Phenomenological Eigenfunctions for Image Irradiance

Peter Nillius Jan-Olof Eklundh

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

Abstract

We present a framework for calculating low-dimensional bases to represent image irradiance from surfaces with isotropic reflectance under arbitrary illumination. By representing the illumination and the bidirectional reflectance distribution function (BRDF) in frequency space, a model for the image irradiance is derived. This model is then reduced in dimensionality by analytically constructing the principal component basis for all images given the variations in both the illumination and the surface material. The principal component basis are constructed in such a way that all the symmetries (Helmholtz reciprocity and isotropy) of the BRDF are preserved in the basis functions. Using the framework we calculate a basis using a database of natural illumination and the CURET database containing BRDFs of real world surface materials.

1. Introduction

The appearance of a surface depends on its shape, material and the illumination. To effectively do many computer vision tasks, including object recognition we must be able to cope with all these aspects.

One step towards dealing with illumination changes is given by Basri and Jacobs [1] and Ramamoorthi and Hanrahan [12]. They realized that a surface with Lambertian reflectance acts as a low-pass filter on the incident light. This makes the image irradiance in practice band-limited and its representation in frequency space is given by a finite number of basis functions. More precisely they suggest that nine basis functions is enough to represent the image irradiance from a Lambertian surface under arbitrary illumination.

In an image, only one hemisphere of the surface normals are visible. This reduces the variability of the images. In [11] Ramamoorthi derives analytically the principal component basis of all images of a Lambertian surface under a varying point light source. His results show that about five basis functions are enough to represent this space, something that corresponds well with earlier empirical work [6, 5]. In [10] we extended the analytic principal component analysis to any illumination distribution.

The described works deal primarily with Lambertian surfaces. For an algorithm to be really useful it must handle more realistic types of surface reflectance. In this paper we consider materials with non-Lambertian reflectance properties. Many materials, although not Lambertian, still act as low-pass filters on the incident illumination. This means that the reflected light can be represented in frequency space with a finite number of basis functions. We derive such a basis. Using this basis we analytically construct the principal components of the images of an object when both the illumination and surface reflectance properties vary. We assume an isotropic BRDF. The principal components are created in such a way that all the symmetries of the BRDF (isotropy and Helmholtz reciprocity) are preserved in the principal components. The goal is to obtain bases that can be used to represent the image irradiance from a wide range of materials under a wide range of illuminations.

2. Image Formation

The image irradiance, E, is proportional to the scene radiance and can be written as (see [7]),

$$E(\theta_r, \phi_r) = \int_{\mathbf{H_{S2}}} L(\theta_i, \phi_i) f(\theta_i, \phi_i, \theta_r, \phi_r) \cos\theta_i d\omega_i, \quad (1)$$

where L is the incident radiance (or the light field) and f is the BRDF, which gives the ratio of reflected radiance in direction (θ_r, ϕ_r) and the irradiance due to incident light from (θ_i, ϕ_i), see Figure 1a.

The coordinates in (1) are defined in the local frame set by the surface normal. To use the light field defined in the global frame we need to rotate (θ_i, ϕ_i) to a global frame. We define the global frame to have its z-axis along the optical axis, pointing towards the camera and the x- and y-axes in the image plane. For a point with surface normal (α, β, γ) in the global frame, see Figure 1b, the incident direction in the global frame can be written as $(\theta, \phi) = R_{\alpha,\beta,\gamma}(\theta_i, \phi_i)$. Now, the image irradiance is

$$E(\alpha, \beta, \gamma, \theta_r, \phi_r) = \int_{\mathbf{H_{S2}}} L(R_{\alpha,\beta,\gamma}(\theta_i, \phi_i)) f(\theta_i, \phi_i, \theta_r, \phi_r) \cos\theta_i d\omega_i \quad (2)$$

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE

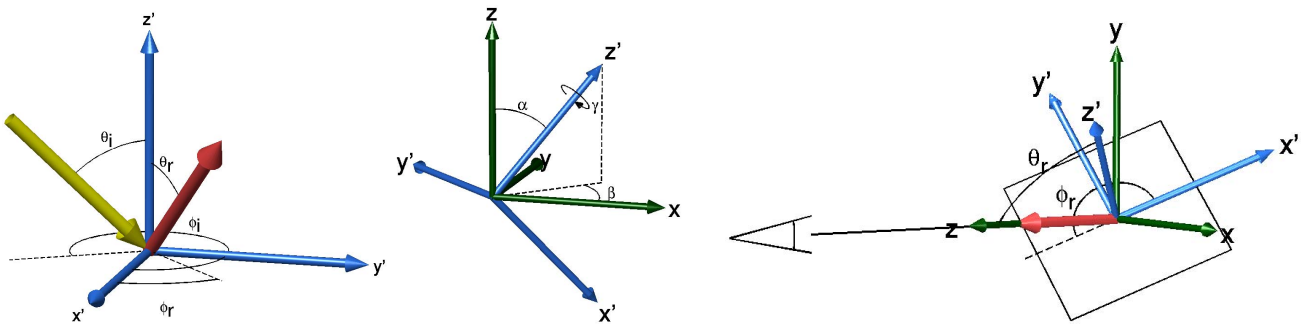

Figure 1: Geometry of reflectance. a) Incident and reflected beam in the local coordinate frame. b) Relative position of global and local frames. c) Assuming orthographic projection the reflected beam aligns with the optical axis leading to $\theta_r = \alpha$ and $\phi_r = \pi$

The parameter γ defines the rotation of the local frame around the surface normal. It should be set so that the local frame aligns with the "direction" of the material on the surface (to correspond with the parameterization of the BRDF). Isotropic materials have no direction. In that case γ can be selected arbitrarily. Hence, we can set it to $\gamma = 0$.

