University of Pennsylvania

ScholarlyCommons

Technical Reports (CIS)   Department of Computer & Information Science

January 1999

Spectral Gradient: A Surface Reflectance Measurement Invariant to Geometry and Incident Illumination

Elli Angelopoulou
University of Pennsylvania

Sang W. Lee
University of Michigan

Ruzena Bajcsy
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/cis_reports

Spectral Gradient: A Surface Reflectance Measurement Invariant to Geometry and Incident Illumination

Abstract

Although photometric data is a readily available dense source of information in intensity images, it is not widely used in computer vision. A major drawback is its dependence on viewpoint and incident illumination. A novel methodology is presented which extracts reflectivity information of the various materials in the scene independent of incident light and scene geometry. A scene is captured under three different narrow-band color filters and the spectral derivatives of the scene are computed. The resulting spectral derivatives form a spectral gradient at each pixel. This spectral gradient is a surface reflectance descriptor which is invariant to scene geometry and incident illumination for smooth diffuse surfaces. The invariant properties of the spectral gradients make them a particularly appealing tool in many diverse areas of computer vision such as color constancy, tracking, scene classification, material classification, stereo correspondence, even re-illumination of a scene.

Comments

Spectral Gradient: A Surface Reflectance Measurement Invariant to Geometry and Incident Illumination

Technical Report MS-CIS-99-02

Elli Angelopoulou
Department of Computer and
Information Science
University of Pennsylvania
elli@grip.cis.upenn.edu

Sang W. Lee
Department of Electrical Engi-
neering and Computer Science
University of Michigan
swlee@eecs.umich.edu

Ruzena Bajcsy
Department of Computer and
Information Science
University of Pennsylvania
rbajcsy@grip.cis.upenn.edu

January 18,1999

Abstract

Although photometric data is a readily available dense source of information in intensity images, it is not widely used in computer vision. A major drawback is its dependence on viewpoint and incident illumination. A novel methodology is presented which extracts reflectivity information of the various materials in the scene independent of incident light and scene geometry. A scene is captured under three different narrow-band color filters and the spectral derivatives of the scene are computed. The resulting spectral derivatives form a spectral gradient at each pixel. This spectral gradient is a surface reflectance descriptor which is invariant to scene geometry and incident illumination for smooth diffuse surfaces. The invariant properties of the spectral gradients make them a particularly appealing tool in many diverse areas of computer vision such as color constancy, tracking, scene classification, material classification, stereo correspondence, even re-illumination of a scene.

1 Introduction

The starting point of most computer vision techniques is the light intensity reflected from an imaged scene. The reflected light is directly related to the geometry of the scene, the reflectance properties of the materials in the scene and the lighting conditions under which the scene was captured. One of the complications which have troubled computer vision algorithms is the variability of an object's appearance as illumination and scene geometry change. Slight variations in viewing conditions often cause large changes in an object's appearance. Consider, for example a yellow car seen in a sunny day, at night, or in dense fog.

Many areas of computer vision are affected by this problem. Maloney and Wandell[19] were the first to

develop a tractable color constancy algorithm. Color constancy is the problem of recognizing colors despite changes in ambient lighting conditions. They modeled both the surface reflectance and the incident illumination as a finite dimensional linear model. This basic principle was further explored in the color constancy work of Forsyth[8], Ho et al.[13], Finlayson et al.[7, 9, 6, 1] and Healey and Slater[11]. At the same time, Swain and Ballard[29] showed that it was possible to correctly identify objects by using color cues only. Healey and Slater[11, 27], Funt and Finlayson[9] and Finlayson[6] made object recognition by color even more widely applicable, by cancelling variations in color appearance caused by illumination changes.

