Efficient Linear Re-rendering for Interactive Lighting Design

Patrick C. Teo
Dept. of Computer Science
Stanford University
Stanford, CA 94305
teo@cs.stanford.edu

Eero P. Simoncelli
Center for Neural Science, and
Courant Inst. of Mathematical Sciences
New York University
New York, NY 10003-6603
eero@cns.nyu.edu

David J. Heeger
Dept. of Psychology
Stanford University
Stanford, CA 94305
heeger@white.stanford.edu

Abstract

We present a framework for interactive lighting design based on linear re-rendering. The rendering operation is linear with respect to light sources, assuming a fixed scene and camera geometry. This linearity means that a scene may be interactively re-rendered via linear combination of a set of basis images, each rendered under a particular *basis light*. We focus on choosing and designing a suitable set of basis lights. We provide examples of bases that allow 1) interactive adjustment of a spotlight direction, 2) interactive adjustment of the position of an area light, and 3) a combination in which light sources are adjusted in both position and direction. We discuss a method for reducing the size of the basis using principal components analysis in the image domain.

1 Introduction

Lighting design plays a major role in theatrical production planning, interior design, and computer graphics animation. This paper describes a system for interactive lighting design. Specifically, the system allows interactive specification of both the positions and radiance distributions of the lights, for a fixed scene geometry and viewing direction.

Numerous advances in graphical rendering techniques have enabled efficient generation of photo-realistic images from model scenes. Recently, several incremental techniques have been proposed that can efficiently re-render a model scene when some aspect of its description is modified [1, 2, 4, 10, 19, 21]. Nevertheless, accurate model scenes are often tremendously complicated, and in such cases these techniques are not fast enough to permit the interactive modification of lighting specifications.

A recently developed alternative approach is to "re-render" the images as linear combinations of a fixed set of previously rendered *basis images* [8, 17, 6, 7]. The validity of this approach rests on a fundamental property of graphical rendering: linearity with respect to light source intensities [3, 12]. Specifically, rendering obeys the principle of superposition: (1) multiplying the intensity of the light source by an arbitrary factor scales the intensities in the rendered image by the same factor, and (2) an image rendered under two light sources is the sum of the two images rendered under each light source independently[1].

The linearity of the rendering operation leads to an efficient method for re-rendering scenes. Given two images of a scene rendered under two different light sources, one need not run a time-consuming rendering program to view the scene with both lights turned on. Instead, one can simply add the two images! More generally, the image of a scene illuminated by a weighted linear combination of *basis lights* may be computed via a linear combination of *basis images*, where each basis image is a rendering of the scene under a corresponding basis light. The creation of the basis images requires a full rendering operation, and is typically quite time-consuming. After this work is done, however, the linear re-rendering procedure is highly efficient, thus allowing interactive manipulation of lighting.

This simple approach to lighting design has several advantages. First, the method depends only on the linearity of the rendering operation and thus is quite general. Linearity holds for scenes with arbitrarily complex geometry and bi-directional surface reflectances, including those containing shadows and complex inter-reflections. Second, the computational

[1] We typically assume incoherent light sources, although the statement also holds for coherent lights sources if one retains both phase and amplitude information.

and storage requirements depend only on the number of basis images and the size (number of pixels) of each image. In particular, the method does not require computation or storage of visibility or other auxiliary information. Third, linear systems theory can be applied in various ways, as we shall show, to design convenient sets of basis lights and to reduce the number of basis images.

The usefulness and flexibility of a lighting design system based on linear re-rendering depends on both the number and the radiance distributions of the light sources. The computational cost of the re-rendering (and pre-rendering) operations is directly proportional to the number of basis light sources. It is thus important that this number be kept as low as possible. The choice of basis lights will determine the lighting design workspace. The re-rendering operation can only produce images corresponding to light lying in the linear subspace spanned by the basis lights. In previous work, basis lights have been chosen to span a rotation-invariant subspace of light sources. In particular, Nimeroff *et al.* [17] use a set of "steerable" area lights on a hemisphere, designed to approximate the illumination effects of daylight. The steerability property allows the representation of a continuum of sun positions. Dobashi *et al.* [6] use basis lights spanning the space of directional spot light sources, each positioned at the same location but aimed in different directions. A spherical harmonic decomposition was used to ensure rotation-invariance.

In this paper, we extend the re-rendering approach to lighting design. We elucidate a design methodology for: (1) directional spot lights whose directions of foci and angular radiance distributions can be continuously varied, (2) area and volumetric light sources whose positions and spatial radiance distributions can be continuously varied, and (3) light sources that are a combination of the first two types (i.e., directional area or volumetric lights whose directions of foci and positions can all be changed). In each of these cases (particularly the last), the size of the basis is a concern. If there are too many basis images, then the lighting design can no longer be interactive. To ameliorate these difficulties, we describe a method for reducing the number of basis images significantly via principal components analysis.

