EGSR 2008

Exploiting Visibility Correlation in Direct Illumination

Petrik Clarberg **Tomas Akenine-Möller**

Lund University
Sweden

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Direct Illumination

Rendering equation: $L_{\text{out}} = \int_{\Omega} L\, B\, V\, d\omega$

[Kajiya86]

Monte Carlo:

$$L_{\text{out}} = \frac{1}{N} \sum LBV$$

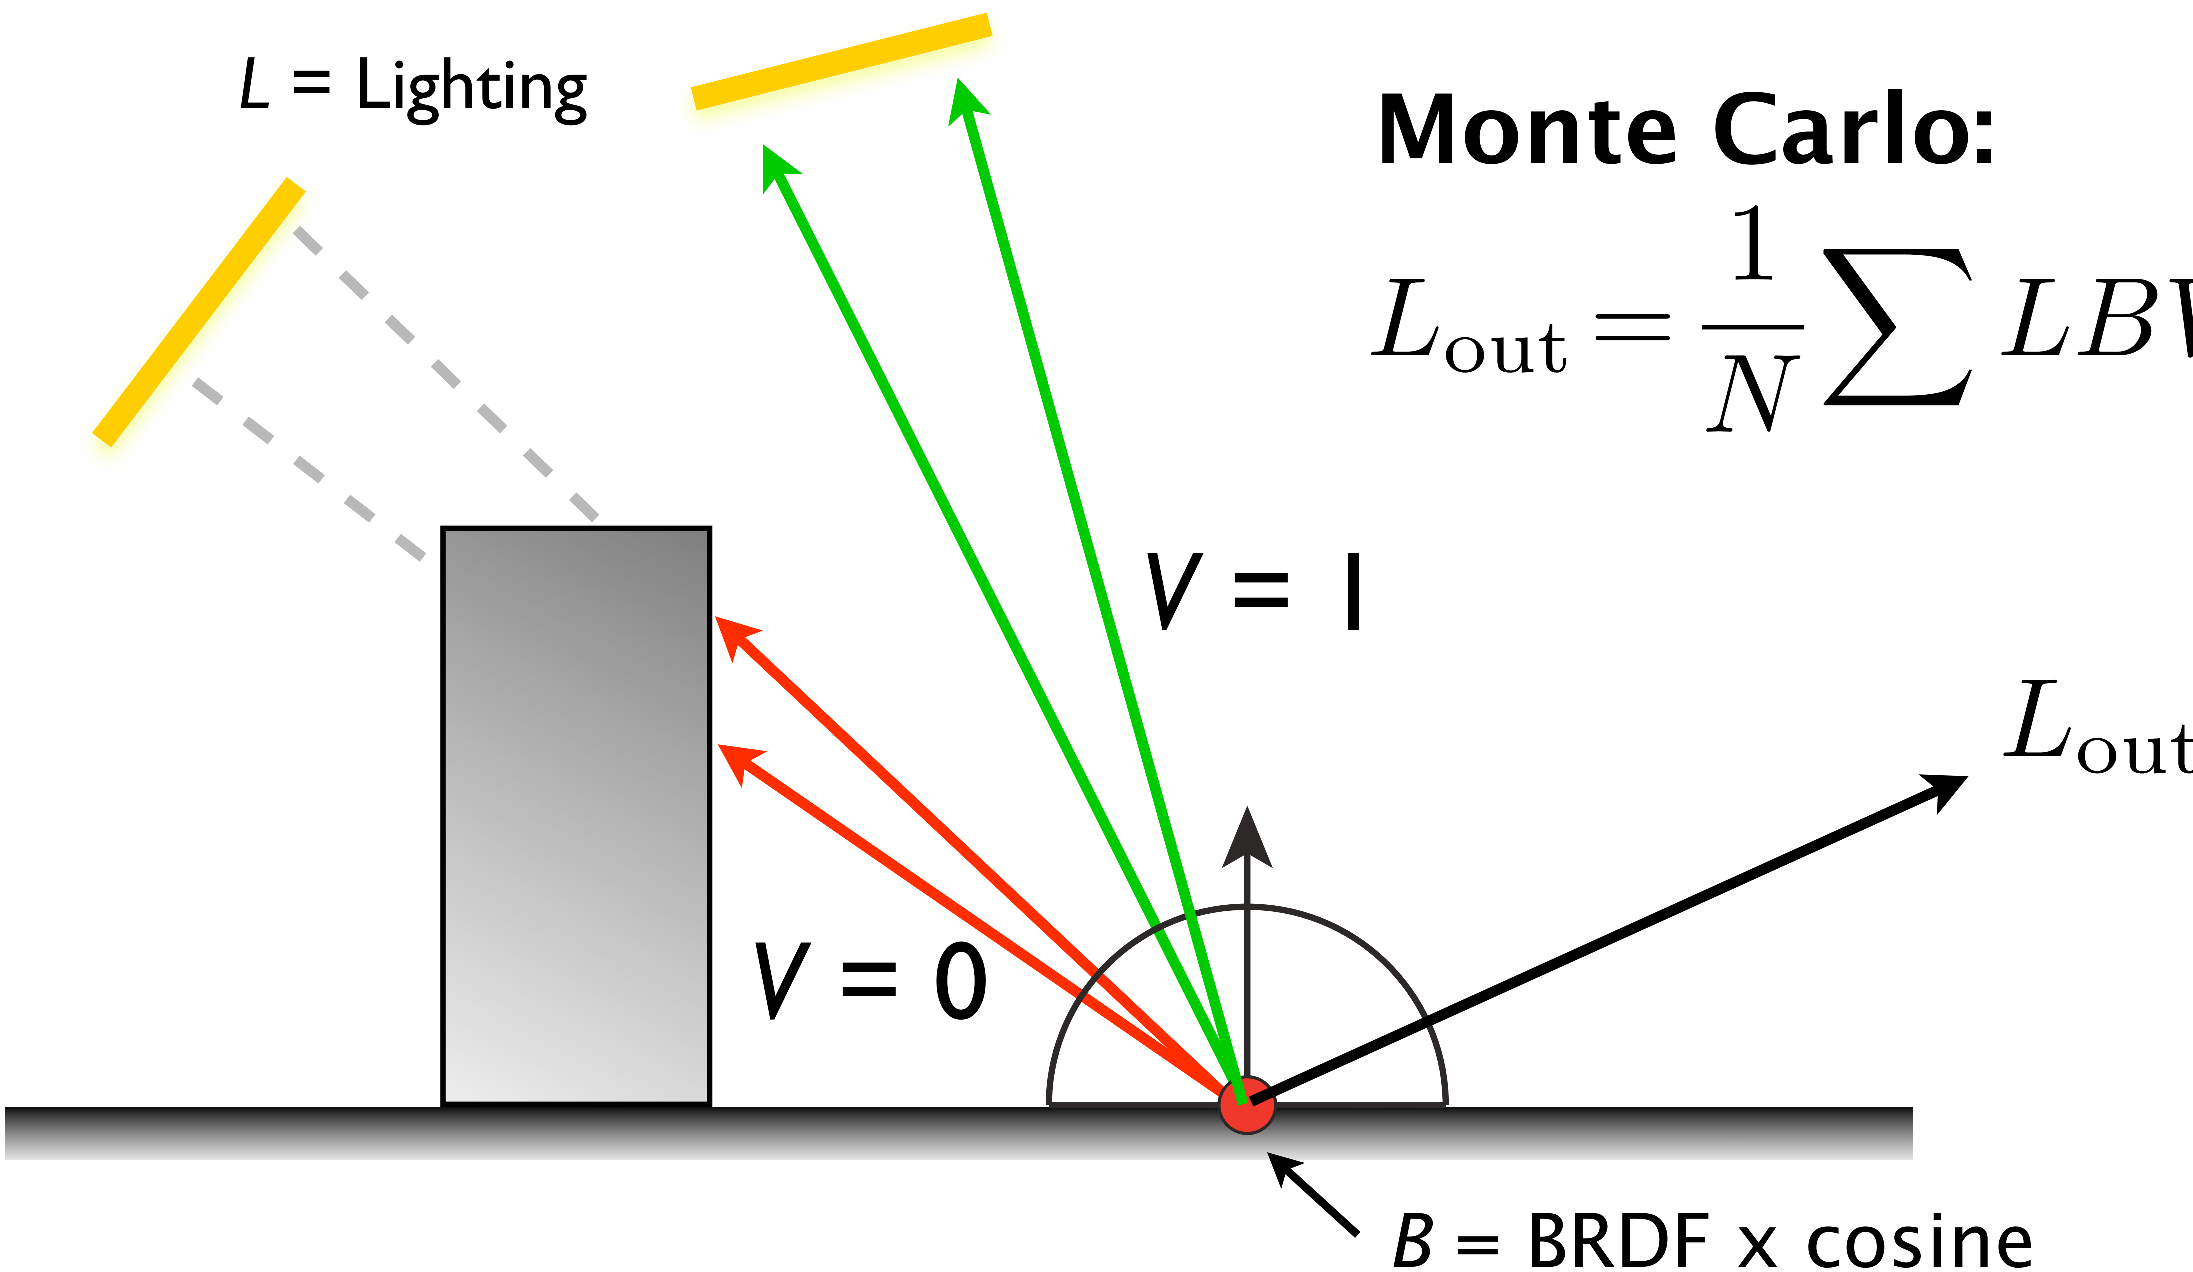

