

Exploiting Spatial and Spectral Image Regularities for Color Constancy

Barun Singh and William T. Freeman
MIT Computer Science and Artificial Intelligence Laboratory

July 2, 2003

Abstract

We study the problem of color constancy—inferring from an image the spectrum of the illumination and the reflectance spectra of all the depicted surfaces. This estimation problem is underdetermined: many surface and illumination spectra can be described as a linear combination of 3 basis functions, giving 3 unknowns per pixel, plus 3 parameters for a global illumination. A trichromatic visual system makes fewer measurements than there are unknowns.

We address this problem by writing the reflectance spectra of small groups of pixels as a linear combination of spatio-spectral basis functions. These aggregated surface reflectance spectra require fewer parameters to describe than the sum of the spectral parameters for the individual surface pixels, giving us more measurements than unknown parameters.

We explore this problem in a Bayesian context, showing when the problem is over or underdetermined based on analyzing the local curvature characteristics of the log-likelihood function. We show that using the spatio-spectral basis functions gives improved reflectance and illumination spectral estimates when applied to real image data.

1 Introduction

Color is important in our understanding of the visual world and provides an effective cue for object detection and recognition. In general, however, the observed color of an object differs from its true color due to factors such as lighting and orientation. Color constancy refers to the ability to perceive the color of an object as approximately constant regardless of the color of the light falling upon it [7]. We are able to reliably use color for a variety of vision tasks largely because of our ability to perform color constancy well.

Inherently, the color constancy problem is underdetermined. The image formation process consists of an illumination reflecting off of a surface, then being passed through a set of sensors. The input to the sensors is the product of the illumination spectrum and the surface's reflectance

spectrum. Different combinations of illumination and reflectance spectra can produce the same spectrum impinging on the sensor, and various sensor inputs can result in the same sensor responses. Due to these ambiguities, it may be impossible to uniquely separate the effect of the illumination and the surface reflectances in an image. The goal of computational color constancy [18, 14, 4, 15, 3, 17, 10, 11, 19, 2, 6], therefore, is to determine an optimal separation of reflectance and illumination under some metric. Given the sensor responses, we seek an estimate of the reflectance and illumination spectra.

This would appear to involve estimating very many numbers. If we specify a 300 nm spectrum at 10 nm intervals, even assuming a single illumination spectrum over the whole image, we would need $30(N + 1)$ numbers to specify the reflectance spectra for N pixels plus one illuminant. In order to simplify the color constancy problem, low-dimensional linear models are used to describe illumination and reflectance spectra [13, 5, 16, 12]. A sufficiently large portion of the energy in these spectra (typically 99%) - can be described using as low as three-dimensional models for both surfaces and illuminants. An additional constraint used by some approaches is that surfaces and illuminants must be physically realizable, e.g. their spectra cannot be negative and surface reflectances must be less than 1 [10, 2, 6].

Even after these simplifications, the color constancy problem remains underdetermined and existing approaches must make further assumptions in order to obtain a solution. Buchsbaum's Gray World algorithm assumes that the mean reflectance of all images is the same, and the illumination is estimated using this mean [4]. Maloney and Wandell's Subspace method requires that a two-dimensional model describe surface reflectances for the case of trichromatic sensors [17]. Gamut mapping methods [9, 8] exploit the observed color gamut to estimate the illuminant spectrum, and can be combined with realizability constraints. Since a Bayesian approach [2] can incorporate the various assumptions above into its prior probabilities, Bayesian decision theory [1] provides a principled framework which we will use for studying the color constancy problem.

A parameter counting argument reveals the underlying problem we address in this work: if we use a 3-dimensional model for the surface reflectance at each pixel, and a 3-dimensional model for the illumination spectrum, for the case of trichromatic color sensors, we have more unknowns than observations. We measure 3 numbers at each position, yet have to estimate 3 numbers at each position, plus the 3 numbers describing the illuminant. The problem is always underdetermined, requiring, in a Bayesian solution, more reliance on the prior probabilities or loss function than on the likelihood function to estimate the best answer. One would prefer the situation where the data likelihood function dominated the other terms.

In this paper, we make the problem overdetermined by introducing linear basis functions that describe both spectral and spatial variations in the surface reflectances. These spatio-spectral basis functions exploit image regularities which allow us to specify surfaces reflectances using fewer numbers, yielding more image measurements than unknown param-

ters.

In the next sections, we introduce the notation and linear basis function analysis, then explore the Bayesian approach. We show with simple examples that the structure of the posterior indicates conditions under which color constancy is still undetermined by the data. We then apply our approach to real images, verifying (using hyperspectral image data) that our surface reflectance estimates improve through using the spatio-spectral basis functions.