2 Steerability of Light Sources

Any light source can be fully described by the spatial and angular distribution of its emitted radiance. We write this function as $L(\boldsymbol{x}, \boldsymbol{\omega})$ where $\boldsymbol{x}$ specifies the position and $\boldsymbol{\omega}$ specifies the angular direction. For example, an isotropic point light source centered at location $\boldsymbol{x}_\circ$ would be described by the distribution $L_{\text{point}}(\boldsymbol{x}, \boldsymbol{\omega}) = \boldsymbol{\delta}(\boldsymbol{x} - \boldsymbol{x}_\circ)$ where $\boldsymbol{\delta}(\boldsymbol{x}') = 1$ when $\boldsymbol{x}' = \mathbf{o}$ and zero otherwise. Throughout this paper, the light sources are continuous functions of $\boldsymbol{x}$ and $\boldsymbol{\omega}$.

In lighting design, each light source is typically parameterized by a set of parameters which the designer continuously adjusts to achieve the desired visual effect. For example, these parameters could be the location of the light source or the direction of focus of a directional light source. A family of parameterized light source distributions is denoted by $\{L(\boldsymbol{x}, \boldsymbol{\omega}; \boldsymbol{p})\}$ where $\boldsymbol{p}$ is a vector of parameters.

A parameterized light source is said to be *steerable* in its parameters $\boldsymbol{p}$ if its radiance distribution can be written as a linear combination of a finite set of basis lights, where the weights involved in the linear combination are functions solely of the parameter vector $\boldsymbol{p}$. Mathematically, $L(\boldsymbol{x}, \boldsymbol{\omega}; \boldsymbol{p})$ is steerable if

$$L(\boldsymbol{x}, \boldsymbol{\omega}; \boldsymbol{p}) = \sum_{i=1}^{N} \alpha_i(\boldsymbol{p}) L_i(\boldsymbol{x}, \boldsymbol{\omega}) \tag{1}$$

where α_i are called the *weighting* or *steering functions* and L_i are the radiance distributions of the basis lights.

Steerable filters have been been developed primarily in the context of image processing and computer vision [14, 15, 9, 22, 18, 16, 11]. The term "steerable" refers to the particular case in which the basis lights each have the same shape (e.g., directional spot lights with identical radiance distributions aimed in different directions). But the more general definition above is useful for the purposes of this paper.

Note that the choice of basis light sources is not unique. Any other set of basis lights spanning the same space can also be used. Two such sets of basis lights are related by an invertible linear transformation. In practice, however, there may be reasons for choosing one set of basis lights over another. One useful choice of basis lights is obtained by sampling $L(\boldsymbol{x}, \boldsymbol{\omega}; \boldsymbol{p})$ with N values of $\boldsymbol{p}$; that is, the basis consists of the light sources, $\{L(\boldsymbol{x}, \boldsymbol{\omega}; \boldsymbol{p}_1), \cdots, L(\boldsymbol{x}, \boldsymbol{\omega}; \boldsymbol{p}_N)\}$. We will refer to this as a *sampled basis set*.

Let $\mathcal{R}$ denote the rendering operator for a fixed model scene from a given viewpoint. This operator takes as input the radiance distribution of the light source L and produces an image $I = \mathcal{R}(L(\boldsymbol{x}, \boldsymbol{w}))$. As explained in the *Introduction*, this operator is *linear*. Combining this notation with the steerability equation gives an expression describing the re-rendering process:

$$\begin{aligned}\mathcal{R}(L(\boldsymbol{x},\boldsymbol{\omega};\boldsymbol{p})) &= \mathcal{R}\left(\sum_{i=1}^{N}\alpha_i(\boldsymbol{p})L_i(\boldsymbol{x},\boldsymbol{\omega})\right)\\ &= \sum\alpha_i(\boldsymbol{p})\mathcal{R}(L_i(\boldsymbol{x},\boldsymbol{\omega}))\\ &= \sum\alpha_i(\boldsymbol{p})I_i,\end{aligned}\tag{2}$$

where $L_i(\boldsymbol{x},\boldsymbol{\omega})$ are the basis lights, I_i are the basis images, and $\alpha_i(\boldsymbol{p})$ are the weighting functions. That is, an image of the model scene with the new light source $L(\boldsymbol{x},\boldsymbol{\omega};\boldsymbol{p})$ may be re-rendered by linearly combining the basis images.

3 Steerable Directional Spot Lights

A directional spot light is a point light source whose emitted angular radiance distribution is anisotropic. Typically, spot lights are rotationally symmetric about their directions of foci; thus, their radiance distributions can be described as:

$$L_{\text{spot}}(\boldsymbol{x},\boldsymbol{\omega};\boldsymbol{p}) = \boldsymbol{\delta}(\boldsymbol{x}-\boldsymbol{x}_{\circ})\ f(\boldsymbol{\omega}\cdot\boldsymbol{p})\tag{3}$$

where $\boldsymbol{p}$ is a unit vector parameterizing the direction of focus of the spotlight, and $\boldsymbol{x}_{\circ}$ denotes the origin of the spot light.

