A Physically Based Colour Model

Robert J Oddy
Philip J Willis

Computer Graphics Group,
School of Mathematical Sciences.
University Of Bath,
Claverton Down,
Bath, Avon.
BA2 7AY

Abstract

We propose an intuitively simple way of representing colour which has the additional virtue that it permits mixing and overlaying of transparent and opaque paints to an arbitrary degree. Our approach is related to the earlier alpha channel model used for compositing. It includes this as a special case but has applications in many other areas, especially animation, paint programs and graphics libraries.

Keywords: colour, colour representation, colour manipulation, animation, paint programs, graphics libraries, rendering, raster image processors.

1 Introduction

Within the field of computer graphics, colour is one of the most important resources available and yet the specification of colour is still historically tied to the physical restraints of computer and screen. To specify "full" colour graphics the colour is most commonly indicated with the RGB colour cube notation. Twenty-four bits of data (eight for each primary) are used to represent all distinguishable colours within the colour cube. Note that the range of colours provided will be some subset of all possible colours (as specified by the CIE colour diagram), depending upon the actual primaries used as the orthogonal axes of the colour cube. This twenty-four bit model is accepted as being a suitably compact and accurate representation for colour output.

Much work on the limitations of colour output has been completed, describing the disparities between CIE and screen (or printer) gamuts and how to transfer images correctly between devices. More recently, investigations have studied the differences between colour as displayed and as perceived [Hall 1989, Pratt 1978]. Many studies are concerned with colour science [Wyszecki & Stiles 1982], or consider light and colour to be synonomous. Our approach is to treat colour as an *input* parameter helping to describe a picture, with light as an output created by the display device. In other words we separate modelling with colour, from rendering to produce light.

For this approach to be workable, users require a simple paradigm by which they can imagine, for example, how layers of colour combine, so that images can be built up in an intuitive manner. This enables pictures to be painted with *materials* rather than light and it is this concept of a colour material that our approach explores.

In [Porter & Duff 1984] it was suggested that an extended colour model involving a fourth channel - the alpha channel - could be used to hold matte information about an image. This additional channel allows images from various sources (or from various programs) to be combined into single new 'composite' images. The paper describes an elegant algebra for the compositing of images, providing a method for automatic edge anti-aliasing techniques. The authors also claimed that the alpha channel model could be used to simulate the opacity of certain objects within the images. The ability to represent opacity was very much a secondary issue, and lacks the rigorous treatment that was given to the other aspects of the model. Our paper defines how certain physical properties, including opacity, can be modelled as part of the colour specification and presents an algebra for use with this new model.

2 A Physical Model

To enable a straightforward visualisation of colour combination, it is important to present colour in a readily understood form. Our approach is to use an idealised physical analogy, which we believe both to be more tractable to explanation and to be more useful in its own right than earlier methods.

To explain how our model works, we first examine two extremes of what is commonly understood as the colour of an object. We then consider how such colours may be combined to give a general representation.

2.1 Colour From Pigments

An opaque object is one which we view entirely by scattered light. We may imagine that light originates from an ambient light source and that this light impinges on the surface of the coloured material. The surface has the property that it absorbs certain wavelengths and scatters those that remain. The eye sees only the scattered light and it is this that we informally refer to as the colour of the object. Of course the perceived colour also depends on the colour of the light source (as well as on subtler environmental factors that shall not be considered).

In order to give this opacity a physical reality, we imagine that a material is heavily loaded with microscopic pigmented *particles*. Each particle has the same colour and the particles are so densely packed that no light can pass through the material (we may also imagine an equivalent arrangement in which particles are constrained to be of one of three primary colours and their relative density determines the scattered colour).

2.2 Colour From Filters

At the other extreme we have diaphanous objects, such as coloured glass. In this case the ambient light passes through the object and emerges from the other side. As with opaque objects some frequencies are absorbed by the material, but now the remaining light is transmitted rather than being scattered. Obviously the light source and viewer must be on opposite sides of the object, whereas for pigmented materials they will be on the same side of the object. As with opaque objects we model filters as a material. This material will have no pigment particles at all but consists entirely of a homogeneous *medium* of the required colour.

