

Flexible Image-Based Photometric Reconstruction using Virtual Light Sources

Simon Gibson, Toby Howard, Roger Hubbard

Advanced Interfaces Group, Department of Computer Science, University of Manchester, UK.

Abstract

Photometric reconstruction is the process of estimating the illumination and surface reflectance properties of an environment, given a geometric model of the scene and a set of photographs of its surfaces. For mixed-reality applications, such data is required if synthetic objects are to be correctly illuminated or if synthetic light sources are to be used to re-light the scene. Current methods of estimating such data are limited in the practical situations in which they can be applied, due to the fact that the geometric and radiometric models of the scene which are provided by the user must be complete, and that the position (and in some cases, intensity) of the light sources must also be specified a-priori. In this paper, a novel algorithm is presented which overcomes these constraints, and allows photometric data to be reconstructed in less restricted situations. This is achieved through the use of virtual light sources which mimic the effect of direct illumination from unknown luminaires, and indirect illumination reflected off unknown geometry. The intensity of these virtual light sources and the surface material properties are estimated using an iterative algorithm which attempts to match calculated radiance values to those observed in photographs. Results are presented for both synthetic and real scenes that show the quality of the reconstructed data and its use in off-line mixed-reality applications.

1. Introduction

The ability to capture reflectance and illumination information from a real environment opens many useful applications in both computer graphics and computer vision. Among other things, performing such *photometric reconstructions* allows for realistic renderings of a real environment under novel lighting conditions, as well as the seamless introduction of correctly illuminated synthetic objects into images of the scene.

Modelling real-world reflectance and illumination data is a complex process. Reflective properties of real surfaces, in the form of *bi-directional reflectance distribution functions* (BRDFs) may be obtained using laboratory measurement equipment^{1,2}. Such equipment cannot, however, be used to capture data *in situ*, and is therefore limited in its practical applicability. The photometric properties of standard luminaires have also been well documented by their manufacturers. In real environments, however, factors such as age, cleanliness and utilization level affect the accuracy of these distributions, reducing their effectiveness in modelling real light sources. Modelling natural illumination is also possible

through the use of analytical models representing the sun and skylight^{3,4}, although fitting the model parameters to match specific illumination conditions is a difficult process, and is not always possible due to weather factors not accounted for by the model.

Image-based approaches to photometric reconstruction attempt to estimate illumination and reflectance data from one or more images of the scene. Current image-based methods, however, impose restrictions on the situations in which they can be applied. A complete geometric model of the scene must often be provided by the user, and radiometric information in the form of photographs of each surface are also required^{5,6,7}. Also, algorithms either place restrictions on the type of illumination being modelled⁸, or require knowledge of the position or spatial intensity distributions of the luminaires^{6,5,7}.

Requirements on the amount of user-supplied data must be reduced if image-based photometric reconstruction algorithms are to be applied in practical real-world situations. Constructing geometric and radiometric models of the scene places constraints on the size and complexity of the prob-

lems which can be tackled, since using computer vision algorithms to reconstruct a complete geometric description is a time-consuming and difficult process, especially in outdoor or large indoor environments. By reducing the amount of geometry we need to capture, the burden on the user is lifted and the reconstruction may be performed more quickly and efficiently. Aside from geometry capture, requiring *a-priori* information on the type and variety of luminaires incurs the greatest limitation. Environments often contain a mixture of unknown artificial light sources and natural illumination from the sun and skylight, complicating the process of modelling their position and photometric properties.

A general reconstruction algorithm will provide the capability to obtain usable estimates of surface reflectance and illumination distribution from incomplete geometric and radiometric models of the environment. This means that reconstructions may be performed in situations where modelling all scene geometry is not possible. A general algorithm will also not restrict the type or mixture of luminaires which can be included, and should not require the user to specify their photometric or geometric properties before the reconstruction can be performed. In this paper, we present such an algorithm and describe how it can be used to estimate the illumination and surface reflectance properties of real environments. Reconstructions are performed using a number of *virtual light sources* positioned around the scene, which model the effect of direct illumination from unknown luminaires, and indirect illumination reflected off unknown geometry. Our algorithm does not require the user to provide any information regarding the position or photometric properties of the luminaires. We will show how reconstructions of illumination and surface reflectance data can be performed in real situations containing both artificial and natural illumination.

The rest of this paper is structured as follows: In Section 2 we briefly review other image-based approaches to photometric reconstruction. Following that, Sections 3 to 5 outline our new algorithm, describe how virtual light sources are positioned and how they are used to estimate surface reflectance and illumination. Results for both synthetic and real images are then presented in Section 6. Finally, we draw conclusions from the work in Section 7.

2. Previous Work

Recent research efforts into photometric reconstruction have attempted to tackle the problem of capturing illumination and reflectance data from sets of images of a real environment. In ^{9,10} illumination data was captured in the form of a high-dynamic range (HDR) omni-directional image, and material properties of surfaces were estimated so that synthetic objects could be merged into photographs of real scenes. With these methods, the illumination is only valid near to where the omni-directional image was captured. In⁸, reflectance properties of architectural structures were estimated from images illuminated by sun and skylight. This

data was then used to re-render the structures under different distributions of skylight.

Recently, Yu *et al.* ⁶ presented methods to recover diffuse, specular, and textured material properties of surfaces from images. A complete geometric model of the environment and knowledge of the position and intensity of the light sources was first specified by the user. After iterative refinement, estimates of an albedo map and parameterized BRDF model for each surface were obtained. HDR images¹¹ were used to remove the effect of non-linearities in the image capture process and provide radiometric measurements for each surface in the scene.