Assuming orthographic projection the viewing geometry has further implications for the parameters of (2). As can be seen in Figure 1c the reflected beam aligns with the optical axis (the z-axis in the global frame) which means that $\theta_r = \alpha$. Furthermore, setting $\gamma = 0$ leads to $\phi_r = \pi$. So, for isotropic materials we can eliminate θ_r and ϕ_r from the equation. Now the image irradiance is

$$E(\alpha, \beta) = \int_{\mathbf{H_{S^2}}} L(R_{\alpha,\beta,0}(\theta_i, \phi_i)) f(\theta_i, \phi_i, \alpha, \pi) \cos\theta_i d\omega_i. \quad (3)$$

2.1. Light Field Representation

As in [1, 12, 11] we represent the light field in spherical harmonics.

$$L(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} L_l^m Y_l^m(\theta, \phi) \quad (4)$$

where L_l^m are the spherical harmonic coefficients of the light field and $Y_l^m(\theta, \phi)$ are the spherical harmonic basis functions. The spherical harmonic basis functions form the Fourier basis on the sphere.

One property of the spherical harmonics is that a rotated basis function can be written as a linear combination of the other basis functions of the same order, see Appendix. The representation of the rotated light field in (3) is therefore

$$L(R_{\alpha,\beta,\gamma}(\theta_i, \phi_i)) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{n=-l}^{l} D_l^{mn}(\alpha, \beta, \gamma) L_l^m Y_l^n(\theta_i, \phi_i), \quad (5)$$

where $D_l^{mn}(\alpha, \beta, \gamma)$ are the rotation reparameterization coefficients for the spherical harmonics.

2.2. Isotropic BRDF Representation

We use the isotropic BRDF representation of Koenderink and Van Doorn [9]. It offers a compact orthonormal representation of the BRDF while incorporating the very general Helmholtz reciprocity

$$f(\theta_i, \phi_i, \theta_r, \phi_r) = f(\theta_r, \phi_r, \theta_i, \phi_i) \quad (6)$$

The BRDF is represented as

$$f(\theta_i, \theta_r, \Delta\phi_{ir}) = \sum_{opq} b_{op}^q I_{op}^q(\theta_i, \theta_r, \Delta\phi_{ir}), \quad (7)$$

where

$$I_{op}^q(\theta_i, \theta_r, \Delta\phi_{ir}) = \frac{1}{2\pi} \sqrt{\frac{(o+1)(p+1)}{(1+\delta_{op})(1+\delta_{q0})}} \times \left(R_o^q(\sqrt{2}\sin\frac{\theta_i}{2}) R_p^q(\sqrt{2}\sin\frac{\theta_r}{2}) + R_p^q(\sqrt{2}\sin\frac{\theta_i}{2}) R_o^q(\sqrt{2}\sin\frac{\theta_r}{2}) \right) \times \cos q\Delta\phi_{ir} \quad (8)$$

$R_o^q(\rho)$ are the Zernike polynomials. We have used the Kronecker delta for a more compact notation of the normalization factor, $\delta_{ij} = 1$ if $i = j$ and 0 otherwise.

The restrictions on the indices are $o \geq p \geq q \geq 0$ and $(o - q)$ and $(p - q)$ are even.

2.3. Image Irradiance Representation

Inserting the representations (5) and (7) into the irradiance expression (3) we get

$$E(\alpha, \beta) = \sum_{\substack{lmn \\ opq}} L_l^m b_{op}^q D_l^{mn}(\alpha, \beta, 0) \times \int_{\mathbf{H_{S^2}}} Y_l^n(\theta_i, \phi_i) I_{op}^q(\theta_i, \alpha, \phi_i - \pi) \cos\theta_i d\omega_i \quad (9)$$

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE

Now consider the integral in equation (9). Inserting the expressions for the spherical harmonic (see (46) in the appendix) and BRDF basis functions (8), we see that it can be separated into an azimuthal integral over ϕ_i and a polar integral over θ_i. The azimuthal integral is easily solved and will eliminate all terms in the sum (9) except when $n = q$.

$$\int_0^{2\pi} \Phi_n(\phi_i) \cos q(\phi_i - \pi) d\phi_i = \delta_{nq} \pi (-1)^q \sqrt{2(1 + \delta_{q0})} \quad (10)$$

