

A Common Framework for Ambient Illumination in the Dichromatic Reflectance Model

Christian Riess Johannes Jordan
Elli Angelopoulou

Chair of Pattern Recognition, Department of Computer Science,
Friedrich-Alexander-University Erlangen-Nuremberg,
Martensstr. 3, 91058 Erlangen, Germany

{riess, jordan, elli}@i5.informatik.uni-erlangen.de

Abstract

The dichromatic reflectance model has been successfully applied on different tasks in color research, such as color constancy and highlight or shadow segmentation. In its original version, it incorporates only one direct illuminant. In this work, we analyze a recently published model, the Bi-Illuminant Dichromatic Reflectance Model (BIDR) proposed by Maxwell et al., that extends the dichromatic reflectance model for a very general ambient term. The proposed method can represent optical phenomena like inter-reflections and inhomogeneous ambient light. We will show that this new model is sufficiently general, in the sense that it encompasses established variations and extensions to the dichromatic reflectance model.

1. Introduction

Understanding scene reflectance provides useful insights for many computer vision tasks. Complex lighting environments introduce difficult problems in higher level vision tasks such as object recognition or tracking. The presence of different illumination conditions, highlights, shadows or interreflections may change the appearance of objects dramatically. With insights from reflectance analysis, a way is sought to make these tasks more robust to a variety of real-world conditions.

Many approaches to reflectance analysis build on machine learning, e.g. [1, 4, 8, 17]. By definition, the success of these methods is closely tied to the employed training set. In comparison, physics-based approaches try to explicitly model reflectance properties according to their physical behavior. One of the most prominent models in physics-based reflectance analysis is the dichromatic reflectance model by Shafer [22]. For a single dominant light source, it assumes that irradiance can be decomposed into a linear combination

of specular and diffuse reflectance components. Due to its relative simplicity and usefulness, many researchers developed algorithms that built on top of this model, e.g. [14, 23].

However, the most widely used physics-based methodologies are difficult to transfer to real-world scenarios: since reflectance phenomena are explicitly modeled, the problem of estimating scene properties is typically underconstrained. One real-world example is the presence of shadows, where a model with only one dominant light source fails. Shadow areas are typically assumed to be illuminated by ambient light, which can be treated as a second, extended light source with special properties. For example in outdoor scenes, the sky acts as a second light source, which shifts the illuminant color approximately towards blue due to Rayleigh scattering [6].

Since shadows are a common error source for higher level computer vision tasks, several groups addressed the problem of shadow segmentation or removal. One approach is to deduce rules for shadow regions from observations. Gershon *et al.* related ambient reflected intensities and diffuse reflected intensities [11]. Funka-Lea and Bajcsy extracted several cues for shadow and penumbra detection [9]. Barnard and Finlayson [3] proposed a methodology for detecting shadow edges from a database of intensity ratios between shadow and non-shadow areas. Shadow removal is done by Finlayson *et al.* with a calibrated camera [7]. Other shadow removal work, like that of Levine and Bhat-tacharyya [17], is based on machine learning. Finally, there is also a large body of work on shadow invariant color features, e.g. by Gevers and Smeulders *et al.* [12], Salvador *et al.* [21] and van de Weijer and Schmid [25].

Multiple illuminants have sometimes been explicitly exploited for the recovering of scene information. For these methods, two images are taken: the first one being only illuminated by ambient light, the second one with an additional flash light that is mounted on the camera. By studying the

relationship of the two differently illuminated scenes, under few additional constraints, the color of the scene illuminant and the albedos of the objects in the scene can be estimated [5, 18]. The core of these methods is the computation of a synthetic, flash-only illuminated image from the two input images. From the different reflectance behavior in the flash-only and the ambient illuminated image, conclusions about the missing scene information can be drawn. Despite the ubiquity of images that involve multiple illumination sources, the single light source model is still extensively used.

In this work, we analyze a recently published extension by Maxwell *et al.* [19] of the dichromatic reflectance model [22], the Bi-Illuminant Dichromatic Reflectance Model (BIDR). It adds a comprehensive ambient term to the original dichromatic reflectance model. The model can also integrate a representation for interreflections in the dichromatic reflectance model. Although interreflections have previously been studied by several authors (see e.g. [2, 16, 10]), most often they are ignored due to their great complexity. We assume that with the BIDR, it is possible to consider their effects within a well-established physics-based reflectance model.

To the best of our knowledge, although Maxwell *et al.* proposed the BIDR, the model has not yet been fully analyzed. In this work, we examine the side effects caused by the addition of a general ambient term. For many typical computer vision tasks, only a single image with a mixture of ambient and direct light is available. Hence, it is crucial to understand the reflectance of an image which is illuminated by ambient and direct light at the same time. As shown in Fig. 1, understanding the properties of the BIDR could help in finding new algorithms for applications such as shadow segmentation, material classification and color constancy.

