Colour is a Medium as well as a Message [1]

Robert J. Woodham [2]

1 Introduction

My research is in "physics-based" vision (Wolff et al. 1992; Healey and Jain 1994). Physics-based vision takes seriously radiometric, as well as geometric, constraints in image analysis. The approach has led to new computational methods to determine scene parameters from images including photometric stereo (Woodham 1980; Woodham 1981) and multiple light source optical flow (Woodham 1990). The idea exploited is that multiple images acquired under the identical viewing geometry but different conditions of illumination are a principled way to obtain additional local radiometric constraint.

The work is both theoretical and experimental. Theoretical results are implemented and tested in real systems. Recently, advances in imaging technology have led to the development of near real-time implementations (Woodham 1994; Siegerist 1996). These implementations, in turn, support further experimentation. Principal applications are to robot vision and to remote sensing.

No connection to colour perception and biological vision is evident. Indeed, David Marr's book (1982, 249) introduced the topic of photometric stereo as follows "Finally, there is a technique for recovering shape from reflectance maps that cannot possibly have any biological significance, but which is so elegant that I cannot resist mentioning it."

I am fond of this quote, especially the word "elegant." At the same time, the conclusion that photometric stereo "cannot possibly have any biological significance" always struck me as too strong. One could accept, on empirical grounds, a final determination that photometric stereo does not have any biological significance. But, robot and biological vision systems both function in the identical physical world. One must always allow for the possible existence of biological systems that use a method successfully employed in robot vision (and vice versa).

Photometric stereo was new at the time of the writing and the subsequent posthumous publication of Marr's book (1982). Much has happened in the years that have followed. This chapter is a retrospective. I do not claim to demonstrate that photometric stereo has biological significance. For the purpose at hand, however, it is sufficient to suggest ways that it might. This is consistent with the inter-disciplinary nature of these conference proceedings. It also is a way to pay homage to David Marr who saw the process of scientific discovery as forever new and on-going.

The objective of this chapter is to document another way to use colour in computational vision and to suggest possible connections to biological vision. The hope is that this will inspire new collaboration between researchers in biological and robot vision. The remainder of Section 1 develops the case that colour is a medium as well as a message. The basics of photometric stereo and multiple light source optical flow are described, including the role the medium of colour plays in current implementations. Section 2 provides an overview of spectral radiometry and the issues that arise in robot vision and in remote sensing. Section 3 bridges to biological systems with particular emphasis on colour vision in aquatic environments. Section 4 concludes with a discussion of implications for colour vision research.

---

[1] Preprint of paper that appeared as, R. J. Woodham, "Colour is a medium as well as a message," in *Colour Perception: Philosophical, Psychological, Artistic and Computational Perspectives* (S. Davis, ed.), vol. 9 of *Vancouver Studies in Cognitive Science*, pp. 117–140, Oxford University Press, 2000.

[2] Laboratory for Computational Intelligence, Department of Computer Science, University of British Columbia, Vancouver, BC, Canada V6T 1Z4. email: woodham@cs.ubc.ca url: `http://www.cs.ubc.ca/~woodham/`

1.1 Colour as Message

1.1.1 Spectral Reflectance

We normally think that the purpose of colour vision is to determine the spectral reflectance of objects. That is, spectral reflectance is the message. Section 2 is more precise about spectral radiometry. For the moment, it suffices to note that spectral reflectance depends upon the intrinsic properties of the surface material and is defined independently of the conditions of illumination and viewing.

This creates two problems for colour perception. First, colour measurement depends not only on the spectral reflectance of objects but also on the spectral characteristics of the illumination. That is, colour the medium confounds colour the message. The ability of humans to perceive an object's colour as constant over a range of illumination conditions is called "colour constancy." Accounting for colour constancy in both robot and biological vision systems remains a significant challenge.

Second, as a physical quantity, spectral reflectance is defined for a continuum of wavelengths. Any colour system that relies on a finite number of channels, regardless of the wavelength response of those channels, will induce colour metamers. That is, there will be different spectral reflectances that, for a given illumination, produce identical, therefore indistinguishable, responses in each of the channels. Conversely, a given spectral reflectance may well be perceived differently by vision systems whose channels differ in number or in response characteristics.

1.1.2 Aerial Perspective

Forests and mountains in the distance appear bluer and hazier. This colour shift and decrease in contrast is called aerial perspective. It is caused by the scattering of light in the air between the viewer and a distant target. Perception of distance is influenced by aerial perspective. Indeed, the degree of bluishness of a dark target provides a direct estimate of its distance from the viewer. In aerial perspective, the message of colour is distance.

1.2 Colour as Medium

1.2.1 Anaglyphic Stereo

Our two eyes acquire images of the world from slightly different viewpoints. We can perceive depth based on differences in the relative position of points in the left image and in the right image as a function of the distance of that point from the viewer. The ability to perceive depth in this way is called stereopsis.

Anaglyphic stereo is an example of colour as medium likely familiar to those with an interest in colour perception. It is a 3D image presentation technique. Two achromatic images forming a binocular stereo pair are displayed in complementary colours. By convention, the image intended for the left eye is red and the image intended for the right eye is cyan (i.e., blue plus green). When viewed through red (left eye) and cyan (right eye) filtered lenses, stereopsis is achieved. Perception is 3D and achromatic. In anaglyphic stereo, colour is not the message. Rather, it is the medium used to carry the binocular stereo data resolved by human stereopsis.

1.2.2 Photometric Stereo

In binocular stereo, if one knows the correspondence between a point in the left image and the matching point in the right image then computation of depth is a simple geometric triangulation. Unfortunately, determining point correspondences is a major computational task.

Photometric stereo uses multiple images acquired from a single, fixed viewpoint. The advantage is that the need to determine point correspondences is avoided completely. The disadvantage is that multiple images, per se, are not of value since, except for measurement noise, each would be identical. Something must be varied in order to obtain useful additional information.

The idea of photometric stereo is that one can vary the illumination. For example, suppose two images are obtained from the same viewpoint but under different conditions of illumination. In general, the two brightness values measured at a given image point will differ. Providing a theoretical account for this difference requires excursion into principles of optics. Indeed, this is exactly what is done in physics-based vision. What emerges is that the brightness values measured allow one to estimate local surface orientation.

There are a variety of ways to represent surface orientation mathematically, including the (unit) surface normal vector and the gradient. The essential observation is that surface orientation has two degrees of freedom. Thus, in principle, images acquired under two conditions of illumination are sufficient to estimate surface orientation locally. In general, however, the underlying equations are non-linear so that, in a two image case, the solution is not unique. It has proven advantageous to implement photometric stereo using three (or more) different conditions of illumination. This makes the local solution unique. It also protects against local degeneracies and improves accuracy in the presence of measurement noise.

