Institut für Computergraphik und Algorithmen
Technische Universität Wien

Karlsplatz 13/186/2
A-1040 Wien
AUSTRIA

Tel: +43 (1) 58801-18601
Fax: +43 (1) 58801-18698

Institute of Computer Graphics and Algorithms
Vienna University of Technology

email:
technical-report@cg.tuwien.ac.at

other services:
http://www.cg.tuwien.ac.at/
ftp://ftp.cg.tuwien.ac.at

Importance Sampling with Hemispherical Particle Footprints

Heinrich Hey Werner Purgathofer
{hey,wp}@cg.tuwien.ac.at

TR-186-2-01-05
January 2001

Abstract

We present a new importance sampling technique for stochastic ray-based global illumination methods. It allows to enhance the efficiency of global illumination calculations in general scenes with complex illumination settings by selecting preferably those sampling or shooting directions which yield a high contribution. The probability density functions for this are generated with a photon map or importance map that represents the expected contribution. A direction for a given point in the scene is selected by making directional footprints of the nearest neighbor particles onto the hemisphere of the point. By selecting the radii of the footprints adaptively according to the directional density of the particles, rays can be shot precisely into highly contributing regions where several particles come from in a small solid angle. We compare our new technique with existing particle map based importance sampling and show that our new technique achieves a considerably lower level of noise and spots in the same computation time.

Importance Sampling with Hemispherical Particle Footprints

Heinrich Hey Werner Purgathofer
Vienna University of Technology
{hey,wp}@cg.tuwien.ac.at

Abstract. We present a new importance sampling technique for stochastic ray-based global illumination methods. It allows to enhance the efficiency of global illumination calculations in general scenes with complex illumination settings by selecting preferably those sampling or shooting directions which yield a high contribution. The probability density functions for this are generated with a photon map or importance map that represents the expected contribution. A direction for a given point in the scene is selected by making directional footprints of the nearest neighbor particles onto the hemisphere of the point. By selecting the radii of the footprints adaptively according to the directional density of the particles, rays can be shot precisely into highly contributing regions where several particles come from in a small solid angle. We compare our new technique with existing particle map based importance sampling and show that our new technique achieves a considerably lower level of noise and spots in the same computation time.

1 Introduction

In scenes that are large or that contain complex illumination settings efficiency demands to concentrate the computational effort during the global illumination simulation and rendering to those parts of the scene that contribute most to the image. For stochastic ray-based rendering and global illumination methods this means that paths shall be shot preferably into directions where their effect is high. Light paths shall be shot preferably into parts of the scene that are visible so that as little work as possible is spent into unnecessarily illuminating invisible parts of the scene. View paths shall be shot preferably into directions where much light comes from so that they contribute most to the image. This problem can be solved with importance sampling where the outgoing direction of a ray of a path is selected stochastically distributed according to a probability density function (PDF) which approximates the contribution of the path into the outgoing direction.

A simple way of doing importance sampling is to select the direction according to the bidirectional scattering distribution function (BSDF) at the scattering point [BLS94,DW94,LF97,Lan91,TN+98]. It does not require to estimate the incoming light or visibility, but this of course makes it less likely that the PDF corresponds to the actual contribution.

Several methods use meshing to store an approximation of the contribution in the scene [DW95,NN+96,SCP99,UT97]. Another solution is to generate the PDFs from incoming radiance which is stored in a 5D tree [LW95], or to store the illumination in the scene in a neural gas structure [Bus97]. Importance sampling can also be done by selecting the outgoing directions evolutionary [Bus97,LB94] instead of stochastically.

Photon map based importance sampling [Jen95], which is used eg. in photon map global illumination [Jen96b], uses a photon map, which is generated in a light pass, as approximation of the illumination to select the scattering directions in a subsequent path tracing pass. This kind of importance sampling can also be used for selecting scattering directions of light paths by usage of an importance map [PP98], which is generated in a preceding view pass and which is the analogue of a photon map (an importon is the analogue of a photon). A PDF is generated for a given scattering point in the scene by inserting the contribution of the k_p (typically 50) nearest particles from the photon or importance map into a grid that is mapped onto the hemisphere of the point [Shi92] (see figure 1). A grid cell is selected accordingly to the accumulated contributions of the cells, and the outgoing direction is selected randomly within this cell. The targeting precision into important directions is therefore limited to the resolution of the grid, and the grid resolution is limited to the number of nearest neighbor photons.

