Beyond Lambertian Shape from Shading

Qing Tian and Carl Mueller-Roemer

April 27, 2012

Abstract

Traditional shape from shading methods are based on the Lambertian surface model. Real images often contain specularities which violate this assumption and lead to undesired results. In [1] a method for color subspace based specularity removal has been proposed. This method only works with a good estimate of the source color given as a vector in RGB space. In this project we present two novel methods for estimating the source color from image data alone. One is global optimization based and the other is local neighborhood analysis based. Experimental results show that using our methods we can get accurate estimates of the source color for the color subspaces based specularity removal algorithm and therefore produce images on which traditional shape from shading methods can work.

1 Introduction

In [1, 2] a method for the removal of specularities based on color subspaces was introduced. This method uses a rotation of the RGB color space into an orthogonal SUV color space, where S is the source color, to obtain a diffuse grayscale image from the U and V components. The images processed by this method can later be used for different methods which usually have trouble dealing with specular images like shape from shading and photometric stereo [1]. One of the key problems of the color subspace based specularity removal is that we need to know the colour of the light source S given in RGB components.

Therefore in our course project we have implemented the method as shown in [1] and developed two distinct methods to estimate the source color from given image data alone. While the method for source color estimation referenced in [1] needs a previous non-trivial segmentation procedure for images containing surfaces with multiple colors our methods both don't need such preprocessing. We have tested the resulting specularity free images using the Horn-Frankot shape from shading algorithm [3, 4] which we have implemented in the second course assignment. Our results show that our algorithms both work reliably to estimate the source color.

2 Background

The method presented in [1, 2] assumes that the surfaces in the image can be described appropriately by the dichromatic model. The dichromatic model describes the BRDF of a surface with constant index of refraction as

$$f(\lambda, \boldsymbol{\theta}) = g_d(\lambda) f_d(\boldsymbol{\theta}) + f_s(\boldsymbol{\theta})$$

where λ is the wavelength, $\boldsymbol{\theta} = (\theta_i, \phi_i, \theta_r, \phi_r)$ is an array containing the incident angles and the reflectance angles. The indices d and s stand for diffuse and specular components. This equation assumes that the specular and diffuse components are additive and that only the diffuse component is reflected in dependence of the wavelength. That is to say, each specularity pixel in the image is composed of a diffuse, wavelength dependent component and a specular, wavelength independent component.

Consider an observed surface point x illuminated from direction $\hat{l}$ where the normal vector is $\hat{n}$. From the dichromatic BRDF model above, the image intensity at this pixel x obtained from a camera sensor can be modeled as I_k.

$$I_k = (D_k f_d(\boldsymbol{\theta}) + S_k f_s(\boldsymbol{\theta}))\hat{\boldsymbol{n}}^T\hat{\boldsymbol{l}},\ k = 1, 2, 3$$

with

$$S_k = \int C_k(\lambda)L(\lambda)d\lambda,\ k = 1, 2, 3$$

and

$$D_k = \int C_k(\lambda)L(\lambda)g_d(\lambda)d\lambda,\ k = 1, 2, 3$$

where $\boldsymbol{\theta}$, as defined above, represent the illumination direction and the viewing direction in the local coordinate system. The function $g_d(\lambda)$, often referred to as the spectral reflectance, is an intrinsic property of the material that controls the diffuse component. $C_k(\lambda)$ is the sensitivity function corresponding to color channel k of the camera sensor which can be considered as a linear device and $L(\lambda)$ is the spectral power distribution (SPD) of the light source.

It is worth noting that different from D_k which has a material-related spectral reflectance $g_d(\lambda)$, S_k, representing the effective source strength as measured by the kth sensor, is independent of the material of the surface being observed. Since $f_s(\boldsymbol{\theta})$ is the same for all three channels, the specular component of the pixel x and the source color are aligned in RGB color space.

