

# Object Recoloring based on Intrinsic Image Estimation

Shida Beigpour, Joost van de Weijer

Centre de Visio per Computador, Computer Science Department

Universitat Autònoma de Barcelona, Edifici O, Campus UAB (Bellaterra), C.P.08193, Barcelona, Spain

{ shida, joost } @cvc.uab.cat

## Abstract

Object recoloring is one of the most popular photo-editing tasks. The problem of object recoloring is highly under-constrained, and existing recoloring methods limit their application to objects lit by a white illuminant. Application of these methods to real-world scenes lit by colored illuminants, multiple illuminants, or interreflections, results in unrealistic recoloring of objects.

In this paper, we focus on the recoloring of single-colored objects presegmented from their background. The single-color constraint allows us to fit a more comprehensive physical model to the object. We show that this permits us to perform realistic recoloring of objects lit by colored lights, and multiple illuminants. Moreover, the model allows for more realistic scene relighting. Recoloring results on images captured by uncalibrated cameras demonstrate that the proposed framework obtains realistic recoloring for complex natural images. Furthermore we use the model to transfer color between objects and show that the results are more realistic than existing color transfer methods.

## 1. Introduction

Recoloring refers to the modification and adjustment of color appearance in images. Object recoloring methods are used in photo montage, image color correction, visual effects in movies, and also to facilitate the industrial design by visualizing the final color appearance of the object before production. In the current work we focus on recoloring of single-colored objects in images of medium quality as typically encountered on the Internet.

One of the most popular color modification applications is the recoloring of a specific object with another color or under different lighting condition (e.g., warm-tone sunset or cold-tone early morning). In many circumstances, it may not be possible to create the object in the desired color or to simulate the desired lighting condition. Another case is when an impossible scenario is desired, for example a blue apple, and here the choices are to either render a 3D model




Figure 1. The first row is an example of the reflectance decomposition achieved by DRM [20]. Using this decomposition, *object recoloring* is performed by changing the *body reflectance* (the second row), and *illuminant recoloring* is achieved by changing the *specular reflectance* (the third row).

of the scene or to simply photograph the object and then *recolor* it. Recoloring should result in physically plausible scenes and should require minimum user interaction.

Here our main objective is to develop a physics-based method to extract the underlying reflectance model of the object and separate the geometric characteristics from the colors of the object and the illuminant. Such physics-based model can then be used in order to generate an image of the object in the same lighting and viewing angles, varying only the object and/or illuminant colors. Fig 1 provides an example of reflectance decomposition as well as object and illuminant recoloring.

Images describing the underlying physical properties of the scene such as reflectance, orientation, and illumination are known as *intrinsic images* and were first introduced by Barrow and Tenenbaum [2]. Intrinsic images are more appropriate for higher-level scene analysis than the original light intensity images, because they are less prone to scene accidental events such as illuminant direction and color changes. The *Dichromatic Reflection Model* (DRM) [20] models the object reflectance using two chromatic coefficients: *body reflectance*  $c^b$ , and *specular reflectance*  $c^s$ :

$$f(\mathbf{x}) = m_b(\mathbf{x})c_b + m_s(\mathbf{x})c_s, \quad (1)$$

where, for each pixel  $\mathbf{x}$ ,  $m^b$  and  $m^s$  are the intrinsic images

describing the interaction between the light and the surface as a function of geometric parameters such as incident angle, viewing angle, and surface normal.

In this paper, we investigate the application of the *single-colored object* constraint to derive the intrinsic images of a scene. We assume a segmented mask of a single-colored object to be given as an input. A user working in a photo-editing environment has multiple segmentation tools to quickly segment objects [11, 19]. This single-colored object constraint greatly simplifies the estimation of intrinsic images. We show that this constraint allows us to further extend the DRM to model more complex scenes with multiple illuminants which proves crucial for outdoor scenes where two illuminants (e.g. the sun and a blue skylight) illuminate the object.

We propose a Multi-illuminant Dichromatic Reflection model (MIDR), and provide an algorithm for solving the case of two illuminants. This algorithm is then embedded in a framework which is capable of recoloring complex objects in the presence of shadows and specularities formed by two *unknown* illuminants (e.g. colored-shadows and interreflections) and achieving physically plausible results for uncalibrated natural scene images. As an additional application we show that our framework applied to color transfer, handles complex objects with specularity and under multiple illuminant better than existing methods.

## 2. Related work

**Intrinsic images.** Several methods have been proposed to compute the intrinsic images of Eq. 1 based on various constraints. A common constraint is to assume Lambertian reflectance ( $m_s = 0$ ). For this case, Weiss [27] shows that for an image sequence assuming  $c_b$  to be constant over time, and using the prior that illumination images give rise to sparse filter outputs, estimation of the intrinsic images is achievable. Tappen et al. [24] show that by assuming that shading and reflectance boundaries do not occur at the same location the intrinsic images can be derived from a single image.