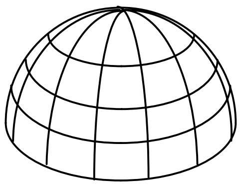

Fig. 1. Grid on hemisphere that is used for the PDF in existing particle map based importance sampling.

Directional importance information can also be represented in a hierarchical data structure, eg. a kd-tree or a hierarchy of spherical triangles, instead of a grid [TN+98]. A hierarchical data structure is useful if the PDF that it represents is estimated by means of a large number of samples, eg. for generating a PDF at a lightsource by means of the contribution of already shot light paths [DW94]. If only a small number of nearest neighbor particles is available to estimate the PDF, as it is the case in particle map based importance sampling, then simply inserting these particles into such a hierarchical data structure results in unnecessary blurring of the borders of highly contributing regions that lie in a small solid angle (see figure 2). An optimal PDF should be able to represent such important regions precisely to allow precise targeting (see figure 3).

In this paper we present a new importance sampling method that features this desired targeting characteristic without increasing the required number of nearest neighbor particles. It also uses a particle map, supports surfaces with general BSDFs and needs no meshing of the scene. For importance sampling of view paths a photon map is used which is generated in a preceding light pass. For importance sampling of light paths an importance map is used which is generated in a preceding view pass.

The outgoing direction of a ray that shall be shot from a given point in the scene is selected by making footprints of the nearest neighbor particles onto the hemisphere of the point (see figure 4). A footprint is located at the incoming direction of its particle, the radius of the footprint corresponds to how many other particles come from a nearby direction, and its volume corresponds to the light power of the photon or importance of the importon and the BSDF of the surface. The PDF according to which the outgoing direction of the ray is selected corresponds to these footprints. This has the advantage over existing particle map based importance sampling that due to the adaptive radii of the footprints an outgoing ray can be shot more precisely into highly contributing regions where several particles come from in a small solid angle.

In section 2 we describe our new importance sampling technique in detail. In section 3 we compare our technique with exisiting particle map based importance sampling in a path tracing application. We finaly present our conclusions in section 4.

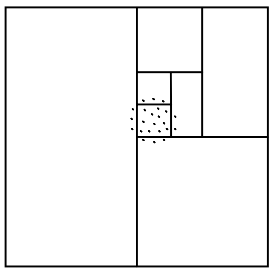

Fig. 2. Several particles lie in a small solid angle and represent a region of high contribution, eg. bright light that comes through a small opening or from a small reflector. If these particles are simply inserted into a hierarchical data structure, eg. a kd-tree, then several particles from this dense region lie in large low density nodes of the data structure. Therefore the region´s border is unnecessarily blurred and not the whole highly contributing region can be targeted precisely.

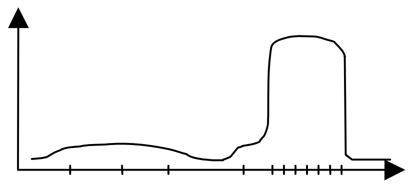

Fig. 3. An optimal PDF (here sketched in 1 dimension of the hemisphere) should be fitting tightly in regions with directionally dense particles to allow precise targeting into these regions where the particles provide dense illumination information. On the other hand it should be fitting loosely in sparse regions, which means that particles in sparse regions should distribute their contribution to the PDF over a wider solid angle to compensate the lower density of illumination information that the particles provide in these regions.

