Importance Driven Path Tracing using the Photon Map

Henrik Wann Jensen*

Dept. of Graphical Communication
Technical University of Denmark

Abstract: This paper presents a new importance sampling strategy for Monte Carlo ray tracing in which a rough estimate of the irradiance based on the photon map is combined with the local reflection model to construct more efficient probability density functions that can be used in an importance sampling scheme.
The algorithm gives unbiased results, handles arbitrary reflection models and it is particularly efficient in scenes with highly non-uniform indirect illumination. Initial results and comparisons with traditional importance sampling strategies indicate a reduction in the noise level of more than 70%

Key Words: Global Illumination, Path Tracing, Importance Sampling, Photon Map

1 Introduction

Photorealistic rendering requires accurate simulation of global illumination and much work has been done in this area in the last 10 years. The problem was actually solved in 1986 by Kajiya [6] using a method called path tracing. This method is basically a brute force Monte Carlo simulation of light interaction with a given model. Path tracing is very general and it can be applied to arbitrarilly complex models. It requires only small amounts of memory and it is very applicable to parallel computers. The rendering time with path tracing is however so enormous that the method — in spite of being general — is not very attractive on current computers architectures.

Several papers have presented improvements to the path tracing method. Ward et al. [19, 22], did a very good job, by introducing a caching scheme in which indirect illumination is stored and reused at ideal diffuse surfaces. This significantly reduces the number of rays necessary in scenes with many flat, ideal diffuse surfaces. Shirley et al. [15] and Ward [20] reduce the number of shadow rays required to compute direct illumination in scenes with many light sources. Arvo et al. [1] described a technique called Russian roulette that eliminates the

* e-mail: igkhwj@unidhp.uni-c.dk url: http://www.gk.dtu.dk/home/hwj/

infinite reflection of light rays without introducing bias in the final solution. Several authors [10, 4] use importance sampling where the reflection characteristics of the surface are used to guide sampling rays into those directions from where the light will contribute mostly. Another approach is adaptive sampling where samples are concentrated in the most interesting parts of the scene. Lafortune et al. [9] constructs a 5d tree representing the overall flux in the scene and uses this tree in the selection of new sampling directions. Adaptive sampling can however easily lead to biased solutions [2] and it must be used with great care. Recently bidirectional path tracing was introduced by Lafortune et al. [8] and Veach et al. [18]. It is a mixture of path tracing from the eye and path tracing from the light sources and it is particularly efficient in scenes with highly non-uniform indirect illumination.

Currently the simulation of global illumination is obtained most efficiently by the two pass methods in which a light pass (ie. radiosity) produces a solution that is visualized with a simplified path tracing algorithm [14, 17, 5, 13, 7]. Another efficient method or collection of methods is presented in the Radiance rendering program by Ward [23].

These two pass methods and the Radiance program do however degrade to pure path tracing in very complex environments where the surfaces are no longer ideal diffuse or ideal specular or where the objects are either too complex or too many to be represented by polygons. Even though some of the two pass methods can let the light pass step work on simplified scenes [13, 7] they still have to visualize the solution using path tracing in order to eliminate artifacts in the light pass step. This path tracing step does not utilize the fact that the model does contain information on the distribution of the light and the importance sampling is therefore still only using the reflection characteristics of the surface to guide new sampling rays.

This is particularly problematic in scenes with highly non-uniform indirect illumination. These situation are handled more efficiently using bidirectional path tracing, but this method does not use the irradiance stored within the model. It would be more appropriate to use the information created in the light pass to guide the path tracing algorithm. This would also permit the use of the method in two pass techniques where path tracing is used only for the first reflection seen by the eye [5, 13, 7].

In this paper we present an importance sampling scheme in which the scene is preprocessed and rough estimates of the incoming light are created everywhere in the model. These rough estimates are used to generate sampling directions based on probability density functions that more closely fits the true light contribution as opposed to standard importance sampling methods which only use the local reflection model.