The polar integral can also be calculated analytically but the expression is complicated and for now we denote it

$$C_{lo}^q = \int_0^{\frac{\pi}{2}} P_l^q(\cos\theta_i) R_o^q(\sqrt{2}\sin\frac{\theta_i}{2}) \cos\theta_i \sin\theta_i d\theta_i. \quad (11)$$

Inserting these results in (9) we get an expression for the irradiance

$$E(\alpha,\beta) = \sum_{\substack{lm \\ opq}} L_l^m b_{op}^q K_{lop}^q D_l^{mq}(\alpha,\beta,0) \times \left(C_{lo}^q R_p^q(\sqrt{2}\sin\frac{\alpha}{2}) + C_{lp}^q R_o^q(\sqrt{2}\sin\frac{\alpha}{2}) \right), \quad (12)$$

where

$$K_{lop}^q = (-1)^q N_l^q \sqrt{\frac{(o+1)(p+1)}{2(1+\delta_{op})}}. \quad (13)$$

Let

$$E_{lop}^{mq}(\alpha,\beta) = K_{lop}^q D_l^{mq}(\alpha,\beta,0) \left(C_{lo}^q R_p^q(\sqrt{2}\sin\frac{\alpha}{2}) + C_{lp}^q R_o^q(\sqrt{2}\sin\frac{\alpha}{2}) \right). \quad (14)$$

Then

$$E(\alpha,\beta) = \sum_{\substack{lm \\ opq}} L_l^m b_{op}^q E_{lop}^{mq}(\alpha,\beta). \quad (15)$$

The restrictions on the indices are $o \geq p \geq q \geq 0$, $(o-q)$ and $(p-q)$ are even, $l \geq q$ and $m = -l, \ldots, l$.

We have derived a basis for the image irradiance from isotropic materials under arbitrary illumination. The image irradiance is represented as a sum of the basis functions, E_{lop}^{mq} multiplied by the coefficients of the light source, L_l^m and the coefficients of the BRDF of the material, b_{op}^q. In Section 3 we derive a low-parameter approximation of this seemingly complicated representation using principal component analysis.

2.4. How many basis functions are required?

Image irradiance from a Lambertian surface has been shown to be well approximated by 9 basis functions [1, 12]. The number of basis functions required for a non-Lambertian surface depends though on the surface's BRDF. A surface with low-order scattering modes (I_{op}^q) acts as a low-pass filter on the incoming illumination making the reflected light band-limited. Hence it can be represented with a limited number of basis functions. A surface with higher order scattering modes passes through higher frequencies of the illumination and thus more basis functions are needed to represent the reflected light.

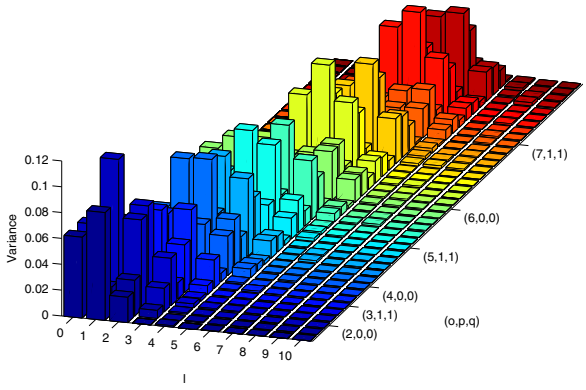

Figure 2: Image variance of basis functions.

We can illustrate this by calculating the contributing signal variance of each basis function E_{lop}^{mq}. To compute the signal variance in the image we need to assume a geometry of the surface, or at least, a surface normal distribution of the surface in the image. This distribution will due to foreshortening be proportional to the cosine of the slant angle, $\cos\alpha$ in our case. The variance of each basis function can then be computed as $\int_{\mathbf{H_S^2}} E_{lop}^{mq}(\alpha,\beta)^2 \cos\alpha d\omega$.

Figure 2 shows the signal variance of each basis function plotted with the illumination modes on one horizontal axis and the surface modes on the other. The variances for each illumination order l, $m = -l, \ldots, l$ have been summed together for a more viewable graph. The first row shows the variances for the Lambertian case. We can again see that orders up to $l = 2$ (9 basis functions) suffice to represent most of the variance. Higher order modes require more basis functions. For instance a surface having modes up to order $o = 4$ requires basis functions with orders of l up to about 5 (464 basis functions). A 7-order surface requires basis functions up to order $l = 7$ (2264 basis functions).

As can be seen, for non-Lambertian surfaces more and more basis functions are required to represent the image irradiance. But these basis functions are not orthogonal in the space of images. To orthogonalize them we will construct the principal component basis.

2.5. Ordering of the basis functions

It is convenient to order the basis functions so that they can be indexed by a single variable. We order them in decreas-

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE

ing variance using the same variance as in the previous section. In single index notation the image irradiance can be written as

$$E(\alpha,\beta) = \sum_s c_s E_s(\alpha,\beta), \tag{16}$$

where $c_s = L_l^m b_{op}^q$ and l, m, o, p and q are given by s due to the ordering of the functions.