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Importance Sampling

Send rays in directions with large *LBV*

But, *V* is unknown so hard to include in sampling

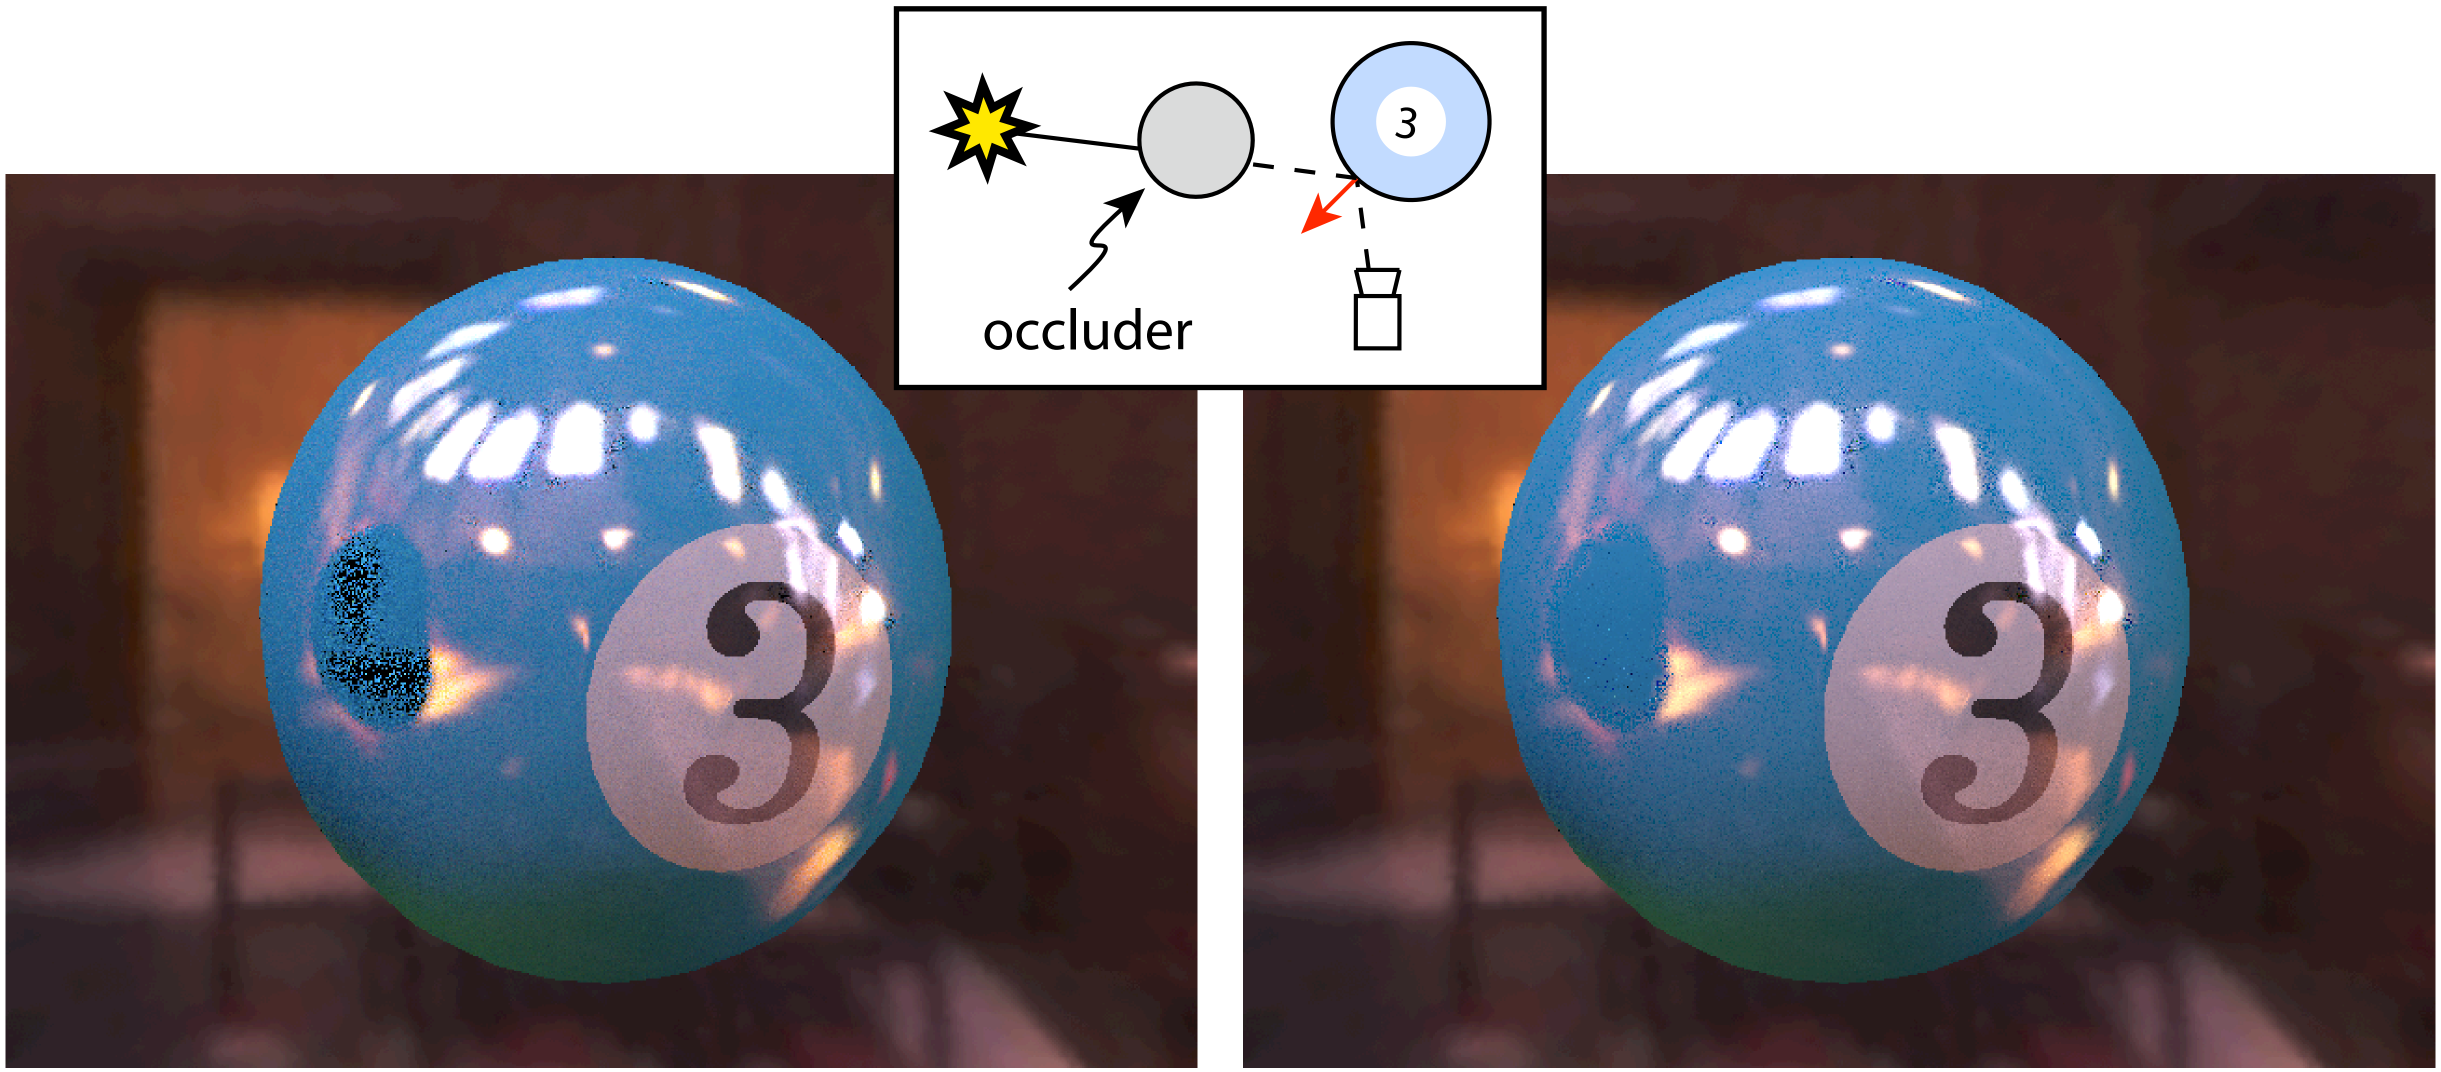

LB (Light x BRDF) sampling

Our algorithm

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Related Work

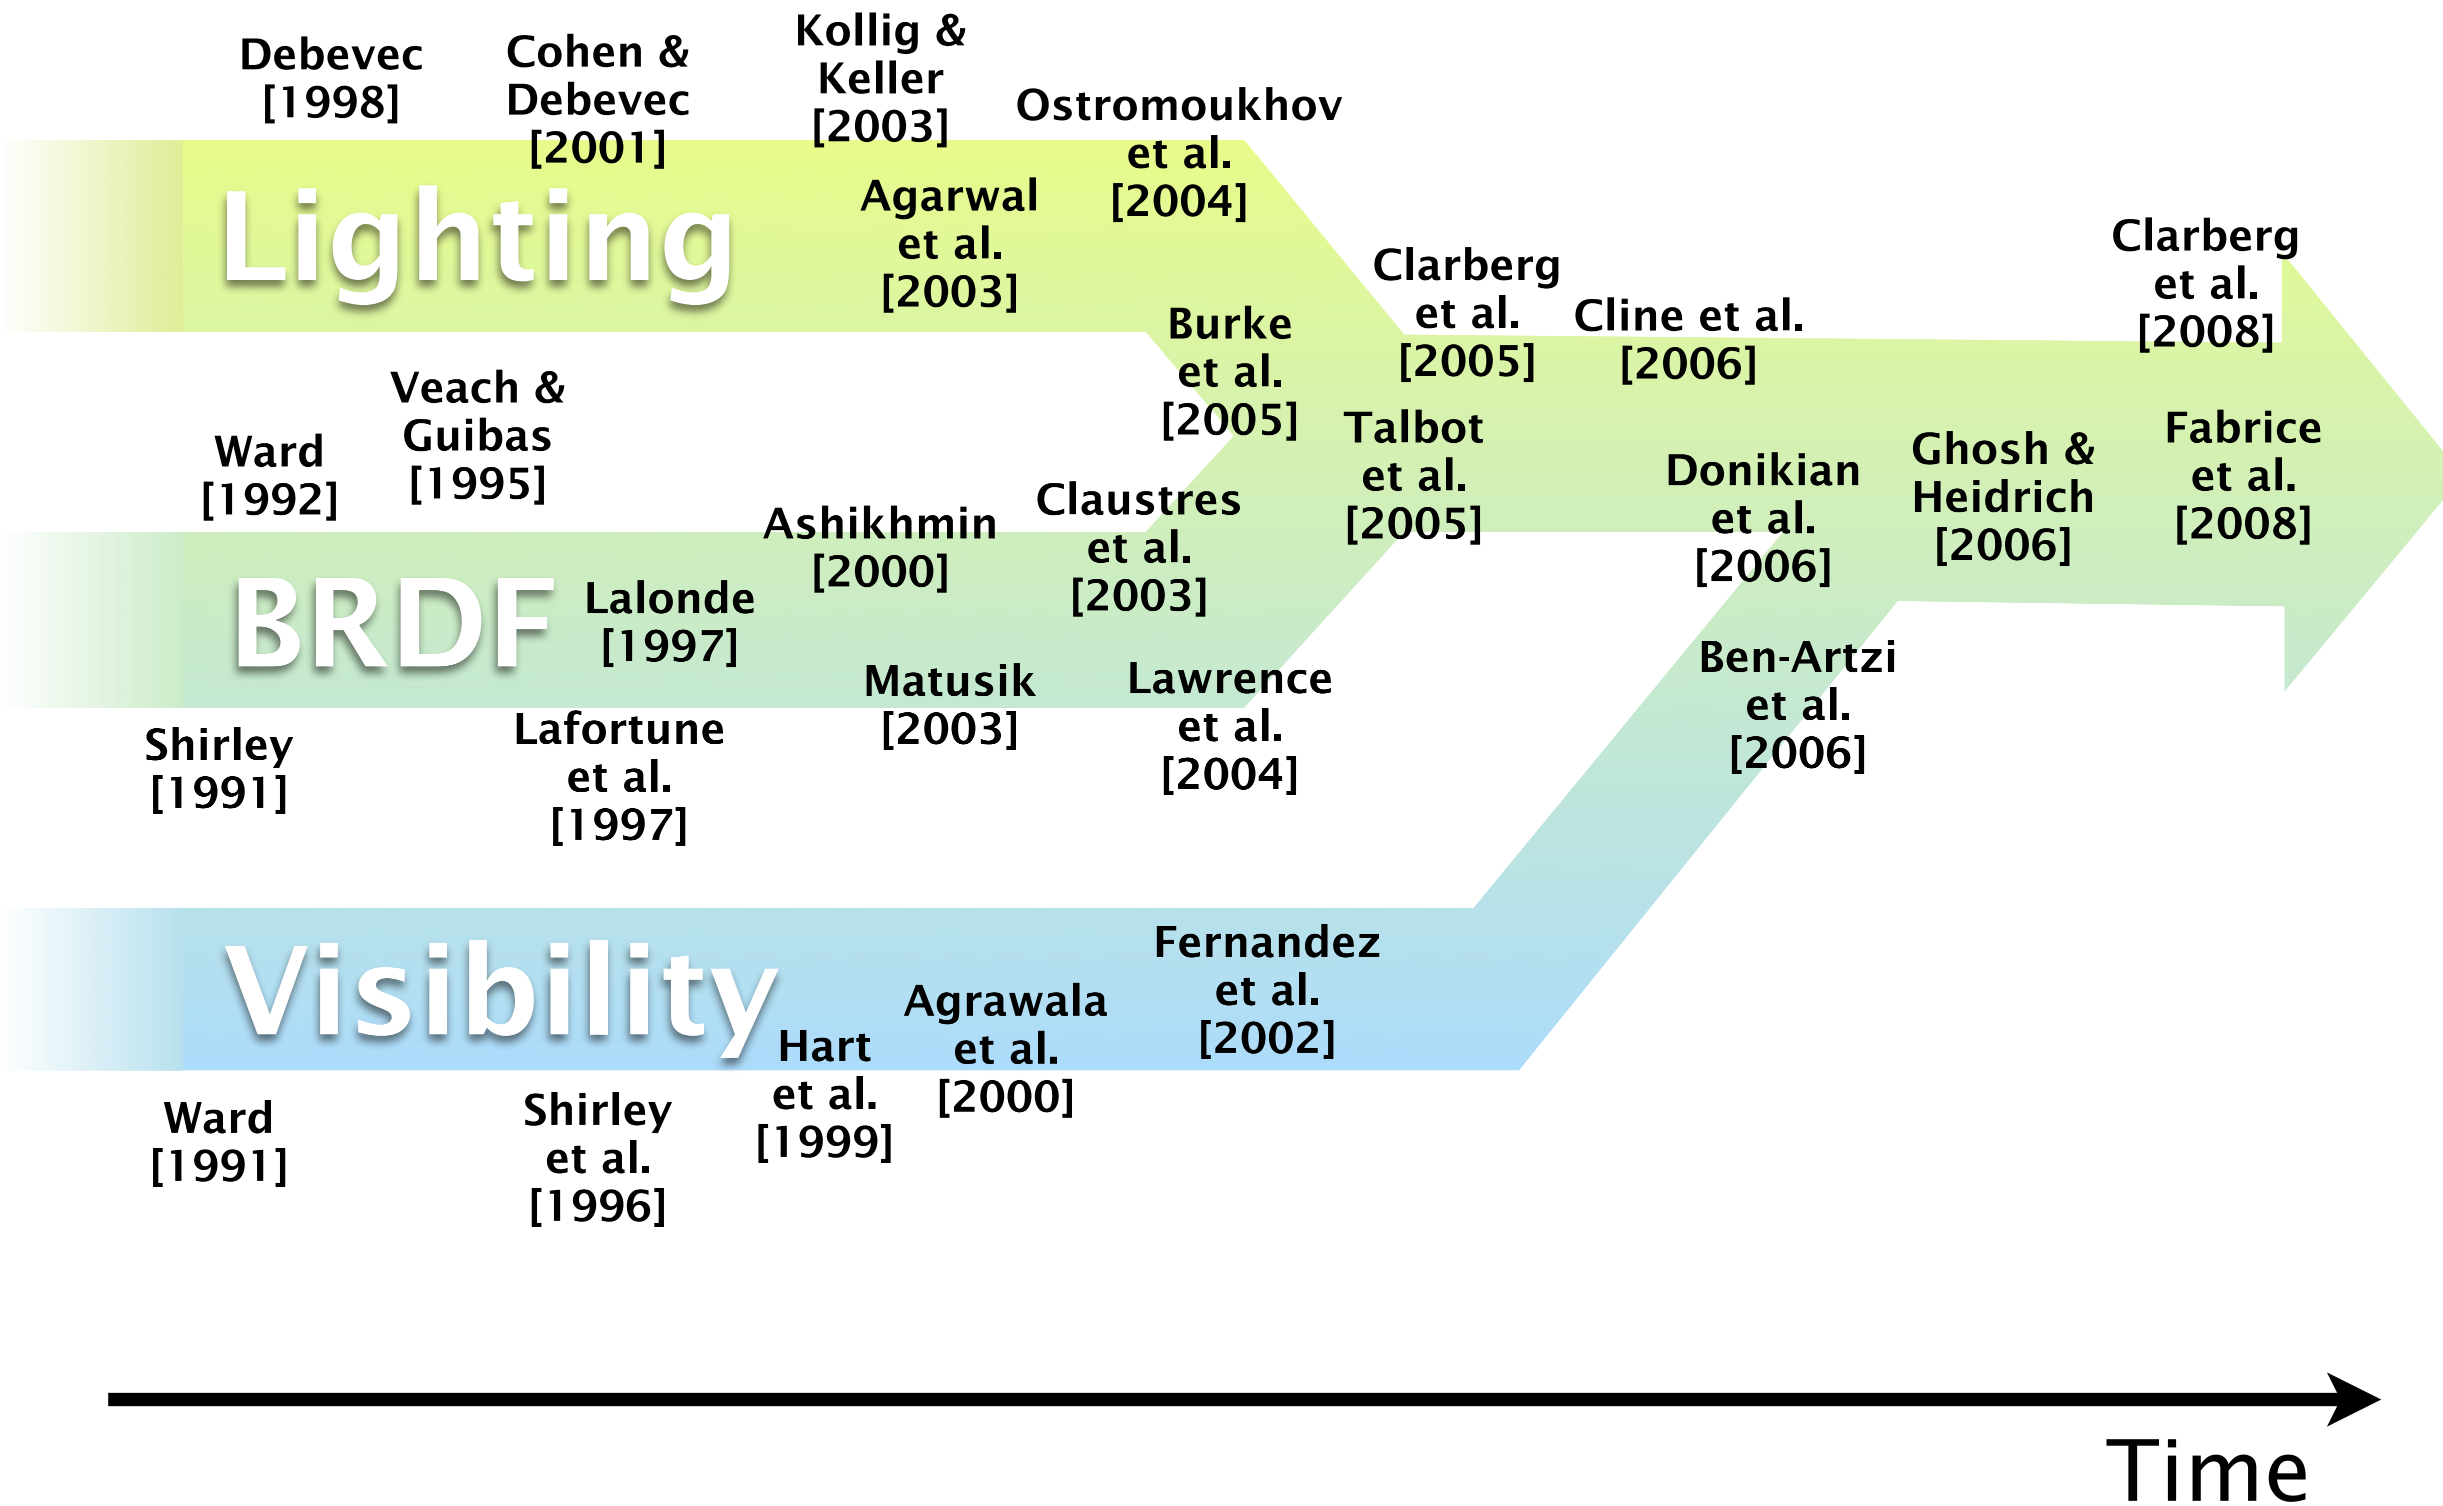

Outline

1. Visibility Measurements

Statistical correlation of the visibility function

2. The Visibility Cache

Similar to radiance caching, but for visibility

Applications:

- Reduce noise in Monte Carlo ray tracing
- Ambient occlusion
- Fast lighting preview (aka PRT)

Petrik Clarberg and Tomas Akenine-Möller, “*Exploiting Visibility Correlation in Direct Illumination*”, Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

