy: They perceive the same color of an object despite large differences in illumination [1], [2]. Various computer vision-related topics like human-computer interaction [3], color feature extraction [4], [5], and color appearance models [6] would benefit from a similar ability.

Many computational color constancy algorithms have been proposed, see [7], [8] for recent overviews. In general, color constancy algorithms can be divided into two groups. The first group consists of algorithms based on low-level image features that can be directly applied to images. One of the first color constancy methods is based on the Retinex-theory by Land and McCann [9]. Examples of algorithms which are derived from this theory include the White-Patch algorithm [10], the Gray-World algorithm [11], and more recently, the Shades-of-Gray algorithm [12] and the Gray-Edge algorithm [13].

The second group consists of algorithms that use information acquired in a learning phase to obtain knowledge about the images, like possible light sources and the distribution of possible reflectance colors to be present in natural scenes. This information is then used to estimate the

illuminant. One of the first algorithms of this type is the gamut mapping algorithm by Forsyth [14]. This algorithm is based on the assumption that in real-world images, for a given illuminant, only a limited number of colors can be observed. Using this assumption, the illuminant can be estimated by comparing the distribution of colors in the current image to a prelearned distribution of colors (called the canonical gamut). Many algorithms have been derived from the original algorithm including Color-by-Correlation [15], the Gamut-constrained illuminant estimation [16], and derivative-based gamut mapping [17]. Other approaches that use a learning phase include probabilistic methods [18], [19] and methods based on genetic algorithms [20].

Because the color constancy problem is an underconstrained problem (see, for example, [14], [15], [8]), all of the above color constancy methods are based on specific assumptions. These assumptions include constrained gamuts (limited number of image colors which can be observed under a specific illuminant) and the distribution of colors that are present in an image (e.g., white patch, gray-world, gray-edge, etc.). As a consequence of the use of different assumptions, no color constancy algorithm can be considered as universal. With the large variety of available methods, the problem is how to select the method that is most suitable for different image settings and scenes. Furthermore, the subsequent question is how to combine the different algorithms in a proper way.

Little research has been published on the selection and fusion of color constancy methods. In [21], fusion is performed by a weighted average of several methods. More recently, a statistics-based method is combined with a physics-based method [22]. Both methods are based on weighting the output of the used color constancy algorithms, where the weights are optimized for a specific data set. However, the combination of the algorithms used still depends on the type of images being processed.

- The authors are with the Faculty of Science, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands.
E-mail: {a.gijsenij, th.gevers}@uva.nl.

Manuscript received 10 Nov. 2008; revised 1 Sept. 2009; accepted 18 Feb. 2010; published online 2 Apr. 2010.

Recommended for acceptance by G. Finlayson.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-2008-11-0781.

Digital Object Identifier no. 10.1109/TPAMI.2010.93.

Therefore, in this paper, the most appropriate color constancy algorithm for an image is selected based on the statistical contents of this image. For instance, even though edge-based color constancy algorithms are, *on average*, preferred over pixel-based color constancy algorithms [13], [17], this does not imply that edge-based color constancy algorithms are the best methods for *all* images. If only a few edges are present in an image, i.e., edge-based methods only have a little information to work with, then it is likely that pixel-based methods will perform better. To this end, the Weibull parameterization is used to express these image characteristics in terms of grain size (texture) and contrast. Then, based on these image characteristics, the proper color constancy algorithm (or best combination of algorithms) is selected for a specific image. As Weibull distributions are derived from higher order image statistics (i.e., image derivatives), the choice of a proper set of different color constancy methods should support this. To this end, the color constancy framework proposed by [13] is used. This framework incorporates higher order statistics (e.g., first and second-order derivatives). Further, it allows us to generate different color constancy algorithms in a systematic way.

This paper is organized as follows: First, in Section 2, color constancy based on low-level image features is discussed. In Section 3, the concepts of natural image statistics and scene semantics are provided. In Section 4, several approaches are given to combine the different color constancy algorithms, and in Section 5, the methods are evaluated on a large data set containing over 11,000 images. Finally, in Section 6, the results and contributions of the paper are discussed in more detail.

2 COLOR CONSTANCY

In this paper, color constancy is defined as correcting an image that was taken under an unknown light source so that it appears to be taken under a canonical (often white) light source. This is computed by first estimating the color of the light source, followed by a transformation of the original image (e.g., [8]). To model this process, in Section 2.1, the image formation model is considered first. Then, in Section 2.2, several algorithms to estimate the color of the light source are presented. Finally, in Section 2.3, the model to transform images from one light source to another is discussed.

2.1 Reflection Model

The image values $\mathbf{f} = (f_R, f_G, f_B)^T$ for a Lambertian surface depend on the color of the light source $I(\lambda)$, the surface reflectance $S(\mathbf{x}, \lambda)$, and the camera sensitivity function $\boldsymbol{\rho}(\lambda) = (\rho_R(\lambda), \rho_G(\lambda), \rho_B(\lambda))^T$, where λ is the wavelength of the light and \mathbf{x} is the spatial coordinate (e.g., [8], [14], [15]):

$$f_c(\mathbf{x}) = m(\mathbf{x}) \int_{\omega} I(\lambda) \rho_c(\lambda) S(\mathbf{x}, \lambda) d\lambda, \quad (1)$$

where ω is the visible spectrum, $m(\mathbf{x})$ is Lambertian shading, and $c = \{R, G, B\}$. It is assumed that the scene is illuminated by one light source and the observed color of the light source \mathbf{e} depends on the color of the light source $I(\lambda)$ as well as the camera sensitivity function $\boldsymbol{\rho}(\lambda)$:

$$\mathbf{e} = \begin{pmatrix} e_R \\ e_G \\ e_B \end{pmatrix} = \int_{\omega} I(\lambda) \boldsymbol{\rho}(\lambda) d\lambda. \quad (2)$$

Color constancy can be achieved by estimating the color of the light source \mathbf{e} , given the image values of \mathbf{f} , followed by a transformation of the original image values using this illuminant estimate. This transformation will leave the intensity of every pixel unaltered as the proposed method will only correct for the chromaticity of the light source. More information on this transformation can be found in Section 2.3. Since both $I(\lambda)$ and $\boldsymbol{\rho}(\lambda)$ are, in general, unknown, the estimation of \mathbf{e} is an underconstrained problem that cannot be solved without further assumptions. Therefore, in practice, color constancy algorithms are based on simplifying assumptions such as restricted gamuts (limited number of image colors which can be observed under a specific illuminant), the distribution of colors that are present in an image (e.g., white patch, gray-world, etc.), and the set of possible light sources. In this paper, the focus is on the distribution of colors that are present in an image as the major assumption. In the next section, a framework is discussed generating different color constancy methods, where each method is based on a specific assumption about the presence of colors and color edges in images.