For example, let f be the linear polynomial: $f(\boldsymbol{\omega}\cdot\boldsymbol{p}) = 1+\boldsymbol{\omega}\cdot\boldsymbol{p}$. Writing the dot-product explicitly gives

$$L_{\text{spot}}(\boldsymbol{x}_{\circ},\boldsymbol{\omega};\boldsymbol{p}) = 1 + p_x\omega_x + p_y\omega_y + p_z\omega_z$$

where the subscripts x, y, z refer to the corresponding components of the vectors $\boldsymbol{\omega},\boldsymbol{p}$. This parameterized light source is steerable, and can be written as a linear combination of four basis light sources:

$$L_{\text{spot}}(\boldsymbol{x}_{\circ},\boldsymbol{\omega};\boldsymbol{p}) = L_1 + p_x L_{\omega_x} + p_y L_{\omega_y} + p_z L_{\omega_z}$$

where L_{ω_x} corresponds to a basis light whose radiance distribution is $L_{\omega_x}(\boldsymbol{x}_{\circ},\boldsymbol{\omega}) = \omega_x$, etc. Because of the linearity of the rendering operation, the same relationship holds for the images:

$$I_{\boldsymbol{p}} = I_1 + p_x I_{\omega_x} + p_y I_{\omega_y} + p_z I_{\omega_z},$$

where $I_{\boldsymbol{p}}$ is the image rendered with $L_{\text{spot}}(\boldsymbol{x}_{\circ},\boldsymbol{\omega};\boldsymbol{p})$, I_{ω_x} corresponds to the basis image rendered using the basis light source L_{ω_x}, etc.

In the above example, the spot lights have rather broad radiance distributions. More generally, we construct a directional spot light of degree N as:

$$L_{\text{spot}}(\boldsymbol{x}_{\circ}, \boldsymbol{\omega}; \boldsymbol{p}) = (1 + \boldsymbol{\omega} \cdot \boldsymbol{p})^N. \tag{4}$$

Expanding this polynomial gives:

$$L_{\text{spot}}(\boldsymbol{x}_{\circ}, \boldsymbol{\omega}; \boldsymbol{p}) = \sum_{n=0}^{N} \sum_{\substack{i+j+k=n, \\ i,j,k \geq 0.}} \alpha^n_{i,j,k}(\boldsymbol{p})\ L_{\omega_x^i \omega_y^j \omega_z^k} \tag{5}$$

where the weighting coefficients are:

$$\alpha^n_{i,j,k}(\boldsymbol{p}) = \frac{N!}{(N-n)!(i)!(j)!(k)!} p_x^i p_y^j p_z^k.$$

Equation 5 appears to indicate that the total number of basis light sources is $\sum_{n=0}^{N} \sum_{i=0}^{n} \sum_{j=0}^{n-i} 1 = (N+1)(N+2)(N+3)/6$. Fortunately, the number of basis lights is much less, because the basis light source distributions $\omega_x^i \omega_y^j \omega_z^k$ are not linearly independent. The linear dependence is evident when you consider that $\omega_x, \omega_y, \omega_z$ are components of a unit vector, i.e., $\omega_x^2 + \omega_y^2 + \omega_z^2 \equiv 1$. The actual number of basis functions required is only $(N+1)^2$, the number of spherical harmonics up to degree N.

To take advantage of this, we need to choose a set of $(N+1)^2$ basis lights, and then derive the weighting functions for the new, reduced-size basis set. One approach is to use the spherical harmonic functions [6], but the weighting functions used to steer spherical harmonics are cumbersome. We use the more straightforward approach described in [17].

Instead of using the monomial basis light sources $L_{\omega_x^i \omega_y^j \omega_z^k}$, use a sampled basis set comprised of the desired spot light aimed in different directions. In particular, choose a set of unit vectors $\boldsymbol{p}_i$, for $1 \leq i \leq (N+1)^2$, distributed on the sphere. Construct the corresponding set of basis lights: $L(\boldsymbol{x}_{\circ}, \boldsymbol{\omega}; \boldsymbol{p}_i) = (1 + \boldsymbol{\omega} \cdot \boldsymbol{p}_i)^N$. Each new (directional spot) basis light can be expressed as a linear combination of the monomial basis lights, given by Equation 5. This sampled basis set can be related to the original monomial basis set via a linear transform:

$$\boldsymbol{M} = \begin{bmatrix} \cdots \alpha^n_{i,j,k}(\boldsymbol{p}_1) \cdots \\ \cdots \alpha^n_{i,j,k}(\boldsymbol{p}_2) \cdots \\ \vdots \\ \cdots \alpha^n_{i,j,k}(\boldsymbol{p}_{(N+1)^2}) \cdots \end{bmatrix}.$$

There are $(N+1)^2$ rows corresponding to the number of new basis lights, and there are $(N+1)(N+2)(N+3)/6$ columns in $\boldsymbol{M}$ corresponding to the total number of monomial basis lights.

To steer the new basis lights, we need to invert $\boldsymbol{M}$. Since $\boldsymbol{M}$ is under-determined, we compute its pseudo-inverse $(\boldsymbol{M}^T\boldsymbol{M})^{-1}\boldsymbol{M}^T$ using the singular value decomposition. Then, we write the weighting functions $\alpha_i'(p)$ of the new, reduced-size, sampled set of basis lights:

$$(\alpha_1'(\boldsymbol{p}) \cdots \alpha_{(N+1)^2}'(\boldsymbol{p})) = [(\cdots \alpha_{i,j,k}^n(\boldsymbol{p}) \cdots)\ (\boldsymbol{M}^T\boldsymbol{M})^{-1}\boldsymbol{M}^T]$$

Since the samples $\boldsymbol{p}_i$ may be chosen arbitrarily, one must check that the matrix $\boldsymbol{M}$ is full rank. This can be verified when computing its pseudo-inverse. If the rank is less than $(N+1)^2$, one can perturb the vectors $\boldsymbol{p}_i$ and try again.

Figure 1 shows images of a model scene illuminated by spot lights of degree $N = 5$. Note that the reflection of the wall in the sphere is brighter when that wall is illuminated by the spot light; linear re-rendering captures all ray interactions. Figure 2 shows a model of a chemistry set illuminated by a spot light of degree $N = 3$. Note the illumination of the test-tube rack that can be seen through the flask. Again, this would not be possible without capturing all ray interactions.

Spot lights with narrower radiance distributions may be obtained by using larger N. However, the higher the degree N, the more basis light sources are needed. Greater efficiency can be realized by using principal components analysis to reduce the number of basis images (see below).

Aside from reducing the number of basis lights/images, there are several practical advantages in using a sampled basis instead of the monomial basis. First, the radiance distributions of some of the monomial basis lights are negative in some directions. Each of the sampled basis lights, on the other hand, corresponds to an actual physical (non-negative) light source, so standard ray-tracers or radiosity programs can be used without modification. Second, for any given direction of focus, the contribution of each monomial basis light is small and about the same as that of the other basis lights. On the other hand, with the sampled basis set, only a few basis lights (typically those aimed near the desired direction of focus) contribute significantly to the linear combination. Hence, the re-rendering operation can be made more efficient by neglecting basis images with insignificant contributions. For the same reasons, a progressive refinement algorithm (see below) produces a high quality image with only the first few basis images.

The previous example illustrates only one family of steerable directional spot lights. There are several ways to generalize the results. The simplest extension is to allow for a general N^{th} degree polynomial; i.e., $L_{\text{spot}}(\boldsymbol{x}_\circ, \boldsymbol{\omega}; \boldsymbol{p}) = \sum_{p=0}^{N} c_p \, (\boldsymbol{\omega} \cdot \boldsymbol{p})^p$ where c_p are arbitrary coefficients. Naively expanding the above polynomial will indicate that $(N+1)(N+2)(N+3)/6$ basis light sources are needed. However, since $\boldsymbol{\omega}$ is a unit vector, only $(N+1)^2$ basis light sources are actually required and the above method of reducing the number of basis light sources can again be used. If the polynomial is a strictly even or odd function, then only $(n+1)(n+2)/2$ basis light sources are required as shown in [9] for steerable filters.

In addition, the approach can be generalized to include non-axis-symmetric radiance distributions. Using the spherical harmonics as the basis set, any distribution on the sphere can be steered following an arbitrary rotation of the coordinate axis of the sphere. Thus, in addition to being able to change the direction of some given axis, one can also rotate the distribution about that axis.

The techniques described in this section to steer directional spot lights can also be applied directly to steer skylight distributions as in [17]. The function that is being steered is identical in the two instances. The difference is that for directional spot lights, light energy emanates from a fixed point while for skylight distributions, light energy from all directions converges, with the same angular distribution, onto each point in the model scene.

Finally, an arbitrary desired angular radiance distribution can be approximated by a steerable basis set, so that it can be steered. If the basis lights are orthogonal, then the least-squares approximation can be computed by projecting the desired light onto the basis lights. Determining the best approximation is more difficult when the basis lights are not orthogonal (like the sampled set of basis lights discussed above). Dobashi *et al.* [6] used Legendre polynomials, which are orthogonal, in conjunction with spherical harmonics to approximate a desired light source. Alternatively, one could sample the radiance distributions of the desired light and the basis lights, then orthogonalize the basis lights numerically (e.g., using Gram-Schmidt) and compute the least-squares projection by matrix multiplication.

4 Steerable Area Lights

In the previous section, we presented techniques for steering the direction of spot lights. In this section, we explain how to shift the position of spatially distributed light sources. We concentrate on two-dimensional area light sources since the extension to three-dimensional volumetric light sources is straightforward.

The radiance distribution of an area light source is a four-dimensional function: two dimensions specify the angular distribution and two dimensions specify the spatial distribution. For the purpose of steering over position, we assume that the function is separable in its angular and spatial dimensions:

$$L_{\text{area}}(\boldsymbol{x}, \boldsymbol{\omega}; \boldsymbol{p}) = f_x(\boldsymbol{x} - \boldsymbol{p})\, f_\omega(\boldsymbol{\omega}) \tag{6}$$

where $\boldsymbol{p}$ is now a two-dimensional vector parameterizing the *position* of the light source. For example, when the area light source is defined over a plane, $\boldsymbol{x}$ is the two-dimensional coordinates on the plane, and $\boldsymbol{p}$ parameterizes the origin of the coordinate system. Since the function f_ω is not involved in the steering, it can be arbitrarily complex. For simplicity of presentation, we will assume that it is unity.

As with directional spot lights, only area light sources with certain spatial distributions f_x can be steered. Hel-Or *et al.* [11] identified these functions to be the product of (possibly complex) exponentials and polynomials; i.e., functions of the form $\{e^{\alpha x_u + \beta x_v} x_u^i x_v^j\}$ where $0 \le i \le n, 0 \le j \le m$ and α, β are complex constants. The variables x_u, x_v are components of the two-vector $\boldsymbol{x}$. In particular, consider the case when $n = m = 0$ and α, β are purely imaginary. For this case,

$$e^{\alpha(x_u - p_u) + \beta(x_v - p_v)} = \underbrace{e^{-(\alpha p_u + \beta p_v)}}_{\text{weighting functions}} \; \underbrace{e^{\alpha x_u + \alpha x_v}}_{\text{basis lights}}.$$

or in its more familiar real form,

$$\sin a(x_u - p_u) \sin b(x_v - p_v) = \underbrace{\begin{pmatrix} \cos a p_u \cos b p_v \\ -\cos a p_u \sin b p_v \\ -\sin a p_u \cos b p_v \\ \sin a p_u \sin b p_v \end{pmatrix}^T}_{\text{weighting functions}} \underbrace{\begin{pmatrix} \sin a x_u \sin b x_v \\ \sin a x_u \cos b x_v \\ \cos a x_u \sin b x_v \\ \cos a x_u \cos b x_v \end{pmatrix}}_{\text{basis lights}} \tag{7}$$

where $\alpha = ai$ and $\beta = bi$. Since sinusoids (and cosinusoids) do not have compact support, it may seem that one would need to have infinitely-extended light sources. Fortunately, since sinusoids (and cosinusoids) are periodic, we can truncate them at integral periods and still be able to steer them perfectly.

Therefore, a light source with an arbitrary spatial distribution (as in Equation 6) can be steered by first approximating it with these sinusoids and then steering these sinusoids.

Mathematically,

$$L_{\text{area}}(\boldsymbol{x}, \boldsymbol{\omega}; \mathbf{o}) \approx \sum_{i,j} \begin{pmatrix} c_{i,j}^0 \\ c_{i,j}^1 \\ c_{i,j}^2 \\ c_{i,j}^3 \end{pmatrix}^T \begin{pmatrix} \sin 2\pi i\ x_u \sin 2\pi j\ x_v \\ \sin 2\pi i\ x_u \cos 2\pi j\ x_v \\ \cos 2\pi i\ x_u \sin 2\pi j\ x_v \\ \cos 2\pi i\ x_u \cos 2\pi j\ x_v \end{pmatrix}$$

where $c_{i,j}^k$ are the coefficients derived from the approximation. Although the desired light source is approximated, for model scenes that do not contain strongly specular surfaces, the error in the re-rendered images is small. Intuitively, this is because the high frequency errors introduced by the approximation are averaged over the hemisphere of incoming light directions.

Since each of the sinusoids on the right of the above equation is steerable, the light source distribution is steerable:

$$L_{\text{area}}(\boldsymbol{x}, \boldsymbol{\omega}; \boldsymbol{p}) \approx \sum_{i,j} \begin{pmatrix} c_{i,j}^0 \\ c_{i,j}^1 \\ c_{i,j}^2 \\ c_{i,j}^3 \end{pmatrix}^T \left[\underbrace{\boldsymbol{A}(\boldsymbol{p})}_{\text{weighting functions}} \underbrace{\begin{pmatrix} \sin 2\pi i\ x_u \sin 2\pi j\ x_v \\ \sin 2\pi i\ x_u \cos 2\pi j\ x_v \\ \cos 2\pi i\ x_u \sin 2\pi j\ x_v \\ \cos 2\pi i\ x_u \cos 2\pi j\ x_v \end{pmatrix}}_{\text{basis lights}} \right]$$

where $\boldsymbol{A}$ is a 4×4 matrix with each row containing the weighting functions for each sinusoid, as specified in Equation 7.

Figure 3 shows a chess piece illuminated by an area light source with a raised-cosine spatial radiance distribution (in each dimension). The raised cosine was first approximated by a sum of sinusoids. Instead of using basis lights made up of sinusoids, however, we used a sampled basis set in which the basis lights were shifted copies of one another.

Figure 5 shows an example of steering both in position and direction. We used lights that were separable in their angular and spatial dimensions (Equation 6); the spatial distributions were raised cosines, and the angular distributions were polynomials of degree $N = 3$. Because the light distribution is separable, we first steered the full set of angular distributions to the desired position. Then we steered them in direction. Note that the shadow is unchanged when the direction of the light is steered (left and middle images), but the shadow does change when the light's position is shifted (right image).

5 Reducing the Basis Set

The time required to re-render an image is proportional to the number of basis images. For interactive lighting design, the number of basis images must be small. Unfortunately, light sources with narrower angular or spatial distributions require a larger number of basis images. We can partially remedy this problem by taking advantage of the fact that the images in our basis set will have different degrees of importance. In particular, principal component analysis can be used to compute a reduced set of basis images best approximating the original set. Specifically, the first k principal components are the best (in least-squares sense) k basis images approximating the original basis set.

Due to the number of pixels in each basis image, it is infeasible to directly compute the principal components using techniques like the singular-value decomposition (SVD). Instead, we apply the SVD to a smaller matrix and then derive the principal components of the original basis images from this intermediate result. This method was used, for example, by Turk *et al.* [23] to compute "eigenfaces" for a face recognition system.

Let $\boldsymbol{M}$ be a matrix where each column corresponds to a single basis image that has been collapsed into a single long vector, i.e., if each basis image contains m pixels and there are a total of n basis images then $\boldsymbol{M}$ is an $m \times n$ matrix. If we could compute the SVD of $\boldsymbol{M}$, it would give us three new matrices, $\boldsymbol{M} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^T$, where $\boldsymbol{U}$ and $\boldsymbol{V}$ are orthonormal and $\boldsymbol{S}$ is $m \times n$ and diagonal. The columns of $\boldsymbol{U}$ are the desired principal components. Since we can not compute the SVD of $\boldsymbol{M}$ directly, we instead work with with $\boldsymbol{M}^T\boldsymbol{M}$, a symmetric $n \times n$ (small) matrix. The SVD of this matrix is: $\boldsymbol{M}^T\boldsymbol{M} = \boldsymbol{V}\boldsymbol{S}^2\boldsymbol{V}^T$. The singular values $\boldsymbol{S}$ are determined by taking the square root of $\boldsymbol{S}^2$. Having obtained $\boldsymbol{V}$ and $\boldsymbol{S}$, we compute $\boldsymbol{U} = \boldsymbol{M}\boldsymbol{V}\boldsymbol{S}^{-1}$. The columns of $\boldsymbol{U}$ corresponding to small singular values in $\boldsymbol{S}$ are then eliminated and the remaining k columns of $\boldsymbol{U}$ constitute the new reduced-size basis.

Determining how many principal components to use is an empirical question. In our current implementation, k principal components are used such that the k largest singular values sum to 90% of the total sum of all the singular values.

Figure 4 shows three renderings of the same image: the first rendered directly with a ray tracer, the next re-rendered using all 81 basis images, and the last re-rendered with only 20 principal components basis images. Only slight artifacts in the detail of the shadow and the texturing of the bishop are visible.

In fact, most of the images in this paper were computed with reduced basis sets, and the reduction in the number of basis images is significant. The images in Figure 1 were re-rendered with 12 principal component basis images, instead of the 36 original basis images. The images in Figure 3 were re-rendered with 20 principal components, instead of the 81 original basis images. For Figure 5, there were 400 original basis images, but only 50 principal components were used.

6 Progressive Refinement

Although all the basis images need to be combined to re-render the new image perfectly, some basis images contribute more than others. Hence, before re-rendering a new image, it is wise to first sort the weighting coefficients in descending order of magnitude, and sequentially add the weighted basis images in that order. After each basis image is added, the partial result can be displayed. It is important, however, to also consider the salience of each basis image. For example, the weight associated with a particular basis image may be large, but the basis image itself may contain only small intensity values. We currently use the average absolute pixel value to quantify salience, and sort the basis images based on the product of the weight times the salience of the basis image.

7 Discussion

One of the advantages of the linear re-rendering approach to lighting design is that it captures all ray interactions (multiple bounces). On the other hand, if only the first bounce (light-backwards from the eye) is needed, a *deep-buffer* method may be more efficient. A deep buffer stores auxilliary information for each pixel. For example, a model scene illuminated by a directional spot light at a fixed point can be rendered for any angular distribution if associated with each pixel is a bit indicating whether the light is visible to the surface under that pixel. This information has to be available at sub-pixel resolution in order to handle anti-aliasing properly. Although this method accurately accounts for only the first bounce of light, its complexity is independent of the angular distribution of the spot light. By contrast, re-rendering with steerable basis lights requires a large number of basis images when the angular distribution of the light source is not smooth or narrow.

This suggests a hybrid of the two schemes, in which the first bounce is computed via the deep buffer and the second and subsequent bounces are computed via steerable re-rendering. These two contributions can subsequently be added to produce the final image to capture all the ray interactions. If the radiance contribution of the second and subsequent ray bounces vary slowly over space, then fewer steerable basis images will be required. Re-rendering with steerable lights works best (i.e., fewer basis images are required) when the light distributions are smooth and the surface reflectances are diffuse. These are exactly the situations in which ray paths are expensive to compute.

Another hybrid scheme involves decomposing a desired light source distribution into the sum of two components: a smooth, steerable component and a narrow, compactly supported second component. Since re-rendering with steerable light sources is most efficient when the light source distribution is smooth, this method can be used only for the first (smooth) component of the light. The contribution of the second (narrow, compact) component can be ray-traced for each new light source position. Since it is compactly supported, only a small number of rays are needed. This approach might be particularly useful for re-rendering model scenes illuminated by skylight since the distribution of skylight comprises a widespread slowly varying component and a strong, narrow component (the sun). In spirit, this hybrid scheme is akin to Chen's multi-pass method for global illumination [5] where the most appropriate method is used to render each type of ray path.

Finally, we note how to combine our re-rendering scheme with goal-based rendering [13, 20]. The scheme in [20] automatically determines the intensity settings of a fixed set of lights such that the image of the model scene is similar to one "painted" by the designer. The scheme in [13] automatically adjusts the intensity and direction of foci of the lights, but it involves incremental rendering. Re-rendering via linear combinations of steerable lights allows one to change the *geometry* of the light sources. By using lights that are constrained to be steerable, goal-based linear re-rendering can allow for intensity and geometry changes in the light sources.

References

[1] J M Airey, J H Rohlf, and F P Brooks. Towards image realism with interactive update rates in complex virtual building environments. *Computer Graphics*, 24(2):41–50, 1990.

[2] C Buckalew and D Fussell. Illumination networks: fast realistic rendering with general reflectance functions. *Computer Graphics*, 23(3):89–98, 1989.

[3] I. W. Busbridge. *The Mathematics of Radiative Transfer.* Cambridge University Press, Bristol, 1960.

[4] S E Chen. Incremental radiosity: an extension of progressive radiosity to an interactive image synthesis system. *Computer Graphics*, 24(4):135–144, 1990.

[5] S E Chen, H E Rushmeier, G Miller, and D Turner. A progressive multi-pass method for global illumination. *Computer Graphics*, 25(4):165–174, 1991.

[6] Y Dobashi, K Kaneda, H Nakatani, and H Yamashita. A quick rendering method using basis functions for interactive lighting design. In *Eurographics '95*, pages 229–240, 1995.

[7] J Dorsey, J Arvo, and D Greenberg. Interactive design of complex time-dependent lighting. *IEEE Computer Graphics and Applications*, 15(2):26–36, 1995.

[8] Julie Dorsey, Francois X. Sillion, and Donald Greenberg. Design and simulation of opera lighting and projection effects. In *Computer Graphics (SIGGRAPH '91 Proceedings)*, volume 25, pages 41–50, July 1991.

[9] W T Freeman and E H Adelson. The design and use of steerable filters. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(9):891–906, 1991.

[10] D W George, F X Sillion, and D P Greenberg. Radiosity redistribution for dynamic environments. *IEEE Computer Graphics and Applications*, 10(4):26–34, 1990.

[11] Y Hel-Or and P Teo. A common framework for steerability, motion estimation and invariant feature detection. Technical Report STAN-CS-TN-96-28, Stanford University, 1996.

[12] James T. Kajiya. The rendering equation. In *Computer Graphics (SIGGRAPH '86 Proceedings)*, volume 20, pages 143–150, August 1986.

[13] J K Kawai, J S Painter, and M F Cohen. Radioptimization - goal based rendering. *Computer Graphics*, 27(2):147–154, 1993.

[14] H Knutsson and G H Granlund. Texture analysis using two-dimensional quadrature filters. *IEEE Comput. Soc. Workshop Comp. Architecture Patt. Anal. Image Database Mgmt.*, pages 388–397, 1983.

[15] J J Koenderink and A J van Doorn. Representation of local geometry in the visual system. *Biological Cybernetics*, 55:367–375, 1987.

[16] M Michaelis and G Sommer. A Lie group-approach to steerable filters. *Pattern Recognition Letters*, 16(11):1165–1174, November 1995.

[17] J Nimeroff, E Simoncelli, and J Dorsey. Efficient re-rendering of naturally illuminated environments. In *5th Eurographics Workshop on Rendering*, 1994.

[18] P Perona. Steerable-scalable kernels for edge detection and junction analysis. *Image and Vision Computing*, 10(10):663–672, 1992.

[19] C Puech, F Sillion, and C Vedel. Improving interaction with radiosity-based lighting simulation programs. *Computer Graphics*, 24(2):51–57, 1990.

[20] C Schoeneman, J Dorsey, B Smits, J Arvo, and D Greenberg. Painting with light. *Computer Graphics*, 27(2):143–146, 1993.

[21] C H Séquin and E K Smyrl. Parameterized ray tracing. *Computer Graphics*, 23(3):307–314, 1989.

[22] E Simoncelli, W Freeman, E Adelson, and D Heeger. Shiftable multiscale transforms. *IEEE Trans. on Information Theory*, 38(2):587–607, 1992.

[23] M Turk and A Pentland. Face recognition using eigenfaces. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 586–591, 1991.

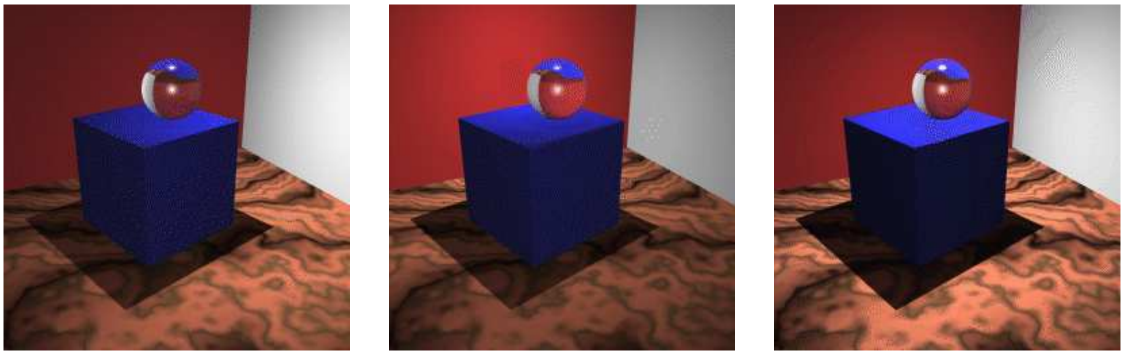

Figure 1: Images of a scene illuminated by a directional spot light. The angular radiance distribution of the light source is a fifth degree polynomial. Each of the images were re-rendered by linearly combining a set of 12 basis images. The left and middle images show the scene re-rendered with the spot light pointed in different directions. The right image shows the scene re-rendered with two spot lights in the same position, but pointing in different directions.