In this way we concentrate the samples in the important parts of the scene. This approach does not give any bias on the final solution as opposed to some of the common adaptive strategies that also tries to put more samples into important parts of the scene.

2 Mathematical Background

Given the integral

$$I = \int g(x)\, dx \quad , \quad x \in D \subset R^n \tag{1}$$

We can evaluate this integral using a Monte Carlo technique known as the sample-mean method [12] by representing the value of the integral as the expected value of any stochastic variable X with p.d.f. $p(x)$, $x \in D$ such that $p(x) > 0$ when $g(x) \neq 0$

$$I = \int \frac{g(x)}{p(x)} p(x)\, dx = E\left\{\frac{g(x)}{p(x)}\right\} \tag{2}$$

An estimate of the integral is obtained by taking N random sample points x_i distributed according to $p(x)$

$$I \approx \frac{1}{N} \sum_{i=1}^{N} \frac{g(x_i)}{p(x_i)} \tag{3}$$

The error on this estimate largely depends on the choice of $p(x)$ and the number of samples but as a general rule the standard deviation is proportional to $1/\sqrt{N}$. That is, in order to halve the error we have to quadruple the number of samples! Careful selection of $p(x)$ can however lower the error. $p(x)$ should be constructed so that more samples are put into those regions where $g(x)$ has the highest absolute value. It can be shown that the optimal choice for $p(x)$ is [12]

$$p(x) = \frac{g(x)}{I} \tag{4}$$

This choice gives a standard deviation of zero always! Unfortunately it also requires knowledge of I which is the value we are trying to compute. But in general we can do better by using a p.d.f. that looks like the function we are integrating instead of just using a uniform distribution.

The problem in global illumination is given in the rendering equation [6]. It expresses the radiance, L_r, reflected from position $\mathbf{x}$ as

$$L_r(\mathbf{x}, \Psi_r) = \int_{\text{all } \Psi_i} f_r(\mathbf{x}, \Psi_r, \Psi_i) L_i(\mathbf{x}, \Psi_i) \cos\theta_i \, d\omega_i = \int_{\text{all } \Psi_i} f_r(\mathbf{x}, \Psi_r, \Psi_i) \frac{d^2\Phi_i(\mathbf{x}, \Psi_i)}{dA\, d\omega_i}\, d\omega_i \tag{5}$$

where Ψ_r and Ψ_i are the direction of the reflected respectively incoming radiance/flux. f_r is the BRDF at $\mathbf{x}$, L_i is the incoming radiance and Φ_i is the incoming flux.

In order to solve this integral light is sampled from a number of discrete directions Ψ_i. Just sampling from random directions is in most situations very inefficient — for a specular surface only light from a small solid angle is important and it would be more efficient to sample within this small solid angle.

Most implementations of path tracing and similar Monte Carlo based algorithms therefore use importance sampling. These implementations use the knowledge of f_r to sample those direction from where incoming light will contribute mostly to the reflected radiance. For specular surface this approach is very efficient since the choice of sampling direction is significantly narrowed down. This is not the case with diffuse surfaces since the entire hemisphere can contribute to the reflected radiance. A p.d.f. based on the ideal diffuse BRDF is particularly inefficient when the incoming radiance L_i is highly non-uniform (ie. the incoming radiance is concentrated in small solid angles). In this situation it would be better to use the optimal p.d.f. $p(\mathbf{x}, \Psi_r, \Psi_i)$

$$p(\mathbf{x}, \Psi_r, \Psi_i) \propto f_r(\mathbf{x}, \Psi_r, \Psi_i) \frac{d^2 \Phi_i(\Psi_i)}{dA \, d\omega_i} \tag{6}$$

Unfortunately this requires knowledge about the incoming flux, Φ_i, which is not available. However, even a crude estimate of Φ_i can be used to create a more optimal p.d.f. and in the following sections we will describe a method in which the photon map is used to obtain this estimate.

