

Saturation-preserving Specular Reflection Separation

Yuanliu Liu [†], Zejian Yuan [†], Nanning Zheng[†], Yang Wu [‡]

[†] Institute of Artificial Intelligence and Robotics
Xi'an Jiaotong University

[‡] Center for Frontier Science and Technology
Nara Institute of Science and Technology

liuyuanliu88@gmail.com, {yuan.ze.jian, nnzheng}@mail.xjtu.edu.cn, yangwu@rsc.naist.jp

Abstract

Specular reflection generally decreases the saturation of surface colors, which will be possibly confused with other colors that have the same hue but lower saturation. Traditional methods for specular reflection separation suffer this problem of hue-saturation ambiguity, producing over-saturated specular-free images quite often. We proposed a two-step approach to solve this problem. In the first step, we produce an over-saturated specular-free image by global chromaticity propagation from specular-free pixels to highlighted ones. Then we recover the saturation based on priors of the piecewise constancy of diffuse chromaticity as well as the spatial sparsity and smoothness of specular reflection. We achieve this through increasing the achromatic component of diffuse chromaticity, while the magnitudes of increments are determined by linear programming under the constraints derived from the priors. Experiments on both laboratory and natural images show that our method can separate the specular reflection while preserving the saturation of the underlying surface colors.

1. Introduction

Highlights are combinations of diffuse and specular reflection components. They are often observable in dielectric inhomogeneous objects [16]. Removing specular component benefits computer vision tasks such as image segmentation and object recognition. In addition, identifying the specular component reveals the roughness of the surface.

In this paper, we address the problem of separating the specular reflection from the diffuse reflection of a single image. This is essentially an under-constrained problem. A typical way is transferring the chromaticity of specular-free pixels to the highlighted pixels [3][19][21], which makes the under-constrained problem solvable. Traditional methods propagate the chromaticity between adjacent pixels [19]. These local propagations will be blocked by color discontinuities. The chromaticity cannot be transferred




Figure 1. Specular reflection separation. (a) Input image. (b) Ground-truth specular-free image. (c-d) The recovered specular-free image and specular reflection by our method. (e) The result of [19]. (f) The result of [21]. (g-h) The result of the first stage of our method: global propagation of diffuse chromaticity among pixels with the same hue.

from the specular-free regions to the isolated highlights, so there will be residual specular reflection left in the recovered specular-free image. See the green patch in the middle of Fig. 1 (e) for an example. Instead we propose to transfer the chromaticity of specular-free pixels to highlighted ones with the same hue. Since the specular reflection retains the hue of the surface color, our global propagation can reach highlighted pixels all over the image including those in isolated regions (e.g., the green and blue patches in the middle of Fig 1 (g)).

When there are surface colors with the same hue but different saturation in the same image, the less saturated ones may be mistaken to be highlighted during hue-based chromaticity propagation. This problem is called the **hue-saturation ambiguity** [6]. See the right half of the ball in Fig. 1 (h) for an example. Accordingly, these pseudo highlight regions will get higher saturations than the ground-truth in the specular-free image as shown in Fig. 1 (d)-(f). We compensate the increments of saturation through raising the achromatic component of the diffuse chromaticity, making sure that the recovered specular reflection is spatially sparse and smooth (as shown in Fig. 1 (d)) while the

diffuse chromaticity is piece-wise constant (Fig. 1 (c)).

1.1. Related Work

The basis of color-based specular reflection separation is the dichromatic reflection model [16]. This model does not separate the chromatic and achromatic part of diffuse reflection, so it cannot solve the problem of hue-saturation ambiguity.

