

Wavelet-based Light Reconstruction from a Single Image

Cameron Upright, Dana Cobzaş, Martin Jägersand
CS, University of Alberta, Canada
{upright,dana,jag}@cs.ualberta.ca

Abstract

In this paper we propose a wavelet-based light representation. This representation is used for recovering illumination and reflectance of a scene with known geometry from a single calibrated image. Previous approaches of light estimation reconstruct either a discrete set of infinite light sources or the projection of light on the spherical harmonics basis. The first approach is more suitable for modeling sharp light effects while the second one works best for diffuse scenes. In contrast, we show that the proposed representation is suitable for both diffuse and specular scenes. We compared our technique with the two previously mentioned approaches and found it superior. In addition to illumination estimation our algorithm also estimates a Phong surface reflection. Experiments with synthetic and real scenes demonstrate the effectiveness of our method.

1. Introduction

Light modeling is an important problem in computer vision. Most shape reconstruction algorithms (e.g; photometric stereo, shape from shading) make simplistic assumptions about the light conditions - usually a single infinite light source with known direction. These assumptions limit the applicability of the reconstruction techniques to carefully controlled laboratory setups. We propose a wavelet-based light representation and an image-based illumination recovery algorithm suitable for scenes with complex light and reflectance.

A variety of techniques have been proposed to recover the illumination from images. They can mainly be divided in two categories: some that recover a discrete collection of point light sources, and the others that recover light as projection on a basis defined on the surface of the illumination hemisphere (usually spherical harmonics basis). The methods from the first category are suitable for modeling sharp light effects, while the ones in the second category work for diffuse, Lambertian scenes. In contrast, our wavelet-based method can recover both diffuse and sharp light effects.

One of the first methods for reconstructing light uses a calibrated diffuse white sphere [15, 16, 2]. On the image of a smooth surface, the change in intensities is maximal when the illumination direction is perpendicular to the surface normal (critical points). The method first reconstructs critical points, that are then grouped together to form cut off curves. One light source direction is estimated inside each region formed by those curves. The approach has several disadvantages, mainly its sensitivity to noise. It also does not account for shadows.

One other approach to represent illumination is to use uniform samples on the illumination hemisphere. Each sample acts like an infinite light source. Hougen and Ahuja [4] integrate shape from shading, stereo and light estimation methods. The light is estimated by solving a linear system of equations for intensities of sampled light positions. An interesting approach is proposed by Sato et al. [11] where a discrete light hemisphere is estimated from shadows. The method works for Lambertian objects with uniform reflectance. For textured objects additional images are needed for capturing the albedo. Nishimo et al. [8] proposed a method that recovers light from specular objects. The method detects specular points on the objects and reflects the ray from the camera to those points in the mirror direction. They rays are intersected with a hemisphere covering the 3D model to generate the illumination hemisphere.

Some recent papers combine the idea of recovering a finite set of lights from critical points and the uniformly sampled light model. Wang and Samaras [14] combine estimation of critical points [16] and estimation of uniformly sampled lights from shadows [11]. Li et al. [5] present a general approach to light recovery suitable for textured objects with general reflectance. The approach combines multiple cues - occluding shadow edges, critical points and specularities. They use consistency among cues to distinguish light changes from texture edges.

Light can also be approximated using a set of basis functions defined on the surface of the illumination hemisphere. Different examples of basis functions have been used in the literature, the most popular one being the spherical harmonics basis due to the result that connects this representation

to rendering diffuse objects [10, 1] (for a Lambertian object, the first 9 coefficient encode $\sim 98\%$ of variability). Marschner and Greenberg [6] were some of the first ones to use the idea of a light basis. Given a photograph of a known object they produce a set of basis images by rendering the object with only one light source (sampled on the illumination hemisphere). The coefficients calculated from the decomposition of the original images on this basis are also the coefficients of the light with respect to a set of light basis sampled on the sphere (correspondent to the lights that were used in the sample renderings). Ramamoorthi and Hanrahan [9] use a spherical harmonics approximation of illumination captured using a light probe for efficient rendering. The light coefficients are used to linearly interpolate corresponding spherical harmonics calculated on the object shape. They showed that the radiance can be efficiently represented as a quadratic polynomial in the Cartesian components of the surface normal. Later they show [7] that representing the illumination hemisphere with coefficients of a wavelet basis gives a better approximation than spherical harmonics coefficients especially for sharp effects (specular highlights, sharp shadows). They use a non-linear approximation and keep only the largest terms in the basis. They did not reconstruct the light, but use the result for efficient rendering. Hara et al. [3] propose a method for simultaneously computing scene illumination and scene reflectance from a single image. The illumination is represented on the surface of a unit sphere using Mises-Fisher distributions by deviating a spherical reflectance model (spherical Torrance-Sparrow). Then they estimate the mixture and the number of distributions. The result is then refined using the original specular reflection model.

