Inverse Global Illumination: Recovering Reflectance Models of Real Scenes from Photographs

Yizhou Yu, Paul Debevec, Jitendra Malik, and **Tim Hawkins** *

Computer Science Division
University of California at Berkeley

ABSTRACT

In this paper we present a method for recovering the reflectance properties of all surfaces in a real scene from a sparse set of photographs, taking into account both direct and indirect illumination. The result is a lighting-independent model of the scene's geometry and reflectance properties, which can be rendered with arbitrary modifications to structure and lighting via traditional rendering methods. Our technique models reflectance with a low-parameter reflectance model, and allows diffuse albedo to vary arbitrarily over surfaces while assuming that non-diffuse characteristics remain constant across particular regions. The method's input is a geometric model of the scene and a set of calibrated high dynamic range photographs taken with known direct illumination. The algorithm hierarchically partitions the scene into a polygonal mesh, and uses image-based rendering to construct estimates of both the radiance and irradiance of each patch from the photographic data. The algorithm computes the expected location of specular highlights, and then analyzes the highlight areas in the images by running a novel iterative optimization procedure to recover the diffuse and specular reflectance parameters for each region. Lastly, these parameters are used in constructing high-resolution diffuse albedo maps for each surface.

The algorithm has been applied to both real and synthetic data, including a synthetic cubical room and a real meeting room. Re-renderings are produced using a global illumination system under both original and novel lighting, and with the addition of synthetic objects. Side-by-side comparisons show success at predicting the appearance of the scene under novel lighting conditions.

CR Categories: I.2.10 [**Artificial Intelligence**]: Vision and Scene Understanding—modeling and recovery of physical attributes I.3.7 [**Computer Graphics**]: Three-dimensional Graphics and Realism—color, shading, shadowing, and texture I.3.7 [**Computer Graphics**]: Three-Dimensional Graphics and Realism—Radiosity I.4.8 [**Image Processing**]: Scene Analysis—Color, photometry, shading

Keywords: Global Illumination, Image-Based Modeling and Rendering, BRDF Models, Reflectance Recovery, Albedo Maps, Radiance, Radiosity, Rendering

*Email:{yyz,debevec,malik,tsh}@cs.Berkeley.edu, yizhouy@acm.org, Website: http://www.cs.berkeley.edu/∼{yyz,debevec,malik,tsh}

1 Introduction

Computer graphics is being increasingly used to visualize real objects and environments. Applications in entertainment, architecture, interior design, virtual reality, and digital museums often require that aspects of the real world be rendered realistically from novel viewpoints and/or under novel illumination. For example, one would want to see how a room in a house would look like with different lighting, or how a statue would look at various times of day in a different wing of a museum. Lastly, one might want to realistically render a film location in different lighting, and add in digital props and characters, with the expectation that the rendered results would be the same as what would have happened had it all been for real.

Work in image-based modeling and rendering e.g. [18, 3, 22, 19, 12, 9, 6, 29]) has shown that photographs of a scene can be used along with geometry to produce realistic renderings of diffuse scenes under the original lighting conditions. However, challenges remain in making modifications to such scenes. Whether it is changing the geometry or changing the lighting, generating a new rendering requires re-computing the interaction of light with the surfaces in the scene. Computing this interaction requires knowing the reflectance properties (diffuse color, shininess, etc.) of each surface. Unfortunately, such reflectance property information is not directly available from the scene geometry or from photographs.

Considerable work (e.g. [32, 16, 5, 27, 21]) has been done to estimate reflectance properties of real surfaces in laboratory settings from a dense set of measurements. However, reflectance properties of real scenes are usually spatially varying, and typically change with use and age, making *a priori* laboratory measurements impractical. It would clearly be preferable to estimate the reflectance properties of an entire scene at once, with the surfaces being illuminated *in situ* rather than as isolated samples, and from a relatively sparse set of photographs. This is difficult for two reasons.

The first is that we wish to use only a sparse set of photographs of the scene, rather than exhaustively photographing every point of every surface from a dense set of angles. With such a set of photographs, we can expect to observe each surface point from only a small number of angles. As a result, there will be too little data to determine fully general bi-directional reflectance distribution functions (BRDFs) for each surface. We address this problem in two ways. First, we limit ourselves to recovering low-parameter reflectance models of the surfaces in the scene. Second, we assume that the scene can be decomposed into areas with related reflectance properties. Specifically, we allow the diffuse reflectance, or *albedo*, of the object to vary arbitrarily over any surface; the estimated albedo is computed as an image called an *albedo map*[1]. In contrast, we require that the directional reflectance properties (such as specular reflectance and roughness) remain constant over each area. In this work, such areas are specified as part of the geometry

[1]The commonly used term *texture map* is sometimes used to refer to this same concept. However, texture maps are also sometimes used to store surface radiance information, which is not lighting-independent.

recovery process.

The second problem we face is that in a real scene, surfaces will exhibit mutual illumination. Thus, the light that any particular surface receives will arrive not just from the light sources, but also from the rest of the environment through indirect illumination. As a result, the incident radiance of an observed surface is a complex function of the light sources, the geometry of the scene, and the as-yet-undetermined reflectance properties of all of the scene's surfaces. In this work, we use radiance data from photographs and image-based rendering to estimate the incident radiances of surfaces in the scene. This allows us to estimate the reflectance properties of the surfaces in the scene via an iterative optimization procedure, which allows us to re-estimate the incident radiances. We refer to this procedure as *inverse global illumination*.

Addressing these two problems makes it possible to robustly recover reflectance parameters from the limited radiance information present in a sparse set of photographs, and the accommodations made are appropriate for a wide variety of real scenes. Even when they are not met, the algorithm will compute the reflectance property parameters that best fit the observed image data, which in many cases can still yield a visually acceptable result.

The input to our algorithm is a geometric model of the scene, a set of radiance maps taken under known direct illumination, and a partitioning of the scene into areas of similar non-diffuse reflectance properties. The algorithm outputs a set of high-resolution albedo maps for the surfaces in the scene along with their specular reflectance properties, yielding a traditional material-based model. This output is readily used as input to traditional rendering algorithms to realistically render the scene under arbitrary lighting conditions. Moreover, modifications to the scene's lighting and geometry and the addition of synthetic objects is easily accomplished using conventional modeling methods.

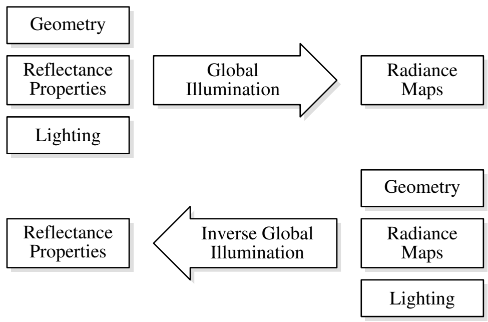

Figure 1: **Overview of the Method** *This figure shows the relationship between global illumination and inverse global illumination. Global illumination uses geometry, lighting, and reflectance properties to compute radiance maps (i.e. rendered images), and inverse global illumination uses geometry, lighting, and radiance maps to determine reflectance properties.*

1.1 Overview

The rest of this paper is organized as follows. In the next section we discuss work related to this paper. Section 3 describes *inverse radiosity*, a stepping stone to the full algorithm which considers diffuse scenes. Section 4 presents a technique for recovering specular reflectance properties for homogeneous surfaces considering direct illumination only. Section 5 describes how these two techniques are combined to produce our inverse global illumination algorithm. Section 6 completes the technical discussion by describing how high-resolution albedo maps are derived for the surfaces in the scene. Section 7 presents reflectance recovery results from both real and synthetic data, a description of our data acquisition, and synthetic renderings which are compared to real photographs. Section 8 presents some conclusions and avenues for future work.