The two-step strategy – recovering the hue first and then the saturation – is widely used in previous work. Bajcsy et al. [3] uses a hue-based segmentation succeeded by a saturation-based segmentation, taking spatial continuity into account. The segmentations rely on color discontinuities, which is sensitive to noise. Further, isolated highlight regions cannot be handled by this method. Mallick et al. [13] proposed a SUV color space [14] to produce a partial separation of diffuse and specular reflection. This partial separation is similar to the over-saturated specular-free image got in the first stage of our method. In the next stage, the diffuse component of the partial separation is removed through evolving a partial differential equation (PDE) that iteratively erodes the specular component at each pixel. These local operations are good at removing specular reflection within small regions. Tan and Ikeuchi [19] generates a pseudo specular-free image through setting the maximum diffuse chromaticity of all the pixels to be a constant. This specular-free image provides a terminate condition for the iterative chromaticity propagation between adjacent pixels. It can also be used as an initial diffuse image for our method, but the remaining specular reflection does not have bounds as tight as those produced by AP. The local propagation may be blocked by color discontinuities generated by noise. Also, it may go across the borders between different surface colors when the color discontinuities cannot be identified by a proper threshold (see the blue blocks in Fig. 1 (e)).

Separating specular reflection from a single image is an under-constrained problem. To make it solvable, many priors have been proposed including the smoothness of diffuse reflection [18], color-texture correlation [17], the smoothness of the derivative of specular reflection, the local constancy and the sparsity of diffuse chromaticity [6]. These priors benefit our process of recovering the saturation. Klinker et al. described the color distribution of the reflection components by skewed T shapes [7]. However these T shapes are hard to extract for natural images due to texture and noise. Another way to solve the under-constrained problem is using multiple images [15][12][11]. We refer to the survey of [2] for details.

Recently, Akashi and Okatani [1] used sparse non-negative matrix factorization to estimate the body colors and separating reflection components simultaneously. This method does not use any spatial priors.

Yang et al. used an efficient bilateral filtering to smooth

the maximum chromaticity of local regions [21]. The smoothing is guided by an approximated maximum diffuse chromaticity to avoid smoothing between different surface colors. The approximation fails in specific cases that the minimum diffuse chromaticity of two pixels are different, in which case the maximum chromaticity may be propagated between different surface colors.

The main contributions of this paper are: (1) we propagate the chromaticity globally from diffuse-only anchor points to highlighted ones, which can deal with isolated highlight regions; and (2) we solve the problem of hue-saturation ambiguity by embedding priors on the piecewise constancy of diffuse chromaticity as well as the spatial sparsity and smoothness of specular reflection into simple linear constraints.

2. Image Formation

According to the dichromatic reflection model [16], an image with both diffuse and specular reflection can be represented as:

$$\mathbf{I}(p) = m_d(p)\mathbf{\Lambda}(p) + m_s(p)\mathbf{\Gamma} \quad (1)$$

where $\mathbf{\Lambda}$ is the chromaticity of diffuse reflection and $\sum_{i \in \{r,g,b\}} \mathbf{\Lambda} = 1$. $\mathbf{\Gamma}$ is the chromaticity of the illumination. Based on the neutral interface reflection (NIR) assumption [10], the color of specular reflection equals the color of the illumination. Without loss of generality, we assume that the illuminations are always white, so $\mathbf{\Gamma}$ is a constant vector $[1/3, 1/3, 1/3]^T$. For images with chromatic illumination, we first normalize the image by a color constancy algorithm such as [20]. m_d and m_s are the intensity of diffuse and specular reflection, respectively. Highlight removal aims at reducing m_s to 0 for each pixel.

Hue-saturation Ambiguity. Specular reflection changes the saturation of the surface color while retaining their hue [19][3]. Thus specular reflection generates a set of hue-equivalent classes \mathcal{H} , i.e., colors with same hue but different saturation. Note that, surface colors with the same hue will fall into the same class with those produced by specular reflection. To reveal the hue-equivalent class, we rewrite Eqn. 1 to be:

$$\mathbf{I}(p) = m_l(p)(t(p)\mathbf{\Lambda}(p) + k(p)\mathbf{\Gamma}) + \tilde{m}_s(p)\mathbf{\Gamma} \quad (2)$$

