

A Truncated Least Squares Approach to the Detection of Specular Highlights in Color Images

Jae Byung Park and Avinash C. Kak
Robot Vision Laboratory
Purdue University
West Lafayette, IN. 47907, U.S.A.
email:{jbpark,kak}@ecn.purdue.edu

Abstract—One of the most difficult aspects of dealing with illumination effects in computer vision is accounting for specularity in the images of real objects. The specular regions in an image are often saturated - which creates problem for all image processing algorithms that use decision thresholds. Such algorithms include those for edge detection, region segmentation, etc. Detecting specularity and whenever possible compensating for it are obviously advantageous. Along these lines, this paper represents a new specularity detection and compensation method which is based on the notion of truncated least-squares approximation to the function that maps the color distribution between two images of an object under different illumination conditions. We also present a protocol for the evaluation of the current method for specularity detection. Our protocol as currently formulated uses human subjects to grade the specularity detection method.

I. INTRODUCTION

The color values in the image of an object surface can vary dramatically as the illumination condition changes. Depending on the location of the illumination sources with respect to the locations of the camera and the object, highly saturated highlights produced by mirror-like reflections from glossy surfaces can appear in an image. Since camera saturation caused by specularity interferes with the image processing algorithms that use decision thresholds, vision researchers from the early days have sought methods for the detection and compensation of specularity[1], [2].

Researchers have proposed a number of specularity detection algorithms[3], [4], [5] that are based on the notion of “well-definedness” of mappings between the chromaticity spaces corresponding to the images taken under different illumination conditions. Using the Dichromatic Reflection Model, Shafer and Klinker[4], [6] have shown that the two components of this model, the diffuse (body) and the specular (surface), form a dog-legged (a skewed “ \perp ” shape) distribution in the dichromatic plane in the RGB color space. Their algorithm fits a convex polygon to the distribution in the dichromatic plane to separate the specular component from the diffuse component. In early 1990, Lee and Bajcsy[5], [7] proposed a specularity detection algorithm called “*spectral differencing*”. Using two color images from different viewing directions, the

spectral differencing approach finds the pixels in one image whose RGB color values do not exist (within a specific tolerance) in the other image; these pixels are declared to be specular. This is done by calculating the Euclidean distance from the RGB coordinates of each pixel in one image to all pixels in the other image. The authors use the notion of *Lambertian Constancy* and the assumption that, under Lambertian reflection, the color coordinates at non-specular pixels are independent of the camera viewpoint. As another closely related work, Drew[8] defines specularities as pixels that disobey the linearity of the Lambertian Model and treats specularities as outliers in the Least-Median-of-Squares (LMedS) regression.

Most of these previous works have been tested either with synthetic images or with a very small number of real images for their ability to detect specularities. In addition, a common constraint in these previous works is that a scene only include Lambertian surfaces. However, many applications in real life environments include shiny surfaces that are not Lambertian; such surfaces cause pixels to become saturated in a camera image. One commonly occurring example of this phenomenon is in the tracking of vehicles in automobile assembly lines where smooth and shiny car surfaces can cause a significant portion of a camera image to saturate and to thus get clipped. As the extent of specularity increases – because a scene has too many surfaces with mirror-like finish – the previously mentioned approaches degrade rapidly. We believe that our proposed method described here overcomes this problem. This method, which we have named as the Truncated Least-Squares Method, constructs a least-squares regression between two images taken under two different illumination conditions, one of these being the image in which the detection of specularities is sought and the other being a reference image substantially free of specularities. The basic idea is to choose candidate thresholds for the detection of specularity, clip the target image using these thresholds, construct a least-squares map from the pixel distribution of the reference image to the retained pixels in the target image, and, finally, use a measure of the quality of this mapping for accepting or rejecting the candidate thresholds.

On the basis of the experimentation we have carried out on a database of 300 images containing both Lambertian and non-Lambertian surfaces, we believe that our proposed approach is sound. No prior information about the illumination condition need be known. Instead, our proposed method requires that the target image has at least some parts that are unsaturated in all camera channels whereas the reference image needs to be substantially specularity-free. Reference images are allowed to contain some specular pixels, as evidenced by the reference images in Figure 1 (m) and (p). Our results show that acceptable specularity detection thresholds can be computed by our method as long as the number of specular pixels in the reference image is less than 5% of the total.