3 The Photon Map

The photon map [7] represents a distribution of photons (particles) throughout the scene and it is created by emitting a large number of photons from each light source into the scene based upon the emissive characteristics of the light source. The technique is similar to particle tracing with the exception that the intersection point of each particle is stored explicitly within the scene.

Each photon is traced through the scene using a strategy similar to path tracing. The first object that the photon hits gives rise to two events: Firstly if the surface of the object is diffuse the photon is stored at the intersection point and secondly Russian roulette is used to determine whether the photon is reflected or absorbed by the object. The new direction of a reflected photon is computed using the BRDF of the surface. We only store the photons representing indirect illumination (ie. photons that have been reflected at least once). The light sources are separated in the sampling process and they should not be part of the irradiance estimate used to compute the p.d.f.

The photon represents a small packet of energy arriving at a surface from a given direction. Since the number of photons can be quite large we have chosen a relatively compact representation that occupies only 36 bytes:

```
struct photon {
    long energy;                     // Packed energy (RGB)
    float position[3];               // Photon position
    float theta,phi;                 // Photon direction
    char normal[3];                  // Surface normal
    char key;                        // kd-tree parameter
    struct photon *left, *right;     // Rest of the kd-tree
    };
```

In this structure the photon energy is packed using a method similar to Wards packed pixels [21]. The energy covers several wavelengths. In other contexts it might be relevant to store energy for individual wavelengths — this would also make the name, photon, more correct.

In the photon map we need to be able to locate the n photons that have the shortest distance to a point $\mathbf{x}$. This can be done quite efficiently using a kd-tree [3].

4 Importance Driven Path Tracing using the Photon Map

In the following discussion we assume that the intersection point $\mathbf{x}$ and the reflected direction Ψ_r are given and they are therefore omitted. As shown in section 2 the optimal choice of $p(\theta_i, \phi_i)$ to select a sampling directions is

$$p(\theta_i, \phi_i) \propto f_r(\theta_i, \phi_i) \frac{d^2\Phi_i(\theta_i, \phi_i)}{dA\, d\omega_i} \tag{7}$$

where $(\theta_i, \phi_i) = \Psi_i$.

The existing importance sampling approaches that use f_r as the p.d.f. can in most situations use standard transformation techniques to generate random numbers with the wanted distribution. In [16] it is described how to create a transformation function T that maps a uniform distribution of points (u, v) on the unit square onto a sampling direction (θ, ϕ). We use the following notation for describing the transformation of (u, v) into a direction

$$(\theta, \phi) = T(u, v) \tag{8}$$

Likewise we use the notation $(u, v) = T^{-1}(\theta, \phi)$ to denote the inverse transformation. As an example, to generate a sampling direction based on the BRDF for an ideal diffuse surface the following transformation should be used

$$(\theta, \phi) = (\mathrm{acos}\sqrt{1-u}, 2\pi v) \tag{9}$$

In order to take Φ_i into account we apply the information from the photon map. Like [7] we locate the N photons that have the shortest distance to $\mathbf{x}$. Each photon p carries the flux $\Delta\Phi_{i,p}$ in the direction (θ_p, ϕ_p) (note that (θ_p, ϕ_p) must be transformed from the global representation stored in the kd-tree into a direction compatible with the local coordinate system used at $\mathbf{x}$). If we assume that all n photons intersected the surface at $\mathbf{x}$ then we can compute the contribution from each photon p to the reflected flux $\Delta\Phi_r$ as

$$\Delta\Phi_r = f_r(\theta_p, \phi_p) \Delta\Phi_{i,p}(\theta_p, \phi_p) \tag{10}$$

To use this information in our generation of sampling directions we construct a discrete p.d.f. on the unit square. Each point (u, v) in this unit square corresponds to a sampling direction $T(u, v)$ and each incoming photon corresponds to