2.3 Combining Pigments And Filters

Our proposal is that, in general, a colour should be specified as a material, in such a way that it permits either of these extremes or any combination in between. To do this we need to define medium and particle colours for the material. We will also need a factor that will determine the relative proportion of particles to medium. When the factor is one the material is completely opaque and scatters light according to the colour of the particles. When the factor is zero the material is transparent and transmits light according to the colour of the medium. For intermediate values the material is translucent and exhibits a mix of these properties (see Figure 1). The mix of particles and medium would realistically be random, following some statistical distribution. However for convenience we may assume that all the particles are massed together as shown, giving a simple split between particles and medium.

We imply that a general material is heterogeneous and consists of a filter-like medium in which are suspended some pigment particles. In general the particles will not be so dense as to prevent all light from passing through the material, but neither will they be so sparse that no light is scattered back. Indeed our model can perhaps be thought of as corresponding to a thin layer of paint, containing a mixture of particles and a diaphanous medium.

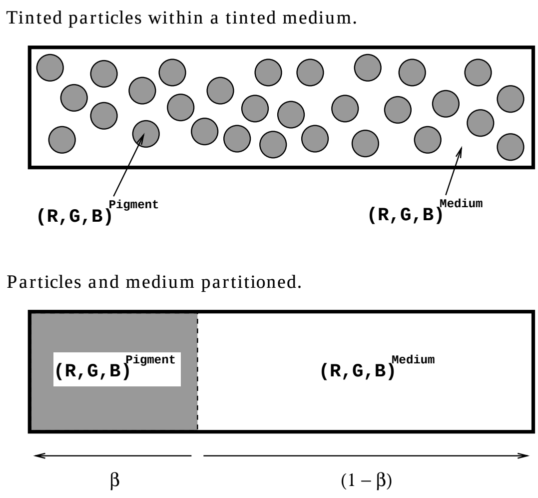

Figure 1: Particles and Medium

3 Implementation: The Beta Model

We will take the alpha channel model as a starting point and then refine it. It is possible to explain the alpha channel model in terms of particles and medium. With the alpha channel, any pixel is represented by four values (R, G, B, α) corresponding to red, green, blue and alpha. Alpha is a measure of coverage for that pixel. In other words, if alpha has the value 0.6, it is assumed that only 0.6 of the pixel is covered and the rest is clear. This allows any underlying layers to show through, weighted by 0.4. In our terms, alpha is a measure of the density of coloured particles in a colourless medium.

3.1 The Basic Beta Model

Now suppose that we redefine the alpha channel and, to clarify the following discussion, we will call this new variant the beta channel. Beta defines the *opacity* of the material. If beta is one then the material is opaque and can only scatter light. If beta is zero then all light passes through the material and it acts as a filter. In either case (R, G, B) gives the brightness and purity of the colour. Values of beta between zero and one correspond to translucency and give a directly proportional effect. With this definition we have a **four** channel method of *modelling* colour (this is an *input* or *object* representation as opposed to an *output* or *pixel* representation) which allows us to render a picture by performing the corresponding mathematics on a data structure representing the picture. Thus we can imagine, as an example, an area of opaque material $(\beta = 1)$ that gradually becomes more translucent as beta decreases, revealing a colour-tinged image of any underlying picture.

3.2 The Generalised Beta Model

Now suppose that we wish the particles to be a different colour to the medium. To accomodate this extra degree of freedom, we must allow separate specification for particles and medium. Thus materials in the beta model are fully defined by:

$$(particle\_colour, \beta, medium\_colour) \Rightarrow (R_P, G_P, B_P, \beta, R_M, G_M, B_M)$$

This provides a **seven** channel model of colour. It will be shown later that all seven channels are required to support a mathematically sound basis for a colour combination algebra. Additionally this enables all of the attributes of the alpha channel model to be simulated as a special case, when the medium is set to be colourless. A summary of the alpha, basic beta and generalised beta models is shown in Figure 2

4 Beta Channel Operations

Given that we have defined a model based on an abstract paint material, it is now appropriate to consider ways of combining these paints. We will consider two linear operations, namely: mixing and over-painting. We have called these operations *plus* and *over* respectively. These are given as useful basic functions, although other operations, whether linear or not, may also be treated in the same manner.