II. EXPERIMENTAL SETUP

To perform our experiments, we employ a Linux system operating on a 850MHz Pentium III and a Matrox Meteor image grabber. All images were taken with a CCD Camera (Jai CV-950) without any gamma correction. As mentioned already, our image database consists of 300 images. These are of 100 widely varying objects, ranging from very dull to very glossy; single colored to multi-colored; with background ranging from simple to complex; etc. Each of these 100 objects was photographed three time under three different illumination conditions: diffuse, ambient, and directed.¹ Figure 1 shows examples of images from the image database.

III. PROPOSED METHOD

Our proposed Truncated Least Squares (TLS) Method is based on discarding the saturated (which is the same thing as clipped) specular pixels from the target image so that a high-quality least-squares mapping can be established between the corresponding pixels in a reference image and what remains of the target image. Eventually, this mapping function also helps us discover the unsaturated specular pixels in the target image. For an image to serve as a *reference image*, it has to be recorded under special illumination conditions, obviously. Additionally, only those pixels in the target image that are free of clipping participate in the mapping function.

The mapping function is derived assuming the dichromatic model for reflectance. Under this model, as suggested by Klinker and Shafer[4], the pixels in an image can be classified into three linear clusters: 1) non-specular (matte) pixels, 2) unsaturated (unclipped) specular pixels, and 3) saturated (clipped) specular pixels. Whereas the first two groups of pixels, referred to as good pixels, are co-planar and define the dichromatic plane, the saturated pixels are not. Since the dichromatic plane does

¹This paper and the entire image database of 300 images are available from the Purdue University Robot Vision Laboratory's web site (http://rvl1.ecn.purdue.edu/RVL/specularity_database).

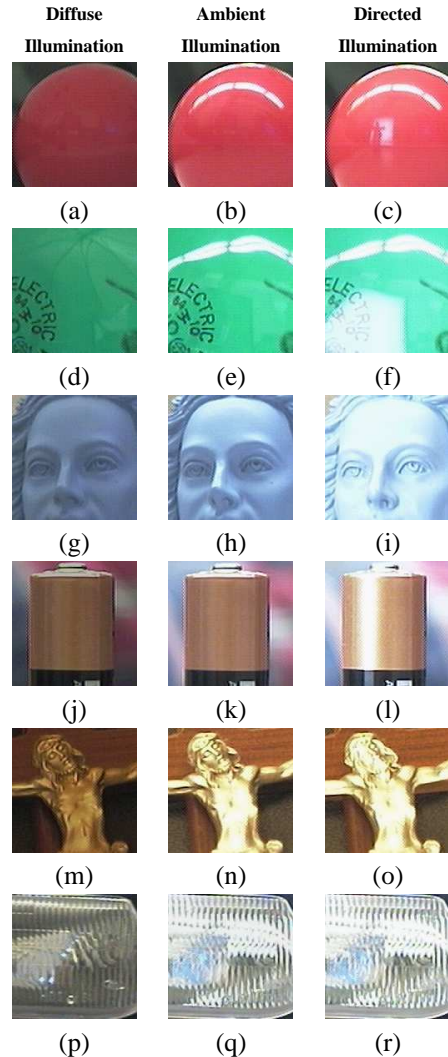


Fig. 1. Example images taken under three different illumination conditions.

not include these saturated specular pixels, obviously the presence of the saturated specular pixels leads to a poor LS (Least-Squares) mapping from one image to another. By discarding the clipped pixels of the target image in this manner, our TLS method gives good mappings even when the target image contains a large number of them.

The TLS method consists of the following three steps:

- Discard the saturated specular pixels in the target image. These are pixels in very bright and saturated areas where the dynamic range of the camera may be exceeded in at least one of the three channels. Since these clipped specular pixels do not generally obey the characteristics of the Dichromatic Reflection Model [4], it is desirable and important to filter these pixels out.
- Calculate a linear regression between the remaining

corresponding pixels in the two images. This will serve as a least-squares mapping from the color values at the good pixels in the reference image to the color values at the good pixels in the target image.

- Detect unsaturated specularities by applying the inverse mapping function to all good pixels of the target image and by comparing the domain values obtained with the distribution of the reference image. If a domain value thus obtained is outside 1.125 times of the standard deviation from the mean of RGB for the reference image, we know that the corresponding pixel in the target image is specular.

