

Low-Dimensional Representations of Shaded Surfaces under Varying Illumination

Peter Nillius

Jan-Olof Eklundh

Computational Vision & Active Perception Laboratory (CVAP)
Department of Numerical Analysis and Computing Science
Royal Institute of Technology (KTH), S-100 44 Stockholm, Sweden

Abstract

The aim of this paper is to find the best representation for the appearance of surfaces with Lambertian reflectance under varying illumination. Previous work using principal component analysis (PCA) found the best sub-space to represent all images of an object under a varying point light source. We extend this to images from any illumination distribution. Specifically we calculate the bases for all configurations of a point plus ambient light source and two point light sources, as well as from a database of captured real world illumination. We also reformulate the optimization criterion used in PCA. The resulting basis, we believe has higher representability and is better for analysing images of shaded objects. The different bases are compared on a database of images to test the representability.

1. Introduction

Computer vision algorithms applicable to real world scenes must be able to cope with variations in appearance caused by changes in the illumination. This is not an easy task. There is an infinite number of possible illuminations which results in an infinite number of possible appearances of an object. In fact, Belhumeur and Kriegman [2] show that the space of images of a convex Lambertian object under all possible lighting conditions is infinite dimensional.

A step towards dealing with this problem is given by Basri and Jacobs [1] and Ramamoorthi and Hanrahan [10]. They show that a Lambertian surface acts as a low-pass filter on the incoming illumination in practice reducing the dimensionality of the reflected light. More precisely, they suggest that a 9D linear sub-space will suffice to represent the reflected light from a Lambertian surface under all possible illuminations.

Additionally, when viewing an object in a single image, only one hemisphere of the surface normals is visible, thereby reducing the dimensionality of the image space further. In [9], Ramamoorthi derives analytically the principal

components of the space of all images of a Lambertian object illuminated with a point light source from all possible directions. He shows that the image space in practice can be represented by approximately five principal components. These results correspond well, both quantitatively and qualitatively with earlier empirical work using PCA on images of objects under varying illumination by Hallinan [6] and Epstein et. al. [5].

In this paper we generalize the work of [9] to include any illumination distribution. We also modify the criterion of PCA slightly. The resulting basis, we believe, is better in terms of representability. Using the developed framework we derive the principal component bases for various illumination distributions including captured real world illumination. The representability of the bases are compared on a set of images. The potential applications are among other things illumination estimation and illumination invariant feature detectors. We demonstrate this by briefly going through the conditions of illumination estimation using the principal component basis, and by deriving the basis for the appearance of apparent contours.

2. Shading in Spherical Harmonics

We use the spherical harmonic representation of image irradiance developed by Basri and Jacobs [1] and Ramamoorthi and Hanrahan [10]. Spherical harmonics are orthonormal basis functions forming the analogy to the Fourier series on the sphere. There are two main advantages using this framework. The first is that we can model any light source. The reflectance distribution function of a Lambertian surface acts as a low-pass filter on the incident illumination making the reflected light, in essence, lie in a lower dimensional space. In most cases, nine coefficients are enough to represent the reflected light.

The second advantage is that attached shadows are taken into account while still having a linear expression. Attached shadows occur when the surface normal points away from the light source.

Let the coordinate frame be located with its x and y -axes in the image plane and the z -axis pointing towards the camera. With the surface normal in spherical coordinates, $\Theta = (\theta, \phi)$ the image irradiance is

$$i(\Theta) = \sum_{l=0}^{\infty} \sqrt{\frac{4\pi}{2l+1}} R_l \sum_{m=-l}^l L_{lm} Y_{lm}(\Theta) \quad (1)$$

where Y_{lm} are the spherical harmonic functions, L_{lm} are the spherical harmonic coefficients of the illumination. R_l are the coefficients for Lambertian reflectance, including attached shadows.

We will also index the spherical harmonics using a single variable. This is commonly done by $s = l^2 + l + m$. The reverse transform to double indices can be done by $l(s) = \text{ceil}(\sqrt{s+1}) - 1$ and $m(s) = s - l(s)^2 - l(s)$. The shading using single index notation is

$$i(\Theta) = \sum_{s=0}^{\infty} \sqrt{\frac{4\pi}{2l(s)+1}} R_{l(s)} L_s Y_s(\Theta) \quad (2)$$

3. The Image Space

We begin by defining the image space. An image of an object depends on the object's shape, the illumination and the surface reflectance function. In this article we only consider Lambertian surfaces.

Let the illumination distribution be described as a distribution of the spherical harmonic coefficients, $p_L(\mathbf{L})$, where \mathbf{L} is the vector containing the coefficients L_s .

Since we are assuming convex objects there are no cast shadows, which means that the image irradiance is uniquely determined by the surface normal. Given the surface normal of the object at each point in the image we can generate all possible images of that object under the illumination distribution we have defined.

