Objective Colour from Multispectral Imaging

Elli Angelopoulou

Computer Science Department, Stevens Institute of Technology
Hoboken, NJ 07030, USA
elli@cs.stevens-tech.edu

Abstract. The light reflected from a surface depends on the scene geometry, the incident illumination and the surface material. One of the properties of the material is its albedo ( ) and its variation with respect to wavelength. The albedo of a surface is purely a physical property. Our perception of albedo is commonly referred to as colour. This paper presents a novel methodology for extracting the albedo of the various materials in the scene independent of incident light and scene geometry. A scene is captured under different narrow-band colour 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 normalized albedo descriptor which is invariant to scene geometry and incident illumination for diffuse surfaces.

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 variations in an object's appearance. Among the most well-known problems is colour constancy, the task of consistently identifying colours, despite changes in illumination conditions. Maloney and Wandell[18] were the first to develop a tractable colour constancy algorithm by modeling both the surface reflectance and the incident illumination as a finite dimensional linear model. This idea was further explored by Forsyth[6], Ho et al.[14], Finlayson et al.[5, 7, 4, 1] and Healey and Slater[11]. colour is a very important cue in object identification. Swain and Ballard[23] showed that objects can be recognized by using colour information alone. Combining colour cues with colour constancy[11, 21, 7, 4] generated even more powerful colour-guided object recognition systems.

Extracting reflectance information is an under-constrained problem. All the aforementioned methodologies had to introduce some additional constraints that may limit their applicability. For example, most colour 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

D. Vernon (Ed.): ECCV 2000, LNCS 1842, pp. 359–374, 2000.
Springer-Verlag Berlin Heidelberg 2000

images for extracting illumination invariant colour information. Furthermore, a considerable body of work on colour assumes that the incident illumination has two or three degrees of freedom. However, Slater and Healey[22] showed that for outdoor scenes, the illumination functions have seven degrees of freedom.

We propose a new technique for extracting colour information that is invariant to geometry and incident illumination. We examine the rate of change in reflected intensity with respect to wavelength over the visible part of the electromagnetic spectrum. For diffuse surfaces, independent of the particular model of reflectance, the only factor that contributes to variations over the wavelength is the albedo of the surface. Thus, what we end up extracting is the reflectivity profile of the surface. 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 eight different colour filters and compute the spectral derivatives of the scene. Unlike many colour constancy methods, we employ multiple colour filters with narrow bandwidth (eight 10nm wide filters as opposed to the typical three 75nm filters). The use of narrow filters increases the discriminatory power of our method. The advantage of the traditional RGB systems is that they resemble the human visual sensor. Unfortunately, as Hal land Greenberg[10] have shown, the employment of only 3 bands of wavelengths introduces significant colour distortions. An additional advantage of our technique over the more traditional band-ratios is that spectral derivatives are used on a per pixel basis. They do not depend on neighbouring regions, an assumption that is common in other photometric methods, which use logarithms and/or narrow-band filters[7].

The collection of spectral derivatives evaluated at different wavelengths forms a spectral gradient. This gradient is normalized albedo descriptor, invariant to scene geometry and incident illumination for smooth diffuse surfaces. Experiments on surfaces of different colours and materials demonstrate the ability of spectral gradients to: a) identify surfaces with the same albedo under variable viewing conditions; b) discriminate between surfaces that have different albedo; and c) provide a measure of how close the colours of the two surfaces are.

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 $\boldsymbol{p} = (x, y, z)$ in the scene depends on the light illuminating the scene, E and the surface reflectance S of the surfaces in the scene:

$$I(\boldsymbol{p}, \quad) = E(\boldsymbol{p}, \quad)S(\boldsymbol{p}, \quad) \tag{1}$$

where , the wavelength, shows the dependence of incident and reflected light on wavelength. The reflectance function $S(\boldsymbol{p}, \quad)$ depends on the surface material, the scene geometry and the viewing and incidence angles.

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

$$E(_i, _i,) = e( )E(_i, _i) \tag{2}$$

where ($_i$, $_i$) are the spherical coordinates of the unit-length light-direction vector and $e( )$ is the illumination spectrum. Note that, the incident light intensity is included in $E(_i, _i)$ and may vary as the position of the illumination source changes. The scene brightness then becomes:

$$I(\boldsymbol{p},) = e(\boldsymbol{p},)E(\boldsymbol{p}, _i, _i)S(\boldsymbol{p},) \tag{3}$$

Before we perform any analysis we simplify the scene brightness equation by taking its logarithm. The logarithmic brightness equation reduces the product of the incident illumination $E(\boldsymbol{p},)$ and the surface reflectance $S(\boldsymbol{p},)$ into a sum:

$$L(\boldsymbol{p},) = \ln e(\boldsymbol{p},) + \ln E(\boldsymbol{p}, _i, _i) + \ln S(\boldsymbol{p},) \tag{4}$$

A colour descriptor which is invariant to viewpoint, scene geometry and incident illumination is albedo. Albedo ( ) is the ratio of electromagnetic energy reflected by a surface to the amount of electromagnetic energy incident upon the surface[20]. A profile of albedo values over the entire visible spectrum is a physically based descriptor of colour. Albedo is one of a multitude of factors that determine the surface reflectance function $S(\boldsymbol{p},)$.

Extracting the absolute albedo values directly from the image brightness without any a-priori knowledge of the materials of the scene or the imaging conditions, is an under-constrained problem. However, an invariant physically based colour descriptor can be computed by taking samples of the reflected light at different wavelengths and measuring the change in scene brightness between wavelengths. One technique for measuring such changes is by calculating the *spectral derivative*, which is the partial derivative of the logarithmic image with respect to wavelength :

$$L (\boldsymbol{p},) = \frac{e (\boldsymbol{p},)}{e(\boldsymbol{p},)} + \frac{S (\boldsymbol{p},)}{S(\boldsymbol{p},)} \tag{5}$$

where $e (\boldsymbol{p},) = e(\boldsymbol{p},)/$ is the partial derivative of the spectrum of the incident light with respect to wavelength and $S (\boldsymbol{p},) = S(\boldsymbol{p},)/$ is the partial derivative of the surface reflectance with respect to wavelength. Our work concentrates on the visible part of the electromagnetic spectrum, i.e. from 400nm to 700nm. Ho, Funt and Drew[14] have shown, that for natural objects the surface spectral reflectance curves, i.e. the plots of $S(\boldsymbol{p},)$ versus , are usually reasonably smooth and continuous over the visible spectrum.

For diffuse objects, the spectral derivative is a measure of how albedo changes with respect to wavelength. As such, it is invariant to incident illumination, scene geometry, viewpoint, and material macrostructure.

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 slowly over small increments of . This means that its derivative with respect to wavelength is approximately zero.

$$e\ (\boldsymbol{p},\ )\ \ 0 \tag{6}$$

An exception to this assumption are the localized narrow spikes that are present in the spectrum of fluorescent light (see fig. 1). The partial derivative with respect to wavelength is undefined at these spikes.

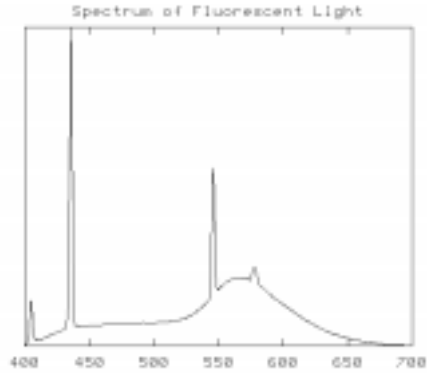

Fig. 1. The emitance spectrum of fluorescent light.

Although different types of fluorescent light exhibit such spikes at different parts of the spectrum, these spikes have always very narrow width. Thus, we can discard in our analysis the wavelengths around which these spikes occur. By discarding these outliers, the assumption of a slowly changing e is valid over most of the visible range. This implies that one can safely assume that in general the partial derivative of the logarithmic image depends only on the surface reflectance:

$$L\ (\boldsymbol{p},\ )\ \ \frac{S\ (\boldsymbol{p},\ )}{S(\boldsymbol{p},\ )} \tag{7}$$

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 describes the behaviour 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 $_i(\boldsymbol{p})$ between the surface normal at point $\boldsymbol{p}$ and the incident illumination changes.

$$S(\boldsymbol{p}, \ ) = \cos \ _i(\boldsymbol{p}) \ \ (\boldsymbol{p}, \ ) \tag{8}$$

where $(\boldsymbol{p}, \ )$ 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. Furthermore, the scene geometry, including the angle of incidence, is independent of wavelength. Therefore, when we take the partial derivative with respect to wavelength, the geometry term vanishes:

$$L \ (\boldsymbol{p}, \ ) \ \ \frac{S \ (\boldsymbol{p}, \ )}{S(\boldsymbol{p}, \ )} = \frac{(\boldsymbol{p}, \ )}{(\boldsymbol{p}, \ )} \tag{9}$$

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

One of the advantages of spectral derivatives is that since the dependence on the angle of incidence gets cancelled out, there is no need for assuming an infinitely distant light source. The incident illumination can vary from one point to another, without affecting the resulting spectral derivative.