2.2 Illuminant Estimation

Two well-established color constancy algorithms, using pixel values, are based on the Retinex Theory proposed by Land and McCann [9]. The White-Patch algorithm [10] is based on the White-Patch assumption, i.e., the assumption that *the maximum response in the RGB-channels is caused by a white patch*. The Gray-World algorithm [11] is based on the Gray-World assumption, i.e., *the average reflectance in a scene under a neutral light source is achromatic*. In [12], these two algorithms are proven to be important instantiations of the Minkowski-norm:

$$\mathcal{L}_c(p) = \left(\int f_c^p(\mathbf{x}) d\mathbf{x} \right)^{\frac{1}{p}} = k e_c, \quad (3)$$

where $c = \{R, G, B\}$ and k is a multiplicative constant chosen such that the illuminant color, $\mathbf{e} = (e_R, e_G, e_B)^T$, has unit length. When $p = 1$ is substituted, (3) is equivalent to computing the average of $\mathbf{f}(\mathbf{x})$, i.e., $\mathcal{L}(1) = (\mathcal{L}_R(1), \mathcal{L}_G(1), \mathcal{L}_B(1))^T$ equals the Gray-World algorithm. When $p = \infty$, (3) results in computing the maximum of $\mathbf{f}(\mathbf{x})$, i.e., $\mathcal{L}(\infty)$ equals the White-Patch algorithm. In general, to arrive at a proper value, p is tuned for the data set at hand. Hence, the value of this parameter may vary for different data sets.

The assumptions of the above color constancy methods are based on the distribution of colors (i.e., pixel values) that are present in an image. The incorporation of higher order image statistics (in the form of image derivatives) is proposed in [13], where a framework is presented that incorporates the well-known methods like (3), as well as methods based on first and second-order statistics:

$$\left(\int \left| \frac{\partial^n f_{c,\sigma}(\mathbf{x})}{\partial \mathbf{x}^n} \right|^p d\mathbf{x} \right)^{\frac{1}{p}} = k e_c^{n,p,\sigma}, \quad (4)$$

where $|\cdot|$ indicates the Frobenius norm, $c = \{R, G, B\}$, n is the order of the derivative, and p is the Minkowski-norm. Further, the derivative is defined as the convolution of the

image with the derivative of a Gaussian filter with scale parameter σ [23]:

$$\frac{\partial^{s+t} f_{c,\sigma}}{\partial x^s \partial y^t} = f_c * \frac{\partial^{s+t} G_\sigma}{\partial x^s \partial y^t}, \quad (5)$$

where $*$ denotes the convolution and $s + t = n$.

Using (4), many different color constancy algorithms can be generated. For instance, algorithms based on zeroth-order statistics like Gray-World, White-Patch, and general Gray-World can be generated by substituting $n = 0$:

1. $e^{0,1,0}$ ($\equiv \mathcal{L}(1)$) is equivalent to the Gray-World algorithm,
2. $e^{0,\infty,0}$ ($\equiv \mathcal{L}(\infty)$) is equivalent to the White-Patch algorithm, and
3. $e^{0,p,\sigma}$ is called the general Gray-World algorithm, where the values for p and σ are dependent on the type of images that are in the data set. In [13], it is shown that for real-world images, $p = 13$ and $\sigma = 2$ are found to produce proper results for a real-world data set, i.e., $e^{0,13,2}$.

Equation (4) extends these instantiations to higher order statistics. For instance, when taking the first-order or second-order derivative, the values for the Minkowski-norm and the smoothing parameter are used to produce different algorithms:

4. $e^{1,p,\sigma}$ is the first-order Gray-Edge. The basic assumption here is that the average edge (i.e., derivative) is gray. The Minkowski-norm p and the smoothing parameter σ are dependent on the images in the data set. In [13], $p = 1$ and $\sigma = 6$ provide good results on real-world data, i.e., $e^{1,1,6}$.
5. $e^{2,p,\sigma}$ is the second-order Gray-Edge. The basic assumption here is that the average second-order derivative is gray. Again, the Minkowski-norm p and smoothing parameter σ depend on the data set used, where for a real-world data set, $p = 1$ and $\sigma = 5$ provide good results [13], i.e., $e^{2,1,5}$.

In conclusion, a wide variety of color constancy algorithms are obtained, corresponding to different instantiations of (4), where each color constancy method has its own basic assumption about the distribution of color values and edges in the image.

2.3 Diagonal Model

The focus of this paper is on estimating the color of the light source. However, in many cases, the color of the light source is of less importance than the appearance of the input image under a reference light (called the canonical light source). Therefore, the aim of most of the color constancy methods is to transform all colors of the input image, taken under an unknown light source, to colors as they appear under this canonical light source. This transformation can be considered to be an instantiation of chromatic adaptation, e.g., [6]. Chromatic adaptation is often modeled using a linear transformation, which, in turn, can be simplified to a diagonal transformation when certain conditions are met [24], [25], [26]. Other possible chromatic adaptation methods include linearized Bradford [27] and CIECAM02 [28].

In this paper, the diagonal transform or *von Kries Model* [29] is used, without changing the color basis [30] or applying spectral sharpening [25]. These latter techniques

are shown to be able to improve the quality of the output image with respect to the diagonal model, i.e., if the color of the light source is known, then these modified algorithms result in more realistic images than the diagonal model. For simplicity, these steps are omitted in this paper, so the diagonal model that is used is given by

$$\mathbf{f}_t = \mathcal{D}_{u,t} \mathbf{f}_u, \quad (6)$$

where \mathbf{f}_u is the image taken under an unknown light source, \mathbf{f}_t is the same image transformed so it appears as if it was taken under the canonical illuminant, and $\mathcal{D}_{u,t}$ is a diagonal matrix which maps colors that are taken under an unknown light source u to their corresponding colors under the canonical illuminant c :

$$\begin{pmatrix} R_c \\ G_c \\ B_c \end{pmatrix} = \begin{pmatrix} d_1 & 0 & 0 \\ 0 & d_2 & 0 \\ 0 & 0 & d_3 \end{pmatrix} \begin{pmatrix} R_u \\ G_u \\ B_u \end{pmatrix}. \quad (7)$$

Even though this model is merely an approximation of illuminant change and might not be accurately able to model photometric changes due to disturbing effects like highlights and interreflections, it is widely accepted as color correction model [24], [25], [26] and underpins many color constancy algorithms (e.g., the gamut mapping [14] and the used framework of methods [13]). The diagonal mapping is used throughout this paper to create output images after correction by a color constancy algorithm, where a perfect white light, i.e., $(\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}})^T$, is used as the canonical illuminant.