2 Background and Related Work

The work we present in this paper has been made possible by previous work in BRDF modeling, measurement and recovery, geometry acquisition, image-based rendering, and global illumination.

In graphics, there is a long history of modeling surface reflectance properties using a small number of parameters. Recent efforts in this direction include models introduced in [14, 32, 25, 17]. These models have been shown to yield reasonable approximations to the reflectance properties of many real materials, and they have been used to produce realistic renderings.

On the other hand, considerable recent work has presented methods for measuring and recovering the reflectance properties of materials using imaging devices. [32] and [16] presented techniques and apparatus for measuring reflectance properties, including anisotropic reflection. [5] measured directional reflectance properties of textured objects. [27] and [21] showed that diffuse and specular reflectance properties could be recovered from multiple photographs of an object under direct illumination. [36] recovered reflectance properties of isolated buildings under daylight and was able to re-render them at novel times of day. [7] estimated material properties of parts of a scene so that they could receive shadows and reflections from synthetic objects. [10, 20] used a model of the scene and forward radiosity to estimate diffuse albedos to interactively modify the scene and its lighting. Although mutual illumination has been considered in the problem of shape from shading [23], it has not yet been fully considered for recovering non-diffuse reflectance properties in real environments. A survey of some of the methods is in Marschner [21].

Certain work has shown that changing the lighting in a scene does not necessarily require knowledge of the surface reflectance properties – taking linear combinations of a large set of basis images [24, 35] can yield images with novel lighting conditions.

Recent work in laser range scanning and image-based modeling has made it possible to recover accurate geometry of real-world scenes. A number of robust techniques for merging multiple range images into complex models are now available [34, 30, 4, 27]. For architectural scenes involving regular geometry, robust photogrammetric techniques requiring only photographs can also be employed. The model used in this research was constructed using such a technique from [9]; however, our basic technique can be used regardless of how the geometry is acquired.

Work in global illumination (e.g. [11, 15, 31, 37]) has produced algorithms and software to realistically simulate light transport in synthetic scenes. In this work we leverage the hierarchical subdivision technique [13, 1] to efficiently compute surface irradiance. The renderings in this paper were produced using Gregory Ward Larson's RADIANCE system [33].

Photographs taken by a camera involve nonlinearities from the imaging process, and do not have the full dynamic range of real world radiance distributions. In this work we use the high dynamic range technique in [8] to solve these problems.

3 Inverse Radiosity

Most real surfaces exhibit specular as well as diffuse reflection. Recovering both diffuse and specular reflectance models simultaneously in a mutual illumination environment is complicated. In this section, we consider a simplified situation where all surfaces in an environment are pure diffuse (Lambertian). In this case, the global illumination problem simplifies considerably and can be treated in

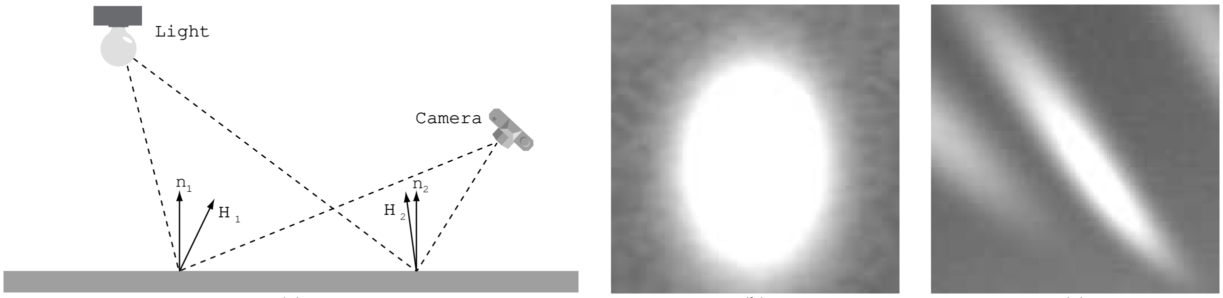

(a) (b) (c)

Figure 2: (a) The lighting and viewing directions at different points on a surface are different with respect to a fixed light source and a fixed viewpoint. This fact can be used to recover a low-parameter BRDF model for the surface from a single image. $\mathbf{n_i}$'s and H_i's are the normals and halfway vectors between lighting and viewing directions at different locations on the surface. We can infer that surface point P_2 with normal n_2 is close to the center of the highlight, and point P_1 with normal n_1 is relatively far away from the center. (b) An example of an isotropic specular highlight, (c) An example of an anisotropic specular highlight.

the radiosity framework [28]. We define *inverse radiosity* as recovering the diffuse albedo at each surface patch in the environment, provided that the geometry, the lighting conditions and the radiance distribution in the scene are known. In the next section we will discuss another simple case — recovering more general reflectance models with specularity considering only direct illumination — and we address the full problem in Section 5.

In the radiosity framework [28], the surfaces in the environment are broken into a finite number of patches. The partitioning is assumed to be fine enough that the radiosity and diffuse albedo of each patch can be treated as constant. For each such patch,

$$B_i = E_i + \rho_i \sum_j B_j F_{ij} \tag{1}$$

where B_i, E_i, and ρ_i are the radiosity, emission, and diffuse albedo, respectively, of patch i, and F_{ij} is the form-factor between patches i and j. The form-factor F_{ij} is the proportion of the total power leaving patch i that is received by patch j. It can be shown that this is a purely geometric quantity which can be computed from the known geometry of the environment [28].

We take photographs of the surfaces, including the light sources, and use a high dynamic range image technique [8] to capture the radiance distribution. Since Lambertian surfaces have uniform directional radiance distributions, one camera position is sufficient for each surface. Then B_i and E_i in Eqn. (1) become known. Form-factors F_{ij} can be derived from the known geometry. Once these are done, $\rho_i = (B_i - E_i)/(\sum_j B_j F_{ij})$. The solution to inverse radiosity is so simple because the photographs capture the final solution of the underlying light transport among surfaces.

4 Recovering Parameterized BRDFs from Direct Illumination

Before tackling the general case of reflectance recovery from photographs of mutually illuminated surfaces with diffuse *and* specular components, we study another special case. Consider a single surface of uniform BRDF which is illuminated by a point light source in known position and photographed by a camera, also in a known geometric position with respect to the surface(Fig. 2). Every pixel in the radiance image provides a measurement of radiance L_i of the corresponding surface point P_i in the direction of the camera, and the known light source position lets us calculate the irradiance I_i incident on that point.

