Removing Shadows from Images of Stained Glass Windows

Shanmugalingam Suganthan, Lindsay MacDonald, and Alfredo Giani
Colour & Imaging Institute, University of Derby, United Kingdom

Abstract

Shadows may be formed on stained glass windows by structural bars supporting the leaded panels, or by external protective wire grilles, or by masonry such as mullions or buttresses, or external objects such as trees. The eye tends to 'discount' such shadow formations when viewing the actual windows even though in the photographic images they are very clearly visible. Digital modelling techniques enable the shadows to be characterised and removed with a reasonable degree of success.

1. A Mixture Model (MM) for Shadows

The stained glass window (SGW) is significantly different from most other subjects because its colour is generated by transmitted, rather than reflected, light. Also it has a wider dynamic range between highlight and shadow areas than most "real world" scenes. In many cases the background is visible through the glass, typically trees, foliage, sky or other buildings. Images of SGW taken with external illumination very often contain shadows cast by structures such as support bars and protective wire grilles (Figure1). These shadowing objects are often unremovable, as they difficult to access, or constitute structural elements of the window. It is therefore necessary to provide a suitable set of Image Processing tools to facilitate the removal of unwanted shadows from images of the windows.

The observed image may be modelled as a combination of a "True Stained Glass image" and a "Grille/bar image". It is our goal to separate this mixture into its original components. Therefore the nature of this combination must be investigated in order to produce a valid model. We consider the image formation model. The RGB sensor responses of the digital camera can be represented as[1]:

$$R = \sum_{n=1}^{N} E(\lambda_n) S(\lambda_n) r(\lambda_n)$$

$$G = \sum_{n=1}^{N} E(\lambda_n) S(\lambda_n) g(\lambda_n) \qquad (1)$$

$$B = \sum_{n=1}^{N} E(\lambda_n) S(\lambda_n) b(\lambda_n)$$

Nick Becket, English Heritage

Figure 1. Shadows of horizontal support bar and wire mesh grille in a stained glass window. The detail (below) shows both the transmitted image of the grille its shadow cast by the sunlight. The window is in the chancel of the Church of St. Mary, Studley Royal, Yorkshire, UK.

where:
$E(\lambda_n)$ is spectral power of illumination at wavelength λ_n ;
$S(\lambda_n)$ is the object's transmittance wavelength λ_n ;
$r(\lambda_n)$ $g(\lambda_n)$ $b(\lambda_n)$ are sensitivities of three camera channels; λ spans the whole visible spectrum;
N is the sampling ratio chosen to represent the continuous function of wavelength with sufficient accuracy.

A shadow is produced when a surface element (grille/bar) acts as a blocker of some of the light emitted by another surface element (light source). The blocker stops some of the radiant energy along its path. The energy is then absorbed or redirected. Because there is usually a gap between the grille/bar and the window, the shadowed region of the glass still receives partial illumination from the non-point source of the sun's disc at the edges (penumbra) and from the diffused skylight overall. This suggests that the response of the camera in the shadowed region should be related to the response in the neighbouring unshadowed region. Since the effect of the blocker is to reduce the amount of illumination $E(\lambda_n)$ uniformly over λ_n (at least to a first approximation), it is clear from Eq. 1 that there is a multiplicative relationship between the pixel value P_S corresponding to a grille line and its closest non-shadowed pixel value P_{NS}:

$$P_S = \alpha(\cdot)\, P_{NS} \qquad \alpha(\cdot) \in [0,1] \tag{2}$$

Note that in Eq. 2 we have assumed that α depends on a not-yet-specified set of parameters. This is because the reduction of illumination may depend on factors such as the distance between the blocker and the glass, on the thickness of the blocker and on the kind of glass. An opalescent glass will tend to scatter light, therefore yielding values of α closer to unity. On the other hand, a very transparent glass will produce a more distinct and possibly darker shadow, therefore yielding values of α closer to zero. We can model the observed image $I_o(x,y)$ as:

$$I_o(x,y) = \alpha(x,y) \cdot I(x,y) \tag{3}$$

Where: x, y are image co-ordinates in pixels; I_o is observed SGW image; I is true SGW image; α is grille/bar shadow image; $\alpha(x,y) \in [0,1]$ if x, y lies on a bar line, else 1.