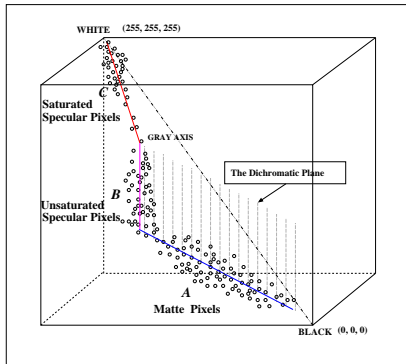


Fig. 2. Three Clusters of pixels distribution in the RGB cube

The following subsections will elaborate on the three steps listed above.

A. Selection of Optimal Thresholds for Removing Saturated Specular Pixels

As mentioned earlier, the Dichromatic Reflection Model [4] describes the color of every pixel in a color image as a mixture of a diffuse (body) reflection and a specular (surface) reflection component. In addition to these two clusters in the dichromatic plane, there are pixels corresponding to saturated specular reflections that do not follow the characteristics of the Dichromatic Reflection Model. These saturated specular pixels do not lie in the dichromatic plane (see Figure 2).

We will now show how a threshold for detecting the saturated specular pixels can be obtained from a Principal Components Analysis (PCA) of the covariance matrix of the color values of a portions of the pixels in a target image. The main idea here is that for each target image we gradually remove saturated pixels by starting from the white point where R, G and B values are all equal to 255 in descending brightness in the RGB cube until all remaining pixels are in the dichromatic plane. Obviously, when all the pixels can be assumed to lie on the dichromatic plane, the color values of the pixels will have only two degrees of freedom, implying that a Principal Components Analysis (PCA) of the color values will yield only two significant

eigenvalues. Additionally, the presence of saturated pixels will cause a significant increase in the “energy” along one of the principal axes of the color values, resulting in a large eigenvalue corresponding to that axis. This implies that if we carry out a PCA of the color values of the pixels of a target image containing specularities, we should expect to see all three eigenvalues, with the largest eigenvalue being of notable size.

Our above observations imply that a search for the optimum threshold for the detection of saturated specularities can be conducted by stepping along each of R, G, and B axes, trying each value as a candidate decision threshold, discarding all the target image pixels whose color values exceed the threshold, and remapping the reference image to the remaining target image pixels. In practice, if the illumination can be assumed to be approximately white, this search can be simplified by stepping along only the gray axis. Our experiments show that this approach to search suffices, at least for the very diverse images of our database. To summarize this search strategy, here are the steps:

- Set the initial value of the threshold for the detection of saturated specularities to the white point in the RGB space.
- Discard all pixels in the target image that have RGB values larger or equal to the current threshold and then apply the PCA to the remaining pixels.
- Reduce the current thresholding values in descending brightness direction along the gray axis in the RGB space.
- Repeat the above steps until the largest eigenvalue decreases sharply.

We have found empirically that these optimal thresholds for removing saturated specular pixels almost always range from the point where R, G, and B values are between 240 and the value of 255 for the white point. Therefore only a few iterations of the steps described above provide the optimal thresholds if any clipping effects are presents. Figure 3 illustrates the changes in the three eigenvalues for the target image shown in Figure 1(e) as the thresholds are reduced one step at a time in descending brightness direction along the gray axis. The second largest eigenvalue (λ_2) and the third largest eigenvalue (λ_3) are relatively smaller compared to the first (largest) eigenvalue (λ_1). The largest eigenvalue is proportional to the distortion introduced in the placement of the dichromatic plane by the presence of saturated pixels. The second and the third eigenvalues correspond to those pixels that form the two linear clusters, marked A and B in Figure 2, in the dichromatic plane. Removing saturated specular pixels significantly decreases the largest eigenvalue.

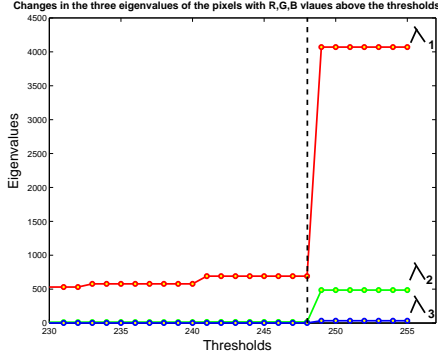


Fig. 3. Changes in eigenvalues as we step through the clipping points along the gray axis. λ_1 , λ_2 , and λ_3 are the three eigenvalues in decreasing order of magnitude. These are obtained by applying PCA to the pixel distribution.