2 Linear models

We assume there is one global illumination in each scene and that the light spectrum leaving each surface is the term-by-term product of the illumination and local reflectance spectrum. For a surface reflectance $S(\lambda)$ and illumination $E(\lambda)$, the response at position x of a photoreceptor with a spectral response of $R_k(\lambda)$ is:

$$y_k^x = \sum_{\lambda} R_k(\lambda) E(\lambda) S^x(\lambda), \quad (1)$$

The illuminant and surface spectra can be written as linear combinations of the illumination basis functions $E_i(\lambda)$ and reflectance basis functions $S_j(\lambda)$, with coefficients e_i and s_j^x , respectively, at position x :

$$E(\lambda) = \sum_{i=1}^L E_i(\lambda) e_i, \quad (2)$$

$$S^x(\lambda) = \sum_{j=1}^L S_j(\lambda) s_j^x, \quad (3)$$

where L is defined as the number of elements in the λ vector (i.e., the number of wavelength samples). Doing so allows us to write the rendering equation as

$$y_k^x = \sum_{\lambda} R_k(\lambda) \sum_{i=1}^L E_i(\lambda) e_i \sum_{j=1}^L S_j(\lambda) s_j^x. \quad (4)$$

Summing over λ , we get a bilinear form,

$$y_k^x = \sum_{i=1}^L \sum_{j=1}^L e_i G_{ij,k} s_j^x, \quad (5)$$

where $G_{ij,k} = \sum_{\lambda} R_k(\lambda) E_i(\lambda) S_j(\lambda)$.

In general, we wish to approximate the true illuminant $E(\lambda)$ and surface spectra $S^x(\lambda)$

using lower dimensional linear models. Thus, we can define

$$\tilde{E}(\lambda) = \sum_{i=1}^{d_E} E_i(\lambda) \mathbf{e}_i, \quad (6)$$

$$\tilde{S}^x(\lambda) = \sum_{j=1}^{d_S} S_j(\lambda) \mathbf{s}_j^x, \quad (7)$$

where d_E is the dimensionality of the illuminant approximation $\tilde{E}(\lambda)$ and d_S is the dimensionality of the surface approximation $\tilde{S}^x(\lambda)$. In addition, we can decompose (5) as:

$$\begin{aligned} y_k^x &= \sum_{i=1}^{d_E} \sum_{j=1}^{d_S} \mathbf{e}_i G_{ij,k} \mathbf{s}_j^x + \sum_{i=1}^{d_E} \sum_{j=d_S+1}^L \mathbf{e}_i G_{ij,k} \mathbf{s}_j^x + \sum_{i=d_E+1}^L \sum_{j=1}^{d_S} \mathbf{e}_i G_{ij,k} \mathbf{s}_j^x + \sum_{i=d_E+1}^L \sum_{j=d_S+1}^L \mathbf{e}_i G_{ij,k} \mathbf{s}_j^x \\ &= \sum_{i=1}^{d_E} \sum_{j=1}^{d_S} \mathbf{e}_i G_{ij,k} \mathbf{s}_j^x + w_k^x, \end{aligned} \quad (8)$$

where w_k^x is a noise term representing the error due to the use of lower dimensional models. A summary of the notation used in this paper is provided in Table 1.

Equation (8) is perhaps best understood by viewing color constancy as an estimation problem. The goal is to estimate $\tilde{S}^x(\lambda)$ and $\tilde{E}(\lambda)$ given y_k^x . Thus, we view the “true” illuminants as having d_E degrees of freedom as defined by the first d_E illuminant basis functions, with remaining basis functions contributing in the form of noise. Analogously for surface reflectance, the first d_S basis functions define the space of “true” reflectances and the remaining basis functions add noise. Note that while the illumination “noise” and reflectance “noise” are orthogonal to the “true” illumination and reflectance, the projections of the “true” data and the “noise” into the sensor response space are not orthogonal.

3 Bayesian Approach to Color Constancy

The Bayesian approach to color constancy utilizes a probabilistic framework to examine the problem. This framework is based on three fundamental probability distributions - the prior, the likelihood, and the posterior. The prior describes the probability that a certain set of parameters is the correct estimate without knowing anything about the data. This distribution is written as $P(\mathbf{e}, \mathbf{s})$, where $\mathbf{e}$ and $\mathbf{s}$ are the global illumination weights and the surface weights for all positions in an image. The likelihood, given by $P(y|\mathbf{e}, \mathbf{s})$, governs the relationship between the parameters being estimated ($\mathbf{e}$ and $\mathbf{s}$) and the observed data (in this case the set of all sensor responses in an image, y). The posterior describes what is known about the parameters being