2. A Physical Model (PM) for Shadows

A prerequisite for the development of a physically based shadow computation is knowledge about the distribution of light in a scene from each of the illuminating sources. Figure 2 illustrates the concept of an illuminating hemisphere.[2,3]

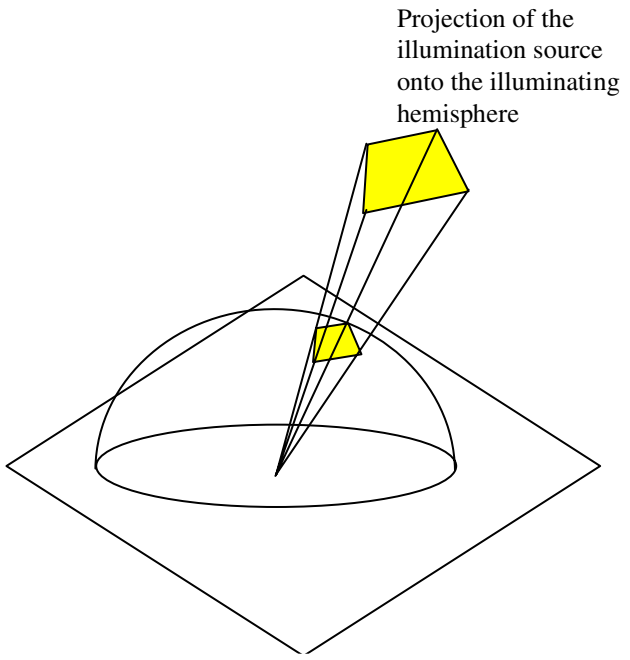

Figure 2. The illuminating hemisphere

The illuminating hemisphere is a notational convenience for describing the illumination *events* above or below a surface. These events, such as light sources or other reflecting surfaces, are projected onto the hemisphere. For computational convenience we assume a hemisphere of radius 1. A solid angle specifies the amount of the hemisphere covered by the projection of the illumination. A differential solid angle is defined as a projected differential surface element on the hemisphere divided by the square of the radius of the hemisphere. The solid angle of an illumination event is determined by integrating the differential solid angle over the bounds of the projection of the event.

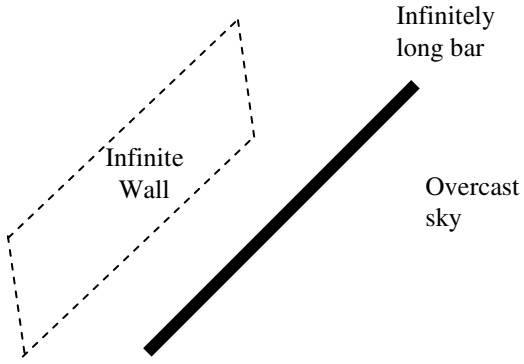

Figure 3. Brightness of the shadow cast by an infinite long bar onto an infinitely long wall, under a uniformly overcast sky.

Physical models can explain the formation of "simple" shadows. The most powerful tool for analysing this problem is to think about *what a source looks like from the surface.* This technique enables us to give a qualitative description of "brightness". An idealised form is the Wall-Bar model, where a bar of infinite length (bar) casts a shadow onto an infinite surface (wall) illuminated by a perfectly diffused light source (overcast sky) with the geometry shown in Figure 3. We wish to know the brightness distribution at the base of an infinitely long bar and infinitely high wall.