We will also show how commonly used simplifications on the ambient term can be incorporated into the extended model. We will demonstrate the effects of several parameters in the extended model, and show a connection between the extended model and the shadow detection or removal algorithms by Barnard and Finlayson [3], Maxwell *et al.* [19] and Lu and Drew [18] as practical ways to overcome the inherent difficulties with the general formula. This imposes also a taxonomic structure on the existing algorithms with respect to the BIDR, as sketched in Fig. 1.

The outline of the paper is as follows: In section 2, we introduce the dichromatic reflectance model and its extension for a comprehensive ambient term. In section 3, we show how the different factors in the extended model influence each other and what conclusions can be drawn within this model. We show how the model can be practically applied by discussing related algorithms [3, 19, 18] in section 4. We present our own experiments in section 5, which show that the ambient and direct light can mix nonlinearly in practice.

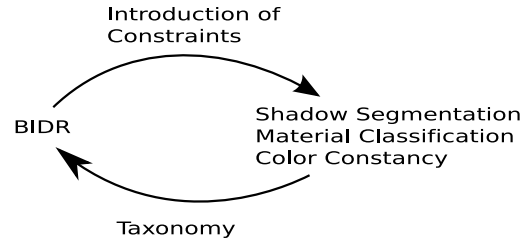


Figure 1. The BIDR serves either as starting point for scene modeling, or as the end point in a taxonomy of reflectance analysis algorithms.

We conclude the paper with perspectives for future research directions in section 6.

2. Reflectance Models

The dichromatic reflectance model assumes a single dominant point light source that uniformly illuminates the scene. It states that on dielectric surfaces, the radiance L is a linear combination of specular and diffuse reflectance:

$$L(\lambda, \mathbf{i}, \mathbf{e}) = L_s(\lambda, \mathbf{i}, \mathbf{e}) + L_d(\lambda, \mathbf{i}, \mathbf{e}) \quad (1)$$

where λ is the wave length, $\mathbf{i} = (\phi_i, \theta_i)$ is the vector of the angles between the surface normal and the incident light, and $\mathbf{e} = (\phi_e, \theta_e)$ is the vector of the angles between the surface normal and the viewing position. L_s is the radiance of the specular reflectance (also known as surface or interface reflectance), and L_d is the body (or diffuse) radiance. L_s and L_d can further be decomposed into the achromatic surface geometry $m_s(\mathbf{i}, \mathbf{e})$ and $m_d(\mathbf{i}, \mathbf{e})$, the incident light $e(\lambda)$ and the specular and diffuse chromatic parts $\rho_s(\lambda)$ and $\rho_d(\lambda)$ of the reflectance function:

$$L(\lambda, \mathbf{i}, \mathbf{e}) = m_s(\mathbf{i}, \mathbf{e})\rho_s(\lambda)e(\lambda) + m_d(\mathbf{i}, \mathbf{e})\rho_d(\lambda)e(\lambda) \quad (2)$$

In RGB-space, pixels that exhibit only diffuse reflectance form a distinct line, and pixels that exhibit a mixture of diffuse and specular reflectance form a slim triangle. As shown in Fig. 2, the axis of elongation of the specular triangle intersects the diffuse line at the higher intensity part of the line [13].

This model has been widely used for specularly based illumination estimation, color segmentation and specularly segmentation (see e.g. [13, 14, 23, 24]). A common limitation for the application in real-world scenes is the assumption of a single light source. If a scene contains either multiple direct illuminants or noticeable ambient lighting, which is typically the case in real-world scenarios, these influences appear as noise in the model. Instead of treating them as noise and trying to work around them, we believe

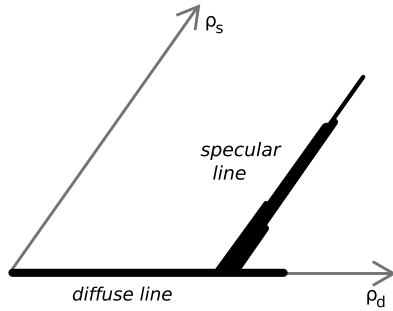


Figure 2. T-Shape according to Klinker and Shafer [13]. Diffuse pixels form a straight line (horizontally displayed), specular pixels form a sharp triangle on top of this line.

that it is helpful to address more complex lighting environments explicitly in a principled manner.

In this work, we follow this approach. Recently, Maxwell *et al.* [19] present a natural extension to the dichromatic reflectance model for a comprehensive ambient term, the Bi-Illuminant Dichromatic Reflectance Model:

$$L(\lambda, \mathbf{i}, \mathbf{e}) = L_{ds}(\lambda, \mathbf{i}, \mathbf{e}) + L_{dd}(\lambda, \mathbf{i}, \mathbf{e}) + L_{as}(\lambda, \mathbf{i}, \mathbf{e}) + L_{ad}(\lambda, \mathbf{i}, \mathbf{e}) \quad (3)$$