The first implementation of photometric stereo (Silver 1980) obtained multiple images by turning light sources on and off. This was the implementation familiar to Marr (1982). More recently, spectral multiplexing has been used. No light sources are turned on and off. Instead, three spectrally distinct light sources continuously illuminate the target objects from three different directions. A suitable RGB color camera acquires three-channel video images that subsequently are treated as three separate B&W images, one corresponding to each condition of illumination. A near real-time (15Hz) implementation for full frame, 512×480, RGB video has been demonstrated (Woodham 1994).

Photometric stereo does need to know the surface reflectance. Indeed, current variants of photometric stereo differ primarily in the way in which reflectance is modeled. My recent implementations model reflectance empirically using measurements obtained from a calibration sphere (Woodham 1994). The empirical approach has the benefit of automatically compensating for the transfer characteristics of the sensor. Calibration results are applicable to the analysis of other objects of different shape but made of the same material as the calibration sphere and illuminated and viewed under the same conditions. In this way, a material with any reflectance characteristic can be handled, provided that the necessary calibration can be done. In some applications, it has been useful to use paint (or other coating) to match reflectance properties between a calibration sphere and the objects to be analyzed.

Photometric stereo provides a direct measurement of surface orientation. The information obtained is about 3D shape, not about surface reflectance. In the spectral multiplexing implementation, colour is not the message. Rather, it is the medium actively used to obtain three independent images simultaneously under three different conditions of illumination.

1.2.3 Multiple Light Source Optical Flow

A point on a moving object will, in general, cause motion at the corresponding image point. Motion is geometric. A camera does not measure motion directly. Rather, it records a time varying pattern of brightness (and colour). The relation between motion and brightness change is not purely geometric because it also depends on radiometric factors, including the illumination and the reflectance properties of the objects in view. In physics-based vision, the apparent motion of the brightness field itself is called the “optical flow.” Optical flow is radiometric.

It is important to note that motion and optical flow do not always coincide. There can be optical flow without motion, and vice versa. Dynamically changing illumination produces optical flow even in the absence of motion. Conversely, a rotating sphere, made of a uniform material without distinct surface markings, is an example of motion without optical flow, since the sphere is self-similar at all speeds of rotation.

Something is needed to link radiometry and geometry. To proceed, it is helpful to introduce some mathematical notation. Let the time varying image brightness be $E(x, y, t)$, where x and y are spatial variables denoting image position and t is time. We measure $E(x, y, t)$ and its partial derivatives E_x, E_y and E_t. This does not determine the 2D motion, expressed as the two total derivatives

$$u = \frac{dx}{dt} \quad \text{and} \quad v = \frac{dy}{dt}$$

Relating partial and total derivatives requires a "conservation law" that leads to a constraint equation. In the standard formulation of optical flow (Horn and Schunck 1981), radiometry and geometry are linked by an assumption of conservation of brightness, namely

$$\frac{dE}{dt} = 0$$

which leads to the constraint equation,

$$E_x\, u \;+\; E_y\, v \;+\; E_t \;=\; 0. \tag{1}$$

This is one equation in the two unknowns u and v. Motion still can not be determined locally owing to what is called the "aperture problem" (Marr 1982, 165–66). Equation (1) determines the component of motion in the direction $[E_x, E_y]$ but provides no constraint on the component of motion in the orthogonal direction (i.e., in the direction along the iso-brightness contour).

For surfaces that exhibit smooth shading, the direction of illumination plays a significant role in determining $[E_x, E_y]$. Thus, it is possible to obtain additional, independent information using brightness values recorded from images acquired simultaneously under different conditions of illumination. Based on this idea, a method has been developed to compute a dense, local representation of optical flow (Woodham 1990).

Because Equation (1) is linear in u and v, two illumination directions are, in principle, sufficient to determine motion locally. As with photometric stereo, it has proven advantageous to use three different conditions of illumination. This overdetermines the solution locally, again avoiding local degeneracies and making the computed solution more robust. More importantly, overdetermining the solution allows one to identify locations in the image where optical flow and motion do (or do not) coincide. The method is not misled, for example, by optical flow arising from moving shadows or from bright objects that, while moving, act as secondary light sources to other parts of the scene.

Recently, a near real-time (3–5Hz) implementation has been demonstrated (Siegerist 1996). In this implementation, spectral multiplexing plays a critical role since it is essential that the three images be acquired simultaneously. Because the method is entirely local, it can determine the point-wise motion of non-rigid objects. This has been demonstrated using the example of an expanding and contracting balloon (Siegerist 1996).

1.3 An Experimental Environment

A calibrated imaging facility (CIF) has been built to control both scene parameters and conditions of imaging. It uses an optical bench and computer actuated motion stages to control the position and motion of cameras, light sources and test objects. Specialized imaging equipment includes a Sony DXC-755 3 CCD 24-bit RGB camera and three Newport MP-1000 Moire (white light) projectors with associated lenses and spectral filters. The light sources are DC powered to avoid fluctuations in output associated with 60Hz AC. The camera, light sources and motion stages are integrated with other vision and robotics equipment in the UBC Laboratory for Computational Intelligence (LCI), including Datacube DigiColor and MaxVideo-200 image processing systems and a parallel network of Texas Instruments C40's.

With appropriate filtration, the three white light projectors become spectrally distinct: one red, one green and one blue. The color separation filters used are the Newport FS-225 set. The filters are manufactured by Corion, Corp. and carry Corion part numbers CA500 (blue), CA550 (green) and CA600 (red). The precise spectral response of each filter has been measured by the manufacturer. There is negligible overlap in the visible spectrum between the red light source and either the green light source or the blue light source. There is a small overlap between the green and the blue light sources for wavelengths in the 500–520nm range.

Much like stage lighting in a theatre production, the three projectors illuminate the work space from three different directions. In a standard configuration, one projector is aligned with the viewing direction, a second illuminates from the upper left and the third illuminates from the upper right. The Sony camera acquires three separate 8 bit-per-pixel images, one each for R, G and B. The goal of spectral multiplexing is to have the camera's R, G and B channels influenced only by illumination from the red light source, green light source and blue light source respectively. When this is true, the three 8-bit RGB channels can be interpreted as three independent achromatic images acquired under different conditions of illumination.

Successful spectral multiplexing also depends on the spectral response of the camera. Even if the light sources themselves are spectrally distinct, overlap will occur if the spectral response of any one of the camera's R, G or B channels is broad enough to be influenced by more than one light source. The precise spectral response of the camera was not provided by the manufacturer. The camera chosen does split the incoming beam into three distinct RGB components. Experiments have determined that the expected spectral separation is attained. Woodham (1994, 3056–57) provides further details of the experimental environment.