4.2 Smooth Diffuse Reflectance Model

In reality there are very few objects that exhibit perfectly Lambertian reflectance. For smooth diffuse objects, the reflected light varies with respect to viewpoint. Wolff[25] introduced a new smooth diffuse reflectance model that incorporates the dependence on viewpoint:

$$\begin{aligned} S(\boldsymbol{p}, \ ) &= \cos \ _i(\boldsymbol{p}) \ \ (\boldsymbol{p}, \ )(1 - F(\ _i(\boldsymbol{p}), n(\boldsymbol{p}))) \\ &\text{x} \ \ 1 - F \ \sin^{-1} \ \frac{\sin \ _r(\boldsymbol{p})}{n(\boldsymbol{p})} \ , \frac{1}{n(\boldsymbol{p})} \end{aligned} \tag{10}$$

where $_i(\boldsymbol{p})$ and $_r(\boldsymbol{p})$ are the incidence and viewing angles respectively, $F()$ is the Fresnel reflection coefficient, and $n(\boldsymbol{p})$ is the index of refraction. The index of refraction does indeed depend on wavelength. However, in dielectrics the refractive index changes by a very small amount over the visible range[2, 3, 8]. Thus, for dielectrics n is commonly treated as a material constant under visible light.

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

$$\begin{aligned} \ln S(\boldsymbol{p}, \quad) \;=\; & \ln\cos \;_i(\boldsymbol{p}) + \ln \;\; (\boldsymbol{p}, \quad) + \ln(1 - F(\;_i(\boldsymbol{p}), n(\boldsymbol{p}))) \\ & + \ln \;\; 1 - F \;\; \sin^{-1} \;\; \frac{\sin \;_r(\boldsymbol{p})}{n(\boldsymbol{p})} \;, \frac{1}{n(\boldsymbol{p})} \end{aligned} \tag{11}$$

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. The dependence on scene geometry including the viewing and incidence angles have been cancelled out. Again, since the spectral derivative is independent of incident illumination, the direction of incident light can vary from one point to the next, without affecting the spectral derivative. The spectral derivative becomes:

$$L \;\; (\boldsymbol{p}, \quad) \;\; \frac{S \;\; (\boldsymbol{p}, \quad)}{S(\boldsymbol{p}, \quad)} = \frac{(\boldsymbol{p}, \quad)}{(\boldsymbol{p}, \quad)} \tag{12}$$

4.3 Generalized Lambertian Model

Of course, not all diffuse surfaces are smooth. Oren and Nayar[19] developed a generalized Lambertian model which describes the diffuse reflectance of surfaces with substantial macroscopic surface roughness. The macrostructure of the surface is modelled as a collection of long V-cavities. (Long in the sense that the area of each facet of the cavity is much larger than the wavelength of the incident light.) The modelling of a surface with V-cavities is a widely accepted surface description[24, 13].

The light measured at a single pixel of an optical sensor is an aggregate measure of the brightness reflected from a single surface patch composed of numerous V-cavities. Each cavity is composed of two planar Lambertian facets with opposing normals. All the V-cavities within the same surface patch have the same albedo, . Different facets can have different slopes and orientation. Oren and Nayar assume that the V-cavities are uniformly distributed in azimuth angle orientation on the surface plane, while the facet tilt follows a Gaussian distribution with zero mean and standard deviation . The standard deviation can be viewed as a roughness parameter. When $= 0$, all the facet normals align with the mean surface normal and produce a planar patch that exhibits an approximately Lambertian reflectance. As increases, the V-cavities get

deeper and the deviation from Lambert's law increases. Ignoring interreflections from the neighbouring facets, but accounting for the masking and shadowing effects that the facets introduce, the Oren-Nayar model approximates the surface reflectance as:

$$\begin{aligned} S(\boldsymbol{p},\ ,\ ) = \frac{(\boldsymbol{p},\ )}{}\cos\ _i(\boldsymbol{p})[C_1(\ ) \\ +\cos(\ _r(\boldsymbol{p}) -\ _i(\boldsymbol{p}))C_2(\ ;\ ;\ _r -\ _i;\ ;\boldsymbol{p})\tan\ (\boldsymbol{p}) \\ +(1-\left|\cos(\ _r -\ _i)\right|)C_3(\ ;\ ;\ ;\boldsymbol{p})\tan\ \frac{(\boldsymbol{p})+\ (\boldsymbol{p})}{2}\ \Big] \end{aligned} \tag{13}$$

where $(\ _i(\boldsymbol{p}),\ _i(\boldsymbol{p}))$ and $(\ _r(\boldsymbol{p}),\ _r(\boldsymbol{p}))$ are the spherical coordinates of the angles of incidence and reflectance accordingly, $(\boldsymbol{p}) = max(\ _i(\boldsymbol{p}),\ _r(\boldsymbol{p}))$ and $(\boldsymbol{p}) = min(\ _i(\boldsymbol{p}),\ _r(\boldsymbol{p}))$. $C_1()$, $C_2()$ and $C_3()$ are coefficients related to the surface macrostructure. The first coefficient, $C_1()$ depends solely on the distribution of the facet orientation, while the other two depend on the surface roughness, the angle of incidence and the angle of reflectance:

$$C_1(\ ) = 1 - 0.5\frac{^2}{^2 + 0.33} \tag{14}$$

$$C_2(\ ;\ ;\ _r -\ _i;\ ;\boldsymbol{p}) = \begin{cases} 0.45\dfrac{^2}{^2+0.09}\sin\ (\boldsymbol{p}) & \text{if } \cos(\ _r(\boldsymbol{p}) -\ _i(\boldsymbol{p})) \quad 0 \\ 0.45\dfrac{^2}{^2+0.09}\ \sin\ (\boldsymbol{p}) - \dfrac{2\ (\boldsymbol{p})}{}^3 & \text{otherwise} \end{cases} \tag{15}$$

$$C_3(\ ;\ ;\ ;\boldsymbol{p}) = 0.125\ \frac{^2}{^2+0.09}\ \ \frac{4\ (\boldsymbol{p})\ (\boldsymbol{p})}{^2}^{\ 2} \tag{16}$$

For clarity of presentation, we set:

$$\begin{aligned} V(\boldsymbol{p},\ ) = C_1(\ ) + \cos(\ _r(\boldsymbol{p}) -\ _i(\boldsymbol{p}))C_2(\ ;\ ;\ _r -\ _i;\ ;\boldsymbol{p})\tan\ (\boldsymbol{p}) \\ +(1-\left|\cos(\ _r -\ _i)\right|)C_3(\ ;\ ;\ ;\boldsymbol{p})\tan\ \frac{(\boldsymbol{p})+\ (\boldsymbol{p})}{2} \end{aligned} \tag{17}$$

The term $V(\boldsymbol{p},\ )$ accounts for all the reflectance effects which are introduced by the roughness of the surface. The angles of incidence and reflectance, as well as the distri-

bution of the cavities affect the value of the function $V(\boldsymbol{p},\ )$. The Oren-Nayar reflectance model can then be written more compactly as:

$$S(\boldsymbol{p},\ ,\ ) = \frac{(\boldsymbol{p},\ )}{}\cos\ _i(\boldsymbol{p})V(\boldsymbol{p},\ ) \tag{18}$$

Once again, when we take the logarithm of the surface reflectance function, we simplify the underlying model, by turning the product into a sum:

$$\ln S(\boldsymbol{p},\ ) = \ln\ (\boldsymbol{p},\ ) - \ln\ + \ln\cos\ _i(\boldsymbol{p}) + \ln V(\boldsymbol{p},\ ) \tag{19}$$