B. Computing TLS Regression

Suppose that after the removal of the saturated pixels in the target image, we are left with only k pixels. We denote the RGB triple for a pixel in the reference image by $\omega_i = [R_i, G_i, B_i]^T$ where R_i , G_i and B_i are the red, green and blue values of the pixel i respectively. In the same manner, we denote the RGB triple for the same pixel in the target image by ω'_i . These k non-discarded pixels in the target image and the corresponding pixels in the reference image are used for the following least-squares minimization:

$$\min \sum_{i=1}^k (r_i)^2 \quad (1)$$

where r_i is the residual that can be defined by the Euclidean distance between the i^{th} pixel in one image and the corresponding pixel in the other image.

In an ideal physics-based reflection model such as the Dichromatic Reflection Model, the tristimulus coordinates (*i.e.*, RGB color coordinates) of the light emanating from the surface patch with surface spectral reflectance function $\rho(\lambda)$ can be found by the following integrals (assuming that the illumination is aimed perpendicularly on the surface):

$$R_{\text{reflected}} \equiv \int_{\nu(R)} P(\lambda) \rho(\lambda) q_R(\lambda) d\lambda \quad (2)$$

$$G_{\text{reflected}} \equiv \int_{\nu(G)} P(\lambda) \rho(\lambda) q_G(\lambda) d\lambda \quad (3)$$

$$B_{\text{reflected}} \equiv \int_{\nu(B)} P(\lambda) \rho(\lambda) q_B(\lambda) d\lambda \quad (4)$$

where $P(\lambda)$ is the spectral power distribution (SPD) of the illuminant and $q_R(\lambda)$, $q_G(\lambda)$, and $q_B(\lambda)$ are the three camera sensor response functions for the R, G, and B channels, respectively. Note also that $\nu(R)$, $\nu(G)$, and $\nu(B)$

are the intervals corresponding to the visible spectrums of R, G, and B respectively[9].

According to the ‘‘factor model’’ proposed by Borges[10], the RGB color values in the reflected light can be approximated by the product of the RGB color components of the illumination e and the RGB components of the surface color s as follows:

$$R_{\text{reflected}} \simeq \frac{e_R s_R}{\sigma_R} \quad (5)$$

$$G_{\text{reflected}} \simeq \frac{e_G s_G}{\sigma_G} \quad (6)$$

$$B_{\text{reflected}} \simeq \frac{e_B s_B}{\sigma_B} \quad (7)$$

where e_R is the R color value of the light source, that is $e_R = \int_{\nu(R)} P_R(\lambda) q_R(\lambda) d\lambda$, (e_B and e_G can be defined accordingly); s_R is the R color value of surface under spectrally white illumination of unit intensity, that is $s_R = \int_{\nu(R)} \rho(\lambda) q_R(\lambda) d\lambda$, (s_G and s_B can also be defined accordingly); and σ_R is the camera scaling term for R channel that is $\sigma_R = \int_{\nu(R)} q_R(\lambda) d\lambda$, (σ_B and σ_G can be defined accordingly). Indeed, a similar approximation was presented by Cowan and Ware[11] and was shown that this approximation works quite well in practice. The expressions above again assume that the illumination is aimed perpendicularly on the surface.

To develop more general versions of the above simplified formulas for the case of slant illumination angles and multiple illumination sources, we can first write the more general version of Eqs. 2 through 4 as follows:

$$R_i = \sum_{l=1}^L \left(\int_{\nu(R)} P_l(\lambda) \rho(\lambda) q_R(\lambda) d\lambda (\bar{a}_l^T \bar{n}_i) \right) \quad (8)$$

$$G_i = \sum_{l=1}^L \left(\int_{\nu(G)} P_l(\lambda) \rho(\lambda) q_G(\lambda) d\lambda (\bar{a}_l^T \bar{n}_i) \right) \quad (9)$$

$$B_i = \sum_{l=1}^L \left(\int_{\nu(B)} P_l(\lambda) \rho(\lambda) q_B(\lambda) d\lambda (\bar{a}_l^T \bar{n}_i) \right) \quad (10)$$

where L is the number of illumination sources, at angles given by the unit normal vectors \bar{a}_l , $l = 1, \dots, L$ and where \bar{n}_i is the surface normal at a point in the scene that corresponds to the i^{th} pixel in the image.