Symbol	Symbol Definition	Symbol	Symbol Definition
N	number of pixels in an image	m	number of pixels along one dimension of a square block being considered within an image
λ	wavelength	γ	index into spatio-spectral basis functions
L	number of wavelength samples (length of λ vector)	K	number of sensor responses available at any pixel
$E(\lambda)$	spectral radiance of illuminant	$E^{\{\}}(\gamma)$	spatio-spectral radiance of illumination in $m \times m$ block
$E_i(\lambda)$	i th basis function of illumination	$\hat{E}_i(\gamma)$	i th spatio-spectral illumination basis function
e_i	scalar weight of i th illumination basis function	$e_i^{\{\}}$	scalar weight of i th spatio-spectral illumination basis function
$\mathbf{e}$	row vector of all e_i	$\mathbf{e}^{\{\}}$	row vector of all $e_i^{\{\}}$
d_E	dimensionality of illumination representation	$\hat{d}_E$	dimensionality of spatio-spectral illumination representation
$\tilde{E}(\lambda)$	approximation to illuminant from d_E dimensional linear model	$\tilde{E}^{\{\}}(\gamma)$	approximation to spatio-spectral illuminant from $\hat{d}_E$ dimensional linear model
$S^x(\lambda)$	surface reflectance at point x	$S^{\{x\}}(\gamma)$	spatio-surface reflectance for pixel block $\{x\}$
$S_j(\lambda)$	j th basis function of surface reflectance	$\hat{S}_j(\gamma)$	j th spatio-spectral surface reflectance basis function
s_j^x	scalar weight of j th surface reflectance basis function at position x	$s_j^{\{x\}}$	weight of j th spatio-spectral surface reflectance basis function for pixel block $\{x\}$
$\mathbf{s}^x$	column vector of s_j^x for all j	$\mathbf{s}^{\{x\}}$	column vector of $s_j^{\{x\}}$ for all j
$\mathbf{s}$	matrix whose columns are given by $\mathbf{s}^x$	$\mathbf{s}^{\{\}}$	matrix whose columns are given by $\mathbf{s}^{\{x\}}$
d_S	dimensionality of surface reflectance representation	$\hat{d}_S$	dimensionality of spatio-spectral surface reflectance representation
$\tilde{S}(\lambda)$	approximation to surface reflectance from d_S dimensional linear model	$\tilde{S}^{\{\}}(\gamma)$	approximation to spatio-spectral surface reflectance from $\hat{d}_S$ dimensional linear model
$R_k(\lambda)$	spectral sensitivity of k th sensor	$\hat{R}_k(\gamma)$	spatio-spectral sensitivity of k th sensor
y_k^x	scalar response of k th sensor at x	$y_k^{\{x\}}$	scalar response of k th sensor for pixel block x
$\mathbf{y}^x$	vector of y_k^x for all k	$\mathbf{y}^{\{x\}}$	vector of $y_k^{\{x\}}$ for all k
$\mathbf{y}$	matrix of $\mathbf{y}^x$ vectors for all x	$\mathbf{y}^{\{\}}$	matrix of $\mathbf{y}^{\{x\}}$ vectors for all x
$G_{ij,k}$	3-d tensor relating illumination and reflectance weights to sensor responses	$\hat{G}_{ij,k}$	3-d tensor relating spatio-spectral illumination and reflectance weights to spatio-spectral sensor responses

Table 1: Symbol Notation

estimated after having observed the data and can be calculated using Bayes' rule as

$$P(\mathbf{e}, \mathbf{s}|y) = \frac{1}{Z} P(y|\mathbf{e}, \mathbf{s}) P(\mathbf{e}, \mathbf{s}), \quad (9)$$

where the normalization constant $\frac{1}{Z}$ is independent of the parameters to be estimated. (In addition to these probabilities, the Bayesian formulation requires that we specify a cost function that defines the penalty associated with a parameter estimate of $(\tilde{\mathbf{e}}, \tilde{\mathbf{s}})$ when the true parameter values are $(\mathbf{e}, \mathbf{s})$. For estimating the relative surface and illumination spectra, independent of any overall scale factor, we find that the loss function has little effect [2], so we use MAP estimation and maximize the posterior probability.)

The underdetermined nature of the general color constancy problem implies that there should be an infinite set of solutions which maximize the posterior. To see that this is the case, we must examine the properties of equation (9).

Assuming independent and identically distributed Gaussian observation noise with variance σ^2 in the photoreceptor responses at each patch, the posterior probability is:

$$P(\mathbf{e}, \mathbf{s}|y) = \frac{1}{Z} \prod_{x,k} e^{-(y_k^x - \sum_{i,j} \mathbf{e}_i G_{ij,k} \mathbf{s}_j^x)^2 / (2\sigma^2)} P(\mathbf{e}, \mathbf{s}). \quad (10)$$

If the posterior is flat along any dimension in this region, there is a family of possible solutions to the minimization problem, whereas if the posterior is not flat along any dimensions, a finite number of local minima must exist. Furthermore, if we assume that the prior probability $P(\mathbf{e}, \mathbf{s})$ varies slowly over the region being considered, the structure of the posterior can be seen by examining the local structure of the log likelihood function,

$$L(y|\mathbf{e}, \mathbf{s}) = \sum_{x,k} \left(y_k^x - \sum_{i,j} \mathbf{e}_i G_{ij,k} \mathbf{s}_j^x \right)^2. \quad (11)$$

3.1 Local Curvature Analysis

We explore how the structure of the posterior probability reveals whether or not the color constancy problem is solvable.

The local structure we are interested in is the curvature of the log likelihood function, which tells us how moving in the parameter space affects our ability to explain the observed data. We will analyze how this structure changes when we introduce low-level spatial information into the model, and use this analysis to motivate the use of such information for color constancy.

The curvature of the log likelihood can be obtained from an eigen decomposition of its Hessian matrix. The Hessian is a matrix of all second order derivatives and thus, for an N -pixel image, is a $(d_E + Nd_S) \times (d_E + Nd_S)$ matrix. The eigenvectors of the Hessian matrix give principal directions of curvature and the corresponding eigenvalues give the extent of curvature in each of these directions.