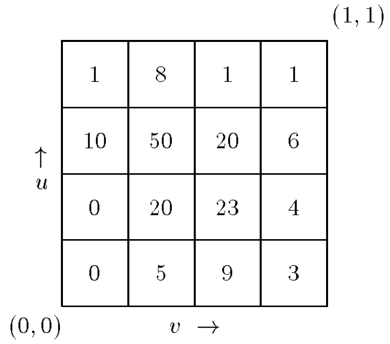

Fig. 1. The unit square is partitioned into distinct regions and the photon contributions in each region are accumulated.

a point $(u_p, v_p) = T^{-1}(\theta_p, \phi_p)$ in the unit square. For every photon we find the corresponding point in the unit square and we insert the photon energy at this position. The unit square is then partitioned into $m \times n$ regions (see fig. 1) and the energy in each region is accumulated. To avoid bias we eliminate all regions with zero energy (probability) by adding a small fraction of the overall energy stored within the unit square to these regions. The result is an estimate of the energy arriving from different sets of directions. This is our p.d.f. and it contains all of the elements found in equation 7.

To use the constructed p.d.f. we create a discrete cumulative probability density function from the information in the unit square. This is illustrated in

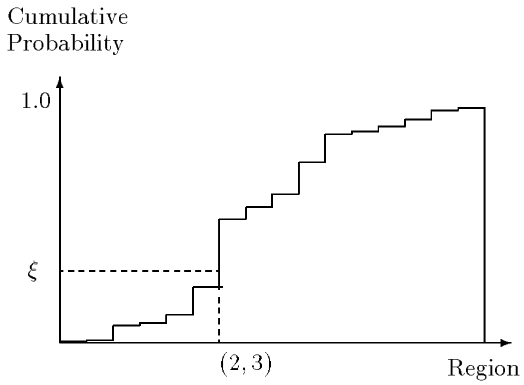

Fig. 2. The accumulated energy is used in the creation of a cumulative probability distribution function and this function is used to select the region that contains the new sampling direction. The dashed line demonstrates how a random value ξ is mapped into a specific region in unit square.

fig. 2 and as demonstrated in the graph this function is used to select a region in the unit square based upon a random value ξ. The chance that a region is selected is proportional to the energy accumulated in that region which means that it is more likely that we select a region (ie. set of incoming directions) that contributes significantly to the reflected radiance. Once a region is found we select a random point (u, v) within the region and our sampling direction becomes $(\theta, \phi) = T(u, v)$.

Due to the non-uniform sampling we have to scale the radiance returned by our sampling ray. The scaling factor s is

$$s = \frac{\text{Total energy in the square}}{(\text{Energy in region}) * (\text{number of regions})} \tag{11}$$

5 Results and Discussion

We have implemented the importance sampling algorithm in a program called `MIRO` on a Pentium PC with 32MB RAM running Linux. The results have been produced using a parallel implementation of the program running on the Linux machine and 31 Silicon Graphics workstations. The implementation currently supports ideal diffuse, ideal specular, rough specular and anisotropic reflection models. The photon map has only been incorporated in the importance sampling algorithm for ideal diffuse reflection.

Two test scenes have been created. Test scene 1 is an empty version of the Cornell box. The walls in the box are white with the exception of the left wall which is red and the right wall which is green. There is a small square shaped light source just below the ceiling. The light source is only illuminating the ceiling and most of the scene is therefore illuminated indirectly.

The purpose of test scene 1 is to test the convergence speed of importance sampling with and without the photon map. We initially computed a reference image using 32000 samples pr. pixel in the resolution 80x60 — we did not use importance sampling for this computation. Due to the strong indirect illumination we had to use a very large number of samples in order to remove all visible noise from the image.

The estimate from the photon map depends on several parameters: The number of photons, N, used in the estimate, the number of photons N_p stored in the kd-tree and the number of regions in the discrete p.d.f. To limit the amount of adjustable parameters we use a fixed partitioning of the p.d.f. We found that 4x16 (4 θ-intervals and 16 ϕ-intervals) gave reasonable results with the values of N and N_p that we tested. In general the number of partitions should be increased as N and N_p become larger.