By replacing the integral parts by the approximations given in Eq. 2, 3, and 4 above, we have

$$R_i \simeq \sum_{l=1}^L (e_R^l s_R / \sigma_R) (\bar{a}_l^T \bar{n}_i) \quad (11)$$

$$G_i \simeq \sum_{l=1}^L (e_G^l s_G / \sigma_G) (\bar{a}_l^T \bar{n}_i) \quad (12)$$

$$B_i \simeq \sum_{l=1}^L (e_B^l s_B / \sigma_B) (\bar{a}_i^T \bar{n}_i) \quad (13)$$

To develop a more compact representation for these formulas, we now introduce two diagonal matrices, S_i and E_i :

$$S_i = \begin{bmatrix} s_R & 0 & 0 \\ 0 & s_G & 0 \\ 0 & 0 & s_B \end{bmatrix} \quad E_i = \begin{bmatrix} E_R & 0 & 0 \\ 0 & E_G & 0 \\ 0 & 0 & E_B \end{bmatrix}$$

where S_i is a diagonal matrix composed of the surface color components s_R , s_G , and s_B at a point in the scene that corresponds to the pixel i , and E_i (called the light matrix), again a diagonal matrix, consists of the following diagonal entries:

$$E_R = \sum_{l=1}^L (e_R^l \bar{a}_i^T \bar{n}_i / \sigma_R) \quad (14)$$

$$E_G = \sum_{l=1}^L (e_G^l \bar{a}_i^T \bar{n}_i / \sigma_G) \quad (15)$$

$$E_B = \sum_{l=1}^L (e_B^l \bar{a}_i^T \bar{n}_i / \sigma_B). \quad (16)$$

Note that each of E_c is a scalar (here the subscript c stands for R , G , or B) and is the same for all pixels under the same illumination condition.

With the help of the matrix notation introduced as above, the three equations (Eq. 11, 12, and 13) can be expressed in the following compact form:

$$\omega_i = S_i E_i \quad (17)$$

by denoting $\omega_i = [R_i, G_i, B_i]^T$ where R_i , G_i and B_i are red, green and blue values of the pixel i respectively.

Since our main goal is to develop a mapping function that tells how the color values in an image taken under non-specular conditions map into the color values taken under specular conditions, we will now examine the above equation assuming that the illumination has changed from E to E' . If the new RGB values at pixel i are denoted as ω'_i , we can write

$$\omega'_i = S_i E'_i = S_i ((E'_i E_i^{-1}) E_i) \quad (18)$$

Substituting Eq. 17, we get

$$\omega'_i = S_i ((E'_i E_i^{-1})) S_i^{-1} \omega_i = S_i \cdot M_i \cdot S_i^{-1} \cdot \omega_i \quad (19)$$

where the 3×3 matrix M_i is given by

$$M_i = E'_i E_i^{-1}.$$

Here notice that matrix E_i is a diagonal matrix, therefore it is invertible except when one or more diagonal entries is 0. And any of these diagonal entries (given in Eqs.

14 through 16) becomes 0 only when the following case becomes true: For all L number of light sources, all $e_c^1, e_c^2, \dots, e_c^L$ are simultaneously 0 (where $c = R, G, \text{ or } B$).

Matrix S_i is also a diagonal matrix, therefore it is always invertible except when one or more diagonal entries (that are RGB color values of surface) becomes 0. That will happen when a scene point consists of an object surface whose color is deeply saturated. For example, when the color of an object surface is pure red, we can expect s_G and s_B to be zero. The matrix S_i will become poorly conditioned when an object surface is deeply black so as to absorb most of the incident illumination. For such surfaces, all three values on the diagonal of S will be close to zero.

Despite these caveats regarding the invertability of E_i and S_i , we will press ahead under the assumption that both these matrices can be inverted. In other words, we will assume that the scene does not contain colors that are deeply saturated or deeply black.

Matrix M_i being diagonal implies that it approximately commutes with S_i so that S_i and S_i^{-1} can be canceled out. Thus the linear equation for modeling illumination changes becomes even simpler:

$$\omega'_\chi \simeq M_i \cdot \omega_\chi. \quad (20)$$

The development so far has considered each pixel individually. The mapping function M_i is to be applied to each pixel separately. We will now generalize this notion and develop the notion of a mapping function that can be applied to all non-saturated pixels in an image. While the individual pixel mapping function was denoted M_i , the mapping to be applied to all the "legal" pixels will be denoted M . The mapping M will be made to be optimal in the least-squares sense for all legal pixels, meaning pixels that are not saturated in the target image.