It is also worth noting the work of Pont and Koenderink [13, 12] on surface illuminance flow. They showed that the light vector can be decomposed into a scalar “normal illuminance” and a vector “surface illuminance”. The scalar generates the familiar shading on objects while the vector component is irrelevant for smooth surfaces but generates texture due to mesorelief. This component is conventionally ignored. They further show that the illumination flow direction (modulo 180°) can be inferred from illumination induced texture. The illumination flow is largely due to global light properties independent of local surface. However because there exist true screening in radiometry the local scene structure has an influence on the local light field. Main causes of those local effects are “vignetting” (local shadows) and multiple “scattering”.

We propose a method that represents light using a wavelet basis. Wavelets are a class of multi-scale basis vectors, best known from their applications in image compression. A very useful property is that they provide local support in both time domain (or spatial domain - here image dimension) and frequency domain. The application

of the basis in the context of light representation is therefore motivated by this property as we can efficiently represent both diffuse effects (limited in frequency domain) and sharp specular effects (limited in spatial domain). Using the wavelet representation we also present an algorithm for recovering light from one image of a known scene. One advantage of capturing the lighting on an object as opposed to introducing a light probe is that we not only recover the scene’s lighting, but also the reflectance parameters of the object. This allows us to render the scene under the same or different lighting conditions, as well as introducing synthetic objects into an augmented reality application. Another advantage of using an object instead of a light probe is that the lighting recovered minimizes the residual in the current scene, allowing a potentially higher quality rendering for moderate view changes. For the moment we have results for objects with uniform albedo (color) but the algorithm can be extended to textured objects if more than one image is available.

The remaining of this paper is organized as follows. We define the inverse light problem in the following section. Sections 3 and 4 present the proposed light representation and its use in the formulation of the rendering (reflectance) equation. Section 5 describes the light reconstruction algorithm. Section 7 presents experimental results with both synthetic and real image as well as a comparison between our light model, the spherical harmonics model and the discretized hemisphere model.

2. Problem Definition

This section presents the general light reconstruction problem and the assumptions we make. The light recovery or the inverse light problem can be stated as follows: Given a calibrated image I of a known uniformly colored object $S = \{\mathbf{X}_j, j = 1 \dots M\}$, with known reflectance, recover the light that illuminates the scene. We consider here the more general problem where the object reflectance is also unknown and it is thus estimated along with the light in an alternate fashion.

There are two aspects that need to be considered when solving the inverse light problem: one is the light representation and the second is the image formation (rendering) model. As described in the introduction, different light models have been considered most common being the discrete light samples and the basis type representation. In this paper we propose a wavelet-based light model that is described in Section 4. The rendering or image formation involves defining a reflectance model. Two reflectance models (Lambertian and specular Phong) are presented in Section 3.

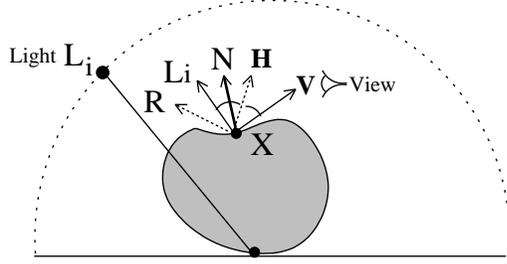


Figure 1. The vectors used in lighting calculations.

3. Image Formation and Reflectance Models

Image formation involves light, object shape and reflectance as well as a camera model. We assume that we have the correspondence between the object points and the image points $I(\mathbf{X})$ (in other words we assume that the camera is calibrated with respect to the object). For now, we consider the light represented with a set of discrete infinite point light sources $L_i, i = 1 \dots N$. We next present two reflectance model one suitable for diffuse objects (the Lambertian model) and the other for specular objects (the Phong model).