However, the criteria we use for deriving PCA depend only on the scalar product of images. Therefore it is not necessary to generate the images as long as the scalar product between images can be calculated. For this purpose only the distribution of surface normals is needed. Given the object's surface normal distribution, $p_{\hat{n}}(\Theta)$, $\Theta \in S^2$, we can calculate the scalar product of images I and J as

$$I \bullet J = \int_{S^2} i(\Theta) j(\Theta) p_{\hat{n}}(\Theta) d\omega_{\Theta}, \quad (3)$$

where $i(\Theta)$ and $j(\Theta)$ are functions returning the image irradiance for a particular surface normal for image I and J .

So, for our purposes the illumination distribution and the surface normal distribution is sufficient to define the image space.

4. Principal Component Analysis

In this section we derive the principal components of the image space. We recreate the results of Ramamoorthi [9] but do the derivation differently. This leads to a generalization where any illumination distribution can be considered as opposed to a single point light source used in [9]. The final expression involves a matrix defining the geometry of the object and a matrix containing the covariance structure of the illumination. These two matrices uniquely determine the principal component basis. Using this expression one can easily experiment with various object geometries and illumination configurations.

4.1. Criteria to Derive the Principal Components

Principal component analysis amounts to finding a coordinate transformation such that the covariance matrix of the random variable (in our case the image) is diagonalized. A geometric interpretation of PCA is that the first eigenvector, here called U_0 , is the direction which has the highest variability. U_0 is the vector that maximizes the variance of the scalar product between U_0 and the images, or mathematically:

$$U_0 = \underset{\|U\|=1}{\operatorname{argmax}} \operatorname{Var}\{I \bullet U\}, \quad (4)$$

where the image I is the random variable. Subsequent eigenvectors are constructed iteratively by maximizing the same variance with the added condition that they are orthogonal to all the previous eigenvectors, [8]. Using this criterion we derive the principal components.

We also use a slightly different criterion. The objective is to find a basis that decomposes the image into a linear combination of coefficients and basis functions, $I = \sum_i c_i U_i$. Typically images are analysed by estimating the coefficients c_i . There are a number of things we can do to make this estimation as robust as possible. First, the basis should be orthogonal. This ensures that the basis is as efficient as possible when truncated. Moreover, it allows us to estimate the coefficients individually using the scalar product, $c_i = I \bullet U_i$, as opposed to estimating them all simultaneously using e.g. least squares.

Secondly, the basis should contain the constant function. The illumination frequently contains an ambient component which can vary. By including the constant function in the basis we make the remaining functions independent of the ambient component.

Thirdly, the components $c_i U_i$ should on average (over the image distribution) have as high variance *in the image* as possible. Components with a high variance have a high signal-to-noise ratio and are therefore more robustly estimated.

To include the constant function in the basis we subtract the image's mean from each image. This forces the basis functions to be orthogonal to the constant function. The constant function is then added to the basis at a later stage. The image mean is calculated by $I \bullet \mathbf{1}^1$, where $\mathbf{1}$ is an image with ones. The variance of each component is simply c_i^2 . We arrive at the following criterion.

$$U_0 = \operatorname{argmax}_{\|U\|=1} E\{((I - (I \bullet \mathbf{1})\mathbf{1}) \bullet U)^2\} \quad (5)$$

In the following sections we display three variants of bases. The first two are created using criterion (4) and are analogous to the bases in [9]. They are referred to as *Non-Centered PCA* and *Standard PCA*. In non-centered PCA the mean is not subtracted when calculating the covariance matrix. The third variant is created using our criterion (5) and is called *Modified PCA*.

4.2. Deriving the Basis

In order to do the analysis we rewrite the eigenfunction in its spherical harmonic representation.

$$U(\Theta) = \sum_s u_s Y_s(\Theta) \quad (6)$$

The scalar product of I and U becomes

$$\begin{aligned} I \bullet U &= \sum_{s,s'} R_{l(s)} \sqrt{\frac{4\pi}{2l(s)+1}} L_s u_{s'} \\ &\times \int_{S^2} p_{\hat{n}}(\Theta) Y_s(\Theta) Y_{s'}(\Theta) d\omega_{\Theta} \end{aligned} \quad (7)$$

We can rewrite this in matrix form.

$$I \bullet U = \mathbf{u}^T \mathbf{M} \mathbf{Q} \mathbf{L}, \quad (8)$$

where \mathbf{L} is the vector containing the spherical harmonic coefficients of the light source L_s , \mathbf{Q} is a diagonal matrix with elements $q_{ii} = R_{l(i)} \sqrt{\frac{4\pi}{2l(i)+1}}$, \mathbf{M} contains the elements

$$m_{ss'} = \int_{S^2} p_{\hat{n}}(\Theta) Y_s(\Theta) Y_{s'}(\Theta) d\omega_{\Theta} \quad (9)$$

and \mathbf{u} is a vector containing the spherical harmonic coefficients of U .

Now, the only random variable in (8) is \mathbf{L} . Let $\operatorname{Covar}\{\mathbf{L}\} = \Sigma_{\mathbf{L}}$. Then

$$\operatorname{Var}\{\mathbf{u}^T \mathbf{M} \mathbf{Q} \mathbf{L}\} = \mathbf{u}^T \mathbf{M} \mathbf{Q} \Sigma_{\mathbf{L}} \mathbf{Q} \mathbf{M} \mathbf{u}. \quad (10)$$

The transpose is dropped on \mathbf{M} and \mathbf{Q} since they are both symmetric.