In object recognition Nayar and Bolle[22], Slater and Healey[25, 26], Lin and Lee[18] and Jacobs et al.[16], among others concentrated in identifying reflectance-based object properties that are invariant to illumination. Another school of thought in object recognition, influenced by the work of Turk and Pentland[30], developed appearance-based models which embed the variability of the imaging conditions to the model of the object[21, 24, 2, 3, 5]. In texture recognition, Healey and Wang[12] developed an illumination invariant distance function for comparing color textures. In real-time tracking Hager and Belheumer[10] adapted the sum of squared differences (SSD) algorithm to handle variations in illumination.

All these systems in order to handle the variations in viewing conditions had to introduce some additional constraints that are often limiting their applicability. For example, most color techniques assume that the spectral reflectance functions have the same degrees of freedom as the number of photoreceptor classes (typically three.) Thus, none of these methods can be used in greyscale images for extracting illumination invariant color information. Furthermore, a considerable body of work on color assumes that the incident illumination has two or three degrees of freedom. However, Slater and Healey[28] showed that for outdoor scenes, the illumination functions have seven degrees of freedom. On the other hand, greyscale object recognition methods[22, 16] take advantage of the invariance in the color distribution on an object and can not handle very well non-textured scenes. As for appearance based approaches, as Mundy et al.[20] pointed out, they do not render themselves in generalizations of identifying objects or materials that should logically belong to the same class but appear different.

We propose a novel method for cancelling variations in geometry and incident illumination by examining the rate of change in reflected intensity with respect to wavelength. The only assumption that we make is that incident illumination remains stable over small intervals in the visible spectrum. It will be demonstrated that this is a reasonable assumption. We take a greyscale image of a scene under three different narrow-band color filters and compute the spectral derivatives of the scene. The resulting spectral gradient is a surface reflectance descriptor, invariant to scene geometry and incident illumination for smooth diffuse surfaces. Experiments on surfaces of different colors and materials demonstrate the accuracy of our method in both: a) identifying materials with the same reflectance under variable viewing conditions and b) discriminating materials with distinct reflectance functions.

2 Spectral Derivative

The intensity images that we process in computer vision are formed when light from a scene falls on a photosensitive sensor. The amount of light reflected from each point $p = (x, y, z)$ in the scene depends on the light illuminating the scene, E and the surface reflectance S of the materials composing the scene:

$$I(p, \lambda) = E(p, \lambda)S(p, \lambda)$$

where λ, the wavelength, shows the dependency of reflected light on wavelength. The reflectance function $S(p, \lambda)$ may depend on the surface material, the geometry of the scene and the viewing and incidence angles.

When the spectral distribution of the incident light does not vary with the position of the light, the geometric and spectral components of the incident illumination are separable:

$$E(\theta_i, \varphi_i, \lambda) = e(\lambda)E(\theta_i, \varphi_i)$$

where (θ_i, φ_i) are the spherical coordinates of the unit-length light-direction vector and $e(\lambda)$ is the illumination spectrum. Note that, the incident light intensity is included in $E(\theta_i, \varphi_i)$ and may vary as the position of the illumination source changes.

The scene brightness then becomes:

$$I(p, \lambda) = e(p, \lambda)E(p, \theta_i, \varphi_i)S(p, \lambda) \qquad (1)$$

By taking the logarithm of the image irradiance equation we alter the multiplicative effect into an additive effect:

$$\mathcal{L}(p, \lambda) = \ln e(p, \lambda) + \ln E(p, \theta_i, \varphi_i) + \ln S(p, \lambda) \quad (2)$$

We are interested in investigating the behavior of the natural logarithm of an image as we vary the wavelength in the visible range, i.e., 400nm to 700nm. Thus, we compute the partial derivative of the logarithmic image with respect to wavelength λ:

$$\mathcal{L}_\lambda(p, \lambda) = \frac{e_\lambda(p, \lambda)}{e(p, \lambda)} + \frac{S_\lambda(p, \lambda)}{S(p, \lambda)} \qquad (3)$$

where $e_\lambda(p, \lambda) = \partial e(p, \lambda)/\partial \lambda$ is the partial derivative of the incident illumination with respect to wavelength and $S_\lambda(p, \lambda) = \partial S(p, \lambda)/\partial \lambda$ is the partial derivative of the surface reflectance with respect to wavelength. Ho, Funt and Drew[13] have shown, that for natural objects the surface spectral reflectance curves, i.e. the plots of $S(p, \lambda)$ versus λ, are usually reasonably smooth and continuous over the visible spectrum, 400nm to 700nm.