where $L_{ds}(\lambda, \mathbf{i}, \mathbf{e}) = L_s(\lambda, \mathbf{i}, \mathbf{e})$ and $L_{dd}(\lambda, \mathbf{i}, \mathbf{e}) = L_d(\lambda, \mathbf{i}, \mathbf{e})$ form the specular and diffuse component of the direct light in accordance with equation (1). $L_{as}(\lambda, \mathbf{i}, \mathbf{e})$ and $L_{ad}(\lambda, \mathbf{i}, \mathbf{e})$ are the new ambient term for specular and diffuse reflectance, respectively. $L_{as}(\lambda, \mathbf{i}, \mathbf{e})$ and $L_{ad}(\lambda, \mathbf{i}, \mathbf{e})$ are defined as

$$L_{as}(\lambda, \mathbf{i}, \mathbf{e}) = \int_{\Omega} m_s(\mathbf{i}, \mathbf{e}) \rho_s(\lambda) e_a(\lambda, \mathbf{i}) d\mathbf{i} \quad (4)$$

$$L_{ad}(\lambda, \mathbf{i}, \mathbf{e}) = \int_{\Omega} m_d(\mathbf{i}, \mathbf{e}) \rho_d(\lambda) e_a(\lambda, \mathbf{i}) d\mathbf{i} \quad (5)$$

Analogously to equation (1), $m_s(\mathbf{i}, \mathbf{e})$ and $m_d(\mathbf{i}, \mathbf{e})$ describe the specular and diffuse geometry and $e_a(\lambda, \mathbf{i})$ the color of the ambient incident light, with respect to the angle of incidence \mathbf{i} : for instance, differently colored interreflections can influence the ambient term from different angles of incidence. Ω is the hemisphere of all possible orientations of the surface normal. The full expansion of the extended dichromatic reflectance model is

$$L(\lambda, \mathbf{i}, \mathbf{e}) = m_s(\mathbf{i}, \mathbf{e}) \rho_s(\lambda) e_d(\lambda) + m_d(\mathbf{i}, \mathbf{e}) \rho_d(\lambda) e_d(\lambda) + \int_{\Omega} m_s(\mathbf{i}, \mathbf{e}) \rho_s(\lambda) e_a(\lambda, \mathbf{i}) d\mathbf{i} + \int_{\Omega} m_d(\mathbf{i}, \mathbf{e}) \rho_d(\lambda) e_a(\lambda, \mathbf{i}) d\mathbf{i} \quad (6)$$

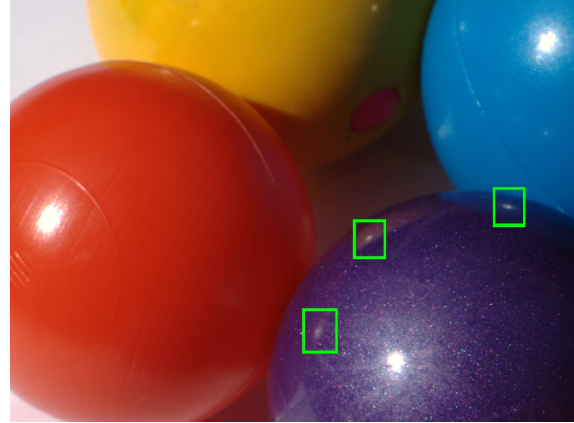


Figure 3. Interreflections between balls under sunlight. The majority of interreflections is diffuse. Specular highlights caused by interreflections are framed in green.

In this formulation, the ambient term acts as a special type of light source. Unlike direct illumination, ambient illumination is the result of interreflections. Light from all directions, and of possibly different color spectra, mixes in a point and forms the color of the ambient illuminant. Therefore, $e_a(\lambda, \mathbf{i})$ additionally depends on the angle of incidence.

Typically, if an ambient term within the dichromatic reflectance model is considered in the literature, it is a diffuse only term. Fig. 3 shows an example for diffuse and specular interreflections. Specular interreflections are framed in green. Most interreflections result from diffuse reflection. By adding a specular ambient term to the model, it is possible to also model specular interreflections. More generally, expressing the ambient term as a special light source (over an integral) implicitly offers a model for interreflections within the dichromatic reflectance model.

In practice, a full exploitation or extraction of the integral is seldomly possible, since the problem is severely under-constrained. However, we further analyze the model to gain insights on how the different terms interact with each other. In order to maximize the practical use of this model, it is especially important to understand how an image under a combination of direct and ambient lighting is formed. In this work, we focus on diffuse pixels, since the interaction of direct and ambient light has the biggest impact under diffuse lighting conditions. Although this examination is somewhat limited, we justify it by claiming that for many applications it is valid to assume that specularities dominate the color of a pixel. Furthermore, for describing the effects of introducing the ambient integral in the next section, it suffices to restrict the presentation to the diffuse terms.

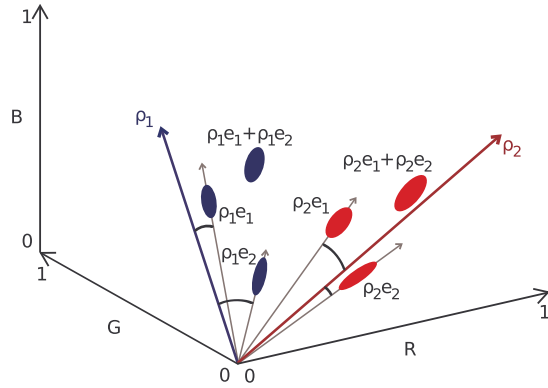


Figure 4. Ambient, diffuse and combined reflectance observed from two objects in RGB-space. Let e_1 the ambient illuminant and e_2 the direct illuminant. Note that the diffuse-only and the ambient-only pixels can (up to geometric variations) be matched on each other via rotation and scaling around the origin. For the mapping, the pixels with combined diffuse and ambient illumination require a translation additionally to rotation and scaling.