The problem of estimating diffuse surface reflectance and light source emission properties for re-lighting real images has been investigated by Loscos *et al.* ^{5,7}. In their approach, diffuse reflectance and light source emissions were roughly approximated using radiosity-based measures similar to the earlier work of Fournier *et al.* ¹². This data was then used in conjunction with a cluster-based hierarchical radiosity algorithm to provide mechanisms to interactively modify the intensity parameters of the light sources, and to introduce synthetic objects into the scene.

The recent work of Sato *et al.* ¹³ took an alternative approach to the problem, by estimating a hemispherical illumination distribution using information contained in shadows falling on a surface. The illumination distribution, along with the diffuse and specular reflectance properties of the shadowed surface were then refined iteratively. This method provides the most flexibility in terms of the type and distribution of luminaires which may be included, since their photometric properties must not be specified *a-priori*. However, in a similar fashion to the omni-directional image techniques of ^{9,10}, the reconstructed data was only valid over a small region of the scene due to the hemispherical nature of the resulting illumination model.

3. Virtual Illumination Surfaces

The novel photometric reconstruction algorithm presented here is capable of estimating illumination and material properties in situations where incomplete geometric and radiometric data are provided, as well as situations where the light is a mixture of unknown artificial and natural illumination. This flexibility is achieved using *virtual light sources* which model the direct illumination from unknown light sources, and indirect illumination reflected off unknown geometry. It is important to understand that we are not attempting to accurately estimate the emission properties of each individual luminaire in the scene, but rather to construct a representation of the illumination which is consistent with the surface radiance values observed in the photographs.

Consider the simple environment shown in Figure 1, which consists of a room containing geometric objects and light sources, and assume that we have only reconstructed




Figure 1: A partial model of a scene with two unknown light sources.

geometric and radiometric data for one corner of the room, visible from a camera position C . The illumination reflected from a point x towards the camera C , can be described as follows:

$$L(x, \vec{\omega}_r) = \int_{\Omega} f(x, \vec{\omega}_i, \vec{\omega}_r) L_i(x, \vec{\omega}_i) \cos \theta_i d\omega_i \quad (1)$$

where $L(x, \vec{\omega}_r)$ is the radiance reflected from the surface point x towards the camera (in direction $\vec{\omega}_r$), Ω is the hemisphere above x , f is the surface BRDF, $L_i(x, \vec{\omega}_i)$ is the radiance seen in direction $\vec{\omega}_i$ from x , and θ_i is the angle between $\vec{\omega}_i$ and the surface normal at x . Note that for clarity, we assume that the surface has a zero emissive component.

The ways in which light reaches point x may be split into four separate transport paths:

- Direct illumination from known light sources;
- Indirect illumination from known surfaces;
- Direct illumination from unknown light sources;
- Indirect illumination from unknown surfaces.

The first two of these transport paths account for light coming from surfaces and light sources for which we have reconstructed geometry and radiometric data from the photographs (the shaded regions in Figure 1). The second two represent transfers from light sources or geometry for which we have no data. If only a small part of the scene has been photographed, these second two transport paths could account for the majority of the incident illumination at x .

We can modify Equation 1 to explicitly separate the contribution due to the first two and second two transport paths as follows:

$$L(x, \vec{\omega}_r) = \int_{\Omega^+} f(x, \vec{\omega}_i, \vec{\omega}_r) L_i(x, \vec{\omega}_i) \cos \theta_i d\omega_i + \int_{\Omega^*} f(x, \vec{\omega}_i, \vec{\omega}_r) L_i(x, \vec{\omega}_i) \cos \theta_i d\omega_i \quad (2)$$

where Ω^+ are directions over the unit hemisphere which intersect reconstructed geometry, and Ω^* are the directions which do not.

Because we have no *a-priori* information regarding the data in directions over Ω^* , we can make no assumptions




Figure 2: A Virtual Illumination Surface can represent the illumination contributions of unknown light sources and geometry.

about the types of luminaire or shapes of surfaces which are illuminating point x from these directions. The missing data can, however, be accounted for using suitably positioned *virtual illumination surfaces* which mimic the effect of this unknown illumination. Virtual illumination surfaces are similar in concept to the virtual walls developed to accelerate radiosity calculations^{14,15}. Each point on these virtual surfaces has an unknown spatial and spectral intensity distribution. The integral over Ω^* can then be transformed into an integral over these virtual surfaces, V_s . This gives a radiance contribution from unknown light sources and geometry of

$$\int_{V_s} f(x, x', \vec{\omega}_r) L_e(x, x') V(x, x') \cos \theta_i dA \quad (3)$$

where V_s is the union of all virtual illumination surfaces, x' is a point on those surfaces, and $V(x, x')$ is a visibility term which is 1 if x and x' are visible, and 0 otherwise. $L_e(x, x')$ is the radiance emitted by the surface at x' towards x , which is the same as the radiance seen by x in direction $x \rightarrow x'$ in the complete environment (Figure 2).