¹Due to our definition of the scalar product we don't need to divide with the number of pixels.

The expression should be maximized under the condition that U is normalized. We obtain the following condition on U 's coefficients.

$$\begin{aligned} U \bullet U &= \sum_{s,s'} u_s u_{s'} \int_{S^2} p_{\hat{n}}(\Theta) Y_s(\Theta) Y_{s'}(\Theta) d\omega_{\Theta} \\ &= \sum_{s,s'} u_s u_{s'} m_{ss'} = \mathbf{u}^T \mathbf{M} \mathbf{u} = 1 \end{aligned} \quad (11)$$

The maximization problem can now be written in terms of U 's spherical harmonic coefficients.

$$\mathbf{u}_0 = \operatorname{argmax}_{\mathbf{u}^T \mathbf{M} \mathbf{u} = 1} \mathbf{u}^T \mathbf{M} \mathbf{Q} \Sigma_{\mathbf{L}} \mathbf{Q} \mathbf{M} \mathbf{u} \quad (12)$$

Applying the coordinate transform

$$\mathbf{v} = \mathbf{M}^{1/2} \mathbf{u}, \quad (13)$$

we obtain

$$\mathbf{v}_0 = \operatorname{argmax}_{\mathbf{v}^T \mathbf{v} = 1} \mathbf{v}^T \mathbf{M}^{1/2} \mathbf{Q} \Sigma_{\mathbf{L}} \mathbf{Q} \mathbf{M}^{1/2} \mathbf{v}. \quad (14)$$

This is a quadratic expression of \mathbf{v} which should be maximized under the condition that \mathbf{v} is normalized. It is well known that the solution is the eigenvector of $\mathbf{M}^{1/2} \mathbf{Q} \Sigma_{\mathbf{L}} \mathbf{Q} \mathbf{M}^{1/2}$ with the largest eigenvalue. The subsequent eigenvectors maximize the expression while being orthogonal to the previous ones. We have in fact performed the PCA. To be sure we check that orthogonality of the vectors \mathbf{v} corresponds to orthogonality in the image space.

$$U_i \bullet U_j = \mathbf{u}_i^T \mathbf{M} \mathbf{u}_j = \mathbf{v}_i^T \mathbf{v}_j \quad (15)$$

The spherical harmonic coefficients for the eigenimages U_i are computed by $\mathbf{u}_i = \mathbf{M}^{-1/2} \mathbf{v}_i$.

In relation to the derivation of [9], subtracting the mean or not is in our case a matter of subtracting the mean or not when calculating $\Sigma_{\mathbf{L}}$.

When constructing the basis according to the second criterion (5) we should subtract the *image* mean. The image mean is

$$\begin{aligned} I \bullet \mathbf{1} &= \sum_s L_s q_{ss} \int_{S^2} Y_s(\Theta) p_{\hat{n}}(\Theta) d\omega_{\Theta} \\ &= \sum_s L_s q_{ss} n_s, \end{aligned} \quad (16)$$

where n_s is the mean of Y_s over the surface normal distribution.

Now,

$$\begin{aligned} (I - (I \bullet \mathbf{1})\mathbf{1}) \bullet U &= \sum_{s,s'} L_s q_{ss} u_{s'} \\ &\times \int_{S^2} (Y_s(\Theta) - n_s) Y_{s'}(\Theta) p_{\hat{n}}(\Theta) d\omega_{\Theta} \\ &= \sum_{s,s'} L_s q_{ss} u_{s'} (m_{ss'} - n_s n_{s'}) \\ &= \mathbf{u}^T (\mathbf{M} - \mathbf{n} \mathbf{n}^T) \mathbf{Q} \mathbf{L} \end{aligned} \quad (17)$$

$$= \mathbf{u}^T (\mathbf{M} - \mathbf{n} \mathbf{n}^T) \mathbf{Q} \mathbf{L} \quad (18)$$

Furthermore

$$E\{((I - (I \bullet \mathbf{1}) \mathbf{1}) \bullet U)^2\} = \mathbf{u}^T (\mathbf{M} - \mathbf{nn}^T) \mathbf{Q} \mathbf{V}_L \mathbf{Q} (\mathbf{M} - \mathbf{nn}^T) \mathbf{u} \quad (19)$$

where $\mathbf{V}_L = E\{\mathbf{L}\mathbf{L}^T\}$.

Doing the same coordinate change as before we get the \mathbf{v} vectors as being the eigenvectors of the matrix

$$\mathbf{M}^{-1/2} (\mathbf{M} - \mathbf{nn}^T) \mathbf{Q} \mathbf{V}_L \mathbf{Q} (\mathbf{M} - \mathbf{nn}^T) \mathbf{M}^{-1/2} \quad (20)$$

and finally the spherical harmonics coefficients of the eigenimages by $\mathbf{u}_i = \mathbf{M}^{-1/2} \mathbf{v}_i$ as before.