3. Principal Component Analysis

Using the basis derived in Section 2.3, we here analytically derive the principal components of the image space when both the illumination and surface reflectance vary.

3.1. Criteria for Deriving 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 principal component, U_0, is the direction which has the highest variability. I.e. U_0 maximizes the variance of the scalar product of itself and the images, or in mathematical terms:

$$U_0 = \underset{||U||=1}{\operatorname{argmax}} \operatorname{Var}\{I \bullet U\}, \tag{17}$$

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

We will 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, $I = \sum_i d_i U_i$. Typically images are analyzed by estimating the coefficients d_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 it is as efficient as possible when truncated. Moreover, it allows the coefficients d_i to be estimated individually using the scalar product, $d_i = I \bullet U_i$, as opposed to estimating them all simultaneously using e.g. least squares.

Secondly, the components $d_i U_i$ 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. The variance of each component $d_i U_i$ is simply $d_i^2 = (I \bullet U_i)^2$.

Thirdly, the basis should contain the constant function. The illumination frequently contains an ambient component which can vary. Including the constant function in the basis makes the remaining functions of the basis independent of the ambient component. The constant function is included in the basis by subtracting 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 of ones. We arrive at the following criterion.

$$U_0 = \underset{||U||=1}{\operatorname{argmax}} \mathrm{E}\{((I - (I \bullet \mathbf{1})\mathbf{1}) \bullet U)^2\} \tag{18}$$

To conclude, the Standard PCA criterion (17) maximizes the variance over illumination and material changes while criterion (18) maximizes the variance in the image. In more detail, the difference between the two criteria is how the data is centered. Standard PCA subtracts the mean image from the dataset. Criterion (18) suggest that this should not be done. Instead each image should be centered individually by subtracting the individual mean. We call the method *Image-Centered PCA*.

3.2. The Image Space

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_l^m. Furthermore, let the variation in an objects material be described as a distribution of the coefficients of the BRDF representation (7), $p_b(\mathbf{b})$, where $\mathbf{b}$ is a vector containing the elements b_{op}^q .

If we assume convex objects there are no cast shadows or interreflections. This means that the image irradiance is uniquely determined by the surface normal. Given that we have the surface normal of an object at each point in the image we can generate all possible images of that object under the illumination and material distribution we have defined. However, the criteria we use for deriving PCA depend only 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, $p_{\hat{n}}(\alpha,\beta), (\alpha,\beta) \in \mathbf{H_{S^2}}$, we can calculate the scalar product of images I and J as

$$I \bullet J = \int_{\mathbf{H_{S^2}}} i(\alpha,\beta) j(\alpha,\beta) p_{\hat{n}}(\alpha,\beta) d\omega, \tag{19}$$

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

So, for our purposes the distribution of the illumination, the material and the surface normals are enough to define the image space.

3.3. Deriving the Basis

We begin by writing the eigenfunction as a sum of the basis functions of the image irradiance. This ensures that the

[1] Due to our definition of the scalar product we need not to divide with the number of pixels

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE

symmetries of the BRDF are preserved in the eigenfunctions.

$$U(\alpha,\beta) = \sum_s u_s E_s(\alpha,\beta) \tag{20}$$

The scalar product of I and U becomes

$$I \bullet U = \sum_{s,s'} c_s u_{s'} \int_{\mathbf{H_{S^2}}} p_{\hat{n}}(\alpha,\beta) E_s(\alpha,\beta) E_{s'}(\alpha,\beta) d\omega \tag{21}$$

We can rewrite this in matrix form.

$$I \bullet U = \mathbf{u}^T \mathbf{M} \mathbf{c}, \tag{22}$$

where $\mathbf{c}$ is the vector containing the products of the coefficients of the light source and the material $L_l^m b_{op}^q$, $\mathbf{M}$ contains the elements

$$m_{ss'} = \int_{\mathbf{H_{S^2}}} p_{\hat{n}}(\alpha,\beta) E_s(\alpha,\beta) E_{s'}(\alpha,\beta) d\omega \tag{23}$$

and $\mathbf{u}$ is a vector containing the coefficients for U.

Now, the only random variable in (22) is $\mathbf{c}$. Let $\mathrm{Covar}\{\mathbf{c}\} = \mathbf{\Sigma_c}$. Then

$$\mathrm{Var}\{\mathbf{u}^T \mathbf{M} \mathbf{c}\} = \mathbf{u}^T \mathbf{M} \mathbf{\Sigma_c} \mathbf{M} \mathbf{u}. \tag{24}$$

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

$$\begin{aligned} U \bullet U &= \sum_{s,s'} u_s u_{s'} \int_{\mathbf{H_{S^2}}} p_{\hat{n}}(\alpha,\beta) E_s(\alpha,\beta) E_{s'}(\alpha,\beta) d\omega \\ &= \sum_{s,s'} u_s u_{s'} m_{ss'} = \mathbf{u}^T \mathbf{M} \mathbf{u} = 1. \end{aligned} \tag{25}$$