To derive this more global mapping, we will define an auto-correlation matrix:

$$C \equiv \sum_{i=1}^k \omega_i \cdot \omega_i^T \quad (21)$$

where i ranges from 1 to k , with k being the total number of pixels remaining in the target image after the saturated pixels are removed. Here we only considered k number of pixels from the image whereas all pixels from the entire image were used in Drew's approach[8].

If we put all k pixels into a $3 \times k$ matrix, denoted Ω , where each column now represents the R, G, and B values at a pixels, then Eq. 21 becomes

$$C \equiv \Omega \cdot \Omega^T \quad (22)$$

where Ω can be written as

$$\Omega = \begin{pmatrix} \omega_1 & \omega_2 & \cdots & \omega_k \end{pmatrix} = \begin{pmatrix} R_1 & R_2 & \cdots & R_k \\ G_1 & G_2 & \cdots & G_k \\ B_1 & B_2 & \cdots & B_k \end{pmatrix} \quad (23)$$

From Eq. 20 and Eq. 22, we can easily retrieve the equation for transformation of the auto-correlation

$$C' = M \cdot C \cdot M^T \quad (24)$$

The only unknown we have now is the 3×3 matrix M that is the best truncated least-squares solution calculated from $\min \|\Omega' - M \cdot \Omega\|^2$. Now we can derive M from the normal equation with the pseudo-inverse

$$M = (\Omega' \cdot \Omega^T)(\Omega \cdot \Omega^T)^{-1} \quad (25)$$

Remaining is just to apply M in Eq. 25 to the two images, Ω and Ω' :

$$\Omega' \simeq M \cdot \Omega \quad (26)$$

As long as M is non-singular, M^{-1} can be obtained. Thus we can also approximately transform the color values in the reference image to the color values in the target image by multiplication with M^{-1} :

$$\Omega \simeq M^{-1} \cdot \Omega' \quad (27)$$

C. Choose Threshold for Detecting Non-saturated Specularities

The basic idea that we use for establishing a threshold for the detection of non-saturated specularities is based on the notion that when a color value is inverse mapped from the target image to the reference image (using the mapping function of the previous subsection), the co-domain point for a specular pixel will fall outside the normal range. That is because the mapping function M applies only to non-specular pixels. In other words, M specifically excludes all specularities – both saturated and non-saturated. (But, of course, we have already eliminated the saturated specularities by the very first step of the processing outlined at the beginning of Section III.) It should therefore be possible to detect the non-saturated specular pixels in the target image on the basis of inverse mapping of their color values. This is the main idea behind the expression we present below for the threshold.

For each RGB channel, we calculate the mean vector $\mu(\Omega)$ and a vector representing the standard deviation $\sigma(\Omega)$ for those pixels of the reference image that form the co-domain of the retained pixels in the target image. In terms of the mean and the standard deviation thus derived, we use the following expression for establishing a decision threshold for detecting non-saturated specular pixels in the target image:

$$U_k = \mu_k(\Omega) + \alpha \sigma_k(\Omega) \quad (28)$$

where k indicates one of R, G, and B and α is a constant factor. A target pixel whose inverse-mapped color values exceed this threshold is declared to be a non-saturated specular pixel. Currently, we set $\alpha = 1.125$ in Eq. 28, which makes the threshold 1.125 times of the standard deviation from the mean. However, we would obviously want this decision threshold to correspond to the human perception of specularity. Toward that end, a further study needs to be carried out to understand how to select superior thresholds using Eq. 28.

IV. EXPERIMENTAL RESULTS

We selected 100 pairs of images from our database, with one image in a pair serving as a reference image and the other image as a target image. The reference image was taken under diffuse illumination and the target image under either directed or under ambient illumination. Since the goal of the research reported here is the detection of specular pixels in an image, it is necessary to have ground-truth information for evaluating our method. The ground truth was obtained by employing an unbiased human examiner and providing him with a graphical tool for delineating the specularities in our database images. The human-entered information was stored as a binary template for each image that contained specularities. A template stored 1's where the human perceived a specularity and 0's elsewhere. The middle column in Figure 4 depicts these templates for the target images in the left column. The right column shows those regions in each image that were declared to be specular by our method. For the images on which we tested our method, 64.5% of the human-delineated specular pixels were detected by our method. Our method also declared 14.6% of the non-specular pixels as specular.