When we compute the partial derivative of this logarithm with respect to wavelength, the only term that remains is the albedo. The angle of incidence $_i(\boldsymbol{p})$ and the constant , are independent of wavelength. None of the terms in the function $V(\boldsymbol{p},\ )$, see equation (17), vary with the wavelength. More specifically, $C_1()$ is only a function of , which is a surface property independent of wavelength. $C_2()$ and $C_3()$ are functions of , the viewing and incidence angles, and of and , which in turn are also functions of the viewing and incidence directions. None of these factors is affected by wavelength. Thus, even for rough diffuse surfaces, the spectral derivative is:

$$L\ (\boldsymbol{p},\ )\ \ \frac{S\ (\boldsymbol{p},\ )}{S(\boldsymbol{p},\ )} = \frac{(\boldsymbol{p},\ )}{(\boldsymbol{p},\ )} \tag{20}$$

5 Spectral Gradient

For diffuse surfaces, independent of the particulars of the reflectance behaviour, the partial derivative with respect to wavelength of the logarithmic image $L\ (\boldsymbol{p},\ )$ is a function of only the surface albedo. More specifically the spectral derivative approximates the normalized partial derivative of the albedo with respect to wavelength $(\boldsymbol{p},\ )/\ (\boldsymbol{p},\ )$. By normalized we mean that the derivative is divided by the magnitude of the albedo itself.

Consider now a collection of spectral derivatives of a logarithmic image at various spectral locations $_k$, $k = 1, 2, 3, \ldots, M$. The resulting *spectral gradient* is an M-dimensional vector $(L\ _1, L\ _2, \ldots, L\ _M)$ 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. Thus, the spectral gradient is a colour descriptor that depends purely on the surface albedo, a physical material property. Although our perception of colour depends on the surface albedo, a colour descriptor, such as the spectral gradient which is purely albedo-based does not convey any perceptual meaning of colour. It remains unaltered by changes in the environment (viewing position, illumination, geometry).

6 Experiments

6.1 Experimental Setup

In order to compute the spectral derivatives we took images of each scene under eight different narrow bandpass filters: a Corion S10-450-F, a Corion S10-480-F, a Corion S10-510-F, a Corion S10-540-F, a Corion S10-570-F, a Corion S10-600-F, a Corion S10-630-F and a Corion S10-660-F. Each of these filters has a bandwidth of approximately 10nm and a transmittance of about 50%. The central wavelengths are at 450nm, 480nm, 510nm, 540nm, 570nm, 600nm, 630nm and 660nm respectively. If one were to assign colour names to these filters, he/she could label them as follows: 450nm=blue, 480nm=cyan, 510=green, 540=yellow, 570=amber, 600=red, 630=scarlet red, 660=mauve.

The use of narrow bandpass filters allowed us to closely sample almost the entire visible spectrum. The dense narrow sampling permitted us to avoid sampling (or ignore samples) where the incident light may be discontinuous (see section 3) Hall and Greenberg[10] have demonstrated that such a sampling density provides for the reproduction of a good approximation of the continuous reflectance spectrum. The images were captured with a Sony XC-77 camera using a 25mm lens (fig. 2.)

Fig. 2. A picture of the experimental setup showing the analog greyscale camera and the filter wheel mounted in front of it.

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 150W bulb and a 200W bulb and two different light positions. One illumination position was to the left of the camera and about 5cm below the camera. Its direction vector formed approximately a 30 angle with the optic axis. The other light-bulb position was to the right of the camera and about 15cm above it. Its direction vector formed roughly a 45 angle with the optic axis. Both locations were 40cm away from the scene.

The imaged objects were positioned roughly 60cm from the camera/filter setup. We tried four different types of materials: foam, paper, ceramic and a curved metallic surface painted with flat (matte) paint. The foam and the paper sheets came in a variety of colours. The foam which was a relatively smooth and diffuse surface came in white, pink, magenta, green, yellow, orange and red samples (see fig. 3 top left). The paper had a rougher texture and came in pink, fuchia, brown, orange, yellow, green, white, blue, and violet colours (see fig. 3 top right). We also took images of a pink ceramic plate. Its surface was extremely smooth and exhibited localized specularities and self-shadowing at the rims (see fig. 3 bottom left). Finally we took images of single albedo curved surfaces (a mug and a painted soda-can) to test the invariance to the viewing and incidence angles (see fig. 3 bottom middle and right respectively).

Fig. 3. Pictures of the objects and materials used in the experiments. Top left: various colours of foam; top right: various colours of paper; bottom left: a pink ceramic plate; bottom center: a white ceramic mug; bottom right: a white spray-painted soda can.