Fewer works have concentrated on solving the case where  $m_s \neq 0$ . Klinker et al. [9] propose a method where segmentation and intrinsic image estimation are iteratively alternated. Within each segment a single DRM is estimated. Hypotheses of possible illuminant and object colors are verified for the segments and neighboring segments. This method is further extended to include multicolored objects in [14, 15]. The main drawback of these approaches is that they face a chicken-and-egg problem: for a good segmentation you need approximately correct DRM parameters, and vice versa. Furthermore, these methods are only evaluated on high-quality images taken in controlled environments, typically without complex backgrounds, which greatly enlarges the hypothesis space to be checked, and limits the

probability of correct convergence.

Several highlight/specularity removal methods have been proposed using the assumption of a *known* illuminant  $c_s$  and that the specular pixels have the same diffuse value as their neighboring diffuse pixels. For example, Robbie Tan et al. [23] proposed an iterative method for reflectance decomposition of textured surfaces. Tan et al. [22] improve the previous methods by adding spacial distribution and texture constraints when available. Mallick et al. [12] uses partial differential equation that iteratively erodes the specular component at each pixel.

**Object recoloring.** Many *colorization* methods have also been used for recoloring. They mainly consist of partial hand-coloring of regions in an image or video and propagating the colored points (known as *color markers* or *hot-spots*) to the rest of the image using an optimization algorithm [10, 3]. Since these algorithms are based on the luminance image they lack the additional color information which allows to separate the Lambertian reflectance and specular reflectance, causing them to fail in the presence of specularities.

*Color transfer* methods extract the color characteristics from a source image and apply it to a target image. Many color transfer methods are based on pixels color distribution [17, 18]. Local color transfer [6, 28] and user-interactive methods [1] try to improve the results by providing more cues. The main issue of the color transfer is that it requires a target scene, while here we solve the case for which no information about the target distribution is given. Furthermore, these methods are generally applied to matte surfaces and do not consider the presence of specularities.

The recoloring embedded in professional photo-editing applications performs by calculating an offset in the hue and saturation between the source and target colors. The source image is adjusted to produce the desired color [7]. This method is fast and capable of producing realistic results. However, as it ignores the underlying physical reflectance, it fails in the case of colored or multiple illuminant.

Omer et al. [16] present an *image specific* color representation robust to color distortion and demonstrated a recoloring example for a Lambertian surface. A more physics-based approach, the closest method to our own, is a DRM based color transfer method [21] in which the object (body) color is estimated and transferred between images. And realistic results on lab conditioned high quality images of objects under single *known* illuminant are presented.

Hsu et al. [8] proposed a novel method to estimate the light mixture in a single image illuminated by two lights specified by user while the reflectance is modeled as solely diffuse. The method achieves good results on white balance and light color change.

### 3. Object Reflectance Modeling

In this section, we describe a physics-based reflectance model for object pixels to achieve a high quality recolored image. We begin with an overview of the DRM and then we extend it for the Multi-illuminant case.

#### 3.1. Dichromatic Reflection Model (DRM)

According to Shafer, pixel values for a set of points on a single colored surface must lie within a parallelogram in the RGB space, bounded by body reflectance  $\mathbf{c}_b$  and the specular reflectance  $\mathbf{c}_s$  [20]. Validity of the DRM has been proven for a variety of inhomogeneous dielectric materials commonly observed in natural scenes [25]. In this paper, we assume that color changes can be modelled by a diagonal model, or Von Kries model, which has been proven a sufficient approximation [4]. We indicate the illuminant color by  $\mathbf{l}$ , and  $\mathbf{L} = \text{diag}(\mathbf{l})$  is its diagonal matrix representation. In this case the DRM can be written as

$$\mathbf{f} = m_b \mathbf{c}_b + m_s \mathbf{c}_s = m_b \mathbf{cL} + m_s \mathbf{l}, \quad (2)$$

where  $\mathbf{f}$  is the RGB triple defining the color of every pixel in the object surface,  $m_b$  and  $m_s$  are the intrinsic images denoting the magnitude of the body and specular reflectance respectively (Fig 1). The body reflectance is a multiplication of the *material reflectance*  $\mathbf{c}$  and the illuminant according to  $\mathbf{c}_b = \mathbf{cL}$ . We assume *neutral interface reflectance*, causing the specular reflectance to have the same chromaticity as the illuminant color  $\mathbf{c}_s = \mathbf{l}$ . This equation can be divided into intrinsic images and the chromaticity of the object and illuminant in matrix notation according to

