

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

Supervised Color Constancy Using a Color Chart

Carol L. Novak and Steven A. Shafer

14 June 1990

CMU-CS-90-140

**School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213**

This research was supported in part by the Defense Advanced Research Projects Agency (DOD) monitored by the Avionics Laboratory, Air Force Wright Aeronautical Laboratories, Aeronautical Systems Division (AFSC), Wright-Patterson AFB, Ohio 45433-6543 under Contract F33615-87-C-1499, ARPA Order No. 4976, Amendment 20 by the Jet Propulsion Laboratory, California Institute of Technology, sponsored by the National Aeronautics and Space Administration under Contract 957989; and by the Air Force Office of Scientific Research under Contract F49620-86-C-0127.

The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of DARPA, the Jet Propulsion Laboratory, NASA or the U.S. Government.

Keywords: Vision and scene understanding

Abstract

One of the key problems in machine vision is color constancy: the ability to match object colors in images taken under different colors of illumination. This is a difficult problem because the *apparent* color will depend upon the spectral reflectance function of the object and the spectral distribution function of the incident light, both of which are generally unknown. Methods to solve this problem use a small number of basis functions to represent the two functions, and some sort of reference knowledge to allow the calculation of the coefficients. Most methods have the weakness that the reference property may not actually hold for all images, or will have too little information to recover enough of the functions to make an accurate determination of what the color should be.

We have developed a method for color constancy that uses a color chart of known spectral characteristics to give stronger reference criteria, and with a large number of colors to give enough information to calculate the illuminant to the desired degree of accuracy. We call this approach "supervised color constancy" since the process is supervised by a picture of a known color chart. We present here two methods for computing supervised color constancy, one using least squares estimation, the other using a neural network.

We show experimental results for the supervised calculation of the spectral power distribution of an unknown illuminant. Once this has been calculated, the color of any object with known reflectance can be reliably predicted. We are developing an extension to allow the prediction of color appearance for an object whose spectral reflectance function is not known. We also propose a method of "incremental color constancy" which determines object color by repeated application of supervised color constancy under changing illumination.

Table of Contents

1. Introduction	1
1.1 Mathematical Formulation of Color Constancy	2
1.1.1 The Traditional Approach to Color Constancy	3
1.1.2 Limitations of the Traditional Approach	4
1.2 Previous Work in Color Constancy	4
1.3 Supervised Color Constancy	6
2. The Color Chart Method for Supervised Color Constancy	8
2.1 Estimation of the Illuminant Using the Color Chart	8
2.1.1 The Least Squares Method	9
2.1.2 The Color Chart	9
2.1.3 Choice of Basis Functions	9
2.1.4 Summary of Estimating the Illumination	11
2.2 Estimation of the Illuminant Using a Neural Network	12
3. Simulated Picture Taking	13
3.1 The Simulation Process	13
3.2 The Spectral Properties of Our Equipment	13
3.3 Evaluation Method	18
3.4 Results Using Least Squares Method	19
3.5 Results Using Neural Network Method	22
3.6 Comparison of Least Squares Method with Neural Network	25
4. A Complete Method for Supervised Color Constancy	29
4.1 Estimation of the Reflectance Function Using Multiple Illuminants	29
4.2 Summary of Proposed Method for Supervised Color Constancy	30
4.3 Incremental Color Constancy	31
5. Conclusions	32

List of Figures

Figure 1: The MacBeth ColorChecker	10
Figure 2: Legendre Polynomials 0-3	11
Figure 3: Spectral reflectances of the first four squares of a hypothetical color chart	14
Figure 4: Spectral transmittances of filters	14
Figure 5: Spectral responsivity of a hypothetical CCD camera	15
Figure 6: Comparison of sensor responsivities with and without aperture balancing	16
Figure 7: Spectral reflectances of six color standard tiles	16
Figure 8: Spectral power distribution of hypothetical indoor lights	17
Figure 9: Spectral power distribution of CIE daylight illuminants	18
Figure 10: Weighting Function $V(\lambda)$	20
Figure 11: Comparison of indoor light with recovered illuminant in noise-free simulation	20
Figure 12: Comparison of outdoor light with recovered illuminant in noise-free simulation	21
Figure 13: Recovery of indoor light with unit basis functions (function with noise has been scaled by 0.01 so that it will fit on the same graph)	21
Figure 14: Noisy simulation of indoor light using 10 basis functions	22
Figure 15: Noisy simulation of outdoor light using 10 basis functions	22
Figure 16: Noisy simulation of indoor light using 5 basis functions	23
Figure 17: Noisy simulation of outdoor light using 3 basis functions	23
Figure 18: Recovery of indoor light using network method	24
Figure 19: Recovery of outdoor light using network method	24
Figure 20: Average error ΔL and $\Delta_w L$ for indoor lights using least squares method	25
Figure 21: Average error ΔL and $\Delta_w L$ for outdoor lights using least squares method	26
Figure 22: Average error E for indoor and outdoor lights using least squares method	26
Figure 23: Average error ΔL and $\Delta_w L$ for indoor lights using network method	27
Figure 24: Average error ΔL and $\Delta_w L$ for outdoor lights using network method	27

List of Tables

Table 2-1: Layout of the MacBeth ColorChecker

1. Introduction

We usually think of color as a property of objects, but in actuality the observed color of an object varies all the time because it depends upon the lighting environment. The confusion comes about because we use the word color loosely, to mean both a physical quantity and a subjective quality. As a physical quantity, color is the spectral characteristic of light emanating from an object or light source. As a subjective quality, color is the way that something appears to us. Subjectively, things appear the same color day after day, indoors and outdoors. Physically, the color varies with every little change in the illumination. Color constancy is the process that brings these two ideas together, allowing us to perceive the color of an object as a constant no matter what the illumination may be.

For computers to identify objects by their color, a computer vision system must be able to automatically correct for different colors of incident light. It is well known that the human visual system can automatically correct for fairly large variations in the incident light. An object will usually appear to us to be the same color under two different lighting conditions even though a physical instrument would measure two different spectral reflectance functions. It is not known precisely how humans accomplish this adjustment, nor even precisely where in the visual system this correction occurs, although several researchers have offered theories [30].

In order to recognize objects by their colors, a computer vision system must be able to make a similar type of correction. We can view the problem as first selecting some object and recording its color under an initial light source. Given a picture of the same object taken under any other light source, we would like to calculate a transform that would change the recorded color into the color the object appears to possess in the new picture. A good transform will yield transformed colors very close to the actual color in the picture, so that we can locate the object in the new picture quickly and accurately. We would like to be able to identify a large number of objects by their color in this way.

However we do not claim that a good method by our standards must work in any way similar to the way that humans do, nor do we wish to limit ourselves only doing as well as humans. We would like our method to succeed even in instances where the human fails. In order to achieve this, we would like to use spectral power distributions of illuminants and spectral reflectance functions of objects because the spectral reflectance curve is an invariant property of the object that is useful for object recognition.

These spectral functions may in general be arbitrary, so their calculation seems formidable. The usual approach is to represent a spectral function in terms of some weighted basis functions and to calculate the coefficients by reference to some *a priori* knowledge about the image in which the object appears. Most often, only one image is used for the calculation; this requires the use of an unreliable reference criterion and provides few constraints for calculating the coefficients. In this paper we propose a new approach that uses an image of a color chart with known properties to estimate the spectral properties of the illumination. The presence of the chart in the image is a reliable referent and yields a large number of constraints used for calculating the coefficients of the illuminant. We call this approach "supervised color constancy" because the calculation is supervised by an image of the reference chart.

We begin the discussion with a mathematical formulation of the color constancy problem. We then present two methods for calculation, least squares estimation and a neural network, and compare the results.

1.1 Mathematical Formulation of Color Constancy

To explain why the problem of color constancy is so formidable, we present a mathematical framework for the problem. By analyzing the mathematics carefully, we can identify several assumptions and limitations that characterize past work in this area. We will then present a method of performing color constancy which does not have these limitations.

In the common approach to color constancy a color image is taken of a scene illuminated by some unknown type of light. In a typical color picture three measurements are made at each point in the scene, yielding a three dimensional measurement vector M for each point. The goal of color constancy is to automatically calculate corresponding measurements M' of the same scene as it would appear if illuminated by a standard light. This allows a computer to perform color matching. A traditional assumption is that a single linear transform will apply everywhere in the scene. In mathematical terms this is equivalent to saying that there should be some transform X such that

$$M' = X M \text{ for all measurements } M \text{ taken of the scene} \quad (1-1)$$

The problem with this simple approach can be seen by examining the physical process by which the measurements are made. The light reflected from a surface is equal to the light incident upon that object $L(\lambda)$ (known as the *illuminant*) times the *spectral reflectance* function of that surface $\rho(\lambda)$. The measurement m obtained from a sensor looking at that surface is the light reflected from the surface times the responsivity function of the sensor $s(\lambda)$, integrated over the range of the spectrum for which $s(\lambda) > 0$ (usually the "visible" wavelengths of light) [19]

$$m = \int L(\lambda) \rho(\lambda) s(\lambda) d\lambda \quad (1-2)$$

If the surfaces being observed are matte, then there is no appreciable specular reflection, so $\rho(\lambda)$ is an invariant property of the surface. Typically there are three sensors or one sensor with three different filters, so that under each type of light, three measurements are made. In our work we use red, green, and blue filters, although other filters may be chosen. Let $s_R(\lambda)$ be the function defined by multiplying the sensor responsivity with a red filter transmittance function, $s_G(\lambda)$ be the sensor responsivity times a green filter transmittance, and $s_B(\lambda)$ be the sensor times a blue filter. Then for each surface the three measurements are

$$\begin{aligned} m_R &= \int L(\lambda) \rho(\lambda) s_R(\lambda) d\lambda \\ m_G &= \int L(\lambda) \rho(\lambda) s_G(\lambda) d\lambda \\ m_B &= \int L(\lambda) \rho(\lambda) s_B(\lambda) d\lambda \end{aligned}$$

In order to represent an arbitrary curve in a finite manner, we can choose to approximate the illuminant $L(\lambda)$ in terms of some some basis functions $b_1(\lambda) \dots b_n(\lambda)$. This way the spectral function of the light can be approximated by the sum of appropriately weighted basis functions:

$$L(\lambda) \approx \sum_{j=1}^n L_j b_j(\lambda)$$

If the basis functions are well chosen and a large enough number are used, then any type of illuminant or reflectance function can be closely approximated in this representation. Using this new representation, the sensor measurements may be rewritten as

$$m_i = \sum_{j=1}^n L_j \int b_j(\lambda) \rho(\lambda) s_i(\lambda) d\lambda$$

where the i subscript refers to one of the three filters.

Looking at the same surface with the same camera but under a standard illuminant $L'(\lambda)$ can be modeled as changing the illumination coefficients L_j to L'_j to produce sensor measurements m'_i . If we let $r_{ij} = \int b_j(\lambda) \rho(\lambda) s_i(\lambda) d\lambda$ then the measurements are simply

$$m_i = \sum_{j=1}^n L_j r_{ij}$$

$$m'_i = \sum_{j=1}^n L'_j r_{ij}$$

If we let M be the vector of the sensor measurements m_i , L be the vector of light coefficients L_j , and R be the matrix with entries r_{ij} , then we have a matrix equation for the measurements under the unknown illuminant

$$M = RL$$

and a similar one for the measurements under the standard illuminant

$$M' = RL'$$

It should be noted that R is not necessarily square, since there may be more measurements than basis functions. Combining M and M' using equation (1-1) we have

$$RL' = XRL$$

and solving for L' yields

$$L' = R^{-1}XRL$$

(1-3)

If R is not square the pseudo-inverse is used rather than the true inverse, where the pseudo-inverse is defined by $(R^T R)^{-1} R^T$. The result is that the coefficients which describe the standard illuminant L' may be expressed in terms of the unknown illuminant L , the matrix R (and its inverse), and the transformation matrix X . Since R is composed from the reflectance of the object as well as the sensor sensitivities and the basis functions, a different R will be needed for each object in the scene. We will assume that L and L' are constant in their respective scenes. Therefore either R should be eliminated from the equation, or X will have to vary throughout the scene, negating the original hypothesis that a single X existed for all points in the scene.

1.1.1 The Traditional Approach to Color Constancy

Equation (1-3) can be simplified if R is square and both R and X are diagonal. This means that R would take the form

$$R = \begin{bmatrix} r_1 & 0 & 0 \\ 0 & r_2 & 0 \\ 0 & 0 & r_3 \end{bmatrix}$$

and similarly

$$X = \begin{bmatrix} x_1 & 0 & 0 \\ 0 & x_2 & 0 \\ 0 & 0 & x_3 \end{bmatrix}$$

so that the transform would be "color balancing" by simply scaling each sensor band. Under these assumptions we would then be able to simplify

$$R^{-1}XRL = X$$

The result is that the matrix which transforms the pixel values M to M' is the same matrix which transforms the light

L to L' . Therefore in order to calculate what color any object will have under a new type of light, it is only necessary to calculate the way in which the light has changed; the spectral reflectance curve of the object itself does not enter into calculating the transform. This justifies the original assumption that a single transform will work everywhere in the image (i.e. for all M) and is the implicit assumption traditionally made by computer vision approaches to color constancy.

The contents of the diagonal transformation matrix X still need to be calculated. To do this some sort of reference property is needed. Most often this reference property is some known values of M and corresponding M' . Since X is traditionally assumed diagonal, it is a simple matter to calculate its entries based upon as little information as one correspondence. For example if a surface is measured as having the RGB values [1 2 3] and is known to correspond to values of [1 1 1] with respect to an ideal illuminant, then the diagonal entries of X will be 1, 1/2, and 1/3.

1.1.2 Limitations of the Traditional Approach

In deriving the traditional approach, we made the assumption that R is square and diagonal, i.e. its contents are

$$r_{ij} = 0 \text{ if } i \neq j$$

Recall that by definition $r_{ij} = \int b_i(\lambda)s_j(\lambda)\rho(\lambda)d\lambda$. So in order to make R diagonal, the basis functions must be chosen so that b_1 is orthogonal to s_2 and s_3 , b_2 is orthogonal to s_1 and s_3 , and so forth. Most often researchers have not chosen basis functions that would satisfy this relation. In particular, the RGB color TV primaries are not a suitable choice for basis functions in this formulation, since their curves do overlap and so cannot be orthogonal. For example, if the NTSC color TV primaries are used as basis functions as well as sensor functions, and the surface in question is an ideal white (flat spectral reflectance function) then

$$R = \begin{bmatrix} 1.000 & 0.166 & 0.002 \\ 0.166 & 0.136 & 0.010 \\ 0.002 & 0.010 & 0.074 \end{bmatrix}$$

where the entries of R have been scaled so that $r_{1,1} = 1$. If the basis functions to represent the illumination are chosen so that they are orthogonal to the sensor functions, they may give an entirely inadequate description of most types of illumination. This is because most sensors are fairly broad band and overlap to some degree.

In the more general case that R is not diagonal, the transformation from one illuminant to another varies at each different surface as shown in Equation (1-3). In this case the original assumption that a single transformation X will map all M s to corresponding M' s must not be correct. Rather than trying to get around these limitations, we propose to abandon the idea of finding a single transformation matrix that works everywhere in the image. Instead, we allow the more general formulation of the reflectance matrix R and present a method to explicitly calculate the illuminant $L(\lambda)$. We also suggest how the surface reflectance function $\rho(\lambda)$ can be calculated.

1.2 Previous Work in Color Constancy

Proposed methods for color constancy can be classified by the basis functions b_i they use for performing the transformation, and the by the reference property that goes into determining the transform X . As discussed above many traditional methods assume that X is diagonal, so they address only the calculation of multiplicative scale factors for the sensor bands.