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

$$\mathbf{u_0} = \underset{\mathbf{u}^T \mathbf{M} \mathbf{u} = 1}{\operatorname{argmax}} \mathbf{u}^T \mathbf{M} \mathbf{\Sigma_c} \mathbf{M} \mathbf{u} \tag{26}$$

Applying the coordinate transform

$$\mathbf{v} = \mathbf{M}^{1/2} \mathbf{u}, \tag{27}$$

we obtain

$$\mathbf{v_0} = \underset{\mathbf{v}^T \mathbf{v} = 1}{\operatorname{argmax}} \mathbf{v}^T \mathbf{M}^{1/2} \mathbf{\Sigma_c} \mathbf{M}^{1/2} \mathbf{v}. \tag{28}$$

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 $\mathbf{M}^{1/2} \mathbf{\Sigma_c} \mathbf{M}^{1/2}$ 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 can convince ourselves that orthogonality of the vectors $\mathbf{v}$ corresponds to orthogonality in the image space.

$$U_i \bullet U_j = \mathbf{u_i}^T \mathbf{M} \mathbf{u_j} = \mathbf{v_i}^T \mathbf{v_j} \tag{29}$$

The coefficients for the eigenimages U_i are computed by $\mathbf{u_i} = \mathbf{M}^{-1/2} \mathbf{v_i}$.

When constructing the basis according to the second criterion (18) we should subtract the image mean. The image mean is

$$I \bullet \mathbf{1} = \sum_s c_s \int_{\mathbf{H_{S^2}}} E_s(\alpha,\beta) p_{\hat{n}}(\alpha,\beta) d\omega = \sum_s c_s e_s, \tag{30}$$

where e_s is the mean of E_s over the surface normal distribution.

Now,

$$\begin{aligned} (I - (I \bullet \mathbf{1})\mathbf{1}) \bullet U &= \sum_{s,s'} c_s u_{s'} \\ &\times \int_{\mathbf{H_{S^2}}} (E_s(\alpha,\beta) - e_s) E_{s'}(\alpha,\beta) p_{\hat{n}}(\alpha,\beta) d\omega \\ &= \sum_{s,s'} c_s u_{s'} (m_{ss'} - e_s e_{s'}) = \mathbf{u}^T (\mathbf{M} - \mathbf{e}\mathbf{e}^T) \mathbf{c} \end{aligned} \tag{31}$$

Furthermore

$$\mathrm{E}\{((I - (I \bullet \mathbf{1})\mathbf{1}) \bullet U)^2\} = \mathbf{u}^T (\mathbf{M} - \mathbf{e}\mathbf{e}^T) \mathbf{V_c} (\mathbf{M} - \mathbf{e}\mathbf{e}^T) \mathbf{u} \tag{32}$$

where $\mathbf{V_c} = \mathrm{E}\{\mathbf{c}\mathbf{c}^T\}$ which corresponds to the covariance matrix $\mathbf{\Sigma_c}$ calculated without subtracting the mean.

Applying the same coordinate change as previously we obtain the $\mathbf{v}$ vectors as the eigenvectors of the matrix

$$\mathbf{M}^{-1/2} (\mathbf{M} - \mathbf{e}\mathbf{e}^T) \mathbf{V_c} (\mathbf{M} - \mathbf{e}\mathbf{e}^T) \mathbf{M}^{-1/2} \tag{33}$$

and finally the coefficients of the eigenimages by $\mathbf{u_i} = \mathbf{M}^{-1/2} \mathbf{v_i}$ as before.

3.4. Calculating the Covariances

It is worth making some comments about calculating the covariance matrices $\mathbf{\Sigma_c}$ and $\mathbf{V_c}$. The two matrices are related by

$$\mathbf{\Sigma_c} = \mathbf{V_c} - \mu_\mathbf{c} \mu_\mathbf{c}^T \tag{34}$$

where $\mu_\mathbf{c}$ is the mean vector of $\mathbf{c}$. Since we can assume that the illumination and the surface reflectance are independent[2] the calculation of the elements of $\mu_\mathbf{c}$ and $\mathbf{V_c}$ can be partitioned into

$$\mu_{c_s} = \mathrm{E}\{L_l^m\} \mathrm{E}\{b_{op}^q\} \tag{35}$$

$$\mathrm{E}\{c_s c_{s'}\} = \mathrm{E}\{L_l^m L_{l'}^{m'}\} \mathrm{E}\{b_{op}^q b_{o'p'}^{q'}\} \tag{36}$$

The indices l, m, o, p and q are given by s due to the ordering of the basis functions.

[2] Though one could argue that e.g. the illumination from an overcast sky correlates with shiny materials in the form of raincoats.

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE

4. Creating a General Purpose Basis

Can we create a basis that can be used to represent the image irradiance from a wide variety of materials under a wide variety of illuminations? Well, at least we can try. In this section we construct a basis using a database of real world surface reflectances, the CURET database [3], and a database of captured illumination, Debevec [4].

To create the basis we need to compute the covariances of the illumination and material coefficients.