$$\mathbf{f} = [m_b(\mathbf{x}) \ m_s(\mathbf{x})] [\mathbf{L} \ \mathbf{c} \ \mathbf{l}]^T = \mathbf{M} \mathbf{C}^T, \quad (3)$$

where  $\mathbf{x}$  is a vector of  $n \times 2$  coordinates,  $\mathbf{f}$  is the  $n \times 3$  matrix of pixels RGB values, and the intrinsic image matrix  $\mathbf{M} = [m_b(\mathbf{x}), m_s(\mathbf{x})]$  is  $n \times 2$  matrix containing intrinsic images. The color characteristics matrix  $\mathbf{C} = [\mathbf{L} \ \mathbf{c} \ \mathbf{l}]$  contains the relevant parameters for scene recoloring. In Section 4 we purpose methods to estimate the model parameters.

#### 3.2. Multi-illuminant Dichromatic Reflection (MIDR) model

Real-world objects often exhibit body and surface reflection under more than just one illuminant. An example of multi-illuminant scenario is an outdoor scene with blue sky and yellow sun, or a scene with indoor lighting combined with outdoor lighting through a window. Conventional methods often ignore the secondary illuminants present in the scene to simplify the modelling. Here we extend the reflectance model to the Multi-illuminant Dichromatic Reflection model(MIDR) to account for the secondary illuminants. The MIDR for  $n$  illuminants is given by

$$\mathbf{f} = [\mathbf{M}^1 \dots \mathbf{M}^n] [\mathbf{C}^1 \dots \mathbf{C}^n]^T = \mathbf{M} \mathbf{C}^T, \quad (4)$$

where  $\mathbf{M}^n$  contains the intrinsic images regarding the  $n^{th}$  illuminant and  $\mathbf{C}^n$  is the corresponding color characteristics matrix. Note that the material reflectance  $\mathbf{c}$  remains constant for all intrinsic color matrices. Due to the high complexity of the model, in Section 5 we solve for a simplified case of the MIDR model.

The dichromatic reflection model has also been extended to include *ambient lighting*. Originally Shafer [20] modelled ambient light as a constant offset over the scene. Later work improved the modelling [13] and showed that the ambient term results in an object color dependent offset. For the matter of simplification, in this work we assume the ambient illuminant to be negligible.

### 4. Dichromatic Reflection Model estimation

Since the introduction of the DRM multiple approaches to solve this model have been proposed [9, 15, 23, 12]. In this paper, we are interested in solving the DRM for the application of recoloring single colored objects. Users interested in object recoloring work within a photo-editing environment, allowing them to quickly segment the object of interest. This *single-colored object* constraint allows us to fit a more realistic illumination model, allowing the object to be lit by multiple illuminants.

A successful object recoloring algorithm has to face several challenges:

- **Uncalibrated images:** Photo-editing software users typically work with uncalibrated, compressed images of medium quality and unknown settings. Most previous methods experiment on high quality calibrated images taken in lab conditions [9, 15], and known illumination [23, 12]. To handle these lower quality images we propose a *robust estimator*(Section 4.1).
- **Complex color distribution:** several existing approaches estimate the illuminant by fitting L and T-shapes to the color distribution [9, 15]. These methods are based on the hidden assumption that the  $m_b$  is assumed constant while  $m_s$  is changing. In real-world images we often face much more complex distribution which rather form a plane. To tackle this problem we use the *illuminant estimation* described in Section 4.2.
- **Complex lighting conditions:** the objects in real-world images are often lit by multiple illuminants, colored shadows, and interreflections. Ignoring these lighting conditions would make the resulting object recoloring look artificial. Therefore, in Section 5, we propose an iterative algorithm to solve for two illuminants.

#### 4.1. Robust Body Reflectance Estimation (RBRE).

For the task of body reflectance color ( $\mathbf{c}_b$ ) estimation on medium quality images we propose the Robust Body Reflectance Estimation (RBRE). Since object pixel values of

the non-specular part ( $m_s = 0$ ) form a line passing through the origin, fitting a line through these pixels allows us to compute  $\mathbf{c}_b = \mathbf{c}\mathbf{L}$ . The fitting error of an object pixel  $\mathbf{x}$  to a line given by the normalized vector  $\hat{\mathbf{c}}_b$  is

$$e(\mathbf{x}) = \left\| \mathbf{f}(\mathbf{x}) - \left( (\mathbf{f}(\mathbf{x}))^T \hat{\mathbf{c}}_b \right) \hat{\mathbf{c}}_b \right\|. \quad (5)$$