3.1. Lambertian Reflectance Model

The Lambertian reflectance model is the simplest model considered here. In a Lambertian lighting model, the recorded intensity of a point is independent of the viewing direction. In this case the BRDF is just a constant, known as the albedo. The intensity of a pixel corresponding to the projection of the object point \mathbf{X} on a uniformly colored object is given by

$$I(\mathbf{X}) = k_d V_X \sum_i^N L_i S_i \langle \mathbf{L}_i, \mathbf{N} \rangle \omega_i \quad (1)$$

where k_d is the object's albedo, and V_X denotes visibility of the object point with respect to the view direction. The discrete lights on the cube map are indexed by i . A binary value S_i represents the shadow and is equal to one when the i 'th light is visible from the point \mathbf{X} , and zero otherwise. The intensity, direction and solid angle of the i 'th discrete light is given by L_i, \mathbf{L}_i and ω_i respectively. From now on we will ignore the solid angle as it can be regarded like a scaling constant. The normal of the surface at point \mathbf{X} is given by \mathbf{N} (see Figure 1). $\langle \dots \rangle$ denotes the dot product.

3.2. Phong Reflectance Model

The Phong reflectance model is a simple general reflectance model. While not physically based, it works very

well in practice and it is one of the most popular reflectance models in the computer graphics community. The intensity of a pixel corresponding to \mathbf{X} illuminated by the Phong model is

$$I(\mathbf{X}) = V_X \sum_i^N L_i S_i (k_d \langle \mathbf{L}_i, \mathbf{N} \rangle + k_s \langle \mathbf{L}_i, \mathbf{R} \rangle^n) \omega_i \quad (2)$$

where k_s is the specular value, \mathbf{R} is the reflection of the viewing direction with respect to the normal, and n is the shininess of the material. A common variant of the Phong model is the Blinn-Phong model. The Blinn-Phong model replaces the $\langle \mathbf{L}_i, \mathbf{R} \rangle$ term with $\langle \mathbf{H}_i, \mathbf{N} \rangle$, where $\mathbf{H}_i = \frac{\mathbf{L}_i + \mathbf{V}}{\|\mathbf{L}_i + \mathbf{V}\|}$ is the bisector between light and view directions. The geometry behind the Lambertian and Blinn-Phong models can be seen in Figure 1.

4. Wavelet-based Light Model

Our method uses a cube map to represent the illumination hemisphere. We chose to use a cube map for its simplicity, and it is commonly used in computer graphics for representing environment maps. Each pixel on the cube map represents the direction of a light at infinity. The hemisphere is divided up into nine faces, one large on top and eight small ones on the sides (See Figure 2). We ignore the lower part of the cube map assuming that the image shading is mostly caused by light coming from above the object. A Haar wavelet basis is then defined over each of these faces. The intensity of a light is found by summing together all wavelets which influence that specific cube map pixel

$$L_i = \sum_{k=1}^m w_k b_{ki} \quad (3)$$

where L_i is the intensity of light i , b_{ki} is the value of the k 'th basis function at location i , and w_k is the corresponding weight for b_k . By substituting equation 3 into the Lambertian reflectance equation 1 we get :

$$I(\mathbf{X}) = k_d V_X \sum_i^N S_i \left(\sum_{k=1}^m w_k b_{ki} \right) \langle \mathbf{L}_i, \mathbf{N} \rangle \quad (4)$$

With a little rearranging, equation 4 becomes the following :

$$I(\mathbf{X}) = \sum_{k=1}^m w_k \left(k_d V_X \sum_i^N S_i b_{ki} \langle \mathbf{L}_i, \mathbf{N} \rangle \right) \quad (5)$$

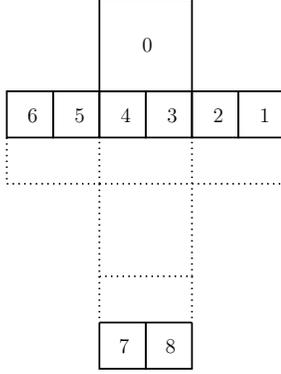


Figure 2. The nine faces of the illumination hemisphere

Similarly, equation 3 can be substituted into equation 2 giving us :

$$I(\mathbf{X}) = \sum_{k=1}^m w_k \left(V_X \sum_i^N S_i b_{ki} (k_d \langle \mathbf{L}_i, \mathbf{N} \rangle + k_s \langle \mathbf{H}_{ij}, \mathbf{N}_j \rangle^n) \right) \quad (6)$$

$$(7)$$

The above equations can be interpreted as a sum of “basis images” and therefore efficiently implemented on graphics hardware. Figure 3 and 5 show two example of such basis images, the first one for a scene illuminated with a single haar scaling function from above and the second one for a scene illuminated with a Haar wavelet.