The solid angle of the bar is determined by projecting it onto the illuminating hemisphere above the surface and integrating the area of the projection. Figure 4 sketches the appearance of the infinitely long bar form at two points A and B on the infinite wall. The bar projected onto the illuminating hemisphere looks like the segment of an orange, converging to a polar point as the bar recedes to infinity. All points on a given horizontal line on the wall see the same input hemisphere, and so must have the same "brightness" but points along a vertical line see different amounts of the input hemisphere (Figure 5).

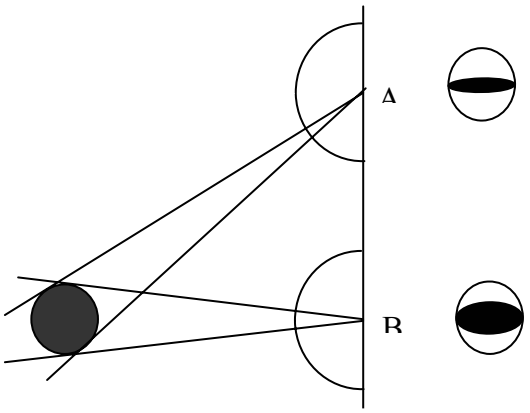

Figure 4. Projection of the bar onto the illuminating hemisphere

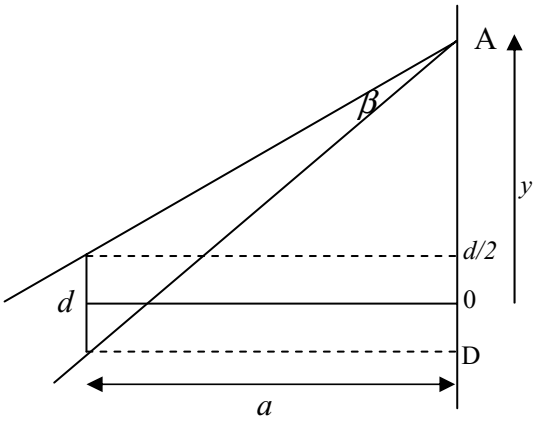

Figure 5. Projected outline of the bar

The angle β defines the projected outline of the bar and can be expressed as a function of the position y on wall as:

$$\beta = \arctan\left(\frac{ad}{a^2 + y^2 - \frac{d^2}{4}}\right) \tag{4}$$

Where a is the distance of the centre of the bar from the wall and d is the bar diameter. Integrating over the bounds of the projected area of the bar, the solid angle is given by:

$$\omega = 2\int_0^{\frac{\pi}{2}}\int_0^{\beta} \sin\theta d\theta d\phi = 2\beta \tag{5}$$

The brightness B is defined as:

$$B = \frac{2\pi - \varpi}{2\pi} = \frac{2\pi - 2\beta}{2\pi} = \frac{\pi - \beta}{\pi} \tag{6}$$

Based on this technique, in Figure 6 shows the typical shadow profile plotting brightness as a function of vertical position on the wall, with the distance a and bar diameter d as parameters. Figure 7 shows measured profile from a series of images captured under controlled illumination, in which only the distance a was changed and all other parameters were kept constant.

The resulting shadow profile is a bell-shaped curve. In order to determine the shape of the curve, it is necessary to know: the distance of the bar from the wall and the diameter of the bar. This simple model provides a useful first approximation, but for a more sophisticated model of transmissive media one should consider other effects, like multiple reflections within the glass, inhomogeneity of the glass itself, and scattering of light by impurities.