3 NATURAL IMAGE STATISTICS AND SCENE SEMANTICS

All methods that comprise the used color constancy framework [13] are based on assumptions on the distribution of colors (edges) that are present in an image. For instance, the Gray-World algorithm assumes that the average color in a scene taken under a neutral light source is achromatic, while the Gray-Edge algorithm assumes that the average *edge* is achromatic. It has also been shown that the incorporation of spatial dependencies between colors (e.g., edges) produces more constrained gamuts improving the accuracy of color constancy in general [17]. This means that the set of possible adjacent color values (i.e., color edges) in real-world images is more restricted than the set of possible pixel values. Hence, the use of local spatial information will provide more stable gamuts than pixel values to compute color constancy. Furthermore, a higher accuracy is obtained when there are a large variety of edges in a scene [17]. The same observation is valid for the Gray-World algorithm in terms of the number of different surfaces, e.g., [32], [33].

Hence, color constancy methods are largely dependent on the distribution of colors and color edges in an image. In the next section, image statistics are used to describe these distributions.

3.1 Spatial Image Structures

Image structures are valuable identification cues in determining which type of scene the image is taken from. In [34], the authors show that the power spectrum (distribution of edge responses) of an image is characteristic for the type of scene. Further, in [35], it is shown that this distribution of edge responses can be modeled by a Weibull distribution. In the context of scene classification, features derived from

the power spectrum and Weibull distributions have been successfully applied [36], [35], [37]. In this paper, we focus on modeling natural image statistics using the *two parameter integrated Weibull distribution* [35]:

$$w(x) = C \exp\left(-\frac{1}{\gamma} \left|\frac{x}{\beta}\right|^\gamma\right), \quad (8)$$

where x is the edge responses in a single-color channel to the Gaussian derivative filter, C is a normalization constant, $\beta > 0$ is the scale parameter of the distribution, and $\gamma > 0$ is the shape parameter. The parameters of this distribution are indicative for the edge statistics of an (natural) image. In fact, the *contrast* of the image is indicated by β (i.e., the width of the distribution) and the *grain size* by γ (i.e., the peakedness of the distribution). Hence, a higher value for β indicates more contrast, while a higher value for γ indicates a smaller grain size (more fine textures).

To fit the Weibull distribution, edge responses are computed by a Gaussian derivative filter. There exists a high correlation between the Weibull parameters that are fitted through the distribution of edges for the first derivative, second derivative, and third derivative. Hence, a single filter type, although measured in different orientations, is sufficient to assess the spatial statistics of images [35].

In Fig. 1, examples are shown of images with their corresponding edge distributions which are approximated by a Weibull-fit. The intensity channel is chosen for the ease of illustration because a six-dimensional edge distribution (i.e., β and γ for each R , G , and B channel) is hard to visualize. The edge distributions and corresponding Weibull-fits computed for separate color channels show similar plots. The images are examples on which the different color constancy algorithms using the corresponding type of information (i.e., pixel values, edges, or second-order transitions) performs best (based on the angular error that is explained in Section 5).

The relationship between the images in Fig. 1 and their corresponding color constancy algorithm becomes clear from the edge distributions that are shown together with the images in Fig. 1. Pixel-based algorithms (i.e., zeroth order) perform better than higher order methods (i.e., first and second order) on images with only little texture. This reflects in an edge distribution that densely sampled around the origin, i.e., many edges with little or zero energy. Higher order methods require more edge information for an accurate illuminant estimate, which is reflected in an edge distribution that is less sharply peaked. At this stage, it is important to note the semantic similarity between images that correspond to the same type of color constancy algorithms. For instance, forest-like scenes show a similar edge distribution in Fig. 1b and are all best solved by a first-order color constancy algorithm. Hence, scene semantics can steer the process of color constancy. Natural image statistics and scene semantics will therefore be used in the next sections to achieve a proper selection of color constancy algorithms.

4 COMBINATION OF ILLUMINANT ESTIMATION METHODS

In this section, a novel strategy is proposed based on natural image statistics to select the color constancy method which performs best for a specific image. To combine and compare different fusion strategies, in Section 4.1, a basic approach is discussed based on using the output of multiple




Fig. 1. Examples of images that can be considered to be characteristic of the corresponding color constancy algorithms, i.e., the corresponding color constancy algorithm will perform best on these types of images. Below each image, the distribution of edges in the intensity channel is plotted. The images come from the data set published in [31]. (a) Zeroth-order method. (b) First-order method. (c) Second-order method.

algorithms. Then, in Section 4.2, natural image statistics are used to identify the most important characteristics of color images. Based on these image characteristics, the proper color constancy algorithm is selected for a specific image. Finally, scene semantics are used to find a category-specific combination of color constancy algorithms.

4.1 Color Constancy Using Standard Fusion

When using the output of multiple algorithms to generate a new estimate of the illuminant, the simplest method is to take the average of the estimates over all algorithms. A straightforward extension is to take the weighted average of the estimated illuminants. If n algorithms are combined, then the weighted average is defined as

$$\bar{\mathbf{e}} = \sum_{i=1}^n w_i \mathbf{e}_i, \quad (9)$$

where $\sum_{i=1}^n w_i = 1$. The average is just a special instance of the weighted average: $w_1 = w_2 = \dots = w_n$. The estimates can also be combined using a nonlinear committee. However, in [21], it is shown that a nonlinear neural network did not produce better results than the weighted

average. In fact, the weighted average outperformed a multilayer Perceptron neural network.

In [22], two algorithms (one statistics-based and one physics-based algorithm) were combined using a similar approach. However, the output of the two used algorithms are somewhat different than the output of a general color constancy algorithm. Both methods produce a vector of probabilities, where each element represents the probability that the corresponding illuminant is the illuminant that was used to create the current image. In the combination phase, the weighted average of these two vectors is determined, after which the illuminant with the highest probability is selected to be correct. Since this method requires the output of the color constancy algorithms to comply to a specific (irregular) form, this approach is not further evaluated here.

4.2 Color Constancy Using Natural Image Statistics

In Section 3, the Weibull distribution is considered as the parameterization of the edge distribution of images. Several characteristics, like the number of edges and the amount of texture and contrast, are captured by this parameterization, i.e., β and γ . In this section, it is proposed to select different color constancy methods based on these statistics. In previous work [39], it is shown that applying the k -means clustering on the Weibull-features, combined with a Gaussian weighting function, provides proper color constancy. In this paper, the k -means approach is generalized to a probabilistic approach, corresponding to a maximum likelihood classifier based on mixture of Gaussians (MoGs). This provides a more principled and probabilistic basis than k -means to relate natural image statistics with color constancy.