where m_l is the illumination, t is intensity of the chromatic part of the diffuse reflection, k is the intensity of the achromatic part of the diffuse reflection, \tilde{m}_s is the **pure specular reflection** without any residual of the achromatic part of diffuse reflection. Note that \tilde{m}_s depends on m_l , but the dependency will not affect the analysis following Eqn. 2, so we do not write it explicitly. Each $\mathbf{\Lambda}$ defines a hue-equivalent class $\mathcal{H}(\mathbf{\Lambda})$ with varying t and k . All the pixels $p \in \mathcal{H}(\mathbf{\Lambda})$

can be represented by the same form of Eqn. 1 after the following transformations:

$$\begin{aligned} m_d &= m_l t \\ m_s &= \tilde{m}_s + m_l k = \tilde{m}_s + m_d r \end{aligned} \quad (3)$$

with $r(p) = k(p)/t(p)$ that measures the ratio between the achromatic and the chromatic part of the diffuse reflection. Note that the specular reflection m_s will contain a portion of diffuse reflection when r is non-zero.

Our goal is to get the pure specular reflection \tilde{m}_s . We achieve this in two stages: (1) find the diffuse chromaticity Λ of each hue-equivalence class through hue-based affinity propagation, which derives the diffuse reflection m_d and specular reflection m_s ; and (2) recover the pure specular reflection \tilde{m}_s through removing the achromatic part of diffuse reflection out of m_s . The details of these two stages will be given in Sec. 3 and 4, respectively.

3. Recovering Diffuse Chromaticity by Hue-based Affinity Propagation

We recover the diffuse chromaticity Λ of each hue-equivalent class through finding an anchor point that is specular-free. According to Eqn. 1, the diffuse chromaticity can be calculated directly by $\Lambda = \mathbf{I}/\|\mathbf{I}\|_1$ if $m_s = 0$.

We use Affinity Propagation (AP) [4] to cluster the pixels into hue-equivalent classes and find their anchor points simultaneously. AP determines the optimal number of clusters automatically, which is suitable for handling images that the number of distinct hues are unknown.

AP requires two types of inputs: the affinity between each pair of pixels and the preference of each pixel to be an anchor point. The affinity is set to be the similarity of hue:

$$\mathbf{A}(p, q) = e^{-\|h(p) - h(q)\|_1} \quad (4)$$

where h is the hue defined in [5].

The preference is defined to be:

$$P(p) = wS(p) \quad (5)$$

where w is a positive weight and S is the saturation defined as follows [5]:

$$\begin{aligned} S &= 1 - \frac{3 \min(I_r, I_g, I_b)}{I_r + I_g + I_b} \\ &= 1 - \frac{3m_d \min(\Lambda_r, \Lambda_g, \Lambda_b) + m_s}{m_d(\Lambda_r + \Lambda_g + \Lambda_b) + m_s} \end{aligned} \quad (6)$$

which is a decreasing function of the specular reflection m_s . In another word, the specular-free pixels often have higher saturation than those highlighted ones. This property can be revealed from Fig. 2 (c) that the highlighted pixels are much darker than their surroundings in the saturation image. As




Figure 2. Affinity Propagation. (a) Input image with red cross indicating anchor points. (b) Hue of input image. (c) Saturation of input image. (d) Hue-equivalent classes. (e) Hue transferred from anchor points of each cluster. (f) Saturation transferred from anchor points of each cluster. We can see that the hue is almost kept the same while the saturation is generally raised more or less.

a result, the preference reflects the probability of each pixel being specular-free.

AP selects a set of anchor points and attaches other points to one of them, which maximizes the summation of the preferences of anchor points and the affinities between anchor points and their belongings. Our affinity matrix (Eqn. 4) encourages pixels with the same hue to be attached to the same anchor point, which forms hue-equivalent classes (see Fig. 2 (d)). At the same time the preference (Eqn. 5) encourages the anchor point of each class to fall at a pixel which is most probably to be specular-free or equivalently, with the highest saturation. Fig. 2 (a) marks all the anchor points of the example image. We can see that none of them locate at highlighted areas. Note that, the affinities will also influence the choice of anchor points, which encourages pixels with common hues to be selected instead of borders of the classes. We refer to the paper of [4] for more details.