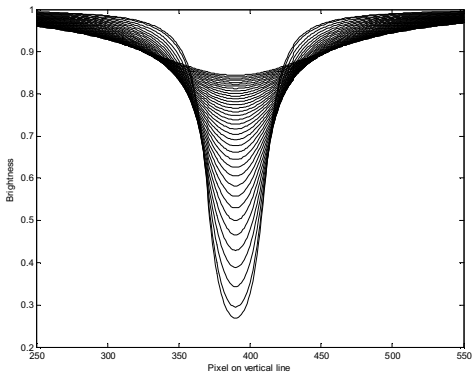

Figure 6. Physical model predictions: family of curves with the distance as a parameter

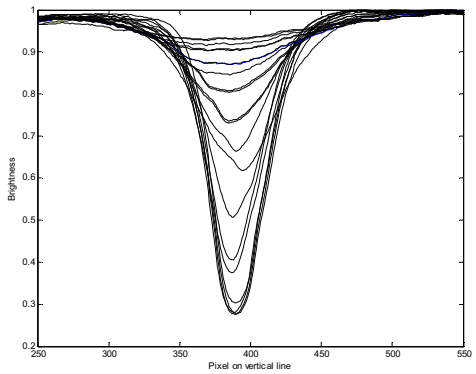

Figure 7. Measured profile on a sample window under controlled illumination: family of curves with the distance as a parameter

3. Removing the Bar Shadows

We captured a test image using a high-resolution digital camera, the Rollei 6008i with a Jenoptik *eyelike MF* digital back,[4] to photograph a Victorian (c.1860) stained glass panel and a metal bar placed on a large light table. A detail from the *RGB* test image was cropped to 600×600 pixels, then transformed into Hue-Saturation-Value (*HSV*) components, using the standard computer graphics model.[5] Figure 8 shows the resulting image components. It is clear that the shadow mainly affects the *V* channel rather than *H* and *S* channels, as shown by the profile of pixel values (Figure 9).

Figure 8. Test image: (a) RGB-image; (b) H-component; (c) S-component; (d) V-component.

In the case of shadows produced by bars across the SGW, if the position and diameter of the bar are roughly known then the PM model can be used to predict the shadow profile. Most common SGW images are photographed under diffuse lighting environments rather than direct light. Here we assume that the SGW panel transmittance process reduces the intensity of illumination by the same amount over all parts of the glass. So we compute the intensity of illumination arriving on the exterior surface of the SGW. Therefore our wall bar model is well suited to this problem. The conjectured mixture model is:

$$I_0(x,y) = \alpha(x,y)\, I(x,y) \quad \alpha(x,y) \in [0,1] \qquad (7)$$

where x,y are the image co-ordinates in pixels, I_0 is the observed SGW image, I is the true SGW image, and α is the profile of the shadow produced by the bar. Using the wall bar model, we estimate the α value then recover I. The algorithm was applied to our test images. Figure 10 shows three results compared to the original image.

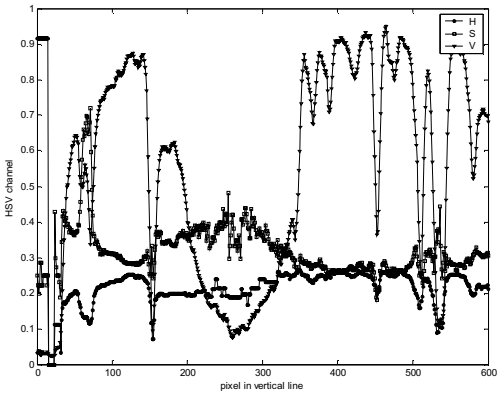

Figure 9. H, S and V profiles of pixels along a vertical line of the test image.

The results show that the wall-bar model developed is at least a good initial step but is over-simplistic. At the centre line of the bar the model gives encouraging results. The shadow is mostly removed, and the features of the glass hidden by the shadow are fully restored. Note the texture of the glass and the paintwork. The periphery of the shadow, however, shows that the model does not fit so well with the actual image profile. The result is an over-compensation of the data, meaning that actual image was brighter than predicted by the model. We conjecture that this is due to the lateral scattering of light occurring within the body of the glass. Another problem may be the uncertainty by which the diameter and distance of the bars are determined. We will investigate whether the model parameters can be fitted to the actual shadow edge profile observed in the image, or whether a more sophisticated blending model might be used.[6]

Figure 10. Results compared to the original image.