5. Illumination Distributions

Let us next derive the covariance matrices for a number of interesting illumination distributions. The derivations will be exemplified with bases calculated using a surface normal distribution of an image of a sphere. When viewing an image of a sphere, the visible area of a surface is proportional to the cosine of the polar angle of the surface normal. Taking this into account, as well as the geometry of the visible hemisphere, the elements of \mathbf{M} are

$$m_{ss'} = \frac{1}{\pi} \int_0^{\frac{\pi}{2}} \int_0^{2\pi} Y_s(\theta, \phi) Y_{s'}(\theta, \phi) \cos \theta \sin \theta d\theta d\phi \quad (21)$$

In the sections below we only derive the elements of \mathbf{V}_L and $\mu_L = E\{\mathbf{L}\}$, since Σ_L and \mathbf{V}_L are related as

$$\Sigma_L = \mathbf{V}_L - \mu_L \mu_L^T. \quad (22)$$

5.1. Rotated Light Sources

First we consider the case of a light field, defined by its spherical harmonic coefficients b_{lm} , rotated to illuminate from all possible directions. The coefficients after rotation can be written as a linear combination of the other coefficients in the same order l

$$L_{lm} = \sum_{n=-l}^l D_{lnm}(\alpha, \beta, \gamma) b_{ln}, \quad (23)$$

where D_{lnm} are the rotation reparameterization functions sometimes referred to as the Wigner D-functions, [3].

The elements of \mathbf{V}_L are the product of two rotated coefficients integrated over the rotation group, $SO(3)$. Because the orthogonality relation

$$\int_{SO(3)} D_{lnm}(R) D_{l'n'm'}(R) dR = \frac{\delta_{l,l'} \delta_{m,m'} \delta_{n,n'}}{2l+1} \quad (24)$$

holds, the off-diagonal elements of \mathbf{V}_L will be zero and the elements on the diagonal are

$$E\{L_{lm}^2\} = \sum_{n=-l}^l \frac{b_{ln}^2}{2l+1}. \quad (25)$$

As $D_{000}(\alpha, \beta, \gamma) = 1$ and due to the orthogonality relation, the means are zero for all L_{lm} except L_{00} and

$$E\{L_{00}\} = b_{00}. \quad (26)$$

These formulas are used in the calculations in the next sections.

5.2. Point Light Source

Here we derive the bases for a single point light source. Since equation (25) already takes into account all rotations of the light source we can select any start position of our light source e.g. $\Theta = (0, 0)$. To find the spherical harmonic coefficients for a light source we integrate the product of the light source and each spherical harmonic function. The point light source is described by a delta function so the coefficients will be a sample of the spherical harmonic function at the position of the light source, i.e. with the point light at $(0, 0)$, $b_{lm} = Y_{lm}(0, 0)$, where

$$Y_{lm}(0, 0) = \begin{cases} \sqrt{\frac{2l+1}{4\pi}} & m = 0 \\ 0 & m \neq 0 \end{cases} \quad (27)$$

Hence

$$E\{L_{lm}^2\} = \frac{2l+1}{4\pi(2l+1)} = \frac{1}{4\pi} \quad (28)$$

$$E\{L_{00}\} = \frac{1}{2\sqrt{\pi}} \quad (29)$$

Table 1 lists the eigenvectors with their eigenvalues and the cumulative sum of eigenvalues, corresponding to the variance accounted for (VAF). Figure 1 shows the first five eigenimages generated from those vectors.

5.3. Point Light and Ambient Source

A more interesting case is the when there is a point light source and an ambient component. As before we position the point light at $(0, 0)$ and use equation (25) to take all rotations into account. The ambient component only contributes to b_{00} . The other b_{lm} remain the same. This means that the top left element of the covariance matrix differs from the previous section. We also allow the level of the ambience, a to vary from zero to a_{max} .

$$E\{L_{00}^2\} = \frac{1}{a_{max}} \int_0^{a_{max}} \left(\frac{1}{2\sqrt{\pi}} + a\right)^2 da \quad (30)$$

and

$$E\{L_{lm}^2\} = \begin{cases} \frac{1}{4\pi} + \frac{a_{max}}{2\sqrt{\pi}} + \frac{a_{max}^2}{3} & l = m = 0 \\ \frac{1}{4\pi} & \text{otherwise} \end{cases} \quad (31)$$

$$E\{L_{00}\} = \frac{1}{2\sqrt{\pi}} + \frac{a_{max}}{2} \quad (32)$$

The eigenvectors, for $a_{max} = 1$, are in Table 2.