Our objective is to use these data (L_i, I_i) to estimate the BRDF of the surface. Since the BRDF is a function of four variables (azimuth and elevation of incident and viewing directions) it is obvious that the 2-dimensional set of measurements for a single camera/light source pairing is inadequate to do this in general. However for many materials it is possible to approximate the BRDF adequately by a parameterized BRDF model with a small number of parameters (e.g. Ward [32], Lafortune [17], He [14] etc). We use Ward's parameterization in which the BRDF is modeled as the sum of a diffuse term $\frac{\rho_d}{\pi}$ and a specular term $\rho_s K(\boldsymbol{\alpha}, \boldsymbol{\Theta})$. Here ρ_d and ρ_s are the diffuse and specular reflectance of the surface, respectively, and $K(\boldsymbol{\alpha}, \boldsymbol{\Theta})$ is a function of vector $\boldsymbol{\Theta}$, the azimuth and elevation of the incident and viewing directions, and parameterized by $\boldsymbol{\alpha}$, the surface roughness vector. For anisotropic surfaces $\boldsymbol{\alpha}$ has 3 components; for isotropic surfaces $\boldsymbol{\alpha}$ has only one component and reduces to a scalar. The precise functional form of $K(\boldsymbol{\alpha}, \boldsymbol{\Theta})$ in the two cases may be found in Appendix 1.

This leads us to the following equation for each surface point P_i,

$$L_i = (\frac{\rho_d}{\pi} + \rho_s K(\boldsymbol{\alpha}, \boldsymbol{\Theta}_i)) I_i \tag{2}$$

where L_i, I_i and $\boldsymbol{\Theta}_i$ are known, and the parameters ρ_d,ρ_s, $\boldsymbol{\alpha}$ are unknowns to be estimated. Depending on whether we are using an isotropic or anisotropic model for the specular term we have a total of 3 or 5 unknown parameters, while there are as many constraining equations as the number of pixels in the radiance image of the surface patch. By solving a nonlinear optimization problem (see Appendix 1 for details), we can find the best estimate of ρ_d,ρ_s, $\boldsymbol{\alpha}$.

There are two important subtleties in the treatment of this optimization problem. One is that we need to solve a weighted least squares problem, otherwise the larger values from the highlight (with correspondingly larger noise in radiance measurements) cause a bias in parameter estimation. The second is the use of color information which needs to be done differently for dielectrics and metals. Both of these issues are discussed in Appendix 1.

To obtain an obvious global minimum for this optimization problem and achieve robust parameter recovery, the radiance image should cover the area that has a specular highlight as well as some area with very low specular component. If the highlight is missing, we do not have enough information for recovering specular parameters, and can only consider the surface to be diffuse.

5 Recovering Parameterized BRDFs in a Mutual Illumination Environment

We are now ready to study the general case when the environment consists of a number of surfaces and light sources with the surface reflectances allowed to have both diffuse and specular components.

Consider a point P_i on a surface patch seen by camera C_v (Fig. 3). The radiance from P_i in the direction of the camera is the re-

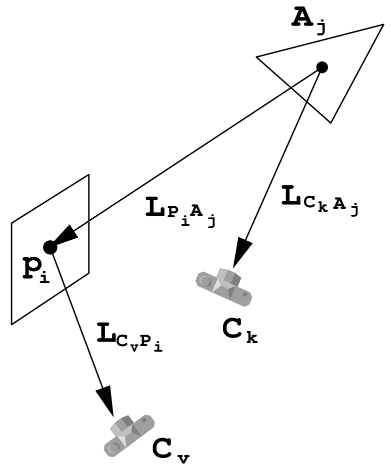

Figure 3: Patch A_j is in the radiance image captured by camera C_k. The specular component at A_j in the direction of sample point P_i is different from that in the direction of camera C_k. The difference is denoted by ΔS.

flection of the incident light contributed by all the light sources as well as all the surrounding surfaces. Eqn. (2) generalizes to

$$L_{C_vP_i} = E_{C_vP_i} + \rho_d \sum_j L_{P_iA_j} F_{P_iA_j} + \rho_s \sum_j L_{P_iA_j} K_{C_vP_iA_j}, \quad (3)$$

where $L_{C_vP_i}$ is the radiance value in the direction of camera C_v at some sample point P_i on the surface, $E_{C_vP_i}$ is the emission in the direction of camera C_v, $L_{P_iA_j}$ is the radiance value along the direction from patch A_j to point P_i on the surface, $F_{P_iA_j}$ is the analytical point-to-patch form-factor [2] between sample point P_i and patch A_j, and $\rho_s K_{C_vP_iA_j}$ is the specular term evaluated at P_i for a viewpoint at camera C_v and a light source position at patch A_j. The arguments, $\boldsymbol{\alpha}$ and $\boldsymbol{\Theta}$, of K have been dropped to simplify notation.

As before, our objective is to estimate ρ_d, ρ_s, and specular roughness parameters $\boldsymbol{\alpha}$. Of the other variables in Eqn. (3), $E_{C_vP_i} = 0$ for nonsources, and $L_{C_vP_i}$ can be measured directly from the radiance image at camera C_v. In general, the radiances $L_{P_iA_j}$ cannot be measured directly but have to be estimated iteratively. Suppose patch A_j in the environment appears in another radiance image taken by camera C_k(Fig. 3). Only if we assume A_j is Lambertian, does $L_{P_iA_j}$ in Eqn. (3) equal $L_{C_kA_j}$, the radiance from A_j to camera C_k. Otherwise, the diffuse components will be equal, but the specular components will differ.

$$L_{P_iA_j} = L_{C_kA_j} + \Delta S_{C_kP_iA_j} \quad (4)$$

Here $\Delta S_{C_kP_iA_j} = S_{P_iA_j} - S_{C_kA_j}$ is the difference between the specular components $S_{P_iA_j}$ and $S_{C_kA_j}$ of the radiances in the two directions. To compute the specular differences $\Delta S_{C_kP_iA_j}$, we need the BRDF of A_j, which is initially unknown. The estimation of ΔS (Section 5.1) therefore has to be part of an iterative framework. Assuming that the dominant component of reflectance is diffuse, we can initialize the iterative process with $\Delta S = 0$ (this sets $L_{P_iA_j} = L_{C_kA_j}$).

To recover BRDF parameters for all the surfaces, we need radiance images covering the whole scene. Each surface patch needs to be assigned a camera from which its radiance image is selected. At least one specular highlight on each surface needs to be visible in the set of images, or we will not be able to recover its specular reflectance and roughness parameters. Each sample point gives an

```
For each camera position C
  For each polygon T
    For each light source O
      Obtain the intersection P between plane of T and line CO'
        (O' and O are symmetric about T);
      Check if P falls inside polygon T;
      Check if there is any occlusion between P and O;
      Check if there is any occlusion between C and any point
        in a local neighborhood of P;
      /* A highlight area is detected if P passed all the above tests.*/
End
```

Figure 4: The specular highlight detection algorithm.

equation similar to Eqn. (3). From these equations, we can set up a weighted least-squares problem for each surface as in Appendix 1. During optimization, we need to gather irradiance at each sample point from the surface patches in the environment. One efficient way of doing this is to subdivide each surface into a hierarchy of patches [13, 1] and link different sample points to patches at different levels in the hierarchy. The solid angles subtended by the linked patches at the sample points should always be less than a prescribed threshold. There is a radiance value from the patch to the sample point and a ΔS associated with each hierarchical link.