4. Removing the Grille Shadows

Removing the shadow cast by the protective wire grille presents a set of challenges somewhat different from the problem previously tackled. Grilles are typically constructed from a mesh of small-diameter iron or copper wire fixed to a frame 10-15 cm away from the glass. A soft shadow is cast by the grille because the scattering of light produces a shadow that is not completely black, but in fact inherits some chromaticity from the colour of the glass onto which the shadow is cast. In other words, the shadow line in the image is "dark greenish" where the interposed glass is green, "dark bluish" where the interposed glass is blue, and so on (see Figure 1). Although exceptions might be found, this can be considered as a general behaviour in typical illumination conditions.

Another characteristic of this kind of shadow is derived directly from the typical periodic structure of a grille. The pattern of horizontal and vertical lines repeats across the window, although its tint changes across the different glass tiles. Figure 11 shows the test image, in which we focus our attention on a horizontal line of 1 pixel width. The luminance profile for pixels on this horizontal line is shown in Figure 12. The grille shadow profile is V-shaped, and its periodicity is clearly evident.

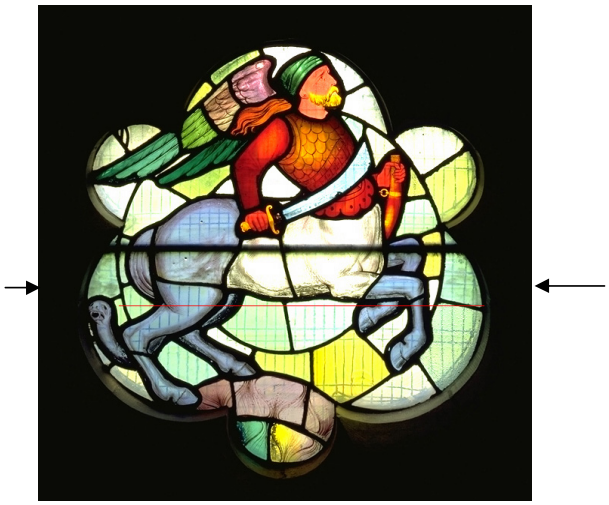

Figure 11. Test image and horizontal sampling line

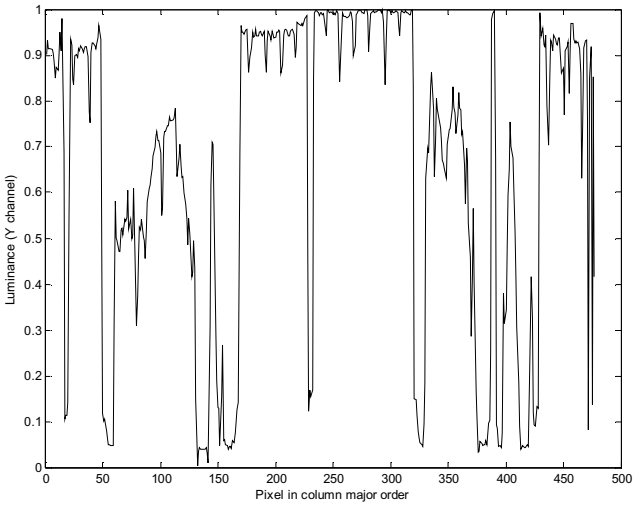

Figure 12. Luminance profile of pixels in horizontal line.

Two convolution windows are implemented (Figure 13) in correspondence with the vertical and horizontal grille lines. The vertical convolution window is the average of column j (say $S_{ij}^{(V)}$) over the averages of the neighbouring columns (say $NS_{ij}^{(V)}$). The quantities are defined on a 5x5 window as shown in Figure 13:

$$S_{ij}^{(V)} = \frac{1}{5}\sum_{k=i-2}^{i+2} f_{kj} \qquad (8)$$

$$NS_{ij}^{(V)} = \frac{1}{20}\sum_{k=i-2}^{i+2}\left[\sum_{l=j-2}^{j-1} f_{kl} + \sum_{l=j+1}^{j+2} f_{kl}\right]$$

where f is the input image.

