Real-time Kd-tree Based Importance Sampling of Environment Maps

Serkan Ergun[1] Murat Kurt[1] Aydın Öztürk[2]

[1]International Computer Institute
Ege University

[2]Department of Computer Engineering
Yaşar University

Spring Conference on Computer Graphics, 2012

Table of Contents

Environment maps

- Environment maps are commonly used for modeling natural lighting to create realistic images.
- Complex real-world illumination can be represented efficiently by environment maps.
- High quality rendering of scenes under image-based lighting requires efficient sampling strategies.

Sampling strategies

- Environment map sampling
- Bidirectional Reflectance Distribution Function (*BRDF*) sampling
- Product sampling
- Multiple Importance Sampling (*MIS*) [33]

Motivation

- Most of the environment map sampling methods can be used in Monte Carlo simulations but they are not suitable for real-time rendering.
- We introduced a new method based on Kd-tree structure that can be used in real-time Monte-Carlo simulations.

Table of Contents

Related Work

Environment map sampling

- Stratified sampling. [2]
- Hierarchical sampling. [7]
- Structural importance sampling. [1]
- Blue noise sampling. [25]
- Interleaved sampling. [17]
- Inversion of the CDF. [31, 19]

Product Sampling

- Bidirectional importance sampling. [3]
- Using wavelet transforms. [5, 12]
- Using spherical harmonics. [13]

Direct illumination from environment maps

- Prefiltered environment maps. [14, 15]
- Irradiance environment maps. [27]
- Frequency space environment map rendering. [28]
- Precomputed radiance transfer [32]
- Non-linear wavelet lighting [21]
- Real-time filtered importance sampling [18]

Inversion method

- Inversion method is based on the observation that cumulative distribution functions (CDFs) range uniformly over the interval $(0, 1)$
- If u is a uniform random number on $(0, 1)$, then using $X = F^{-1}(u)$ generates a random number X from a distribution with specified CDF F.

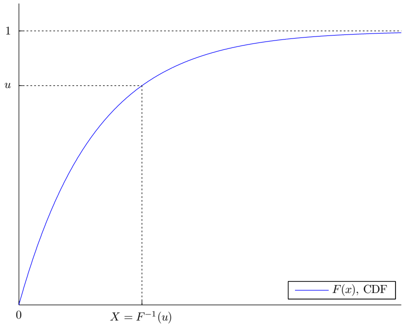

Table of Contents

Overview I

- An environment map defined as a $w \times h$ rectangular block can be treated as a bivariate probability distribution after normalization.

Overview I

- An environment map defined as a $w \times h$ rectangular block can be treated as a bivariate probability distribution after normalization.
- This rectangular block is divided into sub-blocks using a Kd-tree structure.

Overview I

- An environment map defined as a $w \times h$ rectangular block can be treated as a bivariate probability distribution after normalization.
- This rectangular block is divided into sub-blocks using a Kd-tree structure.
- Empirical probabilities corresponding to each sub-block are obtained as sum of the normalized pixel intensities within each sub-block.

Overview I

- An environment map defined as a $w \times h$ rectangular block can be treated as a bivariate probability distribution after normalization.
- This rectangular block is divided into sub-blocks using a Kd-tree structure.
- Empirical probabilities corresponding to each sub-block are obtained as sum of the normalized pixel intensities within each sub-block.
- These empirical probabilities are then sorted in descending order and corresponding block indices are assigned in increasing order.

Overview I

- An environment map defined as a $w \times h$ rectangular block can be treated as a bivariate probability distribution after normalization.
- This rectangular block is divided into sub-blocks using a Kd-tree structure.
- Empirical probabilities corresponding to each sub-block are obtained as sum of the normalized pixel intensities within each sub-block.
- These empirical probabilities are then sorted in descending order and corresponding block indices are assigned in increasing order.

Overview II

- The plots of these probabilities against sub-block indices can be considered as empirical distribution of these sub-block indices.