When compressing an image with wavelets, a good approximation can be made by zeroing wavelets with small coefficients as a post processing step. In the case of light estimation, the areas of the cube map that need more detail are unknown. Therefore the proposed method uses a bottom up approach. The initial reconstruction uses a low resolution basis, which is then iteratively refined by adding higher order wavelet basis functions where needed (see Figure 9 for an example).

In terms of implementation, each face on the cube map is treated as a quad tree. At the root level is the Haar scaling function. Every node has three associated wavelet basis, which are the wavelets required to subdivide that node into four smaller ones. The scaling function and three wavelet functions can be seen in Figure 4.

5. Light Recovery Method

The proposed method for illumination reconstruction uses a single image. With a single image of a known Lambertian object, it is not possible to reconstruct both the lighting and the object albedo. Only their product can be deter-

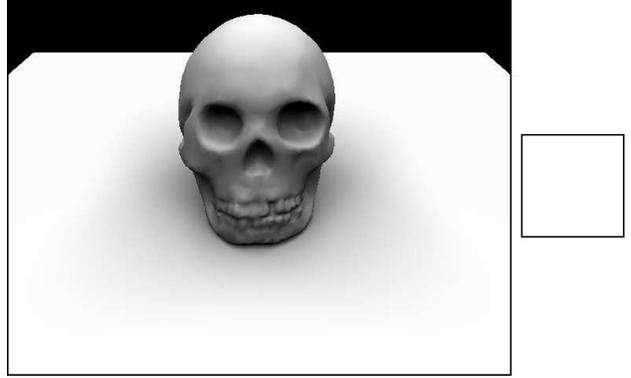


Figure 3. Left: Scene illuminated from above with Haar scaling function. Right: The Haar scaling function.

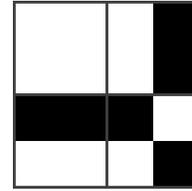


Figure 4. Left to Right, Top to Bottom: The Haar scaling function and three wavelet functions.

mined. As a result, this approach is limited to uniformly colored objects with known reflectance.

Since we are working with a known object, the visibility V_X and the shadow indicator S_i are both known. These are calculated by rendering the scene on graphics hardware, with a stencil buffer shadow algorithm.

When the reflectance properties of a Lambertian or Phong reflectance parameters are known, the problem of illumination reconstruction given the image I takes a simple linear form

$$I_j = I(\mathbf{X}_j) = \sum_{k=1}^m w_k a_{jk} \quad (8)$$

where a_{jk} has the following form in the Lambertian and specular case respectively:

$$a_{jk} = k_d V_j \sum_i^N S_{ij} b_{ki} \langle \mathbf{L}_i, \mathbf{N}_j \rangle \quad (9)$$

$$a_{jk} = V_j \sum_i^N S_{ij} b_{ki} (k_d \langle \mathbf{L}_i, \mathbf{N}_j \rangle + k_s \langle \mathbf{H}_{ij}, \mathbf{N}_j \rangle^n) \quad (10)$$

Non-negativity constraints are needed on the L_i variables to



Figure 5. Left: Scene illuminated from above with a Haar wavelet. Right: The Haar wavelet.

prevent negative lighting

$$\forall i \sum_{k=1}^m w_k b_{ki} \geq 0$$

A similar approach, but without the wavelet basis representation, was used by Sato Ikeuchi [11].

Illumination reconstruction of a uniformly colored object with a known BRDF is a linear least squares problem with linear inequality constraints, which can be formulated as a quadratic programming problem.

$$\min_{\mathbf{w}} (\mathbf{A}\mathbf{w} - \mathbf{I})^T (\mathbf{A}\mathbf{w} - \mathbf{I}) = \min_{\mathbf{w}} (\mathbf{w}^T \mathbf{A}^T \mathbf{A} \mathbf{w} - 2\mathbf{I}^T \mathbf{A} \mathbf{w}) \quad (11)$$

subject to

$$\sum_{k=1}^m w_k b_{ki} \geq 0 \quad (12)$$

where A is the light transport matrix, \mathbf{w} is the vector of basis function coefficients and \mathbf{I} is a vector of the observed pixel values.

Light reconstruction gets slightly more complicated when the specular parameters of the material are unknown. Instead of trying to tackle the whole non-linear optimization problem of finding lighting and parameters simultaneously, we use an iterative method.