3. BIDR-based Scene Analysis

We illustrate the effects of the ambient integral in the BIDR. We believe that the extended model is helpful to the understanding of more complex realistic scenes, where effects from direct and ambient illumination can be observed. In the most general form, the two illuminants are very hard to separate. Therefore, we demonstrate properties of the diffuse part of the extended model. Starting with diffuse reflectance of either solely a direct illuminant or ambient lighting, the discussion progresses towards the full BIDR model step by step.

3.1. Direct Illuminant

Diffuse pixels under a single direct illuminant are known to form a straight line in RGB-space [22]. The radiance is simply

$$L(\lambda, \mathbf{i}, \mathbf{e}) = m_d(\mathbf{i}, \mathbf{e})\rho_d(\lambda)e(\lambda) , \quad (7)$$

such that for a single object under uniform illumination, the geometry term acts as a scaling factor for pixels on a straight line between the origin and $\rho_d(\lambda)e(\lambda)$.

Another point of view is a fixed geometry term and changes in either the illumination $e(\lambda)$ or the reflectance $\rho_d(\lambda)$. These then rotate and scale the object in RGB-space with the center of rotation and scaling at the origin. For two objects in a scene, if either $e(\lambda)$ or $\rho_d(\lambda)$ are fixed, the other factor can be related between the two objects by estimating rotation and scaling.

Fig. 4 shows an example for the rotation and scaling of directly lit pixels. If the direct illuminant is e_2 , the red and blue pixels ρ_1e_2 and ρ_2e_2 can (up to local variations in the

geometry) be mapped on each other by rotating and scaling the pixels with respect to the origin.

3.2. Isotropic Ambient Light

A diffuse ambient illumination is more complex, since geometry and the incident light are bound in the integral over the hemisphere Ω ,

$$L(\lambda, \mathbf{i}, \mathbf{e}) = \int_{\Omega} m_{ad}(\mathbf{i}, \mathbf{e})\rho_d(\lambda)e_a(\lambda, \mathbf{i})d\mathbf{i} . \quad (8)$$

In general, color and energy of the illuminant $e_a(\lambda, \mathbf{i})$ may vary over different spots on the hemisphere. These two factors are additionally coupled with the geometry term, and only $\rho_d(\lambda)$ can be easily factored out. This renders the integral very difficult to handle in the general case, where the ambient points can be almost arbitrarily spread in RGB space. For practical use, the typical assumptions made for simplification are that the illumination is isotropic and therefore uniformly colored over the hemisphere. Under these assumptions m_d and the energy of $e_a(\lambda)$ form a line in RGB-space starting in the origin. This is analogous to the case of direct light as discussed above. As shown in Fig. 4, if e_1 is the ambient illumination, ρ_1e_1 and ρ_2e_1 can be mapped on each other via rotation and scaling around the origin.

3.3. Direct Illuminant and Isotropic Ambient Light

A typical scene contains both direct and ambient light. In the following, we refer to such a scene as the *combined scene*. According to the model, the RGB values of the combined scene are given as the sum of the RGB values of the scene if it would be captured under direct-only light and ambient-only light. As stated above, under the assumption of isotropic and uniformly colored lighting in the ambient term, both illuminants form lines in RGB-space that run through the origin. These lines span a plane in RGB-space that also contains the RGB values of the scene under direct and ambient light.

Fig. 4 illustrates how ambient and diffuse lighting are combined. Both ambient and direct term can significantly influence the combined scene. Since both terms may theoretically scale in an arbitrary fashion, the range of possible RGB values in the combined scene is quite large.

As done previously with only one light source, we now consider the mapping of two objects with different albedo in RGB-space. Even with the constraints of isotropic and homogeneously colored illumination in the ambient term, rotation and scaling around the origin (as in the one-illuminant case) do not suffice to perform the mapping of the two objects under two illuminants. Due to the addition of the second illuminant, translation is required as well as rotation and scaling.

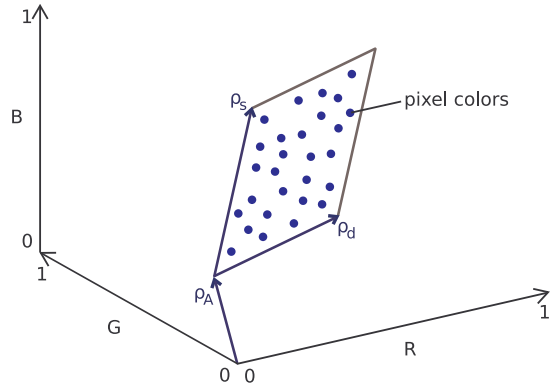


Figure 5. Ambient offset according to Shafer [22]. In this case, diffuse and specular pixels span a plane, which is offset by a constant ambient term A .