Similarly we define the horizontal quantities:

$$S_{ij}^{(H)} = \frac{1}{5}\sum_{k=j-2}^{j+2} f_{ik}$$

$$NS_{ij}^{(H)} = \frac{1}{20}\sum_{k=j-2}^{j+2}\left[\sum_{l=i-2}^{i-1} f_{lk} + \sum_{l=i+1}^{i+2} f_{lk}\right] \quad (9)$$

where $S_{ij}^{(H)}$ is the average value of the five neighbouring pixels on row *i,* and $NS_{ij}^{(H)}$ is the average of a rectangular region surrounding the pixel (*i,j*) (i.e. the average of the neighbouring rows)

The following ratios are defined based on equations 9, 10:

Vertical Luminance Ratio:

$$VIR_{ij} = \frac{S_{ij}^{(V)}}{NS_{ij}^{(V)}} \quad (10)$$

Horizontal Luminance Ratio:

$$HIR_{ij} = \frac{S_{ij}^{(H)}}{NS_{ij}^{(H)}} \quad (11)$$

Firstly we consider a test image as shown in Figure 14. The affected grille/bar line calibre size is 7 pixels. The horizontal convolution window is applied to the test image and HIR_{ij} is estimated. One can see that the shadow profile has been successfully detected, but the edges of the shadow profile are affected. Given $\alpha_{ij} \in [0,1]$ in MM model (Eq. 7), we clip the HIR_{ij} value outside this range. One can see that the profile of α looks similar to the true image luminance profile but it is too narrow. We change the window size by increasing the neighbouring row.

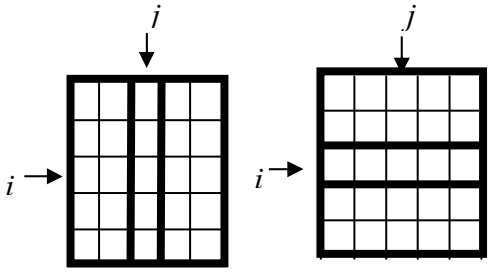

Figure 13. Vertical (left) and horizontal (right) convolution windows

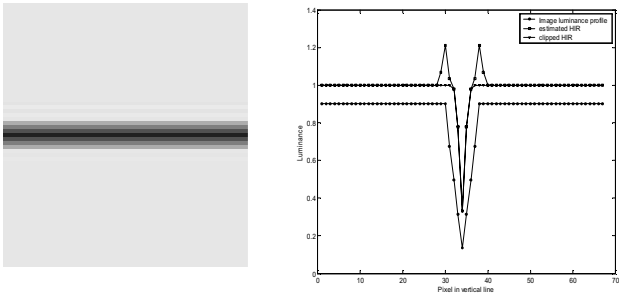

Figure 14. Test image (left); Profile (right) and α value based on 5×5 window

In an ideal case the presence of grille line could be detected by convolving the image with the 1-D convolution windows (Figure 13). Based on this detection, we calculate α by inverting the MM model (Eq. 7). Unfortunately this approach fails because there is another structure that "looks similar" to the shadow, namely the calmes (see Figures 11 and 12).

The calmes are those lead strips holding together the individual tiles of glass forming the SGW. As such, they are not transmissive media, but rather weakly reflective structures, as their metallic nature causes them to reflect the low levels of ambient illumination usually present in cathedrals. The calmes in the image have uniform colour across the whole picture, whereas grille shadows inherit the colour of the glass tile on which they lie. Calmes also usually have much lower intensity values than the grille shadow (Figure 12) and have constant width (calibre) that is different from the grille shadow width. The calmes can be removed from the picture sufficiently well for the purposes of shadow removal by means of a simple thresholding of the colour space. A more thorough separation can be achieved by exploiting the differences in cross-section, with Gabor filters to provide a means to detect line structures of constant width.[7]