3 Invariance to Incident Illumination

Although the spectral distribution of the most commonly used indoor-scene illuminations sources (i.e., tungsten and fluorescent light) is not constant, one can assume that e changes very slowly over small increments of λ. This means that its derivative with respect to wavelength is approximately zero.

$$e_\lambda(\boldsymbol{p}, \lambda) \approx 0$$

Thus, the partial derivative of the logarithmic image depends only on the surface reflectance:

$$\mathcal{L}_\lambda(\boldsymbol{p}, \lambda) \approx \frac{S_\lambda(\boldsymbol{p}, \lambda)}{S(\boldsymbol{p}, \lambda)} \qquad (4)$$

4 Invariance to Geometry and Viewpoint

4.1 Lambertian Model

A very simple model that is often used by both the computer vision community and the graphics community is the Lambertian reflectance model. Lambert's law described the behavior of a perfectly diffuse surface, where the reflected light is independent of viewpoint. For a homogeneous surface, the reflected light changes only when the angle of incidence $\theta(\boldsymbol{p})$ between surface normal and the incident illumination changes.

$$S(\boldsymbol{p}, \lambda) = \cos\theta(\boldsymbol{p})\rho(\boldsymbol{p}, \lambda)$$

where $\rho(\boldsymbol{p}, \lambda)$ is the albedo or diffuse reflection coefficient at point $\boldsymbol{p}$.

Since, by definition, Lambertian reflectance is independent of viewpoint, the spectral gradient is also independent of viewpoint. The scene geometry is independent of wavelength. Therefore, when we take the partial derivative with respect to wavelength, the geometry term vanishes. Thus, for Lambertian surfaces:

$$\mathcal{L}_\lambda(\boldsymbol{p}, \lambda) \approx \frac{S_\lambda(\boldsymbol{p}, \lambda)}{S(\boldsymbol{p}, \lambda)} = \frac{\rho_\lambda(\boldsymbol{p}, \lambda)}{\rho(\boldsymbol{p}, \lambda)} \qquad (5)$$

where $\rho_\lambda(\boldsymbol{p}, \lambda) = \partial\rho(\boldsymbol{p}, \lambda)/\partial\lambda$ is the partial derivative of the surface albedo with respect to wavelength.

4.2 Smooth Diffuse Reflectance Model

In reality there are very few objects that exhibit perfectly Lambertian reflectance. The light that is reflected from a smooth diffuse object varies with respect to viewpoint. Wolff[31] introduced a new smooth diffuse reflectance model that incorporates the dependence on viewpoint:

$$S(\boldsymbol{p}, \lambda) = \cos\theta(\boldsymbol{p})\rho(\boldsymbol{p}, \lambda)(1 - F(\theta(\boldsymbol{p}), n(\boldsymbol{p})))\left(\left(1 - F\left(\sin^{-1}\left(\frac{\sin\varphi(\boldsymbol{p})}{n(\boldsymbol{p})}\right), \frac{1}{n(\boldsymbol{p})}\right)\right)\right)$$

where $\theta(\boldsymbol{p})$ and $\varphi(\boldsymbol{p})$ are the incidence and viewing angles respectively, $\rho(\boldsymbol{p}, \lambda)$ is the surface albedo, $F()$ is the Fresnel reflection coefficient, and n is the index of refraction.