At a maximum of the log likelihood, all eigenvalues of the Hessian must be non-negative (meaning that error cannot decrease as we move away from the local maximum). If there are any eigenvalues equal to zero, this implies that there are directions of zero curvature - e.g., there are ways we can vary the estimated parameters without affecting the likelihood of the data. For a unique maximum, all eigenvalues of Hessian must be positive at the maximum. Furthermore, the larger these eigenvalues are, the more robust is this solution to observation noise or unmodelled image variations.

To help us understand the structure of the color constancy problem for real-world images, we first examine representative special cases in toy examples. Appendix B gives a full derivation of the Hessian for the log likelihood function given by (11).

3.1.1 Underdetermined Case


We constructed a toy problem in which the sensor responses for all points in an image are generated using equation (8) with the noise term set to zero. To perform simulations, we used a hyperspectral data set of 28 natural scenes collected at Bristol University by Brelstaff, et. al. as described in []. Each of these images is 256×256 pixels in size and contains sensor responses to 31 spectral bands, ranging from 400 nm to 700 nm in 10 nm intervals. Each scene also contains a Kodak greycard at a known location with a constant reflectance spectrum of known intensity. The scene illuminant is approximated as the spectrum recorded at the location of the greycard divided by its constant reflectance. Our dataset therefore consists of $256 \times 256 \times 28$ reflectance spectra and 28 illumination spectra.

The first three illumination and reflectance basis functions obtained by applying principal components analysis (PCA) to this data are plotted in Figure 1a and b, respectively. We assume, without loss of generality, a Gaussian model for sensor responses centered at 650 (red), 550 (green), and 450 nm (blue) as shown in Figure 1c.


Generating 10 points using random illumination and reflectance weights and evaluating the Hessian at the correct location in parameter space results in the eigenvalue spectrum shown in Figure 2. We can see that, as expected, there are 3 directions of zero curvature. The three associated eigenvectors give the directions of zero curvature in parameter space. Any linear combination of these eigenvectors is also a direction of zero curvature.

It is interesting to note that there are 3 illumination weights and also three eigenvectors with corresponding zero eigenvalues. Each of these eigenvectors is of dimension $d_E + Nd_S$, with three of those dimensions corresponding to the illuminant weights. The three eigenvectors span the space of the 3-dimensional illuminant space, meaning that for any set of illumination coefficients, there exists some set of reflectance weights such that the likelihood associated with that set of parameters is equal to the likelihood associated with the correct set of parameters. In other words, this means that we can move in any direction in the 3-dimensional space of the illuminant and still find surface reflectance values that render to the observed sensor val-

PSfrag replacements




(a) Illumination Basis Functions



(b) Reflectance Basis Functions

PSfrag replacements



(c) Sensor Responses

Figure 1: Basis functions and sensor responses used in toy examples. Basis functions are obtained from a hyperspectral dataset.

ues. One example of this is a multiplicative scaling up of the illumination and a corresponding scaling down of the surface reflectances.




Figure 2: Eigenvalue spectrum and corresponding curvature of the log likelihood for a toy example of the color constancy problem. The directions of zero curvature and eigenvalues at zero show that the problem is underdetermined

3.1.2 Sufficient Sensors

If there are 4 independent photoreceptor responses available instead of 3, we would expect for a better behaved log-likelihood function. However, there should still be one direction of zero curvature in the Hessian, corresponding to a simple scaling of the illumination and surfaces. The basis functions and sensor responses given in Figure 1 are used for this example as well as a fourth sensor response which is chosen to be random (and thus independent of the other three). We again generate 10 random points and find that the Hessian behaves as expected, as shown in Figure 3. Not only does the Hessian have only one eigenvalue equal to zero, the other eigenvalues have a much larger value than the eigenvalues of the Hessian in the underdetermined case.




Figure 3: Eigenvalue spectrum and corresponding curvature of the log likelihood for a case when the number of sensors exceeds the dimensionality of the surface model. There is now only one direction of zero curvature, corresponding to an unknown scale factor.

3.2 Bayesian Color Constancy in the Underdetermined Case

Although there may be multiple degrees of ambiguity in the likelihood function, in general, the Bayesian approach is still able to produce a unique parameter estimate, because the prior or the loss function will favor one solution over the others [2].

In the ideal case, however, we would like for our parameter estimate to be dominated by the effect of the likelihood function. That is, we would like for the estimate to be based on the data itself, not on prior assumptions about the data or how much we value particular parts of the data.

4 Spatio-spectral Basis Functions

We now introduce an alternative formulation of the finite-dimensional models presented in Section 2 that model both spatial and spectral properties of groups of pixels in natural images. The hypothesis is that by taking characteristic spatial variations of natural images into account, the number of parameters necessary to describe the reflectance of an image will be reduced so that the problem is no longer underdetermined.

Certain physical phenomena, such as interreflections, may generate characteristic spatial and spectral signatures which allow estimation of the desired spectra from image data. Natural images, like foliage, may exhibit characteristic spatial changes in color. The spatio-spectral basis functions allow us to exploit these, or other, regularities in the visual world in order to solve the color constancy problem. We expect this approach will work best in richly textured images, and worst in flat, color Mondrian-type images [18].