In what follows, we have made the following simplifications. Firstly we are working with equal quantities of paints. This is by no means essential and the equations we deduce are easily extended to allow for unequal composition: we omit the weighting factors for clarity. Secondly, where we use the word "Medium", we mean the fraction of ambient light energy that the paint transmits. In a practical implementation therefore, this means one of the three (R_M, G_M, B_M) normalised primaries and the equations will need to be applied three times, once for each primary. An equivalent comment applies to the word "Particle".

4.1 Plus

Plus is the operation which corresponds to taking two paints and mixing them together. Noting that if β is the particle density then $(1 - \beta)$ is the medium density, we may perform this operation as follows:

First we consider the pigment particles. The resultant colour will be the weighted sum of the contributions from the pigments in each paint. We have $\beta_1 \cdot Particle_1$ from the first paint and $\beta_2 \cdot Particle_2$ from the second paint and so the normalization factor is $(\beta_1 + \beta_2)$.

Next we consider the medium. We treat this in exactly the same way but of course the weighting factors and colours are different. We have $(1 - \beta_1) \cdot Medium_1$ from the first paint, $(1 - \beta_2) \cdot Medium_2$ from the second paint and the normalization factor is thus

$$(1 - \beta_1) + (1 - \beta_2) = 2 - (\beta_1 + \beta_2).$$

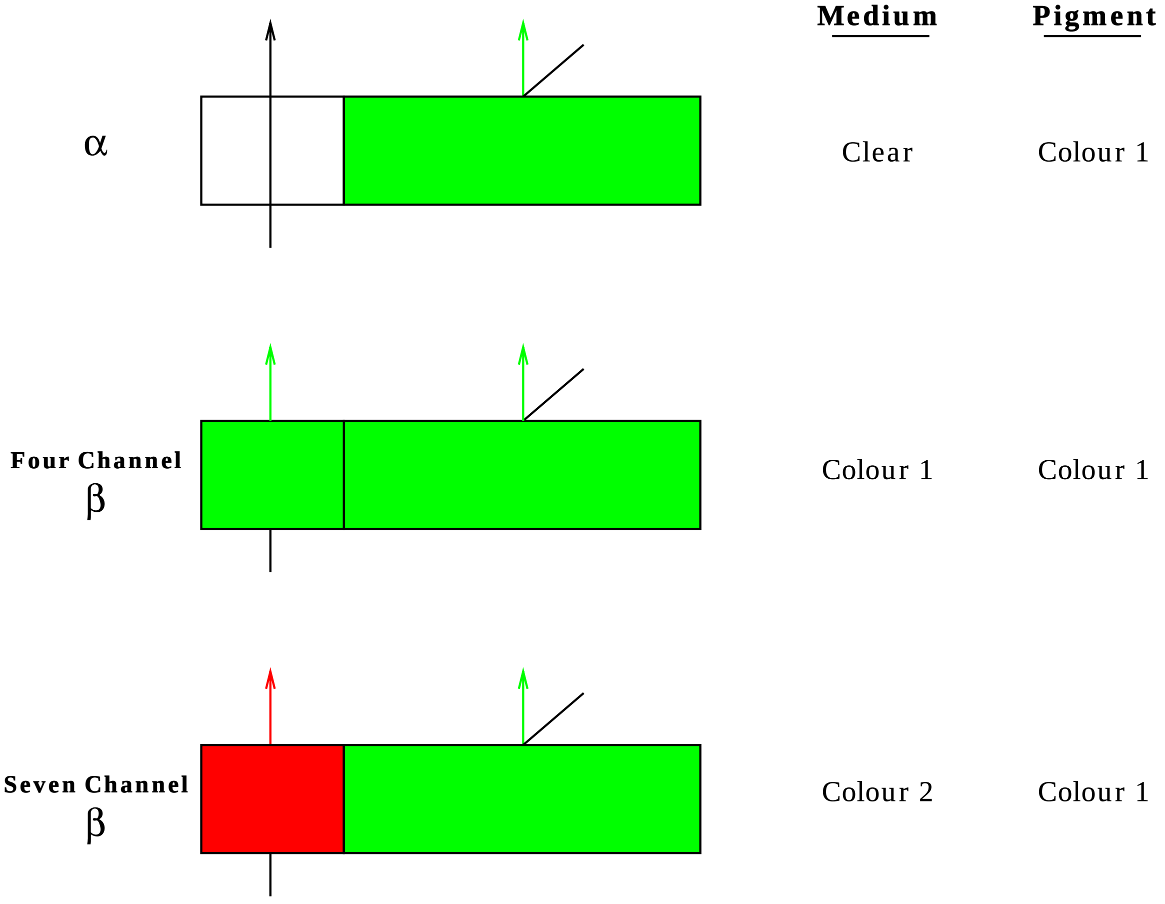

Figure 2: Colour Models

Finally we consider beta. Beta is defined to be the fraction of the paint which is pigment particles and so the new beta is the average of β_1 and β_2. These considerations can be summarised thus:

Definition 4.1

$$Paint_1 \textbf{ plus } Paint_2 = \begin{cases} Particle_{plus} & = \frac{\beta_1 Particle_1 + \beta_2 Particle_2}{\beta_1 + \beta_2} \\ Medium_{plus} & = \frac{(1-\beta_1) Medium_1 + (1-\beta_2) Medium_2}{2-(\beta_1+\beta_2)} \\ \beta_{plus} & = \frac{\beta_1+\beta_2}{2} \end{cases}$$

This definition incorporates normalisation. This is part of the nature of paint mixing as we have defined it.

4.2 Over

Over is the operation in which one layer of (possibly translucent) paint sits on top of another, without mixing taking place. We assume that $Paint_1$ overlays $Paint_2$.

First we consider the pigment particles. Those in $Paint_1$ are directly visible to the viewer and thus contribute $\beta_1 \cdot Particle_1$. Some of those in $Paint_2$ are hidden by the pigment in $Paint_1$. In fact a proportion β_1 will be hidden and a proportion $(1-\beta_1)$ will be visible; that is $(1-\beta_1)\cdot\beta_2$ will be visible. These can only be seen through the medium of $Paint_1$, that is by light filtered by passing through $Medium_1$, scattering from $Particle_2$ and then passing again through $Medium_1$ on the way back to the viewer.

Next we consider the medium. Light which passes through both layers will have been attenuated by both mediums. The resulting attenuation is thus described by the product of the two attenuation coefficients $Medium_1$ and $Medium_2$.

Finally we consider beta. We note that this will be at least β_1, because the top layer $Paint_1$ consists of this proportion. We have already established that $(1-\beta_1)\cdot\beta_2$ particles of $Paint_2$ are visible and so the new beta is the sum of these.

In summary:

Definition 4.2

$$Paint_1 \textbf{ over } Paint_2 = \begin{cases} Particle_{over} & = \beta_1 Particle_1 + \\ & \quad Medium_1 * (1-\beta_1)\beta_2 Particle_2 * Medium_1 \\ Medium_{over} & = Medium_1 * Medium_2 \\ \beta_{over} & = \beta_1 + (1-\beta_1)\beta_2 \end{cases}$$

It is a consequence of definition 4.2 that, when four-channel paints are overlayed, the particles and medium of the result will in general have differing colours, and so the result is no longer representable with only four channels. Therefore the four channel model is not closed when *over* is applied. Its virtue is as a shorthand method of representing those paints with particles and medium of the same colour.

The closure property is important in the alpha channel model, as it allows many different images to be composited together. It is however *essential* in the beta model, since any paint may be used in combination with other paints. A closed algebra ensures that an entire sandwich of paints is equivalent to a single paint, and so we adopt the seven channel beta model.

Further to this it can be seen that normalisation is not included, unlike the resulting paint from a similar *plus* operation. This arises due to the additive properties that *over* is designed to reflect.

4.3 Lighting

One of the advantages of employing the beta model is to allow use of front and back illumination. If we assume that we wish to generate output for a conventional frame buffer, then the lighting equation for each primary in a paint layer is:

$$Pixel = \beta Particle * Light_{fore} + (1-\beta) Medium * Light_{back}$$

where $Light_{fore}$ and $Light_{back}$ specify the illumination colour. Note that the backlight that passes through the medium is attentuated by a factor of $(1-\beta)$, as a proportion β of this is blocked by the particles within the paint. Similarly, only β of the front illumination will be scattered back from the paint.

5 Practical Considerations

It is important to realise that we are proposing this scheme in a context where values are manipulated as part of a rendering process. Our own concern relates to the overlaying of painted cels, as used in an animation system. With the proposed approach, a combination of back-lighting and front-lighting can be used, and indeed varied, to give animated effects which are difficult by any other method. The calculation of the resulting image as an array of pixels we refer to as rendering and the colour outputs from this process are sent to a pixel array (a frame buffer).