By taking the logarithm of the surface reflectance function, we simplify the underlying model, by altering multiplicative terms into additive terms:

$$\ln S(\boldsymbol{p}, \lambda) = \ln \cos \theta(\boldsymbol{p}) + \ln \rho(\boldsymbol{p}, \lambda) + \ln(1 - F(\theta(\boldsymbol{p}), n(\boldsymbol{p}))) + \ln\left(1 - F\left(\sin^{-1}\left(\frac{\sin \varphi(\boldsymbol{p})}{n(\boldsymbol{p})}\right), \frac{1}{n(\boldsymbol{p})}\right)\right)$$

The next step is to compute the partial derivative with respect to wavelength. Once again, all the terms except the albedo are set to zero:

$$\mathcal{L}_\lambda(\boldsymbol{p}, \lambda) \approx \frac{S_\lambda(\boldsymbol{p}, \lambda)}{S(\boldsymbol{p}, \lambda)} = \frac{\rho_\lambda(\boldsymbol{p}, \lambda)}{\rho(\boldsymbol{p}, \lambda)} \tag{6}$$

The index of refraction depends theoretically on wavelength, but in the visible range it changes by a very small amount. Thus, it is commonly treated as a material constant under visible light.

5 Spectral Gradient

For smooth diffuse surfaces the partial derivative with respect to wavelength of the logarithmic image $\mathcal{L}_\lambda(\boldsymbol{p}, \lambda)$ is a function of only the surface albedo. Consider now a collection of spectral derivatives of a logarithmic image at various spectral locations λ_k, $k = 1, 2, 3, ..., M$. The resulting *spectral gradient* is an M-dimensional vector $(\mathcal{L}_{\lambda_1}, \mathcal{L}_{\lambda_2}, ..., \mathcal{L}_{\lambda_M})$ which is invariant to illumination, surface geometry and viewpoint. All it encodes is information at discrete spectral locations about how fast the surface albedo changes as the spectrum changes. It is a profile of the rate of change of albedo with respect to wavelength over a range of wavelengths.

6 Experiments

In order to compute the spectral derivatives we took images of each scene under three different narrow-band filters: a Corion S10-570-F, a Corion S10-600-F and a Corion S10-630-F. Each of these filters has a bandwidth of approximately 10nm and a transmittance of about 50%. The central wavelengths are at 570nm, 600nm and 630nm respectively. The images were captured with a Sony XC-77 camera using a 25mm lens.

The only source of illumination was a single tungsten light bulb mounted in a reflected scoop. For each scene we used four different illumination setups, generated by the combination of two distinct light bulbs, a 100W bulb and a 200W bulb and two different light positions. One illumination position was to the left

of the camera and at about the same height as the camera. Its direction vector formed approximately a 20°
angle with the optic axis. The other light-bulb position was to the right of the camera and about 10" above
it. Its direction vector formed roughly a 55° angle with the optic axis. Both locations were 22" away from
the scene.

The imaged objects were positioned 30" from the camera/filter setup. We tried four different types of
materials: foam, paper, ceramic and plastic. Foam and paper came in a variety of colors. In foam we had
green, magenta, orange, pink, red, white and yellow samples. Our pieces of paper came in brown, green,
orange, pink, red white and yellow. We used two ceramic objects, a pink plate and a white mug, and one
rectangular piece of plastic.

Fig. 1. A small sample of the colors, materials and
shapes that we experimented with.

Fig. 1. shows in the top row and from left to right small samples of green foam, yellow foam and red
paper. On the bottom row are images of the ceramic and plastic objects. These images were taken using the
Corion S10-600-F filter.

6.1 Computing the Spectral Gradient

Once a filtered image was captured, its logarithmic image was generated. In a logarithmic image the
value stored at each pixel was the natural logarithm of the original image intensity. For example,
$\mathcal{L}_{570} = \ln(I_{570})$, where I_{570} was the image of a scene taken with the S10-570-F filter and $\mathcal{L}_{570}$ was its log-
arithmic image.