This novel algorithm aims at combining the estimates of several color constancy algorithms into a single more accurate estimate. To be precise, let M be the set of algorithms that are to be combined, where M_i denotes algorithm i . Further, the accuracy of the estimate of algorithm i on image j (i.e., the performance of algorithm i on image j) is denoted by $\epsilon_i(j)$. The algorithm consists of the following steps:

- First, the image statistics $\omega \in \mathbb{R}^{p \times q}$ for all images are computed, where p is the number of features that are computed and q is the number of images, i.e., ω_{ij} is the i th feature of the j th image. For simplicity, the subscript i is omitted, so ω_j denotes the feature vector representing the image statistics of the j th image.
- Then, all images that are in the training set are labeled. The label y_j of an image j is derived using the performance of the algorithms on image j :

$$y_j = \arg \min_i \{\epsilon_i(j)\}. \quad (10)$$

- Apply the MoG-classifier [40] on the training data. The likelihood of the observed image statistics ω_j for image j given color constancy algorithm y_j is computed as a weighted sum of k Gaussian distributions:

$$p(\omega_j|y_j) = \sum_{m=1}^k \alpha_m G(\omega_j, \mu_m, \Sigma_m). \quad (11)$$

Here, α_m are the positive weights of the Gaussian components (with mean and variance defined as μ_m and Σ_m , respectively) such that $\sum_{m=1}^k \alpha_m = 1$. The

parameters of the model are learned through training using the Expectation-Maximization algorithm.

- Apply the learned MoG-classifier on the test data and assign to the current image j the algorithm that maximized the posterior probability.

Weibull-parameters can be computed for each R , G , and B channel separately. However, these color channels are highly correlated [41]. Therefore, the image is first transformed to a decorrelated color space before computing these parameters. To this end, the opponent color space is used:

$$O_1 = \frac{R - G}{\sqrt{2}}, \quad (12)$$

$$O_2 = \frac{R + G - 2B}{\sqrt{6}}, \quad (13)$$

$$O_3 = \frac{R + G + B}{\sqrt{3}}. \quad (14)$$

The selection of the most appropriate color constancy algorithm for the current image is done by computing the maximum posterior probability of the classifier. The corresponding color constancy algorithm is selected for the current image. The other algorithms are ignored.

In Fig. 2, an example of the output of the algorithm is shown. For this example, the Weibull-features representing the amount of texture and contrast (i.e., γ and β , respectively) that are computed using the gradient of the intensity channel are used, along with three different color constancy algorithms: a zeroth-order algorithm ($e^{0,\infty,0}$), a first-order algorithm ($e^{1,1,1}$), and a second-order algorithm ($e^{2,2,1}$). Again, intensity (i.e., O_3) is used for ease of illustration (instead of the six-dimensional space, i.e., β and γ for O_1 , O_2 , and O_3). Note that the final classifier is computed based on the complete six-dimensional feature space. In this figure, the different color constancy methods that are assigned to the images are visualized by different colors. It can be observed that images on which a certain algorithm *generally* performs well are grouped together.

4.3 Color Constancy Using Scene Semantics

Natural image statistics are known to provide identification cues for the classification of different types of scenes like forest, coast, and street [34]. Further, in [42], it is shown that using high-level visual information in the form of class-specific information improves the accuracy of the estimated illuminant. Van de Weijer et al. [42] assume that an image can be modeled as a mixture of semantic classes. The information on the different classes that are present in an image is used to estimate the color of the light source. In this section, we aim at using scene semantics to find a category-specific combination of color constancy algorithms that optimize the performance of the illuminant estimation.

Before discussing the category-specific combination method, the selection of the most appropriate color constancy algorithm based on scene semantics is examined first. In [38], a data set is provided consisting of eight urban and natural scene categories (e.g., Coast & Beach, Open country, Forest, Mountain, Highway, Street, City center, and Tall building). The corresponding Weibull-parameters of the images of a selection of these categories are shown in Fig. 3, along with the Weibull-parameters of the images that




Fig. 2. A scatter plot of the β and γ of the gradient in the O_3 -channel. Each point represents the Weibull-parameters of one image. The parameters of more than 11,000 images (the real-world set by Ciurea and Funt [31]) are plotted. The differently colored parts in the graphs represent clusters of images that are best solvable by a single-color constancy algorithm. The images are merely used as illustration and are taken from [38].

are derived from the real-world data set [31]. It can be observed that images from the same category have similar edge distributions, resulting in similar Weibull-parameters.

Some categories have a larger variance in edge distribution than others. For instance, most of the images of the category *Highway* have a low value for β and a low value for γ , indicating a low contrast and few edges. Images of the category *Mountain*, on the other hand, generally have a large variance. However, even for this category, it can be

observed that most images have higher values for β and γ , indicating higher contrast and more edges.

From these observations, a supervised selection of a color constancy algorithm for images from all scene categories can be achieved. By classifying an input image as one of these image categories (either supervised by user intervention or unsupervised by a scene recognition system like [36], [37], [43]), the corresponding color constancy algorithm can be applied to the image to obtain a




Fig. 3. Scatter plots of the Weibull-parameters based on O_3 derived from images coming from several categories (defined in [38]), overlaid on the Weibull-parameters of the images that are in the real-world set [31]. From this plot, it can be seen that images from the same scene category have similar Weibull-parameters.




Fig. 4. Examples of images that are in the two data sets that are used in this paper. The first data set consists of images that are generated using surface reflectance spectra combined with illuminant spectra. The second data set consists of real-world images [31]. (a) Mondrian images. (b) Real-world images.

performance that is similar to the proposed automatic selection algorithm.

5 EXPERIMENTS

In this section, the proposed method is evaluated and compared to state-of-the-art methods on large-scale data sets. First, the data sets that are used for training and testing are introduced in Section 5.1 and the performance measure is discussed in Section 5.2. In Section 5.3, the performance of several state-of-the-art algorithms on a large-scale data set is shown, and in Sections 5.4 and 5.5, the performance of the standard and the proposed fusion algorithms is discussed and analyzed. Finally, in Section 5.6, experiments are performed using scene semantics.

5.1 Data Set

A training data set is created based on spectral reflectance data presented in [44]. This data set originally comprises only surface and illuminant spectra, which are first combined into (R, G, B) -values. Then, using these generated pixel colors, several Mondrian-like images are created which all have different properties in the number of edges and the amount of texture and contrast. Since only material surfaces are present in the original data set, shadow gradients are added to several images to enlarge their photometrical variety. Note that the resulting images contain up to tens of different surfaces, and hence, many different transitions, simulating the statistics of real-world images as close as possible. A few examples are shown in Fig. 4a. This data set will be called the Mondrian data set in the remainder of the paper and will be made available together with the source-code of the proposed algorithm (<http://www.colorconstancy.com>).

As a second set, the color constancy data set of Ciurea and Funt [31] is used. In this data set, over 11,000 images are present, extracted from 2 hours of video recorded under a large variety of imaging conditions (including indoor, outdoor, desert, cityscape, and other settings). In total, the images are divided into 15 different clips taken at different locations. Some example images are shown in Fig. 4b and some example results of various algorithms applied to the images in this data set are shown in Fig. 7. The data set provides the ground truth of the color of the illuminant.

This ground truth is acquired by making use of the small gray sphere in the bottom right corner of the images. Note that this gray sphere is omitted while estimating the illuminant in the experiment. Further, since some of the images in this set are rather correlated due to the nature of the set, we make sure that, if applicable, this correlation is taken into account when separating the data set, e.g., for training purposes. Hence, in this way, images, or similar images, that are used for testing a method do not appear in the training set.