If the lighting is known, estimating the specular parameters becomes a non-linear optimization problem. Knowing that:

$$I_j = V_j \sum_i^N L_i S_{ij} (k_d \langle \mathbf{L}_i, \mathbf{N}_j \rangle + k_s \langle \mathbf{H}_{ij}, \mathbf{N}_j \rangle^n) \quad (13)$$

and with the simplified notations

$$\begin{aligned} a_{ij} &= V_j L_i S_{ij} \langle \mathbf{L}_i, \mathbf{N}_j \rangle \\ b_{ij} &= V_j L_i S_{ij} \\ c_{ij} &= \langle \mathbf{H}_{ij}, \mathbf{N}_j \rangle \end{aligned}$$

$$I_j = k_d \sum_i^N a_{ij} + k_s \sum_i^N b_{ij} c_{ij}^n \quad (14)$$

the specular parameters are then estimated as:

$$\{k_s, n\} = \operatorname{argmin}_{k_s, n} \sum_{j=1}^M (I_j - k_d \sum_i^N a_{ij} - k_s \sum_i^N b_{ij} c_{ij}^n)^2 \quad (15)$$

We initially estimate the lighting by solving equations 11 and 12. We start an educated guess of the true specular parameters, which gives us an approximate reconstruction. This solution is then used to estimate the specular parameters by solving equation 15 using the Levenberg-Marquardt algorithm. Given the new specular parameters, the lighting is solved one last time.

6. System and Implementation Details

Input

calibrated image I
 object geometry $\mathbf{X}_j, j=1..M$
 (spec) initial specular params k_s, n
 initial wavelet resolution

Repeat

a fixed number of times
 estimate light coefficients \mathbf{w} (Eq. 11 and 12)
 refine wavelet basis in areas with non-zero light
 (spec) estimate reflectance params k_s, n (Eq. 15)
 estimate lighting one final time.

Figure 6. Overview of light reconstruction method

The input to the light and reflectance estimation method is a calibrated image of a known Lambertian or specular object with uniform reflectance. In the case of the Lambertian scene, the object is assumed white (albedo 1) while for the specular case, we initially guess the specular parameters. Our method starts with a low resolution cube map (only the top hemisphere), shown in Figure 2. The light is reconstructed using a wavelet basis at this resolution by solving the quadratic programming problem from equations 11 and 12. Areas of the cube map with non-zero pixel light values are refined further to the next level of resolution. Nodes which don't influence the reconstructed cube map are removed (all the coefficients are zero). This is then repeated a fixed number of times until the subdivision reaches the finest resolution. In the specular case, the reflectance

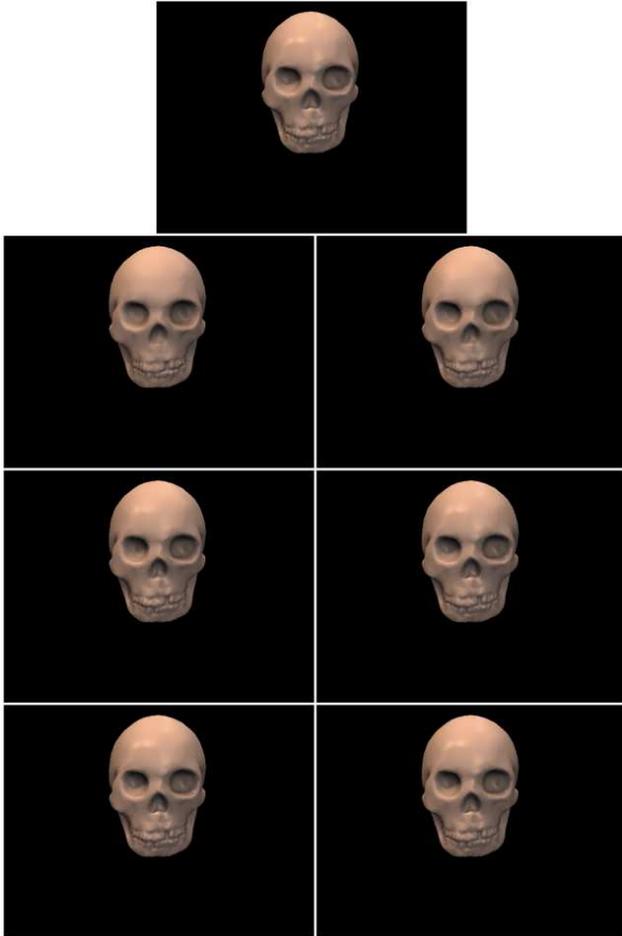


Figure 7. The original synthetic input scene is shown on top, followed by the six increasingly accurate reconstructions.

parameters are now re-estimated using equation 15, and the lighting is then estimated one last time. Figure 6 presents a summary of the algorithm.