As fig. 2 shows, the logarithmic images preserved the overall appearance of the original image. How-
ever, the intensity values were scaled down significantly. From a maximum of 255 in an 8-bit image we
went down to a maximum of 5.54. The images shown in fig. 2 were linearly scaled for display purposes.

Fig. 2. Sample logarithmic images.

The last step was the computation of the spectral derivatives of the logarithmic images. Differentiation was approximated via finite-differencing. Thus, $\mathcal{L}_\lambda$ was computed over the wavelength interval $\delta\lambda = 30nm$ by subtracting two logarithmic images taken under two different color filters which were 30nm apart:

$$\mathcal{L}_{\lambda_1} = \mathcal{L}_{600} - \mathcal{L}_{570} \qquad \mathcal{L}_{\lambda_2} = \mathcal{L}_{630} - \mathcal{L}_{600} \qquad (7)$$

Typically, the resulting derivative images have a median value around 0.2 with minimal variation on smooth materials like glossy ceramic and plastic. The best way to depict this minimal variation was to show the reverse video of the derivative image. Again, this was done for display purposes only. Fig. 3 shows the derivative images of a piece of green foam and the ceramic and plastic objects.

Fig. 3. Sample derivative images.

For each scene taken under the three narrow-band color filters we had two derivative images, $\mathcal{L}_{\lambda_1}$ and $\mathcal{L}_{\lambda_2}$. The spectral gradient at a pixel was the vector $(\mathcal{L}_{\lambda_1}, \mathcal{L}_{\lambda_2})$. This vector was expected to remain constant for materials with the same reflectance function, independent of variations in viewing conditions. At the same time, it should differ significantly for materials that exhibit distinct surface reflectance functions.

6.2 Comparing Spectral Gradients

The desired goal was to determine whether two regions (in the same or different scenes) are depicting objects with the similar or distinct reflectance functions. We performed a pixel by pixel comparison. Let p and p' be two pixels belonging to these two regions and let $(\mathcal{L}_{\lambda_1}, \mathcal{L}_{\lambda_2})$ and $(\mathcal{L}_{\lambda_1}', \mathcal{L}_{\lambda_2}')$ be their respective spectral gradients. The metric we used was the absolute difference vector, after it was normalized for variations in intensity level:

$$(d_1, d_2) = (|\mathcal{L}_{\lambda_1} - \mathcal{L}_{\lambda_1}'|, |\mathcal{L}_{\lambda_2} - \mathcal{L}_{\lambda_2}'|)/I_{avg} \qquad (8)$$

where $I_{avg} = (i_{570} + i_{600} + i_{630} + i'_{570} + i'_{600} + i'_{630})/6$ was the average value of the intensities registered in these two pixels in the original filtered images.

The pixel metric of equation (8) was the basis for comparing regions. The (d_1, d_2) metric was computed for all the corresponding pixels in the two regions, i.e. pixels which had the same coordinates (for instance, both in position x_0, y_0) in the local coordinate system of each region. The median of all the pixel measurements became the region distance metric:

$$(D_1, D_2) = (median \forall d_1, median \forall d_2) \qquad (9)$$

Typically, the values of D_1 and D_2 were very small, mainly because the spectral gradients themselves were small. Hence, subtraction and normalization of such derivative values resulted in values for D_1 and D_2 that ranged in our experiments from 0.001 to 3.898.

6.2.1 Same Material and Color

According to our theory, materials with the same reflectance function should generate the same spectral gradient resulting in a (D_1, D_2) tuple that is almost equal zero. Indeed, we observed that when we were comparing the same material and the same color, independent of illumination conditions and surface orientation, the average of D_1 and D_2 was consistently less than 0.2. More precisely, out of 28 such comparisons, 27 times both $D_1 < 0.2$ and $D_2 < 0.2$. There was a single case, when we compared two regions of orange foam illuminated under two different light intensities (same light position), that $D_1 = 0.2564$ and $D_1 = 0.0057$, but still their average was less than 0.2.