The experimental environment supports the work on photometric stereo and multiple light source optical flow described in Sections 1.2.2 and 1.2.3. In addition, the ability to record an identical scene sequence under different lighting conditions allows researchers to test whether algorithms truly are invariant to the illumination and to other radiometric factors. For example, the CIF has been used to record calibrated motion sequences as part of work on the performance of optical flow techniques (Barron et al. 1994).

2 Overview of Spectral Radiometry

2.1 Modeling the Physics of Imaging

A given spatial arrangement of objects, made of given materials, illuminated in a given way, and viewed from a given vantage point, determine an image according to the laws of optics. Geometric equations determine where each point on a visible surface appears in the image and corresponding radiometric equations determine its brightness and color.

A complete treatment of image formation must take into account a variety of factors including: the intrinsic optical properties of the surface material; the surface microstructure; the surface shape; the position of the sensor; the spatial and spectral distribution and state of polarization of the illumination; and the optical properties of the medium through which the radiant energy is transmitted.

2.1.1 Defining Reflectance

Confusion arises when different disciplines use different terminology to define reflectance. Increasingly, researchers in both computer vision and graphics have adopted a standard nomenclature for reflectance (Nicodemus et al. 1977). The reflectance of an opaque material is specified by its bidirectional reflectance distribution function (BRDF), denoted symbolically as f_r. As defined,

$$f_r(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{dL_r(\theta_i, \phi_i; \theta_r, \phi_r; E_i)}{dE_i(\theta_i, \phi_i)} \qquad [\mathrm{sr}^{-1}]$$

Directions are given in spherical coordinates, (θ, ϕ), defined with respect to the local tangent plane and surface normal. Subscript i denotes quantities associated with the incident radiant flux. Subscript r denotes quantities associated with the reflected radiant flux. E_i is incident irradiance, L_r is reflected radiance and d indicates a differential quantity. The BRDF is a derivative function, relating the irradiance incident from one direction to its contribution to the reflected radiance in another direction. Units are per steradian where steradian is the standard measure of solid angle.

This is not the place to provide a tutorial on the BRDF and its use in computer vision and graphics. It should be noted, however, that the BRDF is not defined simply as the fraction of the incident radiant flux that is reflected. The fraction reflected depends on the incident direction, (θ_i, ϕ_i). More importantly for vision, we are interested not in the total fraction reflected but in the fraction reflected in a specific direction, namely that of the viewer, (θ_r, ϕ_r). The BRDF is a function of four variables, two each for the directions of the incident and the reflected rays.

2.1.2 Spectral Considerations

Selective reflection can alter the spectral distribution of the reflected beam. Spectral reflectance considers the dependence of reflectance on wavelength, λ. This is made explicit as the spectral bidirectional reflectance distribution function (SBRDF),

$$f_r(\theta_i, \phi_i; \theta_r, \phi_r; \lambda)$$

The SBRDF is a function of five variables. If there is interaction between spectral and geometric factors, as can be the case for materials with significant internal scattering, then the geometric distribution also is affected. If there is no interaction between wavelength, λ, and the geometric dependence of reflection then the SBRDF is separable, and we write

$$f_r(\theta_i, \phi_i; \theta_r, \phi_r; \lambda) = f_r(\theta_i, \phi_i; \theta_r, \phi_r)\ f(\lambda)$$

where $f(\lambda)$ is a wavelength dependent attenuation factor. It is common to assume that spectral reflectance is separable, in the sense just described. Unfortunately, there is little data to establish the extent to which natural materials are (or are not) separable.

In the absence of data, it is dangerous to speculate. Nevertheless, it is useful to consider consequences to colour measurement. Imagine holding an object at arm's length and fixating on a particular surface element. By rotating the object in one's hand, the direction of the surface normal at the point of fixation is varied and therefore so are the incident and reflected ray directions (θ_i, ϕ_i) and (θ_r, ϕ_r). If the material's SBRDF is separable, then the brightness of the point might well change but the spectral distribution of the reflected light will not. Conversely, if the material's SBRDF were not separable, then the spectral distribution of the reflected light would change, possibly dramatically.

Some materials do appear to change color with simple rotation as non-separability would suggest. One example is the neck feathers of certain waterfowl that change colour in the blue to green range due to the presence of significant internal scattering. On the other hand, there are relatively few materials whose colour appears to change in this way. It is not clear, however, whether this is to be taken as evidence that the SBRDF of common materials is separable or as evidence of the robustness of human colour perception with respect to the consequences of non-separability.

2.1.3 Problems Modeling Reflectance

The SBRDF is an instantaneous quantity that cannot be measured directly. There are protocols for reflectance measurement but it remains difficult to generalize beyond the particular scale, illumination and viewpoint for which the actual measurements are obtained. For example, there is no scale independent fundamental distinction to be made among optical properties, surface microstructure and gross surface shape.

Even the SBRDF is a simplification. The SBRDF does not take into account significance subsurface penetration (i.e., situations in which a reflected ray emerges at a different surface location from that of the incident ray), interference or diffraction phenomena, polarization and other effects such as fluorescence and phosphorescence.

Reflectance modeling remains difficult because complexities arise at several levels. Local reflectance depends not only on the intrinsic optical properties of a material but also, for example, on its surface roughness. Geometric ray analysis of surface microstructure can be complex. At finer scales, one also needs to take into account the wave nature of light, adding even more complexity. In scenes consisting of multiple objects, geometric analysis at the macroscale also becomes complex, making it difficult to deal effectively with inter-reflections and cast shadows. Despite this, there is considerable work on reflectance models for computer vision and graphics. An edited collection of papers on radiometry and colour provides a good introduction to the relevant literature (Wolff et al. 1992).

2.2 Robot Vision

Robot vision typically assumes that measured brightness depends upon surface shape. This has led to the development of shape–from–shading methods (Horn and Brooks 1989), of which photometric stereo is one example. When illumination and surface material are fixed, it becomes possible to relate measured brightness directly to shape. At the same time, it can be impossible to disambiguate smooth changes in surface shape from either smooth changes in surface material or smooth changes in illumination. If the assumption of a fixed surface material is violated then the perception of shape can be altered. For example, careful application of cosmetic makeup can be used to change the apparent shape of a face.

Shape–from–shading methods, including photometric stereo, are subject to error in the presence of cast shadows and inter-reflection. No purely local technique can succeed since these phenomena are inherently non-local. Inter-reflection typically results in smooth changes in the local illumination.

Recall that photometric stereo can overdetermine the solution locally. This has additional benefits. With three light source photometric stereo, one can can detect cast shadows and inter-reflection locally. Correcting for the effects of inter-reflection is, however, a difficult task (Nayar et al. 1991). Also, if the material's SBRDF is separable, it is possible to use a fourth light source to obtain information both about surface shape and surface material. Using ratios of images, a four light source case becomes equivalent to the three light source case, for the purposes of determining surface shape, with the addition that local information about $f(\lambda)$ also is obtained.