7. Experimental Results

In order to determine the effectiveness of our algorithm, we performed two sets of tests. The first experiment demonstrates the proposed wavelet-based light and reflectance reconstruction on both real and synthetic data. In the second test, we compared the wavelet lighting basis with other alternative basis functions.

7.1. Wavelet Based Reconstruction

In this experiment we tested the performance of our wavelet based method on both synthetic and real scenes.

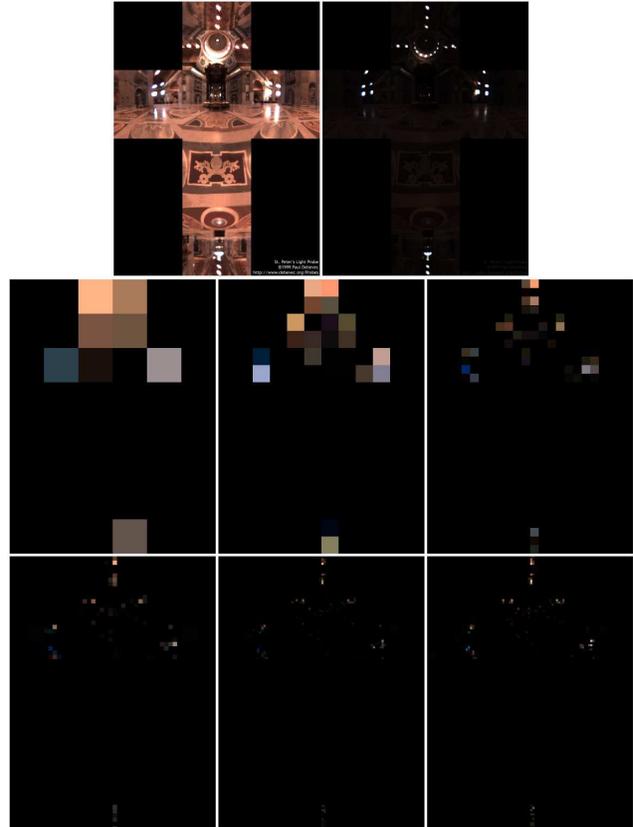


Figure 8. At the top is the original lightmap used to create the synthetic scene, followed by a scaled down version to show bright spots. The remaining 6 images are consecutive lighting reconstructions.

For our synthetic scene we used the skull model from Headus, and the environmental map was courtesy of Debevec. As outlined in Section 6 the proposed light reconstruction method starts with a low resolution cube map that is then refined until the desired reconstruction error is achieved. Figure 8 presents the original light map and six iterations of the lighting refinement. Figure 9 shows the corresponding cubemap subdivisions. Figure 7 presents the original scene and six relit images correspondent to the reconstructed lights from Figure 8.

We performed a similar experiment for a diffuse object and a real scene. Both the synthetic and real experiments were performed with a cubemap face resolution of 32x32. For the real scene, we used an image of a glossy ceramic duck. We obtained the geometric model of the duck using a laser scanner and we registered it with the image using corresponding points (manually selected). The image was calibrated with respect to the camera using the dotted calibration pattern shown in Figure 10 (first image last row).

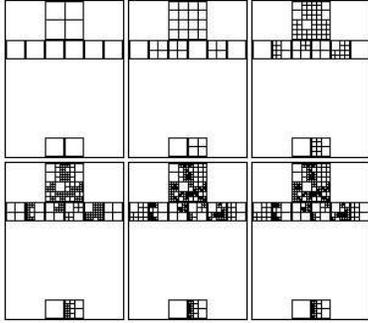


Figure 9. The six successive subdivisions of the reconstructed cubemap in Figure 8.

Comparative results of the original image, the reconstructed scene and the final reconstructed light are shown in Figure 10.