The reference property that different methods use has often been based upon heuristics about real images. A common heuristic is that the average over all pixel values in a picture should be a shade of gray [4]. Which neutral shade between black and white is not specified, but the hues must balance out. Thus the average intensity over the entire picture is calculated independently for each band and then multiplicative constants are computed that will bring the ratios between the bands to 1:1:1. This rule is frequently made by automated film processors. For many

pictures this produces a satisfactory print, but unsatisfactory pictures are easy to find. For example, whenever a picture is *supposed* to consist mainly of a single dominant color, the regions of that color will be printed as less saturated, and regions of other colors will be tinted by the opposing color.

An even more restrictive assumption is that the average of all the pixel values should equal some particular color. The particular color is chosen based upon statistical data gathered about naturally occurring reflectances and illuminants [6]. This method has an advantage over the "any shade of gray" method in that it will correct for different intensities of incident light; the previous method only corrects for biases in hue and saturation as any intensity of gray is acceptable for the average value. However this method suffers from the same flaw, namely that it is easy to find pictures, even outdoor ones, where the assumption will lead to an incorrect answer.

Other models for performing color constancy transformations come from assumption of knowledge about the reflective properties of certain regions in an image. One assumption is that the brightest region in the image is in fact a white surface, that is, a surface that reflects equal amounts of energy at all wavelengths. The extent to which the color of this surface in the image differs from white, yields multiplicative factors for each band. If the RGB values of this brightest surface are [5,3,2] then the green band should be scaled by $5/3$ and the blue band by $5/2$. This method can of course fail whenever the brightest surface in the picture is not in fact white. A more flexible assumption is that the brightest regions measured independently in each band will add up to white [14]. This assumption says that the brightest region in the image need not correspond to a white surface, but that the brightest regions in each band reflect equal amounts of light. Thus if the brightest region of the red band of the picture has a value of 5, the brightest region in the green band a value of 3 and the brightest region in the blue band a value of 2, then the green band should be scaled by $5/3$ and the blue band by $5/2$. Although this assumption is less restrictive than "the brightest region should be white", it too can be violated by everyday images.

It is possible for the transform X from the unknown illuminant to the standard illuminant to have non-zero off-diagonal entries. One method uses correspondences between a large number of regions viewed under the two illuminants. The best transform from one illuminant to the other is then calculated using either least squares estimation [12] or a neural network [11].

All of these methods share the common assumption that the transformation from one color space to another is unique and one-to-one. Unfortunately this is not the case. The existence of metamers, a well known color phenomenon, shows that this assumption will fail on some surfaces [13]. Metamers are two surfaces that appear the same color under one illuminant and yet different colors under a different illuminant. The problem is that spectral reflectance is a continuous function and hence is infinite-dimensional, but it is sampled in only a few dimensions, typically by three broad-band receptors. When three basis functions are picked, two different spectral reflectances may be projected down into the same three values in three dimensions, but when slightly different basis functions are picked, they project into different values [29, 24].

Some researchers have used mathematical constraints to reduce the complexity of the problem. Maloney and Wandell calculate m coefficients for the light by taking measurements at more than m locations in the image [17]. However, the degrees of freedom in the reflectance curves that the surfaces are assumed to have is limited to one less than the number of sensors. An additional constraint is generated by the knowledge that real-world reflectances are never negative [27].

Although most work in color vision has used cameras and RGB values, some researchers have used complete spectral data instead of or in addition to RGB values [10, 26]. Since spectral data is sampled at many points across the spectrum rather than the typical 3 for color cameras, it provides far more constraints for calculating illumination and reflectance functions.

Highlights on inhomogeneous materials have been used by some researchers to find the color of the illumination [4, 8, 26]. According to the Dichromatic Reflection Model, colors observed on an inhomogeneous surface will be the sum of a surface component and a body component [25]. The surface component is usually assumed to have the color of the light only, whereas the body component is assumed to have the color of the light times the body color of the object. A highlight region will have a positive amount of both colors whereas a non-highlight region will have only the light color times the body color. Assuming that highlights can be found along with a non-highlight portion of the same object, the color of the light can be calculated, and then the object's color. While theoretically elegant, methods such as this may be time consuming and be plagued by sensor limitations since highlights are typically far brighter than non-highlight regions. Of course, the method will not work at all if highlights are absent or occur only on non-dielectric materials (i.e. metals). It also depends upon the widely-used simplification that surface reflection is constant across the spectrum. In actuality, surface reflection is a function of wavelength, although for many substances the variation across the spectrum is quite small [25].

The choice of the form and number of basis functions is also an issue since it determines what type of illuminants and surface reflectances may be represented. These spectral functions may be sampled at regular intervals with a spectroradiometer (typically every 10 nanometers within the visible range) [26] or with broad-band receptors such as the ones found in a color camera. Some researchers have represented spectral functions with a set of abstract functions such as Fourier functions [28, 1] or normalized Legendre polynomials [9]. Alternatively, a set of basis functions derived from a principal components analysis of spectral data of natural and man-made objects may be used [6]. Any finite representation will necessarily be inexact, since spectral functions may vary arbitrarily. Typically only the first three basis functions are used. Researchers hope that this small number of whichever set of functions they use will capture most of the variation in reflectance and illumination functions. The extent to which this is true usually depends upon the smoothness of the reflectance and illumination curves. Most surfaces will have smooth reflectance curves; some researchers argue that this is due to the nature of the physical processes which typically cause color [16]. However some illuminants, particularly fluorescent lights, have sharp peaks, making them difficult to represent under most schemes. Recently some researchers have argued that 3 basis functions are insufficient even for surface reflectances [21].

Since most methods generate each constraint from one sensor value, the number of basis functions is limited by the number of sensors. Clearly, to disambiguate as many as colors as possible, a large number of constraints would be desirable. However a practical number of sensors is typically three or four, so we would like to generate more constraints without needing more sensors.

1.3 Supervised Color Constancy

We would like the reference property used to calculate the color constancy transform to be a reliable process rather than a heuristic that works some of the time. One way of getting "ground truth" is to introduce a surface of known reflectance properties into the picture we are interested in, or take a test picture beforehand. If we take a picture of a known white card in some known position in the image under a particular incident lighting, the reflected color measured gives some important information about the color of the light. This method will work identically to the method of assuming that we can find a white surface somewhere in the picture, except that in this case we know that the surface we think is white is in fact white. There is no assumption that can be violated.

We call this approach "supervised" because it involves taking some test data which trains the computer as to what type of light to expect. This type of process is already used in some applications [1]. Illumination changes are common in the real world, and calibrated color standards are easily obtained. At the cost of introducing a known surface into the scene from time to time, we can provide a transform that yields invariant color information as the illumination changes.

In this paper we introduce a method of supervised color constancy using a color chart instead of a simple white card. We use a chart with m color patches with distinct known reflectance functions, so that a three-band color image of the chart will yield $3m$ constraints upon the spectral curve of the incident light. In the general case where there are s sensors, there will be sm constraints. Then we show how these constraints may be solved to give an estimate of the illuminant $L(\lambda)$. Not only is the color chart a reliable reference property, but the large number of constraints it generates allows us to do our calculations in a high dimensional space.

We then present an alternate method that uses a neural network trained to calculate the illuminant vector from the RGB values the camera recorded for the color chart under that illumination. This method has an initial training cost as the network must be shown a number of examples. An advantage over the previous method is that no measurement of camera spectral sensitivity is needed; such measurements can be difficult to obtain.

We show experimental results for these techniques on theoretical data with and without simulated noise added. We also develop quality measures for comparison of our results. These methods of using a color chart for color constancy are useful because they are simple and practical to use and obtain high accuracy in performing color constancy needed in real-world applications.

Finally, our work suggests a way to use the constraint formulation to estimate the reflectance vector ρ of an unknown surface from measurements of that surface taken under one or more different illuminants whose spectral power distributions were estimated by the color chart method. Just as m known color patches will yield $3m$ constraints upon the illuminant in a system with three sensors, observing the surface under k known or estimated illuminants will yield $3k$ constraints upon the unknown reflectance function. In this way multiple observations under different illuminants can be accumulated to provide incremental improvements in the estimate of surface colors.

2. The Color Chart Method for Supervised Color Constancy

It is possible to guarantee that the reference property holds by introducing a known surface, typically white, into the scene. However a single white surface provides a limited amount of information. If measurements are made of it in only three bands, as is typical, it only provides three constraints upon the estimate of the illuminant. Taking measurements in more bands will provide additional constraints, but there is a practical limit on how many different filters may be obtained and how much time can be devoted to taking measurements.

Instead, a color chart with several distinct known reflectance curves may be introduced into the scene and measured by the sensors. The additional surfaces provide more constraints, allowing a more detailed description of the illumination to be calculated by the least squares method. Our method has two parts:

- Estimating the illuminant using the color chart
- Estimating object reflectance using images under several illuminants

In this paper we present the theory used in calculating both parts and experimental results of simulations for the first part.

2.1 Estimation of the Illuminant Using the Color Chart

In order to estimate the illuminant we assume that the sensitivity functions of the sensors are known, as well as the reflectance functions of the surfaces on the color chart. These functions can usually be obtained from the manufacturer or measured directly with a spectroradiometer. The matrix R is calculated by factoring these together with the basis functions. If three sensors are used that are responsive to red/green/blue, the first three rows of R consist of

$$\begin{bmatrix} \int b_1(\lambda) s_R(\lambda) \rho_1(\lambda) d\lambda & \dots & \int b_n(\lambda) s_R(\lambda) \rho_1(\lambda) d\lambda \\ \int b_1(\lambda) s_G(\lambda) \rho_1(\lambda) d\lambda & \dots & \int b_n(\lambda) s_G(\lambda) \rho_1(\lambda) d\lambda \\ \int b_1(\lambda) s_B(\lambda) \rho_1(\lambda) d\lambda & \dots & \int b_n(\lambda) s_B(\lambda) \rho_1(\lambda) d\lambda \end{bmatrix}$$

where $\rho_1(\lambda)$ is the reflectance function of the first surface on the color chart. The next three rows consist of the same calculations for the second surface on the color chart, and so on, so that for m surfaces on the color chart there are $3m$ rows in the matrix. In the general case where there are s sensors there are sm rows in the matrix. Note that all these functions are known so R is constant.

Measurements of the color chart are taken under the unknown illuminant, producing a measurement vector M with $3m$ values. The first three entries in M are the RGB values that the camera measured for the first surface, the second three entries correspond to the second surface, and so on in the same order that the rows of R were calculated. The goal is then to solve the equation

$$M = RL$$

for L . Since there are m surfaces on the color chart, and 3 sensors used to measure the chart, and since we choose to represent the light with n basis functions, M is $3m \times 1$, R is $3m \times n$, and L is $n \times 1$. Typically n is chosen so that $n < 3m$ since it is generally the case that there is some redundancy in the color chart functions. Even if R has $3m$ linearly independent rows, some of the rows may be very close to being dependent. Therefore the least squares method can be used to solve for L , yielding an estimated solution L^* . The estimated illuminant $L^*(\lambda)$ is recovered from L^* by multiplying the coefficients by the appropriate basis functions

$$L^*(\lambda) = \sum_{i=1}^n L^*_i b_i(\lambda)$$

The number of illuminant coefficients n may be chosen empirically or through eigenvector analysis. A calculation of the eigenvectors of $R^T R$ gives a set of vectors that span the solution space for L , while the magnitude

of each eigenvalue tells how significant the corresponding eigenvector is. Because R^TR is positive definite symmetric, the eigenvalues will all be positive. Each eigenvalue is divided by the largest one and sorted in decreasing order. A threshold value is chosen based on the signal-to-noise ratio of the sensors. The number of eigenvalues greater than that threshold is the number of basis functions n that we choose for calculating L^* .

By adding more distinct colors, an increasing number of illuminant coefficients can be accurately calculated. This only works up to the point that it becomes too difficult to formulate enough distinct colors that can be differentiated by the camera. Each additional color will add some information to our constraints, but the amount of new information added by each new color will decrease as the number of colors grows. Eventually, that amount will be below the threshold that can be reliably detected by the camera. However, in practice it is easy to obtain 8 or more coefficients which provide a relatively detailed description of any typical illuminant.

2.1.1 The Least Squares Method

The least squares method is a standard technique for finding the best fit curve to a collection of data. In this case the dependent variable is the value measured in the image; the independent variables are the reflectance curves of the surfaces, the sensors' response curves and the basis functions; and the relation between them is hypothesized to be linear. The mathematical model is

$$m_i = L_1 r_{i,1} + \dots + L_n r_{i,n} + e$$

where m is the value measured in the image, r_{ij} is the matrix calculated from the combination of the surface, sensor, and basis functions, L_j is the desired illumination description, and e is the residual error.

In order for the least squares calculation to produce a good result, several things must be true:

- 1) the hypothesized relationship must closely approximate reality
- 2) there must be more measurements than desired illuminant coefficients
- 3) the independent variables must be independent

We have already shown that our model is a good approximation of the physics of spectral integration. The number of measurements can be manipulated by simply adjusting the number of colored surfaces on the chart. The third condition depends on the nature of the matrix R , so the basis functions and the colors used in the color chart must be chosen carefully.

2.1.2 The Color Chart

We have used the standard MacBeth ColorChecker in our work [15]. The ColorChecker, shown in Figure 1, measures 13.5 inches by 9.25 inches and has 24 colored squares. The Munsell coordinates and common names of the colors are given in Table 2-1. Notice that the last row of the ColorChecker consists of neutral colors which are shades of gray ranging from black to white. This chart is widely used in color work and the spectral reflectance curves of the squares are readily available from the manufacturer.

2.1.3 Choice of Basis Functions

Although many different families of basis functions may theoretically be used, in practice some care must be taken in choosing among them. A simple choice would be a series of unit impulses ranging over the visible wavelengths. For example the first function could be a curve consisting of an impulse at 400nm and zero everywhere else; the second function could have an impulse at 405nm and a zero everywhere else; and so on. The problem with this choice is that in solving the least squares formulation, the assumption that the independent variables are in fact linearly independent may be violated. Many illuminants are fairly smooth across the spectrum; as a result the value of L_2 would be approximately the average of L_1 and L_3 . With such a set of basis functions,

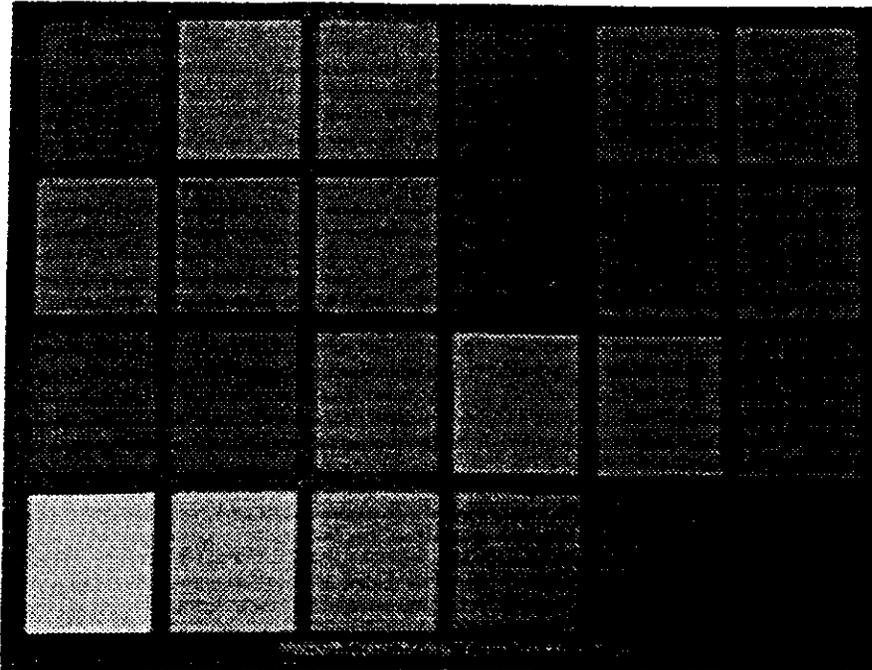


Figure 1: The MacBeth ColorChecker

3YR 3.7/3.2 (Dark Skin)	2.2YR 6.47/4.1 (Light Skin)	4.3PB 4.95/5.5 (Blue Sky)	6.7GY 4.2/4.1 (Foliage)	9.7PB 5.47/6.7 (Blue Flower)	2.5BG 7/6 (Bluish Green)
5YR 6/11 (Orange)	7.5PB 4/10 (Purplish Blue)	2.5R 5/10 (Moderate Red)	5P 3/7 (Purple)	5GY 7.1/9.1 (Yellow Green)	10YR 7/10 (Orange-Yellow)
7.5PB 2.9/12.7 (Blue)	0.25G 5.4/8.65 (Green)	5R 4/12 (Red)	5Y 8/11.1 (Yellow)	2.5RP 5/12 (Magenta)	5B 5.08/8.0 (Cyan)
N 9.5/ (White)	N 8/ (Gray)	N 6.5/ (Gray)	N 5/ (Gray)	N 3.5/ (Gray)	N 2/ (Black)

Table 2-1: Layout of the MacBeth ColorChecker¹

variables would tend to be highly correlated with each other. This state of affairs is known as *multicollinearity* and the result is that very small errors in the measurements can cause very large errors in the calculation of the L_i coefficients. Moreover if one variable is over-estimated, the next is likely to be under-estimated, leading to wild oscillations in the calculated values of the coefficients.