For each sample point, we build hierarchical links to a large number of patches, and gather irradiance from these links. The amount of memory and computation involved in this process limits the number of samples for each highlight area. To make a reasonable tradeoff, we note that irradiance from indirect illumination caused by surrounding surfaces generally has little high-frequency spatial variation. Because of this, it makes sense to draw two sets of samples, one sparse set, and one dense set [2]. For the samples in the sparse set, we build hierarchical links and gather irradiance from the environment as usual. For the samples in the dense set, only their irradiance from light sources is computed explicitly, their irradiance from indirect illumination is computed by interpolation.

We are now ready to state the complete inverse global illumination algorithm. First detect all specular highlight blobs falling inside the radiance images using knowledge of the positions of the light sources, the camera poses, and the geometry (Fig. 4). Set the initial ΔS associated with each hierarchical link to zero. We can then recover an initial estimate of the BRDF parameters for each surface independently by solving a series of nonlinear optimization problems. The estimated specular parameters are used to update all ΔS's and $L_{P_iA_j}$'s associated with the hierarchical links. With the updated incident radiances, we can go back and re-estimate the BRDF parameters again. This optimization and update process is iterated several times to obtain the final solution of the BRDFs for all surfaces. The overall algorithm is shown in Fig. 5.

5.1 Estimation of ΔS

Suppose there is a hierarchical link $l_{P_iA_j}$ between a sample point P_i and a patch A_j which is visible to a camera C_k (Fig. 6). The ΔS for $l_{P_iA_j}$ is defined to be the difference of the specular component in directions $\vec{A_jP_i}$ and $\vec{A_jC_k}$. To estimate this difference, we need to obtain the specular component along these two directions given the BRDF parameters of patch A_j. A one-bounce approximation of ΔS for link $l_{P_iA_j}$ can be obtained by using Monte Carlo ray-tracing [32]. Because of off-specular components, multiple rays

[2] We choose the two sets of samples as follows. We first find the center of the highlight area in the image plane and rotate a straight line around this center to a number of different positions. The dense set of samples is the set of points on the surface corresponding to all the pixels on these lines. We choose the sparse set of samples on each line by separating two consecutive samples by some fixed distance in the object space.

Detect specular highlight blobs on the surfaces.
Choose a set of sample points inside and around each highlight area.
Build hierarchical links between sample points and patches in the environment and use ray tracing to detect occlusion.
Assign to each patch one radiance image and one average radiance value captured at the camera position.
Assign zero to ΔS at each hierarchical link.
For iter=1 to N
  For each hierarchical link,
    use its ΔS to update its associated radiance value;
  For each surface,
    optimize its BRDF parameters using the data
      from its sample points;
  For each hierarchical link,
    estimate its ΔS with the new BRDF parameters.
End

Figure 5: The *Inverse Global Illumination* algorithm.

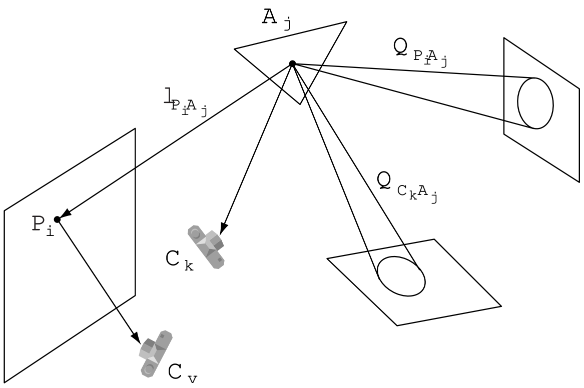

Figure 6: Random rays are traced around the two cones to obtain a one-bounce approximation of ΔS.

should be traced and the direction of the rays is randomized around the mirror directions of $\vec{A_j P_i}$ and $\vec{A_j C_k}$, respectively. For each possible ray direction, the probability density of shooting a ray in that direction is proportional to $K(\boldsymbol{\alpha_j}, \boldsymbol{\Theta})$ where $\boldsymbol{\Theta}$ encodes the incident and outgoing directions. Intuitively, most of the rays fall inside the two cones $Q_{P_i A_j}$ and $Q_{C_k A_j}$ centered at the two mirror directions. The width of each cone depends on the specular roughness parameters $\boldsymbol{\alpha_j}$ of patch A_j. The radiance along each ray is obtained from the patch hit by the ray. Suppose $L_{Q_{P_i A_j}}$ and $L_{Q_{C_k A_j}}$ are the average radiance values of the rays around the two cones, respectively, and $\rho_{s\,A_j}$ is the specular reflectance of patch A_j. Because the average value of Monte Carlo sampling approximates the total irradiance modulated by $K(\boldsymbol{\alpha_j}, \boldsymbol{\Theta})$, ΔS can simply be estimated as $\rho_{s\,A_j}(L_{Q_{P_i A_j}} - L_{Q_{C_k A_j}})$. This calculation could be extended to have multiple bounces by using path tracing [15]; we found that the one-bounce approximation was adequate for our purposes.

5.2 Practical Issues

We do not have a formal characterization of the conditions under which the inverse global illumination algorithm converges, or of error bounds on the recovered BRDF parameter values. In practice, we found it worked well (Section 7). Here we give some heuristic advice on how to acquire images to obtain good performance.

- *Use multiple light sources.* A specular highlight directly caused by one of the light sources should be captured on each surface. Having multiple light sources increases the probability that this can be achieved, and lets the whole scene receive more uniform illumination. This also increases the relative contribution of the diffuse component at any particular sample point P_i, and supports the $\Delta S = 0$ initialization, since highlights from different sources will usually occur at different locations on the surface.

- *Use concentrated light sources.* If the incoming radiance distribution is not very directional, the specular highlights will be quite extended and it will be difficult to distinguish the specular component from the diffuse one.

6 Recovering Diffuse Albedo Maps

In the previous sections, we modeled the reflectance properties as being uniform for each surface. In this section, we continue to do so for specular parameters because a small number of views of each surface does not provide enough information to reliably estimate specular parameters for each point individually. However, we relax this constraint on diffuse albedo and model it as a spatially varying function, an *albedo map*, on each surface. The diffuse albedo for any point x on a surface is computed as:

$$\rho_d(x) = \pi D(x)/I(x) \tag{5}$$

where $\rho_d(x)$ is the diffuse albedo map, $D(x)$ is the diffuse radiance map, and $I(x)$ is the irradiance map.

Suppose there is an image covering the considered surface which gives a radiance map $L(x) = D(x) + S(x)$ where $S(x)$ is the specular radiance map seen from the image's camera position. Then the diffuse radiance map in Eqn. (5) can be obtained by subtracting the specular component from each pixel of the radiance map $L(x)$ using the specular reflectance parameters already recovered. We estimate the radiance due to specular reflection as the sum of specular reflection due to direct and indirect illumination. The specular reflection due to direct illumination is computed from the knowledge of the direct lighting and the estimated reflectance properties, and we estimate the indirect specular reflectance by tracing a perturbed reflected ray into the environment in a manner similar to that in Section 5.1.

The irradiance $I(x)$ can be computed at any point on the surface from the direct illumination and by using analytical point-to-patch form-factors [2] as in previous sections of this paper. For efficiency, we compute the irradiance due to the indirect illumination only at certain sample points on the surfaces, and interpolate these indirect irradiance estimates to generate estimates for all surface points x. Of course, care must be taken to sufficiently sample the irradiance in regions of rapidly changing visibility to the rest of the scene.