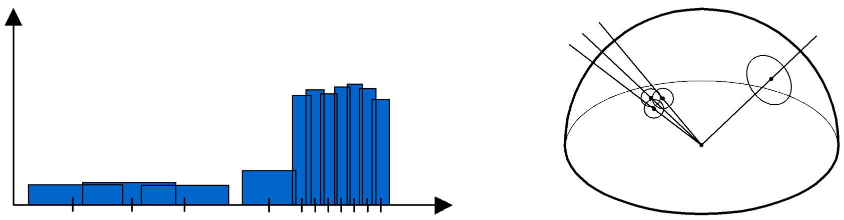

Fig. 4. We realize a PDF with the characteristics from figure 3 as sum of footprints of the nearest neighbor particles plus a small value (not shown here) to avoid biasing in directions without footprints (left: footprints in 1D, right: a few footprints on the hemisphere). Each footprint has an adaptive radius that corresponds to the directional particle density at its particle´s incoming direction.

2 Importance Sampling with Hemispherical Particle Footprints

Selecting an importance sampled outgoing direction at a given scattering point of a path works in the following sequence:

- Get the nearest neighbor particles from the particle map.
- Make a fast and coarse estimation of the directional density of the nearest neighbor particles.
- Select the outgoing direction. This is done
  - by selecting one of the nearest neighbor particles, selecting its footprint radius according to the directional particle density, and selecting the outgoing direction in this particle´s footprint
  - or by selecting the outgoing direction according to the BSDF to avoid bias in directions that are not covered by footprints.

  The decision which of these two methods is used is done stochastically. p_{BSDF} is the user defined probability of selecting the outgoing direction according to the BSDF.
- Weight the outgoing direction according to the value of the PDF in this direction. This PDF value is calculated by means of the footprints of the nearest neighbor particles and the BSDF.

In the following sections we describe these steps in more detail. Note that we do not need to generate the PDF for the whole hemisphere (this could be represented eg. with spherical wavelets [SS95]) because we have to evaluate the PDF only at the generated outgoing direction.

2.1 Particle map

A particle map, which is a photon map [Jen95] for importance sampling of view paths, or an importance map [PP98] for importance sampling of light paths, is generated prior to the pass where the importance sampling shall be done. To perform the importance sampling at a given scattering point x in this pass, we search for the k_p (user defined, typically 50) nearest neighbor particles to x [Jen95] whose contribution

c_p to the given path with the incoming direction Ψ_i at x is not 0. c_p can be 0 eg. if x lies at the frontside of an opaque surface and the particle lies at the backside. If not k_p such particles can be found within a user defined maximum distance to x, then there is not enough information available for importance sampling at x with the particle map, and we therefore have to fall back on importance sampling according to the BSDF.

In the case of a photon, c_p is equivalent to the photon´s reflected flux. In the case of an importon, c_p is equivalent to the importon´s reflected importance.

$$c_p = \Phi_p f_r(x, \Psi_p, \Psi_i) \tag{1}$$

$$c_p = W_p f_r(x, \Psi_p, \Psi_i) \tag{2}$$

Φ_p is the flux that the photon carries, W_p is the importance that the importon carries, Ψ_p is the particle´s incoming direction, and f_r is the BSDF at x from Ψ_p to Ψ_i. According to Peter and Pietrek [PP98] an importon does not store its W_p because it is assumed to be equal for all importons, therefore W_p can be set to 1 in (2). Extending this definition of an importon so that it stores its W_p for each color channel can nevertheless be useful in many scenes, eg. if parts of the scene are seen through a colored glass.

2.2 Directional particle density estimation

Next we perform a fast and coarse estimation of the directional particle density at the hemisphere of x. This estimate is necessary for the selection of the footprints´ radii in the following steps. We estimate the directional particle density by splatting the incoming directions of the nearest neighbor particles onto a grid at the ground plane of the hemisphere (see figure 5). The ground plane coincides with the tangential plane of x. The incoming direction of a particle is projected onto the ground plane of the hemisphere where it falls into a cell of the grid and makes a splat that is centered at this cell. We use splats that are 3x3 grid cells wide, and each splat increases the value g of each of its 3x3 cells in the grid by 1, independently of c_p. For k_p=50 we use a grid with 32x32 cells.