Combining Equations 2 and 3 gives an expression for the light reflected from a point x in direction $\vec{\omega}_r$ in terms of the surfaces visible in the photographs and the virtual illumination surfaces:

$$L(x, \vec{\omega}_r) = \int_{V_s} f(x, x', \vec{\omega}_r) L_e(x, x') V(x, x') \cos \theta_i dA + \int_{\Omega^+} f(x, \vec{\omega}_i, \vec{\omega}_r) L_i(x, \vec{\omega}_i) \cos \theta_i d\omega_i \quad (4)$$

Constructing virtual illumination surfaces like these allows us to abstract away from the shape, position, spatial and spectral intensity distributions of any unknown luminaires. It also allows for abstraction away from the reflectance properties of un-modelled geometry, as light reflected from these surfaces towards x is not considered any differently from light emitted from the luminaires.

The light reaching point x can now be re-classified using three transport paths:

- Direct illumination from known light sources;




Figure 3: The virtual illumination surface can be considered as a collection of virtual light sources, each with different spatial and spectral intensity distributions.

- Indirect illumination from known surfaces;
- Direct illumination from the virtual illumination surfaces.

Note how the last two paths in the earlier classification have been merged into a single direct contribution from the virtual illumination surfaces.

4. Virtual Light Sources

Virtual illumination surfaces can be approximated by a collection of M virtual light sources, each at position x'_i ($i = 0 \dots M - 1$) with area A_i , and spatial intensity distribution $L_{e_i}(x'_i, \vec{\omega}_i)$, as shown in Figure 3. This assumes that the radiance emitted by a virtual illumination surface in a direction $\vec{\omega}$ is constant over a small area around x'_i .

The contribution of the virtual illumination surfaces to the radiance reflected from x in direction $\vec{\omega}_r$ can therefore be approximated using a finite sum over the M virtual light sources:

$$L^*(x, \vec{\omega}_r) = \sum_{i=0}^{M-1} f(x, x'_i, \vec{\omega}_r) L_{e_i}(x'_i, \vec{\omega}_i) V(x, A_i) A_i \cos \theta_i \quad (5)$$

where $\vec{\omega}_i$ is the direction $x'_i \rightarrow x$, and $V(x, A_i)$ is the fraction of area A_i which is visible from point x .

Aside from the virtual illumination surfaces, the remaining part of the radiance which is reflected from x in direction $\vec{\omega}_r$ comes from contributions from the reconstructed surfaces. This can be approximated using Monte-Carlo sampling techniques, whereby a discrete number of rays sample the incident radiance field:

$$L^+(x, \vec{\omega}_r) = \sum_{j=0}^{J-1} f(x, \vec{\omega}_j, \vec{\omega}_r) L_i(x, \vec{\omega}_j) S(x, \vec{\omega}_j) \cos \theta_j \omega_j \quad (6)$$

where J rays are distributed in directions $\vec{\omega}_j$ over the hemisphere above x , $L_i(x, \vec{\omega}_j)$ is the radiance seen from x in direction $\vec{\omega}_j$, and ω_j is the solid angle associated with sampling direction j . Additionally, $S(x, \vec{\omega}_j)$ is the *surface visibility term* which is 1 if the sampling ray hits a surface for which we have radiometric data available from the photographs, and 0 otherwise. Note that we approximate

$L_i(x, \vec{\omega}_j)$ using the HDR image radiances. This implies that we assume the surface intersected by a ray from x in direction $\vec{\omega}_j$ is diffuse.

Combining Equations 5 and 6 then gives an expression for the radiance reflected from point x in direction $\vec{\omega}_r$ in terms of the radiances visible in the photographs and unknown contributions from the virtual illumination surfaces:

$$L(x, \vec{\omega}_r) \approx L^*(x, \vec{\omega}_r) + L^+(x, \vec{\omega}_r) \quad (7)$$

In the following sections, we will show how virtual light sources may be constructed. We will then go on to show how these sources can be used to reconstruct illumination and surface reflectance data for scenes where only partial geometric and radiometric models are available.

An overview of the reconstruction process is as follows: For a point x on the reconstructed scene geometry, each HDR image in which this point is visible gives us an estimate $L(x, \vec{\omega}_k)$ of the radiance reflected from x towards a camera k . This combination of a position and radiance estimate will be referred to as a *surface sample*. From Equation 7, we know that the sample's radiance estimate will be the same as that calculated using the reflectance properties at x , other radiance values obtained from the HDR images, and the radiance distributions of each virtual light source. Assuming we can parametrise the spatial intensity distribution of the sources in a suitable way, and we have more surface samples than we have virtual light source emission coefficients, we can construct an over-determined set of linear equations and solve for the light source radiance distributions. We will show how it is possible to solve for both the reflectance parameters and spatial intensity distributions using standard techniques from linear algebra. The following sections describe in more detail how this process can be initialised and how the reconstruction is performed.

4.1. Surface Sample Positioning

Samples are positioned over the surfaces of the scene by tracing a ray through the centre of each image pixel, intersecting the ray with the scene geometry, and placing a sample at the closest intersection point. For high-resolution images, however, this will generate an unmanageably large number of samples if one is created for each surface intersection. Also, the density of samples will depend on the number of cameras in which each surface is visible, thereby biasing the least-squares solution described in Section 5. For these reasons, we enforce a minimum distance between samples which limits the number created and ensures an even density over each surface. This is achieved by incrementally building a KD-tree using the sample positions, and using the tree to find the distance from each candidate position to the nearest sample which has already been created. If this distance is below a user-specified threshold, the candidate position is ignored. Figure 4 shows a typical set of samples created using this algorithm.




Figure 4: A KD-tree is used to enforce a minimum distance between samples, and ensure an even density over each surface.

4.2. Virtual Light Source Positioning

Once the surface samples have been created, we need to create a suitable set of virtual light sources. This involves solving two problems: firstly, where to position the virtual light sources, and secondly, how to describe their spatial intensity distributions.