Something that complicates estimating diffuse albedos in this manner is that in highlight regions the specular component of the reflectance $S(x)$ will be much larger than the diffuse component $D(x)$. As a result, relatively small errors in the estimated $S(x)$ will cause large relative errors in $D(x)$ and thus $\rho_d(x)$. However, just as a person might shift her view to avoid glare while reading a movie poster, we make use of multiple views of the surface to solve this problem.

Suppose at a point x on a surface, we have multiple radiance values $\{L_k(x)\}_{k=1}^{p}$ from different images. The highest value in this set will exhibit the strongest specular component, so we simply remove this value from consideration. For the remaining values, we subtract the corresponding specular estimates $S_k(x)$ from the radiance values $L_k(x)$, to obtain a set of diffuse radiance estimates $D_k(x)$. We compute a final diffuse radiance component $D(x)$ as a weighted average of the $D_k(x)$, with weights inversely proportional to the magnitude of the estimated specular components $S_k(x)$ to minimize the relative error in $D(x)$. We also weight the $D_k(x)$

values proportionally to the cosine of the viewing angle of the camera in order to reduce the influence of images at grazing angles; such oblique images typically have poor texture resolution and exhibit particularly strong specular reflection. Since we are combining information taken from different images, we smooth transitions at image boundaries using the image blending technique in [9].

Once diffuse albedo maps are recovered, they could be used to separate the diffuse and specular components in the specular highlight areas. This would allow recovering more accurate specular parameters in the BRDF model. In practice, however, we have found good estimates to be obtained without further refinements.

7 Results

7.1 Results for a Simulated Scene

We first tested our algorithm on a simple simulated cubical room with mutual illumination. This allowed us to verify the accuracy of the algorithm and compare its results to ground truth. All the six surfaces of the room have monochromatic diffuse and specular components, but each one has a distinct set of parameters. Each of the surfaces has spatially uniform specularity. We assigned two surfaces to be anisotropically specular and added 10-20% zero mean white noise to the uniform diffuse albedo of two surfaces to simulate spatial variations. We used the RADIANCE rendering system [33] to produce synthetic photographs of this scene. Six of the synthetic photographs were taken from the center of the cube with each one covering one of the six surfaces. Another set of six zoomed-in photographs were taken to capture the highlight areas. The scene was illuminated by six point light sources so that specular highlights could be observed on each surface. These twelve images along with the light source intensity and positions were used to solve the BRDF parameters. The images of the specular highlights are shown in Fig. 2(b)-(c). Some of the highlights are visually very weak, but corresponding parameters can still be recovered numerically. The original and recovered BRDF parameters are given in Table 1. For the last two surfaces with noisy diffuse albedo, the recovered albedo values are compared to the true average values. The total running time for BRDF recovery is about half an hour on a SGI O_2 180MHz workstation.

The numerical errors shown in Table 1 are obtained by comparing the recovered parameters with the original ones. There are three sources of error: BRDF modeling error, rendering error, and BRDF recovery error. BRDF modeling error comes from the inability of a given BRDF model to capture the behavior of a real material. By using the same model for recovery that RADIANCE uses for rendering, BRDF modeling error was eliminated for this test. However, because RADIANCE computes light transport only approximately, rendering error is present. We thus cannot determine the exact accuracy of our BRDF recovery. However, the test demonstrates that the algorithm works well in practice.

7.2 Results for a Real Scene

In this section we demonstrate the results of running our algorithm on a real scene. The scene we chose is a small meeting room with some furniture and two whiteboards; we also decorated the room with colored cards, posters, and three colored metallic spheres[3]. Once the BRDFs of the materials were recovered, we were able to re-render the scene under novel lighting conditions and with added virtual objects.

[3]The spheres were obtained from Baker's Lawn Ornaments, 570 Berlin Plank Road, Somerset PA 15501, (814) 445-7028.

	ρ_d	ρ_s	$\alpha_x(\alpha)$	α_y	γ
True	0.3	0.08	0.6	0.03	0
Recovered	0.318296	0.081871	0.595764	0.030520	-0.004161
Error(%)	6.10	2.34	0.71	1.73	
True	0.1	0.1	0.3		
Recovered	0.107364	0.103015	0.300194		
Error(%)	7.36	3.02	0.06		
True	0.1	0.01	0.1		
Recovered	0.100875	0.010477	0.101363		
Error(%)	0.88	4.77	1.36		
True	0.3	0.02	0.15		
Recovered	0.301775	0.021799	0.152331		
Error(%)	0.59	8.90	1.55		
True	0.2	0.05	0.05		
Recovered	0.206312	0.050547	0.050291		
Error(%)	3.16	1.09	0.58		
True	0.2	0.1	0.05	0.3	45
Recovered	0.209345	0.103083	0.050867	0.305740	44.997876
Error(%)	4.67	3.08	1.73	1.91	

Table 1: Comparison between true and recovered BRDF parameters for the six surfaces of a unit cube. The first and last surfaces have anisotropic specular reflection. They have two more parameters: second roughness parameter α_y and the orientation γ of the principal axes in a local coordinate system. The errors shown are the combined errors from both rendering and recovering stages.

7.2.1 Data Acquisition

We illuminated the scene with three heavily frosted 3-inch diameter tungsten light bulbs. Using high dynamic range photography, we verified that the lights produced even illumination in all directions. A DC power source was used to eliminate 60Hz intensity fluctuations from the alternating current power cycle.

We used a Kodak DCS520 color digital camera for image acquisition. The radiance response curve of the camera was recovered using the technique in [8]. We used a wide-angle lens with a 75 degree field of view so that we could photograph all the surfaces in the scene from a few angles with a relatively small number of shots. Forty high dynamic range radiance images, shown in Fig. 7, were acquired from approximately 150 exposures. Twelve of the images were taken specifically to capture specular highlights on surfaces.

The radiance images were processed to correct for radial light falloff and radial image distortion. Each of these corrections was modeled by fitting a polynomial of the form $1 + ar^2 + br^4$ to calibration data captured with the same lens settings used for the scene images. To reduce glare and lens flare, we shaded the lens from directly viewing the light sources in several of the images. Regions in the images corresponding to the light stands (which we did not model) or where excessive remaining glare was apparent were masked out of the images, and ignored by the algorithm. The thin cylindrical light stands which appear in the synthetic renderings have been added to the recovered model explicitly.

The radiance images were used to recover the scene geometry and the camera positions (Fig. 8) using the Façade [9] modeling system. Segmentation into areas of uniform specular reflectance was obtained by having each polygon of each block in the model (e.g. the front of each poster, the surface of each whiteboard, the top of each table) have its own uniform specular reflectance parameters.

The positions and intensities of the three light sources were recovered from the final three radiance images. During BRDF recovery, the area illumination from these spherical light sources was computed by stochastically casting several rays to each source.

