



Relighting with the Reflected Irradiance Field: Representation, Sampling and Reconstruction

ZHOUCHEN LIN

Microsoft Research, Asia

TIEN-TSIN WONG

The Chinese University of Hong Kong, People's Republic of China

HEUNG-YEUNG SHUM

Microsoft Research, Asia

Received March 19, 2001; Revised December 14, 2001; Accepted March 1, 2002

Abstract. Image-based relighting (IBL) is a technique to change the illumination of an image-based object/scene. In this paper, we define a representation called the *reflected irradiance field* which records the light reflected from a scene as viewed at a fixed viewpoint as a result of moving a point light source on a plane. It synthesizes a novel image under a different illumination by interpolating and superimposing appropriate recorded samples. Furthermore, we study the minimum sampling problem of the reflected irradiance field, i.e., how many light source positions are needed. We find that there exists a *geometry-independent* bound for the sampling interval whenever the second-order derivatives of the surface BRDF and the minimum depth of the scene are bounded. This bound ensures that when the novel light source is on the plane, the error in the reconstructed image is controlled by a given tolerance, regardless of the geometry. We also analyze the bound of depth error so that the extra reconstruction error can also be governed when the novel light source is off-plane. Experiments on both synthetic and real surfaces are conducted to verify our analysis.

Keywords: sampling, BRDF, light field, Lumigraph, plenoptic functions, image-based rendering, relighting

1. Introduction

To achieve realism, traditional geometry-based computer graphics makes use of sophisticated physics-based light transport models, precise geometry and complex rendering algorithms. However, unless carefully tuned, most computer-generated images still can be distinguished from real photographs. As an alternative, IBMR (Image-based Modeling and Rendering) synthesizes images from pre-recorded images, which can be real photographs. It employs a collection of reference images as samples of the plenoptic function (Adelson and Bergen, 1991) so that the rendering pro-

cess simply becomes ray query, interpolation and composition. Photorealistic images can be efficiently synthesized without a complex and long rendering process.

Much of the previous work in IBMR assumes that the lighting condition is fixed and the surface is Lambertian. Obviously, these assumptions cannot fully satisfy computer graphics needs, since illumination modification is a key operation. Image-based relighting (IBL) is a technique to change the illumination of an image-based object/scene. The goal is to modify the illumination in an interactive fashion while preserving correct visual appearance. This presents an important and challenging problem in IBMR.

In this paper, we propose a representation of the plenoptic function called the *reflected irradiance field* for IBL. The reflected irradiance field stores the reflection of surface irradiance when a point light source moves on a plane. With the reflected irradiance field, the relit object/scene can be synthesized simply by interpolating and superimposing appropriate sample reflections. If the geometry is not available, the image can be relit by arbitrary combination of point light sources on the plane. If the geometry is available, the luminaire can be of any kind, including directional light sources.

Like IBMR, IBL is a problem of sampling and reconstructing the plenoptic function. Though for IBMR, the sampling problem of the light field (Levoy and Hanrahan, 1996; Gortler et al., 1996) has been studied by Lin and Shum (2000) and Chai et al. (2000), no previous work has addressed the sampling problem of IBL. Clearly, the sampling criteria of relighting and that of the light field are different. Light field rendering is a view interpolation technique. Therefore, its sampling is more related to geometry, though the complex lighting effect are also recorded by the light field. In contrast, sampling in relighting is more concerned with surface reflectance. In this paper, we present a theoretical analysis on the sampling and reconstruction problems of the reflected irradiance field. We find that there exists a geometry-independent sampling bound that controls the reconstruction error when the light source is on the sampling plane. Finally, we study the reconstruction error when the light source leaves the plane and propose the depth layers for geometry recovery.

The rest of our paper is organized as follows. We first review some related previous work in Section 2. The definition of reflected irradiance field is presented in Section 3. In Section 4, we present a theoretical analysis on the problem “*how many samples are sufficient for on-plane relighting?*” In Section 5, we proceed to study the problem “*how much depth information is sufficient for off-plane relighting?*” Then, experimental results are presented in Section 6. Section 7 extends our analysis to the case of directional light sources. Finally, we draw conclusions and suggest future directions in Section 8.

2. Previous Work

Recently, much work has been done in the area of IBMR. It can be regarded as solving the problem of sampling and reconstructing the plenoptic function (Adelson and Bergen, 1991), as pointed out by

McMillan and Bishop (1995). To reduce the data volume, many researchers have worked on bringing down the dimensionality of the plenoptic function, ranging from five to two dimensions (McMillan and Bishop, 1995; Levoy and Hanrahan, 1996; Gortler et al., 1996; Shum and He, 1999; Chen, 1995).

Although there are advantages to reducing the dimensionality of the plenoptic function, there are also reasons to increase it (Wong et al., to appear). Since the original formulation of the plenoptic function assumes that lighting is constant, the illumination of generated images is also fixed. This limitation hampers the further application of IBMR. However, this limitation can be overcome by increasing the dimensionality of the plenoptic function, using extra parameters to represent the distribution of light sources.

There are several approaches for changing the lighting condition. Haeberli (1992) relit scenes by superimposing images. However, the direction, the type and the number of the light sources are limited to the original lighting setup during the scene capture. Using higher dimensional plenoptic functions, Wong et al. (1997) could both change the lighting (*relight*) and move the viewpoint. Similarly, Debevec et al. (2000) also defined a variant of the extended plenoptic function called the reflectance field to represent the radiance reflected from human face where dense samples were captured and interpolation was ignored. Yu and Malik (1998), Yu et al. (1999) and Sato et al. (1997) endeavored to recover the surface BRDF and approximate geometry. In Fournier et al. (1993), Loscos et al. (2000) and Gibson and Murta (2000), the shadows are more carefully treated, but only Lambertian BRDF is assumed in their global illumination computation. Magda et al. (2001), tried to recover the object geometry, instead of surface BRDF, by taking images of an object when a point light source moves on two concentric spheres enclosing the object. Unlike the above approaches, singular value decomposition (Golub and van Loan, 1989) was applied to extract a set of basis images (Zhang, 1998), where novel images were obtained by linearly combining the basis images.

Unfortunately, most of the previous work ignored the error analysis of sampling and reconstruction. Though the relit images were visually appealing, the correctness was never ensured. As an exception, superposition by steerable functions (Nimeroff et al., 1994) is exact for Lambertian surfaces, but it is limited to empirical cosine illumination models. The study on the related sampling issue is rare. Amanatides (1992) proposed

an anti-aliasing algorithm that based on Phong model for specular highlight by clamping the parameters in the Phong model. Assuming that the object is convex and Lambertian, Belhumeur and Kriegman (1996) have shown that *three* images are sufficient to relight the scene under various directional light sources.

This paper is an extension of our previous work (Lin et al., 2001). We propose the *reflected irradiance field* for IBL. Its concept is close to that of apparent BRDF (Wong et al., 1997) and the reflectance field (Debevec et al., 2000). The apparent BRDF and the reflectance field use a directional light source, while our reflected irradiance field adopts a point light source that moves on a plane. Our rationale is that sampling with point light source can be easily realized in real world. Nevertheless, our analysis can be easily applied to the apparent BRDF and the reflectance field when the light source plane is at infinity. When we study the sampling issues and reconstruction errors in our framework, we try to keep our analysis geometry-independent and to consider general BRDF.

3. The Reflected Irradiance Field

3.1. Tri-Planar Parameterization and its Simplifications

Adelson and Bergen have proposed a multi-dimensional *plenoptic function* (Adelson and Bergen, 1991) for evaluating low-level human vision models. The function describes environment appearance that our eyes observe. Besides the parameters for viewing direction, the center of projection and the wavelength, a time parameter is used to model all other unmentioned factors, such as the change of illumination and the change of environment. When it is constant, the scene is static and the illumination is fixed. Techniques (McMillan and Bishop, 1995; Gortler et al., 1996; Levoy and Hanrahan, 1996) based on this model naturally inherit this rigidity. However, the ability to modify the illumination configuration is very important in computer graphics. To express the change of illumination, we may extract an illumination component from the aggregate time parameter and explicitly specify it (Wong et al., to appear).

In general, the illumination component is a high dimensional function. It is necessary to assume a specific type of light source in order to reduce the dimension. Usually, directional light sources are used, such as the apparent BRDF of pixels in Wong et al. (1997) and

the reflectance field in Debevec et al. (2000). They are simple and convenient for synthetic images. However, they pose practical difficulties in capturing large object/scene because they are usually approximated by positioning strong spotlights at a distance that greatly exceeds the size of the object/scene. A point source, on the other hand, can be reasonably approximated by a tiny yet strong light bulb, such as a halogen bulb, which need not be placed far away. Because of this advantage, we utilize point light sources instead of directional sources like previous approaches (Wong et al., 1997; Debevec et al., 2000). Nonetheless, our analysis on sampling can be extended to the directional-source case. The extension is presented in Section 7.

The major disadvantage of the point-source formulation is that three parameters (rather than just two for the directional-source) are needed. To reduce the dimensionality, we constrain the point light source to lie on a *2D light-source plane*. Hence only two parameters (q, r) are needed. This approach is inspired by the light field (Levoy and Hanrahan, 1996) and Lumigraph (Gortler et al., 1996).

Then our extended plenoptic function can be represented by three planes (Fig. 1). The viewpoint is constrained to lie on the camera plane (uv plane) while the point light source is constrained to be on the light-source plane (qr plane). Together with the object plane (st plane), this representation requires six parameters. Each captured value in the six-dimensional table represents the light reflected through the window (s, t) and

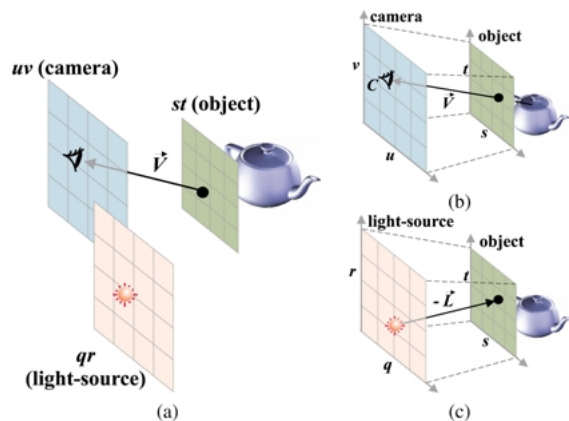


Figure 1. The tri-planar configuration and the dual light slab parameterization. (a) Each captured value represents the radiance reflected through (s, t) and received at a certain (u, v) when the scene is illuminated by a point source positioned at (q, r). (b) The light field is parameterized by the viewing vector. (c) The reflected irradiance field can be conceived as being parameterized by the light vector.

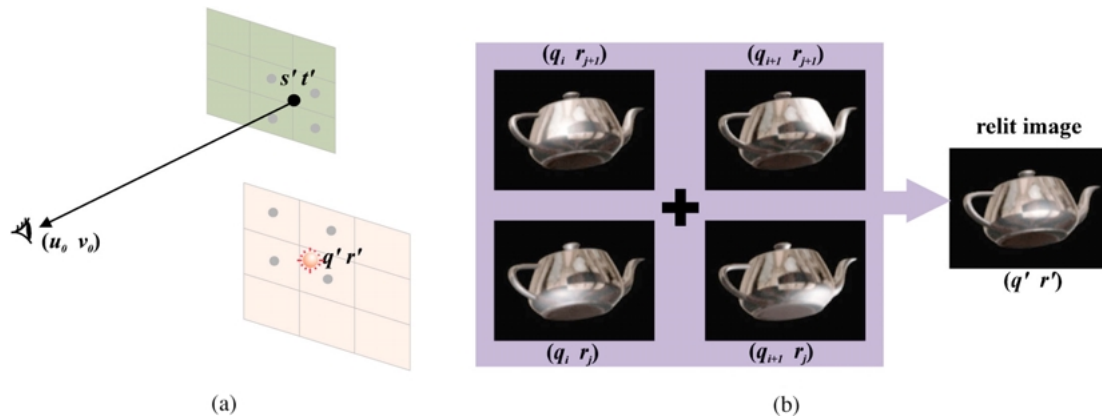


Figure 2. Relighting is done by bilinearly interpolating neighboring radiances in a 4D reflected irradiance field.

received at a certain (u, v) when the scene is illuminated by a point source positioned at (q, r) .

It is apparent that the tri-planar representation reduces to a four-dimensional light field and Lumigraph if the lighting condition is fixed (Fig. 1(b)). On the other hand, if the viewpoint (u, v) is fixed, the tri-planar representation reduces to another four-dimensional function (Fig. 1(c)), which we call the *reflected irradiance field*. It records the reflection of irradiance (w.r.t. the surface element) due to different light source positions.

3.2. Relighting

The rectilinear structure of reflected irradiance field allows us to simplify the relighting process. When the light source is on the qr plane, to determine the radiance received at a given (u_0, v_0) reflected from (s', t') with light at (q', r') , we can linearly interpolate the values of the nearest neighbors in the 4D reflected irradiance field (Fig. 2) in a pixel-by-pixel manner. This interpolation process is similar to that in light field and Lumigraph rendering. Generally speaking, the interpolation is quadri-linear because it is bilinear in both st and qr . If the resolution of the output image is the same as that of the input images, then the interpolation reduces to bilinear interpolation in qr . Real-time relighting (Wong, 2001) can be done by utilizing the bilinear interpolation module in texture-mapping hardware in the similar fashion as in Gortler et al. (1996). Figure 2(b) shows the schematic diagram illustrating the linear interpolation in (q, r, s, t) . Moreover, the image can be relit by various combinations of point light sources on the qr plane different from the one used




Figure 3. The sources of error when the novel light source is off-plane: incorrect ray query, interpolation weight, sample reflections and attenuation.

during the scene capture thanks to the linearity of light transport. The utilization of this superposition property can be also found in Shashua (1997), Nimeroff et al. (1994) and Wong et al. (1997). However, when the light source is not on the qr plane, the squared-distance attenuation of light energy must be considered in the synthesis process, and hence the need of geometry is apparent.

Since the relit image is synthesized from the sampled ones, it is very likely that there is reconstruction error. If the novel light source is on the qr plane, the interpolation error is ensured to be smaller than a given tolerance. This will be studied in detail in Section 4. However, if the novel light source is off-plane, the error will increase due to the lack of scene geometry. Similar phenomenon also occurs in the light field/Lumigraph (Levoy and Hanrahan, 1996; Gortler et al., 1996). Figure 3 illustrates the sources of error. As the actual depth of the scene is often unavailable during relighting, the object is often assumed on a

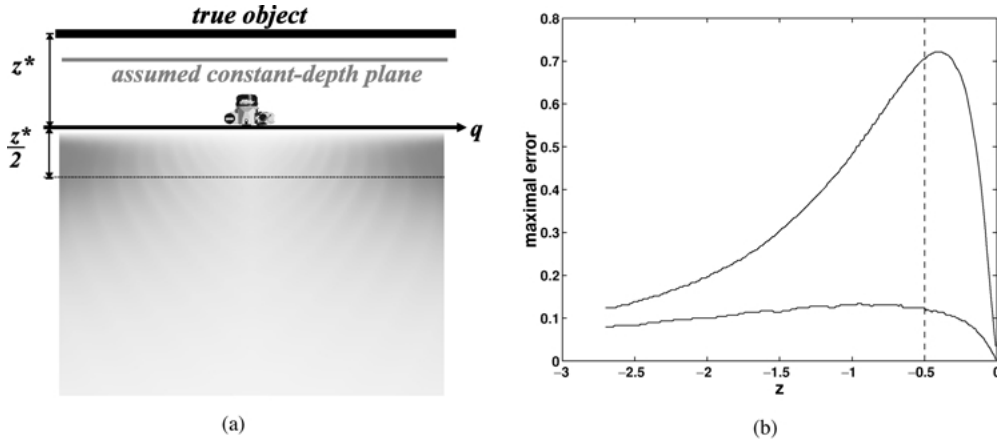


Figure 4. (a) The error map when the novel light source is off-plane. Darker color indicates the error is high while white indicates the error is low. (b) The plots of the errors at the right end (upper curve) and the center (lower curve) of the error map as a function of distance from the qr plane. Note that both curves reach their maximums around $-\frac{1}{2}z^* = -0.5$.

constant-depth plane (Isaksen et al., 2000; Chai et al., 2000; Lin and Shum, 2000). In the simplest case, the constant-depth plane is assumed to be the st plane. Hence the light gray pair of (q, r, s, t) is retrieved for interpolation, even though the correct one should be the dark gray pair. In the case of view interpolation, the error is solely due to this incorrect ray query. However, in the case of relighting, the interpolation weight is also incorrect because the intersection point on the qr plane will be mistakenly calculated as the light gray dot. The correct intersection point should be the dark gray dot. The incorrect intersection point not only induces incorrect interpolation weights, but may also result in wrong samples being chosen for interpolation. Moreover, due to the distance-square falloff of illumination intensity, the attenuation will also be incorrect because the true distance d_{true} from the novel light source is unavailable.

Off-plane relighting using geometry is only an approximation as it does not consider global illumination. Nevertheless, we can still study the behavior of the error when the global illumination is ignored. We call this error the off-plane error. The details will be presented in Section 5. Here we only illustrate the error due to off-plane relighting through a simple example in Fig. 4(a), where a planar object is positioned at z^* w.r.t. to the qr plane and the camera is placed close to the qr plane and facing the scene.

The off-plane error appears only when the assumed constant-depth plane is not coincident with the true object and the novel light source moves out of the qr plane. In Fig. 4(a), the maximum off-plane error at different off-plane light source positions is indicated by gray-

ness. The darker the grayness is, the larger the error is. When the novel light source is very close to the qr plane, the error is very small. As the light source moves farther away, the error increases on the left and right hand sides, but the error behind the camera only increases slightly. The error reaches its maximum around the line of $-z^*/2$ (Fig. 4(b)). Then it decreases when the novel light source moves farther away. The decrease is due to the squared-distance falloff.

4. Sampling Analysis for Relighting

Obviously, the quality of the interpolated image depends on the sampling rate. However, it is difficult to answer the question *how many samples are enough to avoid aliasing?* There are three sampling rates to consider, one for each plane. The sampling rate on st depends on the required resolution of the desired image. The tricky part is the sampling of uv and qr . Lin and Shum (2000) and Chai et al. (2000) have addressed the sampling issue on the uv plane for view interpolation. In this paper, we first address the question of *how many samples are enough to avoid noticeable artifacts in a relit image?* Our goal is to determine the *maximum allowable sampling interval* on the qr plane such that the intensity error in the relit image is smaller than a tolerance. Since incorrect geometry information can also cause aliasing when the light source is not on the qr plane, in the next section we move on to study the problem of *how much depth information is sufficient for correct relighting?*

4.1. Background

We assume a pin-hole camera, with a finite aperture to collect enough light. As the treatment of global illumination, such as cast shadow, interreflection and subsurface scattering (Hanrahan and Krueger, 1993), requires substantial knowledge of geometry and the rendering algorithms are generally complex (Shashua, 1997; Gibson and Murta, 2000; Loscos et al., 2000; Fournier et al., 1993; Yu and Malik, 1998; Yu et al., 1999; Sato et al., 1997), our analysis focuses on *local* reflection. The major assumption in our analysis is that the BRDF (ρ) in the scene must be second-order differentiable and the second-order derivatives are bounded. It is a reasonable assumption because a wide range of reflectance models (both physical and empirical) satisfy this assumption, including the Lambertian model, the Lafortune model (Lafortune et al., 1997), Ward's anisotropic model (Ward, 1992) and Oren and Nayar's non-Lambertian model (Oren and Nayar, 1994). Moreover, an ordinary camera acts as a low-pass filter since its resolution is finite. Hence our assumption may be relaxed to those BRDF models that are only first-order differentiable. To simplify the following analysis, we shall not take the camera resolution into account.

As the camera records exposure instead of radiance directly, we move on slightly to exposure. Referring to Fig. 5, the exposure at the pixel of A is:

$$E = I_i \Delta t,$$

where I_i is the irradiance received through the pixel window associated with A and Δt is the exposure time. As shown in Horn (1986), I_i is related to the intensity of the light source as follow:

$$I_i = \cos^4 \alpha \Delta F f^{-2} I_r, \quad \text{and} \\ I_r = \rho(\ell_1, \ell_2; \nu_1, \nu_2) \ell_2 d^{-2} I_l,$$




Figure 5. A graphical illustration of parameters in the analysis.

where α is the incident angle at C w.r.t. the optical axis of the camera, ΔF is the area of aperture, f is the camera focal length, I_r is the reflected radiance at A towards C , ρ is the surface BRDF, d is the distance between the light source and the surface element, I_l is the intensity of the light source, and (ℓ_1, ℓ_2) and (ν_1, ν_2) are the local directional coordinates of \vec{L} and \vec{V} w.r.t. the local frame $\{\vec{N}, \vec{M}, \vec{P}\}$ at A respectively (\vec{P} is not drawn in Fig. 5 since it is perpendicular to the paper). For clarity, we list the notations used and their meaning in Table 1.

Hence

$$E = c_1 \Phi(\ell_1, \ell_2; \nu_1, \nu_2),$$

where

$$c_1 = I_l \Delta F \Delta t f^{-2}, \\ \Phi(\ell_1, \ell_2; \nu_1, \nu_2) = \rho(\ell_1, \ell_2; \nu_1, \nu_2) \ell_2 d^{-2}.$$

$\cos^4 \alpha$ disappears because most cameras are designed to eliminate the diminishing artifact introduced by it and the field of views (FOVs) of cameras are often small.

As c_1 is constant throughout our analysis, the change of exposure depends on Φ only, which is a function of the light source position (q, r) . Therefore we analyze how Φ changes as the point light source moves on the qr plane.

4.2. Problem Formulation

When a novel point light source is positioned on the qr plane but not at any sample point, the exposure can be interpolated by those recorded values from neighboring qr grid points. Throughout the analysis, bilinear interpolation is assumed to be used for reconstruction since it is simple and can yield satisfactory image quality.

To simplify the discussion, we now look into the 1D interpolation along dimension q (Fig. 6). Then the bilinear interpolation becomes

$$\tilde{E}(q) = \frac{\tilde{I}}{I_l} [\lambda E(q_2) + (1 - \lambda) E(q_1)],$$

where q is the position of the novel light source, \tilde{I} is the emitted radiance from the novel light source, and $\lambda = |q_1 - q| / |q_1 - q_2|$.

Given an error tolerance ε , we want to find the maximum allowable spacing between samples along the

Table 1. Notations.

E	Exposure of a pixel
I_i	Irradiance through a pixel
I_r	Radiance of a surface point
I_l, \tilde{I}_l	Intensity of point and directional light sources
$\Delta t, \Delta F, f$	Exposure time, area of aperture and focal length of the camera
d	Distance between the light source and the surface point
z	Perpendicular distance to the light source plane
α	Incident angle w.r.t. the optical axis of the camera
ρ	BRDF
$\vec{N}, \vec{L}, \vec{V}$	Normalized surface normal, inverse light vector, and viewing vector
$\{\vec{N}, \vec{M}, \vec{P}\}$	The local frame at a surface point
\vec{Z}	Unit vector perpendicular to the light source plane
$(u, v), (s, t), (q, r)$	Coordinates on the camera, object, and light source planes
$(\ell_1, \ell_2, \ell_3), (v_1, v_2, v_3)$	Local coordinates of \vec{L} and \vec{V}
$(n_1, n_2, n_3), (l_1, l_2, l_3)$	Global coordinates of \vec{N} and \vec{L}
Φ	$= \rho \ell_2 d^{-2}$, depends on geometry and light source position
c_1, c_2	Constants for point and directional light sources
B_0, B_1, B_2	Geometry-independent upper bounds of $ d^2 \Phi $, $ d^3 \frac{\partial \Phi}{\partial q} $, and $ \frac{\partial^2 \Phi}{\partial q^2} $, respectively

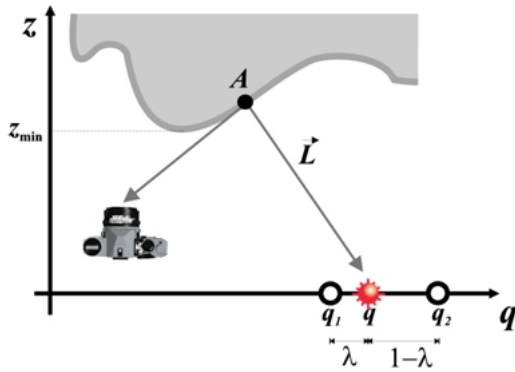


Figure 6. To reconstruct the exposure when a novel light source is positioned at q , linear interpolation (1D case) is done among neighbors.

q -axis such that the difference between every reconstructed exposure and the true exposure is smaller than ε . Obviously E is not band limited as a function of q . Therefore, if the exposure is treated in a continuous manner, it is usually impossible to exactly reconstruct the reflected irradiance field from discrete samples. However, exact reconstruction is unnecessary due to quantization, or the mapping from exposure to discrete greylevels. If the greylevel corresponding to the interpolated exposure is identical to that of the exact exposure (signal-level reconstruction) or the difference

is not visually noticeable (perception-level reconstruction), we may still say that the reflected irradiance field is perfectly reconstructed. Hence the sampling problem can be formulated mathematically as follows:

Find the maximum allowable spacing Δq between the sampling grids, such that

$$|\lambda E(q_1) + (1 - \lambda)E(q_2) - E(q)| \leq \varepsilon, \forall A,$$

and $\forall (q_1, q_2, q)$ satisfying

$$0 \leq q_2 - q_1 \leq \Delta q \quad \text{and} \quad q_1 \leq q \leq q_2.$$

4.3. Geometry-Independent Sample Spacing

We find that there exists a geometry-independent bound for the qr sampling interval whenever the object BRDF is second-order differentiable and the second derivatives are bounded. As long as the sampling interval does not exceed the bound and the novel light source lies on the qr plane, the reflected irradiance field can be correctly reconstructed within prescribed error without knowing any geometric information, such as depth and normal. This is important because surface reflectance depends heavily on the normal and acquiring accurate normal is even harder than acquiring depth.

For a pixel of interest, as we move the light source along the q -axis, the radiance and hence the exposure




Figure 7. An exposure curve can be plotted for each pixel in the image as the light source moves along q . The blue curve is for the pixel marked by the blue box, while the purple one is for the pixel in the purple box.

changes. We can plot the captured exposure as a function of q . Figure 7 shows the plots for two specific pixels in the left image. Given a tolerance ε , one can find an optimal sampling interval Δq_i for each specific curve such that the difference between the linearly-interpolated value and the true value is smaller than ε . For example, in Fig. 7, the optimal sample spacing for pixel ‘a’ is Δq_1 , while it is Δq_2 for pixel ‘b’. The difference in the shape of the curve may due to the surface normal orientation of the surface element visible through the pixel, the distance from the surface element to the light source, and the reflectance difference.

If we find the minimum sample spacing among all possible Δq (i.e. all possible combinations of normal orientation and depth), we can ensure that the reconstruction is correct (within a tolerance) without knowing the geometric details of the scene. All we need to know is the BRDF and the minimum distance z_{\min} (see Fig. 6) between the object and the qr plane. In the following, we show that for BRDFs with bounded second-order derivatives, the bound is positive.

From interpolation theory (Wang, 1999), there exists $\tilde{q} \in [q_1, q_2]$, such that

$$\begin{aligned} & \lambda E(q_1) + (1 - \lambda)E(q_2) - E(q) \\ &= \frac{1}{2}(q - q_1)(q_2 - q) \frac{\partial^2 E(\tilde{q})}{\partial q^2}. \end{aligned}$$

Therefore,

$$\begin{aligned} & |\lambda E(q_1) + (1 - \lambda)E(q_2) - E(q)| \\ & \leq \frac{1}{2} \left[\frac{(q - q_1) + (q_2 - q)}{2} \right]^2 c_1 B_2 = \frac{1}{8} (\Delta q)^2 c_1 B_2, \end{aligned} \quad (1)$$

where

$$B_2 = \text{G.I.} \max \left| \frac{\partial^2 \Phi}{\partial q^2} \right|,$$

or the geometry-independent upper bound of $\left| \frac{\partial^2 \Phi(q)}{\partial q^2} \right|$. Put in other words, B_2 is the maximum value of $\left| \frac{\partial^2 \Phi(q)}{\partial q^2} \right|$ for all possible positions of the light source and all possible positions and orientations of A .

Let us denote $\frac{\partial^2 \Phi}{\partial q^2}$ as Φ'' for simplicity. For a general BRDF, we can prove (see Appendix A), by a hybrid method of numerical computation and analytical deduction, that if the second-order derivatives of the BRDF are bounded, then:

$$\begin{aligned} |\Phi''| &< \frac{1}{z_{\min}^4} \max_{\ell_1, \ell_2, v_1, v_2} \{ |\rho''_{\ell_1 \ell_1}| + 0.77 |\rho''_{\ell_1 \ell_2}| \\ &+ 0.385 |\rho''_{\ell_2 \ell_2}| + 2.046 |\rho'_{\ell_1}| \\ &+ 2.098 |\rho'_{\ell_2}| + 3\rho \}, \end{aligned} \quad (2)$$

where ρ' and ρ'' are 1st-order and 2nd-order partial derivatives of ρ , respectively, and z_{\min} is the minimum depth of the scene w.r.t. the qr plane.

From the above equation, we know that the geometry-independent upper bound of the sampling interval exists. Unfortunately, usually the analytic bound is much larger than the exact bound, except for the Lambertian bound shown below. Therefore, it is better to find the bound by direct numerical computation. In our experiments, we simply performed a full search to find the maximum value of Φ'' under all possible positions and orientations of surface point. One trick is that searching along depth is unnecessary. The exact bound must be attained when the depth of the surface element is z_{\min} .

Substituting Φ'' into Eq. (1), the *sufficient* bound for Δq (sample spacing) is:

$$\Delta q \leq \sqrt{\frac{8\varepsilon}{c_1 B_2}}. \quad (3)$$

It can be seen that the sampling bound is tightly related to the maximum magnitude of the second-order derivatives of the BRDF. With the geometry-independent bound, we can use identical sampling intervals on the qr plane to sample the reflected irradiance field, regardless of the actual geometry. Next, we show the bounds of two popular reflectance models in computer graphics.




Figure 8. Illustration of Φ'' for different surface properties. (a) A specific surface element considered. (b) Φ'' of Lambertian surfaces. (c) Φ'' of Lafortune surfaces.

4.3.1. Lambertian Bound. For Lambertian surfaces, $\rho \equiv \rho_0 \leq 1/\pi$. After some manipulation, we have

$$\Phi'' = \frac{3\rho_0}{d^4} (-\ell_2 - 2n_1l_1 + 5\ell_2l_1^2),$$

where (n_1, n_2, n_3) and (l_1, l_2, l_3) are the directions of \vec{N} and \vec{L} , respectively.

Since

$$\begin{aligned} & -\ell_2 - 2n_1l_1 + 5\ell_2l_1^2 \\ &= n_1(-3 + 5l_1^2)l_1 + n_2(-1 + 5l_1^2)l_2 \\ & \quad + n_3(-1 + 5l_1^2)l_3, \end{aligned}$$

$|\vec{L}| = 1$, and $|\vec{N}| = 1$, using Cauchy's inequality, we have

$$\begin{aligned} & |-\ell_2 - 2n_1l_1 + 5\ell_2l_1^2| \\ & \leq \sqrt{[(-3 + 5l_1^2)l_1]^2 + [(-1 + 5l_1^2)l_2]^2 + [(-1 + 5l_1^2)l_3]^2} \\ & = \sqrt{5l_1^4 - 2l_1^2 + 1} \end{aligned}$$

Therefore

$$|\Phi''| \leq h(q) = \frac{3\rho_0}{d^4} \sqrt{5l_1^4 - 2l_1^2 + 1}.$$

Taking the partial derivative of $h(q)$ shows that

$$\max_q \{h(q)\} = h(q)|_{l_1=0},$$

hence

$$\text{G.I. max}|\Phi''| = 3\rho_0(z_{\min})^{-4}.$$

The equality holds because the right-hand side is attainable, as the example shown below.

To understand the physical meaning of Φ'' , we plot it against q for a specific Lambertian ($\rho_0 = 1/\pi$) surface element. The element is one unit distant from the qr plane and its normal is facing the qr plane (Fig. 8(a)). Intuitively, Φ'' can be regarded as the "acceleration" of radiance. As shown in Fig. 8(b), the radiance accelerates slightly as the light source moves closer to the surface element. It then decelerates quickly as the light source moves to the center of the qr plane. In this example, G.I. $\max|\Phi''|$ is attained when $q = 0$.

4.3.2. Lafortune Bound. Another model of interest is the Lafortune model (Lafortune et al., 1997) because of its capability of modeling a wide range of reflection phenomena, including diffuse, specular, off-specular, non-Lambertian, anisotropic and retro-reflection.

In the Lafortune model,

$$\begin{aligned} & \rho(\ell_1, \ell_2; \nu_1, \nu_2) \\ &= \sum_i (C_{i,1}\ell_1\nu_1 + C_{i,2}\ell_2\nu_2 + C_{i,3}\ell_3\nu_3)^{k_i}, \end{aligned}$$

where (ℓ_1, ℓ_2, ℓ_3) and (ν_1, ν_2, ν_3) are the coordinates of the normalized vectors \vec{L} and \vec{V} in the local frame $\{\vec{M}, \vec{N}, \vec{P}\}$ respectively, i is the index of cosine lobe, and $C_{i,1}$, $C_{i,2}$ and $C_{i,3}$ are the parameters defining the nature of surface reflectance.

An analytic *upper* bound of G.I. $\max|\Phi''|$ in the case of the Lafortune model is:

$$B_2 \leq (z_{\min})^{-4} \sum_i b_i^{k_i} (2.254k_i^2 + 2.305k_i + 1.317), \quad (\text{when } k_i \geq 2), \quad (4)$$

where

$$b_i = \max\{|C_{i,1}|, |C_{i,2}|, |C_{i,3}|\}.$$

The sketch of the proof is presented in Appendix B. It is much more complex than the Lambertian bound.

Again, we plot Φ'' of a Lafortune surface to illustrate its physical meaning. A specular surface (Lafortune, Personal communications) is chosen for examination. The setup of the surface element is the same as in Section 4.3.1. As shown in Fig. 8(c), the reflected radiance from steel changes vigorously and the change concentrates in a very small region. The maximum magnitude of its second-order derivative is much larger than that of a Lambertian. Hence Δq_{\min} is quite small and implies the need for a higher sampling rate.

5. Error Analysis of Off-Plane Relighting

As explained in Section 3.2, off-plane relighting introduces extra error into the reconstructed image if the scene geometry is inexact or unknown. If we want to bound the off-plane error within ε when the novel light source is positioned at *any* position behind the qr plane, then we can prove that the *sufficient* bound of the depth error is:

$$|\Delta a| \leq \frac{27z^3\varepsilon}{4c_1(3B_0 \cos \phi + B_1)}, \quad (5)$$

where $|\Delta a|$ is the distance (along the viewing direction) between the true surface element and the assumed depth, z is the depth of the true surface element w.r.t. the qr plane, $\cos \phi = \hat{V} \cdot \hat{Z}$, $B_0 = \text{G.I. max}|d^2\Phi|$, and $B_1 = \text{G.I. max}|d^3\Phi|$.

The detailed informal proof is given in Appendix C. B_0 and B_1 can be found using a similar method as in finding B_2 . The analytic upper bounds of B_0 and B_1 of the general BRDF, Lambertian model and Lafortune model are listed in Appendix D.

Equation (5) indicates that the tolerable depth error is dependent on the orientation of the viewing ray and the depth of the corresponding surface element. A by-product of the proof of Eq. (5) is that the maximum error occurs when the novel light source is positioned at a distance of around $-z_{\min}/2$ from the qr plane, because the coefficient $\gamma/[(1+\gamma)^2(1+\gamma-\tau)]$ of Eq. (8) in Appendix C reaches the maximum $4/27$ when $\gamma = 1/2$ (ignoring the influence of small τ therein) and the other term in Eq. (8) is nearly independent of γ after

magnification. Though the scene depth may not always be z_{\min} , from Eq. (5) we know that closer objects need more accurate geometry. So close objects dominate the maximum off-plane error when the light source leaves the qr plane. As a result, the error is maximal when the light source moves around the plane $z = -z_{\min}/2$. This explains the error distribution in Fig. 4 and will be further verified by our experiments.

5.1. Depth Layers for Geometry Recovery

Equation (5) suggests that when recovering the local geometry we may actually search the depth along the viewing direction in a discrete manner. Starting from an initial guess a_0 , the relation between successive depths is:

$$\begin{aligned} a_{i+1} &= \frac{\tilde{\varepsilon}}{3B_0 \cos \phi + B_1} (a_{i+1} \cos \phi + z_0)^3 \\ &= a_i + \frac{\tilde{\varepsilon}}{3B_0 \cos \phi + B_1} (a_i \cos \phi + z_0)^3 \end{aligned} \quad (6)$$

where $\tilde{\varepsilon} = 27\varepsilon/(4c_1)$, and z_0 is the depth of the camera w.r.t. the qr plane. To simplify Eq. (6), we introduce what we call the *radiometric depth*:

$$\begin{aligned} \hat{a} &= \left(\frac{\tilde{\varepsilon} \cos \phi}{3B_0 \cos \phi + B_1} \right)^{\frac{1}{3}} (a \cos \phi + z_0) \\ &= \left(\frac{\tilde{\varepsilon} \cos \phi}{3B_0 \cos \phi + B_1} \right)^{\frac{1}{3}} z, \end{aligned}$$

where $z = a \cos \phi + z_0$ is the depth w.r.t. the qr plane. Then Eq. (6) becomes:

$$\hat{a}_{i+1} - \hat{a}_{i+1}^3 = \hat{a}_i + \hat{a}_i^3 \equiv \hat{a}_{i+\frac{1}{2}}.$$

The solution to the above equation is:

$$\hat{a}_{i+1} = \frac{2}{\sqrt{3}} \cos \left(\frac{1}{3} \arccos \left(-\frac{3\sqrt{3}}{2} \hat{a}_{i+\frac{1}{2}} \right) + \frac{4\pi}{3} \right). \quad (7)$$

We can prove that any depth \hat{a} between \hat{a}_i and \hat{a}_{i+1} can be represented by $\hat{a}_{i+\frac{1}{2}}$, i.e., the off-plane error does not exceed ε if \hat{a} is estimated as $\hat{a}_{i+\frac{1}{2}}$. The proof can be found in Appendix E. Equation (7) is valid only when $\hat{a}_{i+\frac{1}{2}} \leq \frac{2}{3\sqrt{3}}$, or $\hat{a}_i \leq \hat{a}^{(\infty)} = (\sqrt[3]{1+\sqrt{2}} + \sqrt[3]{1-\sqrt{2}})/\sqrt{3} = 0.3441\dots$, otherwise \hat{a}_{i+1} does not




Figure 9. Depth layers with different initial guesses on the geometry. The camera center is at $(0, 0.5)$. (a) The initial estimated geometry is a sphere centered at the camera. (b) The initial geometry is a plane parallel to the qr plane.

exist, indicating that the depth layer $[\hat{a}_i, +\infty)$ can simply be represented by $\hat{a}_{i+\frac{1}{2}}$. In this sense, $\hat{a}^{(\infty)}$ could be viewed as the “infinity depth” in the reflected irradiance field.

Actually, it is more convenient to divide the scene according to the radiometric depth, because the *possible maximum* reconstruction error (recall that Eq. (5) is to control the maximum off-plane error in off-plane relighting) between successive radiometric depth layers is constant. If the optical axis of the camera is around the normal of the qr plane, using radiometric depth for depth sweeping is a good choice. However, if some objects’ radiometric depth is near zero, e.g. for oblique light rays ($\cos \phi \ll 1$), the increment between successive \hat{a}_i ’s is very small, so the discrete search may not be efficient enough. A more efficient way is to start searching from approximate geometry for each pixel. The corresponding geometric depth can be calculated accordingly. In this case, the shape of depth layers depends on the initial geometry.

In this way we can divide the scene into several layers of depth and any point in the scene can be “quantized” to its “representative” depth. Figure 9 illustrates the concept of depth layers. In Fig. 9(a), the initial geometry is a sphere; while in Fig. 9(b), the initial guess is a plane parallel to the qr plane.

6. Results

To verify our analysis, we carried out experiments on both synthetic and real images. Before going through the experimental details, we need to choose a reasonable error tolerance ε .

We may choose the tolerance ε to be the smallest difference between two successive greylevels to achieve signal-level reconstruction. However, images captured by CCD cameras usually contain noise of around 9 greylevels. There also exists noise in synthetic images which are generated by stochastic rendering such as ray tracing with jittered samples (Cook et al., 1984). Moreover, the visual difference of a few greylevels is not apparent. Therefore, we suggest choosing a tolerance of about 16 greylevels. However, one may choose other error tolerance according to his/her own practical requirements.

For synthetic experiments, we only consider the surface radiance and hence parameters of a real camera, f , ΔF and Δt , are not taken into account. Since the mapping from computed radiance and greylevel is linear, the error tolerance ε can be set as

$$\varepsilon = \frac{n_e}{255} I_{\max},$$

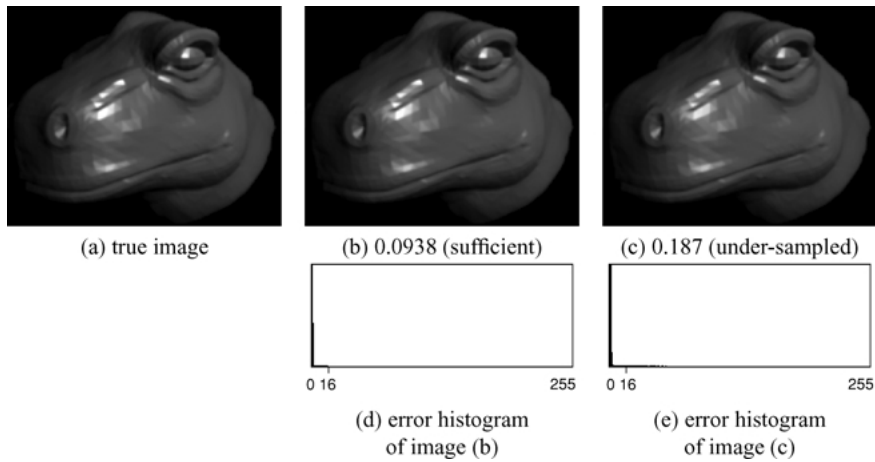


Figure 10. Results of relighting a glossy Lafortune surface. The sampling bound is 0.103. (a) The ground-truth image. (b, d) When the sample spacing is sufficient, the error can be controlled within 16 greylevels. (c, e) When the sample spacing is larger than the sampling bound, the error may exceed 16 greylevels.

where n_e is the tolerable error in greylevel, and I_{\max} is the radiance that corresponds to greylevel 255.

A glossy surface is chosen for testing. It is represented by the Lafortune model (Lafortune, Personal communications).¹ The qr plane is placed at a distance of 2 units from the glossy dinosaur object. In this case, the computed bound of the geometry-independent sampling interval is 0.103 when the error tolerance is 16 greylevels. Figure 10(b) and (c) show the relighting results from images sampled with spacing 0.0938 and 0.187 on the qr plane respectively. Information on sample spacing is also shown at the bottom of each image. The true image is shown on the left (Fig. 10(a)). Figure 10(d) and (e) show the error histograms (histograms of the absolute difference images). The histograms clearly show that the error can be controlled within designated 16 greylevels when the sampling rate is satisfied, and the error exceeds this tolerance when the sampling rate is inadequate.

Our synthetic experiments can also verify the bound in Eq. (5). Unfortunately, as it is only a *sufficient* bound, it only guarantees that the off-plane error is below ε when the depth error does not exceed the bound. It is possible that the off-plane error is still below ε if the depth error is above the bound. However, we found in our experiments that three times of the bound can serve as an exact bound. So we may modify the bound to be three times of its original value. Figure 11(b) and (c) show the relighting results when the novel light source is placed at $-\frac{1}{2}z_{\min} = -0.9015$ behind the qr plane, with depth errors for every pixel being exactly the modified bound and twice the modified bound, respectively.

The error tolerance is again 16 greylevels. The true image is shown in Fig. 11(a). The error histograms again show that the error can be controlled if the depth error is below the modified bound.

Furthermore, the maximum error between Fig. 11(c) and (a) is 22 greylevels when the light source is at $-\frac{1}{2}z_{\min} = -0.9015$. When the depth information does not change but the light source is placed at 0.17, 0.33, 0.56, 0.73, 1.09 (it moves perpendicular to the qr plane), the maximum errors are: 7, 16, 19, 20, and 18 greylevels, respectively. Hence, that the maximum off-plane error occurs at around $-\frac{1}{2}z_{\min}$ is once more verified.

To carry out the experiment on real surfaces, we built a computer-controllable system that precisely moves the point light source over a vertical X-Y table. Figure 12(a) shows the computer-controlled qr plane. A halogen bulb (Fig. 12(c)) is used to simulate a point light source. The experiment setup is shown in Fig. 12(b). The whole capture process took place in a dark room with the halogen bulb as the sole light source, and measures were taken to minimize the interreflection between the object and the environment.

Both the radiometric and geometric calibrations are necessary in real experiments. For radiometric calibration, we applied the algorithm in Debevec and Malik (1997) to recover both the response function of the camera and the high dynamic range radiance maps. The error tolerance ε can be set as:

$$\varepsilon = E_e \cdot E_{\max},$$

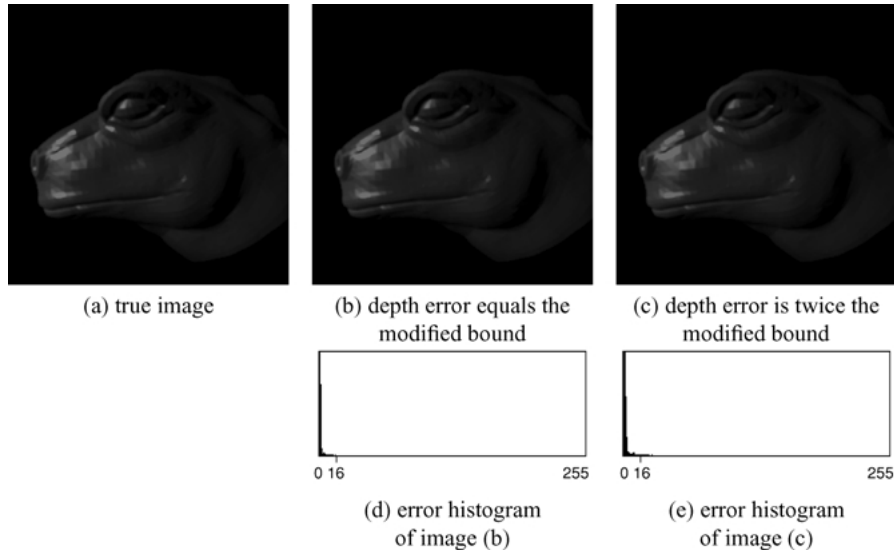


Figure 11. Results of off-plane relighting of a glossy Lafortune surface. (a) The true image. (b, d) When the depth information is sufficient, the error can be controlled with 16 greylevels. (c, e) When the depth error is larger than the modified bound, the error may exceed 16 greylevels.

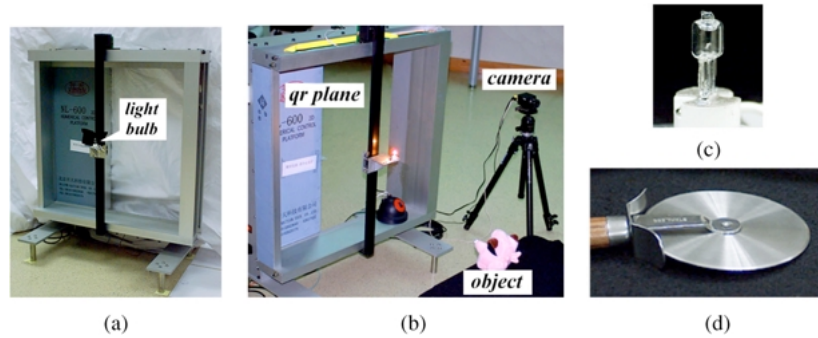


Figure 12. Elements of real experiments. (a) Computer-controlled qr plane. (b) Experiment setup. (c) Halogen bulb. (d) A real surface.

where $E_e \in (0, 1)$ is the tolerable error in exposure that corresponds to tolerable error n_e in greylevel, and is normalized by E_{\max} , and E_{\max} is the exposure that corresponds to greylevel 255.

One may find E_{\max} by shooting the camera at the light source and tuning down both the aperture and the exposure time until the light source produces a greylevel near 255. In this way we only need to record the aperture area and the exposure time, without measuring the focal length because it will be cancelled during computation.

For geometric calibration, we applied the toolbox in Bouguet (2000) to compute the intrinsic parameters and a standard camera calibration algorithm (Zhang, 1999) to compute the homography between the image plane of the camera and the light source plane. Since the camera could not see both the light source and the object due

to the limited FOV, during calibration phase we fixed a thin, light and long straight bar beside the light source so that the camera can see the far-reaching end of the bar. The bar is perpendicular to the qr plane. By taking images of the bar when the light source moves on the qr plane, the transformation matrix can be computed because both the intrinsic parameters of the camera and the positions of the end of the bar are known.

Figure 12(d) shows the real surface we have captured for verification. It is made of steel and is anisotropic. As its geometry is simple, the light source and the viewing directions are known with respect to the normal after the geometric calibration. The Lafortune model of the surface BRDF is fit to the high dynamic range radiance maps by Levenberg-Marquardt method (Press et al., 1997). The surface is by 52 cm away from the qr plane and the geometric independent bound is




Figure 13. Results of relighting steel surfaces. The sampling bound is 1.08 cm. (a) The ground-truth image. (b, e) The sample spacing is just below the sampling bound. (c, d, f, g) The sample spacing is insufficient. Hence the highlight and glare are widened or separated.

found to be 1.08 cm when the error tolerance is 16 greylevels.

Figure 13 shows the result of relighting from images with different sampling intervals. Figure 13(a) is the ground truth. From left to right, the sampling rate is reduced, each time by a factor of 2. As shown in Fig. 13(c) and (d), the highlight band is slightly enlarged and double images of the highlight even appear in (d). Though Fig. 13(b) is synthesized from sufficient samples, the errors of several pixels are still larger than 16 greylevels. This is because of the noise in the captured images and the response function of the camera.

7. Extension to Directional Source

Our analysis can be easily extended to the case when a directional light source is used during scene capture. Instead of determining the maximum allowable sample spacing Δq , we can calculate the maximum allowable sampling angle $\Delta\theta \approx \Delta q/z_{\min}$ (in either elevation or azimuth).

Rewrite Eq. (3) as

$$\begin{aligned} \frac{\Delta q}{z_{\min}} &\leq \frac{1}{z_{\min}} \sqrt{\frac{8\varepsilon}{c_1 B_2}} = \sqrt{\frac{8\varepsilon}{B_2 z_{\min}^2 \cdot I_l f^{-2} \Delta F \Delta t}} \\ &= \sqrt{\frac{8\varepsilon}{\tilde{B}_2 (I_l/z_{\min}^2) c_2}}, \end{aligned}$$

where $\tilde{B}_2 = B_2 z_{\min}^4$.

\tilde{B}_2 is actually the max term in Eq. (2). I_l/z_{\min}^2 is the emitted radiance of the directional light source. We can denote this term by \tilde{I}_l . Making $z_{\min} \rightarrow \infty$, we obtain the angular sample spacing for the directional light source as:

$$\Delta\theta \leq \sqrt{\frac{8\varepsilon}{c_2 \tilde{B}_2}},$$

where $c_2 = \tilde{I}_l f^{-2} \Delta F \Delta t$.

Similarly, we can replace I_l/z^2 in (Eq. (5)) by \tilde{I}_l to determine the error bound when the novel light source is also directional. For relighting with a directional light source, $z \rightarrow \infty$, and hence $|\Delta a| \rightarrow \infty$. This means that the relighting is independent of geometry if *only* parallel light sources are involved, which agrees with the common sense.

8. Conclusions and Future Work

In this paper, we have defined the reflected irradiance field for image-based relighting and studied its sampling and reconstruction problems. The reflected irradiance field is based on a tri-planar representation for the extended plenoptic function that allows the change of both the light source and the viewpoint. For the sampling problem of the reflected irradiance field, we prove that there exists a geometry-independent bound of the

sampling interval, which is analytically related to the BRDF of the scene. For the reconstruction problem, we propose a bound of acceptable depth error and the division of the scene geometry into several depth layers is also presented. Finally, our analysis can also be straightforwardly extended to the case of directional light source. Experiments are conducted to verify our theoretical analysis.

Currently we choose the error tolerance according to a physically-based metric, so we refer to it as the *radiometric tolerance*. In some cases, humans may not notice artifacts even though the radiometric error is large in the reconstructed image. What is interesting and useful for future work is to determine the *photometric tolerance* which takes the response function of human vision into account. Recent work in psychophysical computer graphics (Pattanaik et al., 1998; Ferwerda et al., 1997) provides a psychophysical framework for measuring the sensitivity of human vision to artifacts due to undersampling. Another way to incorporate human perception into our sampling analysis is to analyze the bounds of the psychophysical reflection model that was recently proposed (Pellacini et al., 2000).

In the current work, we compress the reflected irradiance field in the following manner. Firstly, the captured radiance values are rebinned so that all values associated to a pixel are grouped to form an image. Then, each of these images is compressed using 2D DCT. Block-coding technique is also used in Miller et al. (1998). Other coding techniques, such as vector quantization, and wavelet compression could be also used to reduce the data volume. However, random-access and selective decoding must be considered in selecting a coding scheme in order to improve the performance of our relighting system.

Appendix A: Proof of Eq. (2)

By expanding Φ'' , we get,

$$\begin{aligned} \Phi'' = & \frac{1}{d^4} [\rho''_{\ell_1 \ell_1} (m_1 - \ell_1 l_1)^2 \ell_2 \\ & + 2\rho''_{\ell_1 \ell_2} (m_1 - \ell_1 l_1) (n_1 - \ell_2 l_1) \ell_2 \\ & + \rho''_{\ell_2 \ell_2} (n_1 - \ell_2 l_1)^2 \ell_2 + \rho'_{\ell_1} (2m_1 n_1 \\ & - \ell_1 \ell_2 - 2n_1 \ell_1 l_1 - 8m_1 \ell_2 l_1 + 9\ell_1 \ell_2 l_1^2) \\ & + \rho'_{\ell_2} (2n_1^2 - \ell_2^2 - 10n_1 \ell_2 l_1 + 9\ell_2^2 l_1^2) \\ & + \rho(-3\ell_2 - 6n_1 l_1 + 15\ell_2 l_1^2)] \end{aligned}$$

where (m_1, m_2, m_3) , (n_1, n_2, n_3) and (l_1, l_2, l_3) are the coordinates of \vec{M} , \vec{N} and \vec{L} , respectively, and ρ' and ρ'' are 1st-order and 2nd-order partial derivatives respectively.

Direct full search can conveniently produce the coefficients in Eq. (2), though it might be time-consuming. However, mimicking the proof in Section 4.3.1, the coefficients of $\rho''_{\ell_1 \ell_1}$ and ρ can be proven to be exactly z_{\min}^{-4} and $3z_{\min}^{-4}$, respectively.

Appendix B: Proof of Lafortune Bound

Let

$$\begin{aligned} \phi_i &= C_{i,1} \ell_1 v_1 + C_{i,2} \ell_2 v_2 + C_{i,3} \ell_3 v_3, \quad \text{and} \\ \psi_i &= C_{i,1} m_1 v_1 + C_{i,2} n_1 v_2 + C_{i,3} p_1 v_3, \end{aligned}$$

where (m_1, m_2, m_3) and (p_1, p_2, p_3) are the global coordinates of \vec{M} and \vec{P} , respectively, and $\{\vec{N}, \vec{M}, \vec{P}\}$ is the local frame at the surface point.

Then

$$\begin{aligned} & \frac{\partial^2 [\phi_i^{k_i} \ell_2 d^{-2}]}{\partial q^2} \\ &= -(k_i + 3) \phi_i^{k_i} [\ell_2 + 2n_1 l_1 - (k_i + 5) \ell_2 l_1^2] d^{-4} \\ & \quad + 2k_i \psi_i \phi_i^{k_i - 1} [n_1 - (k_i + 3) \ell_2 l_1] d^{-4} \\ & \quad + k_i (k_i - 1) \psi_i^2 \phi_i^{k_i - 2} \ell_2 d^{-4} \\ & \equiv -(k_i + 3) \phi_i^{k_i} G_1 + 2k_i \psi_i \phi_i^{k_i - 1} G_2 \\ & \quad + k_i (k_i - 1) \psi_i^2 \phi_i^{k_i - 2} G_3 \end{aligned}$$

It can be verified that when $k_i \geq 2$,

$$\begin{aligned} |G_1| &= d^{-4} |n_1 [3 - (k_i + 5) l_1^2] l_1 \\ & \quad + n_2 [1 - (k_i + 5) l_1^2] l_2 + n_3 [1 - (k_i + 5) l_1^2] l_3| \\ & \leq h(q) \equiv d^{-4} \sqrt{(k_i + 1)(k_i + 5) l_1^4 - 2(k_i + 1) l_1^2 + 1} \end{aligned}$$

where the Cauchy's inequality is applied. Taking the partial derivative of $h(q)$ w.r.t. q shows that $h(q)$ reaches maximum at q 's such that:

$$\begin{aligned} l_1 &= 0, \quad \text{or} \\ l_1^2 &= \frac{(k_i + 1)(k_i + 10) \pm \sqrt{(k_i + 1)(k_i^3 + 9k_i^2 + 24k_i - 80)}}{6(k_i + 1)(k_i + 5)} \\ &< 0.385. \end{aligned}$$

For $l_1 = 0$, $h(q) \leq (z_{\min})^{-4}$. For the other two l_1 's,

$$\begin{aligned} & (k_i + 1)(k_i + 5)l_1^4 - 2(k_i + 1)l_1^2 + 1 \\ &= \frac{1}{3}[(k_i + 1)(k_i + 4)l_1^2 - k_i] \\ &< \frac{1}{3}[0.385(k_i + 1)(k_i + 4) - k_i] \\ &< [0.359(k_i + 1.223)]^2 \end{aligned}$$

Therefore

$$|G_1| < 0.359(k_i + 1.223)(z_{\min})^{-4}.$$

Similarly,

$$|G_2| < \frac{k_i + 2}{\sqrt{5}}(z_{\min})^{-4} \quad \text{and} \quad |G_3| \leq (z_{\min})^{-4}.$$

Finally, as both $|\phi_i|$ and $|\psi_i|$ cannot exceed $b_i = \max\{|C_{i,1}|, |C_{i,2}|, |C_{i,3}|\}$, we have

$$\left| \frac{\partial^2[\phi_i^{k_i} \ell_2 d^{-2}]}{\partial q^2} \right| < \frac{b_i^{k_i}(2.254k_i^2 + 2.305k_i + 1.317)}{(z_{\min})^4},$$

and thus (4) results.

Appendix C: Proof of Eq. (5)

Suppose the light source L is off-plane (Fig. 14). For ray CA , the exact depth is $a = |CA|$ and the estimated depth is $a' = |CA_1|$. AL and A_1L intersect the qr plane at q and q_1 , respectively. Then the depth error is $\Delta a = a - a'$ and the off-plane error is:

$$\begin{aligned} \mathcal{E} &= \left(\frac{|Aq|}{|AL|} \right)^2 E(q) - \left(\frac{|A_1q_1|}{|A_1L|} \right)^2 E(q_1) \\ &= \left(\frac{1}{1 + \gamma} \right)^2 E(q) - \left(\frac{1 - \tau}{1 + \gamma - \tau} \right)^2 E(q_1) \end{aligned}$$




Figure 14. Incorrect depth information will cause extra error when the novel light source is not on the qr plane.

$$\begin{aligned} &= \left[\left(\frac{1}{1 + \gamma} \right)^2 - \left(\frac{1 - \tau}{1 + \gamma - \tau} \right)^2 \right] E(q_1) \\ &\quad + \left(\frac{1}{1 + \gamma} \right)^2 [E(q) - E(q_1)], \end{aligned}$$

where $\gamma = \frac{|Lq|}{|Aq|}$, $\tau = \frac{\Xi \Delta a}{|Aq_1|}$ and $\Xi = \frac{\vec{V} \cdot \vec{Z}}{L \cdot \vec{Z}}$.

On the other hand,

$$\begin{aligned} E(q) - E(q_1) &= (q - q_1)E'(\tilde{q}), \\ &\quad (\tilde{q} \text{ is between } q \text{ and } q_1) \end{aligned}$$

therefore

$$\begin{aligned} |\mathcal{E}| &= \frac{\gamma}{(1 + \gamma)^2(1 + \gamma - \tau)} \\ &\quad \times \left| \tau \left(2 - \frac{\gamma\tau}{1 + \gamma - \tau} \right) E(q_1) - \Delta a \Gamma E'(\tilde{q}) \right|, \end{aligned} \quad (8)$$

where $\Gamma = \sqrt{\Xi^2 + 1 - 2\Xi(\vec{L} \cdot \vec{C})}$. It can be proven that

$$\Gamma(\vec{L} \cdot \vec{Z}) \leq 1, \quad (9)$$

but we have to omit the proof here.

Since (8) is too complex, it is necessary to simplify it in order to achieve valuable results. When τ is small,

$$\begin{aligned} \frac{\gamma}{(1 + \gamma)^2(1 + \gamma - \tau)} &\approx \frac{\gamma}{(1 + \gamma)^3} \leq \frac{4}{27}, \\ \left| \frac{\tau}{1 + \gamma - \tau} \right| &\leq \frac{1}{\gamma}, \end{aligned}$$

and $|Aq| \approx |Aq_1| \approx |A\tilde{q}|$. Therefore

$$\begin{aligned} |\mathcal{E}| &\leq \frac{4}{27} \left(3|\tau| \frac{B_0}{|Aq_1|^2} + |\Delta a| \Gamma \cdot \frac{B_1}{|A\tilde{q}|^3} \right) \cdot c_1 \\ &\approx \frac{4}{27} \left(3|\tau| \frac{B_0}{|Aq|^2} + |\Delta a| \Gamma \cdot \frac{B_1}{|Aq|^3} \right) \cdot c_1 \\ &= \frac{4c_1}{27|Aq|^3} (3\Xi B_0 + \Gamma B_1) |\Delta a| \\ &= \frac{4c_1(\vec{L} \cdot \vec{Z})^2}{27z^3} [3(\vec{V} \cdot \vec{Z})B_0 + \Gamma(\vec{L} \cdot \vec{Z})B_1] |\Delta a| \\ &\leq \frac{4c_1}{27z^3} [3(\vec{V} \cdot \vec{Z})B_0 + B_1] |\Delta a|, \end{aligned}$$

where B_0 and B_1 are the normal-independent maxima of Φ and Φ' , respectively, or more precisely,

$$B_0 = \text{G.I. max } |d^2\Phi|, \quad B_1 = \text{G.I. max } |d^3\Phi'|.$$

In the last inequality, (9) is applied.

Then, in order that $|\mathcal{E}| \leq \varepsilon$, it is *sufficient* that

$$\frac{4c_1}{27z^3}(3(\vec{V} \cdot \vec{Z})B_0 + B_1)|\Delta a| \leq \varepsilon,$$

or

$$|\Delta a| \leq \frac{27z^3\varepsilon}{4c_1(3B_0 \cos \phi + B_1)},$$

where $\cos \phi = \vec{V} \cdot \vec{Z}$.

Appendix D: Analytic Bounds for B_0 and B_1

For general BRDFs,

$$B_0 = \max\{\rho\}, \text{ and } B_1 \leq \max\{|\rho'_{\ell_1}| + 0.5|\rho'_{\ell_2}| + 2\rho\}.$$

For the Lambertian model,

$$B_0 = \rho_0, \text{ and } B_1 = 2\rho_0.$$

For the Lafortune model,

$$B_0 \leq \sum_i b_i^{k_i}, \text{ and } B_1 \leq \sum_i b_i^{k_i}(1.5k_i + 1).$$

Appendix E: $\hat{a}_{i+\frac{1}{2}}$ Can Represent Every $\hat{a} \in [\hat{a}_i, \hat{a}_{i+1})$

We want to prove that for every $\hat{a} \in [\hat{a}_i, \hat{a}_{i+1})$,

$$\hat{a} - \hat{a}^3 \leq \hat{a}_{i+\frac{1}{2}} \leq \hat{a} + \hat{a}^3.$$

Indeed,

1. if $\hat{a}_i \leq \hat{a} \leq \hat{a}_{i+\frac{1}{2}}$, then

$$\hat{a} - \hat{a}^3 < \hat{a}_{i+\frac{1}{2}} = \hat{a}_i + \hat{a}_i^3 \leq \hat{a} + \hat{a}^3.$$

2. otherwise, if $\hat{a}_{i+1} < \infty$, then it is necessary that $\hat{a}_{i+\frac{1}{2}} \leq \frac{2}{3\sqrt{3}}$, hence from (7), $\hat{a}_{i+1} \leq 1/\sqrt{3}$. Therefore

$$\begin{aligned} \hat{a}_{i+1}^3 - \hat{a}^3 &= (\hat{a}_{i+1} - \hat{a})(\hat{a}_{i+1}^2 + \hat{a}_{i+1}\hat{a} + \hat{a}^2) \\ &\leq (\hat{a}_{i+1} - \hat{a}) \cdot 3\hat{a}_{i+1}^2 \leq \hat{a}_{i+1} - \hat{a}. \end{aligned}$$

$$\text{So } \hat{a} - \hat{a}^3 \leq \hat{a}_{i+1} - \hat{a}_{i+1}^3 = \hat{a}_{i+\frac{1}{2}} < \hat{a} + \hat{a}^3.$$

3. otherwise, $\hat{a}_{i+\frac{1}{2}} > \frac{2}{3\sqrt{3}}$. One may check that

$$x - x^3 \leq \frac{2}{3\sqrt{3}}, \quad \forall x \geq 0.$$

$$\text{Therefore, } \hat{a} - \hat{a}^3 \leq \frac{2}{3\sqrt{3}} < \hat{a}_{i+\frac{1}{2}} < \hat{a} + \hat{a}^3.$$

Acknowledgment

We would like to thank anonymous reviewers for their constructive comments and suggestions and thank Mr. Le Lü for doing the geometric calibration. We also thank Eric. P. F. Lafortune, Cornell Program of Computer Graphics and Q-Panel Laboratory for providing us data at the early stage of this research. Finally, this research is supported in part by Research Grants Council of the Hong Kong SAR, under RGC Earmarked Grants (Project No. CUHK 4186/00E and HKUST6089/99E).

Note

1. Only the data at wavelength 450 nm is used, where $C_{11} = C_{12} = 1.11999$, $C_{13} = 1.01942$, $k_1 = 15.4571$; $C_{21} = C_{22} = 1.08378$, $C_{23} = 0.626672$, $k_2 = 65.2179$; $C_{31} = C_{32} = 1.01529$, $C_{33} = 1.00108$, $k_3 = 195.773$.

References

- Adelson, E.H. and Bergen, J.R. 1991. The plenoptic function and the elements of early vision. In *Computational Models of Visual Processing*, M.S. Landy and J.A. Movshon (Eds.). MIT Press: Cambridge, MA, Ch. 1, pp. 3–20.
- Amanatides, J. 1992. Algorithms for the detection and elimination of specular aliasing. In *Proceedings of Graphics Interface '92*, Canadian Information Processing Society, pp. 86–93.
- Belhumeur, P.N. and Kriegman, D.J. 1996. What is the set of images of an object under all possible lighting conditions? In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR'96)*, pp. 270–277.
- Bouguet, J.-Y. 2000. Camera calibration toolbox for Matlab. available at http://www.vision.caltech.edu/bouguetj/calib_doc/.
- Chai, J.X., Tong, X., Chan, C.-C., and Shum, H.-Y. 2000. Plenoptic sampling. In *Computer Graphics, SIGGRAPH'2000 Proceedings*, ACM SIGGRAPH, pp. 307–318.
- Chen, S.E. 1995. QuickTime VR—an image-based approach to virtual environment navigation. In *Computer Graphics, SIGGRAPH'95 Proceedings*, ACM SIGGRAPH, pp. 29–38.
- Cook, R.L., Porter, T., and Carpenter, L. 1984. Distributed ray tracing. In *Computer Graphics, SIGGRAPH'84 Proceedings*, ACM SIGGRAPH, pp. 137–145.
- Debevec, P.E., Hawkins, T., Tchou, C., Duiker, H.-P., Sarokin, W., and Sagar, M. 2000. Acquiring the reflectance field of a

- human face. In *Computer Graphics, SIGGRAPH'2000 Proceedings*, ACM SIGGRAPH, pp. 145–156.
- Debevec, P.E. and Malik, J. 1997. Recovering high dynamic range radiance maps from photographs. In *Computer Graphics, SIGGRAPH'97 Proceedings*, ACM SIGGRAPH, pp. 369–378.
- Ferwerda, J.A., Pattanaik, S.N., Shirley, P., and Greenberg, D.P. 1997. A model of visual masking for computer graphics. In *Computer Graphics, SIGGRAPH'97 Proceedings*, ACM SIGGRAPH, pp. 143–152.
- Fournier, A., Gunawan, A.S., and Romanzin, C. 1993. Common illumination between real and computer generated scenes. In *Proceedings of Graphics Interface'93*, Toronto, ON, Canada, pp. 254–262.
- Gibson, S. and Murta, A. 2000. Interactive rendering with real-world illumination. In *Eurographics Workshop on Rendering (Rendering Techniques'2000)*, Brno, Czech Republic, pp. 365–376.
- Golub, G. and Loanc, C. van 1989. *Matrix Computations*. The John Hopkins University Press.
- Gortler, S.J., Grzeszczuk, R., Szeliski, R., and Cohen, M.F. 1996. The lumigraph. In *Computer Graphics, SIGGRAPH'96 Proceedings*, ACM SIGGRAPH, pp. 43–54.
- Haeberli, P. 1992. Synthetic lighting for photography. available at <http://www.sgi.com/grafica/synth/index.html>.
- Hanrahan, P. and Krueger, W. 1993. Reflection from layered surfaces due to subsurface scattering. In *Computer Graphics, SIGGRAPH'93 Proceedings*, ACM SIGGRAPH, pp. 43–54.
- Horn, B. 1986. *Robot Vision*. MIT Press.
- Isaksen, A., McMillan, L., and Gortler, S.J. 2000. Dynamically reparameterized light fields. In *Computer Graphics, SIGGRAPH'2000 Proceedings*, ACM SIGGRAPH, pp. 297–306.
- Lafortune, E.P.F. 2000. Personal communications.
- Lafortune, E.P.F., Foo, S.-C., Torrance, K.E., and Greenberg, D.P. 1997. Non-linear approximation of reflectance functions. In *Computer Graphics, SIGGRAPH'97 Proceedings*, ACM SIGGRAPH, pp. 117–126.
- Levoy, M. and Hanrahan, P. 1996. Light field rendering. In *Computer Graphics, SIGGRAPH'96 Proceedings*, ACM SIGGRAPH, pp. 31–42.
- Lin, Z. and Shum, H.-Y. 2000. On the number of samples needed in light field rendering with constant-depth assumption. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR'2000)*, Hilton Head Island, South Carolina, USA, pp. 588–579.
- Lin, Z., Wong, T.-T., and Shum, H.-Y. 2001. Relighting with the reflected irradiance field: Representation, sampling and reconstruction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR'01)*, Kauai, Hawaii, USA, pp. 561–567.
- Loscos, C., Drettakis, G., and Robert, L. 2000. Interactive virtual relighting of real scenes. *IEEE Trans. on Visualization and Computer Graphics*, 6(3):289–305.
- Magda, S., Zickler, T., Kriegman, D., and Belhumeur, P. 2001. Beyond Lambert: Reconstructing surfaces with arbitrary BRDFs. In *Proceedings of the International Conference on Computer Vision (ICCV'01)*, Vancouver, Canada, pp. 391–398.
- McMillan, L. and Bishop, G. 1995. Plenoptic modeling: An image-based rendering system. In *Computer Graphics, SIGGRAPH'95 Proceedings*, ACM SIGGRAPH, pp. 39–46.
- Miller, G.S.P., Rubin, S., and Ponceleon, D. 1998. Lazy decompression of surface light fields for precomputed global illumination. In *Eurographics Workshop on Rendering (Rendering Techniques'98)*, Vienna, Australia, pp. 281–2928.
- Nimeroff, J.S., Simoncelli, E., and Dorsey, J. 1994. Efficient re-rendering of naturally illuminated environments. In *Eurographics Workshop on Rendering (Rendering Techniques'94)*, Darmstadt, Germany, pp. 359–373.
- Oren, M. and Nayar, S.K. 1994. Generalization of Lambert's reflectance model. In *Computer Graphics, SIGGRAPH'94 Proceedings*, ACM SIGGRAPH, pp. 239–246.
- Pattanaik, S.N., Ferwerda, J.A., Fairchild, M.D., and Greenberg, D.P. 1998. A multiscale model of adaptation and spatial vision for realistic image display. In *Computer Graphics, SIGGRAPH'98 Proceedings*, ACM SIGGRAPH, pp. 287–298.
- Pellacini, F., Ferwerda, J.A., and Greenberg, D.P. 2000. Toward a psychophysically-based light reflection model for image synthesis. In *Computer Graphics, SIGGRAPH'2000 Proceedings*, ACM SIGGRAPH, pp. 55–64.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., and Flannery, B.P. 1997. *Numerical Recipes in C*. Cambridge University Press.
- Sato, Y., Wheeler, M.D., and Ikeuchi, K. 1997. Object shape and reflectance modeling from observation. In *Computer Graphics, SIGGRAPH'97 Proceedings*, ACM SIGGRAPH, pp. 379–387.
- Shashua, A. 1997. On photometric issues in 3D visual recognition from a single 2D image. *Int. Journal of Computer Vision*, 21:99–122.
- Shum, H.-Y. and He, L.-W. 1999. Rendering with concentric mosaics. In *Computer Graphics, SIGGRAPH'99 Proceedings*, ACM SIGGRAPH, pp. 299–306.
- Wang, R.H. 1999. *Numerical Approximation (in Chinese)*. Higher Education Press, Beijing.
- Ward, G.J. 1992. Measuring and modeling anisotropic reflection. In *Computer Graphics, SIGGRAPH'92 Proceedings*, ACM SIGGRAPH, pp. 265–272.
- Wong, T.-T. 2001. irraddemo (version 1.0). available at <http://www.cse.cuhk.edu.hk/~ttwong/demo/irraddemo/irraddemo.html>.
- Wong, T.-T., Fu, C.-W., Heng, P.-A., and Leung, C.-S. 2001. The plenoptic illumination function. *IEEE Transactions on Multimedia*, to appear.
- Wong, T.-T., Heng, P.A., Or, S.-H., and Ng, W.-Y. 1997. Image-based rendering with controllable illumination. In *Eurographics Workshop on Rendering (Rendering Techniques'97)*, Saint Etienne, France, pp. 13–22.
- Yu, Y., Debevec, P., Malik, J., and Hawkins, T. 1999. Inverse global illumination: Recovering re-reflectance models of real scenes from photographs. In *Computer Graphics, SIGGRAPH'99 Proceedings*, ACM SIGGRAPH, pp. 215–224.
- Yu, Y. and Malik, J. 1998. Recovering photometric properties of architectural scenes from photographs. In *Computer Graphics, SIGGRAPH'98 Proceedings*. ACM SIGGRAPH, pp. 207–217.
- Zhang, Z. 1998. Modeling geometric structure and illumination variation of a scene from real images. In *Proceedings of the International Conference on Computer Vision (ICCV'98)*, Bombay, India, pp. 1041–1046.
- Zhang, Z. 1999. Flexible camera calibration by viewing a plane from unknown orientations. In *Proceedings of the International Conference on Computer Vision (ICCV'99)*, Kerkyra, Greece, pp. 666–673.