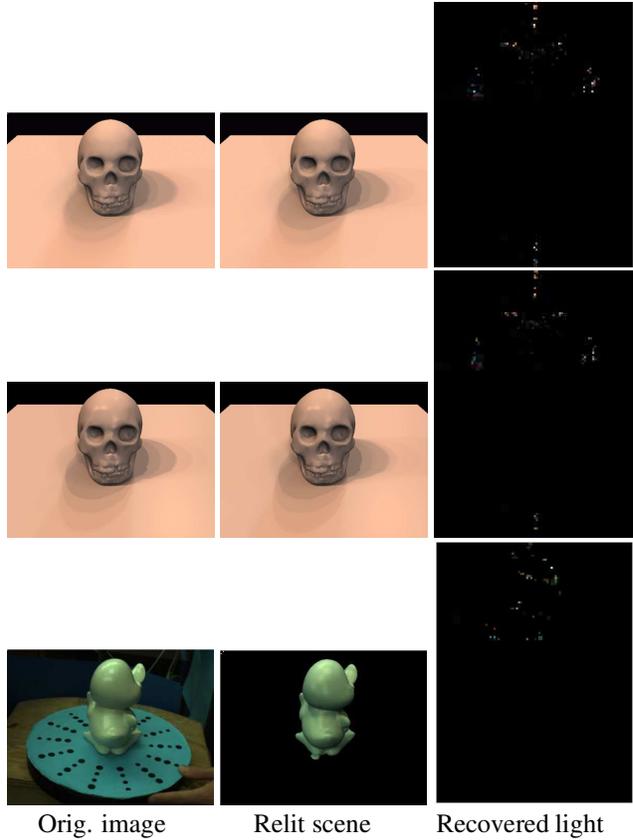
The synthetic scene reconstruction was nearly indistinguishable from the input reference image, while the real scene reconstruction was reasonably close. If one looks carefully at the light map and reconstructed scenes from the shiny duck test, one can notice the three fluorescent lights above the duck and some anomalous blue lights on the horizon. These lights are there to make up for the reflection of the blue calibration pattern, which gives the bottom of the duck a blue appearance.

7.2. Comparative Lighting Experiments

The second test we performed was a comparative experiment, between the proposed wavelet basis method and three other lighting bases: The first basis is a quad tree based lighting basis which uses a subdivision scheme identical to the wavelet based one. Instead of having a wavelet basis over the cubemap, we represent each leaf of the quad tree with its own lighting coefficient. The second basis we implemented used the first 9 spherical harmonics of a global basis as advocated in [10]. Finally, we implemented a simple cubemap where each pixel had its own lighting coefficient (at the finer resolution).

These four methods were applied to the real world scene. Figure 11 shows the original input image, and the four reconstructed scenes, while figure 12 shows the corresponding lightmaps. In order to make a fair comparison, all four methods were run with an 8x8 cubemap.

The final reconstructions for the wavelet, quad tree and simple cubemap lighting bases were quite similar. The set of possible lightmaps are the same for these three methods, so it is understandable that they get such similar results. The spherical harmonic reconstruction was of lesser quality than the other three, simply because sharp lighting effects can't



Orig. image Relit scene Recovered light

Figure 10. Results of the wavelet based light reconstruction algorithm. The first column shows the original image, the middle column shows the scene reconstructed with the final light and the last column shows the final reconstructed light. The first two rows present results for the synthetic experiment (diffuse and specular) and the last row presents the results for the real scene.

be modeled with only nine spherical harmonics. It is possible to get a better reconstruction by using higher order harmonics, but this may produce Gibbs phenomena in the reconstructed lightmap.

It should be noted that although the simple cubemap based method gave results very similar to the wavelet and quad tree methods, it does not scale well. At a higher cubemap resolution of 16x16, the constrained linear least squares solver failed due to poor conditioning.

8. Conclusions

We have presented a light and reflectance reconstruction method that uses a wavelet representation for the light implemented on a cube map. The reflectance is represented using a Phong parametric model. The input to our method is

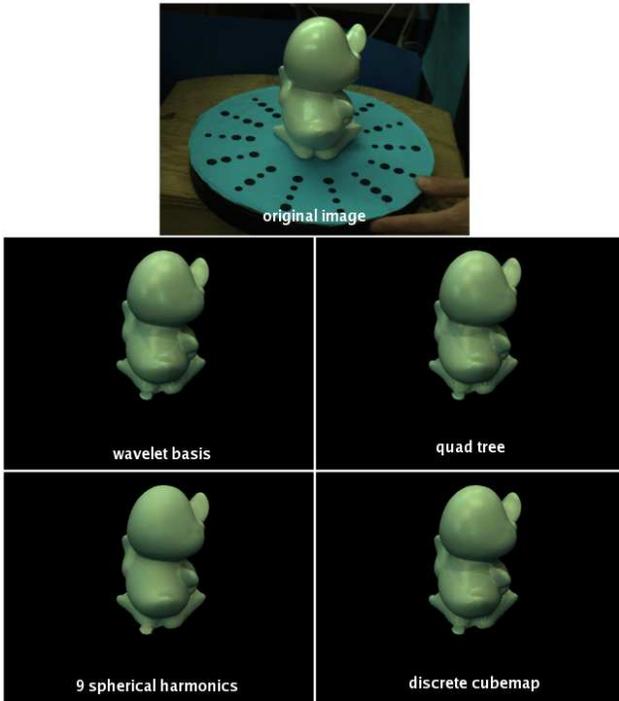


Figure 11. The input scene is shown on the top row. The second row has the wavelet and quad tree based reconstructions. The final row contains the spherical harmonic and simple discrete cubemap results.