The chromaticity of each anchor point will be transferred to the other pixels within the same hue-equivalent class. This process generally keeps the hue of the original image while raising the saturation. These effects can be observed through comparing Fig. 2 (b)-(c) to (e)-(f). Note that, no spatial information is used to measure the affinity between pixels, so pixels far away from each other may also be clustered together. This property ensures that the diffuse chromaticity can be propagated all over the image.

After we got the diffuse chromaticity Λ , Eqn. 1 becomes solvable. The diffuse reflection m_d and specular reflection




Figure 3. Specular-free image from hue propagation. The original image is shown in Fig. 2. (a) The recovered specular-free image. Several regions are over-saturated such as the face of the doll (inside the read circle) and the eye of the fish (the red ring inside the blue circle). (b) The intensity of the diffuse reflectance, which is piece-wise smooth. (c) The specular reflection, which appears not only in highlighted areas but also saturation-enhanced regions. (d) The ratios between specular and diffuse reflection m_r . The saturation-enhanced regions get noticeable non-zero values. (e) The histogram of m_r . While the highest peak locates near zero, saturation-enhanced regions generate lower peaks in the long tail of the distribution. (f)-(g) The histogram of m_r when replacing AP with the specular-free image of [19] and the SUV color space [14], respectively.

m_s can be calculated as follows [19]:

$$\begin{aligned} m_d &= \frac{3I_{max} - \|\mathbf{I}\|_1}{3\Lambda_{max} - 1} \\ m_s &= \|\mathbf{I}\|_1 - m_d \end{aligned} \quad (7)$$

where I_{max} and Λ_{max} are the maximum of the RGB channels of the image \mathbf{I} and the diffuse chromaticity Λ , respectively. Note that Λ_{max} should be larger than $1/3$, so we have to exclude achromatic pixels from the image.

The recovered specular-free image from Eqn. 7 will get lower diffuse reflection than the ground truth. That is because m_d is a decreasing function of the maximum chromaticity Λ_{max} and the propagation of chromaticity from anchor points to highlighted pixels will raise Λ_{max} . Therefore the specular reflection m_s will be higher than the pure specular reflection \tilde{m}_s . These effects can be viewed from Fig. 3. Combining this relation with Eqn. 3 we can get $r \geq 0$, i.e., the ratio between the achromatic and the chromatic part of diffuse reflection is non-negative. Meanwhile, the specular reflection \tilde{m}_s should be always non-negative, so we can get an upper bound $r \leq m_r$, where $m_r = m_s/m_d$. These bounds are useful for inferring r in Sec. 4.

From the histogram of m_r (Fig. 3 (e)) we can see that most m_r are close to 0, but there is also a long tail after the highest peak. These large values may come from those regions whose diffuse reflectance gets lower saturation than the anchor points, e.g., the face of the doll and the eye of the fish in Fig. 3 (d). In Sec. 4, we will recover the saturation of the pixels through inferring r .

The specular-free image produced by some other methods such as [19] and [14] can also be taken as initial diffuse reflection m_r , whose saturation can be recovered in the next step as well. These methods use simple yet effective global transformations directly. However the ratio r produced by these methods do not have tight bounds. [19] sets Λ_{max} (in Eqn. 7) of all the pixels to be a constant. This will break the lower bound of $r \geq 0$ for some pixels, and make the upper bound loose for some other pixels (e.g., Fig. 3 (f)).

Instead AP sets Λ_{max} according to the reference specular-free pixel for each hue-equivalent class separately, which gives a much tighter upper bound while ensuring $r \geq 0$. For the SUV color space [14], "the two-channel diffuse signal suffers severe degradation", as pointed out by its inventors, when "the angle between the image and the source color is less than 10° ". In this case the ratio m_r is sensitive to noise, and the upper bound of r will be loose (e.g., Fig. 3 (g)).

4. Recovering Saturation using Priors on Reflection

Thus far we have got a specular-free image specified by the specular reflection m_s . This specular-free image may contain over-saturated regions due to hue-saturation ambiguity. In this section we further recover the saturation or equivalently, the pure specular reflection \tilde{m}_s based on the piecewise constancy of diffuse chromaticity, the spatial sparsity and the smoothness of specular reflection.

Our analysis is based on the ratio between the achromatic and the chromatic part of the diffuse reflection r . Once we know the value of r , the pure specular reflection can be recovered by

$$\tilde{m}_s = m_s - m_d r \quad (8)$$

according to Eqn. 3 and the pure diffuse reflection can be simply got by $\tilde{m}_d = \|\mathbf{I}_i\|_1 - \tilde{m}_s$.

We formulate the inference of r as the following energy minimization problem:

$$\begin{aligned} \arg \min_r E(r) \\ \text{s.t. } 0 \leq r(p) \leq m_r(p), \forall p \end{aligned} \quad (9)$$

where the objective function is composed of a data term and a regularization term:

$$E(r) = E_{data}(r) + E_{reg}(r) \quad (10)$$

The details of E_{data} and E_{reg} are given in Sec. 4.1 and Sec. 4.2, respectively. The lower bound is derived from the fact




Figure 4. Specular reflectance after correcting the saturation. The original image is shown in Fig. 2. (a) The image of \tilde{m}_r , which is the ratio between specular and diffuse reflection after median filtering. The original ratio is shown if Fig. 3 (d). (b) The inferred r . (c) The specular reflection after saturation recovery.

that we produced an over-saturated specular-free image in the first stage while the upper bound is got from $\tilde{m}_s \geq 0$ as described in Sec. 3.

4.1. Data Term

Spatial Sparsity of Specular Reflection. From Eqn. 3 we can see that, $r = m_r$ when \tilde{m}_s is 0. In natural images, the pure specular reflection \tilde{m}_s is sparsely distributed in the image plane, which means its value is 0 for most pixels within an image patch. The median value of m_r over local patches, denoted by \tilde{m}_r , are expected to fall at specular-free pixels in most of the time. The energy of unary term is modeled as:

$$E_{data}(p) = \|\tilde{m}_r(p) - r(p)\|^2 \quad (11)$$

4.2. Regularization Term

The regularization term is a combination of energies derived from priors on the diffuse chromaticity and the specular reflection:

$$E_{reg}(r) = w_d E_d(r) + w_s E_s(r) \quad (12)$$

where w_d and w_s are weights of the two terms.

Piecewise Constancy of Diffuse Chromaticity. Images with this property are called Mondrian picture, which are originally addressed by the Retinex theory [8] and widely used for intrinsic image decomposition. The piecewise constancy of diffuse chromaticity implies that both the achromatic part k and the chromatic part t are piecewise constant, so their ratio r is piecewise constant. Based on this property, we define the regularizer E_d to be:

$$E_d(r) = \sum_{p,q \in N_p} \alpha(p,q) \|r(p) - r(q)\|^2 \quad (13)$$

where N_p is the neighborhood of p in the image plane. Here the weight α is used to avoid smoothing r across color edges, which is defined as follows:

$$\alpha(p,q) = e^{-\eta \|\Delta h(p,q)\|^2 - \mu \|\Delta m_r(p,q)\|^2} \quad (14)$$

where $\Delta h(p,q)$ is the difference of hue between pixel p and q and $\Delta m_r(p,q)$ is the difference of m_r . η and μ are positive scalars. Small Δm_r implies similar saturation with respect to the maximum channel I_{max} according to Eqn. 7. So the weight α reflects the probability that two neighboring pixels have the same diffuse reflectance.