Although the least squares (LS) orientation estimation would perform well in the case that all pixels belong to the same orientation, in our case in which there are two main orientations ( $\mathbf{c}_b$  and  $\mathbf{l}$ ), the LS estimation will mix the two orientations and give a wrong result. In order to avoid that, a *robust estimator* [26] is constructed:

$$e = \int_{\Omega} \rho(e(\mathbf{x})) dx. \quad (6)$$

In the current work we apply the *Gaussian error norm*:

$$\rho^m(e) = 1 - \exp\left(-\frac{e^2}{2m^2}\right). \quad (7)$$

In a robust estimator, large deviations from the model are considered as outliers, and therefore, they are not taken into account very heavily. While LS estimation is very sensitive to outliers. In our application large deviations from the model are mainly due to the mixing of two different directions,  $\mathbf{c}_b\mathbf{L}$  and  $\mathbf{l}$ . The error, Equation 6, can now be rewritten as (we will omit the spatial arguments):

$$e = \int_{\Omega} \rho^m \left( \sqrt{\mathbf{f}^T \mathbf{f} - \hat{\mathbf{c}}_b^T (\mathbf{f} \mathbf{f}^T) \hat{\mathbf{c}}_b} \right) dx. \quad (8)$$

A Lagrange multiplier is then used for minimization subject to the constraint  $\hat{\mathbf{c}}_b^T \hat{\mathbf{c}}_b = 1$ ,

$$\frac{d}{d\hat{\mathbf{c}}_b} \left( \lambda \left( 1 - \hat{\mathbf{c}}_b^T \hat{\mathbf{c}}_b \right) + e \right) = 0. \quad (9)$$

Using Equation 7 as the error function leads to

$$\eta(\hat{\mathbf{c}}_b) \hat{\mathbf{c}}_b = \lambda \hat{\mathbf{c}}_b, \quad (10)$$

where  $\eta$  is defined according to

$$\eta(\hat{\mathbf{c}}_b) = \int_{\Omega} \mathbf{f} \mathbf{f}^T G^m \left( \sqrt{\mathbf{f}^T \mathbf{f} - \hat{\mathbf{c}}_b^T (\mathbf{f} \mathbf{f}^T) \hat{\mathbf{c}}_b} \right) dx. \quad (11)$$

The main difference with the ordinary LS estimator is that here the matrix  $\eta$  is dependent on  $\hat{\mathbf{c}}_b$ . Eq 10 can be solved by a *fixed point* iteration scheme. We start iteration with the initial estimate  $\hat{\mathbf{c}}_b^0$  given by the LS. Let  $\hat{\mathbf{c}}_b^i$  be the orientation vector estimate after  $i$  iterations. The estimate is updated as the eigenvector  $\hat{\mathbf{c}}_b^{i+1}$  of the matrix  $\eta(\hat{\mathbf{c}}_b^i)$  corresponding to the largest eigenvalue, i.e. we solve

$$\eta(\hat{\mathbf{c}}_b^i) \hat{\mathbf{c}}_b^{i+1} = \lambda \hat{\mathbf{c}}_b^{i+1}. \quad (12)$$

Again, points far away from the line direction  $\hat{\mathbf{c}}_b$  are considered outliers, and therefore, do not corrupt the estimation. Iterative application of Equation 12 yields the estimate of the body reflection,  $\hat{\mathbf{c}}_b$ . The original estimation made by ordinary LS is refined at each iteration by changing the weights leading the method to converge to a robust, and in this case a much better, estimation of the  $\hat{\mathbf{c}}_b$ .

## 4.2. Confined illuminants estimation (CIE)

Having the body reflectance color, there exists a set of possible illuminants which could generate the color distribution of the object. Many of these illuminants are unrealistic. It is shown that the chromaticity of common light sources closely follows the Planckian locus of black-body radiators [5]. We propose to use this constraint to estimate the illuminant.

We sample Planckian colors ( $T \subset 1000 \sim 40000$ ) which vary from orange to yellow to white to blue, resulting in a set of illuminants  $\{\mathbf{l}_1, \dots, \mathbf{l}_m\}$ . We define the reconstruction error of the intrinsic images  $\mathbb{M}$  and intrinsic color characteristics  $\mathbb{C}$  by

$$E_r(\mathbf{f}, \mathbb{M}, \mathbb{C}) = (\mathbf{f} - \mathbb{M}\mathbb{C}^T)^T (\mathbf{f} - \mathbb{M}\mathbb{C}^T). \quad (13)$$

Then, we perform an exhaustive search to find the best matching Planckian light. In other words, we solve Equation 14 by choosing the Planckian light, which minimizes the reconstruction error.

$$\hat{\mathbf{l}} = \arg \min_{\mathbf{l} \in \{\mathbf{l}_1, \dots, \mathbf{l}_m\}} E_r(\mathbf{f}, \mathbb{M}, [\mathbf{c}\mathbf{L} \ \mathbf{l}]). \quad (14)$$

In the next section we will outline the computation of the intrinsic images  $\mathbb{M}$  given  $\mathbb{C}$ , which are needed for the computation of the reconstruction error.