2.3 Remote Sensing

Remote sensing typically assumes that multispectral measurements define a "spectral signature" that depends upon surface material (i.e., ground cover). This has led to the development of spectral classification techniques. Spectral classification has had some success in applications to agriculture and forestry. Not surprisingly, however, difficulties arise when shape and illumination changes confound the measurements. For example, spectral classification techniques alone have had limited success in forestry applications in rugged terrain.

The physics-based approach has been used to decouple geometric effects, associated with elevation, slope and aspect, from spectral effects, associated with surface material (Woodham and Gray 1987). Figure 1 shows the components of image irradiance considered. Irradiance at a ground target includes both direct solar radiance and diffuse sky radiance. Diffuse sky radiance has a component due to solar radiance scattered by the atmosphere. It also has components due to radiance reflected from adjacent targets, both radiance reflected directly to the target and radiance scattered by the atmosphere back to the target. The sensor measures scene radiance from the target and two additional components. Path radiance is radiance scattered to the sensor from the solar beam. Some radiation reflected from adjacent targets also is scattered to the sensor.

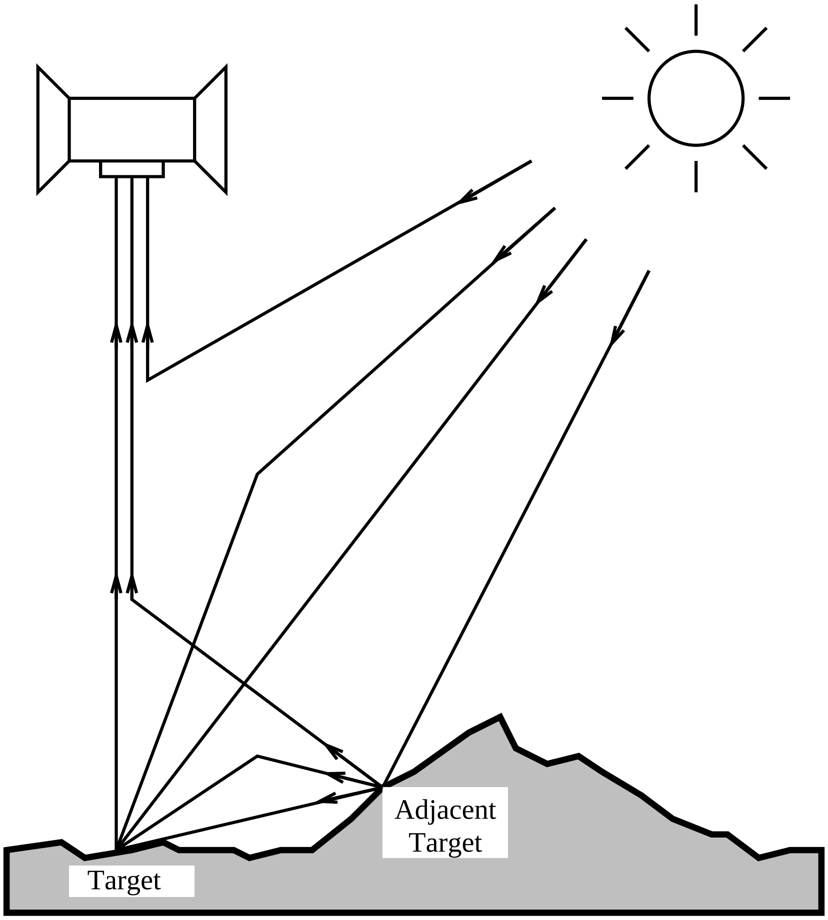

Figure 1: Components of Image Irradiance in Remote Sensing

Adjacent targets give rise to inter-reflection in rugged terrain. Again, this is difficult to model because the effects are non-local. The remote sensing case is further confounded since the atmosphere plays a significant role. Path radiance is the remote sensing equivalent to aerial perspective. Path radiance over dark targets has been measured and shown to be a function of elevation (Woodham and Lee 1985).

An example from remote sensing allows us to revisit the issue of spectral separability. Minnaert reflectance is a phenomenological model that has been used in remote sensing. Its BRDF is

$$f_r = \frac{k+1}{2\pi}\left[\cos(\theta_i)\cos(\theta_r)\right]^{k-1}$$

where k is a free parameter. The SBRDF of a Minnaert surface is not separable if k depends on λ. Remote sensing practitioners have empirically estimated values for k and typically concluded, for the ground covers investigated, that it does depend on λ. Woodham and Gray (1987) repeated these estimates, with similar results, but provided a somewhat different interpretation. First, it was argued that it is inherently implausible to postulate a material with the estimated dependence of k on wavelength since, as suggested in Section 2.1.2, colour measurements would vary dramatically for even small rotations of the surface. Second, it was noted that the application of the Minnaert model only took into account direct solar radiance. A more plausible, indeed expected, explanation for the measured shift towards the blue is that there is a corresponding shift towards the blue in target irradiance as skylight begins to dominate direct solar radiance.

It can be impossible to disambiguate smooth changes in surface reflectance from smooth changes in the spectral characteristics of the illumination. Colour constancy often is seen as the task of accommodating to the spectral characteristics of a single global illuminant. Experience with remote sensing and inter-reflection suggests that the task may be more local, since the spectral characteristics of the illuminant can vary locally.

3 The Ecology of Colour Vision

I am not an expert in biological vision. The material presented in this section largely has been gleaned from the literature. This is risky. Research in biological vision advances rapidly. There always is danger that what an outsider presents as current knowledge is out of date, incomplete or both. With that disclaimer, I proceed.

It is generally accepted that human vision is mediated by three opponent systems, one for luminance and two for hue. The two hue systems differ in more than just their spectral characteristics. Differences begin at the photoreceptors themselves and continue in the pathways up to and including visual cortex. This section highlights aspects of these differences. Attention then is turned to aquatic environments. This is where a link between photometric stereo and biological systems seems possible.

The human fovea contains three types of photosensitive cells, called cones. Cones are active in photopic (i.e., bright-light) vision. The human retina contains a fourth type of photosensitive cell, the rod. Rods are absent from the fovea but are active in scotopic (i.e., dark-adapted) vision. The three types of cones have different spectral responses. Let B-cone, G-cone and R-cone respectively denote the short-wavelength, medium-wavelength and long-wavelength sensitive cones. When considered together, the two longer wavelength cones are referred to as RG-cones.

3.1 B-cones and RG-cones

Evidence from psychophysics, anatomy, electrophysiology, microspectrophotometry (MSP) and molecular biology indicates that B-cones resemble rods in many respects.