Overview II

- The plots of these probabilities against sub-block indices can be considered as empirical distribution of these sub-block indices.
- Representing the empirical distribution of sub-blocks in this way provides a good ground for modeling this empirical distribution by a simple probability density function (pdf).

Overview II

- The plots of these probabilities against sub-block indices can be considered as empirical distribution of these sub-block indices.
- Representing the empirical distribution of sub-blocks in this way provides a good ground for modeling this empirical distribution by a simple probability density function (pdf).
- The resulting estimated model can be used to generate samples for incoming light.

Table of Contents

Kd-tree Construction

- Keep a list of current sub-blocks.

Kd-tree Construction

- Keep a list of current sub-blocks.
- Select a sub-block to split further.

Kd-tree Construction

- Keep a list of current sub-blocks.
- Select a sub-block to split further.
- Choose a splitting plane for this sub-block.

Kd-tree Construction

- Keep a list of current sub-blocks.
- Select a sub-block to split further.
- Choose a splitting plane for this sub-block.
- Split the sub-block and add the resulting sub-blocks into the list.

Kd-tree Construction

- Keep a list of current sub-blocks.
- Select a sub-block to split further.
- Choose a splitting plane for this sub-block.
- Split the sub-block and add the resulting sub-blocks into the list.
- Repeat until a predefined number of sub-blocks are created.

Kd-tree Construction

- Keep a list of current sub-blocks.
- Select a sub-block to split further.
- Choose a splitting plane for this sub-block.
- Split the sub-block and add the resulting sub-blocks into the list.
- Repeat until a predefined number of sub-blocks are created.

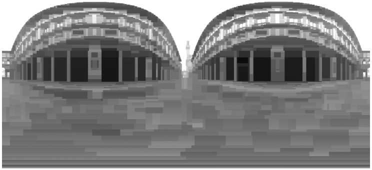

Choosing a sub-block for splitting

- Choosing the most convenient sub-block for splitting needs a special handling.
- We proceed to choose the sub-block having the largest intensity variation first.
- Various measures of variation can be used for this purpose.
- We tested for range, variance and sum of squared error measures respectively.

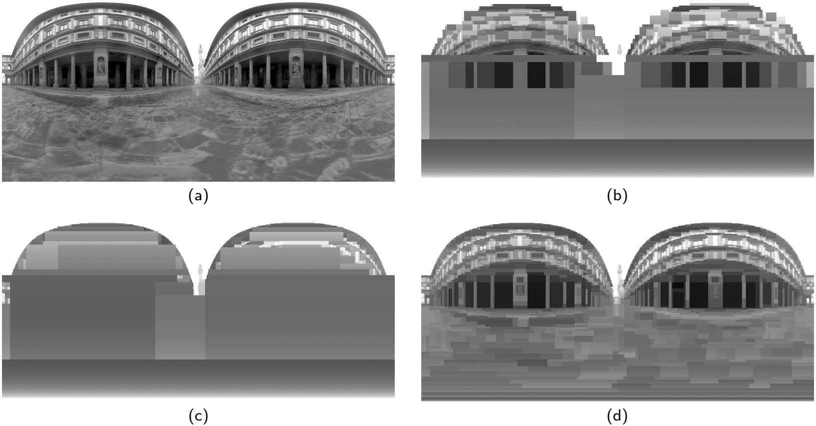

Figure 1 : Rendered environment maps (Uffizi) (a): 2048×1024 resolution environment map requiring 8MBs of memory. (b), (c), (d): Same environment map compressed to 48KB (1:170 compression) using range, variance, and SSE criteria, respectively. Environment maps used in this work are a courtesy of Debevec.

SSE

In this work, we propose to use SSE which is defined by

$$SSE = \sum_{i=1}^{w_b} \sum_{j=1}^{h_b} \left(f_{ij} - \bar{f} \right)^2, \tag{1}$$

as a selection criterion where $\bar{f}$ is the sub block mean, w_b and h_b are the sub-block dimensions.

Splitting a selected sub-block

The splitting plane position is determined in such a way that the pooled variance [16] of the children blocks is minimum. It can be shown that minimization of the pooled variance can be reduced to maximizing the sum of squares of sub-block totals divided by their respective number of pixels.

Table of Contents

Approximating the Empirical Distribution

Empirical block probabilities are sorted in descending order. Therefore, the corresponding pdf is expected to be an exponential type distribution. We approximate this pdf by the following monotonically decreasing function

$$p(x) = \frac{1}{\log\left(1 + \frac{n}{\alpha}\right)(\alpha + x)}, 0 \leq x \leq n, \qquad (2)$$

where n is the total number of sub-blocks in the kd-tree, and α is the parameter of the distribution.

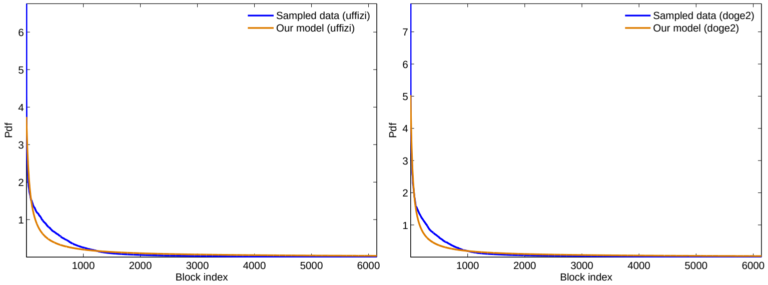

Figure 2 : Empirical pdfs of various environment maps and fitted analytical pdfs for uffizi(left) and doge2(right) environment maps.

Summary of preprocessing steps

- Construct the Kd-tree.
- Average sub-block intensities in the kd-tree are sorted in descending order to obtain the empirical pdf of the block indices.
- The empirical pdf is then approximated by an analytical model.
- The parameter of the pdf and the sub-block bounds are stored for use in the sampling procedure.

Table of Contents

Block selection

- Our importance sampling strategy simply consists of generating sub-block indices first and then generating incoming light direction within this block.
- Sub-block indices can be generated using the well-known probability integral transformation method of obtaining random samples from a known distribution.
- The following inverse function of the cdf is used for generating sub-block indices:

$$x = P^{-1}(\xi) = \alpha \left(\left(1 + \frac{n}{\alpha} \right)^{\xi} - 1 \right), \qquad (3)$$

where ξ is a uniform (0,1) random variable.

Light direction sampling

- We know that the approximated PDF is uniform within the selected sub-block.
- Light direction can easily be generated by uniformly sampling the elevation and azimuth angles.
- Within the bounds of the sub-block, we generate two uniform random variables corresponding to elevation and azimuth angles to obtain a random incoming light direction.

Summary of environment map sampling

- Generate three random variables: ξ_1, ξ_2, ξ_3.
- Select the corresponding block index $x = P^{-1}(\xi_1)$.
- Read the bounds of the selected sub-block.
- Generate elevation and azimuth angles uniformly within the bounds of the selected sub-block using ξ_2, ξ_3.
- The probability of this sample can be computed with $p(x)/Area(\text{sub-block})$.

Table of Contents

- Real-time rendering implementations of our method and inversion method were made using OpenSceneGraph [4], NVIDIA CUDA [24], and Random123 [30] libraries.
- Our method, and the inversion method were also implemented using Physically Based Rendering Toolkit (PBRT) [26] for off-line renderings.
- All programs were executed on an Intel Core i7-920 (2.67 GHz) with 12GBs of RAM and NVIDIA GeForce GTX 480 GPU.

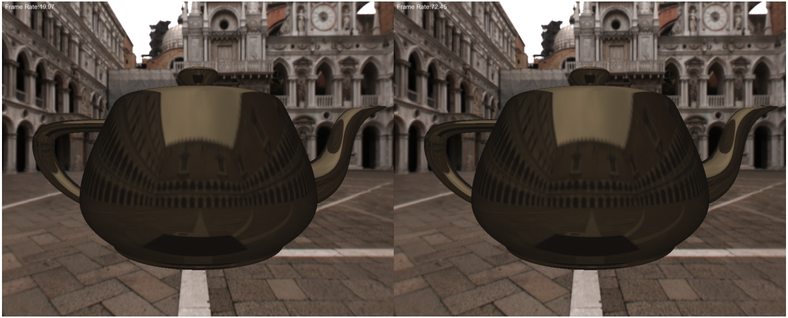

Figure 3 : Comparison of inversion method (left) and our method (right) in real-time rendering. In this scene, the chrome-steel teapot has been rendered with both methods using 16 samples/pixel for testing real-time rendering performance.

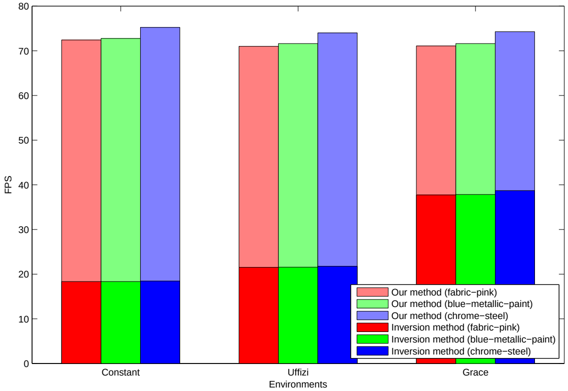

Figure 4 : Comparison of our method and inversion method in real-time rendering. Both of the methods were rendered with 16 samples/pixel. The FPS rates are measured under different environment maps.

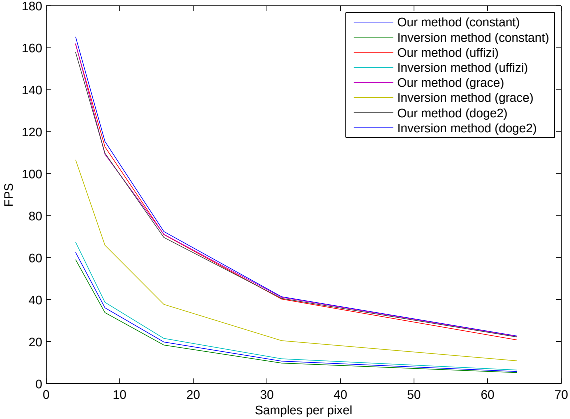

Figure 5 : FPS rates of our method and the inversion method for different sample sizes and environment maps.

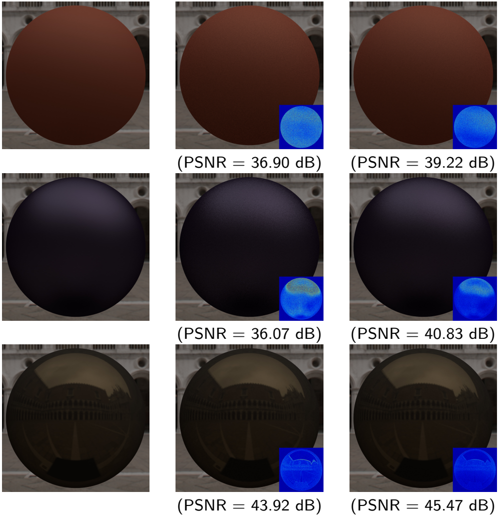

Figure 6 : Rendered spheres based on different materials and different sampling methods. rows show different materials and columns show reference images, inversion method, and kd-tree method respectively. Insets show the scaled difference between the methods and reference images.

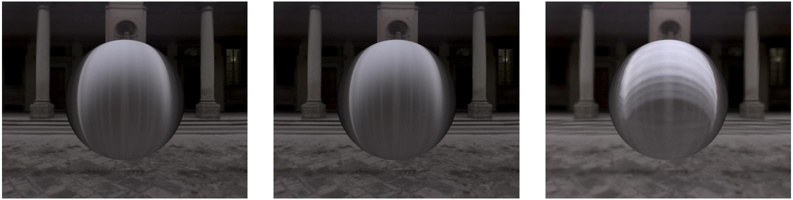

(a) (b) (c)

Figure 7 : Rendered spheres using anisotropic Ward BRDF model with parameters $\alpha_x = 0.5, \alpha_y = 0.001$. (a) Reference image, (b) our kd-tree based importance sampling method (c) Křivánek and Colbert's real-time filtered importance sampling method.