4.1. Calculating the Illumination Covariances

Let the illumination distribution be each of the n_L illumination maps, from the database, rotated over all 3D rotations, $SO(3)$. Let $a^m_{l,k}$ be the spherical harmonic coefficients for illumination map k. The coefficients after rotation are $\sum_{n=-l}^{l} D^{mn}_l(\alpha,\beta,\gamma)a^n_{l,k}$. Because of the orthogonality relation

$$\int_{SO(3)} D^{mn}_l(R)D^{m'n'}_{l'}(R)dR = \frac{\delta_{l,l'}\delta_{m,m'}\delta_{n,n'}}{2l+1} \quad (37)$$

the covariances $\mathrm{E}\{L^m_l L^{m'}_{l'}\}$ will be non-zero only when $l = l'$ and $m = m'$ and

$$\mathrm{E}\{(L^m_l)^2\} = \sum_{k=1}^{n_L}\sum_{n=-l(s)}^{l(s)} \frac{(a^n_{l,k})^2}{n_L(2l(s)+1)} \quad (38)$$

Furthermore, due to $D^{00}_0 = 1$ and (37) the means $\mathrm{E}\{L^m_l\}$ are zero for all elements but L^0_0 and

$$\mathrm{E}\{L^0_0\} = \frac{1}{n_L}\sum_{k=1}^{n_L} a^0_{0,k}. \quad (39)$$

4.2. Calculating the Material Covariances

Calculating the covariances of the material coefficients involves no rotations. Let each of the n_b materials in the database have coefficients $b^q_{op,k}$. The covariances and means are calculated as

$$\mathrm{E}\{b^q_{op}b^{q'}_{o'p'}\} = \frac{1}{n_b}\sum_{k=1}^{n_b} b^q_{op,k}b^{q'}_{o'p',k} \quad (40)$$

$$\mathrm{E}\{b^q_{op}\} = \frac{1}{n_b}\sum_{k=1}^{n_b} b^q_{op,k} \quad (41)$$

4.3. Calculating the Basis

In calculating the basis we need to assume a geometry in the form of a surface normal distribution. We have chosen it to be proportional to the foreshortened area of the tangent plane, a factor $\cos\alpha$.

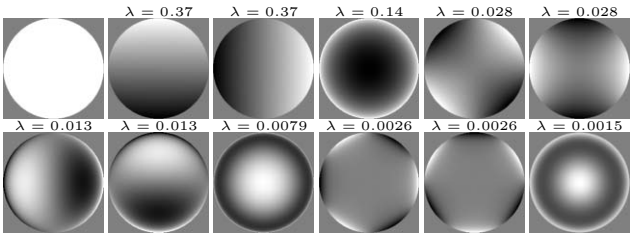

Figure 3: The first 12 basis functions for a sphere created from set of real world materials and illuminations using Image-Centered PCA. The top left image is the constant basis function added to the basis after the PCA is performed.

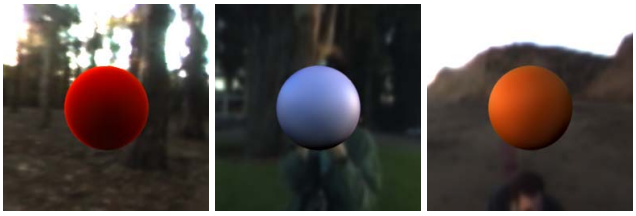

Figure 4: Three of the rendered images. From left to right, velvet in the grove, leather on campus and an orange on the beach. Notice the characteristic bright rim of the velvet sphere and the specularity on the leather.

We chose basis functions, E^{mq}_{lop}, up to order $l = 9$ and $o = 7$. From those we selected the 1000 with the highest variance. A basis was then calculated using the standard PCA criterion and the Image-Centered criterion. Figure 3 shows the first eigenimages for the Image-Centered PCA.

4.4. Testing Representability

To test the basis we rendered images of a sphere of all materials under all illuminations. The rendering was done similarly to [13] although we use the Koenderink, van Doorn basis for the BRDF. The rendering involves summing the contributions from each basis function. By calculating the variance of each component we get a bound for the error and can render to a very high accuracy. Figure 4 shows a few of the rendered images.

The testing was done by fitting the basis to each of the images varying the number of principal components in the basis. The error was chosen to be the variance of the residual divided by the total variance of the image. Figure 5 shows the results. The Standard PCA and Image-Centered PCA bases show similar results. Note that fitting the Standard PCA basis requires subtracting the mean image before fitting the basis. This is not required for the Image-Centered basis.