A more sound choice is a set of basis functions which describe characteristics which are likely to be independent such as the Fourier components of the curve. We use the Legendre polynomials, suggested by Healey and Binford [8]. The definition of the i^{th} Legendre polynomial b_i is

¹Munsell coordinates are usually written as *Hue Value/Chroma*. The *Hue* consists of a letter part and a number part. The letter is one of 5 principal colors, Red(R) Yellow(Y) Green(G) Blue(B) Purple(P), or a combination of two adjacent colors (such as YR for Yellow-Red). The number part ranges from 0 to 10 in such a way that 0R would be a red hue with a lot of purple in it, 10R would be a red hue with a lot of yellow in it, and 5R would be the hue that is usually thought of as "red". The *Value* tells how light the color is, ranging from 0, which is absolute black, to 10, which is absolute white. The *Chroma* tells how strong the color is, ranging from 0 which is gray, to 16 which is very vivid. Chroma is somewhat similar to saturation, but unlike saturation, lightness is factored into the perception of chroma. Neutral colors (grays) are usually written as N *Value/* with the 0 value for the chroma omitted [18].

$$b_i(x) = \sum_{k=0}^{i/2} \frac{(-1)^k (2i-2k)! x^{i-2k}}{2^i k! (i-k)! (i-2k)!}$$

The first few Legendre polynomials are shown in Figure 2. Note that the i^{th} basis function is an i^{th} degree polynomial. These are orthogonal as indicated by the relationship

$$\int_{-1}^{+1} b_i(x) b_j(x) dx = \begin{cases} \frac{2}{2i+1} & \text{if } i=j \\ 0 & \text{otherwise} \end{cases}$$

By scaling the visible wavelengths of light to range over $[-1,+1]$ and using a suitable number of basis functions, we can easily represent the visible portion of any illuminant with a high degree polynomial.

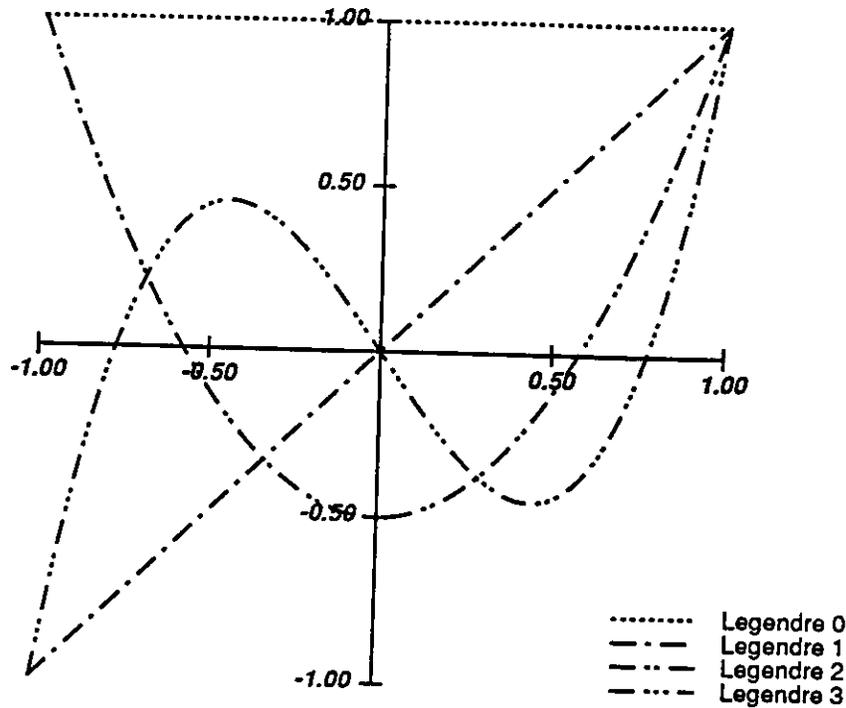


Figure 2: Legendre Polynomials 0-3

2.1.4 Summary of Estimating the Illumination

For three sensors, the elements of the matrix R are calculated by

$$r_{ij} = \int b_j(\lambda) s_{(i \bmod 3)}(\lambda) \rho_{[i/3]}(\lambda) d\lambda$$

where $b_j(\lambda)$ is the j^{th} Legendre polynomial, $s_{(i \bmod 3)}(\lambda)$ corresponds to one of the three sensor functions, and $\rho_{[i/3]}(\lambda)$ corresponds to one of the spectral reflectance curves of the color chart. When a color picture is taken of the chart under a particular illuminant, sensor measurements m are obtained that correspond to each row of R . Using the least squares method, a best-fit estimate for the illuminant is obtained. Because we use the same color chart, basis functions, and sensors each time, the matrix R and its least squares inverse may be calculated off-line once and then used to quickly obtain the illuminant by multiplying the measurements obtained in the experiment by the pre-calculated matrix. Eigenvector analysis of the matrix R obtained from the MacBeth ColorChecker and our sensors shows that 10 eigenvalues are significant, so the illumination is approximated with a 10^{th} degree polynomial. We present results for this method in section 3.4.

2.2 Estimation of the Illuminant Using a Neural Network

Recall that we are trying to solve the equation

$$M = RL$$

for L . Another approach to solving it is to train a neural network to calculate L given M so that it is implicitly calculating the constraints generated by the image of the ColorChecker, or rather their inverse R^{-1} .

Inputs to the network consist of the $3m \times 1$ measurement vector M made of the appearance of the ColorChecker under a particular type of light. Outputs from the network are the n coefficients of that light. Since multicollinearity is not an issue, the basis functions can be simple unit samples, so the illumination coefficients are simply the spectral power measured at each wavelength by a spectroradiometer.

The network configuration we have used has a single hidden layer with n hidden units with complete interconnection between the input and hidden layer, between the hidden and output layer, and between the input and output layer (shortcut connections). The weights were updated using continuous back-propagation [22, 23].

The network is trained using examples of the ColorChecker appearance and the corresponding illumination. Since the constraints generated by the chart reflectances and sensor responses are calculated implicitly by the network, these functions do not have to be measured or obtained from the manufacturer. Measuring the spectral response of a camera can be difficult even with the aid of a spectroradiometer. On the other hand, this method does not require the spectral measurement of the lights used to train the network and this will probably require access to a spectroradiometer. However, once the network has been properly trained, the spectroradiometer will not be needed and the method will calculate the spectral power of further lights very quickly. In effect we have constructed a "poor man's spectroradiometer", obtaining a detailed spectral power distribution function of illuminants using only a color chart and a color camera, with a neural network that "learns" the relationships among them. Results for this method are presented in section 3.5.

3. Simulated Picture Taking

To find out what kind of results we could theoretically expect from our method, we simulated the picture taking process and applied our method to the simulation data.

3.1 The Simulation Process

In Chapter 1 we introduced Equation (1-2) which says that given the spectral power distribution of the illumination, the spectral responsivity of the sensor, and the spectral reflectance of a surface, we should be able to calculate the value that will be obtained by a sensor pointing at that surface. Therefore, if we can somehow obtain the spectral curves of these things, we can multiply them together and calculate the integral to predict the value measured by the sensor.

Accordingly we collected spectral curves of ideal picture-taking equipment, as well as curves for a color chart, some interesting things to look at, and some common types of illumination. Where possible, this data was collected at 5nm intervals from 300nm to 850 nm, a range which extends somewhat beyond the visible wavelengths of light both in the infrared and ultraviolet directions.

The spectral curves were multiplied together and the integrals were determined by calculating the area under the piecewise linear curves. The resulting values were considered the theoretical sensor values of a hypothetical camera looking at a known color chart under an unknown illuminant. Then the calculations described in the previous section were performed to see if the light factored into the sensor values could be recovered. Theoretical sensor values were also used to train the neural network, which was then tested on new theoretical inputs it had not seen before.

Sensor values for a hypothetical object under several different illuminants were also calculated. Using the estimates for the illuminants already calculated by the least squares method, we attempted to recover the spectral curve of the object.

The above calculations were performed on perfect (noise-free) data and on data with simulated noise added. The noisy data was simulated by adding pseudo-random values, weighted by a Gaussian function, to the simulated camera measurements.

3.2 The Spectral Properties of Our Equipment

Although our experiments used simulation and so could have used any kind of spectral curves, we attempted to obtain spectral curves that were fairly close to the kinds of curves associated with the equipment we actually use. Fortunately, spectral data for most camera equipment can be obtained from the manufacturers. While most cameras and filters and the like are not carefully calibrated to meet the standards of the manufacturers, it is usually believed that the equipment comes close to these published standards.

The color squares of the MacBeth ColorChecker are composed of Munsell color papers. The spectral curves of these papers were obtained from the manufacturer and used in the simulation. The curves for the first four squares on the MacBeth chart are shown in Figure 3. These are considered hypothetical values since there is no guarantee that these curves will be realized on an actual MacBeth chart.

We also obtained the theoretical spectral transmittance curves of a red, green and blue filter of the type that we use for taking color images. These filters (Wratten #25 red, #58 green, #47 blue) are widely recommended for color photography and the expected spectral transmittances are published by the manufacturer. In addition we obtained the spectral transmittance of an infrared cut-off filter (Corion FR-400). We commonly use such a filter in picture taking due to our camera's sensitivity to infrared wavelengths [20]. The spectral curves for these four filters are

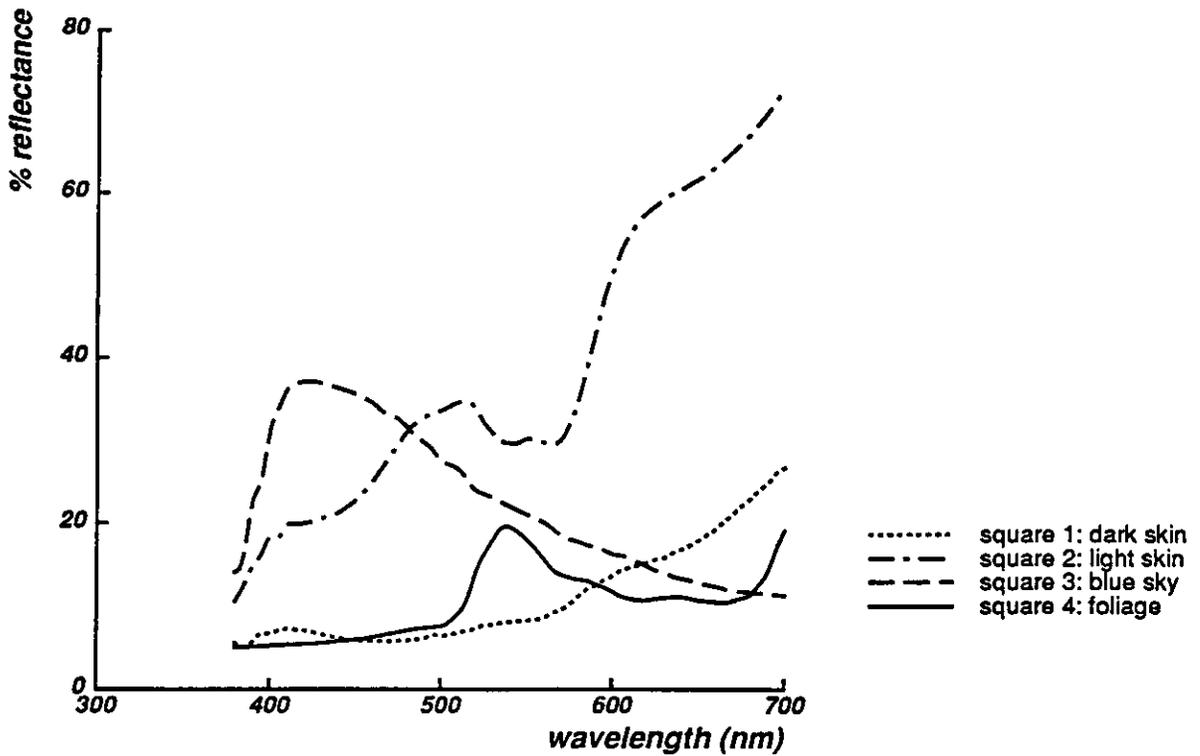


Figure 3: Spectral reflectances of the first four squares of a hypothetical color chart shown in Figure 4. Notice that the effect of the infrared cut-off filter is to limit the responsivity of the imaging system to visible wavelengths of light.

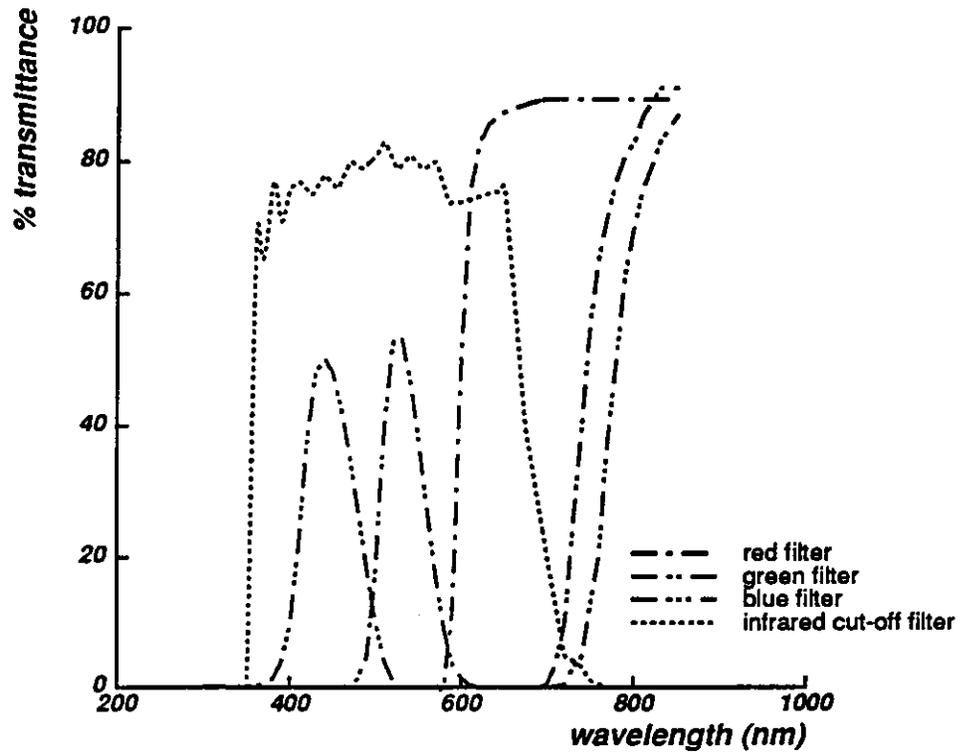


Figure 4: Spectral transmittances of filters

We modeled the spectral responsivity of our simulated camera as being that of a typical published sensitivity for CCD technology [2]. CCD chips are known to exhibit an increase in sensitivity with increasing wavelength, being most sensitive to infrared wavelengths. The sensitivity for our hypothetical camera is shown in Figure 5.