Smoothness of Specular Reflection. The specular reflection tends to be smooth for smooth surfaces. From Eqn. 8 we can get the following linear relationship:

$$m_d(p)r(p) - m_d(q)r(q) = \Delta m_s(p,q) \quad \text{if } \tilde{m}_s(p) = \tilde{m}_s(q) \quad (15)$$

where $\Delta m_s(p,q) = m_s(p) - m_s(q)$. Since \tilde{m}_s is smooth in the image plane, the regularizer E_s is defined as

$$E_s(r) = \sum_{p,q \in N_p} \beta(p,q) \|m_d(p)r(p) - m_d(q)r(q) - \Delta m_s(p,q)\|^2 \quad (16)$$

Here we set $\beta(p,q) = 1 - \alpha(p,q)$ to make this constraint concentrate on pixels pair with different color. Note that, when $m_d(p) = m_d(q)$ and $\Delta m_s(p,q) \approx 0$, Eqn. 16 degrades to a scaled form of Eqn. 13.

A similar prior of smooth variation of specular reflection is used in [6] and [18]. They formalize the prior by limiting the gradient of the ratio of specular reflection $m_s/(m_s + m_d)$. Note that, their objective function includes the diffuse reflection m_d , which is not guaranteed to be smooth across the border of colors where t , Λ and k may change sharply (see Fig. 3 (b) for an example). Our formulation utilizes only the smoothness of pure specular reflection, which is not affected by the change of m_d .

Optimizing Eqn. 10 essentially solves a linear programming problem under least-squares criteria, since all the energy terms (in Eqn. 11, 13 and 16) come from linear constraints on r . This problem can be solved efficiently by off-the-shelf tools, e.g., the Matlab lsqmin function.

5. Experiments

We evaluate our method on high-quality images taken under controlled environments, which are provided by the authors of [19] and [21]. We also test our method on several low-quality images captured by common cameras under uncontrolled environments.

Implementation Details. Using AP to cluster the pixels is computationally expensive, since it constructs a complete graph that takes each pixel as a node. Instead, we first quantize the pixels into several bins by hue. For pixels falling into the same bin, we randomly select n pixels as the candidate anchor points, and assign other pixels to them by solving the Uncapacitated Facility Location problem [9] with AP. In implementation, we set $n = \sqrt{m}$, where m is the total number of pixels within the bin. The parameter w in Eqn. 5 is set to be 1000 times the mean of affinities (the

elements of \mathbf{A}). Both w_d and w_s in Eqn. 12 are set to be 10. The parameters η and μ are set to be 100 and 1, respectively. We raise the upper bound of r (in Eqn. 10) by 0.5 to make it more robust to noise.

Our model is robust to the parameters. We evaluated the robustness by the standard deviation of PSNR (D-PSNR) of the pear and the circle in Fig. 6. We varied w , μ , η , w_d and w_s among 0.1, 0.5, 1, 2 and 10 times its default value. For w , the D-PSNR is 0.0058dB for the pear and 0.23dB for the circle. For 400 combinations of w_d , w_s ($w_d, w_s < 100$), η and μ , the D-PSNRs are 0.33dB and 1.42dB. When we set N_p to be 4, 8, or 16 and the bin width of hue to be 0.01, 0.02, 0.05 or 0.1, the D-PSNRs are 0.13dB and 0.47dB, respectively.

For the 387x353 image in Fig. 5, the running time of our Matlab code on a 3.4GHz Intel Core i7 cpu is 46.8s, where AP takes 11.2s and optimizing Eqn. 10 spends 28.3s.

5.1. Results on High-quality Images

We first test our method on the image of toys presented in Fig. 5. The main challenge comes from that several groups of regions have similar hue but different saturation, e.g., the pink circle around the eye of the fish and the dark red stripes nearby. In addition, the yellow dots (within the red box on the input image) form several isolated highlight regions, where the color discontinuities may block local chromaticity propagation.

We compare our results to those of Tan and Ikeuchi [19], Yang et al. [21] and Kim et al. [6]¹. We can see that all the methods successfully remove most of the highlights. The difference comes from whether the saturation of the surface color is well preserved or not. In this regard, our model achieves the best performance. The method of Tan and Ikeuchi [19] generates several over-saturated regions in the specular-free image, and the recovered specular reflection is non-zero in many areas that are actually specular-free. Their result is quite similar to that of the first stage of our method. This over-saturation effect may be caused by the false chromaticity propagations between surface colors with different saturations. A similar over-saturation problem, although much less severe, can be observed from the results of [21]. In comparison, both the methods of Kim et al. [6] and ours solve the hue-saturation ambiguity quite well. That is mainly due to the usage of constraints on the smoothness of specular reflection and so on. The method of [6] does not handle textures well enough, e.g., the yellow dots on blue background. One reason is that they smooth the ratio between specular and diffuse reflection instead of specular reflection alone, which is risky at the color boundaries where the diffuse reflection may change sharply. Moreover, they use k-means to cluster pixels into a specific

¹The images are cropped from their paper directly. We refer to their paper for a better view.