1. Like rods, B-cones saturate even at low levels of light (Zrenner 1983, 82).

2. B-cones are fewer in number. They are estimated to constitute only about 5% of total cones in retina (Gouras 1985, 187) (Mollon 1990, 66). They are absent in the center of the fovea, appearing only at about 20 minutes arc (Boynton 1979, 326) (Mollon 1982, 177).
3. The absorption curves for R-cones and G-cones are very similar and are separated in wavelength by about 30nm. By comparison, the separation between the B-cone and the G-cone is about 100nm (Mollon 1990, 66).
4. Genes for the R-cone and G-cone lie close together on the X chromosome. The amino acid sequence for the two pigments is 96% identical (Mollon et al. 1990, 119). By comparison, the gene for the B-cone lies on chromosome 7 and the amino acid sequence for its pigment shows only 43% identity with those of the RG-cone (Mollon 1990, 68). The difference in the amino acid sequence between the B-cone and RG-cone pigments is about the same as that between the B-cone and the rod pigment (Mollon et al. 1990, 120).

3.2 B-system and RG-system

The B-cones contribute to a visual pathway called here the B-system. Similarly, The RG-cones contribute to an RG-system. The B-system and the RG-system remain largely separate from each other until relatively late in the processing stream, apparently being combined into an opponent system only at cortical levels (Gouras 1985, 195). The B-system and the RG-system thus appear to operate as parallel visual streams, without significant interaction, up to and beyond the lateral geniculate nucleus (LGN).

1. The B-system has a Weber fraction higher (9%) than that of the RG-system (2%) (Mollon 1982, 177). This means that the smallest perceivable brightness difference for the B-system is greater than that for the RG-system.
2. Spatial integration in the B-system occurs over a larger area resulting in 5–6 times lower spatial resolution than in the RG-system (Williams et al. 1981, 1357) (Williams et al. 1983, 488).
3. Similarly, temporal resolution is lower in the B-system than in the RG-system (Mollon 1982, 177) (Zrenner et al. 1990, 184).
4. The B-system contributes little to the perception of luminance (Zrenner 1983, 85). It is involved almost exclusively with the perception of hue (Boynton 1979, 326) (Mollon et al. 1990, 120). Input from the B-system is critical to the perception of “white” (Zrenner and Gouras 1981, 1609) (Gouras 1985, 190).
5. The B-system can contribute to chromatic border perception, giving rise to “soft” contours (Mollon 1990, 66). But it is two orders of magnitude less effective at producing these contours than is the RG-system (Boynton et al. 1985, 1351).

Evidence suggests that the B-system and the RG-system have been separate for a long time. On the other hand, the splitting of the RG-system into separate luminance and hue components in old world primates is believed to be quite recent, occurring only in the last 30–40 million years (Mollon et al. 1990, 119–20).

3.3 B/Y-system and R/G-system

Evidence of distinct visual pathways with differing spectral characteristics is not, by itself, evidence of colour perception. Colour opponency exists when a neuron is excited by input from one colour pathway and inhibited by input from another, spectrally different, colour pathway. Two systems of colour opponent cells have been observed in primates. The R/G-system receives input from the R-cones and G-cones and the B/Y-system receives input from the B-cones and Y=R+G cones. Compared to the R/G-system, the B/Y-system is not well understood.

1. The R/G system differences the output of two cone types that are similar in properties. The B/Y system differences the output of cone types that differ substantially in their spectral, spatial and temporal properties.
2. R/G cells are numerous throughout the visual pathway. B/Y cells are rare by comparison (Zrenner et al. 1990, 180).
3. Only one type of B/Y cell has been observed conclusively. B/Y cells occur with a B excitatory and a Y inhibitory receptive field (Mollon 1982, 176). Inhibitory B does not seem to occur (Gouras and Zrenner 1981, 1591) (Zrenner 1983, 86–87). All four combinations of R/G center/surround receptive fields are observed.
4. B/Y cells take antagonistic inputs from co-extensive regions of the receptor array (Mollon et al. 1990, 120) (Zrenner et al. 1990, 183). The vast majority of R/G cells, on the other hand, are believed to have distinct center/surround receptive fields (Gouras 1985, 187).
5. B/Y cells have the same neutral point. Neutral points for R/G cells vary over a continuous range from 480nm to 630nm (Zrenner et al. 1990, 184–85). Sensitivity to chromatic boundaries is maximized at a neutral point. Thus, the R/G-system responds effectively to a wide range of chromatic boundaries.

3.4 Luminance and Hue

The ganglion layer of retinal cells contains two intermixed cell types that differ in physical size and function. The small cells are colour opponent cells and contribute to what is called the parvo-system. The large cells do not distinguish cone type and contribute to what is called the magno-system.

The distinction between magno and parvo systems is anatomic. It remains a challenge to relate anatomic differences to visual function, especially at higher levels of visual cortex. It has been hypothesized that the brain has at least three separate visual processing systems, one for perception of shape, one for colour and one for movement, location and spatial organization (Livingstone 1988) (Zrenner et al. 1990, 194–99).

1. The hue system does not contribute to motion perception (Cavanagh et al. 1984). This suggests that motion mainly is a function of the magno-system, not the parvo-system (Zrenner et al. 1990, 198).
2. The hue system does not contribute to the perception of shadows (and hence of depth from shadows). Hue is not used to reject impossible shadow areas. Only luminance differences seem to matter (Cavanagh and Leclerc 1989).
3. In general, the hue system does not contribute significantly to 3D shape perception. Cues such as perspective, relative size of objects, stereo disparity, shading and texture gradients appear colour blind. (Livingstone 1988).

3.5 Aquatic Environments

In aquatic environments, radiant flux is affected by the same optical processes that occur in the atmosphere. But, the effects of scattering and absorption are more pronounced in water than in air. Aquatic environments present an extreme range of conditions for vision and one finds fish living in almost every conceivable photic environment. This provides a unique natural laboratory for vision research.

The refractive index of water relative to air is about 1.33. Because of this, the full hemisphere of incoming light at the air–water interface is compressed into a cone approximately ±48.5 degrees from the vertical (Lythgoe 1979, 13). Thus, light downwelling from the surface is confined to a narrow range of directions and constitutes about 97 degrees of the visual field. The remainder of the visual field has background "spacelight." At the surface, the full spectrum of light in air is represented. As light is transmitted downward, scattering and absorption narrow the spectrum and the vertical light becomes increasingly monochromatic. Figures 2–4, adapted from (Levine and MacNichol Jr. 1982), show typical examples. In clear oceans and lakes light penetrates to considerable depth and becomes blue as the optical path length increases. In fresh water that carries green organic matter light penetrates less and becomes green. In rivers, swamps and marshes that carry products of plant and animal decay light penetrates least and becomes brown. The colour of the background spacelight similarly varies. But, for a given aquatic environment, it has relatively constant colour along every line of sight (McFarland and Munz 1975). Thus, except at depth, illumination in aquatic environments has a two directional components that are spectrally distinguished.