1. Measurements

Large number of random pairs of positions

d = Distance

θ = Normal difference

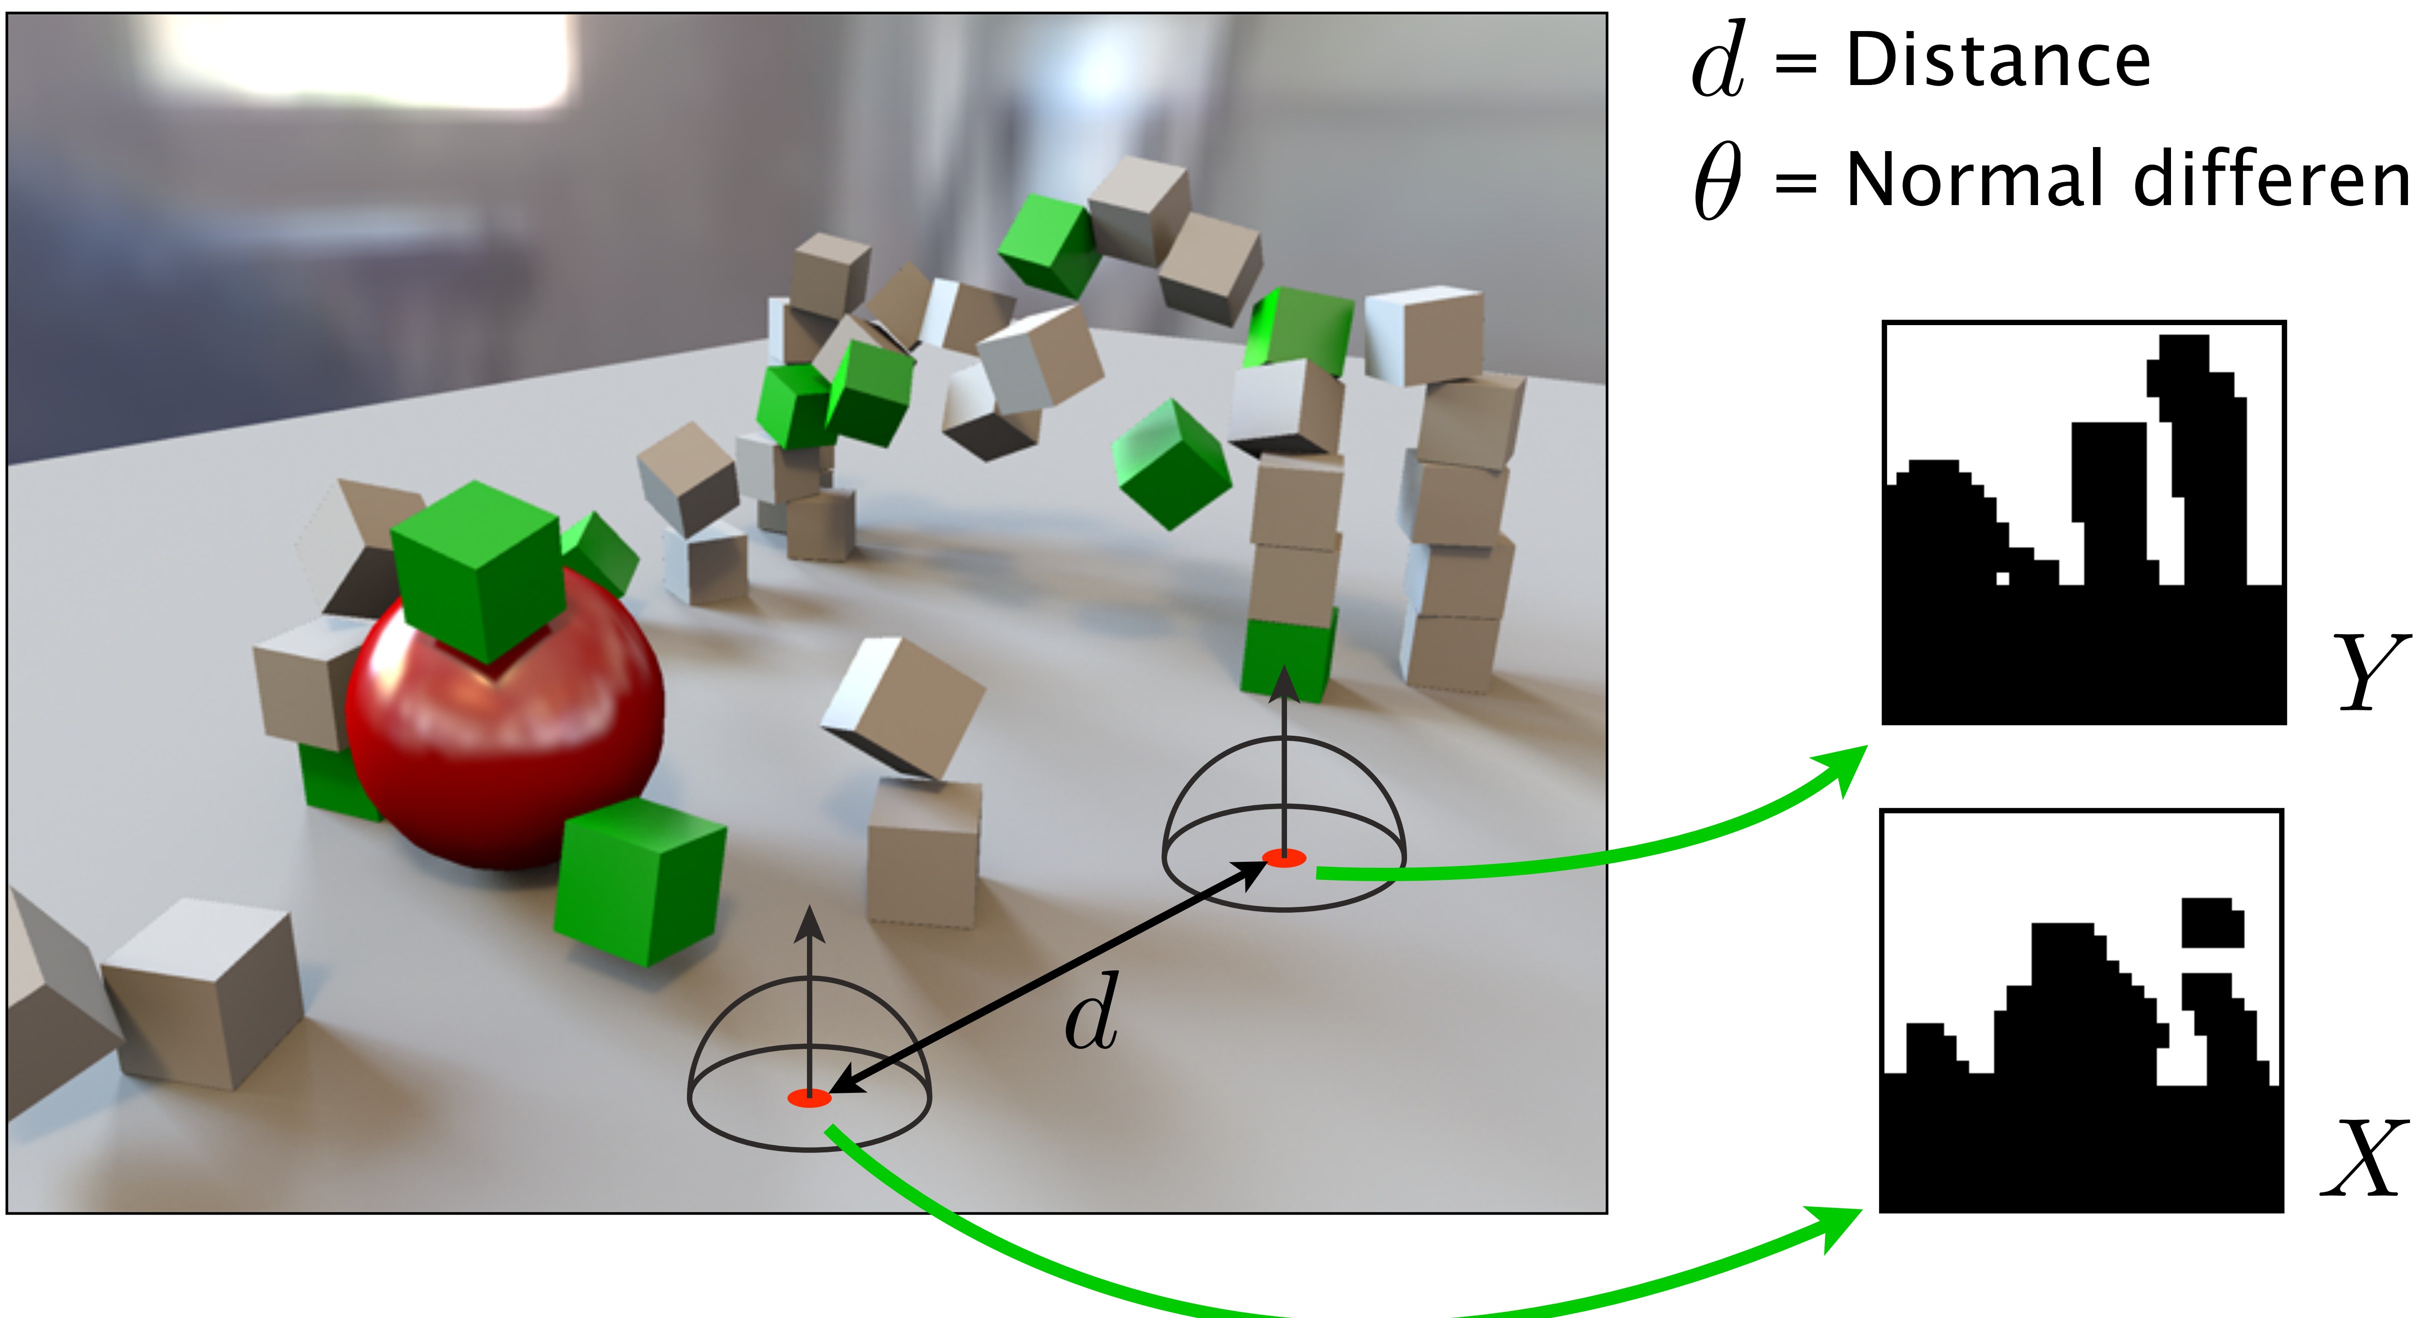

Petrik Clarberg and Tomas Akenine-Möller, “*Exploiting Visibility Correlation in Direct Illumination*”, Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Visibility Correlation

We see the visibility maps as distributions