As described in [1] in more detail one can now use a rotation of the RGB components of each image pixel in order to rotate all of the specular components in the image into one of the color channels, e.g. the first channel and obtain only diffuse components in the two other channels. Since specular components are aligned with the source color, we just need to find a rotation that align one of the RGB axes with the source color vector. Assuming that we know the source color vector and we create two orthogonal vectors U and V in order to span the 3D color space we can find a rotation R. For normed and orthogonal vectors R can be obtained simply from this data.

$$R = \begin{bmatrix} S^T \\ U^T \\ V^T \end{bmatrix}$$

Then the rotation of a pixel $p_{i,j}$ with RGB color component will give us the same pixel in SUV space.

$$[p_{i,j}]_{SUV} = R\,[p_{i,j}]_{RGB}$$

In SUV space according to the theory developed here the S component contains all specular contributions and some of the diffuse contributions, while the U and V component only contain diffuse contributions. Now to use this method on images they have to fulfill certain assumptions. Assuming the dichromatic model is appropriate there had better not be any white surfaces as these will quite likely reflect the source color and make a distinction between specular and diffuse component impossible, also saturation should not occur at the specularities, otherwise the additive nature of the diffuse and specular component will be lost. A diffuse grayscale image of the scene can now be recovered by the vector norm of the second and third component of the U and V channels of each pixel.

$$p^g_{i,j} = \left\| \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} [p_{i,j}]_{SUV} \right\|$$

This method can be applied if the source color vector is known, however it is usually not known, therefore one of the most interesting problems is to recover the source color only from the image data. Previous work on this topic has been done in [5, 6, 7, 8, 9]. In [9] it is noted that color vectors from a dichromatic material seen from different viewpoints lie in the dichromatic plane spanned by the S and D color vectors defined above. [6, 7] show that these vectors often cluster in the shape of a "skewed T" in this dichromatic plane from which one can recover both vectors under the condition that the limbs of the T are sufficiently separated. However, these methods only work for homogeneous, dichromatic surfaces without noisy observations and need a complicated segmentation process to work. Also, further conditions are imposed on the size of the specular lobe which is supposed to be very narrow. Some other methods based on polarization [10] and multiple exposure [11] have also been proposed. Though they are effective in many cases, the complicated image acquisition processes greatly limit their general application. Tominaga and Wandell [12] have proposed an approach to recover the source color that are based on the intersection of dichromatic planes from different dichromatic surfaces, but finding the intersection accurately in the color space is not an easy task itself. [8] uses single homogeneous dichromatic surfaces, but in this method the chromaticity curves of major illuminants must be known beforehand.

3 Source Color Estimation

It is clear from the previous work on source color estimation that most methods are not automatic and effecient enough. They either use complex image acquisition process, identification of multiple surfaces in the image, strict constraints on the homogeneity of surfaces or some previous knowledge about the illuminants. Motivated by this lack of automatic and efficient enough techniques to estimate source color, in this paper we propose two methods for source color estimation from image data. The first approach presented is a global optimization based approach, while the second proposed method is a local neighborhood based approach. Both approaches do not need previous knowledge, complicated segmentation or the use of multiple images. The constraints we impose are largely the same as in [1], with just one additional constraint, that the source color should span all of components of the RGB space, which should not be a major constraint for real world images.

Both of our two approaches needs a pre-processing step in which we identify the areas of the picture that contain the specularities. Under the conditions of the dichromatic reflectance model and an image without white surfaces the highest R, G and B values should be at the specularities, because non-white diffuse parts reflect only a part of the lights source color while at the specularities the diffuse component and the specular component occur together. By multiplication of the R,G and B values and subsequent normalization to the range $[0, 1]$ one can obtain an accurate enough map representing the specularities. The detailed specularity detection process is described as follows:

Let $p_{i,j}$ be one color pixel of the image, we can represent it in RGB space as

$$p_{i,j} = \begin{bmatrix} p^r_{i,j} \\ p^g_{i,j} \\ p^b_{i,j} \end{bmatrix}$$

Then the matrix representing the specularities $\hat{M} = \{\hat{m}_{i,j}\}$ can be obtained by setting

$$\hat{m}_{i,j} = p^r_{i,j} p^g_{i,j} p^b_{i,j}$$

and subsequently normalizing to the desired value range. We call the normalized matrix $M = \{m_{i,j}\}$.

$$M = \frac{\hat{M} - \min(\hat{m}_{i,j})}{\max(\hat{m}_{i,j}) - \min(\hat{m}_{i,j})}$$

The method can often be improved by applying a gamma correction to M, that is by using an exponent value γ to increase the contrast of the image.

$$M_\gamma = \left\{ m^\gamma_{i,j} \right\}$$

Another improvement that can be made is to use a thresholding operation, e.g. we set all values to zero that are below a given value ϑ and remain all other values or set them to 1 if a specularity mask is needed. Figure 1 shows an image of a pear with specular reflections and the specularity mask obtained with this method.