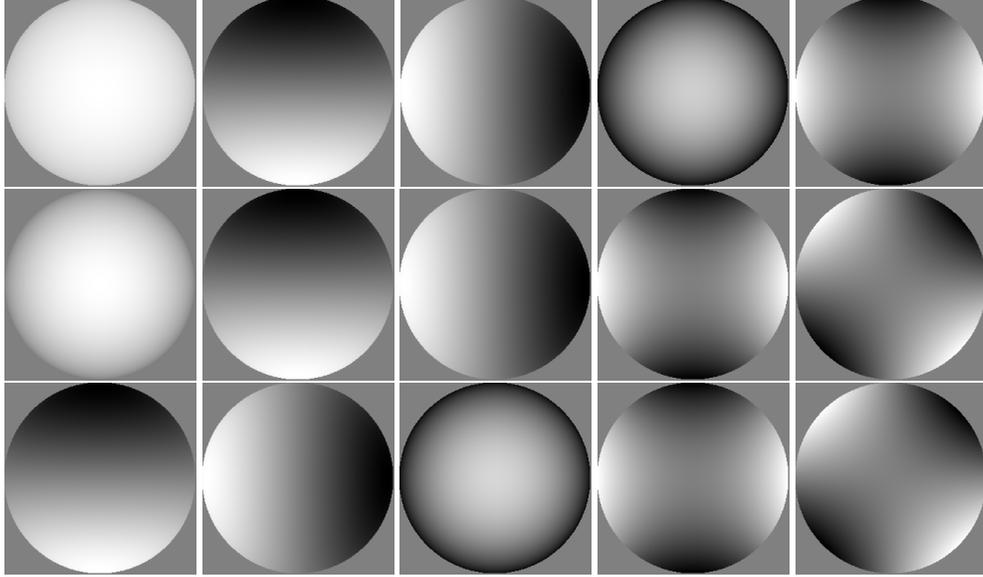


Figure 1: Eigenimages of a sphere under a point light source. The top row shows spheres created using the Non-Centered PCA, the middle row shows spheres from Standard PCA and the bottom row spheres from Modified PCA.

5.4. Two Point Light Sources

We can construct the distribution of all configurations of two point light sources. Position the first light source at $(0, 0)$ and the second at $(\psi, 0)$. Let ψ vary between 0 and π and let equation (25) take care of the rotations. The coefficients before rotation are $b_{lm} = Y_{lm}(0, 0) + Y_{lm}(\psi, 0)$ and

$$E\{L_{lm}^2\} = \sum_{n=-l}^l \frac{\int_0^\pi (Y_{ln}(0, 0) + Y_{ln}(\psi, 0))^2 d\psi}{\pi(2l+1)} \quad (33)$$

$$E\{L_{00}\} = \frac{1}{\pi} \int_0^\pi (Y_{00}(0, 0) + Y_{00}(\psi, 0)) d\psi \quad (34)$$

Table 3 lists the eigenvectors.

5.5. Real World Illumination Maps

Of greater interest is to perform PCA over more realistic illuminations. What is a realistic illumination? Today there exist a number of databases of captured illumination maps from different scenes, so called environment maps or light probe images. These maps are captured with high dynamic range, to capture the whole range of illumination from weak indirect reflections to direct sunlight. We will derive the eigenimages of a sphere illuminated by these maps under all 3D-rotations. We have used the light probe images of Debevec [4] that can be downloaded from <http://www.debevec.org>.

We have n_L light sources. Each light source $k = 1, \dots, n_L$ has spherical harmonic coefficients b_{klm} and a

probability of $P(k)$. Again, starting at (25) and summing over all light sources results in

$$E\{L_{lm}^2\} = \sum_{k=1}^{n_L} \sum_{n=-l}^l \frac{b_{kln}^2 P(k)}{(2l+1)} \quad (35)$$

$$E\{L_{00}\} = \sum_{k=1}^{n_L} b_{k00} P(k) \quad (36)$$

With each illumination map being equally probable $P(k) = \frac{1}{n_L}$. The eigenvectors are in Table 4.

5.6. Discussion

The main difference between the PCA variants is how they behave under variations in the ambient component of the illumination. When the ambient component varies, such as in Table 2 and 4, this gets encoded in the Non-Centered and Standard bases. The Modified basis on the other hand is not affected at all by this. For instance, the Modified bases created from the point light and the point plus ambient light, are identical.

Moreover, when there are variations in the ambient component the first eigenfunction of the Non-Centered basis becomes approximately Y_0 . Subsequent eigenfunctions are orthogonal to the first and hence will be very similar to the Modified PCA basis. This becomes apparent in the next section where the two bases get similar results. However, a Non-Centered basis trained with e.g. a point light source does not behave as well when representing images containing a varying ambient component.