Virtual light sources should be positioned in regions for which radiometric and geometric data is not available from the HDR images. They should also be positioned so that they illuminate the samples which have already been placed over the surfaces in the scene. We achieve this using a Monte-Carlo sampling algorithm, whereby a large number of rays are emitted in a hemispherical distribution from each surface sample. For each ray, we test to see if it intersects the scene geometry. If the ray does intersect a surface, and there is no radiometric data available at the intersection point (i.e. the intersection point is not visible in any of the HDR images), we create a virtual light source at that point. If the ray does not intersect any geometry, we position a virtual light source at the point of intersection between the ray and a bounding sphere positioned around the scene.

Again, using this algorithm *as is* will generate too many virtual light sources, so we enforce a minimum distance between source positions using a KD-tree, but use a larger minimum distance than when positioning surface samples (typically between 5 and 10 times). The area associated with each source, A_i , is calculated as the area of a disc with this distance as its radius. The visibility terms, $V(x, A_i)$ are estimated by jittering a small number of rays over the area of this disc and counting the fraction which do not intersect any scene geometry.

Along with a position, we also store a normal direction with each virtual light source. For sources positioned over surface geometry, this direction is simply the surface normal.

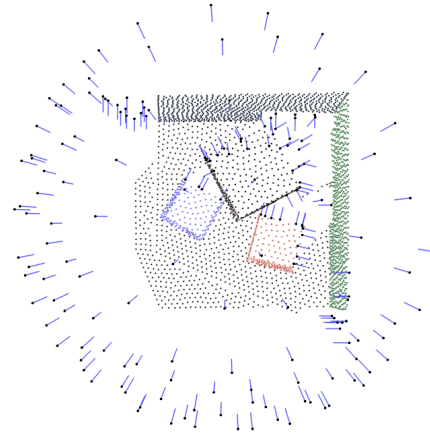


Figure 5: Virtual light sources are positioned around the scene, using a KD-tree to enforce a minimum distance between adjacent source positions.

For other sources, we use the radial direction of the bounding sphere. Figure 5 shows a typical distribution of virtual light source positions and normal directions for the surface sample set shown in the previous section.

Once the positions of the virtual light sources have been selected, we must choose an appropriate representation for their spatial intensity distributions. The unknown radiance distribution which the virtual light source is intended to mimic may have a complex shape, so it is necessary to use as flexible a representation as possible. The more complex the parametrisation is, however, the more variables we must solve for during the reconstruction process. For this reason, we have chosen to represent each virtual light source as the sum of a small number of cosine lobes, each with an associated power exponent and direction. The radiance emitted by virtual light source i from point x'_i in direction $\vec{\omega}$ is therefore:

$$L_{e_i}(x'_i, \vec{\omega}) = \sum_{j=0}^{K-1} E_{ij}(\vec{\omega} \cdot \vec{\omega}_{ij})^{\alpha_i} \quad (8)$$

where E_{ij} is the emittance for lobe $j = 0 \dots K - 1$, $\vec{\omega}_{ij}$ is the direction associated with lobe j , and α_i is the exponent of the lobes associated with virtual light source i . We typically set $K \leq 36$, and $\alpha_i = 20$. These values have been found to give an adequate coverage of the hemisphere while still allowing for relatively fine control over the radiance distribution. For virtual light sources positioned on surfaces, we again assume the unknown surface is diffuse, and create a single emission lobe aligned with the surface normal, setting $\alpha_i = 1$. This has the effect of reducing the number of unknown lobe emissions we must solve for, although the assumption could easily be lifted at the cost of increased solution time. Using this parametrisation in Equation 5 gives an

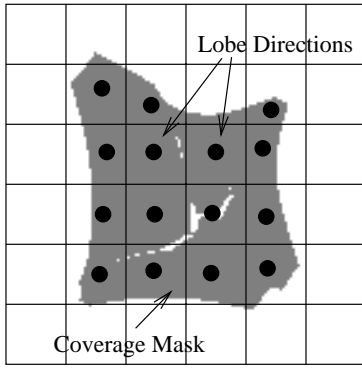


Figure 6: Emission lobe directions are selected using a coverage mask.

expression relating the unknown lobe emittances, E_{ij} , to the radiance reflected from a point x in direction $\vec{\omega}_r$.

In order to choose appropriate directions for each lobe, we build a two dimensional *coverage mask* for each virtual light source (typically of resolution 128×128 or above). Pixels in the mask are mapped uniformly onto directions over the hemisphere associated with the light source normal, and each pixel is set if a ray in that direction intersects surface geometry which is visible in the HDR images (recall that it is these positions over which we have positioned the surface samples). This gives an indication of the directions in which the virtual light source will be illuminating surface samples.

Once a mask has been constructed, it is uniformly subdivided into a small number of regions (typically, 6×6 , allowing for a maximum of $K = 36$ emission lobes per source), and the contents of each region of the mask are examined. If the fraction of pixels in one region is greater than a user-specified fraction (e.g. 10%), a lobe is created in the direction corresponding to the centroid of visible pixels (Figure 6). An example of the union of a set of typical emission lobes is given in Figure 7. After lobes have been created for each virtual light source, we may solve for their unknown emission coefficients. The next section describes how this can be done, assuming the reflectance properties of each surface are known. Following that, an iterative algorithm is given in Section 5.1 that is capable of solving for the unknown lobe emissions and unknown reflectances.