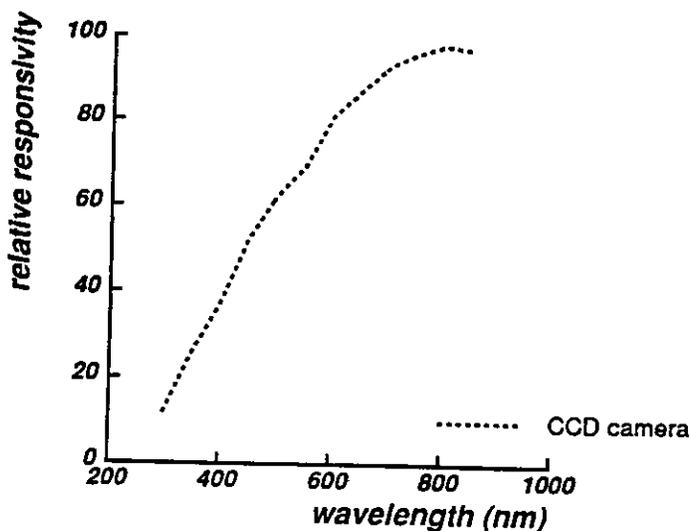


Figure 5: Spectral responsivity of a hypothetical CCD camera

In our practical experience with a CCD camera, its relative insensitivity to shorter wavelengths causes images taken with a green filter to have a lower SNR than those taken with a red filter, and those taken with a blue filter to have a still lower SNR. A simple way to deal with this is to change the aperture of the camera at the same time the filter is changed. We call this method "aperture balancing" [20]. Fortunately it is extremely simple to model: if the aperture is opened up one stop when changing from the red filter to the green filter, the amount of light entering the camera is doubled. If the aperture is opened up an additional stop when changing to the blue filter, the amount of light is doubled again. Figure 6 shows how the aperture ratios of (1:2:4) for the red, green, and blue filters makes the resulting sensor responsivity functions more nearly equal in magnitude. (Recall that the sensor responsivity function is equivalent to the filter transmittance times the camera responsivity.) The resulting sensor responsivities are not identical to NTSC primaries; they represent functions that we believe are close to the way our equipment behaves.

We obtained the typical spectral reflectance curves of the Ceramic Colour Standard tiles which we selected as the "unknown objects" for the simulation. These are standard tiles used for checking spectroradiometers, and their spectral curves are widely published [3]. The spectral reflectance curves for several of these tiles are shown in Figure 7.

Finally, we collected spectral power distribution functions for two classes of hypothetical illuminants. To simulate lights that might be obtained indoors, we used the CIE standard illuminant A with or without hypothetical filters placed in front of it. Illuminant A is supposed to represent light from a full radiator at 2856K; it is similar to what one can expect from incandescent lights. The hypothetical filters were based on filters sold by Edmund Scientific Co. [5] These filters are available in 8 inch by 10 inch sheets which may be placed in front of a spotlight. The spectral curves were obtained from the manufacturer. The hypothetical incandescent light without any filter and two lights that would result from using selected filters are shown in Figure 8.

In order to capture in our simulation the way outdoor lighting varies from day to day, we used theoretical daylight curves generated by the standard CIE method of calculating daylight illuminants with different correlated color temperatures [30]. A daylight power distribution $L_{day}(\lambda)$ may be calculated by

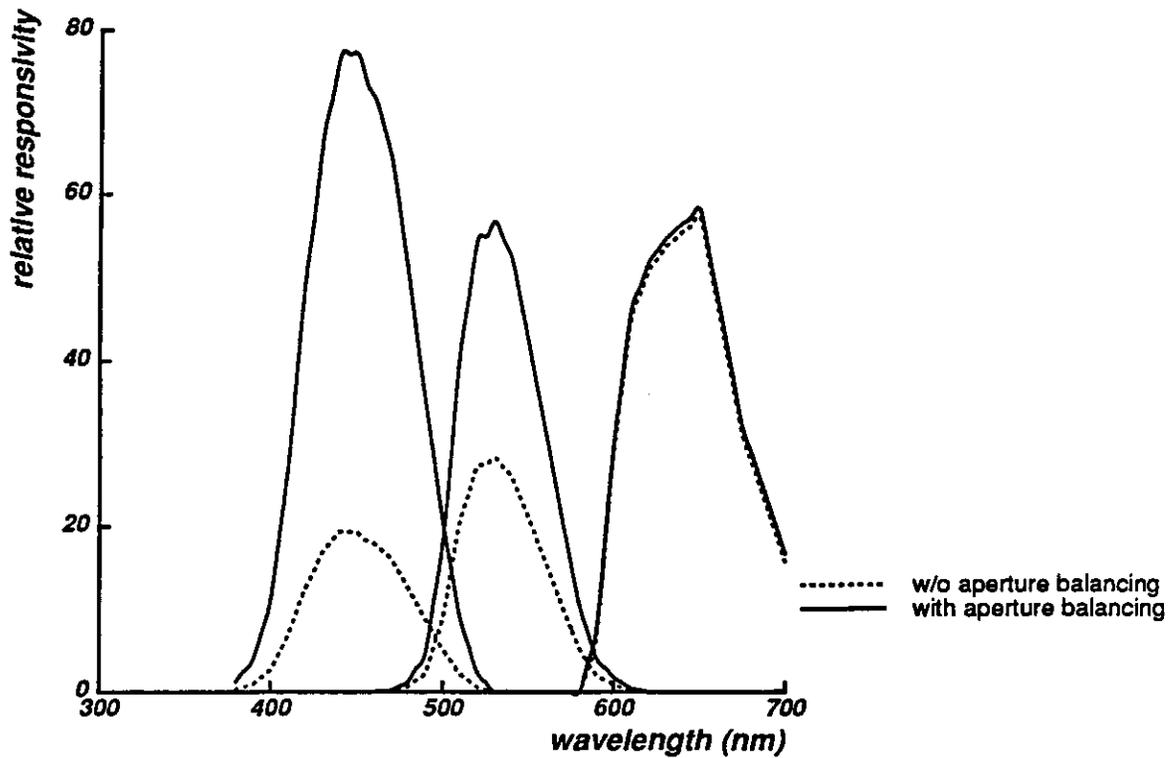


Figure 6: Comparison of sensor responsivities with and without aperture balancing

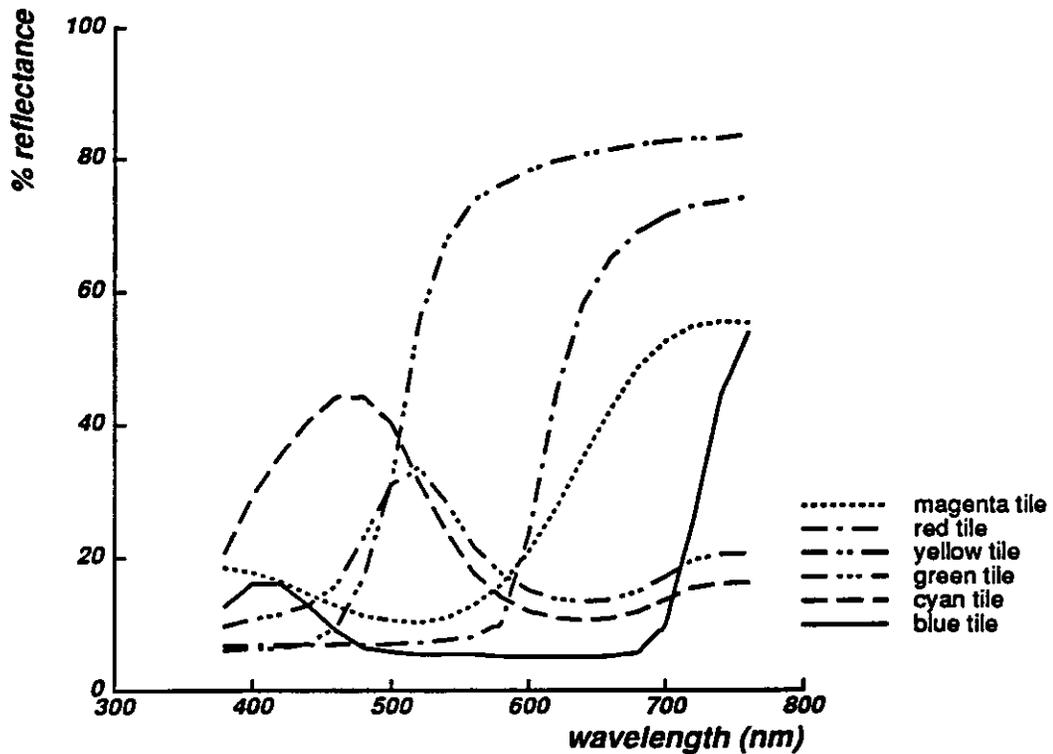


Figure 7: Spectral reflectances of six color standard tiles

$$L_{day}(\lambda) = S_0(\lambda) + M_1 S_1(\lambda) + M_2 S_2(\lambda)$$

where $S_0(\lambda)$ is the calculated mean spectral distribution of several hundred measured daylight samples, and $S_1(\lambda)$

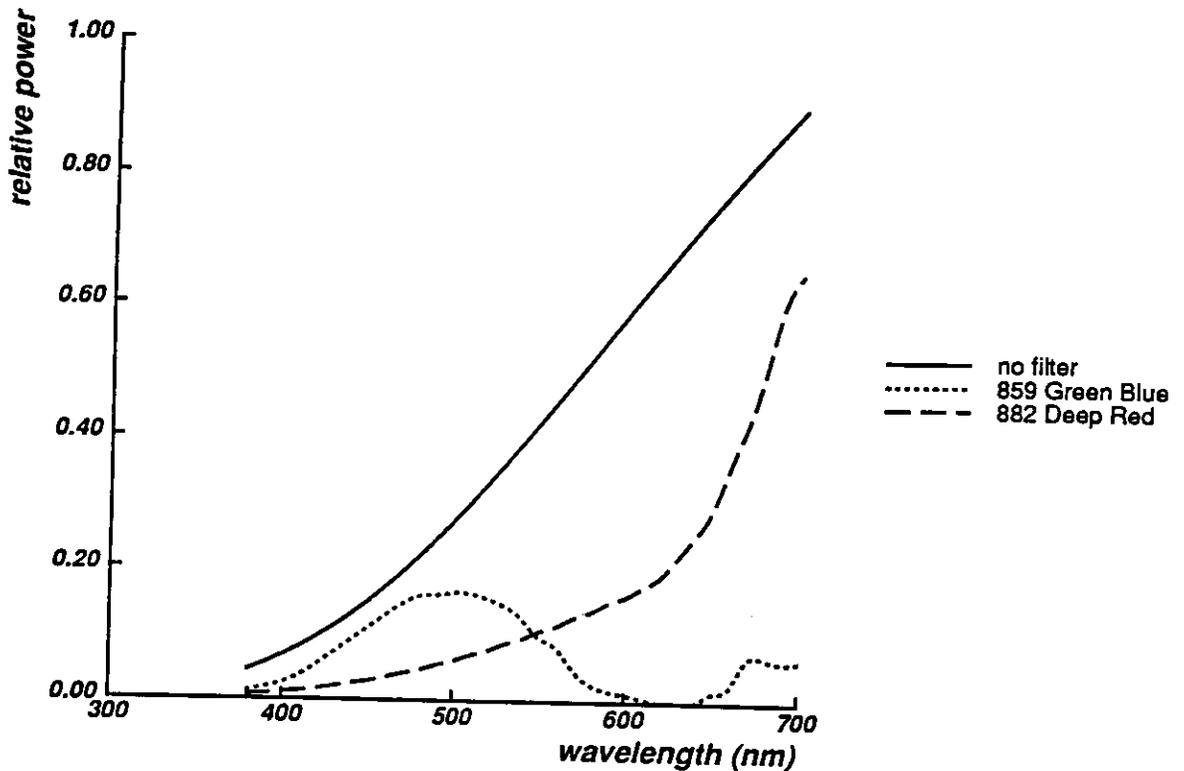


Figure 8: Spectral power distribution of hypothetical indoor lights

and $S_2(\lambda)$ are the two most important eigenvectors of the set. M_1 and M_2 are calculated from the chromaticity coordinates x and y corresponding to the position on the CIE chromaticity chart where the color of the given light would appear.

$$M_1 = \frac{-1.3515 - 1.7703x + 5.9114y}{0.0241 + 0.2562x - 0.7341y}$$

$$M_2 = \frac{0.0300 - 31.4424x + 30.0717y}{0.0241 + 0.2562x - 0.7341y}$$

The chromaticity coordinates in turn can be calculated from a given correlated color temperature T_c by

$$x = -4.6070 \frac{10^9}{T_c^3} + 2.9678 \frac{10^6}{T_c^2} + 0.09911 \frac{10^3}{T_c} + 0.244063 \quad (\text{for } 4000K \leq T_c < 7000K)$$

$$x = -2.0064 \frac{10^9}{T_c^3} + 1.9018 \frac{10^6}{T_c^2} + 0.24748 \frac{10^3}{T_c} + 0.237040 \quad (\text{for } 7000K \leq T_c < 25000K)$$

$$y = -3.000 x^2 + 2.870 x - 0.275$$

These curves are supposed to approximate natural daylight under a wide range of conditions, including direct sunlight, skylight, and daylight under various degrees of cloud coverage. The CIE method of calculation allows us to generate a whole family of curves that approximate daylight illumination using only a single input parameter, namely correlated color temperature. A few of the resulting curves are shown in Figure 9. These curves are *relative* spectral power distributions and so were derived to have the common value of 100 at the 560nm wavelength.