Compute the correlation coefficient: $\rho_{X,Y}$
(Pearson's product-moment coeff.)

Examples:

$\rho\Big($ 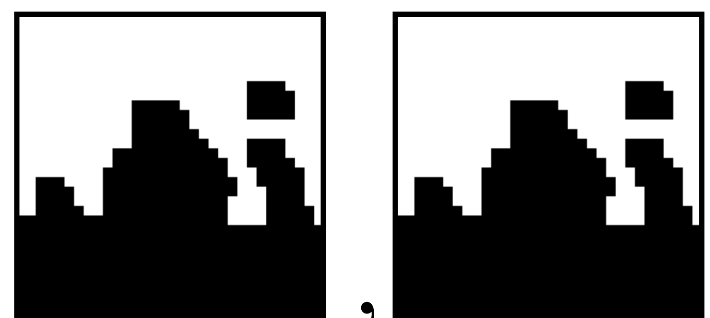 $\Big) = 1$

$\rho\Big($  $\Big) = -1$

$\rho\Big($ 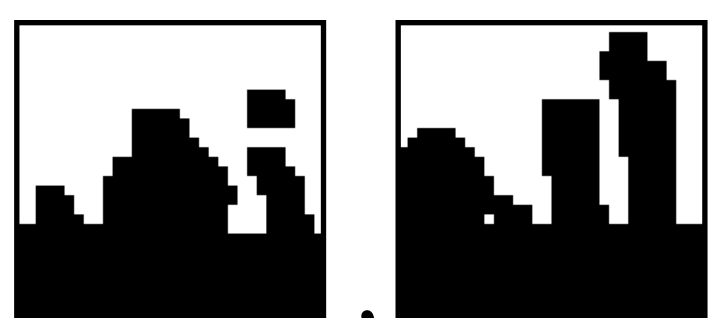 $\Big) = 0.861$

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Scene (a)