7.2.2 BRDF Recovery

Given the necessary input data, our program recovered the surface BRDFs in two stages. In the first stage, it detected all the highlight regions and recovered parametrized BRDFs for the surfaces. In this stage, even if a surface had rich texture, only an average dif-

	ρ_d(red)	ρ_d(green)	ρ_d(blue)	ρ_s(red)	ρ_s(green)	ρ_s(blue)	α
whiteboard	0.5794	0.5948	0.6121	0.0619	0.0619	0.0619	0.0137
roundtable top	0.7536	0.7178	0.7255	0.0366	0.0366	0.0366	0.0976
door	0.6353	0.5933	0.5958	0.0326	0.0326	0.0326	0.1271
wall	0.8543	0.8565	0.8036	0.0243	0.0243	0.0243	0.1456
poster	0.1426	0.1430	0.1790	0.0261	0.0261	0.0261	0.0818
red card	0.7507	0.2404	0.3977	0.0228	0.0228	0.0228	0.0714
yellow card	0.8187	0.7708	0.5552	0.0312	0.0312	0.0312	0.1515
teal card	0.4573	0.5951	0.5369	0.0320	0.0320	0.0320	0.1214
lavender card	0.3393	0.3722	0.4437	0.0077	0.0077	0.0077	0.1144
red ball	0	0	0	0.5913	0.1862	0.3112	0
green ball	0	0	0	0.2283	0.3694	0.3092	0
blue ball	0	0	0	0.2570	0.3417	0.4505	0

Table 2: BRDF parameters recovered for the materials in the test room. All of them are isotropic, and most of them are plastic. The balls are metallic.

fuse albedo was recovered. Surfaces for which no highlights were visible the algorithm considered diffuse. The second stage used the recovered specular reflection models to generate diffuse albedo maps for each surface by removing the specular components.

The running time for each of the two stages was about 3 hours on a Pentium II 300MHz PC. The results show our algorithm can recover accurate specular models and high-quality diffuse albedo maps. Fig. 9 shows how specular highlights on the white board were removed by combining the data from multiple images. Fig. 10 shows the albedo maps obtained for three identical posters placed at different places in the room. Although the posters were originally seen in different illumination, the algorithm successfully recovers very similar albedo maps for them. Fig. 11 shows that the algorithm can remove "color bleeding" effects: colors reflected onto a white wall from the cards on the table do not appear in the wall's diffuse albedo map. Table 2 shows the recovered specular parameters and average diffuse albedo for a variety of the surfaces in the scene. We indicated to the program that all the materials are isotropic, and that the metallic spheres only have ideal specular components[4].

7.2.3 Re-rendering Results

We directly compared synthetic images rendered with our recovered BRDF models to real images. In Fig. 12, we show the comparison under the original lighting conditions in which we took the images for BRDF recovery. In Fig. 13, we show the comparison under a novel lighting condition obtained by removing two of the lights and moving the third to a new location, and adding a new object. There are a few differences between the real and synthetic images. Some lens flare appears in the real images of both figures, which we did not attempt to simulate in our renderings. We did not model the marker trays under the whiteboards, so their shadows do not appear in the synthetic images. In Fig. 13, a synthetic secondary highlight caused by specular reflection from the adjacent whiteboard appears darker than the one in the real image, which is likely due to RADIANCE's approximations for rendering secondary specularities. However, in both figures, real and synthetic images appear quite similar.

Fig 14 shows two panoramic views of the rendered scene. (a) shows the entire scene rendered with novel lighting. The original lights were removed and three track lights were virtually installed on the ceiling to illuminate the posters. Also, a strange chandelier was placed above the spheres on the table. The new lights reflect specularly off of the posters and the table. Since the chandelier contains a point light source, it casts a hard shadow around the midsection of the room. The interior of the chandelier shade is turquoise colored which results in turquoise shadows under the spheres. A small amount of synthetic glare was added to this image. (b) shows the result of adding synthetic objects to various locations in the room, including two chairs, a crystal ball, two metal boxes, and a floating diamond. In addition, a very large orange sculpture, was placed at the back of the room. All of the objects exhibit proper shadows, reflections, and caustics. The sculpture is large enough to turn the ceiling noticeably orange due to diffuse interreflection. The video for this paper shows a fly-through of each of these scenes.

8 Conclusions and Future Work

In this paper we have presented a new technique for determining reflectance properties of entire scenes taking into account mutual illumination. The properties recovered include diffuse reflectance that varies arbitrarily across surfaces, and specular reflectance parameters that are constant across regions. The technique takes as input a sparse set of geometrically and photometrically calibrated photographs taken under calibrated lighting conditions, as well as a geometric model of the scene. The algorithm iteratively estimates irradiances, radiances, and reflectance parameters. The result is a characterization of surface reflectance properties that is highly consistent with the observed radiances in the scene. We hope this work will be a useful step towards bringing visual spaces from the real world into the virtual domain, where they can be visualized from any angle, with any lighting, and with additions, deletions, and modifications according to our needs and imaginations.

There are a few directions for future research. We wish to apply our technique to more general geometrical and photometric data, such as multispectral radiance images and geometry accquired from laser scanners. It would be of significant practical value to be able to calibrate and use existing or natural illumination in recovering reflectance properties. The algorithm should be more robust to errors in the geometric model, misregistration of the photographs, and errors in the light source measurements. It would also be of theoretical value to obtain conditions under which the algorithm converges.

Acknowledgments

The authors wish to thank David Culler and the Berkeley NOW (Network of Workstations, http://now.cs.berkeley.edu/) project, and Tal Garfinkel for his help in using the NOW to render the video sequences. Thanks to Gregory Ward Larson for advice in using RADIANCE and estimating reflectance, Carlo Séquin for providing the sculpture model, and the reviewers for their valuable comments. This research was supported by a Multidisciplinary University Research Initiative on three dimensional direct visualization from ONR and BMDO, grant FDN00014-96-1-1200, the California MICRO program, Phillips Corporation, Interval Research Corporation, Pixar Animation Studios and Microsoft Graduate Fellowship.

[4]For surfaces that have only ideal specular reflection, such as mirrors, there is no diffuse component and the roughness parameter is zero. We can still recover their specular reflectance ρ_s from a single image by noting that the specular reflectance can be computed as the simple ratio between two radiance values. One is the radiance value in the image corresponding to the intersection between the surface and a ray shot from the camera position; the other is the radiance value of the environment along the reflected ray. In practice, we shoot a collection of rays from the camera position to obtain the average reflectance.

Appendix 1. BRDF Model and Parameter Recovery

In this appendix we present more details on the BRDF model, introduced in Section 4, and how its parameters are recovered. We use Ward's [32] model for the specular term in the BRDF, which could be modeled as either isotropic or anisotropic. In the isotropic case,

$$K(\alpha, \mathbf{\Theta}) = \frac{1}{\sqrt{\cos\theta_i \cos\theta_r}} \frac{\exp[-\tan^2\delta/\alpha^2]}{4\pi\alpha^2} \tag{6}$$

where α is a scalar surface roughness parameter, θ_i is the incident angle, θ_r is the viewing angle, and δ is the angle between the surface normal and the halfway vector H between the lighting and viewing directions. θ_i, θ_r are two components (along with ϕ_i, ϕ_r) of the vector $\mathbf{\Theta}$ which represents the incidence and viewing directions.