We tested the different configurations by creating images using $1, 2, 3, \ldots, 500$ samples pr. pixel. We measured the variance of the difference between the computed images and the reference image — also known as the mean-square error, MSE. This value gives a good indication of the amount of noise present in the image and since we use Monte Carlo techniques it provides a good measure of the quality of our images.

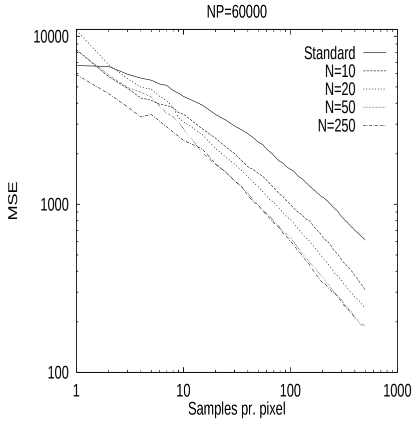

(a)

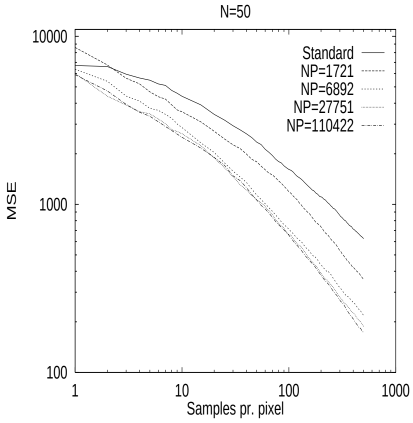

(b)

Fig. 3. The effect of N and N_p on the convergence speed

The graph in fig. 3a demonstrates how an increasing number of photons in the estimate improves the quality of the computed images. Using $N_p \approx 60000$ stored photons we were able to reduce the noise in the images by increasing N to a value of 50 photons. Increasing N beyond this value had only very little effect on the results. The graph also demonstrates how the convergence speed is improved when the information from the photon map is included in the importance sampling scheme. At 500 samples pr. pixel the standard importance sampling method gave a MSE value that was more than 3 times larger than the MSE value obtained when the photon map was used.

We also examined the effect of N_p on the quality of the computed images. Using $N = 50$ we obtained the results shown in the graph in fig. 3b. With the used parameter configuration we found that increasing N_p beyond 30000 did not improve the result.

In fig. 4 (see colour plates) we have shown the computed images corresponding to 100, 200, 300, 400 and 500 samples pr. pixel. The top row contains images computed with standard importance sampling and the bottom row contains images computed using importance sampling based on the photon map. The images demonstrates how the noise level is reduced when the information from the photon map is added.

Our second test scene was created to test the algorithm in a more realistic environment. The desk is illuminated by two light sources: One large in the ceiling and one small light bulb within the lamp shade. The light from the light bulb is scattered diffusely through the lamp shade.

We created two images of test scene 2. In fig. 5 we used standard importance sampling and in fig. 6 we used the photon map with $N_p = 33000$ and $N = 50$. We also created a reference image using 2000 samples pr. pixel and the MSE of fig. 5 was approximately 4 times larger than the value for fig. 6. The reduction in noise is particularly visible at the wall just behind the lamp shade.

The implementation of the photon map has not been optimized. The computation of the photon contribution to the reflected radiance makes heavy use of trigonometric functions. Still the computation time of fig. 6 was only increased with 20% due to the use of the photon map. The time used by the photon map was only affected noticeably by changing N and this is primarily due to the change in the number of evaluations of trigonometric functions.

We tried to reduce the effect of N on the computation time by replacing the photon map with hemicubes stored at discrete positions within the scene. During importance sampling the nearest hemicube would represent the incoming flux. We did however find that this approach suffered from aliasing problems and the results obtained were not very good. Instead we plan to limit the number of directions for a photons to perhaps 65000. In this way we would be able to use lookup tables and completely avoid trigonometric evaluations.