One of the basic assumptions in machine learning is that the distribution of the test data should be similar to the distribution of the training data. Hence, the variety of images used for training the algorithm should be similar to the variety of the images that are used to test it. To this end, we propose two different experiments. The first experiment consists of a training phase using the Mondrian-like images shown in Fig. 4a and a test phase using the real-world data shown in Fig. 4b. In this way, the data sets for training and testing are completely different. This scenario reflects the case when little information on the data set (used for testing the method) is known. The second experiment is based on cross validation: The data will be divided into 15 parts, where we make sure that correlated images are grouped in the same part. This means that images that are taken from the same video shot are *not* separated. Next, the method is trained on 14 parts of the data and tested on the remaining part. This procedure is repeated 15 times, so every images are in the test set exactly once and all images from one shot will either be in the training set or in the test set.

5.2 Performance Measure

For all images in the data set, the correct color of the light source e_l is known a priori. To measure how close the estimated illuminant resembles the true color of the light source, the angular error ϵ is used:

$$\epsilon = \cos^{-1} \left(\frac{e_l \cdot e_e}{\|e_l\| \cdot \|e_e\|} \right), \quad (15)$$

where $e_l \cdot e_e$ is the dot product of the two vectors representing the true color of the light source e_l and the estimated color of the light source e_e and $\|\cdot\|$ indicates the euclidean norm. To measure the performance of an algorithm on a data set, the median angular error is reported as summarizing statistic, as this is considered to be the most appropriate measure [33], [45].

5.3 Color Constancy Algorithms

In Table 1, the results for several (state-of-the-art) algorithms are shown for the real-world data set of Ciurea and Funt [31]. Results of the Gray-World algorithm are comparable to not applying any method at all (i.e., the estimate of the illuminant is set to white), while the White-Patch performs slightly better. Instantiations of the framework that have been previously reported as best-performing on real-world data [13] are also tested on the same data set. The results for $e^{0.13,2}$, $e^{1,1.6}$, and $e^{2,1.5}$ are shown in Table 1 and the three instantiations have similar performance on this data set, outperforming the Gray-World and the White-Patch algorithms. However, choosing a more appropriate set of parameters can significantly improve the results (denoted by “Best . . .” in Table 1). Further, the performance of these more carefully selected parameters confirms the findings of van de Weijer et al. [13], where it is concluded

TABLE 1
Median Angular Errors for Several (Single) Algorithms
on the Complete Data Set of 11,346 Images

Method	Median
Do nothing (<i>i.e.</i> without color constancy)	6.7°
Grey-World ($e^{0,1,0}$)	7.0°
White-Patch ($e^{0,\infty,0}$)	5.3°
General Grey-World ($e^{0,13,2}$)	5.5°
1 st -order Grey-Edge ($e^{1,1,6}$)	5.2°
2 nd -order Grey-Edge ($e^{2,1,5}$)	5.2°
Best 0 th -order algorithm: $e^{0,9,0}$	5.3°
Best 1 st -order algorithm: $e^{1,1,1}$	4.6°
Best 2 nd -order algorithm: $e^{2,1,2}$	4.9°
Gamut mapping	
- trained on: <i>Mondrian data</i>	5.5°
- trained on: <i>real-world data</i>	4.8°

that the first-order Gray-Edge algorithm outperforms other algorithms on the real-world data. Note, however, the difference in performance between the two parameter settings, while both settings are found to be optimal for a certain data set. The fact that, for this data set, a different optimal parameter setting exists than for the data sets used in [13] confirms the need for an algorithm to dynamically tune the parameters for the different instantiations.

For the experiments based on gamut mapping, the same implementation as in [32], [46] is used.¹ The software provides an option to apply segmentation as a preprocessing step, and several different parameter settings are tested (the reported results are computed by using the best parameter settings obtained for data set used). Further, experiments using two different canonical gamuts are performed. The first experiment uses a canonical gamut that is constructed using a large Mondrian data set [44], and the second experiment uses a canonical gamut that is constructed on real-world data.

To avoid overfitting by learning on the test set, for the second experiment, the same setup is used that was described earlier as cross validation. The data set is divided into 15 parts, keeping all images that are extracted from the same clip together. The canonical gamut is constructed on 14 parts, and the remaining part of the data is used for testing. This procedure is repeated 15 times, excluding each scene exactly once from the training data to use it as test data.

As expected, the best results are obtained when using the real-world data to construct the canonical gamut. When using the Mondrian data to calibrate the canonical gamut, the median angular error is 5.5 degrees. This is close to the performance of the White-Patch algorithm. However, by learning the canonical gamut from the real-world data, the error drops to a respectable 4.8 degrees. However, the best-performing algorithm remains the first-order Gray-Edge. Note that the performance of the gamut mapping too is obtained by optimizing the segmentation parameters that can be set using the gamut software provided by [32], [46], and hence, the performance is considered to be best-case.

In Fig. 5, the angular errors of all the methods are shown, together with error bars indicating a confidence interval of 95 percent. These error bars are obtained using resampling:




Fig. 5. Median angular errors for all the methods plotted with a 95 percent confidence interval.

Let ϵ^A be the set of error measurements for a method A on all images of the data set. By taking n bootstrap samples (*i.e.*, samples drawn at random with replacement) from this set of measurements, n sets of error measurements are obtained. For every set of error measurements ϵ_i^A , the mean and median can be computed, resulting in n values for the mean and n values for the median (*i.e.*, distributions for the mean and median). Finally, a p percent confidence interval is obtained from the $p/2$ and $1 - p/2$ quantiles of these distributions [33].

5.4 Color Constancy Using Standard Fusion

To evaluate the different combination strategies, the algorithms of [13] are used as instantiations of combination methods, as described in Section 4. Therefore, the best-performing instantiation is used as a baseline to compare the algorithms with.

A simple average of the output results does not seem to improve the results as compared to the baseline algorithm (see Table 2): The median angular error is slightly worse. Using a weighted average instead of a simple average levels the median angular error with the best-performing single algorithm. The main advantage of this fusion method is that there are less outliers. This results in a slightly better mean angular error (5.5 over 5.7 degrees).

TABLE 2
The Performance of Several Different Methods
to Combine the Output of Single-Color Constancy Methods

Method	Median
Baseline ($e^{1,1,1}$)	4.6°
Simple average	5.0° (+8.7%)
Weighted average	4.6° ($\pm 0\%$)
Proposed: NIS	4.2° (-8.7%)
Proposed: NIS - cross-validation	3.7° (-19.6%)

The color constancy methods used include one zeroth-order method (General Gray-World $e^{0,5,1}$), one first-order method (Gray-Edge $e^{1,1,2}$), and one second-order method (second-order Gray-Edge $e^{2,1,1}$). The entry NIS-Mondrian denotes the proposed algorithm trained on the independent Mondrian data set and tested on the real-world data set, and the entry NIS real world denotes the proposed algorithm evaluated using cross validation on the real-world data set.