5. Illumination Reconstruction

Assuming that the surface BRDFs are known, a linear equation may be obtained from Equations 5 to 8, expressing the HDR image radiance associated with each sample x in terms of the unknown lobe emissions:

$$L(x, \vec{\omega}_r) - L^+(x, \vec{\omega}_r) = \sum_{i=0}^{M-1} \sum_{j=0}^{K-1} E_{ij} f(x, \vec{\omega}_i, \vec{\omega}_r) k_{ij} \quad (9)$$




Figure 7: A typical union of emission lobes for a virtual light source. Lobe directions are shown as lines piercing the surface.

where k_{ij} are geometric coefficients for this sample:

$$k_{ij} = V(x, A_i) (-\vec{\omega}_i \cdot \vec{\omega}_{ij})^{\alpha_i} A_i \cos \theta_i \quad (10)$$

and $\vec{\omega}_i$ is the direction from point x to the virtual light source at x'_i .

Provided we have more samples than unknown lobe emission coefficients, we may solve for all E_{ij} using standard linear algebra techniques. The system of equations will not be satisfied exactly, so we solve in the least-squares sense¹⁶. To ensure that we do not give extra weighting to larger image radiances, we scale each equation by $1/(L(x, \vec{\omega}_r) - L^+(x, \vec{\omega}_r))$. Additionally, we enforce a bound of $E_{ij} \geq 0$ to ensure we obtain a physically plausible solution.

5.1. Reconstruction with Unknown Reflectance

In most real-world situations, the reflectance properties of surface samples will not be known *a-priori*, and so must be estimated at the same time as the illumination distribution. To achieve this, we take a similar iterative approach to¹³. A diagram illustrating the reconstruction process is given in Figure 8. Before the reconstruction is performed, the user groups surfaces of the geometric model into different materials, and specifies each material as being either diffuse or specular.

As a pre-process, we compute and store information to accelerate the calculation of $L^+(x, \vec{\omega}_r)$ for each surface sample. This involves performing a dense hemispherical sampling of the space above each sample position using ray-casting. Each ray is intersected with the environment, and the intersection point is tested to see if it is visible in an HDR image. For points which are visible, we use the radiance from the HDR image to contribute to the incident radiance at the surface sample. Recall that this assumes that the surface at the intersection point is diffuse, although techniques similar to those proposed in⁶ could be used to allow this assumption to be removed. For diffuse surface samples, we store the total irradiance contribution with the sample. This allows $L^+(x, \vec{\omega}_r)$ to be quickly evaluated using different diffuse reflection coefficients. For specular samples, we store a

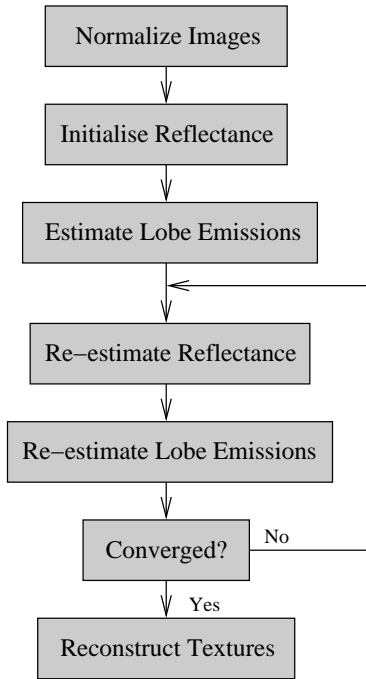


Figure 8: Iterative refinement of both reflectance data and lobe emission coefficients.

record of each radiance contribution, along with the direction it came from, with the surface sample. These radiance records allow Equation 6 to be quickly evaluated for different surface BRDFs without re-sampling the environment.

5.2. Reflectance Initialisation

The radiance values which are observed in the HDR images are a function of the reflectance properties of the surfaces, the emittance values of the illuminants and the camera sensors. Subsequently, we have no way of separating reflectance colour from illuminant colour – a red surface illuminated by white light could look very similar to a white surface illuminated by red light. The iterative reconstruction algorithm described in Section 5.3 is capable of identifying a combination of light source emission coefficients and surface reflectances which, when used to illuminate the surface samples give radiance values matching those observed in the HDR images, but there is no way of knowing that we have found the *correct* combination, as many different combinations could suffice.

As will be explained in more detail, estimates of diffuse surface reflectance are used to initialise an iterative process which alternates between estimating lobe emission coefficients and surface material data. In an attempt to ensure that the reconstruction converges towards an approximately correct combination of reflectance and emission coefficients, we

use a *colour normalization algorithm* during the initial reflectance estimation. This normalization algorithm processes each HDR image in order to remove the effects of illuminant geometry and colour¹⁷.

The initial estimates of sample reflectance are therefore calculated as follows: Firstly, a copy of the HDR image is normalized using the simple iterative normalization process described in¹⁷, whereby alternate steps remove the effects of illuminant intensity and illuminant colour respectively. Once the iterative process converges, we can estimate the reflectance of each diffuse material using an average of the normalized image intensities:

$$\rho_{d_\lambda} = \frac{1}{S} \sum_{i=0}^{S-1} a_{i_\lambda} \quad (11)$$

where ρ_{d_λ} is the diffuse reflectance coefficient of one material for channel $\lambda = [R, G, B]$, S is the number of samples covering surfaces associated with this material which are visible in the HDR images, and a_{i_λ} is the normalized image intensity at the pixel in which sample i is visible.