The existence of two separate colour streams, noted in Section 3.2, is not restricted to primates but occurs in virtually all vertebrates (Jacobs 1981, 177). The reason for the appearance of separate streams is not certain. But, there is increasing support for the theory that separate colour streams first evolved in aquatic environments to enhance contrast between object and background and that this evolutionary trend was given particular impetus by the narrow spectrum of background spacelight (McFarland and Munz 1975, 1079). According to this theory, integration of the separate streams for colour perception occurred much later.

Target detection depends on adequate contrast between the target and the background. Underwater, a dark target ahead or below has maximum contrast if the peak spectral response of the sensor matches that of the background spacelight. A bright target necessarily is lit from above and necessarily reflects a broader range of wavelengths. A sensor with peak spectral response offset from that of the background spacelight will achieve higher contrast since the effect of background spacelight then is minimized (McFarland and Munz 1975; Loew and Lythgoe 1978; Levine and MacNichol Jr. 1982).

Underwater, the amount of available light diminishes with depth. The competing visual goals are sensitivity and contrast enhancement. These goals can be satisfied with two separate processing streams whose spectral responses are tuned to that of the particular photic environment. For best performance, one stream would match the spectral characteristics of the background spacelight and the other would be offset from it.

Underwater spacelight is like aerial perspective, except that the effect on colour and contrast occurs over much shorter distances, including distances that are significant in predator/prey interactions. In this situation, the medium may be colour but the message is contrast.

Fish change their visual environment by migrating, for example, from fresh water to salt water and from near the surface to deep water. Thus, it is reasonable to wonder if the response characteristics of rods and cones are mutable and, if they are, over what time scale. Generally speaking, the visual pigments in fish can change and do tend to be well matched to the local photic environment. Response characteristics are known to change within the lifetime of an individual.

All vertebrate (and most invertebrate) visual pigments consist of a large protein component,

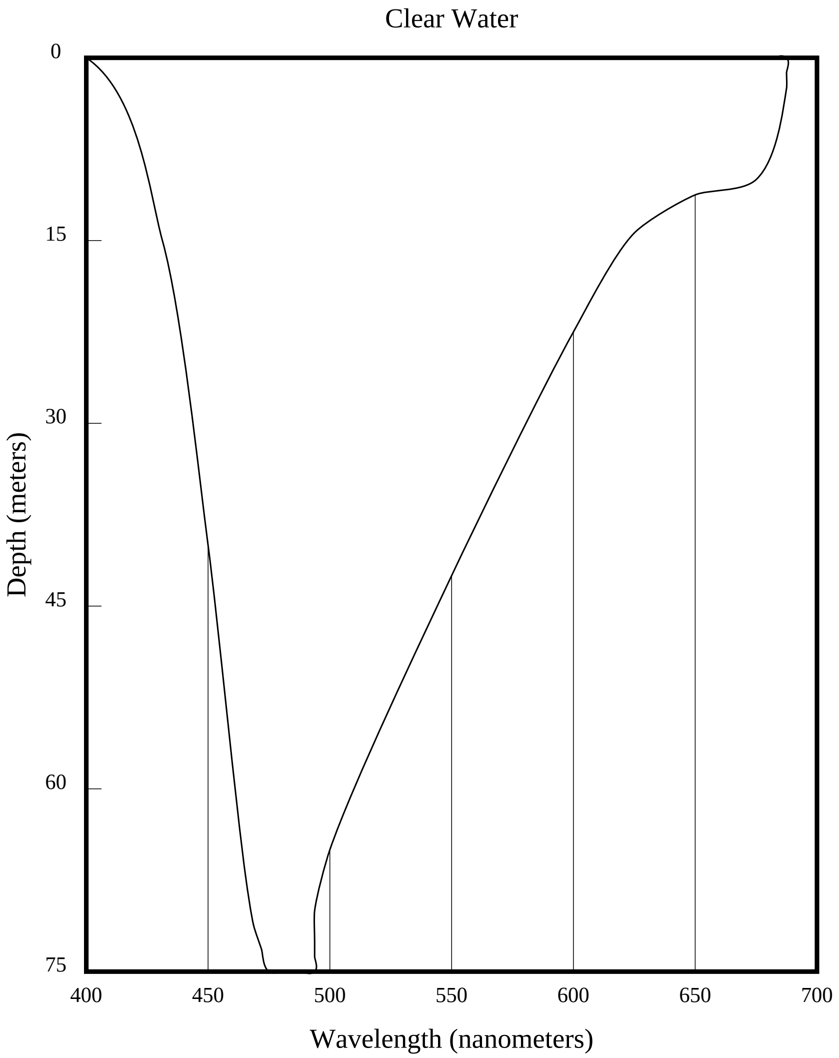

Figure 2: Spectral narrowing: clear oceans and lakes

opsin, which has embedded in it a prosthetic group, the aldehyde of vitamin A, retinal. Since there are two forms of vitamin A, A1 and A2, there are two families of visual pigments rhodopsins, based on retinal from A1, and porphyropsins, based on 3-dehydroretinal from A2 (Bowmaker 1990, 82–83).

Rhodopsins are found throughout the vertebrates whereas porphyropsins are restricted mainly to some teleost fish, amphibians and aquatic reptiles. The wavelength of peak sensitivity is determined both by the amino acid sequence of the opsin and by the type of the retinal. A given opsin can generate a rhodopsin/porphyropsin "pigment pair." A retina can contain all rhodopsin or all porphyropsin, or any proportion of the two, and the proportion can vary as the result of season, temperature and day length (Archer et al. 1987, 1249).