We compared different regions of the ceramic mug, with distinct viewing and incidence angles. The resulting spectral gradients differed by less than 0.1. A similar behavior was observed around the smooth corners of the plastic container. The specular region on the white mug did not affect the stability of spectral gradients. For example, comparing the shiny region in the center of the mug with a region in the right side of the mug generated the following tuple $(D_1, D_2) = (0.0288, 0.0402)$. In general, same color and same material comparisons generated very stable spectral gradients. Some sample comparisons can be found in Table 1.

Table 1: Same Material and Color

Material	(D_1, D_2)
magenta foam	(0.0126, 0.0333)
yellow foam	(0.0010, 0.0036)
red paper	(0.1472, 0.0209)
green paper	(0.0788, 0.0519)
white mug	(0.0262, 0.0558)

6.2.2 Same Material but Different Colors

Surfaces made out of different colors of the same material follow the same reflectance model, but have distinct spectral responses. Thus, their spectral gradients should be distinguishable, at least over some part of the visible spectrum.

We performed 35 comparisons between distinct colors of the same material under similar or different illumination conditions. The spectral gradients were distinguishable 34 out of the 35 times. The average of both D_1 and D_2 was larger than 0.2 in these cases. Examples of (D_1, D_2) for different colors of the same material are shown in Table 2.

Table 2: Same Material but Different Colors

Color 1	Color 2	(D_1, D_2)
magenta foam	red foam	(0.7976, 1.0465)
green foam	orange foam	(0.6755, 0.2663)
red paper	yellow paper	(0.1703, 1.1402)
green paper	orange paper	(0.3806, 1.075)
white ceramic	pink ceramic	(0.2759, 0.5555)

The two times that the spectral gradients were very similar, were when we compared a piece of pink foam with a piece of magenta foam under 200W illumination. These two materials have very similar reflectance spectra that vary the most around 550nm (see fig. 4). The spectral gradients that we used didn't cover that part of the spectrum and the difference between the two materials went undetected.
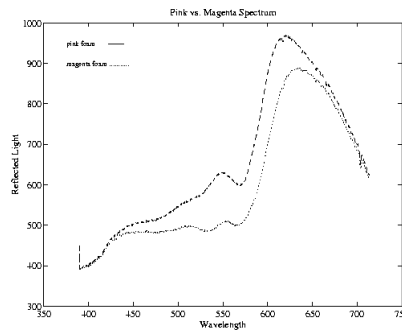

Fig. 4. The spectrum of pink foam versus magenta foam.

6.2.3 Different Materials but Same Color

Spectral gradients are a measurement of the surface reflectance function, independent of viewing conditions. As such, if two distinct materials have similar reflectance, the respective spectral gradients would be similar too. Out of the four materials we tested, the foam and the paper, had very similar reflectance behavior. The ceramic and the plastic samples were also very similar in terns of reflectance behavior. Our paper

was a bit smoother than foam resulting in a reflectance function that was close to the that of the ceramic reflectance. There was also a hue discrepancy between the red paper and the red foam, as well as between the green paper and the green foam.

Overall, when the materials exhibited clearly distinct reflectance functions, the average of both D_1 and D_2 was larger than 0.2. Look for example at pink ceramic versus pink foam in Table 3. Differentiating between foam and paper of very similar color was very difficult. Both materials exhibit an approximately Lambertian reflectance. It was interesting to note that small color variations, like different shades of dark green, were distinguishable across different materials of the same reflectance behavior.

Table 3: Different Materials but Same Color

Color 1	Color 2	(D_1, D_2)
green foam	green paper	(0.2365, 0.6161)
orange foam	orange paper	(0.0120, 0.0765)
pink foam	pink paper	(0.0767, 0.2968)
pink ceramic	pink paper	(0.1359, 0.1243)
pink ceramic	pink foam	(0.0331, 0.4527)

6.3 Error Analysis

Our experimentations showed that spectral gradients achieved a 100% correct identification when it was comparing the same material, independent of the variations in illumination conditions. The success of spectral gradients in discriminating between different colors of the same material was good, about 97%, but it clearly depended on the wavelengths at which we were sampling the partial derivatives. Higher dimensional spectral gradients should provide better discriminatory power. Finally, spectral gradients can discriminate between different materials of the same color, only if their surface reflectance behavior differs.