Figure 6. Results on images with different levels of texture. (a) The input image. (b) The results of [19]. (c) The results of [21]. (d) Our results.

number of classes, which is not suitable for highly textured images. From the third row we can see that some of the small yellow dots are merged into the blue background.

More results are shown in Fig. 6. These images have different levels of texture. We compare to the method of [19] and [21]. All the methods got similar results on the image of pear, which is almost textureless. The image of fish has moderate texture. Our method successfully removed the specular reflections, except for a little residual in the extremely highlighted areas. The image in the last row is highly textured, which is a great challenge for pixel clustering or image segmentation. Our method still gets promising result on this image. Note that, the white paintings are mostly removed by all the methods, since they are hard to be distinguished with a widespread specular reflection.

Evaluation of Step 1. We evaluate the effect of replacing AP in Step 1 with the specular-free image of [19] or a SUV color space [14]. By doing so the PSNR of the circle in Fig. 6 drops about 2dB and 5dB, respectively. Using the specular-free image of [19] cannot remove the specular reflection thoroughly due to the loose upper bound of r . The SUV color space does not perform well on the near-white regions whose colors are close to source color. Their performances on the pear and the synthesized image in Fig. 1 are comparable to our original model.




Figure 5. Results on a toy example. (a) The input image. The middle row in this column shows the specular reflection recovered in the first stage of our method. (b) The results of [19]. (c) The results of [21]. (d) The results of [6]. (e) Our results. The first two rows show the decomposed diffuse and specular reflection, respectively. The area within the red box of the input image (a) is enlarged and shown in the third row.

5.2. Results on Natural Images


The results on natural images are given in Fig. 7. The objects in the images are made of various materials, including plastic, paper, fabric, foliage, wood and ceramic. These images are quite challenging due to noise, overexposure, heavy texture, chromatic illumination, achromatic surfaces and so on. Despite these difficulties, our method achieved promising results on them. The decomposed specular reflection is sparse and smooth as we expected. In the meantime, the diffuse reflection is piece-wise constant, although some detailed texture are lost. In particular, the image of the potted plant is taken from indoor scenes at night, while the illumination is complex. Our method successfully removed the highlights on the leaves while keeping the underlying yellow-green color. The method of [19] produced over-saturated specular-free images for all the examples. The result of [21] didn't keep the saturation of the root of the plant. Our model also recovered the diffuse reflection of surfaces with extremely low saturation, such as the white eyes of the dolls in the top two images and the fin of the sea horse in the third image. The main reason is that our model constrains the specular reflection to be sparse, so the "direct current component" of the specular reflection got from the

first stage will be eliminated.

6. Conclusions and Discussions

We have proposed a new model to separate the specular reflection, while preserving the saturation of surface colors. The underlying methodology is similar, in some respect, to that of shock therapy. To preserve the saturation, we first increase the saturation to the maximum, producing an over-saturated diffuse reflection. Then we recover the saturation by injecting a proper dose of achromatic components back to the diffuse reflection. The advantage of this two stage process is that the specular reflection will be removed completely after the first stage, so we can get tight bounds of the proper dose of achromatic components. Moreover, we can derive useful constraints on the dose of achromatic components from the priors on local reflections.

Separating the specular reflection of natural images is still challenging. For our method, the main issue is the deviation of hue in highlighted areas, which may be caused by chromatic illumination, overexposure or noise. In this case the hue-equivalent classes will break into pieces, some of which will not occupy specular-free pixels. We left these issues to our future work.