The number of components required to represent an image depends on the material. This is illustrated by Figure 6 which shows the errors for each material. Many materi-

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE

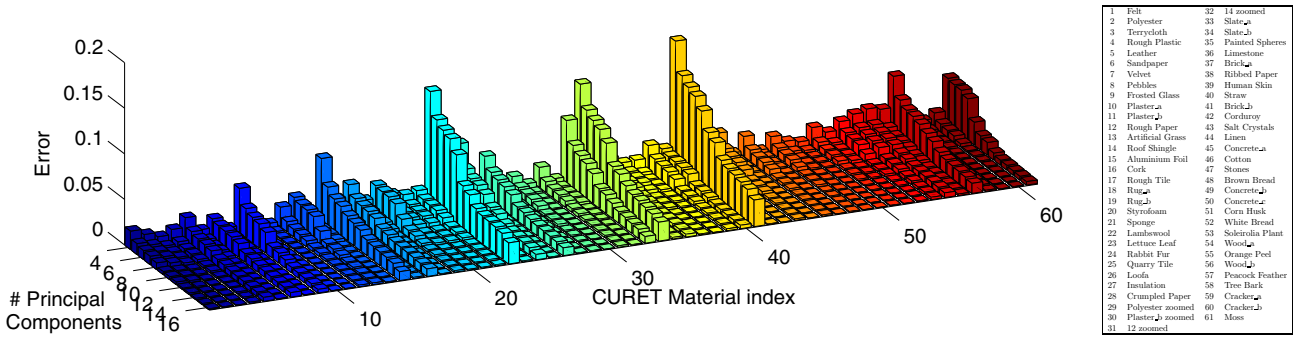

1	Felt	32	14 zoomed
2	Polyester	33	Slate_a
3	Terrycloth	34	Slate_b
4	Rough Plastic	35	Painted Spheres
5	Leather	36	Limestone
6	Sandpaper	37	Brick_a
7	Velvet	38	Ribbed Paper
8	Pebbles	39	Human Skin
9	Frosted Glass	40	Straw
10	Plaster_a	41	Brick_b
11	Plaster_b	42	Corduroy
12	Rough Paper	43	Salt Crystals
13	Artificial Grass	44	Linen
14	Roof Shingle	45	Concrete_a
15	Aluminium Foil	46	Cotton
16	Cork	47	Stones
17	Rough Tile	48	Brown Bread
18	Rug_a	49	Concrete_b
19	Rug_b	50	Concrete_c
20	Styrofoam	51	Corn Husk
21	Sponge	52	White Bread
22	Lambswool	53	Soleirolia Plant
23	Lettuce Leaf	54	Wood_a
24	Rabbit Fur	55	Orange Peel
25	Quarry Tile	56	Wood_b
26	Loofa	57	Peacock Feather
27	Insulation	58	Tree Bark
28	Crumpled Paper	59	Cracker_a
29	Polyester zoomed	60	Cracker_b
30	Plaster_b zoomed	61	Moss
31	12 zoomed		

Figure 6: Error as the number of principal components increases for each material in the CURET database.

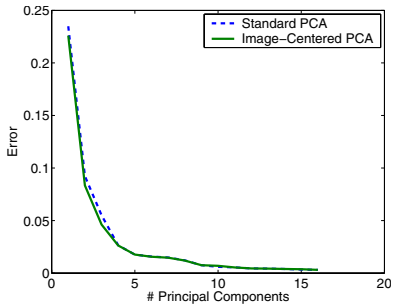

Figure 5: Average error for the whole data-set as the number of principal components increases for the basis created with Standard PCA and Image-Centered PCA.

als require only a small (around five) number of components, e.g. many of the matte materials, but also highly non-Lambertian materials like velvet, see Figure 7a.

More specular materials, such as the leather in the database require more components. The diffuse shading is recreated with a few components but to recreate the specularity 30-50 components are required in this case, see Figure 7b.

5. The Benefits of Analytic PCA

There are many benefits of doing the PCA analytically. For one thing many sources of error are eliminated. Also, by doing the rotation of the illumination analytically we can compute bases for a wide variety of conditions with very little effort. Instead of capturing thousands and thousands of images with varying illumination and material conditions, it is enough to capture a small number of real world illuminations and surface reflectances. Moreover, the shape of the object is easily changed by recomputing the matrix $\mathbf{M}$.

Another important feature is that we can relate the principal component coefficients back to the parameters of the illumination and the surface reflectance. E.g. if $\mathbf{d}$ are the estimated coefficients of an image and $\mathbf{U}$ is the matrix containing the coefficients of the principal components.

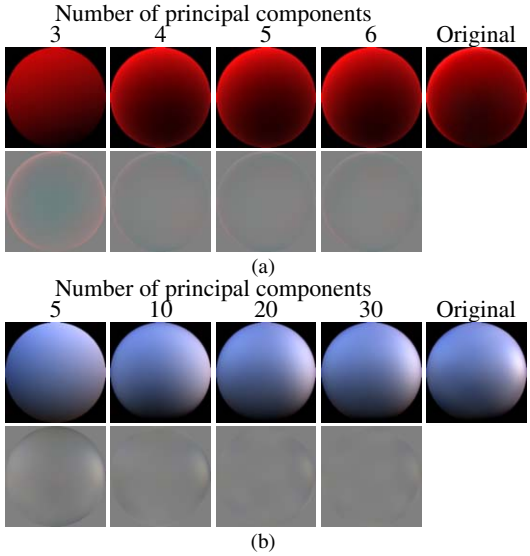

Figure 7: Reconstructed and residual images of a) velvet and b) leather.