3.4. Unconstrained Ambient Illumination

In the most general case of diffuse lighting, we face an even harder problem. Here, the observed RGB values do not lie on a plane anymore. The integral of the ambient term introduces infinitely many unknowns, the ambient lighting may be seen as a combination of infinitely many illuminants or interreflections.

4. Incorporation of Existing Models

The explicit inclusion of the ambient term can be made algorithmically feasible by adding further constraints to the scene. On the converse, existing methodologies can be classified relative to the extended dichromatic reflectance model according to their constraints on the ambient term. As a side effect, this imposes a taxonomy on the existing methodologies and allows to categorize them according to their constraints within a common framework. In the following, we show how different existing approaches for shadow detection and illuminant estimation can be incorporated in the model.

4.1. Constant Ambient Term

The easiest way to incorporate the ambient term is to set it constant (see e.g. [20, 22]). In this case, equation (3) simplifies to

$$L(\lambda, \mathbf{i}, \mathbf{e}) = L_{ds}(\lambda, \mathbf{i}, \mathbf{e}) + L_{dd}(\lambda, \mathbf{i}, \mathbf{e}) + A(\lambda) , \quad (9)$$

which does not change the T-shape according to the dichromatic reflectance model, but shifts the offset of the shape from the origin to $A(\lambda)$, as shown in Fig. 5. The direct pixels and the specular pixels can be separated, e.g. by principal component analysis [20]. Additionally, the intersection of the diffuse part with the RGB cube can be used as a constraint on the color of the ambient illuminant.

4.2. Attenuation of the Direct Illuminant

A straightforward constraint is demonstrated by Barnard and Finlayson [3]: for the detection of shadows, the ratios of color channels between shadowed and non-shadowed areas are trained. The ratios are defined as follows: Let (R_S, G_S, B_S) and (R_N, G_N, B_N) the RGB-representations of a shadowed and non-shadowed pixel, respectively. The computed ratio is

$$r = \left(\frac{R_N}{R_S + R_N}, \frac{G_N}{G_S + G_N}, \frac{B_N}{B_S + B_N} \right) . \quad (10)$$

Effectively, equation (10) models the relative attenuation in every color channel for shadow pixels. This can be expressed in terms of the extended dichromatic reflectance model: in equation (6), the illuminant $e_d(\lambda)$ can be split into a chromatic part $c_d(\lambda)$ and an energy part $E_d(\lambda)$,

$$e_d(\lambda) = c_d(\lambda) \cdot E_d(\lambda) . \quad (11)$$

The ratio between the shadowed and non-shadowed pixels corresponds to a specific scaling of $E_d(\lambda)$, which effectively puts a weight on the direct illuminant:

$$L(\lambda, \mathbf{i}, \mathbf{e}) = E_d(\lambda) \cdot (m_s(\mathbf{i}, \mathbf{e})\rho_s(\lambda)c_d(\lambda) + m_d(\mathbf{i}, \mathbf{e})\rho_d(\lambda)c_d(\lambda)) + L_{as}(\lambda, \mathbf{i}, \mathbf{e}) + L_{ad}(\lambda, \mathbf{i}, \mathbf{e}) \quad (12)$$

Furthermore, for establishing equivalence, it is necessary to neutralize the integrals in the ambient term such that the geometry term can be canceled out by the ratio. This means that both specular components must be excluded, and the diffuse ambient lighting environment needs to be isotropic.

4.3. Simplifying Geometry

As soon as the ambient term is added to the dichromatic reflectance model, geometry adds further complexity to the reflectance, which is difficult to handle. In general, since it is part of the integral, it can not be easily factored out or cancelled. Therefore, application-specific reasonable assumptions have to be made.

A dichromatic shadow removal technique was presented by Maxwell *et al.* [19]. Directly lit pixels exhibit a weighted sum of direct and ambient illumination, full shadows are assumed to be lit by ambient light only. Under the assumption of homogeneous isotropic ambient lighting and Lambertian surfaces, geometry is cancelled in the ambient term. Therefore, the ambient term depends only on the albedo of the object. The idea is to consider the distribution C of diffuse pixels in RGB-space.

$$C = m_d(\mathbf{i}, \mathbf{e})\rho_d(\lambda)(c_d(\lambda) \cdot E_d(\lambda)) + \alpha\rho_d(\lambda) , \quad (13)$$

where $\alpha = \int_{\Omega} m_d(\mathbf{i}, \mathbf{e})e_d(\lambda, \mathbf{i})d\mathbf{i}$ reduces for a single material to a constant under the given restrictions. Note that

$e_d(\lambda)$ is replaced according to equation (11). Considering a single material, the only variable factor in equation (13) is $m_d(\mathbf{i}, \mathbf{e})E_d(\lambda)$. Therefore, after taking the logarithm of the diffuse and ambient component in equation (13), a line in log-space is formed. Pixels with less direct illumination can be projected to the top of this line, which effectively removes shadows.