V. DISCUSSION

In this paper we have presented a method for establishing thresholds for the detection of specularities in images. The method first eliminates the saturated specularities and then finds a least-squares mapping from the color values in a reference image to those in the target image. The reference image is expected to be substantially free of specularities, a condition that can be satisfied when images are specially recorded under diffuse illumination. The decision threshold for the saturated specular pixels is obtained by an iterative process that involves eliminating the target image pixels whose color values exceed a candidate threshold and carrying out a Principal Component Analysis of the remaining pixels. The decision threshold for the non-saturated specular pixels is obtained on the basis of inverse mapping of the target image color values to the reference image color values. A comparative study involving this present method and the previously proposed specular detection methods is underway in our laboratory.

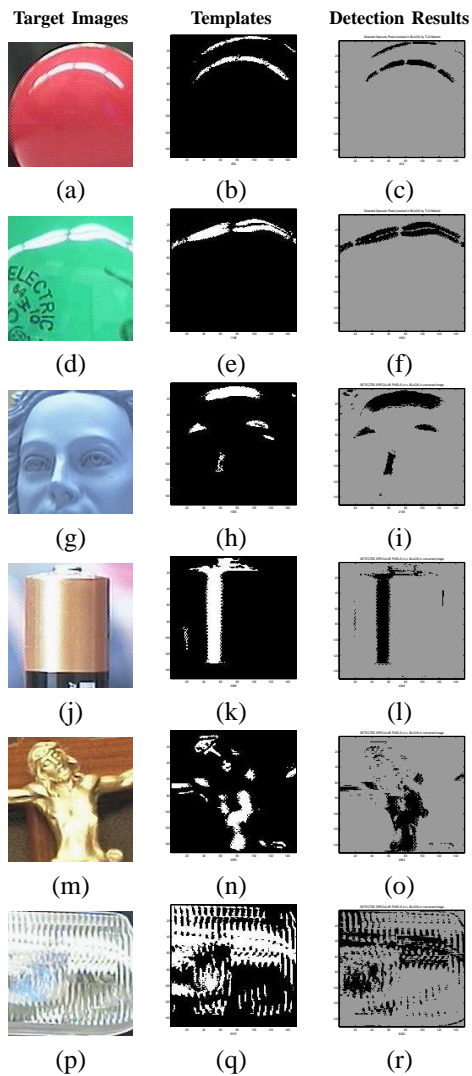


Fig. 4. Example images, templates and results

VI. REFERENCES

- [1] L. T. Maloney, B. A. Wandell, Color Constancy: A Method for Recovering Surface Spectral Reflectance, *Journal of the Optical Society of America*, Vol. A3, pp. 29-33, 1986.
- [2] G. Brelstaff, A. Blake. Detecting Specular Reflections using Lambertian Constraints, *In Proceedings of ICCV*, pp. 297-302, 1988.
- [3] Hsien-Che Lee. Method for computing the scene-illuminant chromaticity from specular highlights, *Journal of the Optical Society of America*, A 3, 1694-1699, 1986.
- [4] G. J. Klinker, S. A. Shafer, T. Kanade, The measurement of highlights on color images. *Int'l Journal Computer Vision*, Vol 2, pp. 7-32, 1988.
- [5] R. Bajcsy, Sang Wook Lee, A. Leonardis, Color

image segmentation with detection of highlights and local illumination induced by interreflections, *IEEE Int'l Conference on Pattern Recognition*, Atlantic City, Vol. 1, pp. 785-790, 1990.

- [6] G. J. Klinker, A Physical Approach to Color Image Understanding, AK Peters, LTD., Wellesley MA, March, 1993.
- [7] S. W. Lee, R. Bajcsy, Detection of Specularity using Color and Multiple Views, *Proceedings of the 2nd European Conference on Computer Vision*, pp. 99-114, Santa Margherita Ligure, Italy, May, 1992.
- [8] Mark S. Drew. Robust Specularity Detection from a Single Multi-illuminant Color Image. *CVGIP(59)*, No 3, pp. 320-327, May 1994.
- [9] G. Wyszecki and W. S. Stiles. Color Science: Concepts and Methods, Quantitative Data and Formulae, 2nd Edition, John Wiley & Sons, New York, 1982.
- [10] C. F. Borges, Trichromatic Approximation method for surface illumination, *Journal Opt. Society of America*, Vol. A: 8, pp. 1319-1323, 1991.
- [11] W. B. Cowan, C. Ware, Tutorial on color perception, In SIGGRAPH, July, 1983.