Figure 1: Original image next to the specularity mask ($\gamma = 1$, $\vartheta = 0.5$)

Now that we have the knowledge of where the specularities lie we can try to find the source color. Our two different approaches to estimate source color are described in the following two subsections.

3.1 Global Optimization Based Source Color Estimation

The first approach of ours is to formulate an optimization problem based on the notion that the source color vector should have a high correlation with the specularities, while it should at the same time have a low correlation with the image as a whole, because the non-white diffuse surfaces reflect only a small part of the color spectrum from the light source in most real world cases. Remembering that in the dichromatic reflectance model the specularities show a combination of the source color and a diffuse component we cannot simply use the specularity color as the source color, this would create black holes in the resulting picture. For an area of the image where the object has only one diffuse color this can be illustrated using the vector

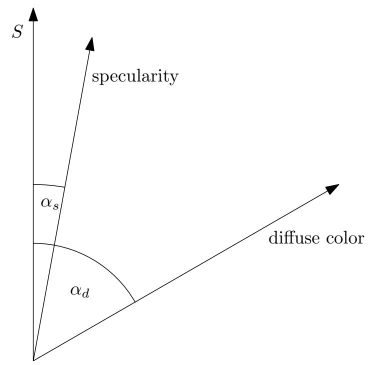

Figure 2: Vector diagram showing how the source color S, the diffuse color and the specular color are related

diagram shown in Figure 2. The specular parts of the image will be close to the source color, but they should contain a diffuse component, which let the specularity be different from the source color vector. We have to find a value for the source color which is close to the color of the specularity while being less close to the diffuse colors found in the image. Naturally this lends to a description as an optimization problem. We write our source color vector as S which has to obey certain constraints to be seen as a color vector and to express that S is a direction not an intensity.

$$S = \begin{bmatrix} s_1 \\ s_2 \\ s_3 \end{bmatrix},\ 0 \leq s_i \leq 1\ \forall i,\ \|S\| = 1$$

The closeness of the source color vector to a color pixel of the image can be expressed using a quadratic form. To measure the closeness of the source color vector to a pixel color we can use the squared cosine of the angle between the source color vector and the pixel's color, expressed through the inner vector product.

$$\left(S^T p_{i,j}\right)^2 = p_{i,j}^T p_{i,j} \cdot \cos^2\left(\alpha_{i,j}\right)$$

Where $\alpha_{i,j}$ is the angle between the source color vector and the pixel color vector. To measure the closeness of the source color vector with the specular parts of the image we can use a weighted sum using our specularity map we obtained above.

$$\frac{\sum_{i,j} M_{\gamma,\vartheta}(i,j)\left(S^T p_{i,j}\right)^2}{\sum_{i,j} M_{\gamma,\vartheta}(i,j) p_{i,j}^T p_{i,j}}$$

In a similar manner we can express the closeness of the source color vector to the the image as a whole.

$$\frac{\sum_{i,j}\left(S^T p_{i,j}\right)^2}{\sum_{i,j} p_{i,j}^T p_{i,j}}$$

To find the true source color vector we can now express this as a single optimization problem by putting these two expressions into one cost function which is to be minimized. As we want to increase the source color vectors correlation with the specular parts of the image and decrease its correlation with the image as a whole, we can add the term to be minimized and subtract the term to be maximized.

$$V = \frac{\sum_{i,j}\left(S^T p_{i,j}\right)^2}{\sum_{i,j} p_{i,j}^T p_{i,j}} - \lambda \frac{\sum_{i,j} M_{\gamma,\vartheta}(i,j)\left(S^T p_{i,j}\right)^2}{\sum_{i,j} M_{\gamma,\vartheta}(i,j) p_{i,j}^T p_{i,j}}, \quad \lambda \in \mathbb{R}^+$$

Some simplifications can be applied here to obtain a simpler expression. As both terms in the cost function are quadratic forms of S we can move S outside. Obtaining a quadratic form depending on a single matrix Γ which can be obtained using the specularity map and the image information, which therefore only has to be computed once.

$$\begin{aligned} V &= S^T \Gamma S \\ \Gamma &= \Gamma_1 - \lambda \Gamma_2 \\ \Gamma_1 &= \frac{\sum_{i,j} p_{i,j} p_{i,j}^T}{\sum_{i,j} p_{i,j}^T p_{i,j}} \\ \Gamma_2 &= \frac{\sum_{i,j} M_{\gamma,\vartheta}(i,j) p_{i,j} p_{i,j}^T}{\sum_{i,j} M_{\gamma,\vartheta}(i,j) p_{i,j}^T p_{i,j}} \end{aligned}$$

It would be false to assume that the solution to this optimization problem is the trivial solution usually obtained for quadratic forms of this type, which is the zero vector. The constraint on S that it's norm is supposed to be 1 doesn't allow this solution, which would make the source color vector meaningless. Instead S only has two degrees of freedom which are further constrained by the lower and upper bound on each value in S. There are nonlinear optimization solvers in most numerical optimization libraries which can handle this kind of constrained nonlinear optimization. Because of the small size of the problem a solution is obtained quickly. The optimization problem can be expressed as

$$\begin{aligned} \min_S V &= S^T \Gamma S \\ \text{s.t. } s_i &\leq 1 \\ s_i &\geq 0 \\ \|S\| &= 1 \end{aligned}$$

Now one of the remaining questions is how to choose the parameterλ representing the tradeoff between our two goals. We found that if λ is chosen to be a small value the image brightness saturates in large areas. Too large values of λ however lead to the specularities turning into dark spots, however the choice of $\lambda = 1$ seems to give good results in most cases. Figure 3 shows the grayscale image of the pear used before and the UV space image after applying the SUV color space rotation [1] .

Applying the algorithm for $\lambda = 1$ on different images leads to the results shown in Figure 4.

This approach can give results quickly for images of any size, because the only size dependent calculations are performed only once and the optimization problem is always of very small scale. Besides this global optimization based approach, we have also developed another approach. Different from our global optimization approach, the second approach is local and focus on individual specularity neighborhoods.

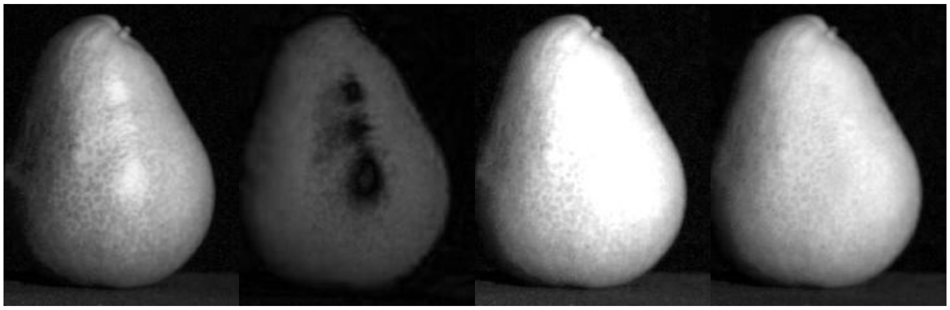

Figure 3: Original image in grayscale and the UV space images for different choices of λ.

Figure 4: Comparison of images with the UV space images for automatic guess of λ

3.2 Local Neighborhood Analysis Based Source Color Estimation

When human beings analyze specularities, the neighborhood is always taken into consideration. For example, just given a small specular region without any context information, most likely people will not even identify this region as specular let alone make some reasonable inference about it. Inspired by this biological intuition, a local specularity neighbourhood analysis approach is proposed to estimate the source color of the scene.

Once the specularity map is obtained, the pixel with the largest intensity of each specular area, which is called the "seed", is selected as a representative of that specularity. Neighborhood analysis is then carried out around each of the detected "seeds". The basic idea of the local neighborhood based approach is to estimate the diffuse component within a local neighborhood first. After obtaining the diffuse component of the "seed" in a specular area, according to the dichromatic model the normalized source color direction can be recovered by subtracting the diffuse component from the mixed color observed and then normalizing the result to unit length. Details of this approach are described in the following subsections.

3.2.1 Neighborhood sampling and outlier removal

Because the genuine diffuse component underlying a specularity seed is in theory unrecoverable, in this paper we will attempt to estimate it from samples of the surrounding diffuse area. In order to sample accurately, general characteristics of specular regions should be analyzed first. After examining many specular regions of real world images, it is revealed that the intensity of each channel in specular regions could be modeled as a hill-like 3D surface and intensity variations within the specular region in each channel are much larger than those in the surrounding matte area. A typical example is envisioned in Figure 5.

Both observations are not difficult to understand. Specular regions stick out due to the additional specular components in each of the three channels while that specular regions have

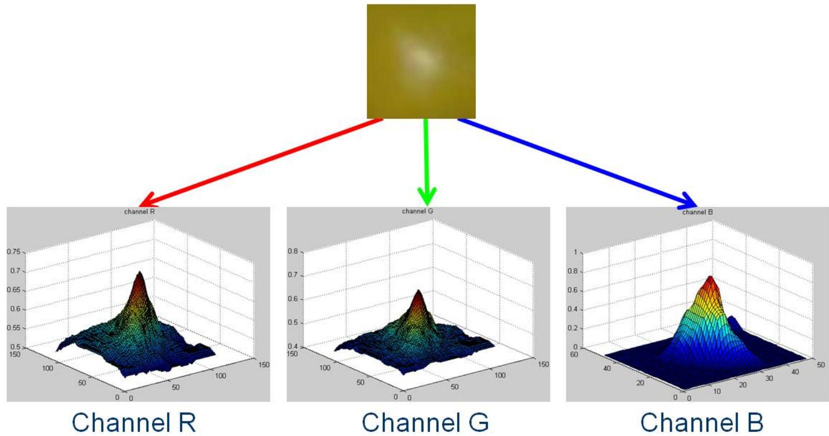

Figure 5: Intensities of specularities could be typically modelled as a "hill" in each channel.

a larger intensity variation is mainly caused by the fact that specular components are very sensitive to changes in viewing direction. Based on these characteristics, a simple downhill algorithm is proposed to descend the specularity "hills". Instead of using some geometrical curve which assumes some dependence among each direction, eight separate detectors are generated outwards from the "peak" of the specularity "hill". For each detector, a sensing range is set and within each step outward the variances of pixels in the sensing range in each of the three channels are calculated. If and only if all three variances become less than a predefined threshold, the approximately planar area is considered to be reached and the detector stops there where the underlying colour is assumed to be a possible diffuse colour. By considering three channels, this downhill algorithm is quite robust: even if the extremely rare case where the specular light is absent in one or two channels happens, this downhill algorithm will still function. Figure 6 shows this method's efficacy.

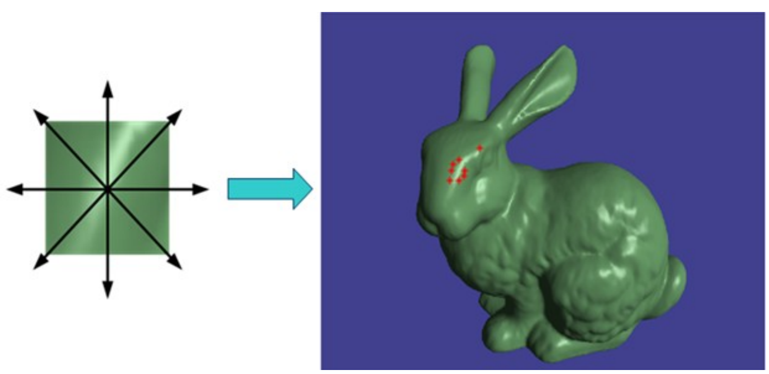

Figure 6: Demonstration of the downhill method

The downhill method generally works very well. However, in some rare cases where the specularity happens to be on the boundary of a surface, some of the detectors will step over the boundary and land on another surface usually with a different diffuse colour. In this case, outlier removal is needed to get rid of these "false" colors. In this paper, the binary tree quantization technique proposed in [13] is adopted to cluster pixels based on their colors. After that, the biggest cluster is selected as the cluster of "correct" diffuse colors.