This directional particle density estimation method requires that we only use the particles from the positive hemisphere or only the particles from the negative hemisphere, because otherwise the directions from both hemispheres would be mixed up at the ground plane. Therefore we stochastically select one of both hemispheres. The probabilities $p_{\Omega+}$ and $p_{\Omega-}$ of selecting the positive Ω_+ or negative hemisphere Ω_- are

$$p_{\Omega+} = \frac{\sum_{p \in \Omega+} c_p}{\sum_{p \in \Omega+ \cup \Omega-} c_p} \tag{3}$$

$$p_{\Omega-} = 1 - p_{\Omega+} . \tag{4}$$

After all nearest neighbor particles from the selected hemisphere have been splatted into the grid, the directional particle density estimate δ for direction Ψ is calculated as

$$\delta(\Psi) = \frac{g(cell_{\Psi})}{\omega(cell_{\Psi})}. \tag{5}$$

$cell_{\Psi}$ is the cell that corresponds to Ψ, and ω is the solid angle of the cell projected onto the hemisphere. ω is precomputed for the whole grid and can be approximated as

$$\omega(cell) \approx \frac{4}{k_g \Psi_{cell} N}. \tag{6}$$

k_g is the number of cells in the grid, Ψ_{cell} is the direction that corresponds to the center of the cell, and N is the surface normal at x.

Note that δ does not need to be very exact, because it is only used to select the radii of the footprints. Even if the footprint radii would be selected arbitrarily, the resulting PDF would still be correct, but of course it would not feature the desired characteristics from figure 3. Note also that we do not use this grid as PDF because it would not have the desired characteristic of fitting tightly in dense regions and fitting loosely in sparse regions.

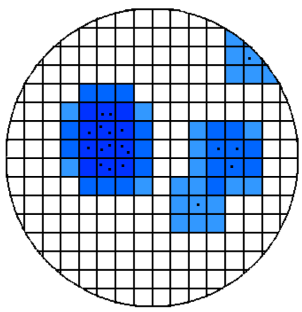

Fig. 5. Directional particle density estimation by making 3x3 splats of the particles into a grid at the ground plane of the hemisphere.

2.3 Footprints

Each nearest neighbor particle of the selected hemisphere distributes its contribution to the PDF in a directional footprint on the hemisphere (see figure 4). We select the footprint radius (see figure 6) of a particle with incoming direction Ψ_p as

$$r = \frac{k_r}{\delta(\Psi_p)} \tag{7}$$

with a user defined scaling faktor k_r. To achieve a valid PDF we have to ensure that all generated footprints lie completely in the selected hemisphere. For a particle with an incoming direction at a low angle the r that we get from (7) results in a footprint that lies partly in the other hemisphere if

$$r > r_{\max} \, . \tag{8}$$

$$r_{\max} = \Psi_p N \tag{9}$$

In such a case we have to resize the footprint so that it fits into the selected hemisphere by setting $r{=}r_{max}$ (see figure 7). The footprint´s solid angle is

$$\omega_f = 2\pi\left(1 - \sqrt{1 - r^2}\right) \tag{10}$$

$$\omega_f \approx r^2\pi \qquad \text{if } r{<<}1. \tag{11}$$

The footprint´s contribution to the PDF is

$$p_f(\Psi) = \frac{(1 - p_{BSDF})4\pi c_p}{\omega_f \sum\limits_{p \in \Omega + \cup \Omega -} c_p} \qquad \text{for all directions } \Psi \text{ inside the footprint} \tag{12}$$

$$p_f(\Psi) = 0 \qquad \text{for all directions } \Psi \text{ outside the footprint.} \tag{13}$$

Ψ is inside the footprint if

$$\left|(\Psi \Psi_p)\Psi_p - \Psi\right| < r \, . \tag{14}$$

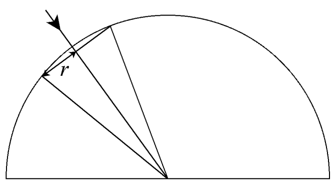

Fig. 6. Fooprint radius.