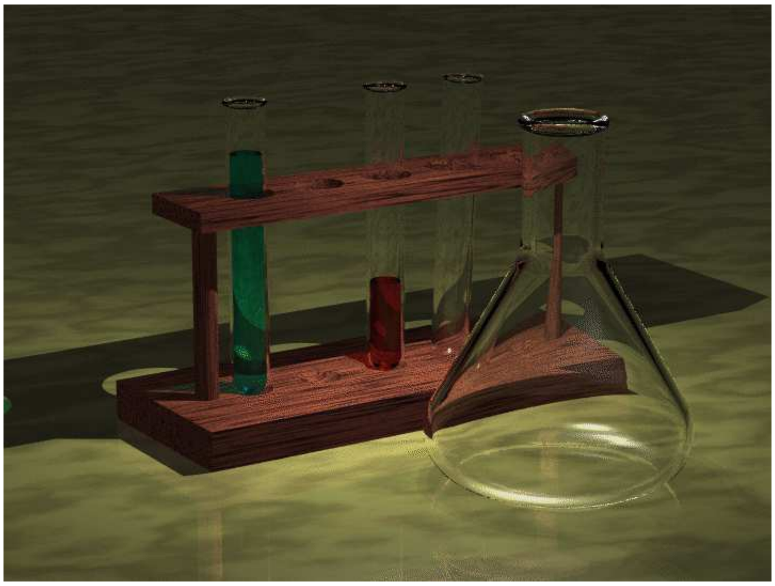

Figure 2: Image of a scene with transparent objects illuminated by a directional spot light. The angular distribution of the light source is a third-degree polynomial. A total of 16 basis images were used to re-render this scene.

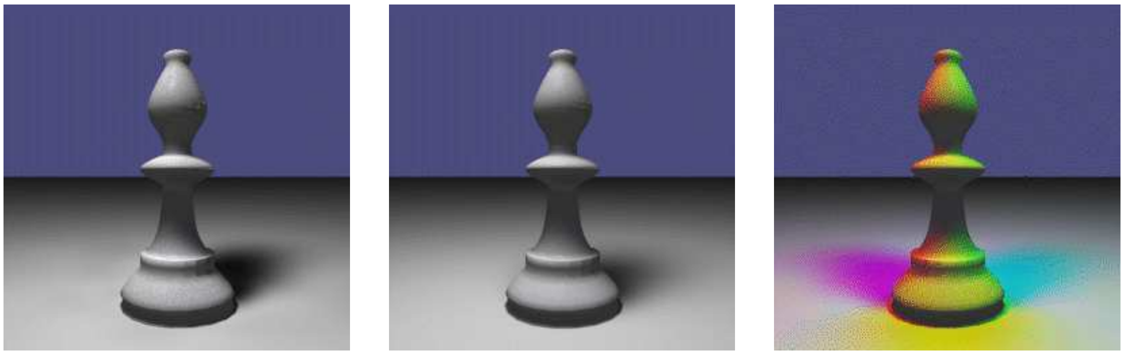

Figure 3: Images of a chess piece illuminated by an area light source. The left image shows the scene re-rendered with the area light source positioned to the front and left of the object. The middle image is a re-rendering of the scene with a broader area light source. The right image shows the object illuminated by three primary colored lights. A total of 20 basis images were used to re-render all three images.

Figure 4: Images of a chess piece illuminated by an area light source. The left image shows the scene rendered using a ray tracer. The middle image is a re-rendering using the full set of 81 basis images. The right image shows a re-rendering using a set of 20 basis images that were chosen with the singular value decomposition.

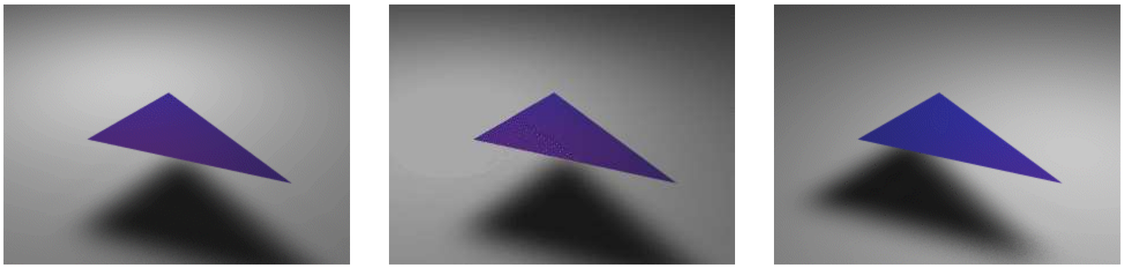

Figure 5: Images of a single polygon illuminated by an area light source that has an anisotropic angular distribution. The left image shows a re-rendering with the light source pointing downwards, and positioned to the rear and left of the object. The middle image shows a re-rendering with the light source in the same position as before but pointing in a different direction. The right image is a re-rendering with the light source centered at a different position. A total of 50 basis images were used to re-render all three images. These basis images were computed using the singular value decomposition; the actual number of basis images required was 400.