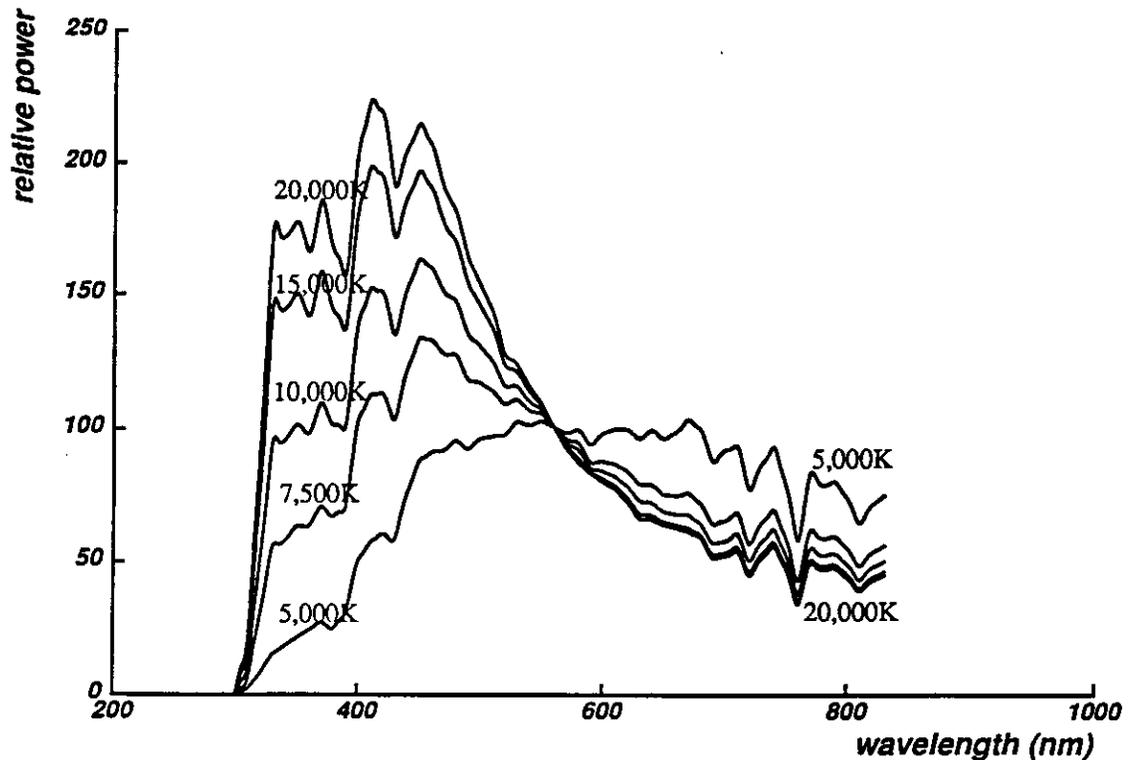


Figure 9: Spectral power distribution of CIE daylight illuminants

3.3 Evaluation Method

To evaluate the success of a color constancy method, researchers need to state what precisely is the goal of the method, and then develop a criterion for measuring how close they have come to that goal. We have examined two related yet different goals and propose a measure of success for each.

One goal is to evaluate the accuracy with which we can predict an object's color under a new type of light. In other words, after estimating the spectral power distribution of the illumination or the spectral reflectance function of an object, how close is the color we predict from these estimates to the true color? Therefore we need a measure of the distance between two colors.

Traditionally "color" is represented by a point in three-space, where the three dimensions may be lightness, hue, saturation; or CIE coordinates X Y Z; or RGB. We shall use RGB to refer to the values measured by our camera through a red filter, green filter, and blue filter respectively. The actual values of RGB may change if we change cameras or filters.

In reporting how far a predicted color value is from the true value, it is possible to use distance between the two points. For a given color value (R, G, B) and its predicted color value (R_p, G_p, B_p) the error E could be measured by

$$E = (|R - R_p|^n + |G - G_p|^n + |B - B_p|^n)^{1/n}$$

for some value n . In particular, if $n = 1$, then this is the sum of the errors in each band. If $n = 2$, then this is the euclidean distance between the two values. If $n = \text{infinity}$ then this is the maximum of the three errors.

The problem with this error measure is that it is biased against bright colors. If we are viewing mostly dark surfaces, the change in their measured values from one light source to another will be small, and so even an ineffective method of color constancy will yield a low error. Therefore we prefer to use a percentage error measure

that takes into account how much room there is for error. We will use the error calculation

$$E = \frac{(|R-R_p|^n + |G-G_p|^n + |B-B_p|^n)^{1/n}}{(|R|^n + |G|^n + |B|^n)^{1/n}}$$

and choose $n = 1$. Note that this is not the same as normalizing the RGB values for intensity and then comparing the chromaticity difference. This error measure points out differences in intensity as well as differences in hue and saturation.

The CIE color difference formula based on $L^*u^*v^*$ coordinates is often used to give a measure of how far apart colors look to humans [7]. However, since we are concerned with computer vision applications, our formula relates directly to the RGB values measured by the computer system.

Another possible goal for our system would be to measure how well it performs as a spectroradiometer since the immediate output of the least squares calculation or the network is a spectral curve. Since the color appearance predictions are calculated from these spectral curves, a more accurate spectral curve should yield a more accurate color. So the error measure of interest is the difference between the predicted spectral curve and the actual one.

A simple difference measure for this criterion is the difference of the spectral power of the two curves at selected wavelengths. These differences can be squared and then added together to compute a total error measure for the predicted curve ΔL .

$$\Delta L = \sum_{\lambda=\lambda_0}^{\lambda_k} (L(\lambda) - L^*(\lambda))^2$$

However, this formulation treats the contribution from all wavelengths equally. If the spectral function is to be used in calculating other quantities, such as color appearance, accuracy at the extremes of the visible spectrum may be much less important than accuracy in the center of the visible spectrum where sensors are much more responsive. Accordingly, we also examine the sum of squared differences weighted by the human luminous efficiency function $V(\lambda)$ shown in Figure 10.

$$\Delta_w L = \sum_{\lambda=\lambda_0}^{\lambda_k} V(\lambda)(L(\lambda) - L^*(\lambda))^2$$

We have found that weighted spectral error corresponds more closely than the unweighted error to the color appearance error measure E .

3.4 Results Using Least Squares Method

Our simulation assumed 24 colored squares and 3 color filters which yield 72 constraints upon the illuminant. However, the eigenvector analysis showed that most of that information was redundant in a camera that can only detect 1 part in 100 reliably (which is typical of our equipment). Eigenvector analysis of the matrix R formed with our hypothetical color chart, filters, and camera showed that at most 10 coefficients are significant if the camera is assumed to have a signal-to-noise ratio of 100:1. If the camera had a SNR of 1000:1 then 21 coefficients could be reliably calculated. If the aperture balancing method is not used, only 8 coefficients are considered reliable in a camera with SNR of 100:1. This coincides with our real world observation that images taken without aperture balancing tend to have highly unreliable values in the blue band and somewhat unreliable values in the green band.

Figures 11 and 12 show a comparison of the recovered illuminant with the original illuminant in a simulation made with noise-free data and using the least squares method with 10 basis functions. One of the lights in the simulation was an indoor light (with a green-blue filter) and the other was an outdoor light. All the spectral curves

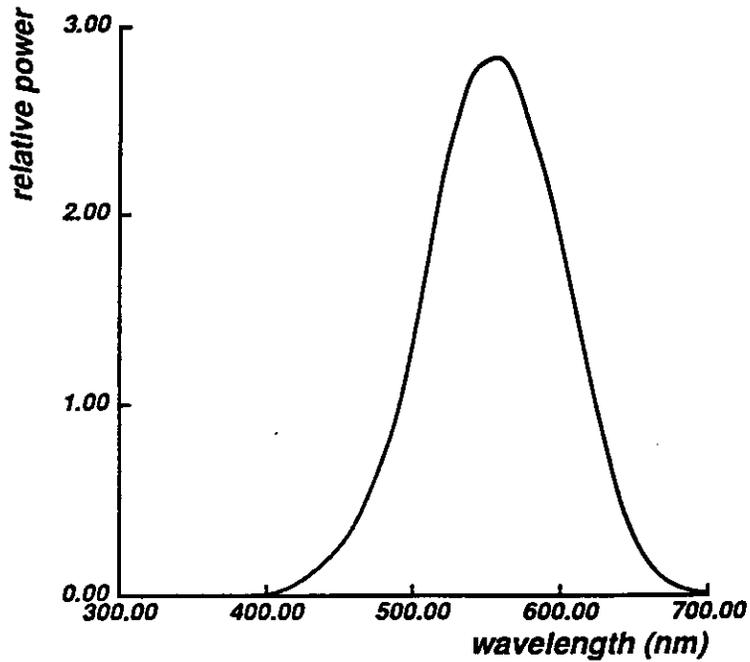
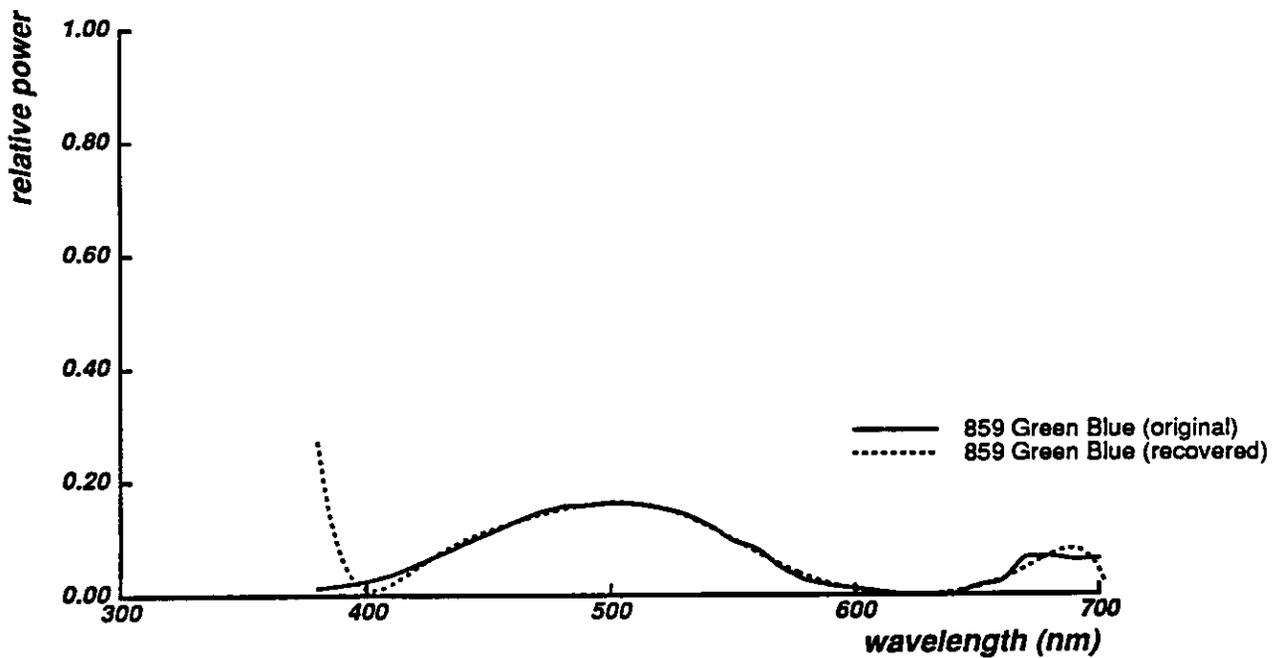


Figure 10: Weighting Function $V(\lambda)$

corresponding to the original light functions have been scaled to be between 0 and 1.



Least Squares Method (10 basis functions)

Figure 11: Comparison of indoor light with recovered illuminant in noise-free simulation

Figure 13 shows why unit basis functions would be a disastrous choice with the least squares method. As long as there is no noise in the simulation, the method will recover the indoor light with no filter perfectly. However, if noise is added to the simulation, the resulting predicted spectral function will oscillate widely. This is due to multicollinearity, as discussed in section 2.1.3, and is a typical result when least squares estimation is used with

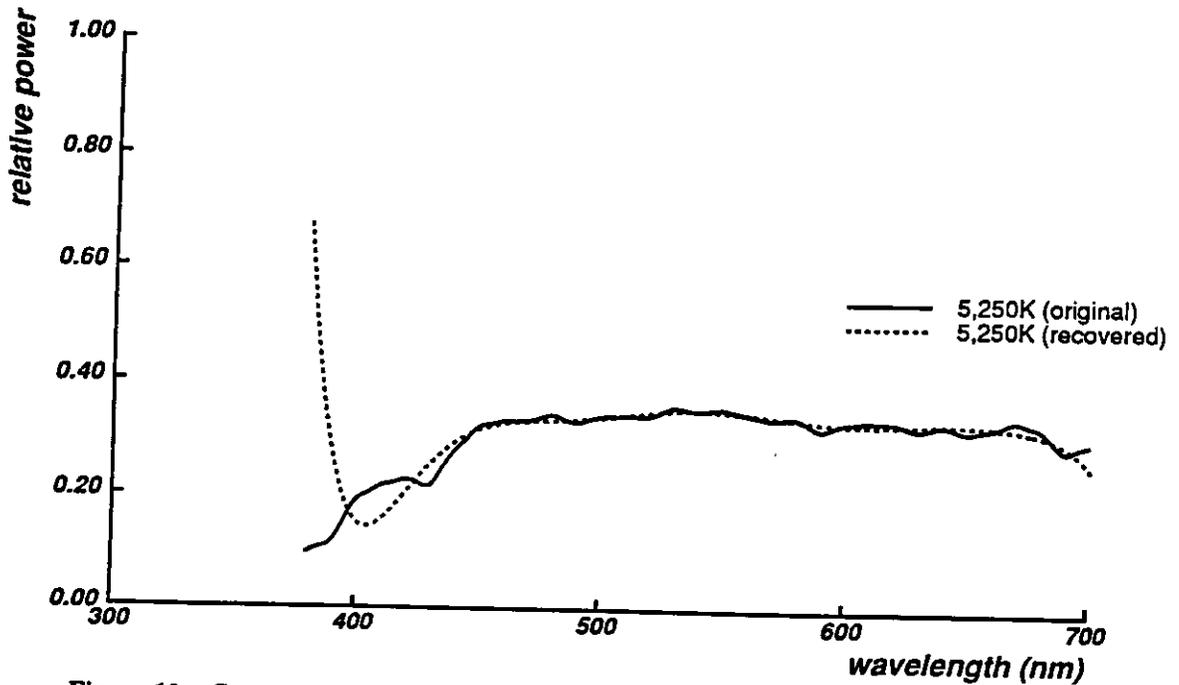


Figure 12: Comparison of outdoor light with recovered illuminant in noise-free simulation correlated basis vectors.

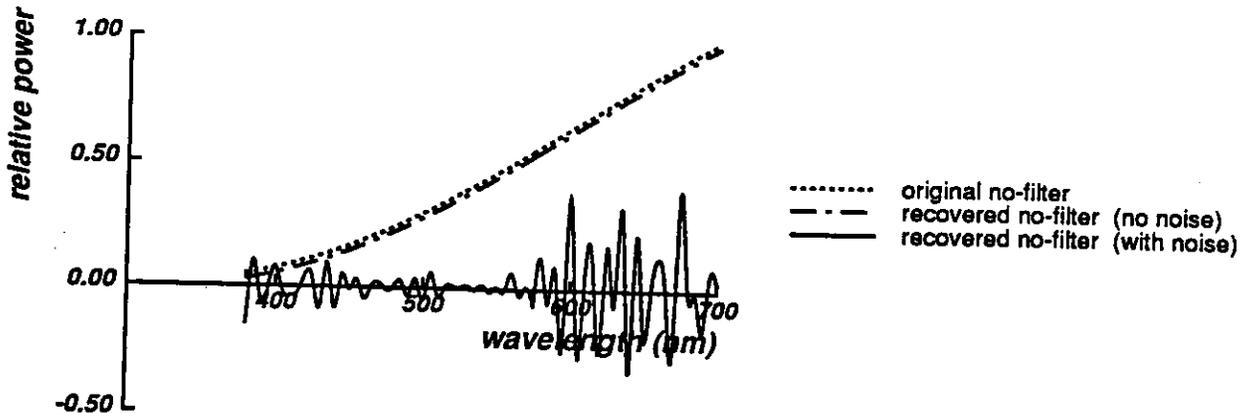


Figure 13: Recovery of indoor light with unit basis functions (function with noise has been scaled by 0.01 so that it will fit on the same graph)

Figures 14 and 15 show the recovery of an indoor and outdoor light using the least squares technique with 10 basis functions that consist of Legendre polynomials. In each case we added Gaussian noise to the hypothetical camera measurements. Each measurement was calculated according to Equation (1-2) and then added to a randomly generated noise value with mean value 0 and standard deviation equal to the maximum camera value divided by the hypothetical SNR (100:1). Figures 16 and 17 show recovery with an empirically derived number of basis functions. That is, we tried the least squares calculation repeatedly, varying the number of basis functions, to find the number that would give the lowest weighted error $\Delta_w L$. The empirical number was 3 for the outdoor lights (these originally came from 3 weighted functions) and 5 for the indoor lights. Surprisingly, the smaller number of basis functions does better than a larger number when noise is introduced into the calculation. The empirically derived numbers also give the smallest color difference E . (The unweighted spectral differences ΔL give a different experimental answer which leads to a large color difference.)