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Correlation

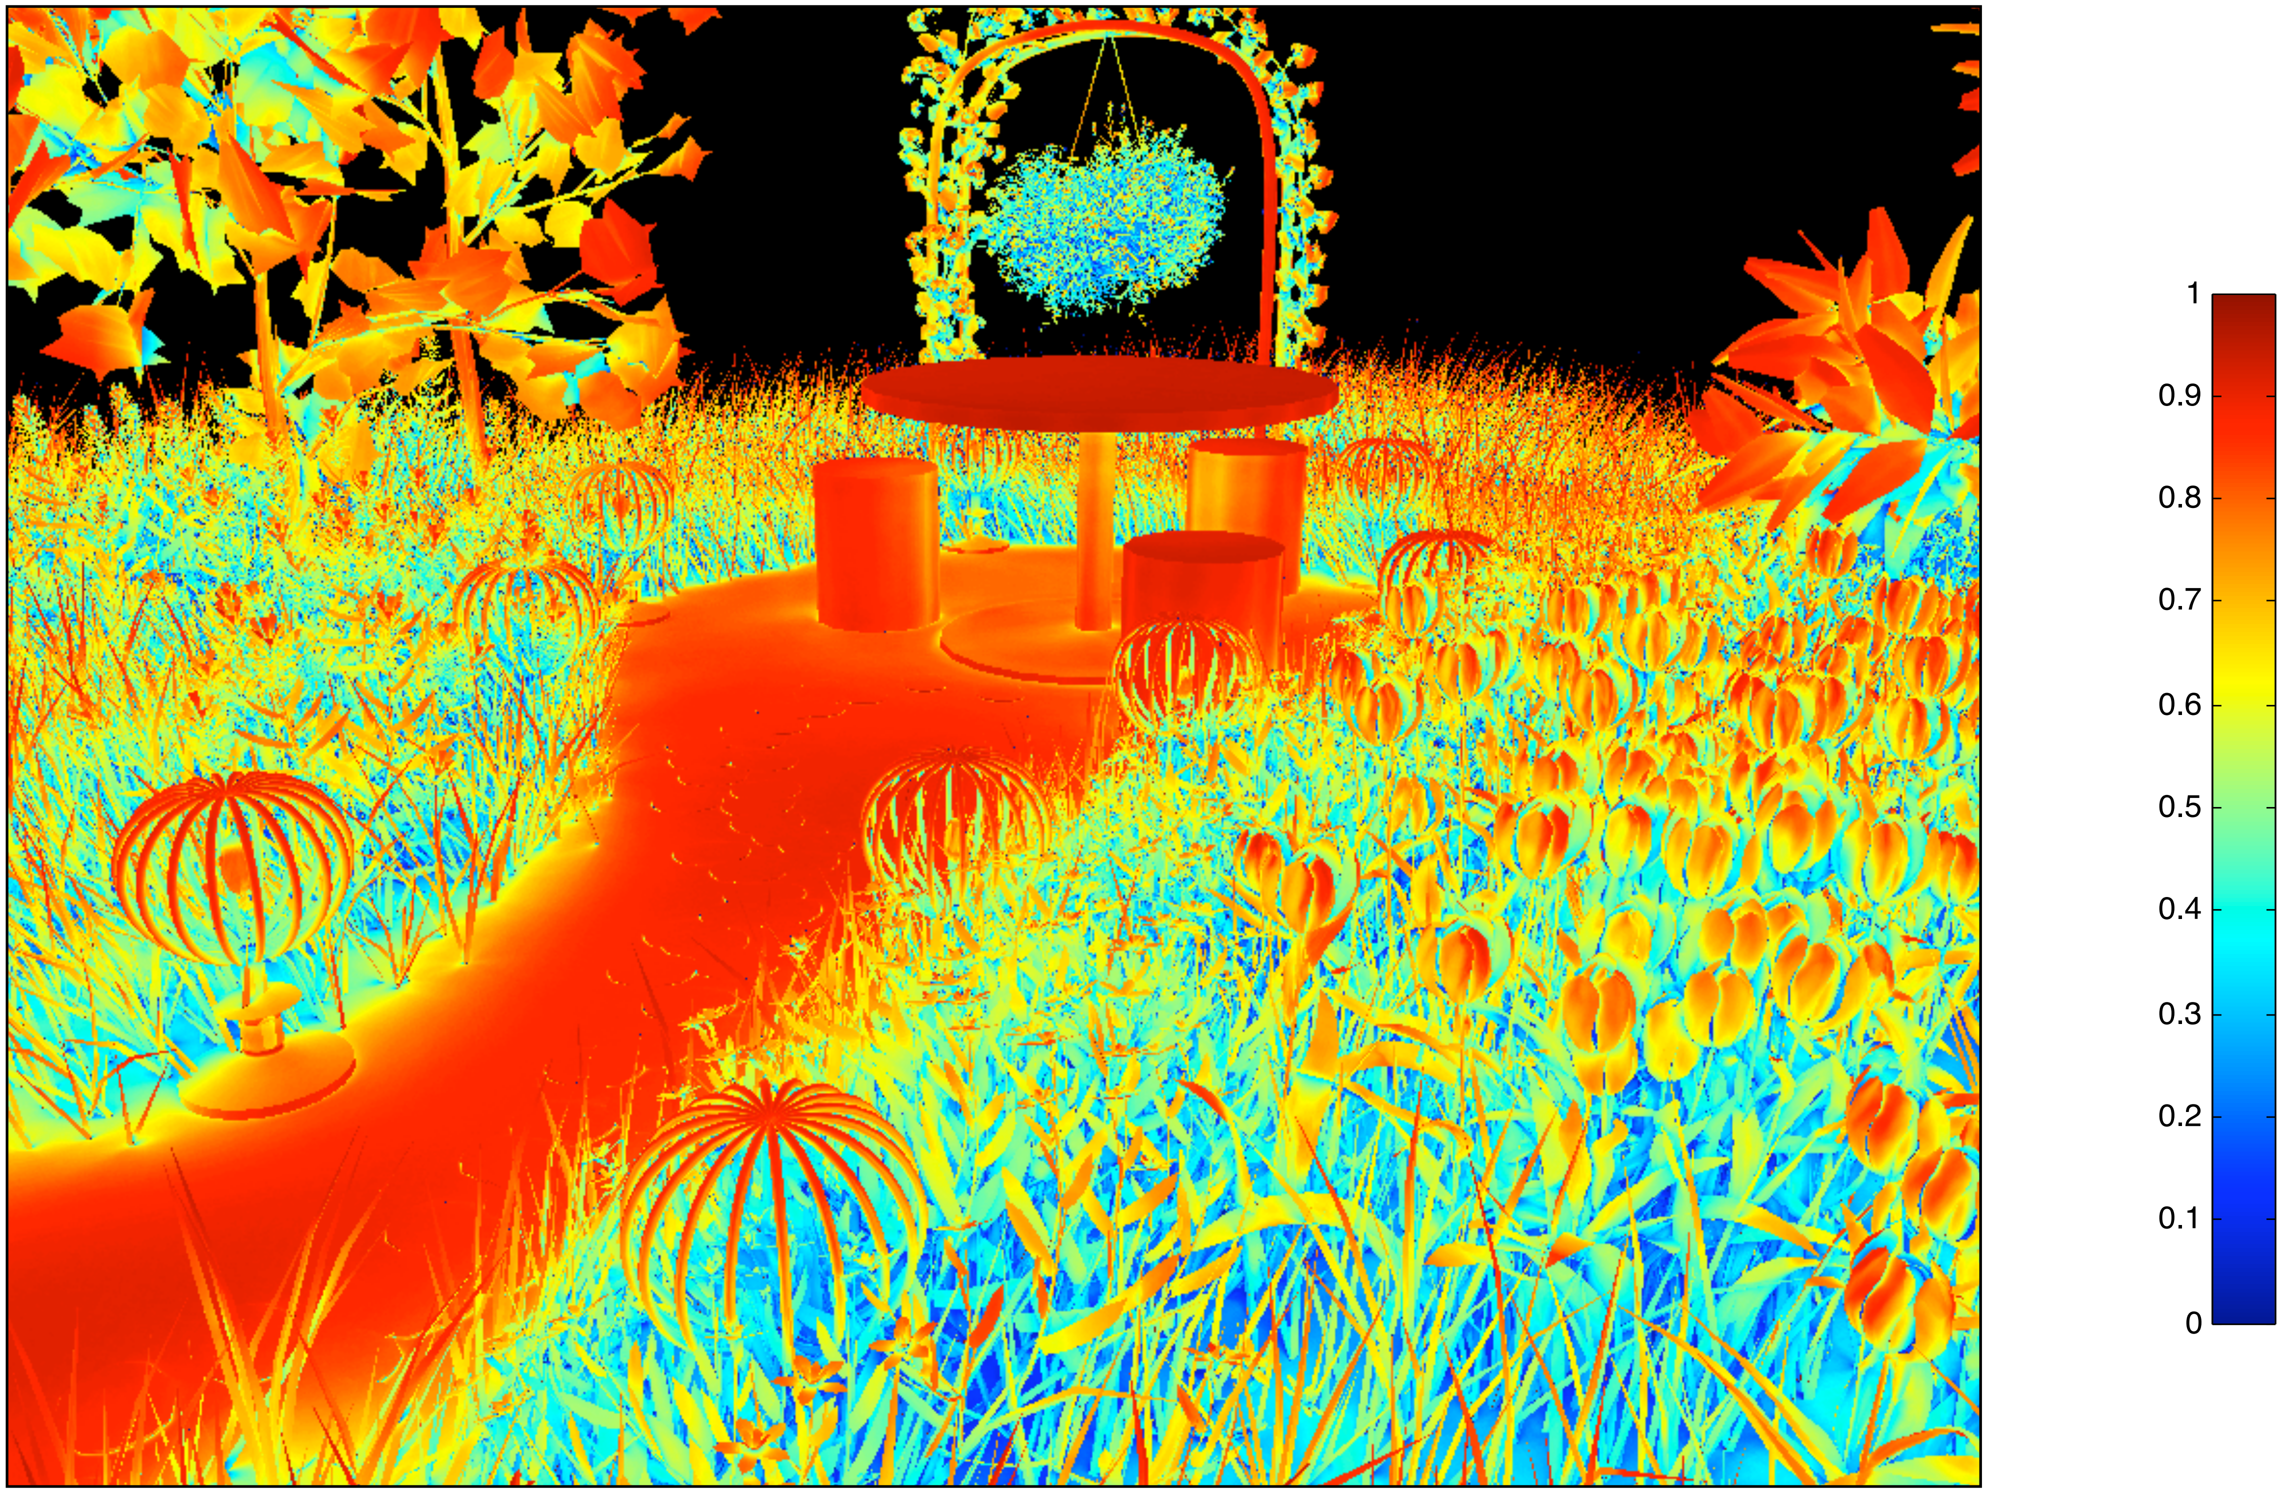

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