For glossy surfaces, we use a modified Phong reflection model¹⁸:

$$f(x, \vec{\omega}_i, \vec{\omega}_r) = \frac{\rho_d}{\pi} + \frac{n+2}{2\pi} \rho_s \cos^n \gamma \quad (12)$$

where ρ_d and ρ_s are the diffuse and specular reflectivities, n is the specular exponent, and γ is the angle between the ideal specular direction and $\vec{\omega}_r$. ρ_d is initialised in the same way as for diffuse materials, and ρ_s is set to zero. After material reflectances have been initialised, initial estimates of the lobe emission coefficients are found, as described above (Equation 9).

5.3. Iterative Refinement

Once initial values for the lobe emission coefficients have been found, we use them to update the estimates of surface reflectance. For diffuse materials, we find reflection coefficients for each colour channel satisfying:

$$\rho_d = \frac{\pi L(x_s \vec{\omega}_r)}{G_s} \quad (13)$$

for all s samples associated with the material in the least-squares sense, where G_s is the total incident radiance at x_s :

$$G_s = I_s + \sum_{i=0}^{M-1} \sum_{j=0}^{K-1} E_{ij} k_{ij} \quad (14)$$

and I_s is the irradiance contribution from surfaces visible in the HDR images which is stored with sample s , E_{ij} are the virtual light source emission coefficients, and k_{ij} are the geometric coefficients associated with sample s . Reflectances are clamped to the range $[0, 1]$ if necessary.

For each glossy material, we must find both diffuse and

specular reflectivity coefficients, as well as the specular exponent. Considering Equation 9, we can derive a cost function to be minimised using non-linear optimisation with unknown variables ρ_d , ρ_s and specular component n . This cost function is the sum (over all surface samples associated with the material) of differences between the left and right-hand sides of Equation 9, where $L^+(x, \vec{\omega}_r)$ is calculated using the pre-stored radiance records. We use a subspace-searching simplex algorithm¹⁹ to perform this optimisation. The sum of diffuse and specular reflectances is constrained to lie between zero and one, and the specular exponent constrained to be positive to ensure a physically plausible solution. These constraints are applied during the optimisation using penalty functions, whereby the error associated with a candidate set of parameters which do not satisfy the constraints is artificially raised.

5.4. Convergence

In order to halt the iterative refinement process shown in Figure 8, we need a measure of convergence. We use the total RMS error in surface sample radiance, when compared to the radiances measured from the HDR images (i.e. the sum of the squares of the differences between the left and right-hand sides of Equation 7). Once the change in RMS error between iterations falls below a user-specified threshold, re-estimating sample reflectance coefficients and lobe emissions is no longer improving the quality of the reconstructed solution. It has been noted previously, however, that RMS error is not a good indication of the visual differences between two images²⁰. Further research is required to develop a suitable measure of error and convergence for photometric reconstruction, and metrics such as Daly's Visible Differences Predictor²¹ should be investigated.

5.5. Textured Surfaces

Most surfaces in the real world do not have uniform reflectance properties. Following⁶, we assume that textured materials have uniform specular properties, but we allow for non-uniform diffuse reflectance coefficients. One approach to account for surface textures would be to estimate separate diffuse reflectances for each surface sample using Equation 13, rather than single coefficients for each material. We have found, however, that this raises the problem of illumination gradients in the HDR images being accounted for by changes in texture reflectance, rather than by changes in incident illumination.

To combat this problem, we estimate diffuse reflectance coefficients for each material as normal, and only reconstruct textures for each surface after the iterative refinement has converged. A texture map is constructed for each polygon by applying Equation 13 to each texel, thereby estimating its diffuse reflectance coefficients.




Figure 9: Reconstructed geometry (shown in wire-frame), overlaid on one of the images from which it was reconstructed.

6. Results

In this section, we present results of the new reconstruction algorithm when applied to both synthetic and real data. In all cases (both synthetic and real) geometric data was obtained using projective reconstruction techniques, and then updated to a Euclidean framework using self-calibration algorithms²². Once the camera positions were correctly calibrated, a geometric model was obtained by matching polyhedral primitives to features in each image²³. HDR images for the reconstructions were obtained using the techniques developed by Debevec *et al.*¹¹. A wire-frame representation of a reconstructed model is shown in Figure 9, overlaid on one of the images from which it was reconstructed.

6.1. Synthetic Data

The performance of the reconstruction algorithms was first tested on synthetic images, generated using the RADIANCE global illumination software²⁴. Pairs of images from two renderings of the same scene are shown in Figure 10. The left-hand pair of images show the scene with uniform diffuse surfaces, illuminated by three coloured light sources. The pair on the right shows the same scene with a specular floor polygon and two procedural textures on the walls. In the second pair, natural light from a hemispherical sky is illuminating the scene through a window, in addition to the artificial light sources. The left-hand image of each pair shows one of the reference HDR images which was used as input to the reconstruction algorithm. Six materials were created: one for each box, and one for each visible wall, and one for the floor.

The techniques described in this paper were then used to estimate the illumination distribution and reflectance properties of the visible surfaces. For these experiments, 3400 surface samples were positioned over the scene, along with 130 virtual light sources, containing just under 600 emission lobes. The time spent on solving the 3400×600 system of




Figure 10: Comparison of reference images against those rendered using reconstructed illumination and material data for two synthetic scenes. For each pair, the reference HDR image is shown on the left, and a rendering of the reconstructed scene on the right.

Material	Estimated	Reference
Box 1	0.72, 0.54, 0.70	0.70, 0.70, 0.70
Box 2	0.71, 0.18, 0.12	0.70, 0.20, 0.10
Box 3	0.12, 0.19, 0.79	0.10, 0.20, 0.70
Wall 1	0.20, 0.64, 0.46	0.20, 0.70, 0.30
Wall 2	0.29, 0.37, 0.66	0.30, 0.50, 0.70
Floor	0.80, 0.53, 0.62	0.70, 0.60, 0.50

Table 1: Diffuse reflectance coefficients estimated by the reconstruction algorithm for a simple synthetic environment.

bounded linear equations for this example was under 4 minutes for each colour channel, using the *DBOLS* software¹⁶ on a 250 MHz MIPS R10000 processor. A similar amount of time was required to estimate specular reflectance parameters at each iteration. The estimated diffuse reflectance coefficients are compared to reference values in Table 1. Note that although the values are not exact, approximately correct colours have been obtained even though the reconstruction was performed with partial geometric and radiometric models of the scene. We believe the main source of error is the initial image normalization, which assumes a ‘grey-world’ colour constancy model¹⁷.

Once the photometric properties of the scene were estimated, Monte-Carlo ray-tracing was used to re-render the environment. The right-hand image of each pair in Figure 10 shows the results. Note that features such as shadow boundaries, surface colours and intensity gradients are perceptually similar to the original algorithm, indicating that the algorithm is capable of providing a visually accurate photometric reconstruction.




Figure 12: Normalized RMS error in radiance for each colour channel as the number of iterations increases.

6.2. Real Data

The photometric reconstruction algorithms have also been tested using images of a real environment. In these cases, a small number of images were taken with a Kodak DC290 digital camera. Three images were used to reconstruct a partial geometric model of the environments, and one HDR image was captured to provide partial radiometric data for the scene. For the environments shown in Figures 11, 13 and 15, four diffuse materials were specified before the reconstructions were performed: one for the floor, one for the cardboard boxes, and two materials for the book.

Figure 11 shows a simple scene lit by artificial illumination. On the left-hand side is one of the original HDR images from which the illumination and material data was reconstructed. In the middle is a completely synthetic image, which was rendered using the reconstructed data (without surface textures). This image was obtained after 7 iterations using 1800 surface samples, and 950 virtual light source emission lobes. After surface textures were constructed, the image was re-rendered, and the result is shown on the right-hand side. Plots of the RMS error over the course of the reconstruction are given in Figure 12.




Figure 11: Comparison between real and synthetic images for a scene illuminated by artificial light. The reference HDR image is shown on the left, with the untextured synthetic model in the middle, and a synthetic rendering with textures on the right.




Figure 13: Results for a similar scene to Figure 11, this time lit by natural light.

Figure 14 shows the difference between the calculated and target radiance (measured from the HDR images) for each surface sample. The top graph was generated after a single iteration of the algorithm (i.e. after illumination was reconstructed using the initial reflectance estimates). The lower graph shows the result after the seventh iteration. The spikes in each graph can be attributed to the Monte-Carlo nature of the reconstruction and the non-uniform reflectance of the surfaces in the HDR images. In Figure 13, similar results are presented for the same scene, lit by natural light.

Once a photometric reconstruction has been performed, the data can be used in mixed-reality applications. Such examples are given in Figure 15. The left-hand image shows a synthetic image of the reconstructed scene rendered from a novel viewpoint. Additional synthetic light-sources have been introduced into the scene in the middle image, and used to illuminate the surfaces. In the right-hand image, several synthetic objects have been introduced into the model, and the shadows cast by these objects onto the reconstructed surfaces can be clearly seen.

7. Conclusion

We have presented a novel algorithm for reconstructing the illumination distribution and surface reflectance properties of a scene from a set of images. Our approach is capable of performing such reconstructions in situations when only a partial geometric and radiometric model of the scene is available. This is achieved by positioning virtual light sources

around the scene which mimic the effect of direct illumination from unknown light sources, and indirect illumination reflected from unknown geometry. Once these virtual sources have been positioned, their emission properties are obtained by solving a set of bounded linear equations. An iterative algorithm for simultaneously estimating surface reflectance properties has also been described.

Future research will examine suitable perceptual metrics for measuring the error and convergence properties of the reconstruction process. Also, further work is needed to improve the estimation of specular reflectance parameters for surfaces which reflect objects and light sources not included in the reconstruction.

Acknowledgements

The authors would like to thank their colleagues in the Advanced Interfaces Group for helpful discussions regarding this work. This research was supported by EPSRC grant number GR/M14531, entitled 'REVEAL: Reconstruction from Video of Environments with Accurate Lighting'.

References

1. G. J. Ward, "Measuring and Modeling Anisotropic Reflection", in *Computer Graphics (ACM SIGGRAPH '92 Proceedings)*, vol. 26, pp. 265–272, (July 1992).
2. K. F. Karner, H. Mayer, and M. Gervautz, "An image based measurement system for anisotropic reflection", *Computer Graphics Forum*, 15(3), pp. 119–128 (1996).
3. CIE-110-1994, "Spatial distribution of daylight – luminance distributions of various reference skies", tech. rep., International Commission on Illumination, (1994).



Figure 15: A reconstructed scene rendered from a novel viewpoint (left), with synthetic light sources illuminating the surfaces (middle), and with additional synthetic objects (right).

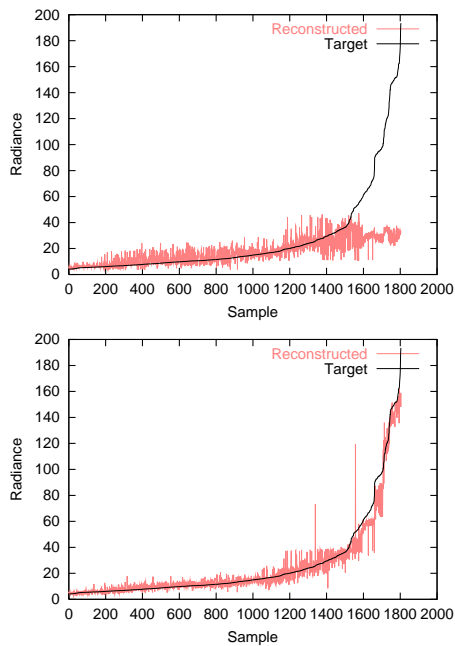


Figure 14: The difference between the reconstructed and target radiance values for each surface sample. Graphs for the red channel are shown after iteration 1 (top), and iteration 7 (bottom). Similar results were obtained for the green and blue channels.

4. A. J. Preetham, P. Shirley, and B. E. Smits, "A practical analytic model for daylight", in *Computer Graphics Proceedings, Annual Conference Series (Proc. SIGGRAPH '99)*, pp. 91–100, (August 1999).
5. C. Loscos, M. C. Frasson, G. Drettakis, B. Walter, X. Grainer, and P. Poulin, "Interactive virtual relighting and remodeling of real scenes", in *Rendering Techniques '99*, (New York, NY), pp. 329–340, Springer Wien, (1999).
6. Y. Yu, P. Debevec, J. Malik, and T. Hawkins, "Inverse global illumination: Recovering reflectance models of real scenes from photographs", *Proceedings of SIGGRAPH 99*, pp. 215–224 (1999).
7. C. Loscos, G. Drettakis, and L. Robert, "Interactive virtual relighting of real scenes", *IEEE Transactions on Visualization and Computer Graphics*, **6**(4), (2000).
8. Y. Yu and J. Malik, "Recovering photometric properties of architectural scenes from photographs", *Proceedings of SIGGRAPH 98*, pp. 207–218 (1998).
9. P. Debevec, "Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography", in *Computer Graphics (ACM SIGGRAPH '98 Proceedings)*, pp. 189–198, (1998).
10. I. Sato, Y. Sato, and K. Ikeuchi, "Acquiring a radiance distribution to superimpose virtual objects onto a real scene", *IEEE Transactions on Visualization and Computer Graphics*, **5**(1), pp. 1–12 (1999).
11. P. E. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs", *Proceedings of SIGGRAPH 97*, pp. 369–378 (1997).
12. A. Fournier, A. S. Gunawan, and C. Romanzin, "Common Illumination Between Real and Computer Generated Scenes", in *Proceedings of Graphics Interface '93*, (San Francisco, CA), pp. 254–262, Morgan Kaufmann, (May 1993).
13. I. Sato, Y. Sato, and K. Ikeuchi, "Illumination distribution from brightness in shadows: Adaptive estimation of illumination distribution with unknown reflectance properties in shadow regions", *Proceedings of IEEE ICCV'99*, (1999).
14. B. Arnaldi, X. Pueyo, and J. Vilaplana, "On the Division of Environments by Virtual Walls for Radiosity Computation", in *Proceedings of the Second Eurographics Workshop on Rendering* (P. Brunet and F. W. Jansen, eds.), (New York, NY), pp. 198–205, Springer-Verlag, (1994).
15. R. van Liere, "Divide and Conquer Radiosity", in *Proceedings of the Second Eurographics Workshop on Rendering* (P. Brunet and F. W. Jansen, eds.), (New York, NY), pp. 191–197, Springer-Verlag, (1994).
16. R. J. Hanson, "Linear least squares with bounds and linear constraints", Tech. Rep. SAND82-1517, Sandia Laboratories, (August 1982). *Software available from <http://lib.stat.cmu.edu/cmlib/src/bocls>*.
17. G. Finlayson, B. Schiele, and J. Crowley, "Comprehensive colour image normalization", *Proceedings of the European Conference on Computer Vision*, (1998).
18. E. P. Lafortune and Y. D. Willems, "Using the Modified Phong BRDF for Physically Based Rendering", Tech. Rep. Technical Report CW197, Department of Computer Science, Katholieke Universiteit Leuven, Leuven, Belgium, (November 1994).
19. T. Rowan, *Functional Stability Analysis of Numerical Algorithms*. PhD thesis, Department of Computer Science, University of Texas at Austin, (1990). *Software available from <http://www.netlib.org/opt>*.
20. A. McNamara, "Visual perception in realistic image synthesis", in *Eurographics 2000 (State of the Art Report)*, (Interlaken, Switzerland), (August 2000).
21. S. Daly, "The visible differences predictor: An algorithm for the assessment of image fidelity", *Digital Image and Human Vision*, pp. 179–206 (1993).
22. R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge University Press, (2000).
23. S. Gibson and T. L. J. Howard, "Interactive reconstruction of virtual environments from photographs, with application to scene-of-crime analysis", *ACM Symposium on Virtual Reality Software and Technology (VRST 2000)*, pp. 41–48 (2000).
24. G. J. Ward, "The RADIANCE Lighting Simulation and Rendering System", in *Computer Graphics Proceedings, Annual Conference Series, 1994 (ACM SIGGRAPH '94 Proceedings)*, pp. 459–472, (1994).