Rendering is thus the means of using object (input) colours and geometry to calculate the required pixel (output) colours and their spatial distribution in the buffer. Our proposal in intended to be used for object colours, in contrast to the pixel operations of the alpha channel model.

For computational reasons similar to those put forward for the alpha channel model, each of the (R, G, B) particle values can be pre-multiplied by beta. Then, given a material with opacity beta, the particle colour is stored as $(\beta R, \beta G, \beta B)$. Similar considerations can be made with medium values.

Finally we remind the reader that we are in essence calculating a new paint from a combination of two others; that any pair of combined paints can be replaced with a single paint because of the closure property; that this operation may be continued indefinitely; and that the resulting paint is a transmission-reflection coefficient to be used with back and front lighting in order to produce the final output colours.

6 Applications

In this section we will show the usefulness of the beta model in three applications. The pictures given in the following examples use colour values that have been calculated employing the formulae from Definitions 4.1 and 4.2.

6.1 Animation

Classic works on traditional animation techniques [Thomas & Johnstone 1981] imply a need for translucency, transparency, and filtering. We note that the alpha channel model has no concept of a filter. Conventional two-dimensional animation is constructed by overlaying a number of hand-drawn acetate sheets, or cels, to build up a composite image. These sheets are then illuminated and a frame of film exposed. A new arrangement of cels is then created and another exposure made. In this way the complete sequence is built up.

We are concerned in this paper with colour rather than movement, so our example is one in which the geometry is fixed but where the lighting is varied with time. In Figure 3, we depict two versions of a simple house. The left hand version is in daylight and so the frontal illumination is bright white. Each of the sky, ground, wall and roof are represented as fully opaque and thus appear in their natural colours. The window areas are represented as almost totally transparent, with a neutral medium. They thus appear black because there are no particles to scatter the front illumination.

When the night time colours were computed, the arrangement was somewhat different. Here the front illumination was a dim blue-grey, to give a moonlit effect. The lighting can be seen to be dim is apparent from the overall dark appearance, whilst the blue bias is most obvious on the neutral grey walls. Moreover we have added a strong yellow backlight to give the effect of light within the house. This shines through the transparent windows and makes the windows much brighter than anything else in the scene.

6.2 Paint And Retouch Systems

Artists have striven to produce images containing translucent objects, whether it be through the use of watercolours or via air-brushes. Clearly the artist is aiming to tint the background scene with a secondary colour. Essentially colours are opaque, or act like filters, or perform some role inbetween.

Figure 3: Daylit/Nightlit House

Our next example shows the graphical capability of the colour model. Figure 4 is a simple example of two surfaces at right angles, which should be interpreted as a vertical wall and horizontal ground. At the top left we have placed a white "light", which is assumed to be partly shaded so that the full strength of it forms a pool of light on the two surfaces. The rest of the scene is much less strongly lit. The basic geometry of the picture has been created by hand, using a paint program. The colouring is however calculated using the beta model.

First we assume that the frontal light is bright white and use this to calculate the brightest areas of the scene. Next we assume that a neutral density filter (no pigment, grey medium) covers the darker parts of the scene and calculate those colours accordingly. The contrast is clearly visible.

In Figure 5 we have covered the entire scene with an additional layer of mist. This is represented as a clear medium containing a relatively low density of white particles. The effect on the picture is quite marked. Firstly, the black background becomes fogged with grey. Secondly, the brightly illuminated mist gives a much better impression of light being thrown in a particular direction. Finally the colours of the wall and ground are diluted because of the presence of the white mist. The dilution effects both bright and dark areas in the correct proportion, increasing the perception that part of the scene is brightly lit.

6.3 Compositing

As a final example, suppose that a black and white portrait photograph has been digitised into a pixel array. We may use this as the basis of a composition designed to resemble an Edwardian vignette. First we overlay a filter (no particles, brown-tinged medium) which will produce a sepia effect. Next we construct a soft mask. We arrange that it has a central disk which is entirely clear, but that around this it becomes progressively more an opaque and white (i.e. increasing the amounts of white particles in a clear medium). Applying this mask we achieve a gradual fading of the sepia portrait into a white background, in the desired manner.