4.4. Exploiting the Relationship between two Illuminants

With some constraints and little additional information of the scene, single factors can be isolated from equation (6). As an example, Lu and Drew proposed a method for illuminant estimation that is based on taking two images, one with flash light and one without [18]. The illumination of the no-flash image can be expressed to be ambient. Furthermore, Lambertian surfaces are assumed and specular pixels are ignored. Effectively, the ambient image L_a can, under these assumptions, be expressed as direct light diffuse-only:

$$L_a(\lambda, \mathbf{i}, \mathbf{e}) = m_d(\mathbf{i}, \mathbf{e})\rho_d(\lambda)e_a(\lambda) , \quad (14)$$

The image taken with a flash L_{fa} is according to the extended dichromatic reflectance model a sum of the direct and the ambient term

$$L_{fa}(\lambda, \mathbf{i}, \mathbf{e}) = m_d(\mathbf{i}, \mathbf{e})\rho_d(\lambda)e_d(\lambda) + L_a(\lambda, \mathbf{i}, \mathbf{e}) . \quad (15)$$

Therefore, by taking the difference $L_f = L_{fa} - L_a$, an artificial flash image L_f without the ambient term can be computed. Implicitly, this separates the sum in equation (15) into single components. Every component is a single product, which makes it possible to separate the factors, e.g. by taking the logarithm of each product. Since the geometry term $m_d(\mathbf{i}, \mathbf{e})$ and $\rho_d(\lambda)$ are equal if the images are registered, the difference reduces to the different illuminants. Ultimately, an illuminant estimate can be obtained by training the difference of the artificial flash image and the ambient image with Planckian light sources.

5. Experiments

A static scene was set up and captured under three different lighting conditions, namely ambient light, direct light and a combined image with ambient and direct light, as shown in Fig. 6. All data presented below was extracted from these images. The capturing device was a Hitachi HV-F22 3CCD camera with a Fujinon 1 : 1.4/25mm lens. Images were taken in a resolution of 1280×960 pixels with Gamma set to 1. Afterwards, a 5×5 pixel Gaussian filter was applied. For the ambient light only condition, the scene was illuminated by indirect sunlight in a room with white walls and ceiling and an array of windows on one side. The

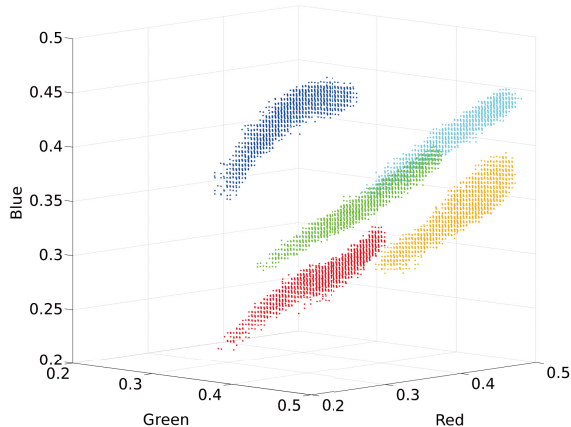


Figure 7. Observed pixel values in RGB-space for chalk of different color under ambient illumination. The point colors in the plot reflect the object colors, whereas cyan and orange denote white and yellow chalk, respectively.

window side of the scene area was obstructed by white material to establish a scene where ambient light is reflected from all sides. For the combined ambient and diffuse image, the scene was additionally illuminated by a halogen lamp, positioned at the top right with a 45 degree from the horizon. Finally, for the direct-only illuminated image, the room was completely darkened, and only the halogen lamp was switched on. All camera parameters were kept fixed while the three images were taken.

The examined objects in the scene consist of chalk, which is widely accepted as a good approximation of a material with Lambertian reflectance (see e.g. [15]). In Fig. 7, pixels from chinks of various color under ambient illumination are plotted in RGB-space. Material and capturing conditions in this setup are much closer to homogeneous illumination and Lambertian surfaces than typical real-world scenes. However, in RGB-space the chalk exhibits a non-linear shape for almost all colors. Therefore, it is important to consider the various factors that can distort the shape of the pixels under ambient illumination, namely geometry and inhomogeneous illumination.

When analyzing natural scenes from a single image, it is typically not possible to decompose the illumination in an ambient and a direct term. Therefore, we aim at understanding the formation of the shape of the combined image under ambient and direct illumination in RGB-space. Fig. 8 shows the observed pixels from orange chalk in RGB-space. The three gray shapes result from the image under ambient, direct and combined ambient and direct illumination. The shapes are formed by all orange chalk pixels from the respective image.