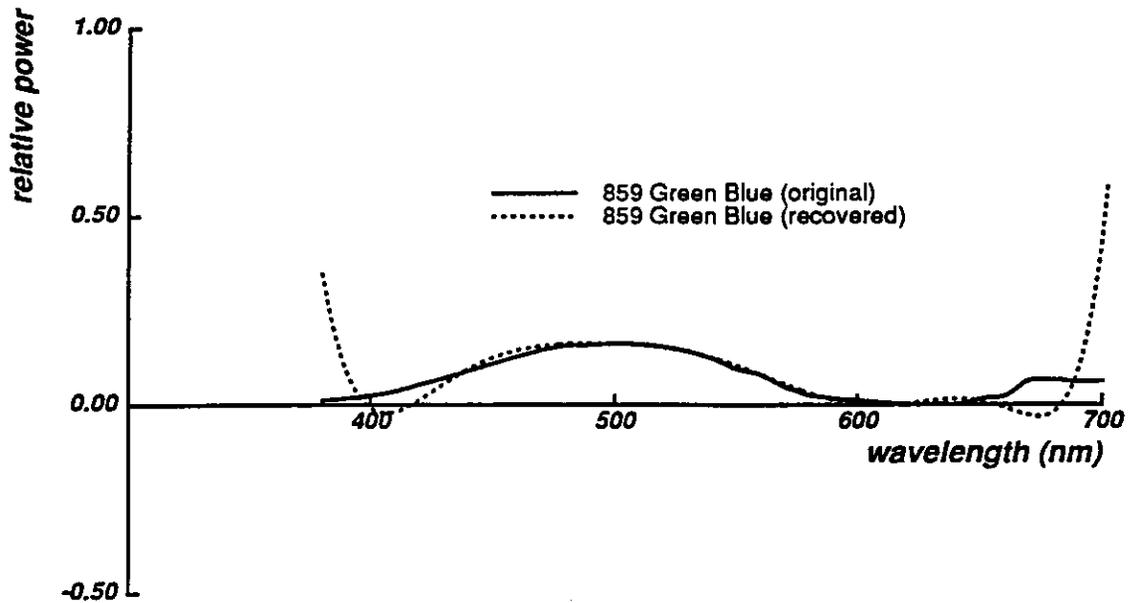


Figure 14: Noisy simulation of indoor light using 10 basis functions

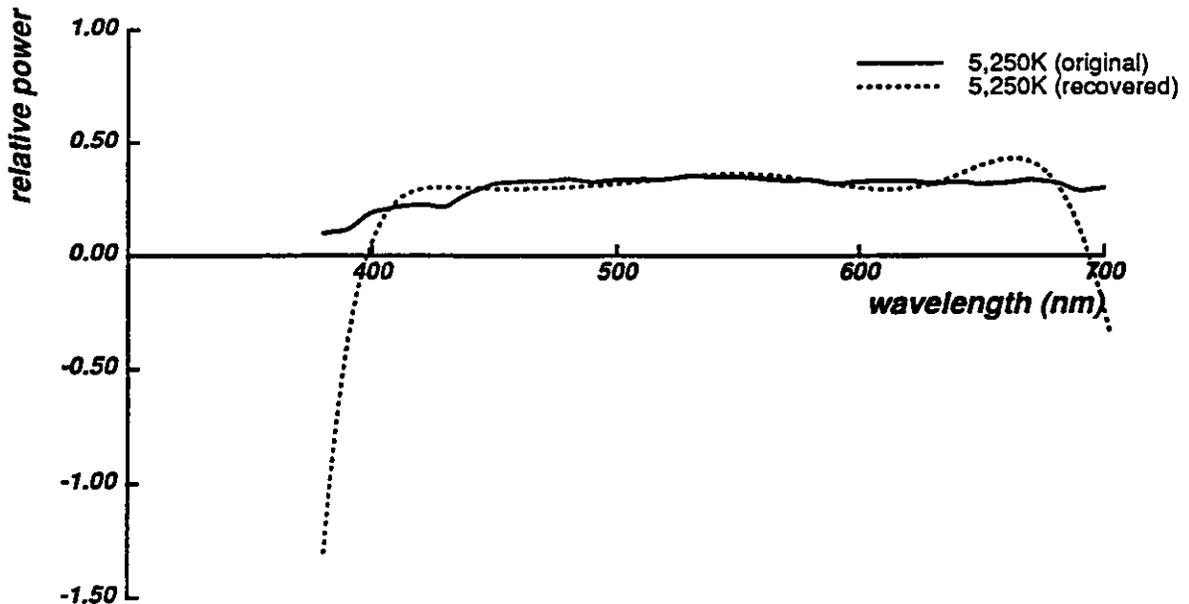


Figure 15: Noisy simulation of outdoor light using 10 basis functions

Our explanation for this less-is-more result is that the larger number of basis functions are magnifying the errors in the simulated camera measurements. Perhaps a different set of basis functions would not suffer so much from this problem. The greatest errors occur at the extremes of the visible range where the differences may be less important for many applications.

3.5 Results Using Neural Network Method

To train the neural network we used 70 outdoor light samples or 47 indoor light samples. The network was then tested on a small number of light samples it had not seen before, although of the same type (indoor or outdoor). The number of hidden units was varied in the different experiments, including some trials with no hidden layer at all.

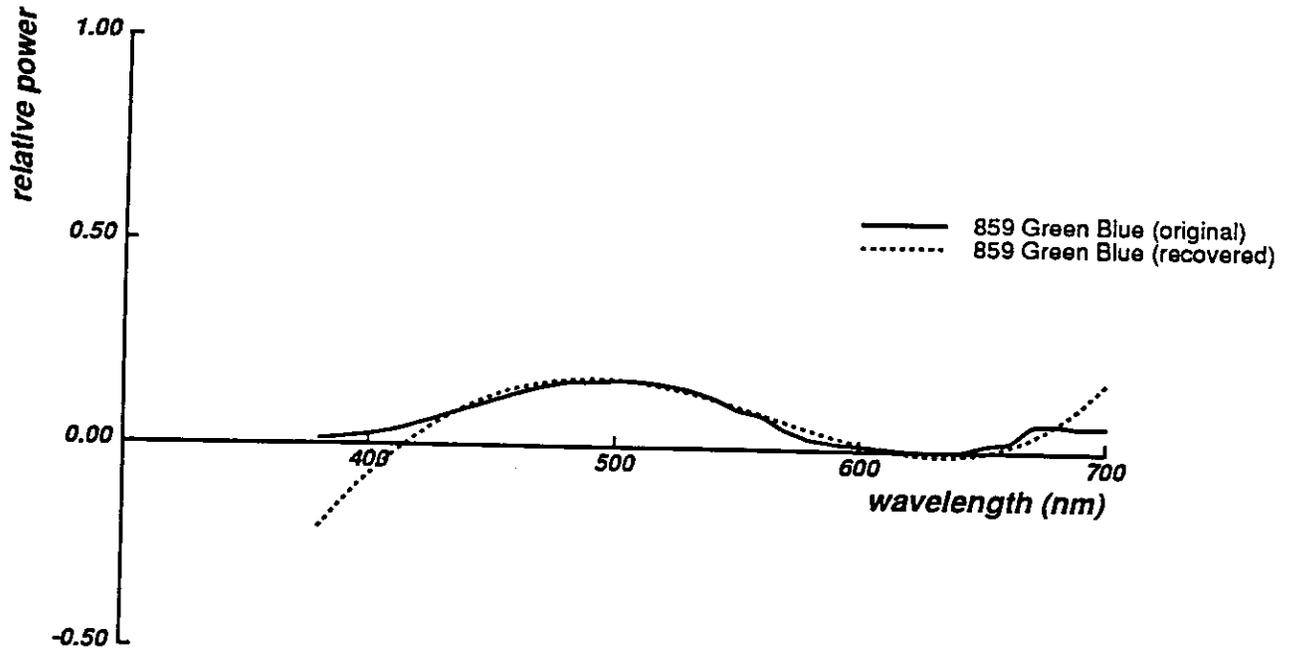


Figure 16: Noisy simulation of indoor light using 5 basis functions

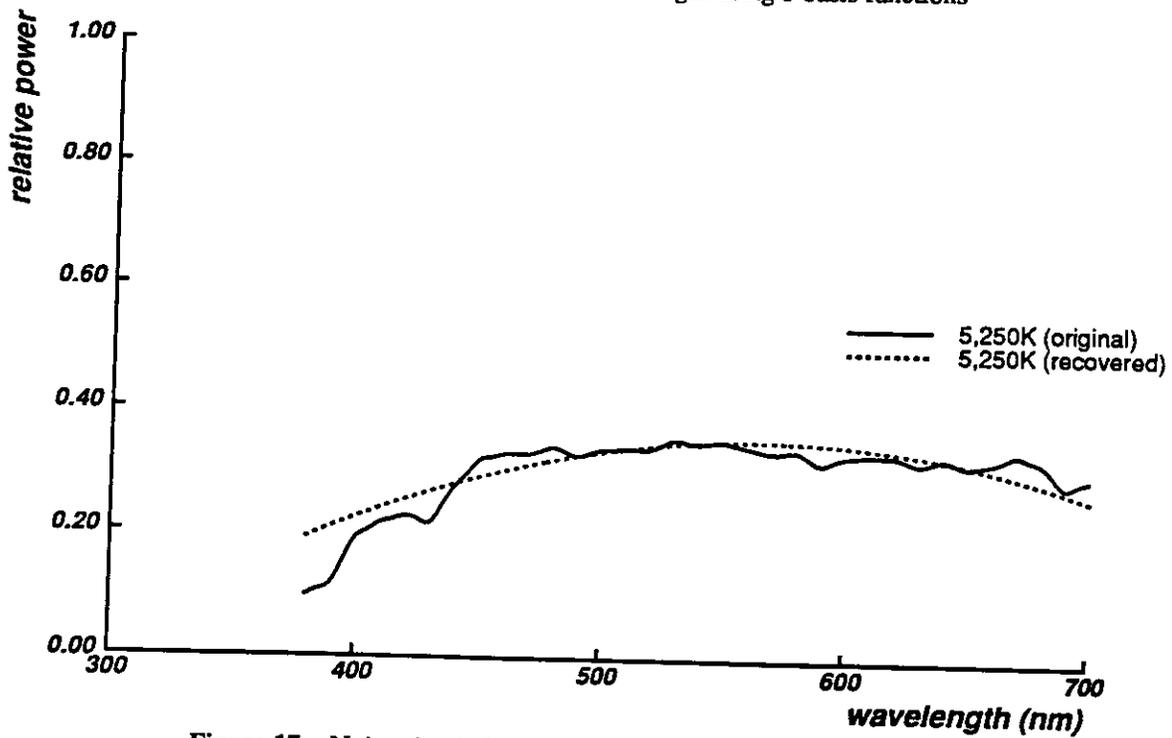


Figure 17: Noisy simulation of outdoor light using 3 basis functions

All networks were trained with simulated camera measurements with noise added.

Figures 18 and 19 show the results from the network, comparing the original light curves with the recovered lights. The results are shown for experiments where no hidden layer was used, and where a hidden layer with 16 units was used. For the outdoor lights, the hidden layer did not make much difference in the accuracy of the network in recovering spectral curves. With the indoor lights, the hidden layer gave a big advantage.

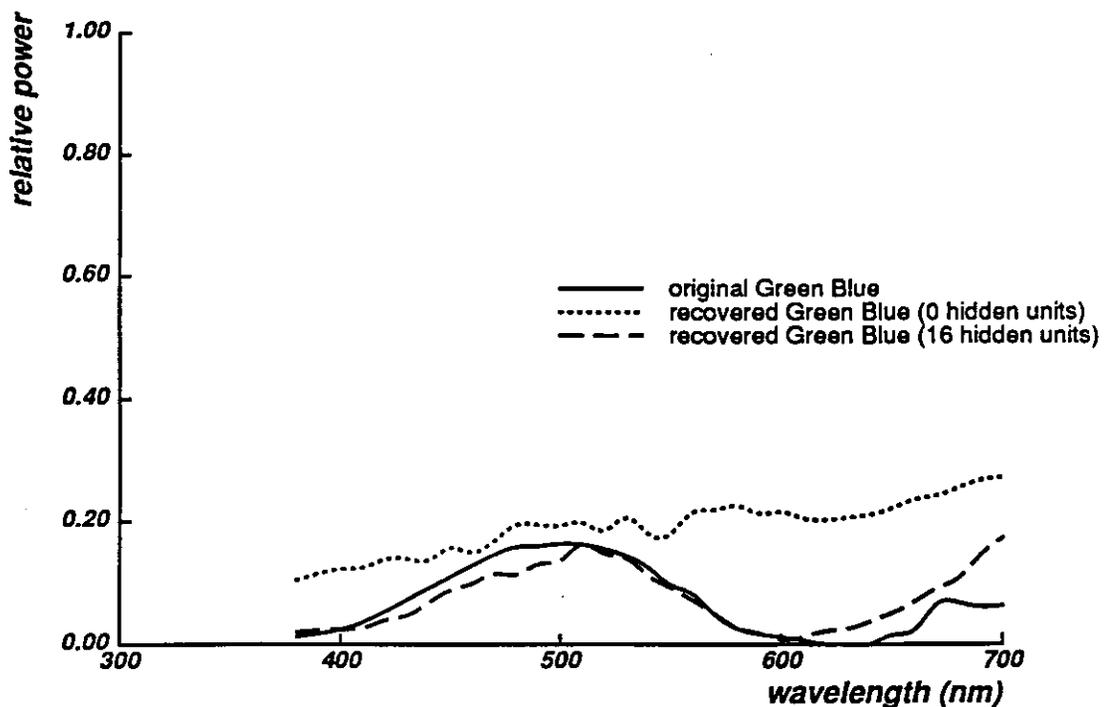


Figure 18: Recovery of indoor light using network method

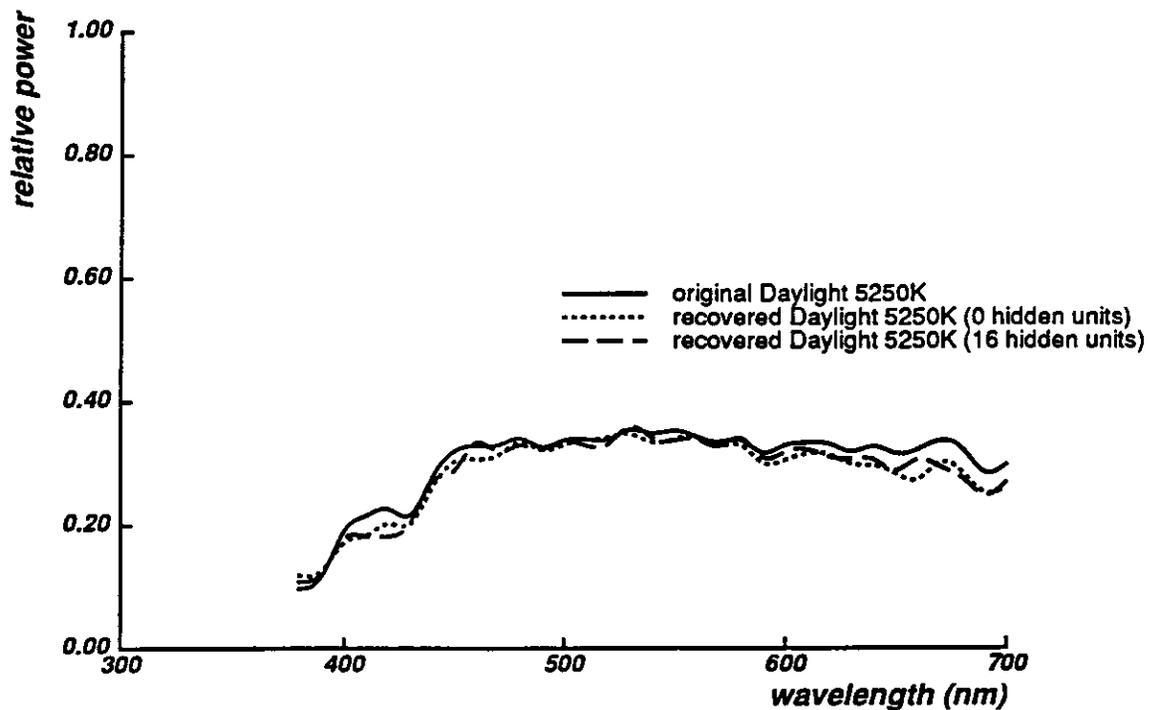


Figure 19: Recovery of outdoor light using network method

3.6 Comparison of Least Squares Method with Neural Network

Figures 20 and 21 show the overall average error in recovering illuminants as a function of the number of basis functions used in the least squares method. This is shown using both the unweighted and weighted error measurements ΔL and $\Delta_w L$, average for all the illuminants used in the experiment. As can be seen, the lowest error occurs at a different number of basis functions for the two types of lights. After this optimum number of basis functions, adding more functions does not help, and in fact hurts the recovery of the illumination. The correct number of functions to be used depends upon the application. If the goal is to recover the illuminant that is closest to the original one, then the unweighted error would be a good measure; it would prescribe that 2 basis functions be used for outdoor lights and 3 for indoor lights. If however the goal is to predict color appearance based upon the calculated illumination, then the weighted error might be a more appropriate indicator. As is seen in Figure 22, the method comes closest to predicting the correct color when 3 basis functions are used for outdoor lights and 5 for indoor lights.