6 Conclusion and Future Directions

In this paper we have demonstrated how the use of photon maps in an importance sampling scheme can improve the quality of images computed using the path tracing algorithm. Based on a rough estimate of the incoming flux we concentrate our samples in those directions that contributes mostly to the reflected radiance. This improves the convergence speed without adding bias to the final solution. By using the photon map we were able to reduce the noise level in the computed images with more than 70% compared with traditional importance sampling approaches without increasing the number of sampling rays.

Future enhancements include further investigation in the parameters for the method. By increasing the number of photons used we should be able to obtain even better results. Combining this with a more efficient usage of the information in the photon map would make the method very suitable for the visualization step in the existing two pass techniques for global illumination.

Acknowledgments

The author wishes to thank GK for providing excessive computational resources and to the reviewers and Eric Lafortune for their helpful comments to the paper.

References

1. Arvo, James and David Kirk, "Particle Transport and Image Synthesis". *Computer Graphics* **24** (4), 53-66 (1990).
2. Arvo, James and David Kirk, "Unbiased Sampling Techniques for Image Synthesis". *Computer Graphics* **25** (4), 153-156 (1991).
3. Bentley, Jon Louis and Jerome H. Friedman, "Data Structures for Range Searching". *Computing Surveys* **11** (4), 397-409 (1979).

4. Blasi, Philippe; Bertrans Le Saëc and Christophe Schlick: "An Importance Driven Monte-Carlo Solution to the Global Illumination Problem". *Eurographics Workshop on Rendering* **5**, Darmstadt (1994).
5. Chen, Shenchang Eric, Holly E. Rushmeier, Gavin Miller and Douglas Turner, "A Progressive Multi-Pass Method for Global Illumination". *Computer Graphics* **25** (4), 165-174 (1991).
6. Kajiya, James T.: "The Rendering Equation". *Computer Graphics* **20** (4), 143-149 (1986).
7. Jensen, Henrik Wann and Niels Jørgen Christensen: "Photon maps in Bidirectional Monte Carlo Ray Tracing of Complex Objects". *Computers and Graphics* **19** no. 2, (1995)
8. Lafortune, Eric P.; Yves D. Willems: "Bidirectional Path Tracing". *Proceedings of CompuGraphics*, 95-104, june (1993)
9. Lafortune, Eric P.; Yves D. Willems: "A 5D Tree to Reduce the Variance of Monte Carlo Ray Tracing". *Eurographics Workshop on Rendering* **6**, Dublin (1995)
10. Lange, Brigitta: "The Simulation of Radiant Light Transfer with Stochastic Ray-Tracing". *Eurographics Workshop on Rendering* **2**, Barcelona (1991)
11. Nicodemus, F. E.; J. C. Richmond; J. J. Hsia; I. W. Ginsberg; T. Limperis: *"Geometrical Considerations and Nomenclature for Reflectance"*. National Bureau of Standards, okt. 1977
12. Rubinstein, Reuven Y.: *"Simulation af the Monte Carlo Method"*. John Wiley & Sons, (1981)
13. Rushmeier, Holly; Charles Patterson and Aravindan Veerasamy: "Geometric Simplification for Indirect Illumination Calculations". *Proceedings of Graphics Interface '93*, 227-236 (1993)
14. Shirley, Peter, "A Ray Tracing Method for Illumination Calculation in Diffuse-Specular Scenes". *Proc. Graphics Interface*, 205-212 (1990).
15. Shirley, Peter and Changyaw Wang, "Direct Lighting Calculation by Monte Carlo integration". *Eurographics Workshop on Rendering* **2**, Barcelona (1991).
16. Shirley, Peter: "Nonuniform Random Point Sets via Warping". In *Graphics Gems III*, David Kirk (ed.), Academic Press, 80-83 (1992).
17. Sillion, François X., James R. Arvo, Stephen H. Westin and Donald P. Greenberg, "A Global Illumination Solution for General Reflectance Distributions". *Computer Graphics* **25** (4), 187-196 (1991).
18. Veach, Eric and Leonidas Guibas: "Bidirectional Estimators for Light Transport". *Eurographics Workshop on Rendering* **5**, Darmstadt (1994)
19. Ward, Gregory J.; Francis M. Rubinstein and Robert D. Clear, "A Ray Tracing Solution for Diffuse Interreflection". *Computer Graphics* **22** (4), 85-92 (1988).
20. Ward, Gregory J., "Adaptive Shadow Testing for Ray Tracing". *Eurographics Workshop on Rendering* **2**, Barcelona (1991).
21. Ward, Greg: "Real pixels". In *Graphics Gems II*, James Arvo (ed.), Academic Press, 80-83 (1991).
22. Ward, Gregory J. and Paul S. Heckbert: "Irradiance Gradients". *Eurographics Workshop on Rendering* **3**, Bristol (1992).
23. Ward, Gregory J.: "The RADIANCE Lighting Simulation and Rendering System". *Computer Graphics* **28** (4), 459-472, (1994).