4.1 Modified Linear Models

Instead of using linear models to describe the reflectance of individual pixels, we will now use these models to describe groups of pixels. Without loss of generality, we can group pixels into $m \times m$ blocks and use the same basic formulation as was developed in Section 2. In order to do so, it is necessary to convert blocks of pixels into vector format. We do this by rasterizing the pixels within the block. The reflectance of a block of pixels is defined as a vector of length m^2L consisting of the reflectance of each pixel within the block stacked on top of each other in raster order. The same process is used to describe the sensor response of the block as a vector of length m^2K and the illumination as a vector of length m^2L .

The basis functions used to model the reflectance of blocks of pixels are now referred to as *spatio-spectral* reflectance basis functions, since they describe both the spectral and spatial characteristics of a block of pixels.

We shall denote a group of pixels by $\{x\}$, so that the illumination and reflectance of a block

of pixels is given by:

$$E^{\{\}}(\gamma) = \sum_{i=1}^{m^2 L} \hat{E}_i(\gamma) \mathbf{e}_i^{\{\}}, \quad (12)$$

$$S^{\{x\}}(\gamma) = \sum_{j=1}^{m^2 L} \hat{S}_j(\gamma) \mathbf{s}_j^{\{x\}}, \quad (13)$$

where $\hat{E}_i(\gamma)$ and $\hat{S}_j(\gamma)$ are the spatio-spectral illumination and reflectance basis functions, $\mathbf{e}_i^{\{\}}$ and $\mathbf{s}_j^{\{x\}}$ are the weights associated with these basis functions, $E^{\{\}}(\gamma)$ is the illumination of all blocks in the image, and $S^{\{x\}}(\gamma)$ is the reflectance of the block of pixels $\{x\}$. Note that the elements of the scene are now written as a function of γ rather than λ . This is due to the fact that the spatio-spectral representation contains information about both the frequency and spatial characteristics of the scene. Approximating these models with fewer dimensions, we can define

$$\tilde{E}^{\{\}}(\gamma) = \sum_{i=1}^{\hat{d}_E} \hat{E}_i(\gamma) \mathbf{e}_i^{\{\}}, \quad (14)$$

$$\tilde{S}^{\{x\}}(\gamma) = \sum_{j=1}^{\hat{d}_S} \hat{S}_j(\gamma) \mathbf{s}_j^{\{x\}}, \quad (15)$$

where $\tilde{E}^{\{\}}(\gamma)$ is the approximate illumination of all blocks in an image, constructed using a $\hat{d}_E$ dimensional model, and $\tilde{S}^{\{x\}}(\gamma)$ is the approximated reflectance for the block $\{x\}$, constructed using a $\hat{d}_S$ dimensional model.

We define an $m \times m$ block of pixels as having $m^2 K$ sensor outputs, with K sensor outputs per pixel. Thus, we define the sensor responses of the group of pixels as the block diagonal matrix

$$\hat{R}(\gamma) = \begin{bmatrix} R(\lambda) & 0 & \dots & 0 \\ 0 & R(\lambda) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & R(\lambda) \end{bmatrix} \quad (16)$$

with m^2 blocks along the diagonal, and

$$R = [R_1(\gamma) \quad R_2(\gamma) \quad \dots \quad R_k(\gamma)].$$

We let $\hat{R}_k(\gamma)$ refer to the k th column of the matrix $\hat{R}(\gamma)$.

Following a derivation analogous to that presented in Section 2, we can write the sensor

output in a bilinear form with noise as:

$$y_k^{\{x\}} = \sum_{i=1}^{\widehat{d}_E} \sum_{j=1}^{\widehat{d}_S} e_i^{\{x\}} \widehat{G}_{ij,k} s_j^{\{x\}} + w_k^{\{x\}}, \quad (17)$$

where $y_k^{\{x\}}$ is the sensor response of the block of pixels $\{x\}$, $\widehat{G}_{ij,k}$ is defined as

$$\widehat{G}_{ij,k} = \sum_{\gamma} \widehat{R}_k(\gamma) \widehat{E}_i(\gamma) \widehat{S}_j(\gamma), \quad (18)$$

and $w_k^{\{x\}}$ is the noise introduced from using a lower-dimensional linear model.

4.2 The Advantage of Using Spatial Information

The advantage of incorporating spatial information into the linear models used to describe scenes is that it provides more descriptive power using a fewer number of parameters. It is straightforward to show that in even in the worst case, i.e. when incorporating spatial information provides no benefit and individual pixels are independent, the spatio-spectral method cannot do worse than the traditional method. The amount of improvement in descriptive power provided by the new method can be tested by finding a linear reflectance basis using real images for blocks of various sizes (including $m = 1$, which corresponds to the traditional case of examining each pixel independently) and calculating how much of the variance of a set of test images is described in each case. To do this, we use the dataset of hyperspectral images described earlier in Section 3.1.1

The amount of variance described by linear models of varying dimensionality when considering block sizes of $m = 1, 2, 3$, and 4 is shown in Figure 4. The figure plots the number of basis functions per pixel (the number of basis functions used to describe the $m \times m$ block divided by m^2) against the amount of variance described when using that many basis functions. The linear model was obtained using approximately half of the data in the 28 scenes (selected randomly) and tested on the entire data set.