(a) (b) (c) (d) (e)

Figure 7. Typical results on natural images. (a) The input image. (b) The results of [19]. (c) The results of [21]. (d) The results of the first stage of our model (Sec. 3). (e) Our final results.

Acknowledgement. This work was supported in part by the State Key Program of National Natural Science Foundation of China (Grant No. 61231018), the National Basic Research Program of China (Grant No. 2015CB351703) and the 111 Project (Grant No. B13043).

References

- [1] Y. Akashi and T. Okatani. Separation of reflection components by sparse non-negative matrix factorization. In *Asian Conference on Computer Vision*, 2014. 2
- [2] A. Artusi, F. Banterle, and D. Chetverikov. A survey of specular removal methods. *Computer Graphics Forum*, 30(8):2208–2230, 2011. 2
- [3] R. Bajcsy, S. W. Lee, and A. Leonardis. Detection of diffuse and specular interface reflections by color image segmentation. *International Journal of Computer Vision*, 17(3):249–272, 1996. 1, 2
- [4] B. J. Frey and D. Dueck. Clustering by passing messages between data points. *Science*, 315:972–976, 2007. 3
- [5] R. Gonzalez and R. Woods. *Digital Image Processing*. Addison-Wesley, 1993. 3
- [6] H. Kim, H. Jin, S. Hadap, and I. Kweon. Specular reflection separation using dark channel prior. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2013. 1, 2, 5, 6, 7
- [7] G. J. Klunker, S. A. Shafer, and T. Kanade. A physical approach to color image understanding. *International Journal of Computer Vision*, 4(1):7–38, 1990. 2
- [8] E. H. Land and J. J. McCann. Lightness and retinex theory. *Journal of the Optical Society of America*, 61(1):1–11, 1971. 5
- [9] N. Lazic, B. Frey, and P. Aarabi. Solving the uncapacitated facility location problem using message passing algorithms. In *AISTATS*, 2010. 5
- [10] H. Lee, E. Breneman, and C. Schulte. Modeling light reflection for computer color vision. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 12(4):402–409, 1990. 2
- [11] S. Lin, Y. Li, S. B. Kang, X. Tong, and H.-Y. Shum. Diffuse-specular separation and depth recovery from image sequences. In *ECCV*, 2002. 2
- [12] S. Lin and H.-Y. Shum. Separation of diffuse and specular reflection in color images. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2001. 2
- [13] S. Mallick, T. Zickler, P. Belhumeur, and D. Kriegman. Specularity removal in images and videos: A pde approach. In *ECCV*, pages 550–563, 2006. 2
- [14] S. P. Mallick, T. Zickler, D. J. Kriegman, and P. N. Belhumeur. Beyond lambert: reconstructing specular surfaces using color. In *CVPR*, 2005. 2, 4, 6
- [15] Y. Sato and K. Ikeuchi. Temporal-color space analysis of reflection. *JOSA*, 11(11):2990–3002, 1994. 2
- [16] S. Shafer. Using color to separate reflection components. *Color Res. Appl.*, 10(4):210–218, 1985. 1, 2
- [17] P. Tan, S. Lin, and L. Quan. Separation of highlight reflections on textured surfaces. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2006. 2
- [18] P. Tan, S. Lin, L. Quan, and H.-Y. Shum. Highlight removal by illumination-constrained inpainting. In *IEEE International Conference on Computer Vision*, 2003. 2, 5
- [19] R. T. Tan and K. Ikeuchi. Separating reflection components of textured surfaces using a single image. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(2):178–193, 2005. 1, 2, 4, 5, 6, 7, 8
- [20] R. T. Tan, K. Nishino, and K. Ikeuchi. Color constancy through inverse intensity chromaticity space. *JOSA A*, 21(3):321–334, 2004. 2
- [21] Q. Yang, S. Wang, and N. Ahuja. Real-time specular highlight removal using bilateral filtering. In *Proceedings of the European conference on Computer vision*, pages 87–100, 2010. 1, 2, 5, 6, 7, 8