The retina of the American eel (_Anguilla rostrata_) contains rods with a P501/P523 rhodopsin/porphyropsin pigment pair. As a yellow eel, during the fresh water stage of its life cycle, the porphyropsin, P523, dominates. The eel migrates to the sea and undergoes metamorphosis to become a silver eel. After metamorphosis, the rhodopsin, P501, dominates. This shift towards shorter wavelengths is consistent with migration from river to sea. At the same time, a new rhodopsin, P482, appears that is

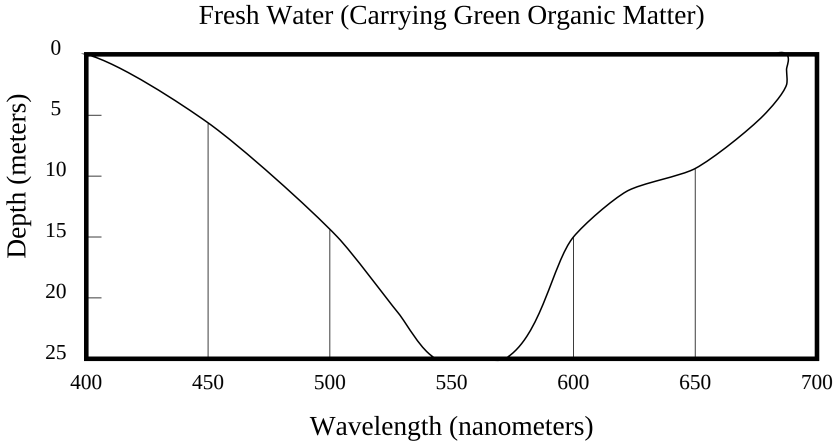

Figure 3: Spectral narrowing: fresh water carrying green organic matter

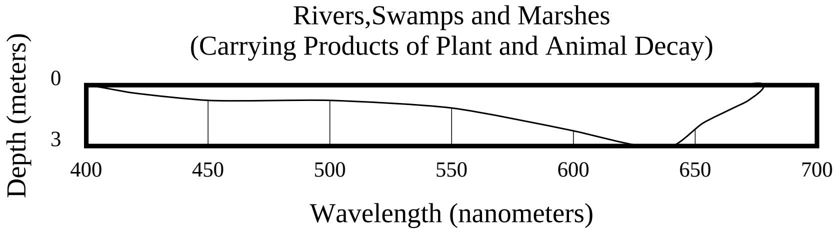

Figure 4: Spectral narrowing: rivers, swamps and marshes

maximally sensitive to the downwelling blue of the clear deep ocean. The P482 rhodopsin appears to be based on an opsin gene switched on during metamorphosis (Beatty 1975).

Being from British Columbia, I should mention that changes also occur in several species of salmon. At sea, Pacific salmon (_Oncorhynchus_) have rods dominated by rhodopsin P503. As they migrate upstream to their breeding grounds, the visual pigment becomes predominantly porphyropsin P527. P503/P527 are a rhodopsin/porphyropsin pigment pair (Beatty 1966).

Microspectrophotometric (MSP) studies have shown that some species of fish have as many as four cone types (Douglas and Hawryshyn 1990, page 402). But, it is difficult to determine how these cone mechanisms might interact to produce hue sensitivity, if indeed this is their purpose. Some fish may have evolved ultraviolet receptors for the detection of plane polarized light (Hawryshyn et al. 1988, 115). Polarized light may mediate navigational behaviour during migration.

Finally, there are examples of intraretina differences in visual pigments. Levine and MacNichol Jr. (1982) describe the guppy (_Poecilia reticulata_). The guppy lives near the surface. Its upper retina, which receives light from below, has visual pigments suited to colour discrimination. Colourfulness of the male is important to the female in her choice of mate. To attract a female, the male displays his colouration from a position ahead and slightly below that of his intended mate. By comparison, the lower retina, which receives light from above, has photoreceptors containing the same green-sensitive pigment. The green sensitive portion of the retina is most suited to detecting dark bits of food silhouetted against the predominantly green vertical spacelight.

4 Implications for Colour Vision Research

Spectral radiometry defines the physical basis for what we loosely term colour. But, there has never been a direct relationship established between local measurements of reflectance and colour perception. Accounting for human colour perception, including the phenomenon of colour constancy, remains a challenge. Yet, there is danger in thinking of colour only in terms of spectral reflectance and human perception. Vision research now suggests that humans are not the best colour perceivers. There are fish, birds and other animals that distinguish more colours and colours more finely. The electromagnetic spectrum is a rich medium. Colour, including the ultraviolet and the near infrared, is used biologically for purposes other than estimating spectral reflectance of surfaces.

It seems impossible, in general, to disambiguate smooth changes in spectral reflectance, smooth changes in surface shape and smooth changes in the spectral characteristics of the illumination. Colour constancy is more than the task of accommodating to the spectral characteristics of a single global illuminant. We need measurements of spectral reflectance to determine what assumptions about the separability of an SBRDF are warranted. We already know that inter-reflection can alter the spectral characteristics of the illumination locally.

In humans, the R/G-system seems well suited to detecting chromatic boundaries at high resolution over a broad range of wavelengths. It is not well suited to making an absolute determination of colour value. Thus, the R/G-system might be insensitive to the precise spectral characteristics of the illuminant and to any smooth changes in spectral reflectance owing to any wavelength dependencies in local surface shading. The B-system is involved in the absolute determination of colour value since, as mentioned, it is critical to the perception of “white.” The B-system, however, has low spatial resolution and thus is integrating information over a larger spatial area. This also provides a degree of insensitivity to local changes in the spectral characteristics of the illuminant and to smooth changes in surface reflectance. Taken together, this suggests that human colour perception is not based on the spectral properties of a point but rather is a property of an extended region.

It has been useful to consider ways that the technique of photometric stereo might have biological significance. Photometric stereo relies on images acquired from a single viewpoint but under different conditions of illumination. This has proven to be an effective way to determine the shape of objects. Implementation has evolved from the use of a B&W camera, with white lights turned on and off, to the use of a colour camera and directional illumination that is spectrally distinct. Colour becomes the medium to carry additional information about object shape.

In aquatic environments, illumination has directional components that are spectrally distinguished. Underwater, colour is a medium that conveys useful information about object identity because it can be used to enhance contrast between object and background. This is possible because of the narrow spectrum of background spacelight. It is fascinating that separate colour pigments may first have evolved in a fishlike vertebrate precisely for this purpose.

Spectrally and spatially distinguished illumination does have biological significance. Biological systems may not have evolved to exploit this for the determination of shape, as in photometric stereo. But, they might have. Alas, David Marr is not available to offer a rebuttal. I like to think that he would still regard photometric stereo as elegant. At the same time, he might be willing to reconsider his earlier view that the technique cannot possibly have any biological significance.

Acknowledgements

David Marr was a member of my Ph.D supervisory committee. Thanks to Patrick Cavanagh, Max Cynader and Ron Rensink for encouragement to explore connections between my work and biological vision. Special thanks to Ron Rensink who provided extensive reviews of the biological vision literature and who first suggested I consider aquatic environments. Thanks to Brian Funt for inviting me to participate in the SFU Conference on Colour Perception. John Endler, an attendee at the conference, provided additional pointers to the biological vision literature. Craig Hawryshyn contributed further insight into colour vision in fish. Support for Woodham's computational vision research was provided by the Institute for Robotics and Intelligent Systems (IRIS), one of the Canadian Network of Centres of Excellence, by the Natural Science and Engineering Research Council of Canada (NSERC) and by the Canadian Institute for Advanced Research (CIAR).