The pixels from the image under combined illumination exhibits an unexpectedly non-linear shape. In order to un-

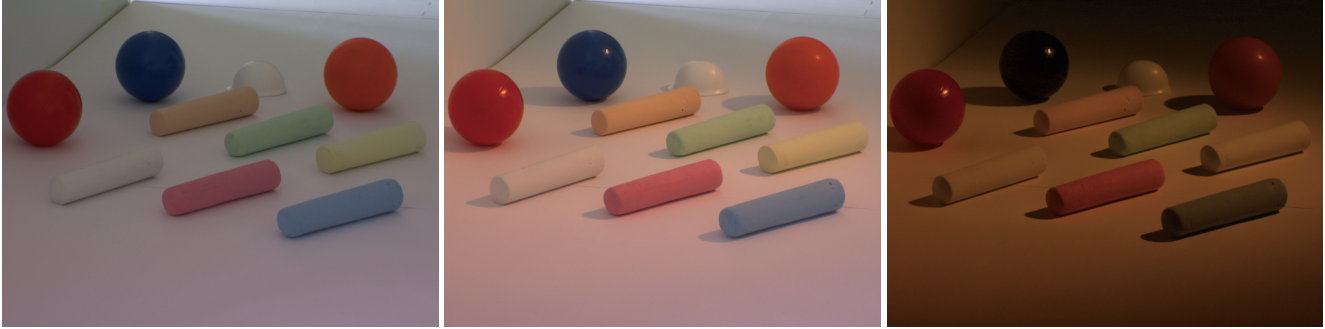


Figure 6. Scene used in the experiments. Left: only ambient illumination. Middle: combined ambient and direct illumination. Right: only direct illumination.

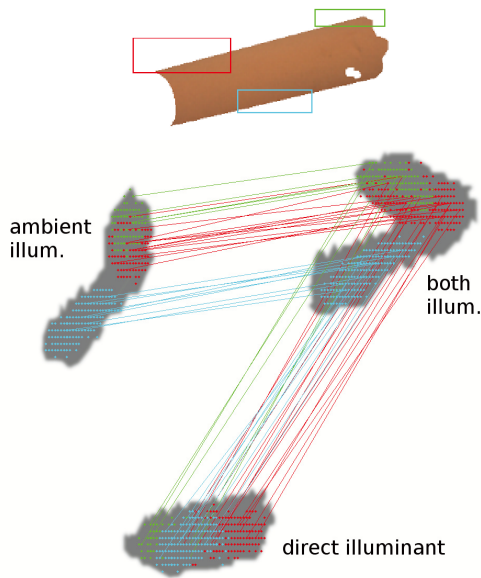


Figure 8. Top: Segment of the studied images (orange chalk). Bottom: Observed pixel values in RGB-space (rotated and scaled for the purpose of presentation). Values from the regions marked in red, green and cyan are plotted for the ambient, diffuse and combined ambient plus diffuse illuminants. The colored lines connect pixels from the same locations in the three images. The color-coding reveals which regions of the ambient and diffuse shapes are added, in order to understand the shape of the combined image.

Understand the formation of this shape, we selected three subregions on the chalk surface, as shown on top of Fig. 8. For all three images, the pixels from these regions are connected in RGB-space and correspondingly color-coded. The connecting lines show a pointwise relationship of the images under ambient, diffuse and combined illumination. The relative amount of light received and reflected differs for these regions according to their respective light source.

The points inside the red and the green region reflect a

high amount of ambient light. However, in the green region, considerably less light is received from the direct light source. The reason is that this region lies on the border of the light cone and has a larger angle between light source and surface normal compared to the red region. The cyan region is also a special case here: It receives more light from the direct light source than expected due to interreflections from the ground.

Therefore, spatially coincident pixels exhibit a different reflectance behavior relative to the centers of the shapes under ambient and direct illumination. This forms the remarkably curved shape under combined illumination. Geometrically, this finding is easily explained. Consider a replacement term for the ambient integral. This replacement can be seen as a pixel-specific second direct light source, whose color and angle of incidence corresponds to the mixed influences of the ambient integral. As a consequence, the scene illumination can per-pixel be modeled as a mixture of two mostly independent illuminants. Since the angles of incidence can be different for these illuminants, there are some surface points where the direct illuminant is relatively brighter than the ambient illuminant and vice versa.

Practically, this justifies the consideration of the BIDR as an extension to the dichromatic reflectance model. The observed nonlinearities can not be expressed by the original model, and must therefore be treated as noise, although these distortions can be quite significant, even in this relatively clean experimental setup.

6. Conclusions

We illustrated and characterized the effects that occur when the dichromatic reflectance model is extended to the Bi-Illuminant Dichromatic Reflectance Model (BIDR). We showed that the geometry parameters and an inhomogeneous lighting environment add much complexity to the analysis. Although the design of algorithms that exploit the possibilities of this representation is more challenging, we believe that several color research areas can benefit from

this more detailed representation: the inclusion of an ambient term enables us to model important optical effects like shadows or interreflections.

At the same time, the BIDR serves as a common framework for a taxonomy of many physics-based algorithms. By relating known methodologies to the BIDR, it is possible to group them according to the additional constraints they impose on the BIDR.

This is preliminary work. In the future, we expand on the examination and exploitation of the geometric effects. Furthermore, with the BIDR we have a model at hands that allows us to represent interreflections within a common framework. This adds another level of complexity, since interreflections have to be treated on specular and diffuse pixels.

