

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

Zhouchen Lin[†] Tien-Tsin Wong[‡] Heung-Yeung Shum[†]
zhoulin@microsoft.com ttwong@acm.org hshum@microsoft.com
[†] Microsoft Research, China [‡] Chinese University of Hong Kong

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 reflection from an object surface irradiated by a point light source that moves on a plane. This representation is dual to that of the light field. 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 point light sources are needed during sampling. 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 the error in the reconstructed image is controlled by a given tolerance, regardless of the geometry. Experiments on both synthetic and real surfaces are conducted to verify our analysis.

1 Introduction

Image-based Modeling and Rendering (IBMR) synthesizes realistic images from pre-recorded images without a complex and long rendering process as in geometry-based computer graphics. 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. Therefore image-based relighting (IBL) receives much attention in recent years. Its goal is to modify the illumination in an interactive fashion while preserving correct visual appearance. This presents an important and challenging problem in IBMR.

There are several approaches for changing the lighting condition. Wong *et al.* [23, 22] proposed to extend the dimensionality of the plenoptic function [1] in order to change the lighting and move the viewpoint. Yu and Malik [24] and Sato *et al.* [20] endeavored to recover the surface BRDF and rough geometry. Alternatively, Debevec *et al.* [6] 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. At the same time, singular value

decomposition was also applied to extract a set of basis images [2, 25], where novel images were obtained by linearly combining the basis images. All the above approaches, however, ignored the issues of sampling and reconstruction error. Though the relit images were visually appealing, their correctness was never ensured.

In this paper, we propose a representation of the plenoptic function called the *reflected irradiance field* for IBL and study its sampling problem. The reflected irradiance field stores the reflection of surface irradiance as an illuminating 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. Although the scene is captured with an illuminating point light source, the image can be relit by any kind of luminaire, including directional light sources.

Like other IBMR problems, IBL is a problem of sampling and reconstructing the plenoptic function. Though for IBMR, the sampling problem of the light field [14, 10] has been studied by Lin and Shum [15] and Chai *et al.* [4], no previous work has addressed the sampling problem of IBL. Clearly, the sampling criterion of relighting and that of the light field are different. Light field rendering is a view interpolation technique. Therefore, its sampling is related to geometry. In contrast, sampling in relighting is more concerned with surface reflectance. In this paper, a theoretical analysis on the sampling problem of the reflected irradiance field is presented. We find that there exists a geometry-independent sampling bound that eliminates artifacts, such as incorrect intensity, position and size of highlights.

The rest of our paper is organized as follows. The definition of reflected irradiance field is presented in Section 2. In Section 3, we present a theoretical analysis on the problem “*how many samples are sufficient for relighting?*” Experimental results are presented in Section 4. Section 5 extends our analysis to the case of directional light sources. Finally, conclusions and future directions are drawn in Section 6.

2 The Reflected Irradiance Field

2.1 Tri-planar Parameterization

Adelson and Bergen proposed a multi-dimensional *plenoptic function* [1] for evaluating low-level human vision models. The function describes environment appearance that our eyes observe. Besides the six parameters for the posi-

tion and direction of the viewpoint and the wavelength of the light, the time parameter models 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 [16, 10, 14] 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.

In general, the illumination component is a high dimensional function. It is necessary to assume a special type of light source in order to reduce the dimension. Usually, directional light sources are used, such as the apparent BRDF of pixels in [23] and the reflectance field in [6]. They are simple and convenient for synthetic images. However, they pose practical difficulties in capturing large object/scene because they can only be 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 do not use directional sources like previous approaches [23, 6]. Instead, we utilize point light sources. Nonetheless, our analysis on sampling can be extended to the directional-source case. The extension is presented in Section 5.

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 [14] and Lumigraph [10].

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 radiance reflected through the window (s, t) and received at a certain (u, v) when the scene is illuminated by a point source positioned at (q, r) . In other words, this six-dimensional representation is parameterized by the light vector (q, r, s, t) and the viewing vector (s, t, u, v) . Here we would like to emphasize that the parameterization is in ray space. The object plane may differ from the scene object and the volume enclosed by the qr and st planes need not be coincident with the one enclosed by the uv and st planes.

2.2 Reflected Irradiance Field

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

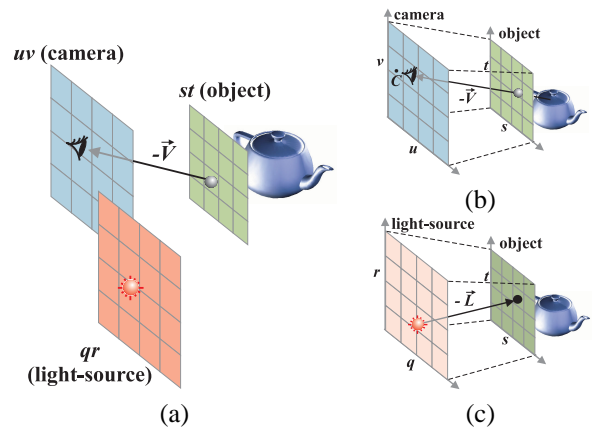


Figure 1: The tri-planar configuration and the dual light slab parameterization. (a) Each captured value represents the radiance reflected through (s, t) and received at a certain (u, v) when the scene is illuminated by a point source positioned at (q, r) . (b) The light field is parameterized by the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

another four-dimensional function (Fig. 1(c)). It seems that any light vector passing through the qr and st planes can be represented by the quadruple (q, r, s, t) . However, this four-dimensional function does *not* directly represent the irradiance incident on object surface, since it is the radiance that is seen by the observer. Instead, it stores the *reflected irradiance* from pixel (s, t) when there exists a point source at (q, r) . Hence, we call this four-dimensional function of (q, r, s, t) the *reflected irradiance field*.

The concept of the reflected irradiance field is close to that of the reflectance field [6]. However, the most general reflectance field is eight dimensional, while the reflected irradiance field is only four dimensional. Moreover, the 8D reflectance field is currently impractical because of the huge data volume and the difficulty of acquiring the incident light field. As a result, in [6] apart from fixed viewpoint, the incident light field has to be reduced to two dimensional and the illuminating light source is directional, which is roughly approximated by a spot light at a mild distance. In the reflected irradiance field, the viewpoint is also fixed but a point light source is used to illuminate an object. In addition, both our representation and analysis can be easily extended to the simplified 4D reflectance field.

2.3 Relighting

The rectilinear structure of reflected irradiance field allows us to simplify the relighting process. 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 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 bi-linear in both st and qr . If the resolution of the output image




Figure 2: (a) A graphical illustration of parameters in the analysis. (b) To reconstruct the exposure when a novel light source is positioned at q , linear interpolation (1D case) is done among neighbors.

is the same as that of the input images, then the interpolation reduces to bilinear interpolation in qr . Real-time relighting can be done by utilizing the bilinear interpolation module in texture-mapping hardware in the same fashion as [10]. Moreover, the image can be relit by various kinds of light sources different from the one used during capturing thanks to the linearity of light transport.

Since the relit image is synthesized from sampled ones, it is inevitable 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 great detail in Section 3. 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 [14, 10]. Incorrect geometry may result in incorrect ray query, incorrect interpolation weights and incorrect intensity attenuation. The property of this off-plane error requires further investigation.

3 Sampling Analysis for Relighting

Obviously, the quality of the interpolated image depends on the sampling rate. In this section, we determine the *maximum allowable sampling interval* on the qr plane such that the intensity error in the relit image is smaller than a user-specified tolerance, hence eliminates noticeable artifacts in the relit images.

3.1 Background

We assume that the camera is pin-hole and its aperture is finite in order to collect enough light. Our analysis focuses on *local* reflection. Refraction, shadowing, interreflection and subsurface scattering [11] are all ignored. The major assumption in our analysis is that the BRDF (ρ) in the scene must be second-order differentiable and its second-order derivatives are bounded. It is a reasonable assumption because a wide range of reflectance models satisfy this assumption, including the Lambertian model, Lafortune model [13], Ward’s anisotropic model [21] and Oren and Nayar’s non-Lambertian model [17].

As the camera records exposure instead of irradiance directly, we move on slightly to exposure. Referring to Fig. 2(a), the exposure at the pixel of A is:

$$E = I_v \Delta t,$$

where I_v is the irradiance at the pixel of A and Δt is the exposure time. I_v is related to the intensity of the light source via [12]:

$$I_v = \cos^4 \alpha \Delta F f^{-2} I_r, \quad I_r = \rho(\xi, \zeta; \mu, \nu) \zeta d^{-2} I_l,$$

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 radiance of the light source, and (ξ, ζ) and (μ, ν) are the local directional coordinates of \vec{L} and \vec{V} w.r.t. local frame $\{\vec{N}, \vec{M}, \vec{P}\}$ at A , respectively.

Hence

$$E = \Phi(\xi, \zeta; \mu, \nu) \kappa_1,$$

where

$$\Phi(\xi, \zeta; \mu, \nu) = \rho(\xi, \zeta; \mu, \nu) \zeta d^{-2}, \quad \kappa_1 = I_l \Delta F \Delta t f^{-2}.$$

$\cos^4 \alpha$ disappears because most cameras are designed to eliminate the diminishing artifact introduced by it.

As κ_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.

3.2 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, we assume that bilinear interpolation is 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. 2(b)). Then the bilinear interpolation becomes linear.

Given an error tolerance ε , we want to find the maximum allowable spacing between samples along the q -axis such that the difference between every reconstructed exposure and the true exposure is smaller than ε . Obviously E is not band limited. 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

$$|\lambda E(q_1) + (1 - \lambda)E(q_2) - E(q)| \leq \varepsilon, \quad \forall A, \text{ and } \forall (q_1, q_2, q) \\ \text{satisfying } 0 \leq q_2 - q_1 \leq \Delta q \text{ and } q_1 \leq q \leq q_2,$$

where $\lambda = |q_1 q| / |q_1 q_2|$.




Figure 3: An exposure curve can be plotted for each pixel in the image as the light source moves along q . Curve a is for the pixel marked ‘ a ’, while curve b is for the pixel ‘ b ’.

3.3 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 without knowing any geometric information, neither depth nor normal. This is important as surface reflectance depends heavily on the normal and acquiring accurate normal is even harder than acquiring depth.

For a pixel of interest, we can plot the curve of its exposure as a function of q as the light source moves along the q -axis. Fig. 3 shows the plots for two specific pixels in the left image. The difference in the shape of the curve is due to the surface normal orientation of the surface element visible through the pixel, and the distance from this surface element to the light source. 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. 3, the optimal sample spacing for pixel ‘ a ’ is Δq_1 , while it is Δq_2 for pixel ‘ b ’.

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. Fig. 4 illustrates this idea graphically. All we need to know is the BRDF and the minimum distance z_{\min} (see Fig. 2(b)) 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, there exists $\tilde{q} \in [q_1, q_2]$, such that,

$$\mathcal{E} \equiv \lambda E(q_1) + (1-\lambda)E(q_2) - E(q) = \frac{1}{2}(q-q_1)(q_2-q)E''_{qq}(\tilde{q}),$$

where f''_{ij} stands for $\frac{\partial^2 f}{\partial i \partial j}$ (similar notations are used in the sequel for simplicity). Therefore,

$$|\mathcal{E}| \leq \frac{1}{2} \left[\frac{(q-q_1)+(q_2-q)}{2} \right]^2 \kappa_1 B_2 = \frac{1}{8}(\Delta q)^2 \kappa_1 B_2, \quad (1)$$

¹ \tilde{P} is not drawn since it is perpendicular to the page.




Figure 4: Since the surface element (indicated by spheres) can be located anywhere (indicated by the color of the spheres, where darker color means larger depth) behind z_{\min} (see Fig. 2(b)) and may have any normal orientation (indicated by the arrows), the minimum sample spacing Δq_{\min} is found among all possible normal orientations and all possible depths.

where

$$B_2 = \text{G.I.} \max |\Phi''_{qq}(q)|,$$

or the geometry-independent upper bound of $|\Phi''_{qq}(q)|$. In other words, B_2 is the maximum value of $|\Phi''_{qq}(q)|$ for all possible positions of the light source and all possible positions and orientations of A .

For a general BRDF, we can prove, by a hybrid method of numerical computation and analytical deduction, that if the second-order derivatives of the BRDF are bounded, then:

$$B_2 < (z_{\min})^{-4} \max_{\xi, \zeta, \mu, \nu} \{ |\rho''_{\xi\xi}| + 0.77|\rho'_{\xi\xi}| + 0.385|\rho''_{\zeta\xi}| + 2.046|\rho'_{\zeta\xi}| + 2.098|\rho'_{\zeta\xi}| + 3\rho \}, \quad (2)$$

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. Due to space limitations, we omit the details of the proof here. However, the complete proof for the Lambertian case is given below.

From Equations 1 and 2, we know that the geometry-independent upper bound of the sampling interval exists. Unfortunately, the analytic bound provided by the right hand side of Equation 2 is usually much larger than B_2 , 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 full search to find B_2 . One trick is that searching along depth is unnecessary because B_2 must be attained when the depth of the surface element is z_{\min} .

With B_2 computed and by making the right hand side of Equation 1 less than ϵ , the *sufficient* bound for Δq (sample spacing) is:

$$\Delta q \leq \sqrt{\frac{8\epsilon}{\kappa_1 B_2}}. \quad (3)$$

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, namely Lambertian and Lafortune models.



(a) 0.0938 (sufficient) (b) true image (c) 0.187 (insufficient)

Figure 5: Results of relighting a glossy Lafortune surface. The sampling bound is 0.103. (a) When the sample spacing is sufficient, no observable error can be detected. (b) The true image. (c) When the sample spacing is larger than the sampling bound, the reconstruction error is apparent. Note that the highlight is blurry and dim.

3.3.1 Lambertian Bound

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

$$\Phi''_{qq} = (3/\pi)d^{-4}(-\zeta - 2n_1l_1 + 5\zeta l_1^2),$$

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

Since

$$\zeta + 2n_1l_1 - 5\zeta l_1^2 = n_1(3 - 5l_1^2)l_1 + n_2(1 - 5l_1^2)l_2 + n_3(1 - 5l_1^2)l_3,$$

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

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

Therefore

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

Taking the partial derivative of $h(q)$ shows that $\max_q \{h(q)\} = h(q)|_{l_1=0}$, hence

$$B_2 \leq (3/\pi)(z_{\min})^{-4}.$$

It can be seen that for some special surface elements the equality holds.

3.3.2 Lafortune Bound

Another model of interest is the Lafortune model [13] 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,

$$\rho(\xi, \zeta; \mu, \nu) = \sum_i (C_{i,1}\xi\mu + C_{i,2}\zeta\nu + C_{i,3}\eta\omega)^{k_i},$$

where (ξ, ζ, η) and (μ, ν, ω) are the local coordinates of the normalized vectors \vec{L} and \vec{V} , respectively, i is the index of cosine lobe, and $C_{i,1}$, $C_{i,2}$, $C_{i,3}$ and k_i are the parameters defining the nature of surface reflectance.

An analytic bound of B_2 in the case of the Lafortune model is:

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

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

The basic idea for deriving the Lafortune bound is the same as that for the Lambertian case, but the details are more complex and are skipped due to page limit.

4 Experimental 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 . It can be chosen to be the smallest exposure 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 generated by renderers utilizing stochastic sampling [5]. Moreover, the visual difference of a few greylevels is not apparent. Therefore, we suggest choosing a tolerance of about 16 greylevels.

For synthetic experiments, we only consider the surface radiance and hence those 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:

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

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. 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. Fig. 5(a) 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 in the middle (Fig. 5(b)) for better visual comparison. It can be seen that there is no observable difference between the sufficiently sampled one in Fig. 5(a) and the true image, while the image in Fig. 5(c) exhibits blurry and dimmed highlights. The blurriness and dimness are due to the incorrect blending of reference images with spatially-different highlights.

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. The experiment setup is shown in Fig. 6(a) where the left part is the computer-controlled qr plane and a halogen bulb is used to simulate a point light source. 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. Moreover, the camera aperture and exposure time are carefully tuned in order to minimize the saturation.

Both radiometric and geometric calibration are necessary in real experiments. For radiometric calibration, we applied the algorithm in [7] to recover the response function of the camera. The error tolerance ε can be set as:

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

where $E_e \in (0, 1)$ is the tolerable error in exposure that corresponds to tolerable error n_e in greylevel, and is normal-




Figure 7: Results of relighting steel surfaces. The sampling bound is 1.01 cm. Upper row: Surface 1. Lower row: Surface 2. (a) & (e): The sample spacing is just below the sampling bound. Note that highlight and glare are well preserved. (b) to (d) & (f) to (h): The sample spacing is insufficient. Hence the highlight and glare are separated.




Figure 6: Elements of real experiments. (a) Experiment setup. (b) Real surfaces tested in the experiment.

ized by E_{\max} . Exposure E_{\max} corresponds to the greylevel 255. It can be found 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 and the exposure time, without measuring the focal length because it will be cancelled during computation.

For geometric calibration, we applied the toolbox in [3] to compute the intrinsic parameters and the algorithm described in [8] to compute the transform between the coordinates of the camera and the light source plane. Since the FOV of the camera confined the camera to viewing both the light source and the object, during calibration phase we fixed a thin, light and long straight bar beside the light source and perpendicular to the qr plane so that the camera can see its far-reaching end. By taking images of the bar when the light source was moved on the qr plane, the transformation matrix could be computed because both the intrinsic parameters of the camera and the positions of the end of the bar were known.

Fig. 6(b) shows the two real surfaces we have captured for verification. Both surfaces 1 and 2 were made of steel and are anisotropic. They were both by 52cm away from the

qr plane. With the Lafortune coefficients we found for the surface BRDF, the geometry-independent bound is 1.01cm when the error tolerance is 16 greylevels.

Fig. 7 shows the result of relighting from image sets with different sampling intervals. Results of surfaces 1 and 2 are shown in the upper and lower rows respectively. The leftmost images (Fig. 7(a) & (e)) are the results of sufficiently sampled image sets. From left to right, the sampling rate is reduced, each time by a factor of 2. As shown in Fig. 7(b) to (d), the highlight band incorrectly enlarges and double images of the highlight appear in (c) & (d) due to the blending of undersampled images. Moreover, the glare at the edge (left hand side) of the disk becomes separated except for the sufficiently sampled one in Fig. 7(a). Similarly, from Fig. 7(e) to (h), double images of the highlight appears when the scene is undersampled. Note that the glare due to highlight is also separated in Fig. 7(f) to (h).

5 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 latitudinal or longitudinal direction).

Rewrite Equation 3 as:

$$\frac{\Delta q}{z_{\min}} \leq \sqrt{\frac{8\epsilon}{B_2 z_{\min}^2 \cdot I_l f^{-2} \Delta F \Delta t}} = \sqrt{\frac{8\epsilon}{\tilde{B}_2 (I_l / z_{\min}^2) \kappa_2}},$$

where $\tilde{B}_2 = B_2 z_{\min}^4$. In fact, \tilde{B}_2 is the max term in Equation 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:

$$\Delta\theta \leq \sqrt{\frac{8\epsilon}{\kappa_2 \tilde{B}_2}},$$

where $\kappa_2 = \tilde{I}_l f^{-2} \Delta F \Delta t$.

6 Conclusions and Future Work

In this paper, we have defined the reflected irradiance field for image-based relighting and studied its sampling problem. 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 its sampling problem, we prove that there exists a geometry-independent bound of the sampling interval, which is analytically related to the BRDF of the scene. Our analysis can be straightforwardly extended to the case of directional light sources. 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 [18, 9] 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 [19].

In the current work, we have not embarked on the compression problem of the reflected irradiance field. In our experiments, all the data are stored in uncompressed form and the volumes vary from 240MB to 320MB. Some compression techniques, such as vector quantization, JPEG, and wavelet transform, can be utilized to reduce the data volume. Random-access and selective decoding must be considered in selecting compression schemes to improve the performance of our relighting system.

Acknowledgment

The authors would like to thank Mr. Le Lü for doing the geometric calibration. This research is supported in part by Research Grants Council of the Hong Kong SAR, under RGC Earmarked Grants (Project No. CUHK 4186/00E).

References

- [1] E. H. Adelson and J. R. Bergen. The plenoptic function and the elements of early vision. In *Computational Models of Visual Processing*, chapter 1, pages 3–20. MIT Press, 1991.
- [2] P. N. Belhumeur and D. J. Kriegman. What is the set of images of an object under all possible lighting conditions. In *IEEE CVPR'96*, 1996.
- [3] J.-Y. Bouguet. Camera calibration toolbox for matlab. available at <http://www.vision.caltech.edu/bouguetj/calibdoc/>.
- [4] J.-X. Chai, X. Tong, C.-C. Chan, and H.-Y. Shum. Plenoptic sampling. In *SIGGRAPH 2000 Proceedings*, pages 307–318, July 2000.
- [5] R. L. Cook, T. Porter, and L. Carpenter. Distributed ray tracing. In *SIGGRAPH '84 Proceedings*, pages 137–145, 1984.
- [6] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *SIGGRAPH 2000 Proceedings*, July 2000.
- [7] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH'97 Proceedings*, pages 369–378, August 1997.
- [8] Olivier Faugeras. *Three-Dimensional Computer Vision: a Geometric Viewpoint*. MIT Press, 1993.
- [9] J. A. Ferwerda, S. N. Pattanaik, P. Shirley, and D. P. Greenberg. A model of visual masking for computer graphics. In *SIGGRAPH'97 Proceedings*, pages 143–152, 1997.
- [10] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen. The lumigraph. In *SIGGRAPH'96 Proceedings*, pages 43–54, August 1996.
- [11] P. Hanrahan and W. Krueger. Reflection from layered surfaces due to subsurface scattering. *SIGGRAPH '93 Proceedings*, pages 165–174, August 1993.
- [12] B. Horn. *Robot Vision*. MIT Press, 1986.
- [13] E. P.F. Lafortune, S.-C. Foo, K. E. Torrance, and D. P. Greenberg. Non-linear approximation of reflectance functions. In *SIGGRAPH'97 Proceedings*, July 1997.
- [14] M. Levoy and P. Hanrahan. Light field rendering. In *SIGGRAPH'96 Proceedings*, pages 31–42, August 1996.
- [15] Z. Lin and H.-Y. Shum. On the number of samples needed in light field rendering with constant-depth assumption. In *CVPR 2000*, pages 588–579, June 2000.
- [16] L. McMillan and G. Bishop. Plenoptic modeling: An image-based rendering system. In *SIGGRAPH'95 Proceedings*, pages 39–46, August 1995.
- [17] M. Oren and S. K. Nayar. Generalization of lambert's reflectance model. In *SIGGRAPH'94 Proceedings*, 1994.
- [18] S. N. Pattanaik, J. A. Ferwerda, M. D. Fairchild, and D. P. Greenberg. A multiscale model of adaptation and spatial vision for realistic image display. In *SIGGRAPH'98 Proceedings*, pages 287–298, 1998.
- [19] F. Pellacini, J. A. Ferwerda, and D. P. Greenberg. Toward a psychophysically-based light reflection model for image synthesis. In *SIGGRAPH 2000 Proceedings*, pages 55–64, 2000.
- [20] Y. Sato, M. D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *SIGGRAPH'97 Proceedings*, pages 379–387, August 1997.
- [21] G. J. Ward. Measuring and modeling anisotropic reflection. In *Computer Graphics Proceedings, Annual Conference Series, SIGGRAPH'92*, pages 265–272, July 1992.
- [22] T.-T. Wong, C.-W. Fu, and P.-A. Heng. Interactive relighting of panoramas. *IEEE Computer Graphics & Applications*, 21(2):32–41, March-April 2001.
- [23] T.-T. Wong, P.-A. Heng, S.-H. Or, and W.-Y. Ng. Image-based rendering with controllable illumination. In *Eighth Eurographics Workshop on Rendering 1997*, pages 13–22, Saint Etienne, France, June 1997.
- [24] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: recovering reflectance models of real scenes from photographs. In *SIGGRAPH 99 Proceedings*, July 1999.
- [25] Z. Zhang. Modeling geometric structure and illumination variation of a scene from real images. In *Proceedings of ICCV'98*, Bombay, India, January 1998.