Table of Contents

Our kd-tree based importance sampling method:

- is more suitable than inversion method for real-time rendering;
- can handle every type of material including anisotropic materials;
- can be extended to multi-dimensional functions such as BRDFs.

Thank You!

Questions?

[1] Sameer Agarwal, Ravi Ramamoorthi, Serge Belongie, and Henrik Wann Jensen.
Structured importance sampling of environment maps.
ACM Trans. Graph., 22:605–612, July 2003.

[2] James Arvo.
Stratified sampling of 2-manifolds.
In *State of the Art in Monte Carlo Ray Tracing for Realistic Image Synthesis, SIGGRAPH 2001 Course Notes*, volume 29, August 2001.

[3] David Burke, Abhijeet Ghosh, and Wolfgang Heidrich.
Bidirectional importance sampling for direct illumination.
In *Proceedings of the Eurographics Symposium on Rendering*, pages 147–156. Eurographics Association, 2005.

[4] Don Burns and Robert Osfield.
Open scene graph a: Introduction, b: Examples and applications.
In *Proceedings of the IEEE Virtual Reality 2004*, page 265, Washington, DC, USA, 2004. IEEE Computer Society.

[5] Petrik Clarberg, Wojciech Jarosz, Tomas Akenine-Möller, and Henrik Wann Jensen.
Wavelet importance sampling: efficiently evaluating products of complex functions.
ACM Trans. Graph., 24:1166–1175, July 2005.

[6] Paul Debevec.
Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography.
In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '98, pages 189–198, New York, NY, USA, 1998. ACM.

[7] Paul Debevec.
A median cut algorithm for light probe sampling.
In *ACM SIGGRAPH 2005 Posters*, SIGGRAPH '05, New York, NY, USA, 2005. ACM.

[8] Frédéric Devernay.
C/C++ Minpack, January 2012.
http://devernay.free.fr/hacks/cminpack/cminpack.html.

[9] George S. Fishman.
Monte Carlo: Concepts, algorithms, and applications.
Springer Series in Operations Research. Springer-Verlag, New York, 1996.

[10] Ned Greene.
Environment mapping and other applications of world projections.
IEEE Comput. Graph. Appl., 6:21–29, November 1986.

[11] Wolfgang Heidrich and Hans-Peter Seidel.
Realistic, hardware-accelerated shading and lighting.
In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '99, pages 171–178, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.

[12] Hao-da Huang, Yanyun Chen, Xing Tong, and Wen-cheng Wang.
Incremental wavelet importance sampling for direct illumination.
In *Proceedings of the 2007 ACM symposium on Virtual reality software and technology*, VRST '07, pages 149–152, New York, NY, USA, 2007. ACM.

[13] Wojciech Jarosz, Nathan A. Carr, and Henrik Wann Jensen.
Importance sampling spherical harmonics.
Computer Graphics Forum, 28(2):577–586, 2009.

[14] Jan Kautz and Michael D. McCool.
Approximation of glossy reflection with prefiltered environment maps.
In *In Graphics Interface*, pages 119–126, 2000.

[15] Jan Kautz, Pere-Pau Vázquez, Wolfgang Heidrich, and Hans-Peter Seidel.
Unified approach to prefiltered environment maps.

In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000*, pages 185–196, London, UK, 2000. Springer-Verlag.

[16] P. R. Killeen.
An alternative to null-hypothesis significance tests.
Psychol Sci, 16(5):345–353, 2005.