1. http://kobus.ca/research/programs/colour_constancy.

5.5 Color Constancy Using Natural Image Statistics

The proposed fusion algorithm, learned on the Mondrian data set and tested on the real-world data set (denoted by NIS in Table 2), performs significantly better than the baseline and the weighted average: An improvement of the median angular error of nearly 10 percent is obtained. In fact, in this experiment, test data are unknown and correspond to the most generic approach of the proposed method. No learning step is required for a new data set since the resulting classifier is independent of the test data.

Results of the second experiment, denoted by NIS cross validation in Table 2, correspond to the situation where the circumstances under which the system is used are known a priori. Training is done using part of the real-world data, and testing is done using an independent part, where images that are similar are either in the training set or in the test set. Hence, no similarities exist between the training data and test data. Results show a significant improvement over the first case (with no a priori knowledge): The median angular error drops to 3.7 degrees, an improvement of nearly 20 percent over the baseline method. In this experiment, test data are known, making the proposed method less generic. Hence, a learning step is required for each new data set.

The proposed algorithm requires data with enough variety to train the method. However, from the experiments, it can be concluded that the proposed method can be trained using a completely independent training set, and still perform significantly better than the baseline algorithm. Also, note that this baseline algorithm is the best-case performance of the single methods since it was manually selected among the single algorithms as the color constancy method with the highest performance on the current data set. Still, the worst-case performance of the proposed method (NIS-Mondrian in Table 2) is significantly higher.

SNR performance analysis. In order to analyze the signal-to-noise sensitivities of pixel-based, edge-based, and higher order statistics-based methods, the following experiment is conducted: A Mondrian-like image is created using eight patches with different colors, keeping the difference between two consecutive colors fixed. After that, an illuminant is selected at random from the set of light sources proposed by Barnard et al. [44]. Based on the diagonal model, this illuminant is used to transform the image so that it appears to be taken under a different light source. Finally, the three-color constancy algorithms are used to estimate the illuminant and the performance of these algorithms is evaluated with respect to the signal-to-noise ratio (SNR). The signal-to-noise ratio is defined by the average difference between two consecutive patches a_i and a_j :

$$\text{SNR} = \frac{a_i - a_j}{\sigma_{\text{noise}}}, \quad (16)$$

where σ_{noise} is kept fixed.

In Fig. 6, the sensitivities of the zeroth-order, the first-order, and the second-order algorithm (i.e., White-Patch, first-order Gray-Edge, and second-order Gray-Edge, respectively) are shown with respect to the SNR. As expected, it can be derived that for a higher signal-to-noise ratio, on average, the accuracy of all algorithms is improved. However, it is also shown that the performance of the pixel-based method degrades less than edge-based and higher order statistics-based methods when the SNR decreases. For higher values of the SNR, the first and second-order Gray-Edge algorithms outperform the pixel-based method; the performance of the




Fig. 6. Experiment monitoring the sensitivity of the different color constancy algorithms to the signal-to-noise ratio. It can be derived that all methods perform better as the SNR increases. Further, first-order and second-order Gray Edges have a higher error than the zeroth-order method for low signal-to-noise ratios, but a lower error for higher signal-to-noise ratios.

three algorithms converge when the SNR reaches very high values.

To conclude, the Gray-Edge algorithms (first and second order) provide better performance than the pixel-based one, but they need a medium to high SNR. Pixel-based methods obtain better performance for a low SNR. In other words, edge-based is preferred when the image contains a reasonable number of (sufficiently) contrasted edges. Otherwise, pixel-based color constancy is preferred. This principle is reflected by the Weibull-distribution parameters contrast β and grain size γ , and forms the basis of the proposed algorithm. The MoG classifier learns the correlation between these parameters.

5.6 Color Constancy Using Scene Semantics

Color constancy using scene semantics is evaluated on the real-world data set as follows: Images from the real-world data set are annotated to obtain images that can be classified as coming from the same scene category. In total, 75 images from five clips of the complete data set (15 images per category) are annotated as *open country* and 70 images from seven clips (10 per category) as *street*, *indoor*, and *forest*. During annotation, the selected images are different enough to avoid the overfitting problem. The performance of the proposed method using scene semantics is compared to the combination method using natural image statistics.

In Table 3, the results are shown. In this table, *scene semantics* refers to applying a one-color constancy algorithm to all images of a specific scene category. This one-color constancy algorithm is selected in a supervised manner, based on the Weibull-parameters of images from that scene category. It can be observed that for some categories, the performance is similar to the method using natural image statistics, even though only one-color constancy method is applied to all images of the same scene category.

6 DISCUSSION

Natural image statistics are indicative for the performance of the different color constancy algorithms. For instance, the Gray-Edge algorithm performs better on images with a



Fig. 7. Examples of images that are in the data set used for evaluation in Section 5 and results of some color constancy algorithms. The corresponding angular errors are shown in the lower right corner of the images. (a) Original images. (b) Ideal correction. (c) Correction using natural image statistics. (d) Correction using Gray-Edge (baseline). (e) Correction using White-Patch. (f) Correction using Gray-World.

reasonable number of edges than on images with only a few edges. Moreover, if the contrast of these edges is rather low, then the accuracy of the estimate, in general, will decrease. The effect of the SNR is analyzed in Section 5.5. From this analysis, it is derived that if the SNR is high, then it is beneficial to use methods based on higher order statistics. However, the performance of such methods decreases as the SNR decreases, resulting in a preference for methods based on first-order statistics for images with a medium SNR. For a very low SNR (i.e., no edge information), pixel-based methods are preferred.

This principle is also reflected in the resulting classifier that is learned on the Mondrian data set. From the analysis of this classifier, it is observed that the method based on second-order statistics is preferred, provided that there is enough information to give an accurate estimate. This means that the distribution of the second-order statistics is balanced and not easily biased toward noisy measurements. Since the available information of second-order statistics decreases rapidly, depending on the content of an image, the performance of color constancy methods based on second-order statistics decreases also. Hence, color constancy

TABLE 3
Median Angular Error for Several Different Methods
on Images from Four Categories

Open Country	Median
Entire image (NIS-Mondrian)	5.5°
Proposed: Scene Semantics	6.7°
Street	Median
Entire image (NIS-C)	2.7°
Proposed: Scene Semantics	3.4°
Indoor	Median
Entire image (NIS-C)	5.1°
Proposed: Scene Semantics	5.5°
Forest	Median
Entire image (NIS-C)	5.6°
Proposed: Scene Semantics	5.5°

methods using edge information are preferred for images with a medium amount of information, and pixel-based methods are preferred when the amount of information becomes low. The amount of information is dependent on the content of an image, and natural image statistics parameterize the amount of available information. For instance, a high contrast of colors (i.e., the β of the edge distributions of the O_1 and O_2 channels) combined with only a few edges (i.e., a high value for γ) indicates that there are only a few edges with enough energy and that the estimate that is based on this information is likely to be biased toward the color of the high contrast. Consequently, in such situations, it is better to select pixel-based color constancy methods. On the other hand, when the value of γ is low, this means that there are relatively many edges with a strong response, making the estimates of methods based on first and second-order statistics more reliable.