Table 1: Eigenfunctions of Sphere under Point Light Source

Non-Centered PCA			Standard PCA			Modified PCA		
eigenfunction	λ	VAF	eigenfunction	λ	VAF	eigenfunction	λ	VAF
$0.88Y_0 + 0.48Y_2 + 0.039Y_6$	0.62	0.62	$0.99Y_2 + 0.1Y_6$	0.43	0.43	$0.99Y_1 + 0.17Y_5$	0.39	0.39
$0.99Y_1 + 0.17Y_5$	0.15	0.78	$0.99Y_1 + 0.17Y_5$	0.24	0.68	$0.99Y_3 + 0.17Y_7$	0.39	0.78
$0.99Y_3 + 0.17Y_7$	0.15	0.93	$0.99Y_3 + 0.17Y_7$	0.24	0.92	$-0.79Y_0 + 0.59Y_2 + 0.2Y_6$	0.13	0.9
$-0.82Y_0 + 0.51Y_2 + 0.28Y_6$	0.035	0.96	Y_8	0.024	0.95	Y_8	0.038	0.94
Y_8	0.015	0.98	Y_4	0.024	0.97	Y_4	0.038	0.98
Y_4	0.015	0.99	$-0.59Y_2 + 0.81Y_6$	0.019	0.99	$-0.78Y_3 + 0.63Y_7$	0.009	0.99
$-0.78Y_3 + 0.63Y_7$	0.003	1	$-0.78Y_3 + 0.63Y_7$	0.005	0.99	$0.78Y_1 - 0.63Y_5$	0.009	1
$-0.78Y_1 + 0.63Y_5$	0.003	1	$-0.78Y_1 + 0.63Y_5$	0.005	1	$0.65Y_0 - 0.7Y_2 + 0.3Y_6$	0.001	1
$0.65Y_0 - 0.7Y_2 + 0.3Y_6$	0.000	1	$0.65Y_0 - 0.7Y_2 + 0.29Y_6$	0	1	Y_0	0	1

Table 2: Eigenfunctions of Sphere under Point and Ambient Light Source

Non-Centered PCA			Standard PCA			Modified PCA		
eigenfunction	λ	VAF	eigenfunction	λ	VAF	eigenfunction	λ	VAF
$1Y_0 + 0.059Y_2 + 0.0042Y_6$	0.9	0.9	$0.89Y_0 + 0.46Y_2 + 0.037Y_6$	0.63	0.63	$0.99Y_1 + 0.17Y_5$	0.39	0.39
$0.99Y_1 + 0.17Y_5$	0.039	0.94	$0.99Y_1 + 0.17Y_5$	0.15	0.78	$0.99Y_3 + 0.17Y_7$	0.39	0.78
$0.99Y_3 + 0.17Y_7$	0.039	0.98	$0.99Y_3 + 0.17Y_7$	0.15	0.93	$-0.79Y_0 + 0.59Y_2 + 0.2Y_6$	0.13	0.9
$-0.79Y_0 + 0.52Y_2 + 0.2Y_6$	0.012	0.99	$-0.82Y_0 + 0.51Y_2 + 0.28Y_6$	0.035	0.96	Y_8	0.038	0.94
Y_8	0.003	0.99	Y_8	0.015	0.98	Y_4	0.038	0.98
Y_4	0.003	1	Y_4	0.015	0.99	$0.78Y_3 - 0.63Y_7$	0.009	0.99
$0.78Y_1 - 0.63Y_5$	0.000	1	$-0.78Y_3 + 0.63Y_7$	0.003	1	$0.78Y_1 - 0.63Y_5$	0.009	1
$-0.78Y_3 + 0.63Y_7$	0.000	1	$0.78Y_1 - 0.63Y_5$	0.003	1	$0.65Y_0 - 0.7Y_2 + 0.3Y_6$	0.001	1
$0.65Y_0 - 0.7Y_2 + 0.3Y_6$	0.000	1	$0.65Y_0 - 0.7Y_2 + 0.3Y_6$	0.000	1	Y_0	0	1

Table 3: Eigenfunctions of Sphere under Two Point Light Sources

Non-Centered PCA			Standard PCA			Modified PCA		
eigenfunction	λ	VAF	eigenfunction	λ	VAF	eigenfunction	λ	VAF
$0.97Y_0 + 0.26Y_2 + 0.025Y_6$	0.71	0.71	$0.99Y_2 + 0.13Y_6$	0.42	0.42	$0.98Y_1 + 0.21Y_5$	0.38	0.38
$0.98Y_1 + 0.21Y_5$	0.11	0.82	$0.98Y_1 + 0.21Y_5$	0.24	0.66	$0.98Y_3 + 0.21Y_7$	0.38	0.76
$0.98Y_3 + 0.21Y_7$	0.11	0.93	$0.98Y_3 + 0.21Y_7$	0.24	0.91	$-0.79Y_0 + 0.57Y_2 + 0.24Y_6$	0.13	0.89
$-0.81Y_0 + 0.52Y_2 + 0.29Y_6$	0.032	0.97	Y_8	0.028	0.94	Y_8	0.044	0.93
Y_8	0.013	0.98	Y_4	0.028	0.96	Y_4	0.044	0.98
Y_4	0.013	0.99	$-0.59Y_2 + 0.81Y_6$	0.022	0.99	$0.78Y_3 - 0.63Y_7$	0.01	0.99
$-0.78Y_1 + 0.63Y_5$	0.003	1	$-0.78Y_3 + 0.63Y_7$	0.006	0.99	$-0.78Y_1 + 0.63Y_5$	0.01	1
$-0.78Y_3 + 0.63Y_7$	0.003	1	$-0.78Y_1 + 0.63Y_5$	0.006	1	$0.65Y_0 - 0.7Y_2 + 0.3Y_6$	0.001	1
$0.65Y_0 - 0.7Y_2 + 0.3Y_6$	0.000	1	$0.65Y_0 - 0.7Y_2 + 0.29Y_6$	0	1	Y_0	0	1