## 4.3. Intrinsic images

The estimation of the intrinsic images, given an estimation of  $\hat{\mathbb{C}}$ , is based on the convex optimization problem:

$$\begin{aligned} & \underset{\mathbb{M}}{\text{minimize}} E_r(\mathbf{f}, \mathbb{M}, \hat{\mathbb{C}}) \\ & \text{subject to } m_b(\mathbf{x}) \geq 0, m_s(\mathbf{x}) \geq 0. \end{aligned} \quad (15)$$

Fig 2 demonstrates an example of intrinsic images recovered for an object. Note that the specular reflectance is correctly separated from the body reflectance.

## 5. Two-illuminant MIDR model estimation

Many real-world objects are lit by multiple illuminants. Here we propose an algorithm to estimate the case of two illuminants. Since the problem is highly underconstrained, we need further assumptions: Firstly, we assume one illuminant to be Planckian and demonstrate specularities; Secondly, specularities of the secondary illuminant to be negligible. We use this as an additional constraint ( $m_s^2(\mathbf{x}) = 0$ ).



Figure 2. An example of intrinsic images recovered for an object. (a) Original image; Intrinsic images: (b) Body reflectance and (C) Specular reflectance; (d) An example recoloring result.

Note that we make no assumption on the chromaticity of the secondary illuminant. Hence the model is given by

$$\mathbf{f} = m_b^1 \mathbf{c} \mathbf{L}^1 + m_s^1 \mathbf{l}^1 + m_b^2 \mathbf{c} \mathbf{L}^2. \quad (16)$$

An iterative algorithm to solve this MIDR model is given in Algorithm 1. First we will assume pixels to be illuminated by only one of the two illuminants  $m_b^1(\mathbf{x})m_b^2(\mathbf{x}) = 0$  and  $m_s^1(\mathbf{x})m_b^2(\mathbf{x}) = 0$ . In the final Step we remove this restriction to allow for pixels being lit by both illuminants at the same time. Here, we also use the *diag*-function to convert vectors to diagonal matrices and vice versa. First an initial estimation is made based on all pixels on the object (Steps 1-4) which gives us the initial values for the dominant illuminant and object color. Based on this model pixels which could be described by this model with affordable error are separated from the rest (Step 5) which are indicated by the *Mask*. At each iteration the estimations and separation mask are refined. We estimate a Lambertian reflectance model for the pixels outside the Mask (Steps 10 and 11). Iteratively the illuminant color estimations are refined until convergence (Step 13). The final model estimation is then given by the object material reflectance color  $\mathbf{c}$ , the two illuminant colors  $\mathbf{l}_1$  and  $\mathbf{l}_2$ , and the corresponding intrinsic images  $m_b^1$ ,  $m_b^2$ , and  $m_s$ .

Although the algorithm gives good estimates for  $\mathbf{c}$ ,  $\mathbf{l}^1$  and  $\mathbf{l}^2$ , the constraint that pixels can only be illuminated by a single illuminant results in artificial edges in the  $m_b^1$  and  $m_b^2$  estimates. In reality there are regions where both lights illuminate the object. To solve this, Step 14 finalizes the algorithm by keeping  $\mathbf{c}$ ,  $\mathbf{l}^1$  and  $\mathbf{l}^2$  and  $m_s^1$  constant in Eq 16, and estimates  $m_b^1$  and  $m_b^2$  constraining them to be positive.

In Fig 3 we show the results of the algorithm on an outdoor car image. The car is illuminated by a white outside lighting as well as a greenish light caused by the light coming from the grass field. The mask is given for several iterations of the algorithm. The algorithm correctly separates the two illuminants. In the last row the intrinsic images show the estimates of the body and specular reflection.

## 6. Experimental results

In the experimental section we analyze our proposed algorithm for MIDR estimation on synthetic images. Additionally we show some results on challenging real-world