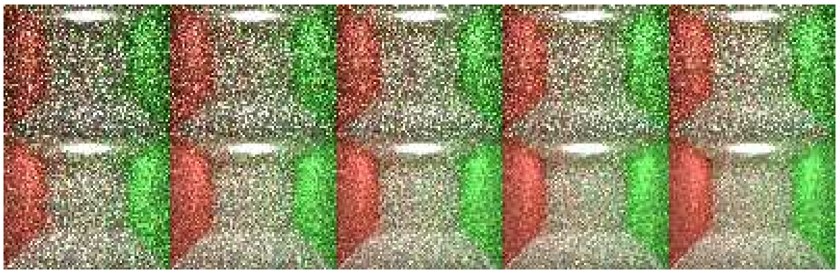

Fig. 4. Test scene 1 sampled using 100,200,300,400 and 500 samples pr. pixel. The top row shows the results using standard importance sampling and the botton row shows the results when the information from the photon map is added

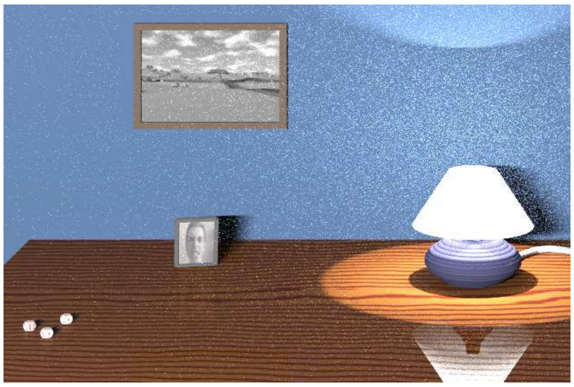

Fig. 5. Test scene 2 sampled with 50 samples pr. pixel using standard importance sampling

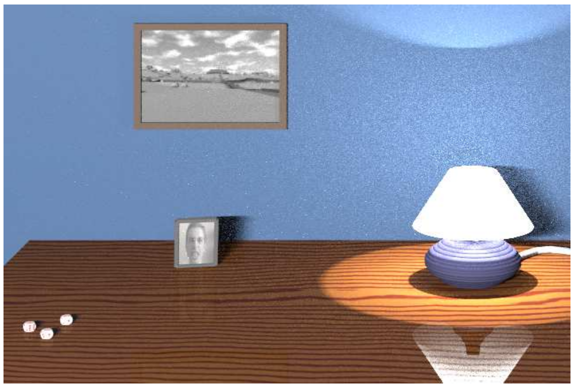

Fig. 6. Test scene 2 sampled with 50 samples pr. pixel using importance sampling enhanced with the photon map