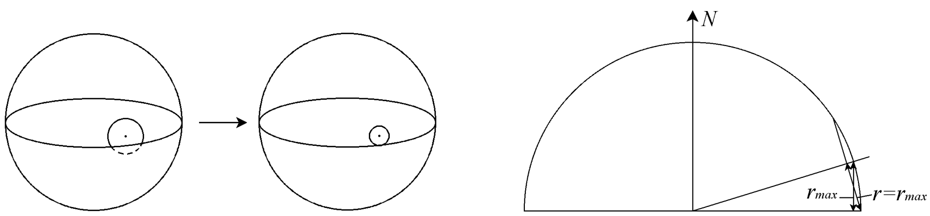

Fig. 7. Resizing the radius of a footprint so that it fits into the hemisphere.

2.4 Generation of an importance sampled direction

After the directional particle density estimation we decide stochastically according to p_{BSDF} if we select the outgoing direction Ψ according to the BSDF or with the footprints. If we decide to select it with the footprints then we stochastically choose one of the nearest neighbor particles of the selected hemisphere Ω. The probability of choosing a particle with contribution c_p is

$$p_P = \frac{c_P}{\sum_{p \in \Omega} c_P} . \tag{15}$$

We select Ψ in this particle´s footprint by drawing random numbers $u,v \in [0,r)$ until $u^2+v^2<r^2$. This gives us

$$\Psi = \left(u, v, \sqrt{1-u^2-v^2} \right) \tag{16}$$

in a local coordinate system with Ψ_p as third axis. After we have generated Ψ this way, or according to the BSDF with PDF contribution p_b, we have to weight its contribution to the path with $1/p$ to avoid bias. The value of the PDF for Ψ is

$$p = p_b(\Psi) + \sum_{p \in \Omega} p_f(\Psi) . \tag{17}$$

3 Implementation and Results

We compare our footprint importance sampling technique with classic photon map based importance sampling [Jen95] by applying these two methods to importance sampling of the view paths in a simple path tracer [Kaj86]. Despite the huge computation times that simple path tracing needs for simulating global illumination (see table 1), we have used it because it allows a straightforward comparison of these importance sampling techniques.

Both methods use a photon map which is generated in a preceding light pass. The memory consumption is not optimized in this implementation, because we have implemented it upon an existing rendering system[1], and for sake of simplicity we have used its default data types (it supports several different multi-spectral color representations, for our tests we have nevertheless used rgb colors). Therefore each photon contains several double values which results in 96 bytes per photon. Memory consumption was not a critical factor in our tests as we needed at maximum 49.6 Mb per photon map. Nevertheless the photon map could as well be stored in a more compact format [Jen95,Jen96a]. We have used p_{BSDF}=0.3 and k_p=50 for both

[1] www.artoolkit.org

importance sampling methods, k_g=32x32=1024, and k_r=6. We use these photon map importance sampling techniques at diffuse or soft glossy surfaces, whereas we use importance sampling according to the BSDF at specular or strong glossy surfaces.

In the first test scene the camera is located in a room which is only indirectly illuminated by the light that is reflected from another room, which contains a lightsource, through the small gap of a door. The difficulty in this scene is to guide the view paths through the small gap, which requires to target the gap precisely.

The second test scene is illuminated by a desk lamp. Direct illumination from the lightsource exists only inside the bright hard bordered spotlight on the desk and at the inside of the lampshade. The rest of the desk and the scene receives only indirect illumination. A major part of the indirect illumination comes from light which is reflected at the small lampshade. The difficulty in this scene is to guide the view paths to this small lampshade. The soft glossy desk surface contains a BSDF component with the Blinn reflection model, which contributes the major part of light that is reflected at the desk.

Footprint importance sampling can target the small regions, where most of the illumination in both scenes comes from, more precisely than classic photon map based importance sampling. Therefore, in the same computation time, footprint importance sampling results in less noise and in fewer spots, as can be seen in figure 8 and 9.