7. Acknowledgements

The authors gratefully acknowledge funding of the Erlangen Graduate School in Advanced Optical Technologies (SAOT) by the German National Science Foundation (DFG) in the framework of the excellence initiative.

References

- [1] V. Agarwal, A. V. Gribok, A. Koschan, and M. A. Abidi. Estimating Illumination Chromaticity via Kernel Regression. In *International Conference on Image Processing*, pages 981–984, 2006.
- [2] R. Bajcsy, S. W. Lee, and A. Leonardis. Detection of Diffuse and Specular Interface Reflections and Inter-Reflections by Color Image Segmentation. *International Journal of Computer Vision*, 17(3):241–272, 1993.
- [3] K. Barnard and G. Finlayson. Shadow Identification using Colour Ratios. In *IS&T/SID 8th Colour Imaging Conference: Colour Science, Systems and Applications*, pages 97–101, Nov. 2000.
- [4] V. C. Cardei, B. Funt, and K. Barnard. Estimating the Scene Illumination Chromaticity Using a Neural network. *Journal of the Optical Society of America A*, 19(12):2374–2386, Dec. 2002.
- [5] J. M. DiCarlo, F. Xiao, and B. A. Wandell. Illuminating illumination. In *IS&T/SID 9th Colour Imaging Conference: Colour Science, Systems and Applications*, pages 27–34, Nov. 2001.
- [6] R. W. Ditchburn. *Light*. Dover Publications, 1991.
- [7] G. D. Finlayson, S. D. Hordley, and M. S. Drew. Removing Shadows from Images. In *European Conference on Computer Vision*, pages 823–836, May 2002.
- [8] G. D. Finlayson, S. D. Hordley, and I. Tastl. Gamut Constrained Illuminant Estimation. *International Journal of Computer Vision*, 67(1):93–109, 2006.
- [9] G. Funke-Lea and R. Bajcsy. Combining color and geometry for the active, visual recognition of shadows. In *International Conference on Computer Vision*, pages 203–209, 1995.
- [10] B. V. Funt and M. S. Drew. Color Space Analysis of Mutual Illumination. *Pattern Analysis and Machine Intelligence*, 15(12):1319–1326, Dec. 1993.
- [11] R. Gershon, A. D. Jepson, and J. K. Tsotsos. Ambient illumination and the determination of material changes. *Journal of the Optical Society of America A*, 3(10):1700–1707, Oct. 1986.
- [12] T. Gevers and A. W. M. Smeulders. Color-based object recognition. *Pattern Recognition Letters*, 32(3):453–464, Mar. 1999.
- [13] G. J. Klinker, S. A. Shafer, and T. Kanade. The Measurement of Highlights in Color Images. *International Journal of Computer Vision*, 2(1):7–32, Jan. 1988.
- [14] G. J. Klinker, S. A. Shafer, and T. Kanade. A Physical Approach to Color Image Understanding. *International Journal of Computer Vision*, 4(1):7–38, 1990.
- [15] J. J. Koenderink, A. J. van Doorn, K. J. Dana, and S. Naray. Bidirectional Reflection Distribution Function of Thoroughly Pitted Surfaces. *International Journal of Computer Vision*, 31(2-3):129–144, Apr. 1999.
- [16] M. S. Langer. When Shadows Become Interreflections. *International Journal of Computer Vision*, 34(2-3):193–204, Aug. 1998.
- [17] M. D. Levine and J. Bhattacharyya. Removing shadows. *Pattern Recognition Letters*, 26(3):251–265, Feb. 2004.
- [18] C. Lu and M. S. Drew. Practical Scene Illuminant Estimation via Flash/No-Flash Pairs. In *IS&T/SID 14th Colour Imaging Conference: Colour Science, Systems and Applications*, Nov. 2006.
- [19] B. A. Maxwell, R. M. Friedhoff, and C. A. Smith. A Bi-Illuminant Dichromatic Reflection Model for Understanding Images. In *Computer Vision and Pattern Recognition, IEEE Conference on*, pages 1–8, June 2008.
- [20] S. Nadimi and B. Bhanu. Moving Shadow Detection Using a Physics-based Approach. In *Pattern Recognition, 16th International Conference on*, volume 2, pages 701–704, Aug. 2002.
- [21] E. Salvador, A. Cavallaro, and T. Ebrahimi. Cast shadow segmentation using invariant color features. *Computer Vision and Image Understanding*, 95(2):238–259, Aug. 2004.
- [22] S. A. Shafer. Using Color to Separate Reflection Components. *Color Research Application*, 10:210–218, 1985.
- [23] R. Tan and K. Ikeuchi. Separating Reflection of Textured Surfaces using a Single Image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27:178–193, 2005.
- [24] R. Tan, K. Nishino, and K. Ikeuchi. Color Constancy through Inverse-Intensity Chromaticity Space. *Journal of the Optical Society of America A*, 21(3):321–334, 2004.
- [25] J. van de Weijer and C. Schmid. Coloring Local Feature Extraction. In *European Conference on Computer Vision*, pages 334–348, May 2006.