---

### Algorithm 1 Two-illuminant MIDR model estimation

---

- 1: Consider the whole object segment as *Mask*
  - 2: Estimate  $\mathbf{c}_b$  using RBRE for the pixels  $\mathbf{x} \in \text{Mask}$
  - 3: Estimate the Planckian illuminant  $\mathbf{l}_1$  using CIE method
  - 4:  $\mathbf{c} \leftarrow \text{diag}(\mathbf{c}_b \mathbf{L}_1^{-1})$
  - 5: Initiate *Mask* to only include the pixels  $\mathbf{x}$  for which  $E_r(\mathbf{f}(\mathbf{x}), \mathbf{M}^1, \mathbf{C}^1) < \text{Threshold}$
  - 6: **repeat**
  - 7:   Estimate  $\mathbf{c}_b^1$  using RBRE for the pixels  $\mathbf{x} \in \text{Mask}$
  - 8:   Estimate the Planckian illuminant  $\mathbf{l}_1$  using CIE method
  - 9:    $\mathbf{c} \leftarrow \text{diag}(\mathbf{c}_b^1 \mathbf{L}_1^{-1})$
  - 10:   Estimate  $\mathbf{c}_b^2$  using RBRE for the object pixels  $f(\mathbf{x}) \notin \text{Mask}$
  - 11:    $\mathbf{L}_2 \leftarrow \text{diag}(\mathbf{c}_b^2) / \text{diag}(\mathbf{c})$  (using the  $\mathbf{c}$  from Step 9).
  - 12:   Update *Mask* to only include the pixels  $\mathbf{x}$  for which  $E_r(\mathbf{f}(\mathbf{x}), \mathbf{M}^1, \mathbf{C}^1) < E_r(\mathbf{f}(\mathbf{x}), \mathbf{M}^2, \mathbf{C}^2)$
  - 13: **until**  $\mathbf{L}_1$  and  $\mathbf{L}_2$  estimates converge
  - 14: Recalculate the  $m_b^1$  and  $m_b^2$  using the previous estimates for  $\mathbf{c}$ ,  $\mathbf{l}_1$ ,  $\mathbf{l}_2$  and  $m_s^1$ .
- 



Figure 4. The first four images are examples of the synthetic images. The last three images are the  $m_b^1$ ,  $m_b^2$ , and  $m_s^1$  ground truth.

images. Here we assume images are in sRGB format; and do *gamma correction*. Further applications of the model are discussed in the end of the section.

Please refer to the supplementary video for more examples of the results in real-world images.

### 6.1. Synthetic Images

Here we test our algorithm on synthetic images which satisfy the assumptions, namely they are lit by two lights, one of which is Planckian. The groundtruth intrinsic images  $m_b^1$ ,  $m_b^2$  and  $m_s^1$  are given (Fig 4). With these we generate a set of 60 test images by varying the illuminants and the object color. Some examples are given in Fig 4. The soundness of our algorithm has been verified on synthetic test im-




Figure 3. An example of the MIDR algorithm performance: (a) Original image; (b) Object mask as the initial *Mask* for the illuminant  $l_1$ ; (c) The Mask after 1st iteration; (d) The Mask at 3rd (final) iteration; (e) estimated  $m_b^1$ ; (f) estimated  $m_b^2$  (the interreflection area,  $l_2$ , has been correctly detected); (g) estimated  $m_s^1$ ; (h) An example recoloring (the interreflection is preserved).

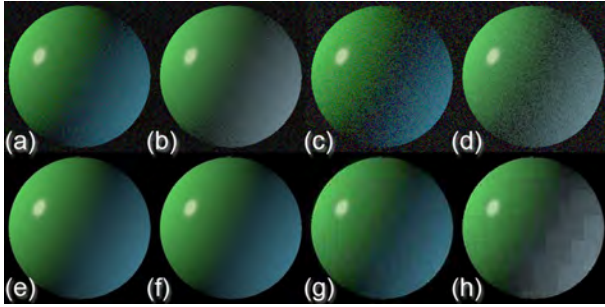


Figure 5. Effect of noise and JPEG compression: (a) and (c) examples of applying noise by sigma 4.0 and 9.0; (b) and (d) their corresponding reconstructions; (e) and (g) examples of applying JPEG compressions of 20% and 80%; (f) and (h) their corresponding reconstructions.

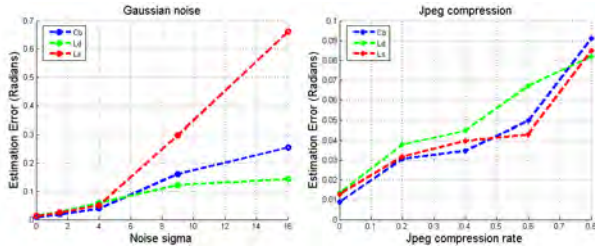


Figure 6. Median angular error (in radian) as a function of: Gaussian noise sigma (left) and JPEG compression rate (right) for  $c_b$ ,  $l_1$  and  $l_2$  estimates.

ages on which the intrinsic image estimation performs with an error close to zero even though a large part of the object is lit by both lights simultaneously.