It can be seen that using a block size of $m = 2$ dramatically improves performance. For example, the data shows that the error rate obtained when describing each pixel with 3 basis functions can be achieved by using only 1.5 basis functions per pixel when describing an image with 2×2 blocks (i.e., 6 basis functions for the block). Increasing the block size beyond $m = 2$ does not show much of an improvement in reconstruction error. This is because the majority of the advantage of using blocks is to take advantage of local flatness in an image, which is already done when 2×2 blocks are used. Figure 5 plots the first 12 spatio-spectral reflectance basis functions for $m = 2$. These basis functions were extracted from the hyperspectral data using the calibration gray card in each image of the dataset, in the same manner as described in Section 3.1.1. It can be seen that the first three basis functions are essentially the same as




Figure 4: The squared reconstruction error using linear models of varying dimensions for block sizes of $m = 1, 2, 3$, and 4. Using block sizes of $m = 2$ or larger greatly reduces the number of coefficients needed to describe each pixel.

the standard single pixel bases repeated at all 4 pixels.

To see the advantage of using spatial information in the context of Bayesian color constancy, we must once again analyze the local curvature of the posterior probability. Following the derivations in Section 3, we once again make the assumption of a locally flat prior distribution, allowing us to simplify the problem and examine the log likelihood function, now written as

$$L(y^{\{x\}} | e^{\{x\}}, s^{\{x\}}) = \sum_{\{x\}, k} \left(y_k^{\{x\}} - \sum_{i,j} e_i^{\{x\}} \hat{G}_{ij,k} s_j^{\{x\}} \right)^2. \quad (19)$$

5 Experiments with Natural Images

We look at groups of 40 pixels, grouped into 10 sets of 2x2 pixel blocks. The 10 sets were selected from randomized positions in the image.

In order to get the priors, we project our dataset of reflectances and illuminations onto the basis functions given in Figures 1a (for illumination) and 5 (for spatio-spectral reflectance). We fit truncated Gaussians to the data and use these as our priors.

Our goal is now to find the set of illumination and reflectance coefficients that minimize the posterior probability. In order to do this, we need to search over the space of all possible illumination and surface reflectance coefficients. Since the number of surface coefficients scales with the number of pixels in the image, this quickly becomes an intractable problem.

Following [2], in order to make the problem more feasible, we limit our search to look over the space of illumination coefficients and solve for the surfaces at each iteration using a deterministic relationship. In the underdetermined case, where the number of sensor responses at each pixel is equal to the number of surface coefficients at each pixel, we solve for sur-




Figure 5: The first 12 spatio-spectral basis functions for 2x2 pixel blocks. Each basis function is plotted for wavelength values from 400 to 700 nm on the x-axis. The first 3 spatio-spectral basis functions, (a), show only spectral variations and no spatial variation. The next 3 basis functions, (b), correspond to spatial derivatives. The final 6 indicate a more complex relationship of spatial and spectral variations, (c) and (d).

faces that produce the sensor responses exactly given an illumination (this is done by solving equation (8) assuming no sensor noise through a simple matrix inversion). This is equivalent to constraining our search to locations of maximum probability in the likelihood. In the case where we have more sensor responses than surface coefficients, we take the pseudo-inverse to find the surface coefficients corresponding to a given illumination.

Unfortunately, applying this technique of constraining the search space to real images gives numerical problems. In these images, the noise term in equation (8), which corresponds to the surface and illumination components not included in a low-dimensional model, causes the surface coefficients that maximize the likelihood term given the true illuminant to be very different from the true surfaces coefficients. This makes it hard to properly analyze the posterior probability corresponding to a given illuminant.

To avoid the numerical issues mentioned above, we test our algorithm on images constructed by first projecting the true surface reflectances down to a number of dimensions equal to the number of sensor responses (3 for 1 pixel at a time, 12 for 2x2 blocks), multiplying the reflectance image by the true illuminant, and passing the result through the Gaussian sensors shown in Figure 1. Figure ? shows that there is very little difference between the image obtained using the true reflectance and the image obtained after projecting the reflectance onto the lower-dimensional subspace. We need to be clear: these are pre-processed images, not “real” images, although they are tantalizingly close to real images. We show the benefits of using spatio-spectral basis functions in solving color constancy using the pre-processed images, and expect that extensions to real images will follow in work by us or by others.




Figure 6: The effect of projecting reflectance to lower dimensions is very small.

Figure 7 shows results for our approach using spatio-spectral basis functions in a Bayesian approach with the standard Bayesian approach and the Gray World approach. The particular illuminant we wish to estimate is not fit well by the prior, as seen in the figure. Each algorithm is run on 15 random groups of 40 pixels chosen from the image. Despite the low prior probability of the illuminant, the Bayesian approach using spatio-spectral basis functions allows

us to locate the correct illuminant for each draw. This illustrates that the likelihood term does dominates our estimate, as we would expect. The standard Bayesian approach, examining one pixel at a time, and the Gray World approach both result in incorrect estimates.