Note that we have tried a variety of filters with different bandwidths. Since we were approximating a sampling function (Dirac delta function), the narrower filters gave more consistent results. We also experimented with various $\delta\lambda$ over which to perform the finite differencing approximation to partial derivatives. Again, as expected, the smaller $\delta\lambda$ performed better, as long as the two filters did not have overlapping bandwidths.

7 Conclusions

We developed a surface reflectance measurement that is invariant to changes in illumination and scene geometry. We made no assumptions about the nature of incident light, other than that its spectrum does not

change with its position. We showed that spectral gradients can be used on a pixel basis and do not depend on neighboring regions, an assumption that is common in other photometric methods which use logarithms and/or narrow-band filters[9, 22]. The effectiveness of spectral gradients as a surface reflectance descriptor was demonstrated on various empirical data.

The invariant properties of spectral gradients together with their ease of implementation and the minimalism of assumptions, make this methodology a particularly appealing tool in many diverse areas of computer vision. They can be used in grey-scale color constancy, or in tracking different regions under variable illumination. They can also be an additional feature used in stereo correspondence.

We believe that spectral gradients are a powerful tool that should be further investigated. Additional experimentation under different types and colors of light sources is under way. Simultaneous use of multiple lights is another issue that is being examined. It is also very important to study more extensively the behavior of spectral gradients in areas with specular highlights. Finally, we would like to study the behavior of spectral gradient on rough surfaces.

References

[1] Barnard, K., Finlayson, G. and Funt, B. "Color Constancy for Scenes with Varying Illumination," *Computer Vision and Image Understanding*, Vol. 65, No. 2, February 1997, pp. 311-321.

[2] Belhumeur, P. N., Hespanha, J. P. and Kriegman, D. J., "Eigenfaces vs. Fisherfaces: Recognition Using Class-Specific Linear Projection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 19, No. 7, July 1997, pp. 711-720.

[3] Belheumer, P. N. and Kriegman, D. J. "What Is the Set of Images of an Object under all Possible Illumination Conditions?," *International Journal of Computer Vision*, Vol. 28, No. 3, 1998, pp. 245-260.

[4] Berwick, D. and Lee, S. W., "A Chromaticity Space for Specularity, Illumination Color- and Illumination Pose-Invariant 3-D Object Recognition," *Proceedings International Conference on Computer Vision*, 1998, pp. 165-170.

[5] Camps, O. I., Huang, C. Y. and Kanungo, T., "Hierarchical Organization of Appearance Based Parts and Relations for Object Recognition," *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 1998, pp. 685-691.

[6] Finlayson, G. D., "Color in Perspective," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 18, No. 10, October 1996, pp. 1034-1038.

[7] Finlayson, G. D., Drew, M. S. and Funt, B. V., "Color Constancy: Generalized Diagonal Transforms Suffice," *Journal of the Optical Society of America A*, Vol. 11, 1994, pp. 3011-3019.

[8] Forsyth, D., "A Novel Algorithm for Color Constancy," *International Journal of Computer Vision*, Vol.

5, No.1, 1990, pp. 5-36.

[9] Funt, B.V., and Finlayson, G. D, "Color Constant Color Indexing," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 17, 1995, pp. 522-529.

[10] Hager, G. D. and Belhumeur, P. N., "Real-Time Tracking Of Image Regions With Changes In Geometry And Illumination," *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 1996, pp. 403-410.

[11] Healey, G. and Slater, D. "Global Color Constancy: Recognition of Objects by Use of Illumination-Invariant Properties of Color Distribution," *Journal of the Optical Society of America A*, Vol. 11, No. 11, November 1994, pp. 3003-3010.