Since we want to apply our method to standard Internet images, we further investigated its robustness to both Gaussian noise and JPEG compression (Fig 6). The comparison is made using the *Angular Error* ( $E_a$ ) in radians between the ground-truth ( $\hat{c}_{gt}$ ) and estimated ( $\hat{c}_{est}$ ) colors as de-

finied below,

$$E_a = \arccos(\hat{c}_{gt} \cdot \hat{c}_{est}). \quad (17)$$

As can be seen the algorithm is sensitive to Gaussian noise but relatively robust to JPEG compression (angular error of all estimations for 60% compression is under 0.07 radian). To better interpret the results in the graphs we also provide the reconstruction results on one synthetic object for several noise and JPEG compression settings in Fig 5.

## 6.2. Real-world Images

Fig 8 compares MIDR-based recoloring with the one done by hue-saturation shift method. The secondary illuminant (greenish interreflection) is correctly preserved by MIDR while wrongfully changed to blue by the professional photo-editor. In Fig 9 the MIDR and DRM has been compared for the accuracy of their recoloring results. The secondary illuminant (bluish shadow) is well preserved by MIDR while lost in the case of DRM. Note that here we only modeled two illuminants and therefore the third illuminant (the small brownish interreflection on the back of the car) is lost.

Note that theoretically the method fails to correctly make the intrinsic image decomposition in the case object and illuminant colors are collinear. Also having no Planckian illuminant confuses the CIE estimator. The latter is shown in the example of Fig 7.

## 6.3. Other Applications of MIDR

Here we show two other interesting applications for the proposed framework, namely *Physics-based Color Transfer* and *Photo Fusion*.

**Physics-based Color Transfer.** A popular photo-editing task is transferring the color characteristics of an image to another. Even though color transfer methods are often successful in transferring the atmosphere of one image onto the



Figure 8. Comparing the MIDR method performance with a professional photo-editor: (a) Original image (containing complex interreflection); (b) Recoloring result by MIDR (the secondary illuminant, green interreflection, has been preserved); (c) Recoloring result using the hue-saturation shift method (the green interreflection is wrongfully changed to blue).



Figure 9. Comparing the methods based on MIDR and DRM: (a) Original image; (b) Recoloring result by MIDR (zoomed area: blue shadows have been preserved); (c) Recoloring result using DRM (missed the colored-shadows).



Figure 7. An example failure case: Here the planckian light assumption is violated by having a purple light. Since purple is not Planckian, the method wrongfully picked white as the illuminant and purple as the object color. The recoloring shows that even though the object itself looks realistic it does not match the scene illumination.



Figure 11. An example of photo montage: The interreflection of the green grass (zoomed area) in the original image is re-lighted by the red color of the carpet to match the target scene.

other, they make unrealistic assumptions (e.g, Gaussian distribution, Lambertian objects). These shortcomings become apparent when applied to the object color transfer. Fig 10 compares the physics-based color transfer performed using MIDR and DRM models with the methods from [17, 18]. We apply the color transfer only to the presegmented objects. After inferring the object color and two illuminants, MIDR successfully transfers the object color. Note that the methods of [17, 18] mixing the illuminants and object colors resulted in unrealistic images. Furthermore, the resulting objects exhibit different colors than the target objects.

**Photo Fusion.** Fig 11 is an interesting example made possible by our method. The car in Fig 3 is copied into another scene. Here the object is recolored using the estimated intrinsic images. But to match the target scene, the interreflection caused by the grass is re-illuminated using the color of the carpet simply by changing the second illuminant color to the red of the carpet resulting in a more realistic scene where the red carpet is reflected in the side of the car.

## 7. Conclusion and future work

We have presented a method for recoloring single-colored objects based on intrinsic image estimation. The single-color constraint allows us to fit more complex reflectance models which better describe real-world images. Whereas most existing recoloring methods assume a white illuminant, we presented a method to recolor objects taken under colored illuminants, and the more complex case of multiple illuminants. Results on synthetic images demonstrate that our algorithm correctly estimates the intrinsic parameters of the scenes. Further we show that the proposed method is able to achieve physically realistic recoloring results in challenging real-world images. In addition we present how our method improves other photo-editing applications like *Color Transfer* and *Photo Fusion*.

As future research, we will investigate further extensions of the dichromatic reflection model, such as the bi-illuminant reflection model recently proposed by Maxwell [13]. This model allows for the modeling of ambient light which we believe could improve the quality of the




Figure 10. Comparing the *Color Transfer* results by DRM, and [18, 17]. (a) and (f) Original images; (b) and (g) MIDR results; (c) and (h) DRM results; (d) and (i) results by [18]; (e) and (j) results by [17]. Note that the secondary illuminants (interreflections) on the side of the car and the plane wing are lost in (c) and (h), wrongfully transformed in (d), (e), (i), and (j), while being preserved in (b) and (g).

recoloring for the low luminance regions of the image.