Figure 7: Results of color constancy applied to a natural image. Each algorithm was run 15 times, with results shown in dotted lines. The mean illumination from the prior is given in the dashed line as reference, and the true illuminant is marked by circles.

In our experiments, we also found that the positivity constraint restricted the set of possible illuminants more when using 2×2 blocks than when considering 1 pixel at a time, so that the results for 2×2 blocks using 12 coefficients (underdetermined) were better than the results for the standard underdetermined case.

6 Conclusions

We have studied the effect of using spatio-spectral surface reflectance basis functions to find a low-dimensional estimate of surface reflectance characteristics.

These basis functions can convert the color constancy problem from an under-determined problem to an over-determined one. We can see this in the curvature characteristics of the

log-likelihood function at parameter local maxima.

For natural images, pre-processed to restrict the surface reflectance variations to be explained by an average of 3 basis functions per pixel, we find that using spatio-spectral basis functions allows for very accurate estimates of the global illumination spectrum (and therefore of surface reflectance spectra, too). For the same images, using purely spectral basis functions results in an underdetermined problem, with much higher estimation error for the optimal Bayesian estimate.

Acknowledgments

We thank David Brainard for many helpful conversations.

APPENDIX: The Hessian of the Log Likelihood Functions with Bilinear Forms

$$H_{\mathbf{e}_l \mathbf{e}_m} = \frac{\partial^2 L}{\partial \mathbf{e}_l \partial \mathbf{e}_m}$$

$$\frac{\partial L}{\partial \mathbf{e}_l} = \frac{\partial}{\partial \mathbf{e}_l} \sum_{x,k} (y_k^x - \sum_{i,j} \mathbf{e}_i G_{ij,k} \mathbf{s}_j^x)^2 \quad (\text{A-1})$$

$$= \sum_{x,k} \frac{\partial}{\partial \mathbf{e}_l} (y_k^x - \sum_{i,j} \mathbf{e}_i G_{ij,k} \mathbf{s}_j^x)^2 \quad (\text{A-2})$$

$$= -2 \sum_{x,k} \left((y_k^x - \sum_{i,j} \mathbf{e}_i G_{ij,k} \mathbf{s}_j^x) \left(\frac{\partial}{\partial \mathbf{e}_l} \sum_{i,j} \mathbf{e}_i G_{ij,k} \mathbf{s}_j^x \right) \right) \quad (\text{A-3})$$

$$= -2 \sum_{x,k} \left((y_k^x - \sum_{i,j} \mathbf{e}_i G_{ij,k} \mathbf{s}_j^x) \sum_j G_{lj,k} \mathbf{s}_j^x \right) \quad (\text{A-4})$$

$$\frac{\partial^2 L}{\partial \mathbf{e}_l \partial \mathbf{e}_m} = -2 \frac{\partial}{\partial \mathbf{e}_m} \sum_{x,k} \left((y_k^x - \sum_{i,j} \mathbf{e}_i G_{ij,k} \mathbf{s}_j^x) \sum_j G_{lj,k} \mathbf{s}_j^x \right) \quad (\text{A-5})$$

$$= -2 \sum_{x,k} \frac{\partial}{\partial \mathbf{e}_m} \left((y_k^x - \sum_{i,j} \mathbf{e}_i G_{ij,k} \mathbf{s}_j^x) \sum_j G_{lj,k} \mathbf{s}_j^x \right) \quad (\text{A-6})$$

$$= 2 \sum_{x,k} \left(\left(\sum_j G_{lj,k} \mathbf{s}_j^x \right) \left(\sum_j G_{mj,k} \mathbf{s}_j^x \right) \right) \quad (\text{A-7})$$

$$H_{s_l^n s_m^p} = \frac{\partial^2 L}{\partial s_l^n \partial s_m^p}$$

$$\frac{\partial L}{\partial s_l^n} = \frac{\partial}{\partial s_l^n} \sum_{x,k} (y_k^x - \sum_{i,j} e_i G_{ij,k} s_j^x)^2 \quad (\text{A-8})$$

$$= \frac{\partial}{\partial s_l^n} \left(\sum_{x \neq n} \sum_k (y_k^x - \sum_{i,j} e_i G_{ij,k} s_j^x)^2 + \sum_{x=n} \sum_k (y_k^x - \sum_{i,j} e_i G_{ij,k} s_j^x)^2 \right) \quad (\text{A-9})$$

$$= \frac{\partial}{\partial s_l^n} \sum_k (y_k^n - \sum_{i,j} e_i G_{ij,k} s_j^n)^2 \quad (\text{A-10})$$

$$= -2 \sum_k \left((y_k^n - \sum_{i,j} e_i G_{ij,k} s_j^n) \left(\frac{\partial}{\partial s_l^n} \sum_{i,j} e_i G_{ij,k} s_j^n \right) \right) \quad (\text{A-11})$$

$$= -2 \sum_k \left((y_k^n - \sum_{i,j} e_i G_{ij,k} s_j^n) \left(\sum_i e_i G_{il,k} \right) \right) \quad (\text{A-12})$$