In summary, the three basic indicators for the performance of color constancy methods based on pixel and edge information are:

- Image features. The set of possible adjacent color values (i.e., color derivatives) in real-world images is more restricted than the set of possible pixel values. Hence, the use of derivatives provides more stable gamuts than pixel values alone to compute color constancy. The order of best performance in terms of image features is then second-order statistics first, followed by first-order and zeroth-order statistics.
- Number of features. A higher color constancy accuracy is obtained when there is a large variety of different edges in an image [17]. The same argument holds for the Gray-World algorithm for the number of different surfaces in a scene, e.g., [32], [33]. From [17], it is derived that if the image contains a limited number of edges, pixel-based color constancy is preferred. In case of sufficient edges (e.g., more than eight different edges), edge-based color constancy is preferred. The order of the best performance in terms of the number of edges (from low to high) is then zeroth-order statistics first, followed by first-order and second-order statistics.
- SNR. Pixel-based methods obtain a better performance for a low SNR. The order of best performance

in terms of the amount of SNR (from low to high) is then zeroth-order statistics first, followed by first-order and second-order statistics.

The above principles are reflected by the Weibull-distribution parameters contrast β and grain size γ . The amount of texture and the number of edges and contrast are captured by this parameterization. The proposed algorithm exploits this information to select and combine the different order color constancy methods as follows:

- Image features. The value γ corresponds to grain size and is related to the amount of texture.
- Number of features. The value γ corresponds grain size and is related to the number of edges.
- SNR. The value β corresponds to contrast and is related to SNR.

The MoG classifier computes the correlation between the values of the image features, number of features, and SNR and the values of the Weibull-distribution parameters contrast β and grain size γ .

7 CONCLUSION

In this paper, the Weibull parameterization (texture and contrast) has been used to identify the most important characteristics of color images. It is shown that the Weibull parameterization (grain size and contrast) is related to the number of edges, amount of texture, and signal-to-noise ratio to which the used color constancy methods are sensitive. An MoG-classifier has been used to learn the correlation and weighting between the Weibull-parameters (grain size and contrast) and these image attributes (number of edges and amount of texture and SNR). The output of the classifier is the selection of the best performing color constancy method for a certain image.

Experimental results show a large improvement over state-of-the-art single algorithms. On a data set consisting of more than 11,000 images, the best-performing single algorithm is found to be the first-order Gray-Edge. Comparing the median angular error of this algorithm with our proposed algorithm, an increase of nearly 20 percent can be obtained when the circumstances under which the algorithm will be used are known a priori. When this is not the case, i.e., when no information about the test images is known, a prelearned system can be used that is trained on the Mondrian set. This system only needs to be trained once and is completely independent of the test data. Using this classifier, the median angular error still decreases significantly (nearly 10 percent) compared to the best-case performance of the single algorithms. Finally, we showed that for certain scene categories, one specific color constancy algorithm can be used instead of the classifier considering several algorithms.

REFERENCES

- [1] P.B. Delahunt and D.H. Brainard, "Does Human Color Constancy Incorporate the Statistical Regularity of Natural Daylight?" *J. Vision*, vol. 4, no. 2, pp. 57-81, 2004.
- [2] D.H. Foster, K. Amano, and S.M.C. Nascimento, "Color Constancy in Natural Scenes Explained by Global Image Statistics," *Visual Neuroscience*, vol. 23, nos. 3/4, pp. 341-349, 2006.
- [3] J. Yang, R. Stiefelhagen, U. Meier, and A. Waibel, "Visual Tracking for Multimodal Human Computer Interaction," *Proc. SIGCHI*, pp. 140-147, 1998.