The photon maps that have been used for both importance sampling techniques are shown in figure 10. A more uniform photon distribution and therefore a smaller number of photons could be achieved by using density control [SW00].

Table 1. Statistics for figure 8-10, performed on a 900 MHz PC.

scene	no. photons	light pass	path tracing pass with footprint IS	path tracing pass with classic photon map IS
door	276266	78.2 s	7011 s	7332 s
desk	187065	7.0 s	39138 s	41904 s

4 Conclusion

We have presented a particle map based importance sampling technique that uses footprints of the nearest neighbor particles on the hemisphere to select an outgoing direction. In our results we have shown the efficiency of this new method, which is caused by the ability to target precisely into small highly contributing regions.

Future work includes the application of this technique to advanced global illumination methods that allow considerably faster computation times, and an adaptive control of some of the technique´s parameters.

Acknowledgments

This work has been supported by the Austrian Science Fund (FWF) project no. P13600-INF.

References

[BLS94] P. Blasi, B. Le Saëc, C. Schlick. An Importance Driven Monte-Carlo Solution to the Global Illumination Problem. Eurographics Workshop on Rendering 1994 pp. 173-183

[Bus97] E. Bustillo. A neuro-evolutionary unbiased global illumination algorithm. Eurographics Workshop on Rendering 1997

[DW94] P. Dutré, Y. D. Willems. Importance-driven Monte Carlo Light Tracing. Eurographics Workshop on Rendering 1994 pp. 185-194

[DW95] P. Dutré, Y. D. Willems. Potential-driven Monte Carlo Particle Tracing for Diffuse Environments with Adaptive Probability Functions. Eurographics Workshop on Rendering 1995 pp. 339-348

[Jen95] H. W. Jensen. Importance Driven Path Tracing using the Photon Map. Eurographics Workshop on Rendering 1995 pp. 359-369

[Jen96a] H. W. Jensen. Rendering Caustics on Non-Lambertian Surfaces. Graphics Interface 1996 pp. 116-121

[Jen96b] H. W. Jensen. Global Illumination using Photon Maps. Eurographics Workshop on Rendering 1996

[Kaj86] J. T. Kajiya. The Rendering Equation. SIGGRAPH 1986 pp. 143-150

[LW95] E. P. Lafortune, Y. D. Willems. A 5D Tree to Reduce the Variance of Monte Carlo Ray Tracing. Eurographics Workshop on Rendering 1995 pp. 11-20

[LF97] P. Lalonde, A. Fournier. Generating Reflected Directions from BRDF Data. EUROGRAPHICS 1997 pp. 293-300

[Lan91] B. Lange. The Simulation of Radiant Light Transfer with Stochastic Ray-Tracing. Eurographics Workshop on Rendering 1991

[LB94] B. Lange, M. Beyer. Rayvolution: An Evolutionary Ray Tracing Algorithm. Eurographics Workshop on Rendering 1994

[NN+96] A. Neumann, L. Neumann, P. Bekaert, Y. D. Willems, W. Purgathofer. Importance-driven Stochastic Ray Radiosity. Eurographics Workshop on Rendering 1996

[PP98] I. Peter, G. Pietrek. Importance Driven Construction of Photon Maps. Eurographics Workshop on Rendering 1998

[SS95] P. Schröder, W. Sweldens. Sherical Wavelets: Efficiently Representing Functions on the Sphere. SIGGRAPH 1995 pp. 161-172

[Shi92] P. Shirley. Nonuniform Random Point Sets via Warping. Graphic Gems III pp. 80-83, Academic Press 1992

[SW00] F. Suykens, Y. D. Willems. Density Control for Photon Maps. Eurographics Workshop on Rendering 2000 pp. 23-34

[SCP99] László Szirmay-Kalos, Balázs Csébfalvi, Werner Purgathofer. Importance driven quasi-random walk solution of the rendering equation. Computers & Graphics 23(2) 1999 pp. 203-211