$$\frac{\partial^2 L}{\partial s_l^n \partial s_m^p} = 0 \quad (\text{A-13})$$

$$\frac{\partial^2 L}{\partial s_l^n \partial s_m^n} = \frac{\partial}{\partial s_m^n} \left(-2 \sum_k \left((y_k^n - \sum_{i,j} e_i G_{ij,k} s_j^n) \left(\sum_i e_i G_{il,k} \right) \right) \right) \quad (\text{A-14})$$

$$= 2 \sum_k \left(\left(\frac{\partial}{\partial s_m^n} \left(\sum_{i,j} e_i G_{ij,k} s_j^n \right) \right) \left(\sum_i e_i G_{il,k} \right) \right) \quad (\text{A-15})$$

$$= 2 \sum_k \left(\left(\sum_i e_i G_{im,k} \right) \left(\sum_i e_i G_{il,k} \right) \right) \quad (\text{A-16})$$

$$H_{e_l s_m^p} = \frac{\partial^2 L}{\partial e_l \partial s_m^p}$$

$$\frac{\partial^2 L}{\partial e_l \partial s_m^p} = \frac{\partial}{\partial e_l} \left(-2 \sum_k \left((y_k^p - \sum_{i,j} e_i G_{ij,k} s_j^p) \left(\sum_i e_i G_{im,k} \right) \right) \right) \quad (\text{A-17})$$

$$= -2 \sum_k \left(\left(\frac{\partial}{\partial e_l} (y_k^p - \sum_{i,j} e_i G_{ij,k} s_j^p) \right) \left(\sum_i e_i G_{im,k} \right) + \right. \\ \left. (y_k^p - \sum_{i,j} e_i G_{ij,k} s_j^p) \left(\frac{\partial}{\partial e_l} \sum_i e_i G_{im,k} \right) \right) \quad (\text{A-18})$$

$$= 2 \sum_k \left(\left(\sum_j G_{lj,k} s_j^p \right) \left(\sum_i e_i G_{im,k} \right) - (y_k^p - \sum_{i,j} e_i G_{ij,k} s_j^p) G_{lm,k} \right) \quad (\text{A-19})$$

References

- [1] J. O. Berger. *Statistical decision theory and Bayesian analysis*. Springer, 1985.
- [2] D. H. Brainard and W. T. Freeman. Bayesian color constancy. *Journal of the Optical Society of America A*, 14(7):1393–1411, 1997.
- [3] D H Brainard and B A Wandell. Analysis of the retinex theory of color vision. *Journal of the Optical Society of America A*, 3(10):1651–1661, 1986.

- [4] G. Buchsbaum. A spatial processor model for object colour perception. *J. Franklin Inst.*, 310:1–26, 1980.
- [5] CIE. Colorimetry. Bureau Central de la CIE, 1986.
- [6] M. D’Zmura, G. Iverson, and B. Singer. Probabilistic color constancy. In R. D. Luce et. al., editor, *Geometric representations of perceptual phenomena: papers in honor of Tarow Indow’s 70th birthday*. Lawrence Erlbaum, Hillsdale, NJ. in press.
- [7] R. M. Evans. *The perception of color*. Wiley, New York, 1974.
- [8] G. Finlayson and S. Hordley. Improving gamut mapping color constancy. *Image and Vision Computing*, 9:1774–1783, 2000.
- [9] D. Forsyth. A novel approach to color constancy. *International Journal of Computer Vision*, 5:5–36, 1990.
- [10] D. A. Forsyth. A novel approach to colour constancy. In *Proc. 1st Intl. Conf. Computer Vision*, pages 9–18. IEEE, 1988.
- [11] R. Gershon and A. D. Jepson. The computation of color constant descriptors in chromatic images. *Color Res. Appl.*, 14:325–334, 1989.
- [12] T. Jaaskelainen, J. Parkkinen, and S. Toyooka. A vector-subspace model for color representation. *Journal of the Optical Society of America A*, 7:725–730, 1990.
- [13] D. B. Judd, D. L. MacAdam, and G. W. Wyszecki. Spectral distribution of typical daylight as a function of correlated color temperature. *J. Opt. Soc. Am.*, page 1031, 1964.
- [14] E H Land. The retinex theory of color vision. *Scientific American*, 237(6):108–128, 1977.
- [15] E. H. Land. Recent advances in retinex theory and some implications for cortical computations: color vision and the natural image. *Proc. Nat. Acad. Sci. USA*, 80:5163–5169, 1983.
- [16] L. T. Maloney. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *Journal of the Optical Society of America A*, 3(10):1673–1683, 1986.
- [17] L T Maloney and B A Wandell. Color constancy: a method for recovering surface spectral reflectance. *Journal of the Optical Society of America A*, 3(10):29–33, 1986.
- [18] J. J. McCann, S. P. McKee, and T. H. Taylor. Quantitative studies in retinex theory: a comparison between theoretical predictions and observer responses to the ‘color mon-drian’ experiments. *Vision Research*, 16:445–458, 1976.
- [19] H. J. Trussell and M. J. Vrhel. Estimation of illumination for color correction. In *Proc. ICASSP*, pages 2513–2516, 1991.