Fig. 4 shows samples of the actual images that we captured. All the images in these figure are taken using the Corion S10-600-F filter.The light illuminating the scene is a 200W tungsten bulb positioned at the lower left corner (as can be seen from the highlights). On the top left are the coloured samples of foam, while on the right are the multiple coloured samples of paper. By comparing the stripes in fig. 4 with those in fig. 3, one can tell which colours reflect around 600nm (pink, orange, yellow and white in the case of the paper)

.

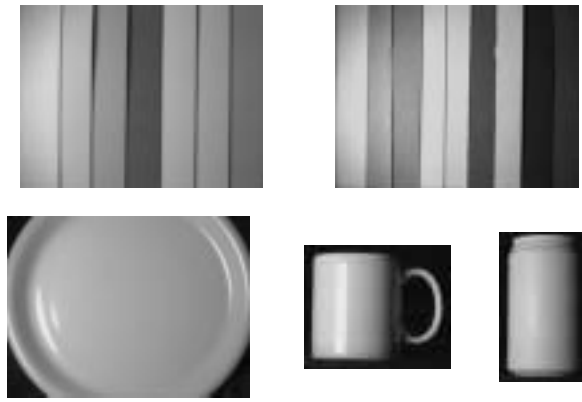

Fig. 4. Sample filtered images of the objects shown in fig. 3. All the images in this figure were taken with the red Corion S10-600-F filter.

6.2 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:

$$L_w = \ln(I_w) \qquad \text{where } w = 450, 480, 510, 540, 570, 600, 630, 660 \tag{21}$$

where I_w was the image of a scene taken with the S10-*W*-F filter and L_w was its logarithmic image. The last step involved the computation of the spectral derivatives of the logarithmic images. Differentiation was approximated via finite-differencing. Thus, each L_k was computed over the wavelength interval $= 30nm$ by subtracting two logarithmic images taken under two different colour filters which were 30nm apart:

$$L_k = L_{w+30} - L_w \tag{22}$$

where $k = 1, 2, \ldots, 7$ and $w = 450, 480, 510, 540, 570, 600, 630$ accordingly. In our setup the spectral gradient was a 7-vector:

$$(L_1, L_2, L_3, L_4, L_5, L_6, L_7) = (L_{480} - L_{450}, L_{510} - L_{480}, \ldots, L_{660} - L_{630}) \tag{23}$$

This vector was expected to remain constant for diffuse surfaces with the same albedo profile, independent of variations in viewing conditions. At the same time, the

spectral gradient should differ between distinct colours. Furthermore, the more distant the colours are, the bigger the difference between the respective spectral gradients should be.

The following figures show the plots of the spectral gradient values for each colour versus the wavelength. The horizontal axis is the wavelength, ranging from the $_1 = 480 - 450$ interval to the $_7 = 660 - 630$ interval. The vertical axis is the normalized partial derivative of albedo over the corresponding wavelength interval. Fig. 5 shows the plots of different colours of paper fig. 5(a) and of different colours of foam fig. 5(b). Within each group, the plots are quite unique and easily differentiable from each other.

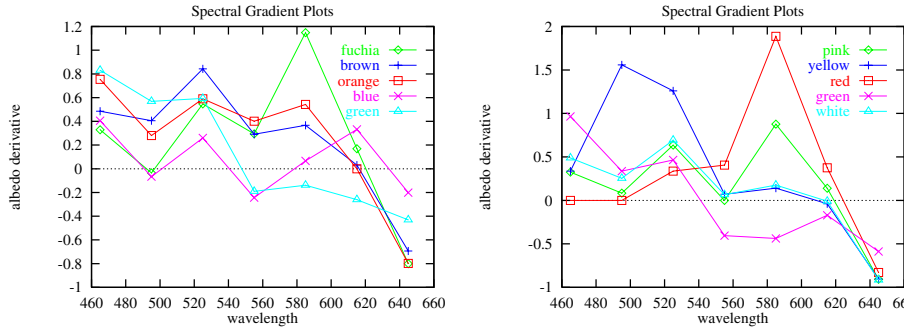

Fig. 5. Spectral gradients of different colours of (a) paper and (b) foam under the same viewing conditions (same illumination, same geometry).

On the other hand, the spectral gradients of different surfaces of the same colour, generate plots that look almost identical. Fig. 6 shows the gradient plots for the white paper, the white foam, the white mug, and the white painted soda can.

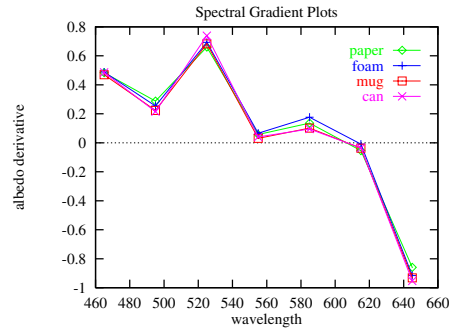

Fig. 6. Spectral gradients of different white surfaces (foam, paper, ceramic mug, diffuse can) under the same viewing conditions (same illumination, same geometry).

In a similar manner, when we have similar but not identical colours, the spectral gradient plots resemble each other, but are not as closely clustered. Fig. 7 shows the spectral gradients of various shades of pink and magenta. The closer the two shades are, the more closely the corresponding plots are clustered.

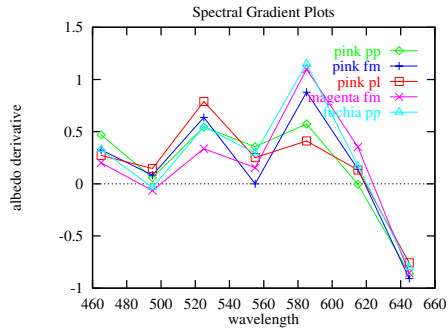

Fig. 7. Spectral gradients of different shades of pink and different materials: pink foam, magenta foam, pink paper, fuchia paper and pink ceramic plate. All images were taken under the same viewing conditions (same illumination, same geometry).

The next couple of figures demonstrate that the spectral gradient remains constant under variations in illumination and viewing. This is expected as spectral gradients are purely a function of albedo. The plots in fig. 8 were produced by measuring the spectral gradient for the same surface patch while altering the position and intensity of the light sources.

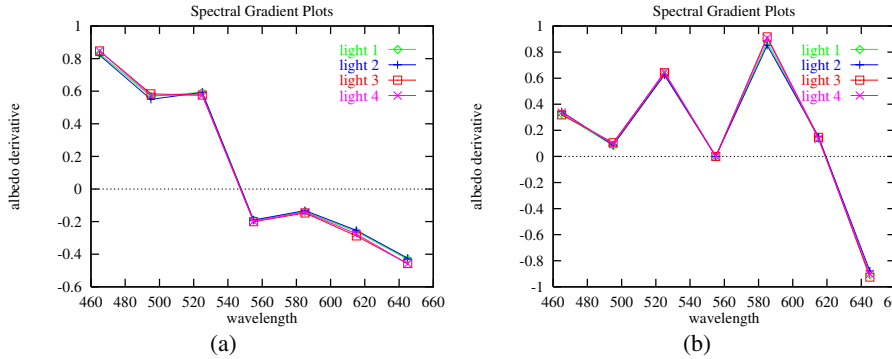

Fig. 8. Spectral gradients of the same colour (a) green and (b) pink under varying illumination. Both the position and the intensity of illumination is altered, while the viewing position remains the same.

In order to demonstrate the invariance to the viewing angle and the surface geometry, we show the plots of the spectral gradients produced by different patches of the

same curved object (the painted soda can in this case). At least one patch has its tangent roughly parallel to the image plane, while the other patches are at mildly to quite oblique angles to the viewer and/or the incident light. As fig. 9 shows, the spectral gradient plots still remain closely clustered.

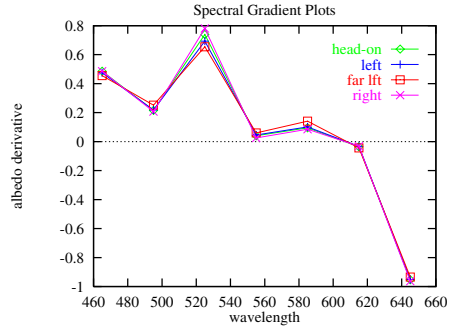

Fig. 9. Spectral gradients of white colour at different angles of incidence and reflectance. The spectral gradients at different surface patches of the white soda can are shown. The surface normals for these patches vary from almost parallel to the optic axis, to very oblique.