3.2.2 Source colour estimation and specularity removal

Now that a bunch of colour samples in the matte area are obtained, how to estimate the final diffuse colour underlying the specularity seed from these samples? A naive answer would be taking the average. However, the best result is not guaranteed by just assigning each colour vector with equal weight. Instead, this paper takes an optimization approach. The best diffuse vector is defined as the one that has the highest overall correlation with the sample colour vectors. Since there are several such vectors, a balance should be stricken among all the individual correlations. Thus, the best final diffuse vector should be the one that maximizes the following σ^2, which involves a sum of the squared cosines that measure correlation.

$$\begin{aligned}\sigma^2 &= \frac{1}{N_{sample}} \sum_i (clr_i^T v)^2 \\ &= \frac{1}{N_{sample}} (v^T C C^T v)^2\end{aligned}$$

In above equation, clr_i is the colour vector of one sample in the matte area, which is normalized to prevent the bias caused by vectors with higher magnitudes. C is a matrix whose columns are the bunch of colour vectors of the samples in the matte area. N_{sample} is the number of samples. The best unit vector v that maximizes this variance can be found by solving the following constrained optimization problem.

$$\begin{aligned}v &= \arg\max_v (v^T C C^T v + \lambda(1 - v^T v)) \\ &= \arg\max_v (v^T \left[C C^T - \lambda I \right] v + \lambda)\end{aligned}$$

where the Lagrange multiplier must be determined so as to enforce the constraint that v is a unit vector. After taking the derivative of above formula with respect to v and equating the result to 0, the problem turns out to be an eigenvalue problem in the following form.

$$C C^T v = \lambda v$$

The unit lenth eigenvector associated with the eigenvalue one is chosen as the final diffuse direction. By multiplying it with the magnitude that is estimated as average brightness, the diffuse component is achieved. According to the dicromatic model, the specular component, which has the same direction with the source color, can then be acquired by just subtracting the estimated diffuse component from the mixed colour observed. As an example, the source colour direction of the bunny image is estimated using the local approach and then input to the color subspaces based algorithm [1] to removal specularities. The resulting specularity-free image along with the estimated source color direction is shown in Figure 7.

As can be seen from Figure 7, our local approach also works quite well. All the specularities on the bunny have been removed using the estimated source color direction. Since this approach is local, it can also deal with the situation where different parts of the scene are illuminated by different light sources. Last but not least, this approach does not need much time and memory, which is of great importance to real-time tasks.

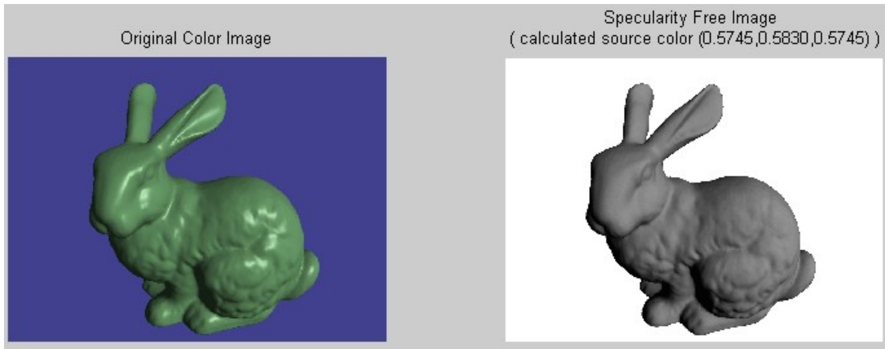

Figure 7: An example result of source color estimation and specularity removal

4 Application to Shape from Shading

In the experiment part, we test the images obtained through our algorithms by using a shape from shading to recover the shape of the surfaces in an image. For this purpose we made an artificial image of a vase with directional lighting, orthogonal projection and additive specularities simulated by use of the Blinn-Phong model. An artificial image is used because we need an object with uniform albedo for the shape from shading algorithm to work satisfactorily. And we need a sufficiently simple shape so that the effect of specularities is more visible. The image we made is shown in Figure 8.