a single calibrated image of a known scene. The reconstruction starts with a low resolution cube map that is then refined a fixed number of times. We demonstrate good results for both diffuse and specular objects on real and synthetic scenes. We compared our representation with other popular light models (spherical harmonics and discrete hemisphere) and found it superior. Furthermore, it can represent complex lightmaps in a both accurate and compact way.

References

[1] R. Basri and Jacobs. Lambertian reflectance and linear subspaces. *PAMI*, 25(2), 2003.
 [2] C. Bouganis and M. Brookes. Multiple light source detection. *PAMI*, 26(4):509–514, April 2004.
 [3] K. Hara, K. Nishino, and K. Ikeuchi. Multiple light sources and reflectance property estimation based on a mixture of spherical distributions. In *ICCV05*, pages II: 1627–1634, 2005.
 [4] D. Houghen and N. Ahuja. Estimation of the light source distribution and its use in integrated shape recovery from stereo and shading. In *ICCV*, pages 29–34, 1993.
 [5] Y. Li, S. Lin, H. Lu, and H.-Y. Shum. Multiple-cue illumination estimation in textured scenes. In *ICCV '03: Proceedings of the Ninth IEEE International Conference on Computer Vision*, page 1366, 2003.

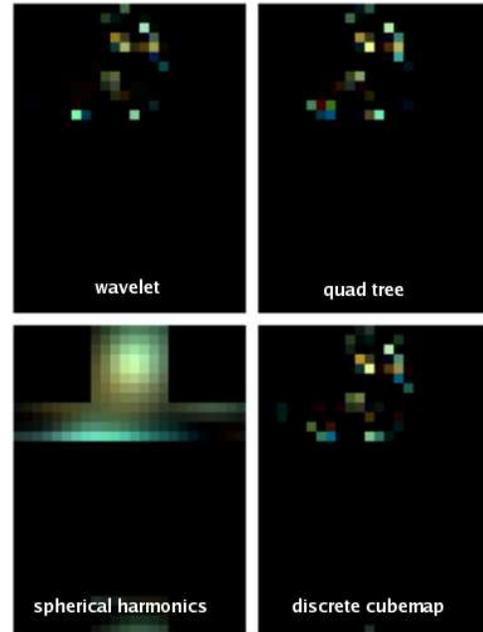


Figure 12. The reconstructed lightmaps corresponding to the four scenes in figure 11. From top left to bottom right - wavelet basis, quad tree, spherical harmonics, discrete cubemap

[6] S. R. Marschner and D. P. Greenberg. Inverse lighting for photography. In *Proceedings of the Fifth Color Imaging Conference, Society for Imaging Science and Technology*, 1997.
 [7] R. Ng, R. Ramamoorthi, and P. Hanrahan. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Trans. Graph.*, 22(3):376–381, 2003.
 [8] K. Nishino, Z. Zhang, and K. Ikeuchi. Determining reflectance parameters and illumination distribution from a sparse set of images for view-dependent image synthesis. In *ICCV*, 2001.
 [9] R. Ramamoorthi and P. Hanrahan. An efficient representation for irradiance environment maps. In *Siggraph*, 2001.
 [10] R. Ramamoorthi and P. Hanrahan. On the relationship between radiance and irradiance: Determining the illumination from images of a convex lambertian object. *J. Opt. Soc. Am. A*, 18(10), 2001.
 [11] I. Sato, Y. Sato, and K. Ikeuchi. Illumination from shadows. *PAMI*, 25(3):290–300, 2003.
 [12] P. S.C. and K. J.J. Illuminance flow. In *Computer analysis of images and Patterns*, pages 90–97, 2003.
 [13] P. S.C. and K. J.J. Surface illuminance flow. In *3DPTV*, 2004.
 [14] Y. Wang and D. Samaras. Estimation of multiple directional light sources for synthesis of augmented reality images. *Graph. Models*, 65(4):185–205, 2003.
 [15] Y. Yang and A. Yuille. Sources from shading. In *CVPR*, pages 534–439, 1991.
 [16] Y. Zhang and Y. Yang. Illuminant direction determination for multiple light sources. In *CVPR00*, pages I: 269–276, 2000.