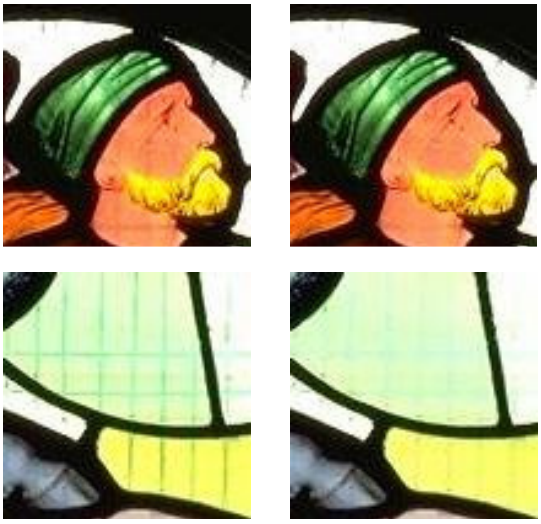

Figure 15. Result of the grille shadow removal algorithm, showing the filtered image (right) compared to the original image (left).

The grille shadow removal algorithm was applied to the test image, after excluding the calmes. Figure 15 shows the results on two details from Figure 11 compared to the original image. Although faint traces of the grille shadow are still perceptible, its visual impact is greatly reduced.

5. Conclusions

Removing shadows from stained glass windows presents specific image processing challenges. We characterised it as a mixture problem, with a true SGW image mixed with a grille/bar image. We firstly developed a simple physically based model to explain shadow formation. Based on this model we removed the bar shadow. To improve this way of removing the shadow needs more real world information such as bar position, its thickness, distance between the bar and glass, glass property, SGW size, position, etc. A more robust method was found for grille removal. The shadowed areas are considered as function of the neighbouring areas. The results are satisfactory, although more generalisations must be sought by exploiting the geometrical structure of the grille in terms of a vertical/horizontal periodic structure. This may be done by using spectral techniques based on the autocorrelation function of a region or on the power distribution in the Fourier transform domain in order to detect grille periodicity.

References

1. L.W. MacDonald and W. Ji, "Colour Characterisation of a High-Resolution Digital Camera", *The First IS&T European Conf. on Colour in Graphics, Imaging and Vision (CGIV)*, 2002.
2. D.F. Rogers, *Illumination and Color in Computer Generated Imagery*, Springer-Verlag, 1989.
3. D.A. Forsyth and J. Ponce, *Computer Vision: A Modern Approach*, Prentice-Hall, 2002.
4. L.W. MacDonald, W. Ji, S. Bouzit, and K. Findlater, "Characterisation of a Digital Camera for Heritage Photography", *Proc. Intl. Congress of Imaging Science (ICIS'02)*, pp.345-346, Tokyo, May 2002.
5. J.D. Foley, A. Van Dam, S.K. Feiner, and J.F. Hughes, *Computer Graphics – Principles and Practice*, 2nd Edition, pp. 590-592, Addison-Wesley, 1990.
6. A. Giani and L.W. MacDonald, "Blending Images of Stained Glass Windows", *Proc. IS&T/SID 11th Color Imaging Conf.*, Scottsdale, 2003.
7. A. Giani, L.W. MacDonald, C. Machy and S. Suganthan, "Image Segmentation of Stained Glass", *IS&T/SPIE 15th Annual Symp. on Electronic Imaging - Color Imaging VIII: Processing, Hardcopy, and Applications*, January 2003, Santa Clara, California.

Biography

Shanmugalingam Suganthan received his B.Sc. Honors degree in Computer Science from University of Jaffna, Sri Lanka in 2001. He also worked in University of Jaffna as an Assistant lecturer for one year. He is currently studying towards a PhD in Image Processing at the Colour and Imaging Institute, University of Derby, UK.