Correlation

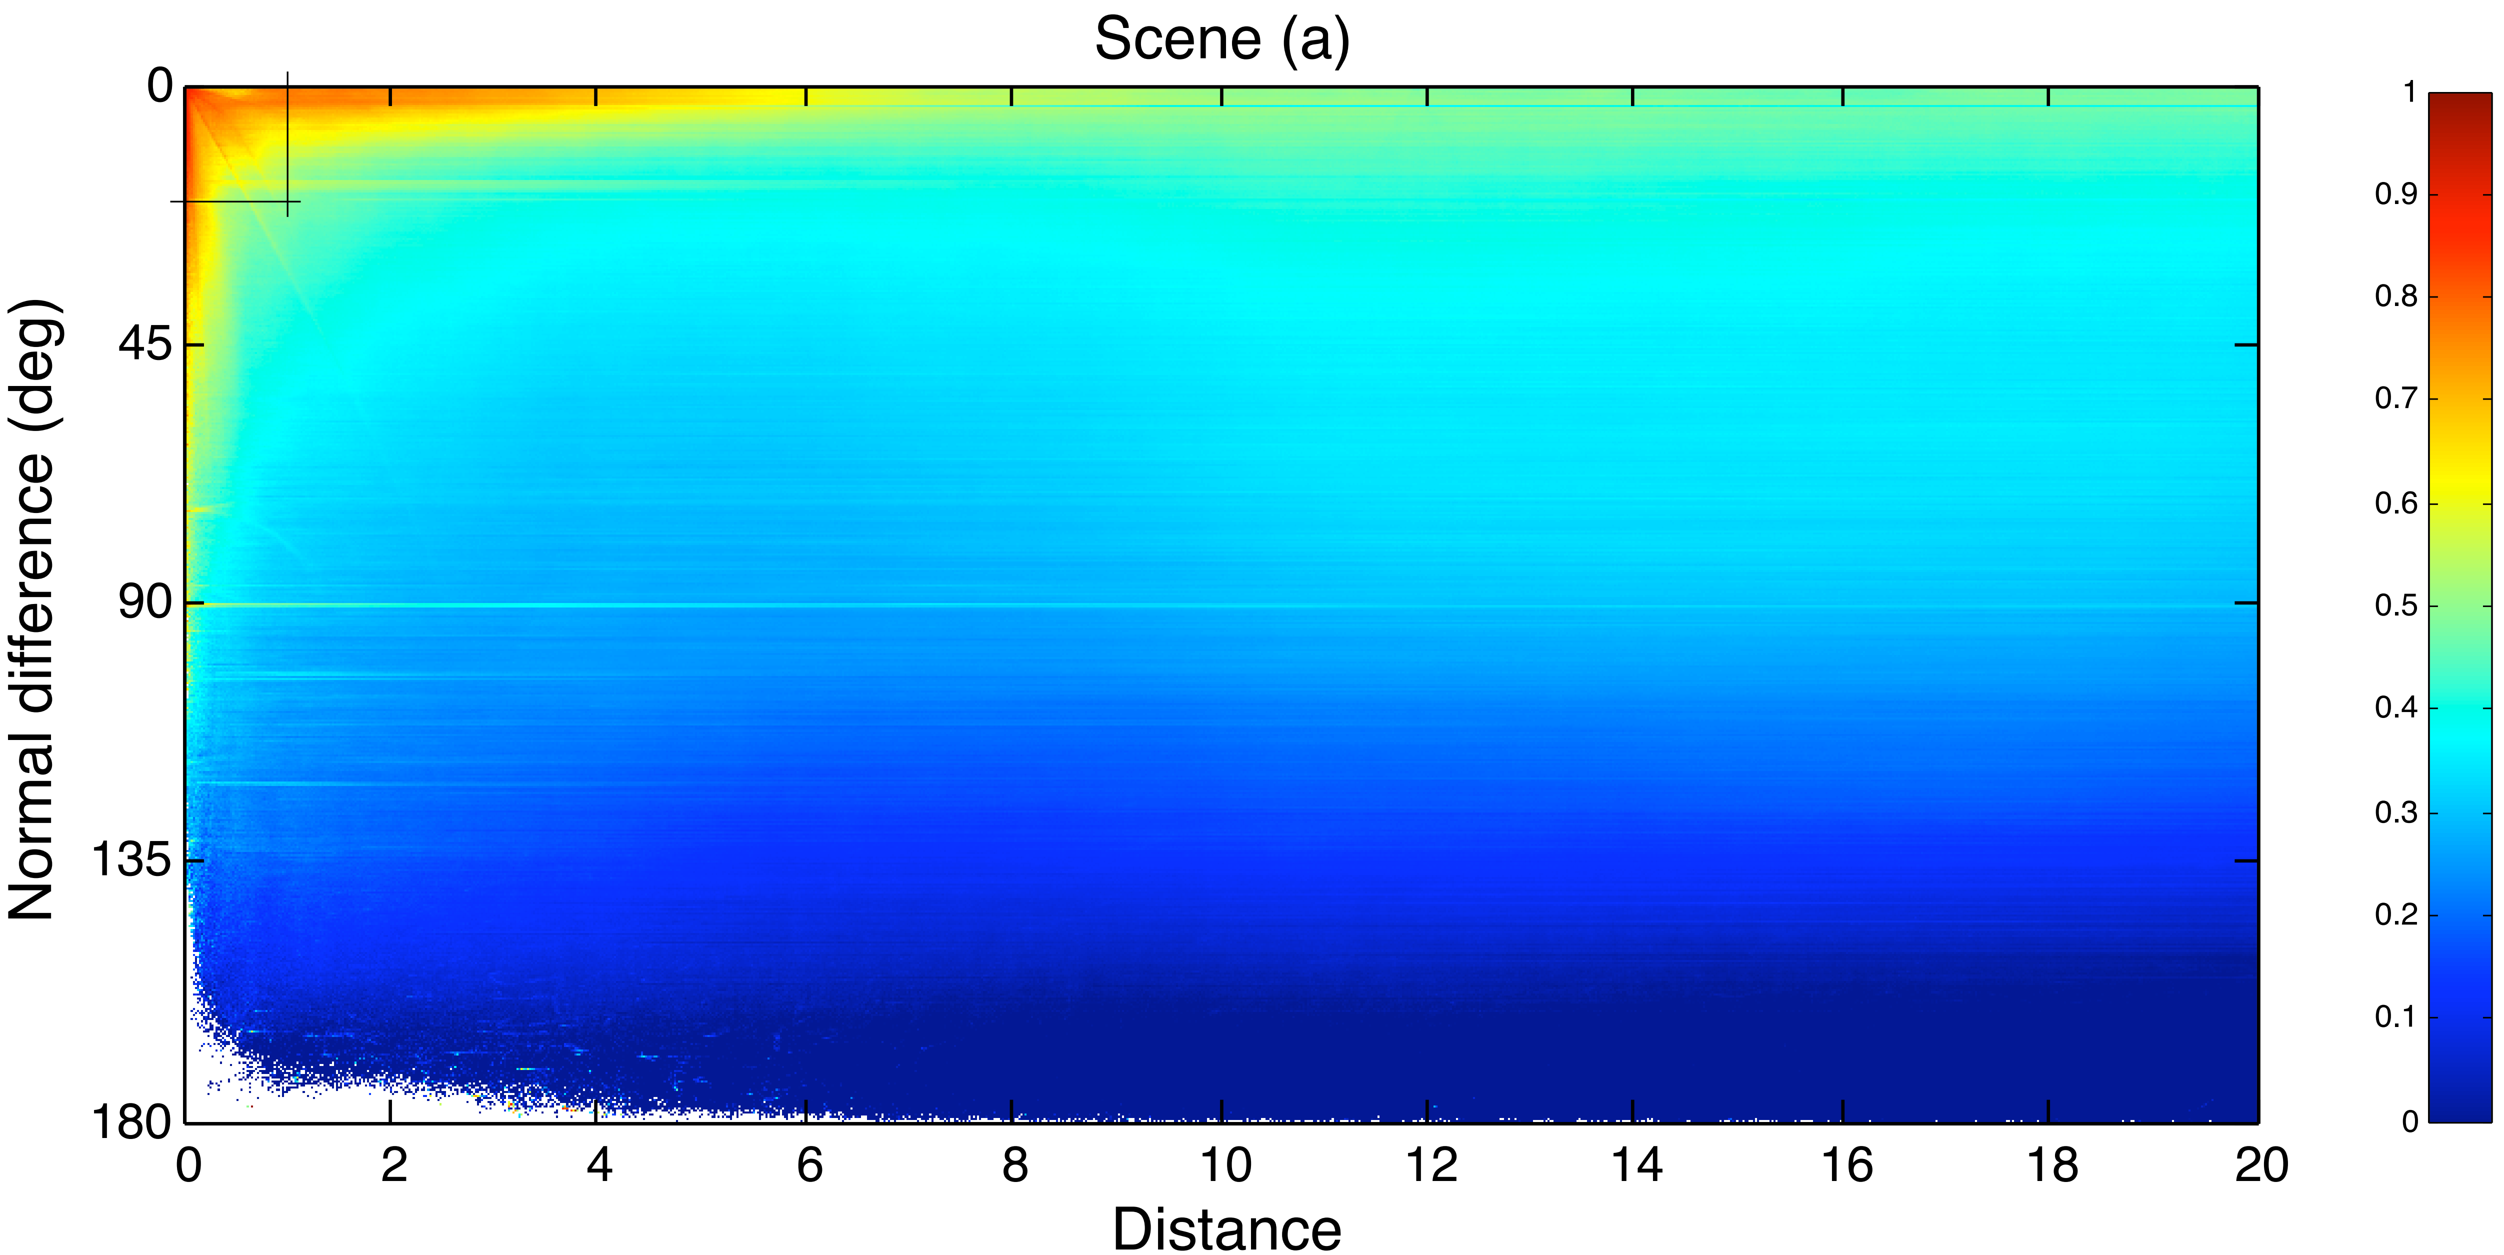

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Scene (c)