This last series of tests, showed also a limitation of our technique. For very oblique angles of incidence, where $_i > 60$ the spectral gradients do not remain constant. Deviations at large angles of incidence are a known physical phenomenon[17]. Oren and Nayar[19] also point out that in this special case, most of the light that is reflected from a surface patch is due to interreflections from nearby facets. In fig. 10 we show the spectral gradient plots obtained from patches of the painted can which are at large angles of incidence. (The light is almost grazing the surface). For comparison purposes we included also two plots produced from patches with smaller angles of incidence.

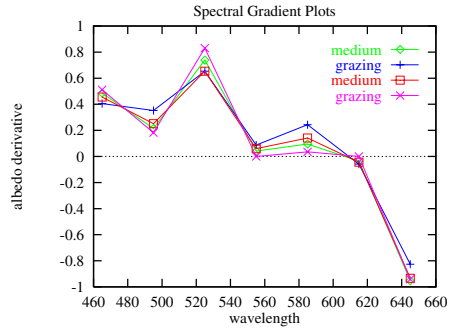

Fig. 10. Spectral gradients of white colour at large angles of incidence. The spectral gradients at different surface patches of the white soda can are shown. For all these patches the incident light is almost grazing the surface.

As a last note, we would like to point out that although some of our test objects exhibit specular highlights, we did not analyze any of the data from the specular areas. Our methodology assumes diffuse reflectance and we limited our tests to the diffuse parts of the images. There is currently ongoing investigation of how specularities and interreflections affect the spectral gradient.

7 Conclusions and future work

We developed a technique for extracting surface albedo information which is invariant to changes in illumination and scene geometry. The spectral information we extract is purely a physical property. Hence the term objective colour. 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 neighbouring regions. The effectiveness of spectral gradients as a colour 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 material classification, grey-scale colour constancy, or in tracking different regions under variable illumination.

Spectral gradients and the underlying dense spectral sampling provide a rich description of surface reflectance behaviour. We are already studying the phenomenon of interreflections at large angles of incidence. We are also examining the behaviour of highlights under this fine multispectral sampling. Another topic that we are working on is the simultaneous employment of multiple illumination sources. Finally a topic that we are also planning to address is the effect of changes in the spectrum of the incident illumination.

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. Bass, M., ed., *Handbook of Optics: Fundamentals, Techniques and Design*, 2nd ed., Vol. I, McGraw-Hill, 1995.
3. Born, M. and Wolf, E., *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*, 5th ed., Pergamon Press, 1975
4. Finlayson, G. D., "Color in Perspective," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 18, No. 10, October 1996, pp. 1034-1038.
5. 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.
6. Forsyth, D., "A Novel Algorithm for Color Constancy," *International Journal of Computer Vision*, Vol. 5, No.1, 1990, pp. 5-36.
7. 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.

8. Garbuny, Max, *Optical Physics*, Academic Press, 1965.
9. 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.
10. Hall, R. A. and Greenberg, D. P., "A Testbed for Realistic Image Synthesis," *IEEE Computer Graphics and Applications*, Vol. 3, No. 8, 1983, pp. 10-19.
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. Hering, R. G. and Smith, T. F., "Apparent Radiation Properties of a Rough Surface," *AIAA Progress in Astronautics and Aeronautics*, Vol. 23, 1970, pp. 337-361.
14. 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.
15. Horn, B. K. P., "Understanding Image Intensities," *Artificial Intelligence,* Vol. 8, No. 2, 1977, pp. 1-31.
16. Horn, B. K. P. and Brooks, M. J., *Shape from Shading*, MIT Press, 1989.
17. Kortum, G., *Reflectance Spectroscopy*, Springer-Verlag, 1969.
18. 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.
19. Oren, M. and Nayar, S. K., "Generalization of the Lambertian Model and Implications for Machine Vision," *International Journal of Computer Vision,* Vol. 14, No.3, April 1995, pp. 227-251.
20. Sabins, F. F., *Remote Sensing - Principles and Interpretation*, 3rd ed.W. H. Freeman and Co., 1997.
21. 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.
22. 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.
23. Swain, M. J. and Ballard, D. H., "Color Indexing," *International Journal of Computer Vision*, Vol. 7, No.1, 1991, pp. 11-32.
24. Torrance, K. and Sparrow, E., "Theory for Off-Specular Reflection from Rough Surfaces," *Journal of the Optical Society of America*, Vol. 67, September 1967, pp. 1105-1114.
25. 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.