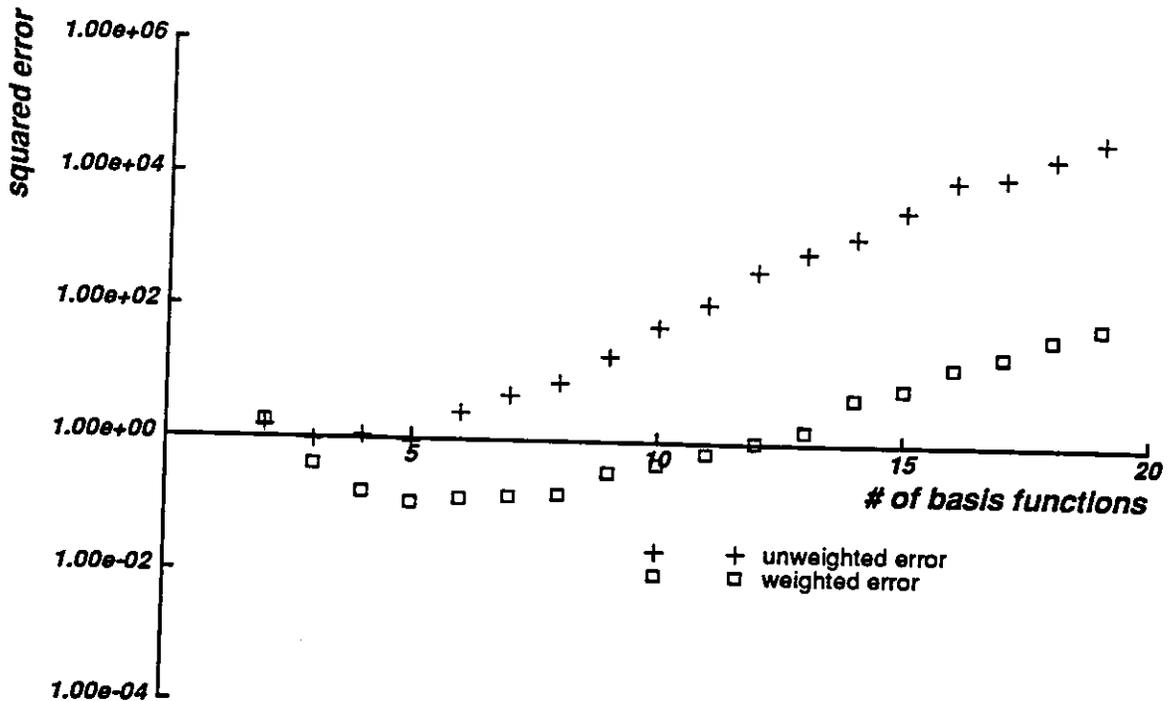


Figure 20: Average error ΔL and $\Delta_w L$ for indoor lights using least squares method

Figures 23 and 24 show the average error in recovering illuminants as a function of the number of hidden units used in the neural network. Again, both the unweighted and weighted squared errors are shown, although for the network the two measures are very similar. The error in recovering outdoor lights is very low and the addition of a hidden layer makes very little difference. The network does a better job with the outdoor lights than the least squares method using any number of basis functions.

For the indoor lights, a hidden layer gives a significant improvement. Although we found the lowest error with 16 hidden units, the addition of units beyond 4 gave only a very slight improvement. The best least squares method (with 5 basis functions) performed slightly better than the network method in terms of weighted squared error. If unweighted error is used, then the network performed significantly better than all the least squares methods used.

The neural network method has several advantages over the least squares method, which help explain its success. The outputs from a network are constrained to lie between 0 and 1. This property is highly useful for calculating values that correspond to real-world measurements that cannot be negative. Since spectral power distribution

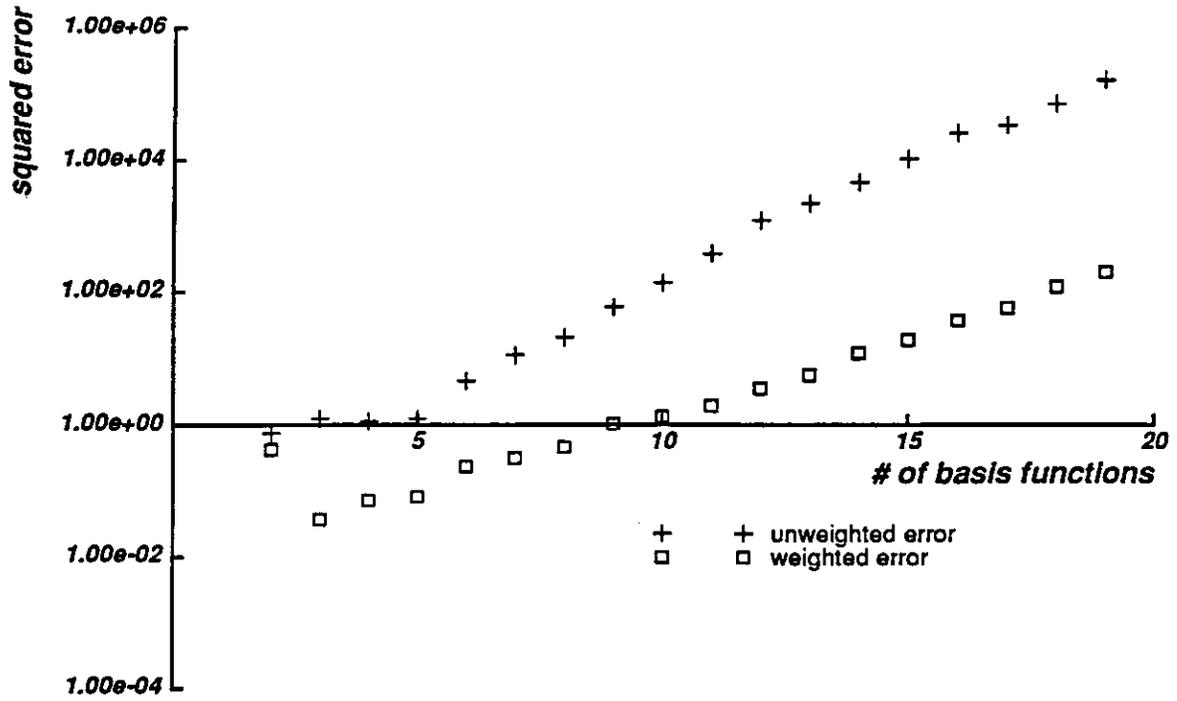


Figure 21: Average error ΔL and $\Delta_w L$ for outdoor lights using least squares method

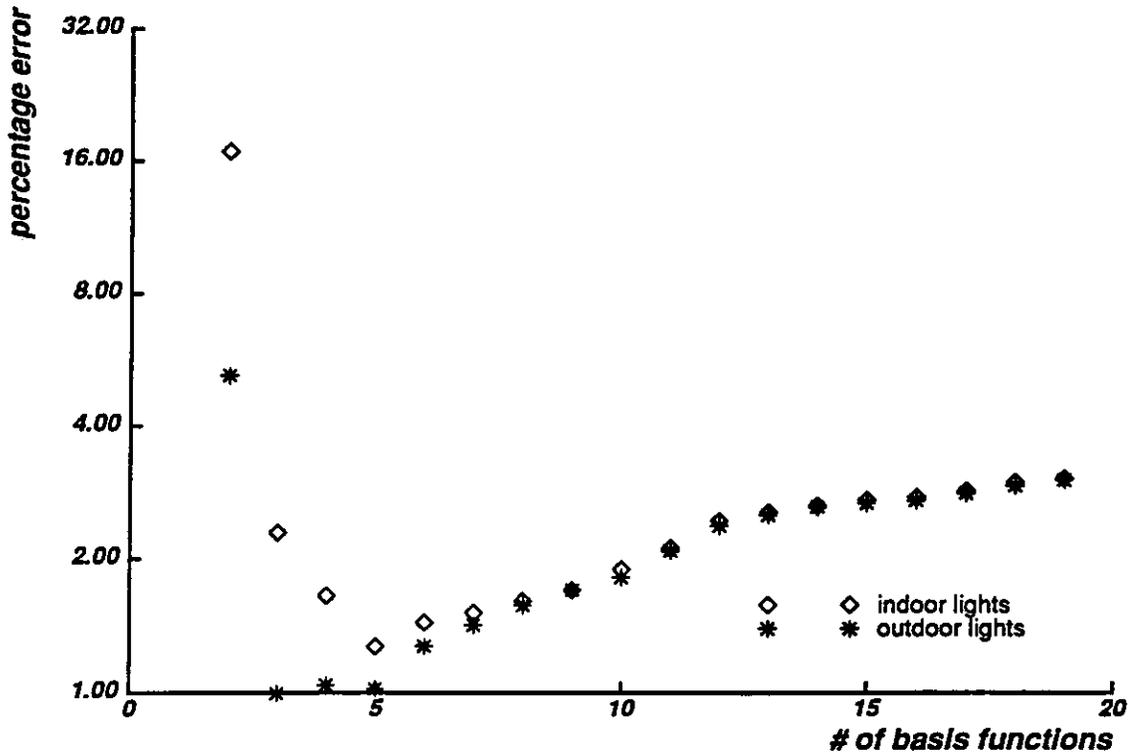


Figure 22: Average error E for indoor and outdoor lights using least squares method

functions cannot ever have negative values, this gives the network an advantage over the least squares method formulation which can and does compute negative values. The network's output constraints also prevent it from calculating the spurious "tails" seen in the least squares result. These "tails", especially dramatic in Figures 14 and 15, are large deviations from the original curve seen at the extremes of the visible range. The tails are most likely an

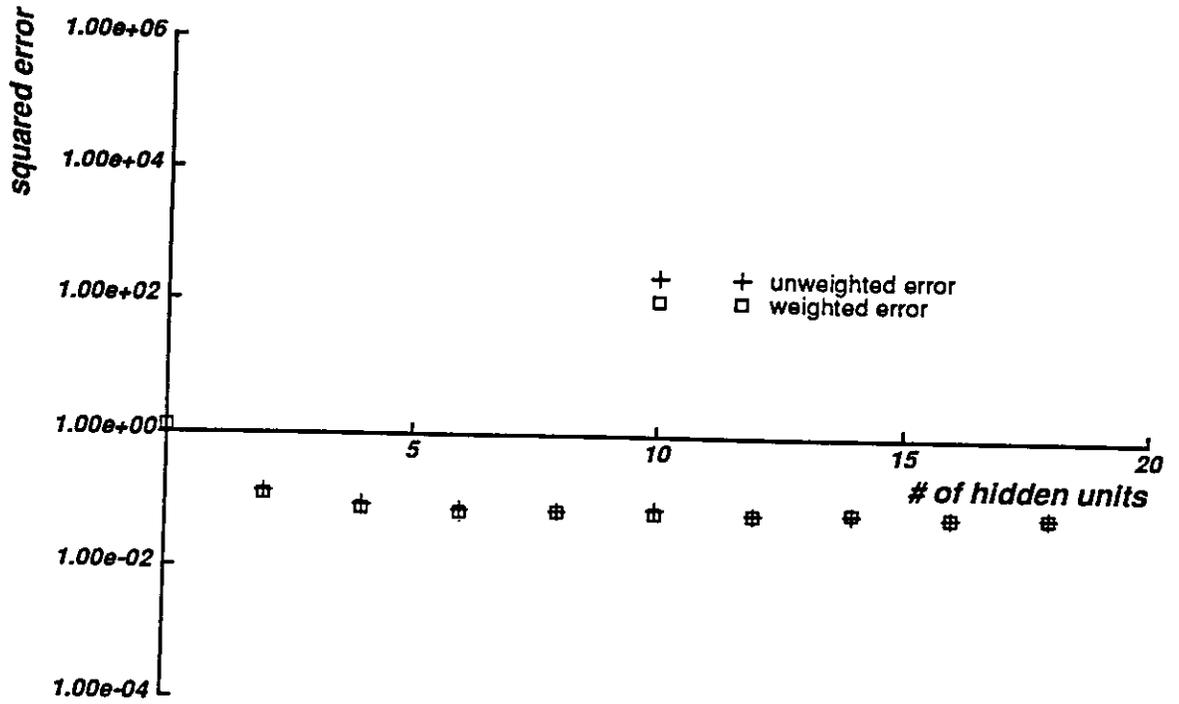


Figure 23: Average error ΔL and $\Delta_w L$ for indoor lights using network method

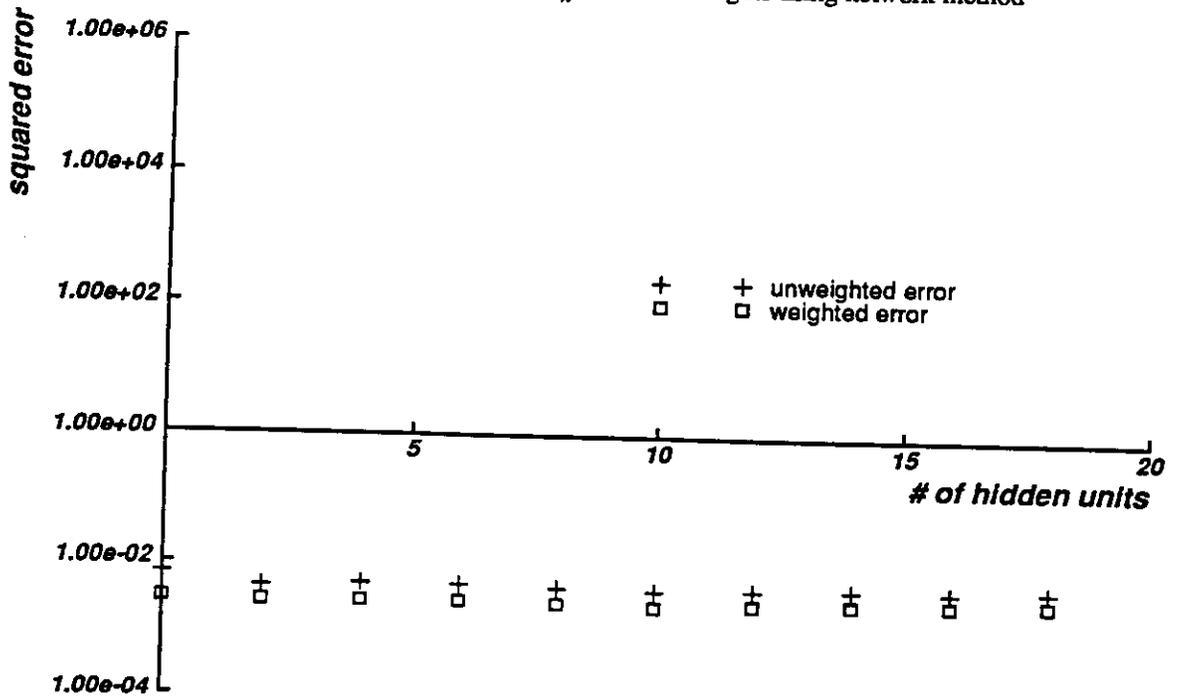


Figure 24: Average error ΔL and $\Delta_w L$ for outdoor lights using network method

artifact of using higher power Legendre polynomials. The values at the extremes remain +1 or -1, but the values in the center of the spectrum become progressively smaller.