[TN+98] R. F. Tobler, L. Neumann, M. Sbert, W. Purgathofer. A new Form Factor Analogy and its Application to Stochastic Global Illumination Algorithms. Eurographics Workshop on Rendering 1998

[UT97] C. Ureña, J. C. Torres. Improved Irradiance Computation by Importance Sampling. Eurographics Workshop on Rendering 1997

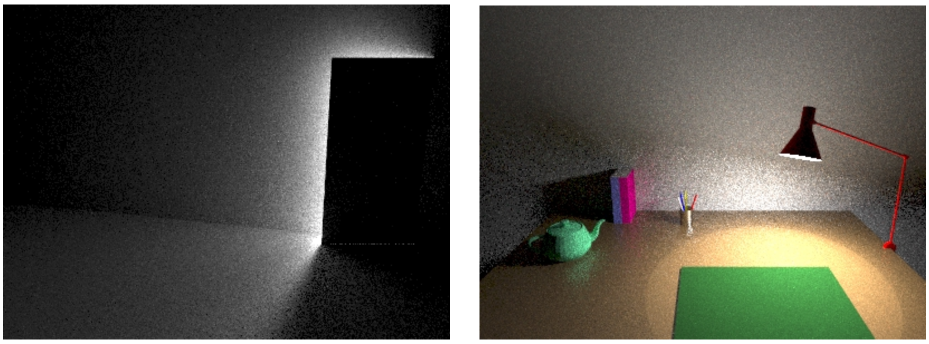

Fig. 8. Images rendered with path tracing with hemispherical photon footprint importance sampling.

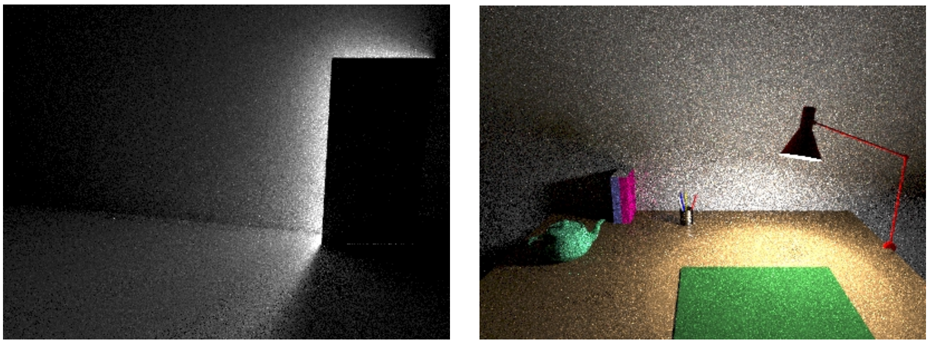

Fig. 9. Images rendered with path tracing with classic photon map importance sampling in the same computation time as figure 8.

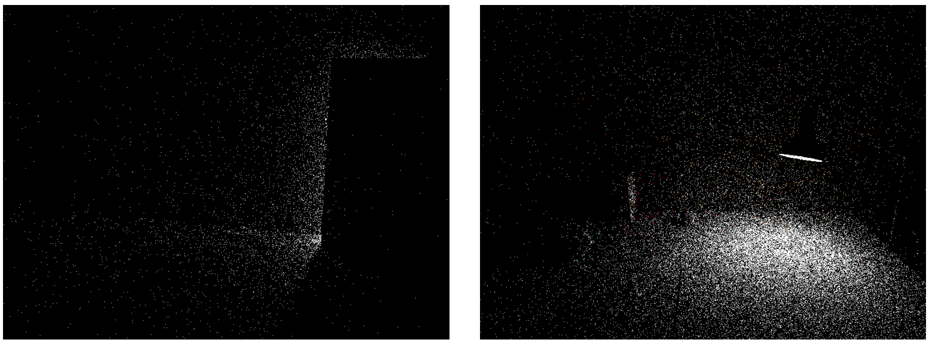

Fig. 10. Photon maps which were used for importance sampling in figure 8 and 9.