[17] Thomas Kollig and Alexander Keller.
Efficient illumination by high dynamic range images.
In *Proceedings of the 14th Eurographics workshop on Rendering*, EGRW '03, pages 45–50, Aire-la-Ville, Switzerland, Switzerland, 2003. Eurographics Association.

[18] Jaroslav Křivánek and Mark Colbert.
Real-time shading with filtered importance sampling.
Computer Graphics Forum, 27(4):1147–1154, 2008.

[19] Jason Lawrence, Szymon Rusinkiewicz, and Ravi Ramamoorthi.
Adaptive numerical cumulative distribution functions for efficient importance sampling.
In Oliver Deussen, Alexander Keller, Kavita Bala, Philip Dutré, Dieter W. Fellner, and Stephen N. Spencer, editors, *Proceedings of the Eurographics Symposium on Rendering*, pages 11–20. Eurographics Association, 2005.

[20] Michael D. McCool and Peter K. Harwood.
Probability trees.
In *Graphics Interface'97*, pages 37–46, 1997.

[21] Ren Ng, Ravi Ramamoorthi, and Pat Hanrahan.
All-frequency shadows using non-linear wavelet lighting approximation.
ACM Trans. Graph., 22:376–381, July 2003.

[22] Addy Ngan, Frédo Durand, and Wojciech Matusik.
Experimental analysis of brdf models.

In *Proceedings of the Eurographics Symposium on Rendering*, pages 117–226. Eurographics Association, 2005.

[23] F E Nicodemus, J C Richmond, J J Hsia, I W Ginsberg, and T Limperis.
Geometrical considerations and nomenclature for reflectance.
Science And Technology, 160(October):1–52, 1977.

[24] NVIDIA.
CUDA introduction page, January 2012.
http://developer.nvidia.com/category/zone/cuda-zone.

[25] Victor Ostromoukhov, Charles Donohue, and Pierre-Marc Jodoin.
Fast hierarchical importance sampling with blue noise properties.
ACM Trans. Graph., 23:488–495, August 2004.

[26] Matt Pharr and Greg Humphreys.
Physically Based Rendering, Second Edition: From Theory To Implementation.
Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2nd edition, 2010.

[27] Ravi Ramamoorthi and Pat Hanrahan.
An efficient representation for irradiance environment maps.
In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '01, pages 497–500, New York, NY, USA, 2001. ACM.

[28] Ravi Ramamoorthi and Pat Hanrahan.
Frequency space environment map rendering.
ACM Trans. Graph., 21:517–526, July 2002.

[29] Iain E. Richardson.
Video Codec Design: Developing Image and Video Compression Systems.
John Wiley & Sons, Inc., New York, NY, USA, 2002.

[30] John K. Salmon, Mark A. Moraes, Ron O. Dror, and David E. Shaw.

Parallel random numbers: as easy as 1, 2, 3.
In *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '11, pages 16:1–16:12, New York, NY, USA, 2011. ACM.

[31] Adrian Secord, Wolfgang Heidrich, and Lisa Streit.
Fast primitive distribution for illustration.
In *Proceedings of the 13th Eurographics workshop on Rendering*, EGRW '02, pages 215–226, Aire-la-Ville, Switzerland, Switzerland, 2002. Eurographics Association.

[32] Peter-Pike Sloan, Jan Kautz, and John Snyder.
Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments.
ACM Trans. Graph., 21:527–536, July 2002.

[33] Eric Veach.
Robust monte carlo methods for light transport simulation.
PhD thesis, Stanford, CA, USA, 1998.
AAI9837162.

[34] Rui Wang and Oskar Åkerlund.
Bidirectional importance sampling for unstructured direct illumination.
Computer Graphics Forum, 28(2):269–278, 2009.

[35] Rui Wang, Rui Wang, Kun Zhou, Minghao Pan, and Hujun Bao.
An efficient gpu-based approach for interactive global illumination.
ACM Trans. Graph., 28(3):91:1–91:8, July 2009.

[36] Gregory J. Ward.
Measuring and modeling anisotropic reflection.
In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '92, pages 265–272, New York, NY, USA, 1992. ACM.