References

Archer, S. N., J. A. Endler, J. N. Lythgoe, and J. C. Partridge (1987). Visual pigment polymorphism in the guppy *Poecilia reticulata*. *Vision Research*, 27: 1243–1252

Barron, J. L., D. J. Fleet, and S. S. Beauchemin (1994). Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1): 43–77

Beatty, D. D. (1966). A study of the succession of visual pigments in Pacific salmon (*Oncorhynchus*). *Canad. J. Zool.*, 44: 429–455

Beatty, D. D. (1975). Visual pigments of the American eel *Anguilla rostrata*. *Vision Research*, 15: 771–776

Bowmaker, J. K. (1990). Visual pigments of fish. In R. H. Douglas and M. B. A. Djamgoz (eds.) *The Visual System of Fish*, 81–107. Chapman and Hall

Boynton, R. M. (1979). *Human Color Vision*. New York: Holt, Rinehart and Winston

Boynton, R. M., R. T. Eskew Jr., and C. X. Olson (1985). Blue cones contribute to border distinctness. *Vision Research*, 25: 1349–1352

Cavanagh, P., and Y. G. Leclerc (1989). Shape from shadows. *Journal of Experimental Psychology: Human Perception and Performance*, 15: 3–27

Cavanagh, P., C. W. Tyler, and O. E. Favreau (1984). Perceived velocity of moving chromatic gratings. *Journal of the Optical Society of America, A*, 1: 893–899

Douglas, R. H., and C. W. Hawryshyn (1990). Behavioral studies of fish vision: An analysis of visual capabilities. In R. H. Douglas and M. B. A. Djamgoz (eds.) *The Visual System of Fish*, 373–418. London: Chapman and Hall

Gouras, P. (1985). Colour coding in the primate retinogeniculate nucleus. In D. Ottoson and S. Zeki (eds.) *Central and Peripheral Mechanisms of Colour Vision*, 183–197. London: MacMillan

Gouras, P., and E. Zrenner (1981). Color coding in the primate retina. *Vision Research*, 21: 1591–1598

Hawryshyn, C. W., M. G. Arnold, W. N. McFarland, and E. R. Loew (1988). Aspects of color vision in bluegill sunfish (*Lepomis macrochirus*): ecological and evolutionary relevance. *Journal of Comparative Physiology A*, 164: 107–116

Healey, G., and R. Jain (1994). Physics-based machine vision. *Journal of the Optical Society of America, A*, 11: 2922. (Introduction to special issue.)

Horn, B. K. P., and M. J. Brooks (eds.) (1989). *Shape from Shading*. Cambridge, MA: MIT Press

Horn, B. K. P., and B. G. Schunck (1981). Determining optical flow. *Artificial Intelligence*, 17: 185–203

Jacobs, G. H. (1981). *Comparative Color Vision*. New York: Academic Press

Levine, J. S., and E. F. MacNichol Jr. (1982). Color vision in fish. *Scientific American*, 246(2): 140–149

Livingstone, M. S. (1988). Art, illusion and the visual system. *Scientific American*, 258(1): 78–85

Loew, E. R., and J. N. Lythgoe (1978). The ecology of cone pigments in teleost fishes. *Vision Research*, 18: 715–722

Lythgoe, J. N. (1979). *The Ecology of Vision*. New York: Oxford University Press

Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W.H. Freeman

McFarland, W. N., and F. W. Munz (1975). The evolution of photopic visual pigments in fishes. *Vision Research*, 15: 1071–1080

Mollon, J. (1990). The tricks of colour. In H. Barlow, C. Blakemore, and M. Weston-Smith (eds.) *Images and Understanding*, 61–78. Cambridge: Cambridge University Press

Mollon, J. D. (1982). Colour vision and colour blindness. In H. B. Barlow and J. D. Mollon (eds.) *The Senses*, 165–191. Cambridge: Cambridge University Press

Mollon, J. D., O. Estévez, and R. C. Cavonius (1990). The two subsystems of colour vision and their roles in wavelength discrimination. In C. Blakemore (ed.) *Vision: Coding and Efficiency*, 119–131. Cambridge: Cambridge University Press

Nayar, S. K., K. Ikeuchi, and T. Kanade (1991). Shape from interreflections. *International Journal of Computer Vision*, 6(3): 173–195

Nicodemus, F. E., J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis (1977). Geometrical considerations and nomenclature for reflectance. NBS Monograph 160, National Bureau of Standards, Washington, DC

Siegerist, C. E. (1996). Parallel near-real-time implementation of multiple light source optical flow. In *Vision Interface 96*, 167–174, Toronto, Ontario

Silver, W. M. (1980). Determining shape and reflectance using multiple images. Ms thesis, MIT Dept. Electrical Engineering and Computer Science, Cambridge, MA

Williams, D. R., R. J. Collier, and B. J. Thompson (1983). Spatial resolution of the short-wavelength mechanism. In J. D. Mollon and L. T. Sharpe (eds.) *Colour Vision: Physiology and Psychophysics*. London: Academic Press

Williams, D. R., D. I. A. MacLeod, and M. M. Hayhoe (1981). Punctate sensitivity of the blue-sensitive mechanism. *Vision Research*, 21: 1357–1375

Wolff, L., S. Shafer, and G. Healey (eds.) (1992). *Physics-Based Vision: Principles and Practice (Vol. I Radiometry, Vol. II Color and Vol. III Shape Recovery)*. Boston: Jones and Bartlett

Woodham, R. J. (1980). Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19: 139–144

Woodham, R. J. (1981). Analysing images of curved surfaces. *Artificial Intelligence*, 17: 117–140

Woodham, R. J. (1990). Multiple light source optical flow. In *Proc. 3rd International Conference on Computer Vision*, 42–46, Osaka, Japan

Woodham, R. J. (1994). Gradient and curvature from the photometric stereo method, including local confidence estimation. *Journal of the Optical Society of America, A*, 11: 3050–3068

Woodham, R. J., and M. H. Gray (1987). Analytic method for radiometric correction of satellite multispectral scanner data. *IEEE Transactions on Geoscience and Remote Sensing*, 25: 258–271

Woodham, R. J., and T. K. Lee (1985). Photometric method for radiometric correction of multispectral scanner data. *Canadian Journal of Remote Sensing*, 11: 132–161

Zrenner, E. (1983). *Neurophysiological Aspects of Color Vision in Primates*. Berlin: Springer-Verlag

Zrenner, E., I. Abramov, M. Akita, A. Cowey, M. Livingstone, and A. Valberg (1990). Color Perception. In L. Spillmann and J. S. Werner (eds.) *Visual Perception: The Neurophysiological Foundations*, 163–204. San Diego: Academic Press

Zrenner, E., and P. Gouras (1981). Characteristics of the blue sensitive cone mechanisms in primate retinal ganglion cells. *Vision Research*, 21: 1605–1609