In the anisotropic case, we need two distinct roughness parameters α_x, α_y for two principal axes on the surface and an azimuth angle γ to define the orientation of these principal axes on the surface relative to a canonical coordinate system. Then, the parameter vector $\boldsymbol{\alpha}$ actually has three components $(\alpha_x, \alpha_y, \gamma)$ and we have:

$$K(\boldsymbol{\alpha}, \mathbf{\Theta}) = \frac{1}{\sqrt{\cos\theta_i \cos\theta_r}} \frac{\exp[-\tan^2\delta(\cos^2\phi/\alpha_x{}^2 + \sin^2\phi/\alpha_y{}^2)]}{4\pi\alpha_x\alpha_y} \tag{7}$$

where δ is the same as in the isotropic case, and ϕ is the azimuth angle of the halfway vector H projected into the local 2D coordinate system on the surface patch defined by the two principal axes. To compute ϕ, γ, which relates this coordinate system to the canonical coordinate system, is necessary.

Now to parameter recovery. We wish to find ρ_d, ρ_s and $\boldsymbol{\alpha}$ that minimize the squared error between the measured and predicted radiance,

$$e(\rho_d, \rho_s, \boldsymbol{\alpha}) = \sum_{i=1}^{m} (L_i - \frac{\rho_d}{\pi} I_i - \rho_s K(\boldsymbol{\alpha}, \mathbf{\Theta}_i) I_i)^2 \tag{8}$$

where L_i is the measured radiance and I_i is the irradiance (computable from the known light source position) at sample point p_i on the surface, and m is the number of sample points.

Note that given a guess of $\boldsymbol{\alpha}$, $K(\boldsymbol{\alpha}, \mathbf{\Theta}_i)$ becomes a known quantity, and minimizing the error e reduces to a standard linear least-squares problem for estimating ρ_d and ρ_s. Plugging in these values in the right hand side of Eqn. (8) lets us compute e as a function of $\boldsymbol{\alpha}$. The optimization problem thus simplifies to a search for the optimum value of $\boldsymbol{\alpha}$ to minimize $e(\boldsymbol{\alpha})$. This is either a one-dimensional or three-dimensional search depending on whether an isotropic or anisotropic model of the specular term is being used. We use golden section search [26] for the isotropic case, and the downhill simplex method [26] in the anisotropic case. It is convenient that neither method requires evaluating the derivative $e'(\boldsymbol{\alpha})$, and both methods are fairly robust.

To deal with colored materials, we estimate both diffuse and specular reflectance in each of the red, green, blue color channels. The specular roughness parameters $\boldsymbol{\alpha}$ are the same for all color channels. The nonlinear optimization is still over 1 or 3 parameters, since given $\boldsymbol{\alpha}$, ρ_d and ρ_s estimation for each channel remains a linear least squares problem.

To make the parameter estimation additionally robust, we make two simple extensions to the basic strategy derived above. The first is to solve a weighted least squares problem instead of the vanilla version in Eqn. (8). Radiance measurements from the highlight area have much larger magnitude than those from the non-highlight area. Correspondingly the error in those measurements is higher both because of noise in imaging as well as error in the BRDF model. Giving all the terms in (8) equal weight causes biased fitting and gives poor estimation of the diffuse reflectance. From a statistical point of view, the correct thing to do is to weight each term by the reciprocal of the variance of expected error in that measurement. Not having a good model for the error term, we chose a heuristic strategy in which the weight w_i for the i-th term in the summation in Eqn. (8) is set to $\frac{1}{K(\boldsymbol{\alpha_c}, \mathbf{\Theta}_i)}$ where $\boldsymbol{\alpha_c}$ is some *ad hoc* or iteratively improved roughness vector. Since the roughness of most isotropic materials is less than 0.2, we used an initial value between 0.1 and 0.2 for scalar α_c.

The second refinement to improve parameter recovery is to use specular color information. For instance, specular highlights on dielectric and plastic materials have the same color as the light source, while the color of specular highlights on metals is the same as their diffuse components, which is the color of the light modulated by the diffuse albedo. For plastic objects, there would be one distinct variable ρ_d for each color channel, but the same variable ρ_s for all color channels. For metallic objects, there would be one variable ρ_d for each channel and a common ratio between the specular and diffuse reflectance in all channels. Thus, we can reduce the degree of freedom from $2N$ to N+1 where N is the number of color channels. For plastic, we can still obtain both analytic and numerical linear least-squares solutions for the N+1 variables provided the other parameters are fixed. The program performs a heuristic test to determine whether a material should be estimated with the metal or plastic specular reflectance model. Our program first solves for the specular reflectance of each color channel separately and then checks to see if they are larger than the estimated diffuse components. If they are larger, then the material is considered metallic. Otherwise, the plastic model is used. Then the smaller number of parameters corresponding to these material types are solved.

References

[1] Aupperle, L., and Hanrahan, P. A hierarchical illumination algorithm for surfaces with glossy reflection. In *SIGGRAPH 93* (August 1993), pp. 155–164.

[2] Baum, D. R., Rushmeier, H. E., and Winget, J. M. Improving radiosity solutions through the use of analytically determined form factors. In *SIGGRAPH 89* (1989), pp. 325–334.

[3] Chen, E. QuickTime VR - an image-based approach to virtual environment navigation. In *SIGGRAPH '95* (1995).

[4] Curless, B., and Levoy, M. A volumetric method for building complex models from range images. In *SIGGRAPH '96* (1996), pp. 303–312.

[5] Dana, K. J., Ginneken, B., Nayar, S. K., and Koenderink, J. J. Reflectance and texture of real-world surfaces. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.* (1997), pp. 151–157.

[6] Debevec, P., Yu, Y., and Borshukov, G. Efficient View-Dependent Image-Based Rendering with Projective Texture-Mapping. In *9th Eurographics Workshop on Rendering*, (1998), pp. 105-116.

[7] Debevec, P. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH 98* (July 1998).

[8] Debevec, P. E., and Malik, J. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH 97* (August 1997), pp. 369–378.

[9] Debevec, P. E., Taylor, C. J., and Malik, J. Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach. In *SIGGRAPH '96* (August 1996), pp. 11–20.

[10] Drettakis, G., Robert, L., and Bougnoux, S. Interactive common illumination for computer augmented reality. In *8th Eurographics workshop on Rendering, St. Etienne, France* (May 1997), J. Dorsey and P. Slusallek, Eds., pp. 45–57.

[11] Goral, C. M., Torrance, K. E., Greenberg, D. P., and Battaile, B. Modeling the interaction of light between diffuse surfaces. In *SIGGRAPH '84* (1984), pp. 213–222.

[12] Gortler, S. J., Grzeszczuk, R., Szeliski, R., and Cohen, M. F. The Lumigraph. In *SIGGRAPH '96* (1996), pp. 43–54.

[13] Hanrahan, P., Salzman, P., and Aupperle, L. A rapid hierarchical radiosity algorithm. In *SIGGRAPH 91* (1991), pp. 197–206.

[14] He, X. D., Torrance, K. E., Sillion, F., and Greenberg, D. P. A comprehensive physical model for light reflection. In *SIGGRAPH 91*, (August 1991).

[15] Kajiya, J. The rendering equation. In *SIGGRAPH '86* (1986), pp. 143–150.

[16] Karner, K. F., Mayer, H., and Gervautz, M. An image based measurement system for anisotropic reflection. In *EUROGRAPHICS Annual Conference Proceedings* (1996).

[17] Lafortune, E.P.F., Foo, S., Torrance,K.E., and Greenberg, D.P. Non-Linear Approximation of Reflectance Functions. In *SIGGRAPH 97*, (1997), pp.117-126.