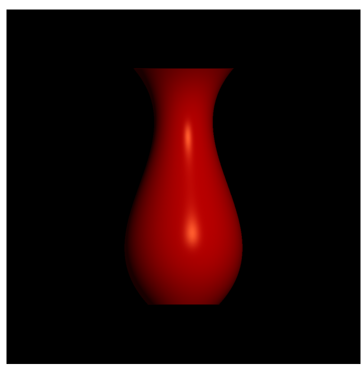

Figure 8: Simulated shaded image of a vase

An artificial image also has the benefit that the light color S is known.

$$S = \begin{bmatrix} 0.58 \\ 0.73 \\ 0.36 \end{bmatrix}$$

Our global approach estimates the source color vector as $\hat{S}_{global}$.

$$\hat{S_{global}} = \begin{bmatrix} 0.57 \\ 0.74 \\ 0.37 \end{bmatrix}$$

Our local approach estimates the source color vector as $\hat{S}_{local}$.

$$\hat{S}_{local} = \begin{bmatrix} 0.60 \\ 0.72 \\ 0.36 \end{bmatrix}$$

As can be seen that both of our two approaches have a very close estimate of the source color vector. Since the difference is negligible, in the following only the specularity-free image generated by the global approach is used and it is worth noting that the results are almost the same by using the local approach. The grayscale image and the specularity-free diffuse image in UV space are shown in 9. Note that because the images are all normalized to have

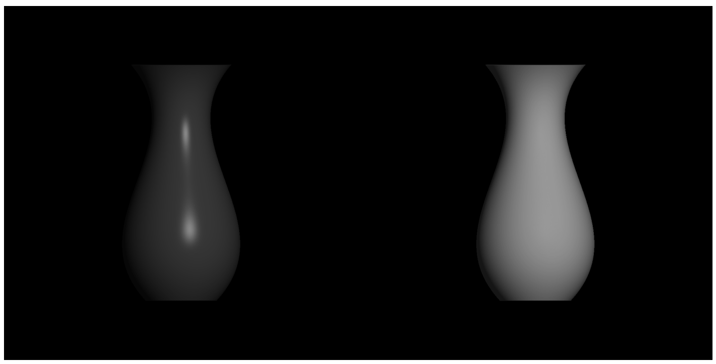

Figure 9: Original grayscale and specularity-free image of the vase

the brightest pixel set to one, the specularity-free image in UV space is brighter than the grayscale image. We use the shape from shading algorithm described in [14], which is based on the techniques from [3] and [4]. The results of shape from shading on both the original specular image and our processed specularity-free image are shown in Figure 10.

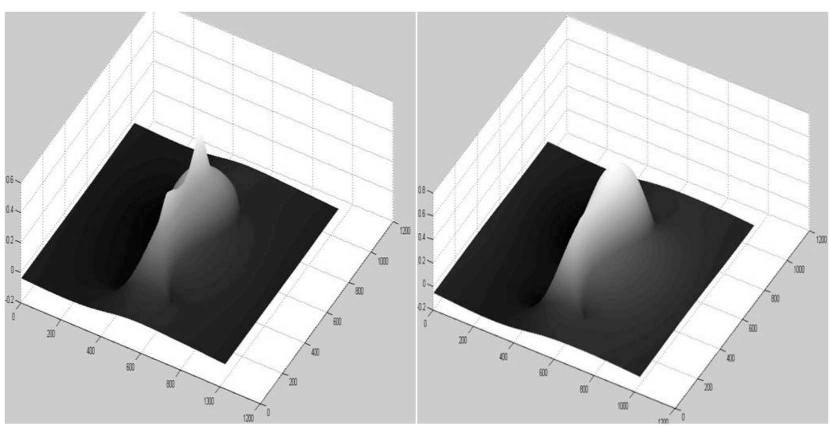

Figure 10: Shape from shading results with (right) and without (left) specularity removal as a preprocess

As can be seen from above results, two bulges on the body of the vase caused by specularities have been successfully removed using images produced by our algorithm, which demonstrates our algorithm's efficacy for shape from shading.

5 Conclusion

In this paper, we have implemented the basic idea of the color subspaces based specularity removal algorithm in [1]. Motivated by one great limitation of this algorithm that the source color should be known as a prior, we have proposed two new approaches for source color estimation which do not need such preprocesses as image segmentation. Both methods have been shown to work well on various images. Using our algorithms we could produce specularity-free images automatically from just raw image data without any human intervention and thus have generalized the application of shape from shading to more complex specular scenes. Experimental results show that the quality of shape from shading results has been improved.