[12] Healey, G. and Wang, L. "Illumination-Invariant Recognition of Texture in Color Images," *Journal of the Optical Society of America A*, Vol. 12, No. 9, September 1995, pp. 1877-1883.

[13] Ho, J., Funt, B. V. and Drew, M. S., "Separating A Color Signal Into Illumination And Surface Reflectance Components: Theory And Applications," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 12, No. 10, October 1990, pp. 966-977.

[14] Horn, B. K. P., "Understanding Image Intensities," *Artificial Intelligence,* Vol. 8, No. 2, 1977, pp. 1-31.

[15] Horn, B. K. P. and Brooks, M. J., *Shape from Shading*, MIT Press, 1989.

[16] Jacobs, D. W., Belheumer, P. N. and Basri, R. "Comparing Images Under Variable Illumination," *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 1998, pp. 610-617.

[17] Land, E. H. and McCann, J. J., "Lightness and the Retinex Theory," *Journal of the Optical Society of America*, Vol. 61, No. 1, January 1971, pp. 1-11.

[18] Lin, S. and Lee, S. W., "Using Chromaticity Distributions and Eigenspaces for Pose-, Illumination-, and Specularity Invariant 3D Object Recognition" *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 1997, pp. 426-431.

[19] Maloney, L. T. and Wandell, B. A. "A Computational Model of Color Constancy," *Journal of the Optical Society of America A*, Vol. 3. No. 1, 1986, pp. 29-33.

[20] Mundy, J. L., Liu, A., Pillow, N., Zisserman, A., Abdallah, S., Utcke, S., Nayar, S. and Rothwell, C.A., "An Experimental Comparison of Appearance and Geometric Model Based Recognition," *Proceedings of the International Workshop on Object Representation in Computer Vision II*, Springer-Verlag 1996, pp. 247-270.

[21] Murase, H. and Nayar, S. K. "Visual Learning and Recognition of 3-D Objects from Appearance," *International Journal of Computer Vision*, Vol. 14, 1995, pp. 5-24.

[22] Nayar, S. K. and Bolle, R. "Reflectance Based Object Recognition," *International Journal of Com-*

puter Vision, Vol. 17, No. 3, March 1996, pp. 219-240.

[23] Oren, M. and Nayar, S. K., "Generalization of the Lambertian Model," *Proceedings Image Understanding Workshop 1993*, Morgan Kaufmann Publishers, 1993, pp. 1037-1048.

[24] Shapiro, L. G. and Costa, M. S. "Appearance-Based 3D Object Recognition," *Proceedings of the International Workshop on Object Representation in Computer Vision*, Springer-Verlag 1995, pp. 51-64.

[25] Slater, D. and Healey, G., "Using A Spectral Reflectance Model For The Illumination-Invariant Recognition Of Local Image Structure," *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 1996, pp. 770-775.

[26] Slater, D. and Healey, G., "Object Recognition Using Invariant Profiles," *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 1997, pp. 827-832.

[27] Slater, D. and Healey, G. "The Illumination-Invariant Recognition of 3D objects Using Local Color Invariants," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 18, No. 2, February 1996, pp. 206-210.

[28] Slater, D. and Healey, G., "What Is the Spectral Dimensionality of Illumination Functions in Outdoor Scenes?," *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 1998, pp. 105-110.

[29] Swain, M. J. and Ballard, D. H. "Color Indexing," *International Journal of Computer Vision*, Vol. 7, No.1, 1991, pp. 11-32.

[30] Turk, M. and Pentland, A. "Face Recognition Using Eigenfaces," *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 1991, pp. 586-591.

[31] Wolff, L. B. "Diffuse-Reflectance Model for Smooth Dielectric Surfaces," *Journal of the Optical Society of America A*, Vol. 11, No. 11, November 1994, pp. 2956-2968.