Another advantage the network has is that it is free to calculate its own representation of the data, whereas the least squares method is constrained to use the basis functions we chose. Although the network was tested on

illuminants on which it was not trained, the test illuminants were of a similar nature to the training illuminants.

This is particularly true of the simulated outdoor lights since they were all generated using a mean spectral distribution and two eigenvectors. Although these are supposed to be representative of the range of outdoor illumination, this probably refers more to the general spectral trend (high in blue vs. high in red); real outdoor illumination probably varies more from one wavelength to the next. The indoor lights were based on real spectral curves of commercially available filters. Although these indoor lights have smoother curves than the simulated outdoor lights, they are probably not so easily captured in a small number of basis functions. It is possible that the least squares method might do significantly better if we first examined the training data used by the network, calculated the most significant eigenvectors of the spectral data, and used those as basis functions rather than the Legendre polynomials.

A disadvantage of the neural network method is of course the time and effort it takes to acquire the training data and the amount of time it takes to train the network adequately. However, once this has been done the network will operate very quickly.

4. A Complete Method for Supervised Color Constancy

In the experiments just presented, we assumed that the reflectance functions of objects of interest were already known to us. Thus, given an accurate estimate of the illumination, we can predict the appearance of an object under that type of illumination by multiplying the estimated illumination by the known sensor functions and object reflectance and integrating over the visible wavelengths. However in general vision scenarios the reflectances of interesting objects in the scene will not be known in advance. Therefore it is necessary to have a way to estimate object reflectances in addition to the illumination.

4.1 Estimation of the Reflectance Function Using Multiple Illuminants

We propose that surface reflectance functions might be estimated in much the same way that we estimated the illuminant. In the methods described for calculating the illumination, constraints upon the illuminant were generated by observations of the illuminant's effect upon known surfaces on the color chart (or by learning the relationship between spectral power of the light and chart appearance, in the case of the network). With such estimates of the illuminant in hand, constraints upon a surface reflectance function can be generated by observations of the surface illuminated by known or estimated illuminants.

As before, we start with the equation describing the measurement made by the sensor, but this time using the estimated illuminant in the calculation

$$m = \int L^*(\lambda) \rho(\lambda) s(\lambda) d\lambda \quad (4-1)$$

This time we choose to represent the unknown surface reflectance function, rather than the illuminant, in terms of some basis functions. Although we use the Legendre polynomials for representing surface reflectances just as we did for illuminants, different basis functions could be chosen [6].

$$\rho(\lambda) = \sum_{j=1}^n \rho_j b_j(\lambda)$$

and substituting into Equation (4-1):

$$m_i = \sum_{j=1}^n \rho_j \int b_j(\lambda) L^*(\lambda) s_i(\lambda) d\lambda$$

where the i subscript refers to one of the three filters. If we group together the known quantities, letting $\Lambda_{ij} = \int b_j(\lambda) L^*(\lambda) s_i(\lambda) d\lambda$ then the equation can be rewritten simply as

$$m_i = \sum_{j=1}^n \rho_j \Lambda_{ij}$$

We let M be the vector of the sensor measurements m_i under one type of light and let Λ be the matrix with entries of Λ_{ij} , yielding a simple matrix equation

$$M = \Lambda \rho$$

which is then solved for ρ , the vector of reflectance coefficients ρ_j .

If there are only three sensors and one light source then the equation only provides three constraints upon the estimate of the reflectance coefficients ρ . In the previous sections we described how we were able to increase our knowledge about an unknown light source by looking at its influence upon several known surfaces. Now we can turn this around and increase the number of constraints upon an unknown surface reflectance by looking at the influence of several known, or estimated, illuminants.

We construct Λ with the first three rows consisting of

$$\begin{bmatrix} \int b_1(\lambda) s_R(\lambda) L^*_1(\lambda) d\lambda & \dots & \int b_n(\lambda) s_R(\lambda) L^*_1(\lambda) d\lambda \\ \int b_1(\lambda) s_G(\lambda) L^*_1(\lambda) d\lambda & \dots & \int b_n(\lambda) s_G(\lambda) L^*_1(\lambda) d\lambda \\ \int b_1(\lambda) s_B(\lambda) L^*_1(\lambda) d\lambda & \dots & \int b_n(\lambda) s_B(\lambda) L^*_1(\lambda) d\lambda \end{bmatrix}$$

where $L^*_1(\lambda)$ is the estimate of the first illuminant under which the unknown surface is viewed. The next three rows of Λ are calculated with the estimate for the second illuminant, and so on, so that for m estimated illuminants there are $3m$ rows in the matrix. In the general case where there are s sensors there are sm rows in the matrix.

The measurement vector M is constructed in a similar fashion. The first three entries in M are the RGB values measured by the camera looking at the surface under the first illuminant, the second three entries correspond to the second illuminant, and so on in the same order that the rows of Λ were calculated.

As before, the number of basis functions n used in the representation is chosen so that $n < 3m$ since there may be some redundancy in the illuminants. Eigenvector analysis of Λ will yield an approximate value for n . The least squares method is used to solve for ρ yielding an estimated solution ρ^{**} . (The double star indicates that the estimate was calculated using data that had also been estimated. If the surface reflection was estimated using known illuminants rather than estimated ones, it could be designated ρ^* .) The estimated surface reflection function $\rho^{**}(\lambda)$ is then recovered by multiplying the coefficients by the appropriate basis functions

$$\rho^{**}(\lambda) \approx \sum_{j=1}^n \rho_j^{**} b_j(\lambda)$$

4.2 Summary of Proposed Method for Supervised Color Constancy

The entire process of estimating illuminants and object reflectances is outlined here. Given

1. Color chart of known properties
2. Several illuminants of unknown properties
3. Object(s) of unknown properties

Then for each illuminant

1. Take picture of color chart
2. Calculate $L^*(\lambda)$ using either least squares or neural network method
3. Take picture of object(s)

and for each object

1. Form vector of accumulated measurements and matrix of estimated illuminants
2. Calculate $\rho^{**}(\lambda)$

Object reflectance estimates can be calculated repeatedly as new pictures of the objects under new light sources are acquired.

If the experiment is taking place indoors, then it may be difficult to acquire enough picture under different types of illumination. One way might be to use colored filters or combinations of filters in front of a spotlight. However, in an outdoor setting the lighting can change from day to day and from moment to moment. Our proposed method not only can accommodate the changing illumination, but actually uses it to incrementally improve its knowledge of unknown reflectance functions. Changes in illumination thus need no longer be considered a source of noise but rather a useful source of information about the scene.

4.3 Incremental Color Constancy

The color constancy method described above has an annoying limitation: the color chart must be re-introduced into the scene every time the light changes. This is because the color chart is used to estimate the illuminant. However, if there are objects in the scene whose reflectances have been previously calculated with the method above, these objects might also be used as reference objects for estimating the illuminant.

An application where this idea might be used is a vehicle moving outdoors. As the vehicle moves along, new objects will come into view, but at the same time some objects that have been seen before will still be present. The outdoor illumination may also be changing as the vehicle drives along, so the estimate of the illuminant will need to be revised continuously. This can be done by using those object present in a number of images.

The procedure would consist of taking pictures at frequent intervals. For each picture the algorithm would:

1. Determine which objects visible in this image have been seen previously. These objects will already have estimated spectral reflectance functions that were calculated from previous observations.
2. Use least squares estimation to simultaneously update estimated spectral reflectances of these reference objects and calculate estimated spectral power distribution of the illuminant at this moment in time.
3. For objects that have not been previously seen, use the estimated illuminant to calculate an initial estimate of the object's reflectance.

We call this process "incremental color constancy", which adapts to both a (slowly) changing illuminant and a (slowly) changing scene. The color chart would be needed only at the start, to get an initial estimate of the illuminant and to derive initial object reflectance estimates.

5. Conclusions

In this paper we have presented a mathematical formulation of color constancy that shows the limitations of traditional approaches to color constancy. We also classified modern color constancy methods according to the reference criterion chosen and the number of basis functions used and argued that the traditional criterion may be too weak and the typical number of basis functions too few to give an adequate representation of spectral functions.

We propose a new approach that we call "supervised color constancy" that uses a color chart of known properties to provide information about illuminants. The presence of the chart in an image provides an extremely reliable reference criterion, and the large number of colored squares provide enough constraints to calculate more basis coefficients than other color constancy methods.

We have shown simulated results for two methods of calculating the illuminant: least squares estimation and neural networks. Both these methods were tested on data with simulated noise to try to give a more realistic scenario. Under these conditions the neural network which was trained on similar illuminants tended to perform better than the least squares estimation technique which used Legendre polynomials as basis functions.

This paradigm might be used to accomplish color constancy by calculating a progressively better estimate of the spectral reflectance of unknown objects under changing illumination. This method would use the illumination estimates calculated by the above methods to impose constraints upon the reflectance functions of objects. The implementation and testing of this method remains the subject of future work. Taken together, these methods comprise a powerful new approach for reliable and accurate color recognition in robot vision.

Acknowledgments

Thanks to Barak Pearlmutter for his advice and the use of his neural network simulator and to Dave Touretzky for his advice and encouragement.

References

- [1] R. Bajcsy, S. W. Lee, and A. Leonardis.
Image Segmentation with Detection of Highlights and Inter-reflections Using Color.
MS-CIS 89-39, University of Pennsylvania, Philadelphia, June, 1989.
- [2] W. Budde.
Optical Radiation Measurements. Volume 4: Physical Detectors of Optical Radiation.
Academic Press, New York, 1983.
- [3] F. J. J. Clarke and P. R. Samways.
The spectrophotometric properties of a selection of ceramic tiles".
Metrology Centre Report MC2, National Physical Laboratory, Teddington, England, August, 1968.
- [4] D'Zmura, M. and P. Lennie.
Mechanisms of color constancy.
J. Opt. Soc. Am. 3(10):1622-1672, October, 1986.
- [5] Edmund Scientific Co.
No. 9081 Spectral Curves.
Barrington, NJ, .
- [6] R. Gershon.
The Use of Color in Computational Vision.
PhD thesis, University of Toronto, 1987.
- [7] F. Grum and C. J. Bartleson (editors).
Optical Radiation Measurements.
Academic Press, New York, 1980.
- [8] G. Healey and T. O. Binford.
The Role and Use of Color in a General Vision System.
In *Proceedings of the ARPA Image Understanding Workshop*, pages 599-613. DARPA, 1987.
- [9] G. Healey and T. O. Binford.
A Color Metric for Computer Vision.
In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10-17. IEEE, New York, 1988.
- [10] J. Ho, B. V. Funt, and M. S. Drew.
Disambiguation of Illumination and Surface Reflectance from Spectral Power Distribution of Color Signal: Theory and Applications.
CSS/LCCR 88-18, Simon Fraser University, Burnaby, B.C., 1988.
- [11] A. C. Hurlbert and T. A. Poggio.
Learning a Color Algorithm from Examples.
Proceedings of Neural Info. Processing Systems , 1987.
- [12] A. C. Hurlbert and T. A. Poggio.
Synthesizing a Color Algorithm from Examples.
Science 236:482, January, 1988.
- [13] R. M. Johnston.
Color Theory.
Pigment Handbook.
John Wiley & Sons, Inc., New York, 1974, Chapter III-D-b.
- [14] Land, Edwin H.
The Retinex Theory of Color Vision.
Scientific American 237(6):108-128, December, 1977.

- [15] Macbeth.
Macbeth ColorChecker.
trademarked product.
1977
- [16] Maloney, Laurence T.
Evaluation of linear models of surface spectral reflectance with small numbers of parameters.
J. Opt. Soc. Am. 3(10):1673-1683, October, 1986.
- [17] L. T. Maloney and B. A. Wandell.
Color constancy: a method for recovering surface spectral reflectance.
J. Opt. Soc. Am. A 3(1):29-33, January, 1986.
- [18] Munsell Color.
Munsell Book of Color.
MacBeth Division of Kollmorgen Corporation, Baltimore, 1976.
- [19] C. L. Novak and S. A. Shafer.
Color Vision.
Encyclopedia of Artificial Intelligence.
John Wiley & Sons, Inc., New York, 1990.
- [20] C. L. Novak, S. A. Shafer, and R. G. Willson.
Obtaining Accurate Color Images for Machine Vision Research.
Proc. SPIE , 1990.
- [21] J. P. S. Parkkinen, J. Hallikainen, and T. Jaaskelainen.
Characteristic spectra of Munsell colors.
J. Opt. Soc. Am. A 6(2):318-322, February, 1989.
- [22] D. E. Rumelhart, G. E. Hinton, J. L. McClelland.
A general framework for parallel distributed processing.
Parallel Distributed Processing: Explorations in the microstructures of cognition. Volume 1: Foundations.
Bradford Books/MIT Press, Cambridge, MA, 1986.
- [23] D. E. Rumelhart, G. E. Hinton, R. J. Williams.
Learning internal representations by error propagation.
Parallel Distributed Processing: Explorations in the microstructures of cognition. Volume 1: Foundations.
Bradford Books/MIT Press, Cambridge, MA, 1986.
- [24] S. A. Shafer.
Describing light mixtures through linear algebra.
J. Opt. Soc. Am. 72(2):299-300, February, 1982.
- [25] S. A. Shafer.
Using Color to Separate Reflection Components.
Color Research and Application 10(4):210-218, Winter, 1985.
- [26] S. Tominaga and B. A. Wandell.
Standard surface-reflectance model and illuminant estimation.
J. Opt. Soc. Am. A 6(4):576-584, April, 1989.
- [27] S. Tominaga and B. A. Wandell.
Component estimation of surface spectral reflectance.
J. Opt. Soc. Am. A 7(2):312-317, February, 1990.
- [28] B. A. Wandell.
The Synthesis and Analysis of Color Images.
IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-9(1):2-13, January, 1987.
- [29] G. Wyszecki.
Evaluation of metameric colors.
J. Opt. Soc. Am. 48:451-454, 1958.

- [30] G. Wyszecki and W. S. Stiles.
Color Science: Concepts and Methods, Quantitative Data and Formulae.
Wiley, New York, 1982.