It is worth noting that most of the work are done cooperatively. However, the local neighborhood analysis based approach is more credited to Qing Tian while the global optimization based approach is more credited to Carl Mueller-Roemer.

6 Acknowledgement

Firstly, we would like to express our deep and sincere gratitude to our course instructor, Professor Kaleem Siddiqi whose wide knowledge and logical way of thinking and teaching have inspired us and have laid the foundation for this paper as a course project report. Also, we would like to acknowledge our TA Mr. Emmanuel Piuze who has given us very precious suggestions on the shape from shading assignment. Finally, some experimental images are collected from the Internet. Thanks are also due to their providers.

References

[1] S. Mallick, T. Zickler, D. Kriegman, and P. Belhumeur, "Beyond Lambert: Reconstructing Specular Surfaces Using Color," in *CVPR*, 2005. [Online]. Available: http://www.cs.virginia.edu/~{}mjh7v/bib/Mallick05.pdf

[2] T. Zickler, S. Mallick, D. Kriegman, and P. Belhumeur, "Color Subspaces as Photometric Invariants," *International Journal of Computer Vision*, vol. 79, no. 1, pp. 13–30, 2008. [Online]. Available: http://dx.doi.org/10.1007/s11263-007-0087-3

[3] B. K. P. Horn and M. J. Brooks, "The variational approach to shape from shading," *Computer Vision, Graphics, and Image Processing*, vol. 33, no. 2, pp. 174–208, 1986.

[4] R. T. Frankot, R. Chellappa, and S. Member, "A method for enforcing integrability in shape from shading algorithms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 10, pp. 439–451, 1988.

[5] G. D. Finlayson and G. Schaefer, "Constrained dichromatic colour constancy," in *ECCV (1)*, ser. Lecture Notes in Computer Science, D. Vernon, Ed., vol. 1842. Springer, 2000, pp. 342–358.

[6] R. Gershon, *The use of color in computational vision*, ser. Technical reports on research in biological and computational vision at the University of Toronto. Dept. of Computer Science, University of Toronto, 1987. [Online]. Available: http://books.google.ca/books?id=DDKgSgAACAAJ

[7] G. J. Klinker, S. A. Shafer, and T. Kanade, "The measurement of highlights in color images," 1988.

[8] H.-C. Lee, "Method for computing the scene-illuminant chromaticity from specular highlights," *J. Opt. Soc. Am. A*, vol. 3, no. 10, pp. 1694–1699, Oct 1986. [Online]. Available: http://josaa.osa.org/abstract.cfm?URI=josaa-3-10-1694

[9] S. A. Shafer, "Color," G. E. Healey, S. A. Shafer, and L. B. Wolff, Eds. USA: Jones and Bartlett Publishers, Inc., 1992, ch. Using color to separate reflection components, pp. 43–51. [Online]. Available: http://dl.acm.org/citation.cfm?id=136809.136817

[10] L. Wolff and T. Boult, "Constraining object features using a polarization reflectance model," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, pp. 635–657, 1991.

[11] S. K. Nayar, X.-S. Fang, and T. Boult, "Separation of reflection components using color and polarization," *Int. J. Comput. Vision*, vol. 21, no. 3, pp. 163–186, Feb. 1997. [Online]. Available: http://dx.doi.org/10.1023/A:1007937815113

[12] S. Tominaga and B. A. Wandell, "Standard surface-reflectance model and illuminant estimation," *J. Opt. Soc. Am. A*, vol. 6, no. 4, pp. 576–584, Apr 1989. [Online]. Available: http://josaa.osa.org/abstract.cfm?URI=josaa-6-4-576

[13] M. Orchard and C. Bouman, "Color quantization of images," *Signal Processing, IEEE Transactions on*, vol. 39, no. 12, pp. 2677 –2690, dec 1991.

[14] E. Trucco and A. Verri, *Introductory Techniques for 3-D Computer Vision*. Upper Saddle River, NJ, USA: Prentice Hall PTR, 1998.