It is worth noting at this point that the sepia filter and the mask can be combined to produce a single *filter-mask* which produces the same net effect and that this can be retained for use with other photographs. This combined paint layer will have sepia medium everywhere and white pigment in the appropriate distribution.

7 Concluding Remarks

Throughout this paper we have assumed that colours are defined within an (R, G, B) coordinate system and that the calculations on these coordinates are entirely independent. These conditions can be varied in order to achieve the maximum benefit from the intuitive basis of our approach.

For simplicity we have so far assumed that *Particle* and *Medium* are essentially scalar quantities ie. the red, green and blue channels are independently calculated. In fact these items can usefully be *any* function of the three input channels. For example, if they are matrices then the output red (say) can

Figure 4: Spotlight Scene

Figure 5: Spotlight Scene With Mist

be any linear weighted sum of the three input channels. Consider the following matrix equation.

$$[R'G'B'] = [RGB] \begin{bmatrix} 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 \end{bmatrix}$$

This has the property of producing a medium (or particles) which transmit (or reflect) a monochrome image of the full colour input. This is conceptually useful (even if *physically* difficult to realise). Alternatively it would not be difficult to implement such a filter using scalars, if the output functions were defined with parameters in an $(Intensity, Hue, Saturation)$ space.

Matrices may also be used to simulate fluorescence, in which the incoming broad-spectrum light energy is concentrated into a limited range of the spectrum, and other properties of this general type. At this level the operations move from graphics generation to image processing.

In order to enable matrix functionality to be included into the model, Definition 4.2 shows the term

$$Particle_{over} = \beta_1 Particle_1 + Medium_1 * (1 - \beta_1)\beta_2 Particle_2 * Medium_1$$

instead of

$$Particle_{over} = \beta_1 Particle_1 + \beta_2 Particle_2 * (1 - \beta_1) Medium_1^2$$

Our main aim in this paper has been to show that, given a careful definition of the paint material that we employ, it is possible to combine such paints (to an arbitrary degree) by mixing and overlaying and in such a manner that can be both easily imagined and simply rendered. The closure property is central to this and guarantees that the user can safely ignore the mathematical details.

It is interesting to note how differing applications affect the method for provision of similar facilities. Transparency is quite common in ray-traced three-dimensional scenes but far less so in two dimensional pictures where the emphasis is on colour as an opaque pigment. This paper has combined these two aspects of colour into a single strand. In doing so we hope we have produced a way of modelling with colour which is substantially more general than the direct (R, G, B) manipulation in current use. In particular we believe that our model is a better way of thinking about colour and colour-related effects. There is a strong intuitive feel in the way it is used.

All of our examples use two dimensional models. The resulting two dimensional graphics, which we consider to be a relatively neglected area of research, serve to show that this way of modelling colour has a wide variety of applications. Indeed, as shown in Figures 4 and 5 it is perfectly possible to produce images of apparently three dimensional scenes. The use of layers of picture, with appropriate translucency and filtering effects, seems to be an especially effective way of building up a composition. While this is not itself original and also has an obvious analogy with cel animation, we believe that this two and a half dimensional approach has much to commend it for graphics in general. Finally, our method gives an emphasis to "graphics" (the visual effect) rather than to "computer" (the means), which must surely be a step forward.

8 Acknowledgements

The authors of this paper wish to thank the Science and Engineering Research Council for supplying the Research Studentship which supports R.J.Oddy.

References

[Hall 1989] Roy Hall. *Illumination and Color in Computer Generated Imagery.* Springer-Verlag, 1989. ISBN 0-387-96774-5.

[Porter & Duff 1984] Thomas Porter and Tom Duff. Compositing digital images. *Computer Graphics*, **18**(3):253–259, July 1984. Siggraph (ACM), 11 West 42nd Street, New York, NY 10036.

[Pratt 1978] W.K. Pratt. *Digital Image Processing.* Wiley, 1978.

[Thomas & Johnstone 1981] Frank Thomas and Ollie Johnstone. *Disney Animation: The Illusion of Life.* Abbeville, Walt Disney Productions, Burbank, California, 1981. ISBN 0-89659-233-2.

[Wyszecki & Stiles 1982] G. Wyszecki and W.S. Stiles. *Color Science.* Wiley, 2nd edition, 1982.