- [4] T. Gevers and A. Smeulders, "Color-Based Object Recognition," *Pattern Recognition*, vol. 32, no. 3, pp. 453-464, 1999.
- [5] T. Gevers and A.W.M. Smeulders, "Pictoseek: Combining Color and Shape Invariant Features for Image Retrieval," *IEEE Trans. Image Processing*, vol. 9, no. 1, pp. 102-119, Jan. 2000.
- [6] M.D. Fairchild, *Color Appearance Models*, second ed. John Wiley & Sons, 2005.
- [7] M. Ebner, *Color Constancy*. John Wiley & Sons, 2007.
- [8] S.D. Hordley, "Scene Illuminant Estimation: Past, Present, and Future," *Color Research and Application*, vol. 31, no. 4, pp. 303-314, 2006.
- [9] E.H. Land and J.J. McCann, "Lightness and Retinex Theory," *J. Optical Soc. of Am. A*, vol. 61, pp. 1-11, 1971.
- [10] E.H. Land, "The Retinex Theory of Color Vision," *Scientific Am.*, vol. 237, no. 6, pp. 108-128, Dec. 1977.
- [11] G. Buchsbaum, "A Spatial Processor Model for Object Colour Perception," *J. Franklin Inst.*, vol. 310, no. 1, pp. 1-26, July 1980.
- [12] G.D. Finlayson and E. Trezzi, "Shades of Gray and Colour Constancy," *Proc. IS&T/SID Color Imaging Conf.*, pp. 37-41, 2004.
- [13] J. van de Weijer, T. Gevers, and A. Gijsenij, "Edge-Based Color Constancy," *IEEE Trans. Image Processing*, vol. 16, no. 9, pp. 2207-2214, Sept. 2007.
- [14] D.A. Forsyth, "A Novel Algorithm for Color Constancy," *Int'l J. Computer Vision*, vol. 5, no. 1, pp. 5-36, 1990.
- [15] G.D. Finlayson, S.D. Hordley, and P.M. Hubel, "Color by Correlation: A Simple, Unifying Framework for Color Constancy," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 23, no. 11, pp. 1209-1221, Nov. 2001.
- [16] G.D. Finlayson, S.D. Hordley, and I. Tastl, "Gamut Constrained Illuminant Estimation," *Int'l J. Computer Vision*, vol. 67, no. 1, pp. 93-109, 2006.
- [17] A. Gijsenij, T. Gevers, and J. van de Weijer, "Generalized Gamut Mapping Using Image Derivative Structures for Color Constancy," *Int'l J. Computer Vision*, vol. 86, nos. 2/3, pp. 127-139, 2010.
- [18] D.H. Brainard and W.T. Freeman, "Bayesian Color Constancy," *J. Optical Soc. of Am. A*, vol. 14, pp. 1393-1411, 1997.
- [19] M. D'Zmura, G. Iverson, and B. Singer, "Probabilistic Color Constancy," *Geometric Representations of Perceptual Phenomena*, pp. 187-202, Lawrence Erlbaum Assoc., 1995.
- [20] M. Ebner, "Evolving Color Constancy," *Pattern Recognition Letters*, vol. 27, no. 11, pp. 1220-1229, 2006.
- [21] V.C. Cardei and B.V. Funt, "Committee-Based Color Constancy," *Proc. IS&T/SID Color Imaging Conf.*, pp. 311-313, 1999.
- [22] G. Schaefer, S.D. Hordley, and G.D. Finlayson, "A Combined Physical and Statistical Approach to Colour Constancy," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 148-153, 2005.
- [23] W.T. Freeman and E.H. Adelson, "The Design and Use of Steerable Filters," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 9, pp. 891-906, Sept. 1991.
- [24] G. West and M.H. Brill, "Necessary and Sufficient Conditions for Von Kries Chromatic Adaptation to Give Color Constancy," *J. Math. Biology*, vol. 15, no. 2, pp. 249-258, 1982.
- [25] G.D. Finlayson, M.S. Drew, and B.V. Funt, "Spectral Sharpening: Sensor Transformations for Improved Color Constancy," *J. Optical Soc. of Am. A*, vol. 11, no. 5, pp. 1553-1563, 1994.
- [26] B.V. Funt and B.C. Lewis, "Diagonal versus Affine Transformations for Color Correction," *J. Optical Soc. of Am. A*, vol. 17, no. 11, pp. 2108-2112, 2000.
- [27] K.M. Lam, "Metamerism and Colour Constancy," PhD dissertation, Univ. of Bradford, 1985.
- [28] N. Moroney, M.D. Fairchild, R.W.G. Hunt, C. Li, M.R. Luo, and T. Newman, "The Ciecarn02 Color Appearance Model," *Proc. IS&T/SID Color Imaging Conf.*, pp. 23-27, 2002.
- [29] J. von Kries, "Influence of Adaptation on the Effects Produced by Luminous Stimuli," *Sources of Color Vision*, D. MacAdam, ed., pp. 109-119, MIT Press, 1970.
- [30] H.Y. Chong, S.J. Gortler, and T. Zickler, "The Von Kries Hypothesis and a Basis for Color Constancy," *Proc. Int'l Conf. Computer Vision*, pp. 1-8, 2007.
- [31] F. Ciurea and B.V. Funt, "A Large Image Database for Color Constancy Research," *Proc. IS&T/SID Color Imaging Conf.*, pp. 160-164, 2003.
- [32] K. Barnard, V.C. Cardei, and B.V. Funt, "A Comparison of Computational Color Constancy Algorithms; Part I: Methodology and Experiments with Synthesized Data," *IEEE Trans. Image Processing*, vol. 11, no. 9, pp. 972-984, Sept. 2002.
- [33] S.D. Hordley and G.D. Finlayson, "Reevaluation of Color Constancy Algorithm Performance," *J. Optical Soc. of Am. A*, vol. 23, no. 5, pp. 1008-1020, May 2006.
- [34] A. Torralba and A. Oliva, "Statistics of Natural Image Categories," *Network: Computation in Neural Systems*, vol. 14, no. 3, pp. 391-412, 2003.
- [35] J.-M. Geusebroek and A.W.M. Smeulders, "A Six-Stimulus Theory for Stochastic Texture," *Int'l J. Computer Vision*, vol. 62, nos. 1/2, pp. 7-16, 2005.
- [36] A. Torralba, "Contextual Priming for Object Detection," *Int'l J. Computer Vision*, vol. 53, no. 2, pp. 169-191, 2003.
- [37] J.C. van Gemert, J.M. Geusebroek, C.J. Veenman, C.G.M. Snoek, and A.W.M. Smeulders, "Robust Scene Categorization by Learning Image Statistics in Context," *Proc. IEEE CVPR Workshop Semantic Learning Applications in Multimedia*, June 2006.
- [38] A. Oliva and A. Torralba, "Modeling the Shape of The Scene: A Holistic Representation of the Spatial Envelope," *Int'l J. Computer Vision*, vol. 42, no. 3, pp. 145-175, 2001.
- [39] A. Gijsenij and T. Gevers, "Color Constancy Using Natural Image Statistics," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 1-8, June 2007.
- [40] C.M. Bishop, *Neural Networks for Pattern Recognition*. Oxford Univ. Press, 1996.
- [41] D.L. Ruderman, T.W. Cronin, and C.C. Chiao, "Statistics of Cone Responses to Natural Images: Implications for Visual Coding," *J. Optical Soc. of Am. A*, vol. 15, no. 8, pp. 2036-2045, 1998.
- [42] J. van de Weijer, C. Schmid, and J.J. Verbeek, "Using High-Level Visual Information for Color Constancy," *Proc. Int'l Conf. Computer Vision*, 2007.
- [43] J.-M. Geusebroek, "Compact Object Descriptors from Local Colour Invariant Histograms," *Proc. British Machine Vision Conf.*, vol. 3, pp. 1029-1038, 2006.
- [44] K. Barnard, L. Martin, B.V. Funt, and A. Coath, "A Data Set for Color Research," *Color Research and Application*, vol. 27, no. 3, pp. 147-151, 2002.
- [45] A. Gijsenij, T. Gevers, and M. Lucassen, "A Perceptual Analysis of Distance Measures for Color Constancy Algorithms," *J. Optical Soc. of Am. A*, vol. 26, no. 10, pp. 2243-2256, 2009.
- [46] K. Barnard, L. Martin, A. Coath, and B.V. Funt, "A Comparison of Computational Color Constancy Algorithms; Part II: Experiments with Image Data," *IEEE Trans. Image Processing*, vol. 11, no. 9, pp. 985-996, Sept. 2002.



Arjan Gijsenij received the MSc degree in computer science from the University of Groningen in 2005, and the PhD degree from the University of Amsterdam. He is currently a postdoctoral researcher in the Intelligent Systems Lab, University of Amsterdam. His main research interests include color in computer vision and psychophysics, (color) image processing, visual categorization, and statistical pattern recognition. He has served on the program committees of several major conferences and workshops. He is a member of the IEEE.



Theo Gevers is an associate professor of computer science at the University of Amsterdam, The Netherlands, where he is also the teaching director of the MSc in Artificial Intelligence. He currently holds a VICI Award (for research excellence) from the Dutch Organisation for Scientific Research. 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. He is the chair for various conferences and is an associate editor for the *IEEE Transactions on Image Processing*. Further, he is a program committee member for a number of conferences, and an invited speaker at major conferences. He is a lecturer delivering postdoctoral courses given at various major conferences (CVPR, ICPR, SPIE, and CGIV). He is a member of the IEEE.