Then the illumination and surface reflectance coefficients $c_s = L_l^m b_{op}^q$ can be computed as

$$\mathbf{c} = \mathbf{U}\mathbf{d} \tag{42}$$

For e.g. rendering it is useful to do the reverse transformation $\mathbf{d} = \mathbf{U}^T\mathbf{M}\mathbf{c}$ (as $\mathbf{U}^{-1} = \mathbf{U}^T\mathbf{M}$).

6. Conclusions

We have derived a basis for the image irradiance with the illumination in spherical harmonic representation and the BRDF represented by the Koenderink, van Doorn basis for isotropic surface reflectance. Using this basis we have shown how to analytically construct the principal component basis of the image space of an object when there are variations in both the illumination and the surface reflectance properties. The variation is expressed as covariances and means of the coefficients of the illumination and

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE

the BRDF. The goal has been to create a framework for constructing low-parameter bases that can be used to represent the image irradiance from a wide variety of materials under a wide variety of illuminations.

We demonstrate the framework by constructing a PCA basis from two databases containing captured illuminations and reflectance properties of real materials.

A. Rotating Real Spherical Harmonics

The spherical harmonics basis functions are defined as

$$y_l^m(\theta,\phi) = N_l^m P_l^m(\cos\theta) e^{im\phi}, \tag{43}$$

where $N_l^m = \sqrt{\frac{(2l+1)(l-m)!}{4\pi(l+m)!}}$ is a normalization factor and $P_l^m(z)$ are the associated Legendre functions.

A spherical harmonic basis function rotated can be described as a linear combination of the spherical harmonic basis functions of the same polar order.

$$y_l^m(R_{\alpha,\beta,\gamma}(\theta',\phi')) = \sum_{n=-l}^{l} d_l^{mn}(\alpha,\beta,\gamma) y_l^n(\theta',\phi'), \tag{44}$$

where

$$d_l^{mn}(\alpha,\beta,\gamma) = P_l^{mn}(\cos\alpha) e^{im\beta+in\gamma} \tag{45}$$

and P_l^{mn} are the generalized associated Legendre polynomials. For their explicit form and more details on rotating spherical harmonics see [2].

Real spherical harmonics can be defined as follows.

$$Y_l^m(\theta,\phi) = N_l^m P_l^m \Phi_m(\phi), \tag{46}$$

where

$$\Phi_m(\beta) = \begin{cases} \sqrt{2}\cos m\beta & m > 0 \\ 1 & m = 0 \\ \sqrt{2}\sin m\beta & m < 0 \end{cases} \tag{47}$$

This can be written in a form more suitable for rotation

$$Y_l^m(\theta,\phi) = \begin{cases} \frac{1}{\sqrt{2}}\left(y_l^m(\theta,\phi) + (-1)^m y_l^{-m}(\theta,\phi)\right) & m > 0 \\ y_l^m(\theta,\phi) & m = 0 \\ \frac{1}{i\sqrt{2}}\left(y_l^m(\theta,\phi) - (-1)^m y_l^{-m}(\theta,\phi)\right) & m < 0 \end{cases} \tag{48}$$

Using (44) we can now rotate these functions. The result will be in terms of complex spherical harmonics. Converting these back to real spherical harmonics we get the D_l^{mn} functions for the real harmonics. Depending on whether m and n are positive, negative or zero we get nine cases. Only some of those are relevant to this paper ($\gamma = 0$ and $n \geq 0$). For those cases

$$D_l^{mn}(\alpha,\beta,0) = \frac{P_l^{mn}(\cos\alpha) + (-1)^n P_l^{m,-n}(\cos\alpha)}{\sqrt{2(1+\delta_{n0})}} \Phi_m(\beta). \tag{49}$$

$\Phi_m(\beta)$ is the same as (47).

Acknowledgements

This work was done within the EU-IST project IST-2000-29688 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] G. S. Chirikjian and A. B. Kyatkin. *Engineering Applications of Noncommutative Harmonic Analysis*. CRC Press, 2001.

[3] K.J. Dana, B. van Ginneken, S.K. Nayar, and J.J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, January 1999.

[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. Computer Vision and Pattern Recognition*, pages 995–999, 1994.

[7] B.K.P. Horn. *Robot Vision*. McGraw-Hill, 1986.

[8] R. A. Johnson and D. W. Wichern. *Applied Multivariate Statistical Analysis, 4:th ed.* Prentice-Hall, 1998.

[9] J.J. Koenderink and A.J. van Doorn. Phenomenological description of bidirectional surface reflection. *J. Optical Soc. of Am. A*, 15(11):2903–2912, November 1998.

[10] P. Nillius and J.O. Eklundh. Low-dimensional representations of shaded surfaces under varying illumination. In *Proc. Computer Vision and Pattern Recognition*, pages II:185–192, 2003.

[11] 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):1322–1333, October 2002.

[12] R. Ramamoorthi and P. Hanrahan. On the relationship between radiance and irradiance: determining the illumination from images of a convex lambertian object. *J. Optical Soc. of Am. A*, 18(10):2448–2458, October 2001.

[13] R. Ramamoorthi and Hanrahan P. Frequency space environment map rendering. In *SIGGRAPH*, 2002.

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE