



Aalborg Universitet

AALBORG UNIVERSITY
DENMARK

A Comparison of Techniques for Approximating Full Image-Based Lighting

Madsen, Claus B.; Laursen, Rune Elmggaard

Published in:

Proceedings: 15th Danish Conference on Pattern Recognition and Image Analysis, Copenhagen, Denmark

Publication date:

2006

Document Version

Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Madsen, C. B., & Laursen, R. E. (2006). A Comparison of Techniques for Approximating Full Image-Based Lighting. In S. I. Olsen (Ed.), Proceedings: 15th Danish Conference on Pattern Recognition and Image Analysis, Copenhagen, Denmark (pp. 53-60)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- ? Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- ? You may not further distribute the material or use it for any profit-making activity or commercial gain
- ? You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

A Comparison of Techniques for Approximating Full Image-Based Lighting

Claus B. Madsen and Rune E. Laursen
Laboratory of Computer Vision and Media Technology
Aalborg University, Aalborg, Denmark
cbm@cvmt.aau.dk
www.cospe.dk

Abstract

Image-Based Lighting (IBL) has become a very popular approach in computer graphics, especially for special effects, such as insertion of virtual objects into real imagery (Augmented Reality). In essence IBL is based on capturing the illumination conditions in a scene in an omnidirectional image, such that the image describes intending radiance at a point from all directions. Such an omnidirectional measurement of radiance is typically called a light probe image. Using the illumination information from such an image virtual objects can be rendered with consistent shading including global illumination effects such as color bleeding.

Rendering with light probe illumination is extremely time consuming. Therefore a range of techniques exist for approximating the intending radiance described in a light probe image by a finite number of directional light sources. We describe 4 such techniques from the literature and perform a comparative evaluation of them in terms of how well they each approximate the final irradiance as a function of how many sources they are allowed to use in the approximation. We demonstrate that for relatively low numbers of sources (e.g., less than 100 sources) one particular method performs significantly better than the three other techniques.

1 Introduction

Image-based approaches have gained widespread popularity in computer graphics because of the inherent problems with purely model-based approaches, [12]. Image-based techniques have been used for 3D modeling of real scenes (Image-Based Modeling), for rendering from a bank of images with no 3D model whatsoever (Image-Based Rendering), and for modeling the complex illumination conditions in real scenes (Image-Based Lighting). The latter has especially been applied for special effects, i.e. rendering virtual objects into real imagery for movies, commercials, etc.

Image-Based Lighting (IBL) has become an extremely frequently used technique since it is an intuitively simple

manner allows you to render a virtual object with illumination conditions that perfectly match those of a real scene. The idea is simply that you use a camera to measure the light arriving at some point in the scene, the point at which you want to insert a virtual object. In practice people most often use a polished steel ball, place it somewhere in a scene, and take an image of it with a tele-lens from some distance away. After cropping away everything which is not a projection of a point on the sphere the image now contains information about how much light arrives at the position of the ball from all possible directions. In other words the ball image contains information about the incident radiance (field radiance) at the ball location. Figure 1 illustrates this concept.

Using reflective spheres for acquisition of light probes is the standard approach and employed by most researchers in the field for its simplicity. Most of the light probes shown in this paper have been downloaded from Debevec's probe gallery, [2], and have been acquired by merging two views of a reflective sphere in order to avoid the reflections of the camera and the photographer in the final light probe, and in order to avoid the problem with a small "blind region" behind the sphere.

When we acquire our own light probes we use a Sigma 8mm 180 degree field-of-view fish eye lens. By taking two such images in opposing direction we can merge them together to a complete spherical image using the HDR-Shop program, [7]. With this approach we get much higher resolution light probes, and avoid the smearing of detail that an imperfect mirror ball can result in. Figure 2 shows an example of two such hemi-spherical images.

Regardless of whether the light probe is acquired with a reflective sphere or with multiple views with a fish-eye lens it is still very important to handle the dynamic range of the light in the scene. This problem is handled by acquiring the same view at multiple different exposures, gradually lowering the exposure time until no pixels in the image are saturated. These multiple exposure are then fused into a single High Dynamic Range (HDR) floating point image, [6].

Once a light probe has been acquired at some position in some scene it can be used for many purposes. The light probe is a map of the incident radiance at the acquisition



Figure 1: Left: the light probe image is based on a cropped image of a reflective sphere. Right: light probe images are omni-directional, i.e., cover the entire sphere around the light probe position. Here we have remapped the light probe image to longitude-latitude format, where the full 360 degrees are represented along the horizontal (longitude) axis, and 180 degrees are represented along the vertical (latitude) axis. For construction and remapping of light probe images we use HDRShop 2.09, [7].



Figure 2: Left and center: two hemi-spherical images acquired with a Sigma 180 degree field-of-view fish-eye lens in opposing directions. Right: the two semi-spherical images merged and mapped as a longitude-latitude light probe using HDRShop 2.09, [7].

point. Each pixel in the map corresponds to a certain direction and solid angle, and together all pixels cover the entire sphere around the acquisition point. A light probe can thus also be called a radiance map, or an environment map. With this information virtual objects can be rendered into the scene with scenario consistent illumination e.g., [3, 4, 8]. Light probes can also be used to estimate the reflectance distribution functions of surfaces from images, as demonstrated in [15, 14]. For a review of illumination models in mixed reality see [9].

Actually using light probes for rendering is computationally very heavy. Using image-based lighting for a full global illumination rendering with path tracing is extremely time consuming in order to reduce the noise level in the final rendering, simply because the light probe has to be treated as a spherical area light source enclosing the entire scene. To get a noise free estimate of the irradiance at a certain point requires thousands and thousands of samples of this area source.

To combat this problem several approaches have been proposed which take a light probe and attempt to approxi-

mate its illumination by a relatively low number of directional light sources. That is, the idea of these approaches is to find directions and the radiances of some number, say 64, directional light sources, such that the combined illumination from these sources approximate the combined illumination from the entire light probe.

With such a directional light source approximation to a light probe, Image-Based Lighting using light probes can also be implemented in real-time applications taking into account that each source causes shadows to be cast.

The aim of the present paper is simply to test the performance of these approximation techniques in terms of how well they actually approximate the light probe for a given number of sources.

The paper is organized as follows. First section 2 gives a very brief overview of the approach and results in the paper. As a foundation for subsequent technique descriptions section 3 introduces concepts and terminology. In section 4 we give a brief description of four different approaches to using N directional light sources to approximate the illumination in a light probe image. Section 5

then tests these four techniques in terms of their relationships between approximation error and number of sources used. Conclusions and directions for future research are given in section 6.

2 Overview of the idea of this work

Figure 3 shows a light probe together with the result from one of the approximation techniques we study in this paper. In this particular case the technique has been allocated 8 directional sources which it has then distributed across the light probe longitude-latitude map in an attempt to capture the radiance distribution of the original light probe. Naturally, the accuracy of the approximation depends on the number of sources allocated. The original light probe is simply W times H directional sources, where W is the number of pixels in the longitude direction, and H is the number of pixels in the latitude direction.

We then run any given technique on some light probe image to produce approximations with 2, 4, 8, 16, etc. light sources. Given these sets of approximated sources we compute what the resulting error in irradiance is compared to ground truth, which in this case is the irradiance computed by using the radiance from all pixels in the light probe.

Before we proceed with the actual techniques and their performances we need to establish a small theoretical basis.

3 Terminology

Above we have used the term “directional light source” a few times. A directional light source is in reality identical to an infinitely distant point light source. As the radiometric property radiance, L , is measured in $W/(m^2 \cdot Sr)$ it is somewhat meaningless to speak about the radiance of a point light source, let alone of a directional light, since they have no area.

Phar and Humphreys, [13], have no trouble speaking about the radiance of directional light sources, though. They even go so far as stating that the energy received by the scene is LA , where L is the radiance of the source, and A is the total area of the scene which receives light from the source. At best this “energy” is measured in W/Sr . So, there simply is no physically correct way to treat directional light sources.

Point light sources can be treated by considering source’s radiant power, Φ (in W), and then distributing this energy evenly (assuming an isotropic point source) across the surface of a sphere with radius equal to the distance from the scene to the source using $A = 4\pi r^2$. The “flux density” at the receiving surface would then be, $E = \cos(\theta)\Phi/A$, where θ is the angle between surface

normal and direction vector to point light, [10], (note that Jensen’s book has an error regarding received irradiance from a point source, an error which has been acknowledged in the list of corrections on his homepage). This flux density is an irradiance measure in W/m^2 .

But using the point source concept requires a known, finite distance to the source. The very idea behind rendering with Image-Based Lighting using light probes is that the environment captured in the light probe is assumed to be infinitely distant, [3]. Therefore we will stick with the concept of directional light sources, and the purpose of the remainder of this section is to establish a physically correct terminology.

3.1 The radiometry of light probes

In formal terms we can say that the light probe image is a spatially discrete measurement of the continuous function describing the incident radiance (measured in $W/(m^2 \cdot Sr)$), which in turn is a function of the incident direction. Let \vec{n} be the normal of a differential area surface, and let $\Omega_{\vec{n}}$ be the hemi-sphere defined by this normal. By integrating the incident radiance, $L(\vec{\omega})$, from the direction $\vec{\omega}$ over the hemi-sphere the total irradiance, $E(\vec{\omega})$, can be computed:

$$E(\vec{n}) = \int_{\Omega} L(\vec{\omega})(\vec{n} \cdot \vec{\omega})d\vec{\omega} \quad (1)$$

which then is measured in W/m^2 . The term $d\vec{\omega}$ signifies the differential solid angle $d\omega = |d\vec{\omega}|$ in the direction $\frac{d\vec{\omega}}{|d\vec{\omega}|}$.

For computational purposes it is beneficial to formulate these matters in terms of standard spherical coordinates. A direction in space is then written as $\vec{\omega}(\theta, \phi) = [\sin(\theta) \cos(\phi), \sin(\theta) \sin(\phi), \cos(\theta)]$, where θ is the angle the direction vector makes with the coordinate system z -axis (latitude), and ϕ is the angle the projection of the vector on the xy -plane makes with the x -axis. The irradiance from Eq. 1 then becomes:

$$E(\vec{n}) = \int \int_{(\theta, \phi) \in \Omega_{\vec{n}}} L(\theta, \phi)(\vec{n} \cdot \vec{\omega}(\theta, \phi)) \sin(\theta)d\theta d\phi \quad (2)$$

In this paper we will exclusively use the latitude-longitude mapping (LL mapping) of light probe images. Let the resolution of the LL light probe image be W by H pixels, and let u and v represent pixel coordinates in an image coordinate system with origin in the top left corner of the LL map, and v -axis oriented downwards. Thus the *middle row* in the image corresponds to the equator of the unit sphere, i.e, corresponds to $\theta = \pi/2$, the top row corresponds to $\theta = 0$ and the bottom row corresponds to $\theta = \pi$. Moreover $\phi = 0$ corresponds to the *leftmost column*. Each light probe pixel, $P(u, v)$, represents the radiance in $W/(m^2 \cdot Sr)$ (if the light probe acquisition is radiometrically calibrated) from the direction



Figure 3: Result from running the Median Cut approximation technique using 8 directional sources on Galileo’s Tomb light probe. The light probe is obtained from [2]. Each rectangular region contains a red dot. This red dot marks the chosen direction for a particular directional source, and all the combined radiance from the region has been transferred to this particular source direction.

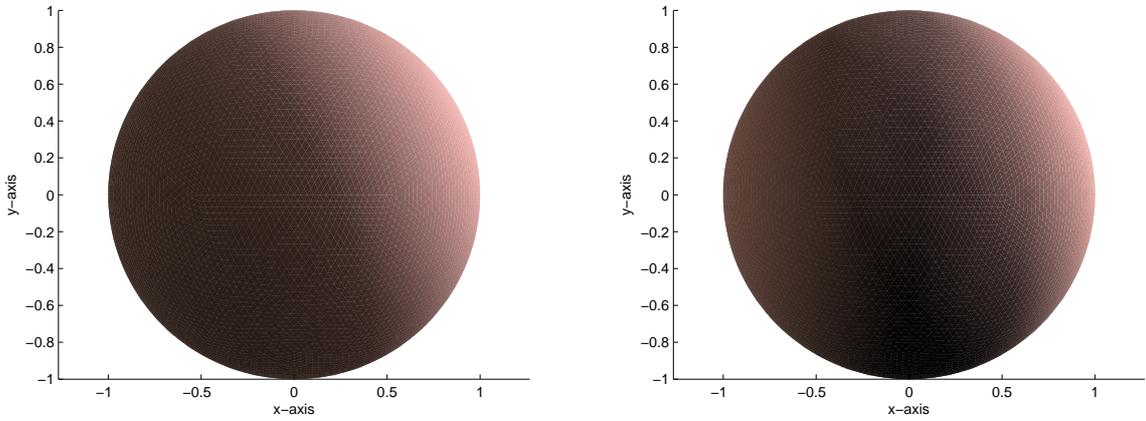


Figure 4: Left: ground truth irradiance for around 20000 normals distributed evenly on a sphere computed for the Galileo’s Tomb light probe, figure 3. Right: irradiances resulting from running the Median Cut source approximation technique to produce 8 directional sources. On print the difference may be visually subtle, but the average error is actually around 25 percent, and the maximum error is more than 75 percent.

given by $\vec{\omega}(u, v) = \vec{\omega}(\theta(v), \phi(u))$, where $\theta(v) = v\Delta_\theta$ and $\phi(u) = u\Delta_\phi$, where $\Delta_\theta = \pi/H$ and $\Delta_\phi = 2\pi/W$. The discrete version of Eq. 2 then becomes:

$$E(\vec{n}) \approx \sum_u \sum_v P(u, v)(\vec{n} \cdot \vec{\omega}(v, u)) \sin(\theta(v)) \Delta_\theta \Delta_\phi \quad (3)$$

where the summations are subject to the constraint that $(\theta(v), \phi(u)) \in \Omega_{\vec{n}}$, i.e., that the combinations of u and v represent pixels inside the region corresponding to the hemi-sphere defined by the surface normal \vec{n} .

From Eq. 3 it is evident that if every pixel, $P(u, v)$, in the LL map is scaled with $\Delta_\theta \cdot \Delta_\phi = 2\pi^2/(W \cdot H)$ and weighted by $\sin(\theta(v))$, we get a very simple summation. We therefore produce a new LL map, where each pixel $Q(u, v) = 2\pi^2 P(u, v) \sin(\theta(v))/(W \cdot H)$. The irradiance for a given normal is then simply computed as:

$$E(\vec{n}) \approx \sum_u \sum_v Q(u, v)(\vec{n} \cdot \vec{\omega}(v, u)) \quad (4)$$

where the summations again are subject to the constraint that $(\theta(v), \phi(u)) \in \Omega_{\vec{n}}$.

To recapitulate in a different way: Each pixel in the LL map acts as a small area light source subtending a solid angle of $A_p = 2\pi^2/(W \cdot H)$ [Sr/pixel]. By weighting each pixel by $\sin(\theta(v))$ we achieve “permission” to treat all pixels equally in the sense that we cancel out the effect of the non-uniform sampling density of the LL mapping (poles are severely over-sampled). By subsequently scaling by A_p we convert the solid angle domain from steradians to pixels. I.e., each $Q(u, v) = 2\pi^2 P(u, v) \sin(\theta(v))/(W \cdot H)$ measures the radiance in $W/(m^2 \cdot pixel)$, such that by performing a simple cosine weighted sum of pixels we directly get the irradiance contributed by the pixels involved in the sum (Eq. 4). Another way of putting it is: each pixel $Q(u, v)$ is an area light source contributing $Q(u, v)(\vec{n} \cdot \vec{\omega}(v, u))$ irradiance to the differential area surface with normal \vec{n} .

3.2 Physically correct directional light sources

As described above there is no such thing as a physically correct directional light source. We will use the term extensively in the remainder of the paper, though, for simplicity. In this context we take the meaning of a directional light source to be a very small area light source (there are normally a lot of pixels in a light probe image). The direction to such a source is taken to be the direction to its center, and it is assumed that for each such source we know its radiance and its area.

4 Light probe approximation techniques

As mentioned previously there exist a number of approaches to finding a set of directional light sources which approximate a full radiance map in the form of a light probe. We have found four different techniques, three of which are closely related and operate directly in the radiance space image domain of the light probe, in particular on the longitude-latitude mapping. The last technique is quite different in that it operates in *irradiance* space.

Below we briefly describe the four different techniques, starting with the radiance space techniques. The common idea behind the radiance space techniques is that they divide the light probe LL map image into a number of regions. Subsequently the sum of all the radiance within a region is transferred to the centroid pixel of the region.

4.1 Lightgen

A description of this technique will be included in a future version of this paper. The normal citation given for the approach is [1].

4.2 Median cut

The median cut technique, [5], is conceptually wonderfully simple. The idea is to recursively split the LL map into regions of approximately equal summed radiance. Since the method splits all regions K times the technique produces 2^K sources, i.e., 2, 4, 8, 16, 32, etc. Figure 3 illustrated the result of running the technique on a light probe. The algorithm is as follows:

1. Add the entire light probe image to the region list as a single region.
2. For each region in the list, subdivide along the longest dimension such that its light energy is divided evenly.
3. If the number of iterations is less than K , return to step 2.

4. Place a light source at the centroid of each region, and set the light source radiance to the sum of the pixel values within the region.

The strength of this approach is that it is so straightforward, computationally light and easy to implement. The problems with this approach lie in two issues. The first issue is that it subdivides all regions at each iteration and depending on K there can be a large jump in the number of sources, which may be disadvantageous for real-time rendering with the approximated sources, where one would like as many sources as possible, but at the same time there is a performance limit in the graphics hardware. The second issue relates to step 4, where, for small K , and thereby large regions, a lot of radiance is transferred quite large distances over the sphere without any cosine weighting. This transfer could be done physically correct, but only for a single known surface normal. For arbitrary normals there is no alternative to just transferring the radiance and hope the regions are small enough that it does not constitute a grave error. This is naturally an invalid assumption for very small K .

4.3 Adaptive median cut

A description of this technique will be included in a future version of this paper. We have developed this technique which is heavily based on the original Median Cut technique, but our version can produce any number of sources, not just a power of 2.

4.4 Irradiance Optimization

The Irradiance Optimization technique by Madsen et al., [11], is significantly different from the first three. While the first three all operate entirely on a pixel level in the light probe image, i.e., operate in radiance space, the Madsen method operates in *irradiance* space.

The method is based on first using the original light probe image to compute the ground truth irradiance for a large number (M) of normal directions uniformly distributed across the unit sphere using Eq. 4. These M irradiance values constitute the goal vector in an optimization to estimate the parameters of N directional sources. Each source is defined by five parameters (RGB radiances and two angles for direction).

Given an estimate of these 5 times N parameters it is possible to compute the approximated irradiances for the M normals. Let L_i be the radiance of the i th source, and let $\omega(\theta_i, \phi_i) = [\sin(\theta_i) \cos(\phi_i), \sin(\theta_i) \sin(\phi_i), \cos(\theta_i)]$ be the direction vector to the i th source. Furthermore, let A be a fixed, small area (in steradians) of each light source to accommodate physical correctness. An disc area light source of 1 degree has an area of $2.392 \cdot 10^{-4}$ steradian. Finally, let \vec{n}_k be the k th normal. The irradiance on a dif-

ferential area surface with normal \vec{n}_k is then:

$$E(\vec{n}_k) = \sum_{i=1}^N L_i \text{Amax}(0, (\omega(\theta_i, \phi_i) \cdot \vec{n}_k) \quad (5)$$

By comparing the approximated irradiances to the ground truth irradiances we obtain an error vector, which can be converted to a parameter update vector. The source estimation process is thus an iterative, non-linear optimization process based on Newton’s iterative method, since the Jacobian can be expressed analytically.

5 Comparative evaluation

In a future version of this paper we will have compiled more extensive evaluations. At present we have only tested two methods (Median Cut and Irradiance Optimization), and they have just been evaluated on one light probe image. In this section we describe the results from these initial evaluations, and offer some observations based on them.

5.1 Tested light probes

The evaluation documented in this paper is based on the light probe shown in figure 5.

5.2 Performances

We ran the Median Cut and the Irradiance Optimization methods on the test light probe, and produced directional source approximations with 2, 4, 8, 16, 32, 64, and 128 sources. The Irradiance Optimization technique can be produce any number of sources, but was constrained to the source number cases which where also feasible for the Median Cut approach. We were unable to obtain a convergence on a 128 source solution with the Irradiance Optimization technique. This will be discussed later.

The evaluation is based on computing the irradiances resulting from the estimated set of sources for a large number of surface normal evenly distributed on a unit sphere, and comparing them to the ground truth irradiances. For each color channel when then compute the mean and the maximum of the absolute differences between estimated and ground truth irradiances. Figure 6 shows curves representing average and maximum error for each of the two methods as a function of the number of light sources used. The errors are an average over RGB.

5.3 Discussion

Figure 6 clearly shows that the Irradiance Optimization techniques performs much better than the Median Cut method. Generally the Median Cut method requires 2 to

3 times as many sources to achieve the same error as the Irradiance Optimization technique. For rendering this is very important from a computational point of view, since it will always be an advantage to use as few sources as possible.

In this test it was seen that the Irradiance Optimization technique could not converge when the number of sources comes above some threshold (64). This is a general tendency we have noticed, and it is strongly believed to be related to the fact that when number of sources grows too high there is too little energy (irradiance) for some sources to latch on to. Very quickly in the iterations the dominant sources become stable, leaving ever smaller amounts of energy to distribute among the rest of the sources being estimated. At the same time the sources tend to repel each other when they distribute across the sphere, because each source has a semi-spherical “footprint” in the irradiances, so each source is like shining a torch on a sphere, and sources will be reluctant to overlap footprints too much.

6 Conclusions

We have demonstrated that there is significant differences in the performance of the available techniques to approximate light probes radiance maps with a set of directional light sources. Test so far clearly demonstrate that the Irradiance Optimization technique requires much less sources to achieve the error level as the other techniques.

In terms of rendering speed it will always be an advantage to have as few light sources as possible and still achieve visibly acceptable performance. If a rendering is solely for visual purposes it may not be crucial whether the irradiance at a point is 1 percent or 5 percept wrong, but renderings can are also used in more radiometrically challenging contexts such as for inverse methods, aiming at estimating surface reflectance parameters from images, [15, 14]. For inverse problems the accuracy in the approximation can be very important.

Future work includes several straight forward issues, plus one somewhat more complicated. Primarily we need to test all the techniques, and we need to do it on more qualitatively different light probes. In this regard we are thinking about both indoor and outdoor scenarios. So far we have only tested the techniques in terms of resulting irradiance. Future experiments will attempt to test not only irradiance but also the spatial distribution of incident radiance. We plan to evaluate this issue by rendering scenes with the approximated sources and test the reflected radiance in and around cast shadows and compare them to shadows rendered with Monte Carlo path tracing.

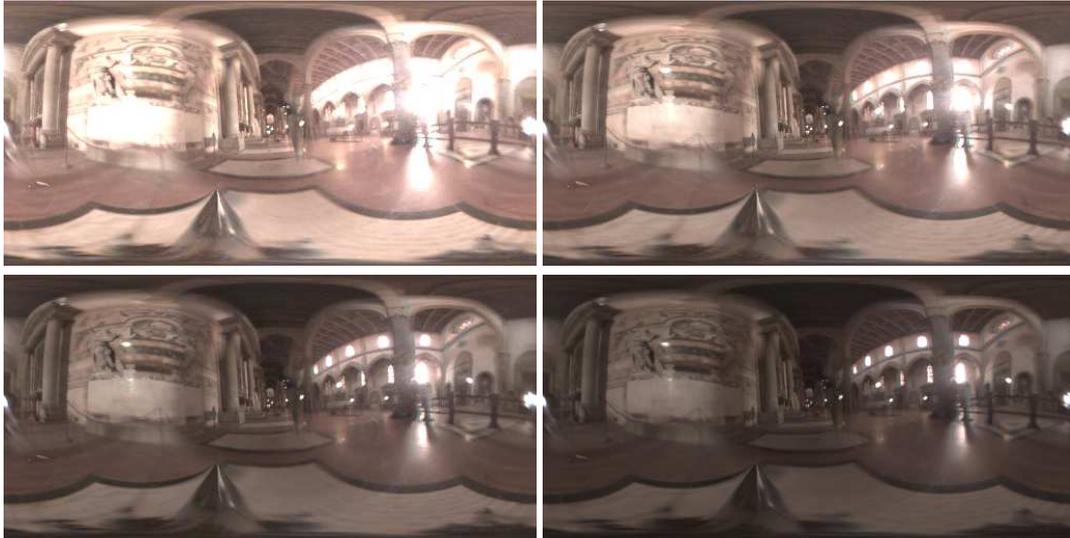


Figure 5: The tested light probe: Galileo's tomb in Florence, Italy. Acquired from [2]. Here shown in four different exposures to illustrate dynamic range.

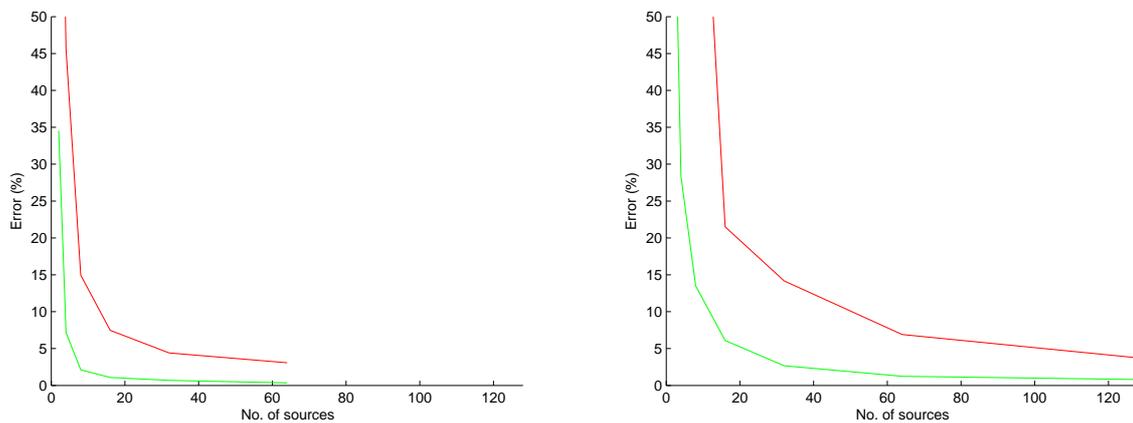


Figure 6: Mean and max irradiance error in percent for the two tested techniques. Left: Irradiance Optimization. Right: Median Cut.

Acknowledgments

This research is funded by the CoSPE project (26-04-0171) under the Danish Research Agency. This support is gratefully acknowledged.

References

- [1] J. M. Cohen and P. Debevec. *The LightGen HDRShop plugin*, 2001. www.hdrshop.com/main-pages/plugins.html.
- [2] P. Debevec. www.debevec.org/probes.
- [3] P. Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings: SIGGRAPH 1998, Orlando, Florida, USA, July 1998*.
- [4] P. Debevec. Tutorial: Image-based lighting. *IEEE Computer Graphics and Applications*, pages 26 – 34, March/April 2002.
- [5] P. Debevec. A median cut algorithm for light probe sampling. In *Proceedings: SIGGRAPH 2005, Los Angeles, California, USA, August 2005*. Poster abstract.
- [6] P. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Proceedings: SIGGRAPH 1997, Los Angeles, CA, USA, August 1997*.

- [7] P. Debevec et al. *www.hdrshop.com*.
- [8] S. Gibson, J. Cook, T. Howard, and R. Hubbard. Rapid shadow generation in real-world lighting environments. In *Proceedings: EuroGraphics Symposium on Rendering, Leuven, Belgium*, June 2003.
- [9] K. Jacobs and C. Loscos. State of the art report on classification of illumination methods for mixed reality. In *EUROGRAPHICS*, Grenoble, France, September 2004.
- [10] Henrik W. Jensen. *Realistic Image Synthesis Using Photon Mapping*. A. K. Peters, 2001.
- [11] C. B. Madsen, M. K. D. Sørensen, and M. Vittrup. Estimating positions and radiances of a small number of light sources for real-time image-based lighting. In *Proceedings: Annual Conference of the European Association for Computer Graphics, EUROGRAPHICS 2003, Granada, Spain*, pages 37 – 44, September 2003.
- [12] M. M. Oliveira. Image-based modelling and rendering: A survey. *RITA - Revista de Informatica Teorica e Aplicada*, 9(2):37 – 66, October 2002. Brasillian journal, but paper is in English.
- [13] M. Phar and G. Humphreys. *Physically Based Rendering – From Theory to Implementation*. Elsevier, 2004.
- [14] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *Proceedings: SIGGRAPH 1999, Los Angeles, California, USA*, pages 215 – 224, August 1999.
- [15] Y. Yu and J. Malik. Recovering photometric properties of architectural scenes from photographs. In *Proceedings: SIGGRAPH 1998, Orlando, Florida, USA*, pages 207 – 217, July 1998.