## 8. Acknowledgments

This work has been supported by the EU project ERGTS-VICI-224737; the Spanish Research Program Consolider-Ingenio 2010: MIPRCV (CSD200700018); and the Spanish project TIN2009-14173. Joost van de Weijer acknowledges the support of a Ramon y Cajal fellowship.

## References

- [1] X. An and F. Pellacini. User-controllable color transfer. *Eurographics*, 2010. 2
- [2] H. Barrow and J. Tenenbaum. Recovering intrinsic scene characteristics from images. In *A. Hanson and E. Riseman, editors, Computer Vision Systems*, 1978. 1
- [3] M. Drew and G. Finlayson. Realistic colorization via the structure tensor. *ICIP*, 2008. 2
- [4] G. Finlayson, M. Drew, and B. Funt. Spectral sharpening: Sensor transformations for improved color constancy. *JOSA A.*, 11(5), 1994. 3
- [5] G. D. Finlayson and G. Schaefer. Solving for colour constancy using a constrained dichromatic reflection model. *IJCV*, 42, 2002. 4
- [6] D. Freedman and P. Kisilev. Object-to-object color transfer: Optimal flows and smp transformations. In *CVPR*, 2010. 2
- [7] R. Gonsalves. Method and apparatus for color manipulation. *United State Patent 6,351,557*, Feb 26, 2002. 2
- [8] E. Hsu, T. Mertens, S. Paris, S. Avidan, and F. Durand. Light mixture estimation for spatially varying white balance. *SIGGRAPH '08*, 2008. 2
- [9] G. Klinker and S. Shafer. A physical approach to color image understanding. *IJCV*, 4, 1990. 2, 3
- [10] V. Konushin and V. Vezhnevets. Interactive image colorization and recoloring based on coupled map lattices. *Graphics*, 2006. 2
- [11] A. Levin, D. Lischinski, and Y. Weiss. A closed-form solution to natural image matting. *PAMI*, 30, 2008. 2
- [12] S. P. Mallick, T. Zickler, P. N. Belhumeur, and D. J. Kriegman. Specularity removal in images and videos: A pde approach. In *ECCV*, 2006. 2, 3
- [13] B. Maxwell, R. Friedhoff, and C. Smith. A bi-illuminant dichromatic reflection model for understanding images. In *CVPR*, 2008. 3, 7
- [14] B. Maxwell and S. Shafer. Physics-based segmentation of complex objects using multiple hypothesis of image formation. *CVIU*, 65, 1997. 2
- [15] B. Maxwell and S. Shafer. Segmentation and interpretation of multicolored objects with highlights. *CVIU*, 77, 2000. 2, 3
- [16] I. Omer and M. Werman. Color lines: Image specific color representation. *CVPR*, 2, 2004. 2
- [17] F. Pitié, A. C. Kokaram, and R. Dahyot. Automated colour grading using colour distribution transfer. *CVIU*, 107, 2007. 2, 7, 8
- [18] E. Reinhard, M. Ashikhmin, B. Gooch, and P. Shirley. Color transfer between images. *IEEE Computer Graphics and Applications*, 21, 2001. 2, 7, 8
- [19] C. Rother, V. Kolmogorov, and A. Blake. "grabcut": interactive foreground extraction using iterated graph cuts. *ACM SIGGRAPH*, 23, 2004. 2
- [20] A. Shafer and D. Lischinski. Using color to separate reflection components. *Color Research and Application*, 10, 1985. 1, 3
- [21] H. Shen and J. Xin. Transferring color between three-dimensional objects. *Applied Optics*, 44(10), 2005. 2
- [22] P. Tan, L. Quan, and S. Lin. Separation of highlight reflections on textured surfaces. In *CVPR*, 2006. 2
- [23] R. T. Tan and K. Ikeuchi. Separating reflection components of textured surfaces using a single image. In *PAMI*, 2003. 2, 3
- [24] M. F. Tappen, W. T. Freeman, and E. H. Adelson. Recovering intrinsic images from a single image. *PAMI*, 27, 2005. 2
- [25] S. Tominaga and B. Wandell. Standard surface-reflectance model and illuminant estimation. *JOSA*, 6(4), 1989. 3
- [26] J. van de Weijer and R. van den Boomgaard. Least squares and robust estimation of local image structure. *IJCV*, 64(2-3), 2005. 4
- [27] Y. Weiss. Deriving intrinsic images from image sequences. *ICCV*, 2001. 2
- [28] Y. wing Tai, J. Jia, and C. keung Tang. Local color transfer via probabilistic segmentation by expectation-maximization. In *CVPR*, 2005. 2