Table 4: Eigenfunctions of Sphere under Natural Illumination

Non-Centered PCA			Standard PCA			Modified PCA		
eigenfunction	λ	VAF	eigenfunction	λ	VAF	eigenfunction	λ	VAF
$0.99Y_0 + 0.13Y_2 + 0.0056Y_6$	0.83	0.83	$0.86Y_0 + 0.51Y_2 + 0.024Y_6$	0.64	0.64	$1Y_3 + 0.097Y_7$	0.41	0.41
$1Y_1 + 0.097Y_5$	0.069	0.9	$1Y_1 + 0.097Y_5$	0.15	0.79	$1Y_1 + 0.097Y_5$	0.41	0.82
$1Y_3 + 0.097Y_7$	0.069	0.97	$1Y_3 + 0.097Y_7$	0.15	0.95	$-0.78Y_0 + 0.62Y_2 + 0.12Y_6$	0.12	0.94
$-0.79Y_0 + 0.6Y_2 + 0.13Y_6$	0.018	0.99	$-0.81Y_0 + 0.56Y_2 + 0.17Y_6$	0.029	0.98	Y_8	0.024	0.96
Y_8	0.004	0.99	Y_8	0.009	0.99	Y_4	0.024	0.99
Y_4	0.001	1	Y_4	0.009	0.99	$-0.77Y_3 + 0.64Y_7$	0.006	0.99
$-0.77Y_3 + 0.64Y_7$	0.001	1	$0.77Y_1 - 0.64Y_5$	0.002	1	$0.77Y_1 - 0.64Y_5$	0.006	1
$0.77Y_1 - 0.64Y_5$	0.001	1	$-0.77Y_3 + 0.64Y_7$	0.002	1	$0.65Y_0 - 0.7Y_2 + 0.3Y_6$	0.001	1
$0.65Y_0 - 0.7Y_2 + 0.3Y_6$	0.000	1	$0.65Y_0 - 0.7Y_2 + 0.3Y_6$	0.000	1	Y_0	0	1



Figure 2: Spheres rendered with groove, beach and stpeters illumination maps of [4].

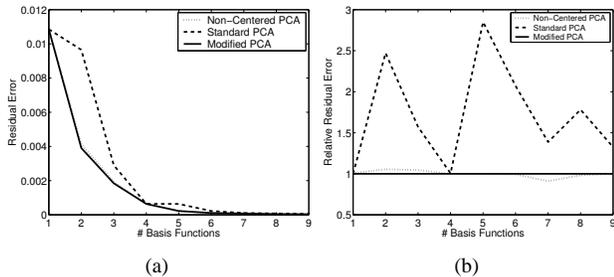


Figure 3: Comparison of PCA variants. a) error as the number of basis functions increases. b) errors relative to the Modified PCA basis.

6. Testing Representability

In order to compare the different bases, we rendered images of spheres using the illumination maps of Debevec. Naturally, it is more desirable to use real images. However, by using synthetic images we rule out many potential sources of errors. More importantly it gives us access to the geometry of the scene, needed for the experiments. The images were rendered using importance sampled Monte Carlo integration and by using real world illumination maps we get images with the full complexity of natural lighting. Figure 2 shows three of the rendered images.

Each basis was tested by calculating how well it represents the rendered images. The mean squared residual was used as a measure. For each basis the residual was calculated using one to all of the basis functions. The Standard PCA basis was completed with the constant function Y_0 which is proportional to the mean image. The Modified PCA basis was also completed with Y_0 as intended. Figure 3a shows the residual errors for bases of the three PCA variants constructed from the illumination map distribution. The Modified PCA basis exhibits the best representability. When plotting relative errors, with the Modified basis as a reference, we see that Standard PCA has around twice the error for the same number of components, Figure 3b.

When comparing bases constructed from different illumination distributions the differences are smaller. Figure 4 shows the relative residual errors. The basis constructed with a point light plus ambient source was not included since it is identical to the basis constructed with only a point light. The illumination map basis was constructed in a cross-validation sense so that the map used to render the image never was in the “training” set. Still it shows the best representability for the most part.

7. Light Field Estimation from a Single Image and Known Geometry

Light field estimation from several view points and known geometry has already been demonstrated using the spherical

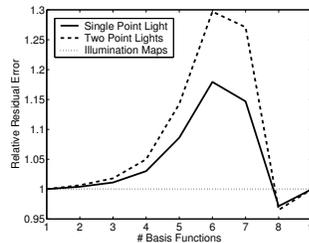


Figure 4: Comparison of bases constructed using different illumination distributions. Graph shows relative errors of bases constructed using the Modified PCA.

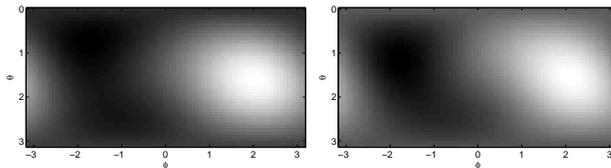


Figure 5: Low-pass filtered version of the groove light field (left) and the estimated light field (right). It was estimated using five principal components.

harmonic representation, [10]. When estimating the light field from a single image we need to take into account that the image lies in a subspace of lower dimensionality than the 9D space required by the spherical harmonic representation. The problem should be regularized. The natural way to do this is to estimate only the components with the highest variance in the image. Components with a low variance will “drown” in the noise, making the estimates error prone. The principal components are created to have as high variance as possible and are therefore ideal in this situation.