[18] Laveau, S., and Faugeras, O. 3-D scene representation as a collection of images. In *Proceedings of 12th International Conference on Pattern Recognition* (1994), vol. 1, pp. 689–691.

[19] Levoy, M., and Hanrahan, P. Light field rendering. In *SIGGRAPH '96* (1996), pp. 31–42.

[20] Loscos, C., Frasson, M.-C., Drettakis, G., Walter, B., Granier, X., and Poulin, P. Interactive Virtual Relighting and Remodeling of Real Scenes. Technical Report, iMAGIS-GRAVIR/IMAG-INRIA, (May 1999), http://www-imagis.imag.fr/Membres/Celine.Loscos/relight.html.

[21] Marshner, S. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University, August 1998.

[22] McMillan, L., and Bishop, G. Plenoptic Modeling: An image-based rendering system. In *SIGGRAPH '95* (1995).

[23] Nayar, S. K., Ikeuchi, K., and Kanade, T. Shape from interreflections. *International Journal of Computer Vision 6*, 3 (1991), 173–195.

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

[25] Oren, M., and Nayar,S.K., "Generalization of Lambert's Reflectance Model", *Computer Graphics Proceedings, Annual Conference Series*, pp.239-246 (1994).

[26] Press, W., Flannery, B., Teukolsky, S., and Vetterling, W. *Numerical Recipes in C*. Cambridge Univ. Press, New York, 1988.

[27] Sato, Y., Wheeler, M. D., and Ikeuchi, K. Object shape and reflectance modeling from observation. In *SIGGRAPH '97* (1997), pp. 379–387.

[28] Sillion, F. X., and Puech, C. *Radiosity and Global Illumination*. Morgan Kaufmann Publishers, San Francisco, 1994.

[29] Szeliski, R., and Shum, H.-Y. Creating full view panoramic image mosaics and environment maps. In *SIGGRAPH 97* (1997), pp. 251–258.

[30] Turk, G., and Levoy, M. Zippered polygon meshes from range images. In *SIGGRAPH '94* (1994), pp. 311–318.

[31] Veach, E., and Guibas, L. J. Metropolis light transport. In *SIGGRAPH 97* (August 1997), pp. 65–76.

[32] Ward, G. J. Measuring and modeling anisotropic reflection. In *SIGGRAPH '92* (July 1992), pp. 265–272.

[33] Ward, G. J. The RADIANCE lighting simulation and rendering system. In *SIGGRAPH '94* (July 1994), pp. 459–472.

[34] Y.Chen, and Medioni, G. Object modeling from multiple range images. *Image and Vision Computing 10*, 3 (April 1992), pp.145–155.

[35] Wong T.-T., Heng P.-A., Or S.-H. and Ng W.-Y. Image-based Rendering with Controllable Illumination. In *8th Eurographics Workshop on Rendering*, (June 1997), pp.13–22.

[36] Yu, Y., and Malik, J. Recovering photometric properties of architectural scenes from photographs. In *SIGGRAPH 98* (July 1998), pp. 207–217.

[37] Yu, Y., and Wu, H. A Rendering Equation for Specular Transfers and its Integration into Global Illumination. Eurographics'97, In *J. Computer Graphics Forum*, 16(3), (1997), pp. 283-292.

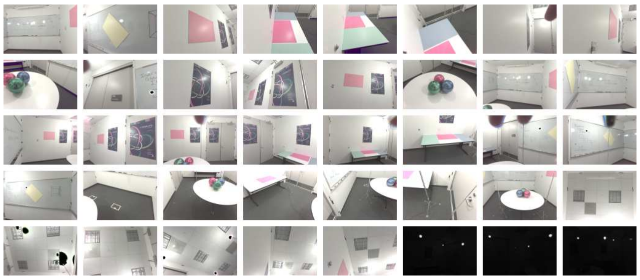

Figure 7: The complete set of forty radiance images of the room used to recover reflectance properties. Except for a few small areas, every surface in the room was seen in at least one radiance image. Each radiance image was constructed from between one and ten digital pictures depending on the dynamic range of the particular view. Black areas indicate regions which were saturated in all input images, and are not used by the recovery algorithm. The last three radiance images, reproduced ten stops darker than the rest, intentionally image the light bulbs. They were used to recover the positions and intensities of the sources. Full-resolution images are available in the electronic version of this paper.

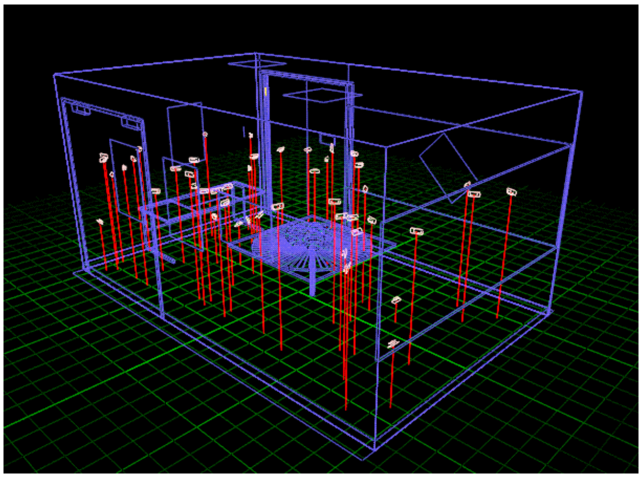

Figure 8: The model of the room, photogrammetrically recovered from the photographs in Fig 7. The recovered camera positions of the forty photographs are indicated.

Figure 9: The top picture is a radiance image of a whiteboard, showing strong specular highlights. The bottom picture shows the diffuse albedo map of the whiteboard recovered from several images. Unlike the radiance image, the diffuse albedo map has a nearly uniform background, and is independent of the illumination.

Figure 10: The diffuse albedo maps of three posters with the same texture. The posters were placed at different locations in the scene and thus received significantly different illumination. Nonetheless, the recovered albedo maps are nearly the same. For identification purposes, a small yellow square was placed in a different location on the lower right of each poster.

Figure 11: The top image shows a part of a wall that becomes noticeably colored from light reflecting from the cards placed on the table below, an effect known as ”color bleeding”. The bottom image shows the recovered albedo map of the same part of the wall. It is nearly uniform, showing that the color bleeding was properly accounted for. The black line has been added to indicate where the table top abutted the wall.

Figure 12: A comparison between real images (top) and synthetic renderings of our room with the recovered reflectance parameters (bottom). The simulated lighting is the same as in the original pictures, and the synthetic viewpoints have been matched to the recovered camera positions of the real images. The images show that good consistency was achieved.

Figure 13: A comparison between real and virtual, this time with novel lighting. Two of the lights were switched off and the third was moved to a new location. In addition, a real mirrored sphere was placed on the red card. The scene was photographed from two locations and these real views are shown in the top row. To render the bottom row, we recovered the camera positions and light source position in the top views, estimated the reflectance properties and position of the sphere, and added a virtual sphere to the model. The main noticeable difference is lens flare; however, some inaccuracies in the model (e.g. the whiteboard marker tray was not modeled) are also apparent. Otherwise, the illumination of the scene and appearance and shadows of the synthetic object are largely consistent.

(a) Synthetic rendering of room under novel illumination.

(b) Synthetic rendering of room with seven virtual objects added.

Figure 14: Panoramic renderings of the room, with various changes to lighting and geometry.