The procedure is as follows. Using e.g. least squares the principal component coefficients are estimated so that the image is written as a linear combination of the principal components.

$$I = \sum_{i=0}^N c_i U_i \quad (37)$$

Now, let \mathbf{c} be a vector containing the estimated coefficients and \mathbf{U} the matrix containing the column vectors of spherical harmonics coefficients of $U_i, i = 0, \dots, N$. Then the light field coefficients are given by

$$\mathbf{L} = \mathbf{Q}^{-1} \mathbf{U} \mathbf{c}, \quad (38)$$

where \mathbf{Q} is as in (8). Since the surface acts as low-pass filter on the illumination, only the lower frequencies of the light field is recoverable. Figure 5 shows the low-pass filtered light field of the groove illumination map and its estimated counterpart. The light field was estimated using five principal components plus the constant function (Y_0).

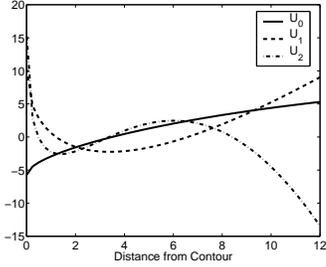


Figure 6: Eigenfunctions for apparent contours.

8. The Appearance of Apparent Contours

Using the same framework we can derive the principal component basis for apparent contours. We model the intensity profile of the surface near the apparent contour, along a line perpendicular to the contour. Assume we have spherical geometry in the cross-section of the object near the apparent contour. Since we are modeling a line as opposed to an image we set $\phi = 0$. θ varies from some value, depending on how far away from the contour we stop measuring, and $\pi/2$ (the value at the occluding contour). Say that we look at the intensity profile from 0 to d pixels distance from the contour and that the radius of the spherical cross-section is R then the elements of \mathbf{M} will be

$$m_{ss'} = k \int_{\theta_0}^{\pi/2} Y_s(\theta, 0) Y_{s'}(\theta, 0) \cos \theta d\theta, \quad (39)$$

where $\theta_0 = \arcsin(R - d)/R$ and k is a constant for normalizing the integral. With $d = 12$ and $R = 60$ the first three eigenfunctions have eigenvalues 0.97, 0.037 and 0.0002. They eigenfunctions are plotted in Figure 6. The functions show significant similarity to the eigenfunctions by Huggins et. al. [7] who created their basis from empirical data.

9. Conclusion

In [9] Ramamoorthi derives the principal component analysis on the space of images of a convex Lambertian object under all possible single point light sources. We have generalized his results to any illumination distribution. Our results show that the principal component basis is determined by two matrices. One depending on the objects shape and the other being the covariance matrix of the spherical harmonic coefficients of the illumination. We derive bases for a number of illumination distributions of interest.

Furthermore, we modify the criterion to derive the PCA. The resulting basis, we believe, is better for representing the image space. This basis is orthogonal to the constant function. When completing the basis with the constant function

this should allow for a better representation when using a truncated basis. Experimental data supports this.

Acknowledgements

This work was done within the EU-IST project Insight2+. The support is gratefully acknowledged.

References

- [1] R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 25(2):218–233, February 2003.
- [2] P. N. Belhumeur and D. J. Kriegman. What is the set of images of an object under all possible illumination conditions? *Int. Journal of Computer Vision*, 28(3):245–260, 1998.
- [3] G. S. Chirikjian and A. B. Kyatkin. *Engineering Applications of Noncommutative Harmonic Analysis*. CRC Press, 2001.
- [4] P. Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH*, pages 189–198, 1998.
- [5] R. Epstein, P.W. Hallinan, and A.L. Yuille. 5+/-2 eigenimages suffice: An empirical investigation of low-dimensional lighting models. In *IEEE Workshop Physics-Based Modeling in Computer Vision*, pages 108–116, 1995.
- [6] P. Hallinan. A low-dimensional representation of human faces for arbitrary lighting conditions. In *Proc. CVPR*, pages 995–999, 1994.
- [7] P.S. Huggins, H.F. Chen, P.N. Belhumeur, and S.W. Zucker. Finding folds: On the appearance and identification of occlusion. In *Proc. CVPR*, pages II:718–725, 2001.
- [8] R. A. Johnson and D. W. Wichern. *Applied Multivariate Statistical Analysis, 4:th ed.* Prentice-Hall, 1998.
- [9] R. Ramamoorthi. Analytic pca construction for theoretical analysis of lighting variability in images of a lambertian object. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 24(10), October 2002.
- [10] R. Ramamoorthi and Hanrahan P. On the relationship between radiance and irradiance: determining the illumination from images of a convex lambertian object. *JOSA-A*, 18(10):2448–2458, October 2001.