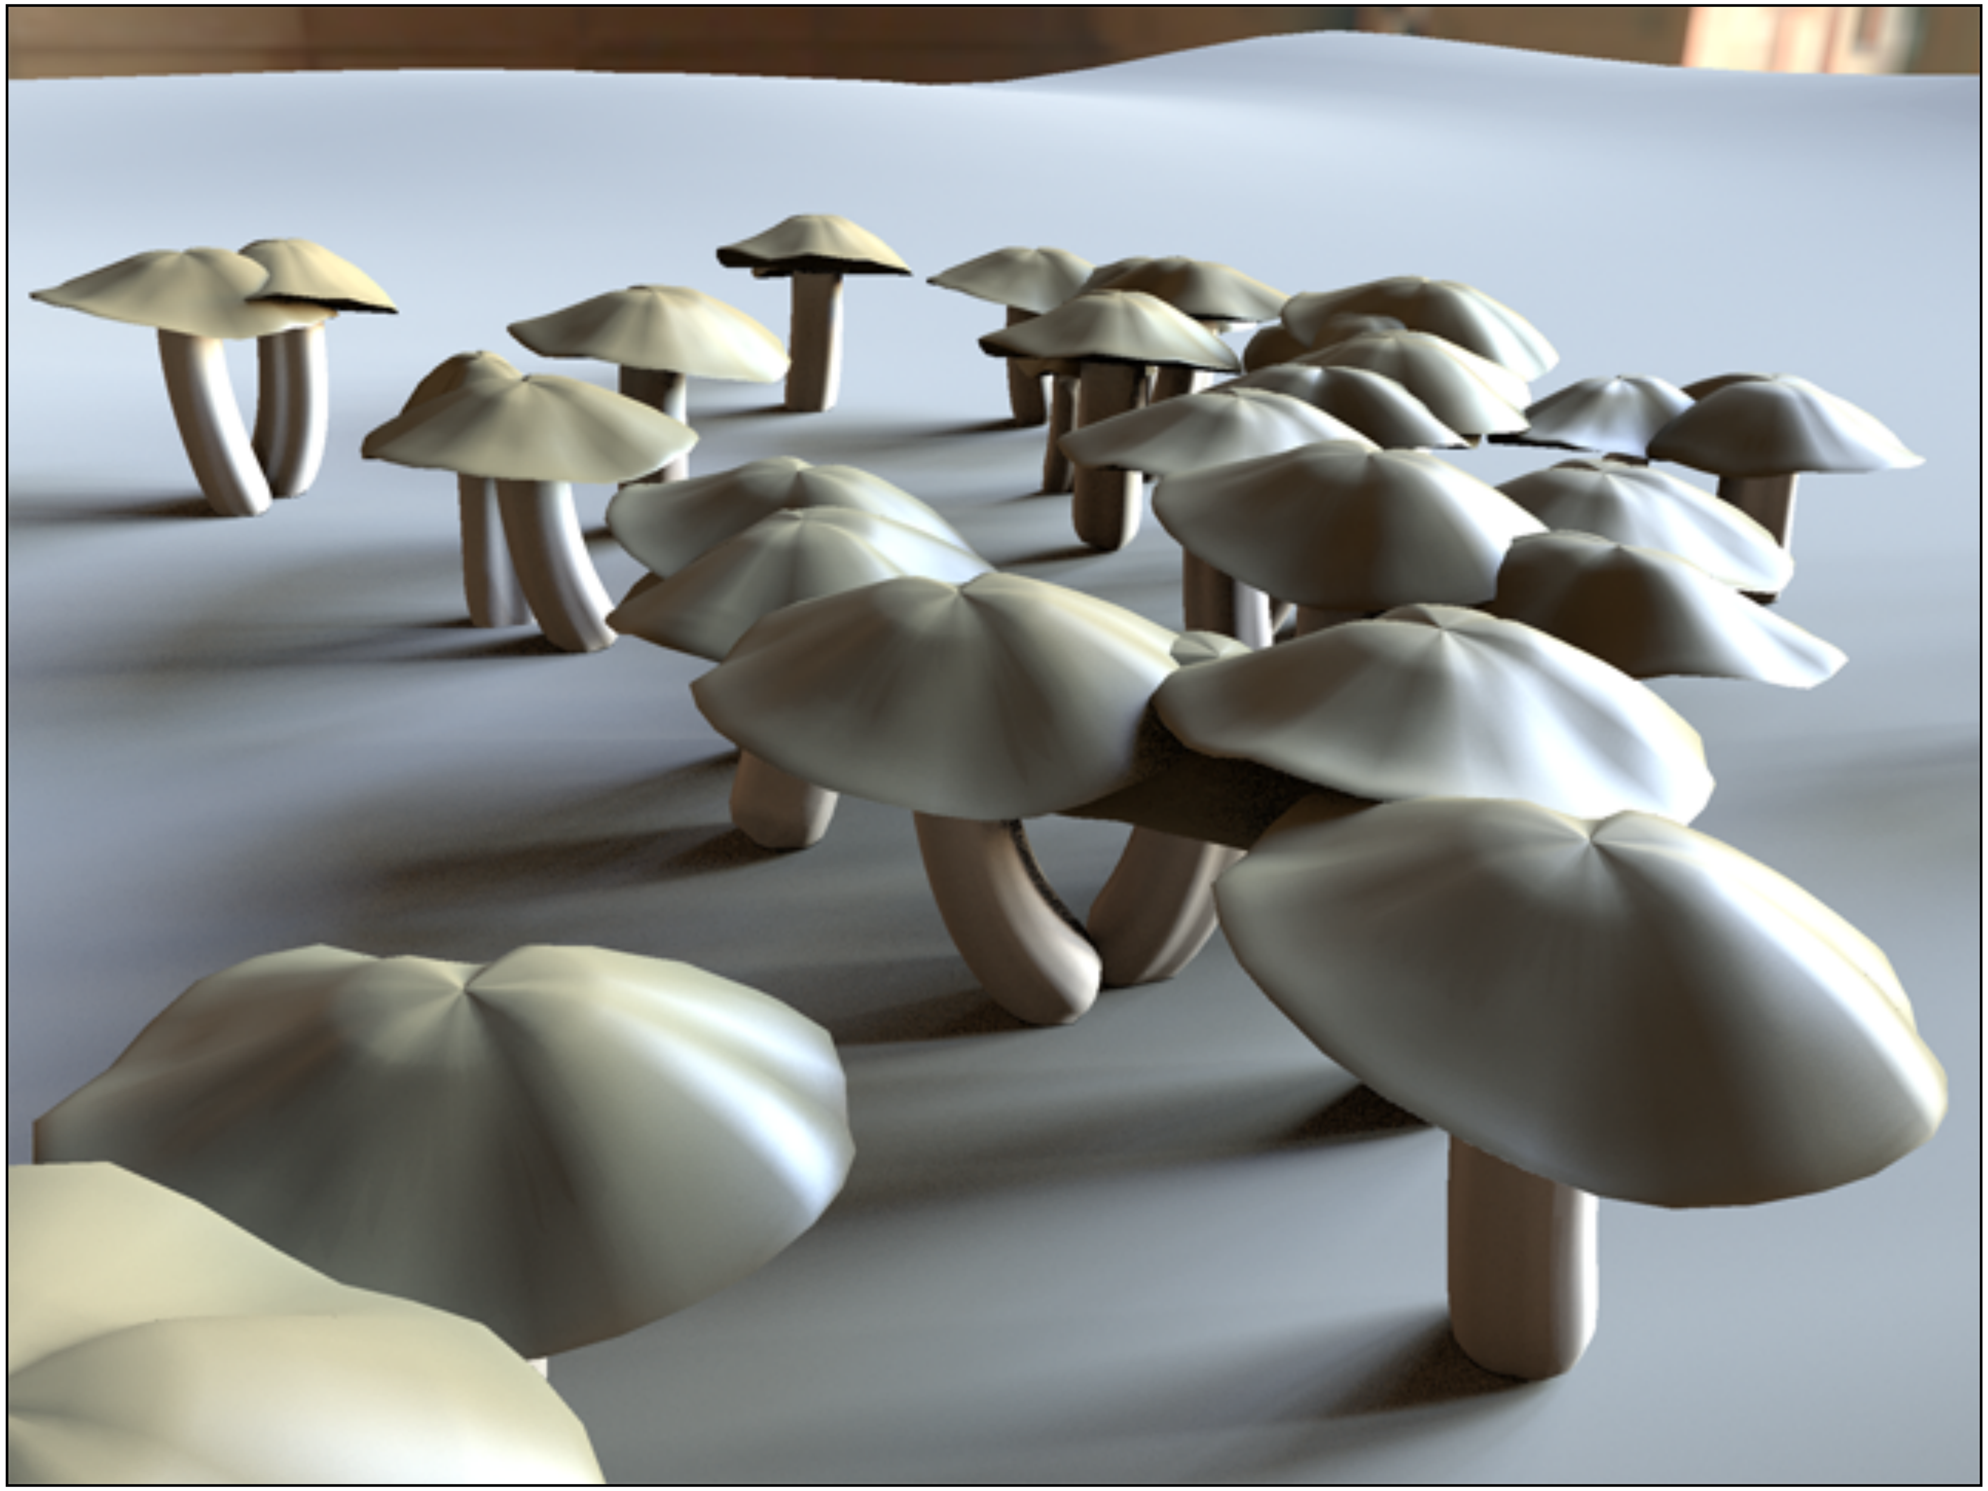

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Correlation

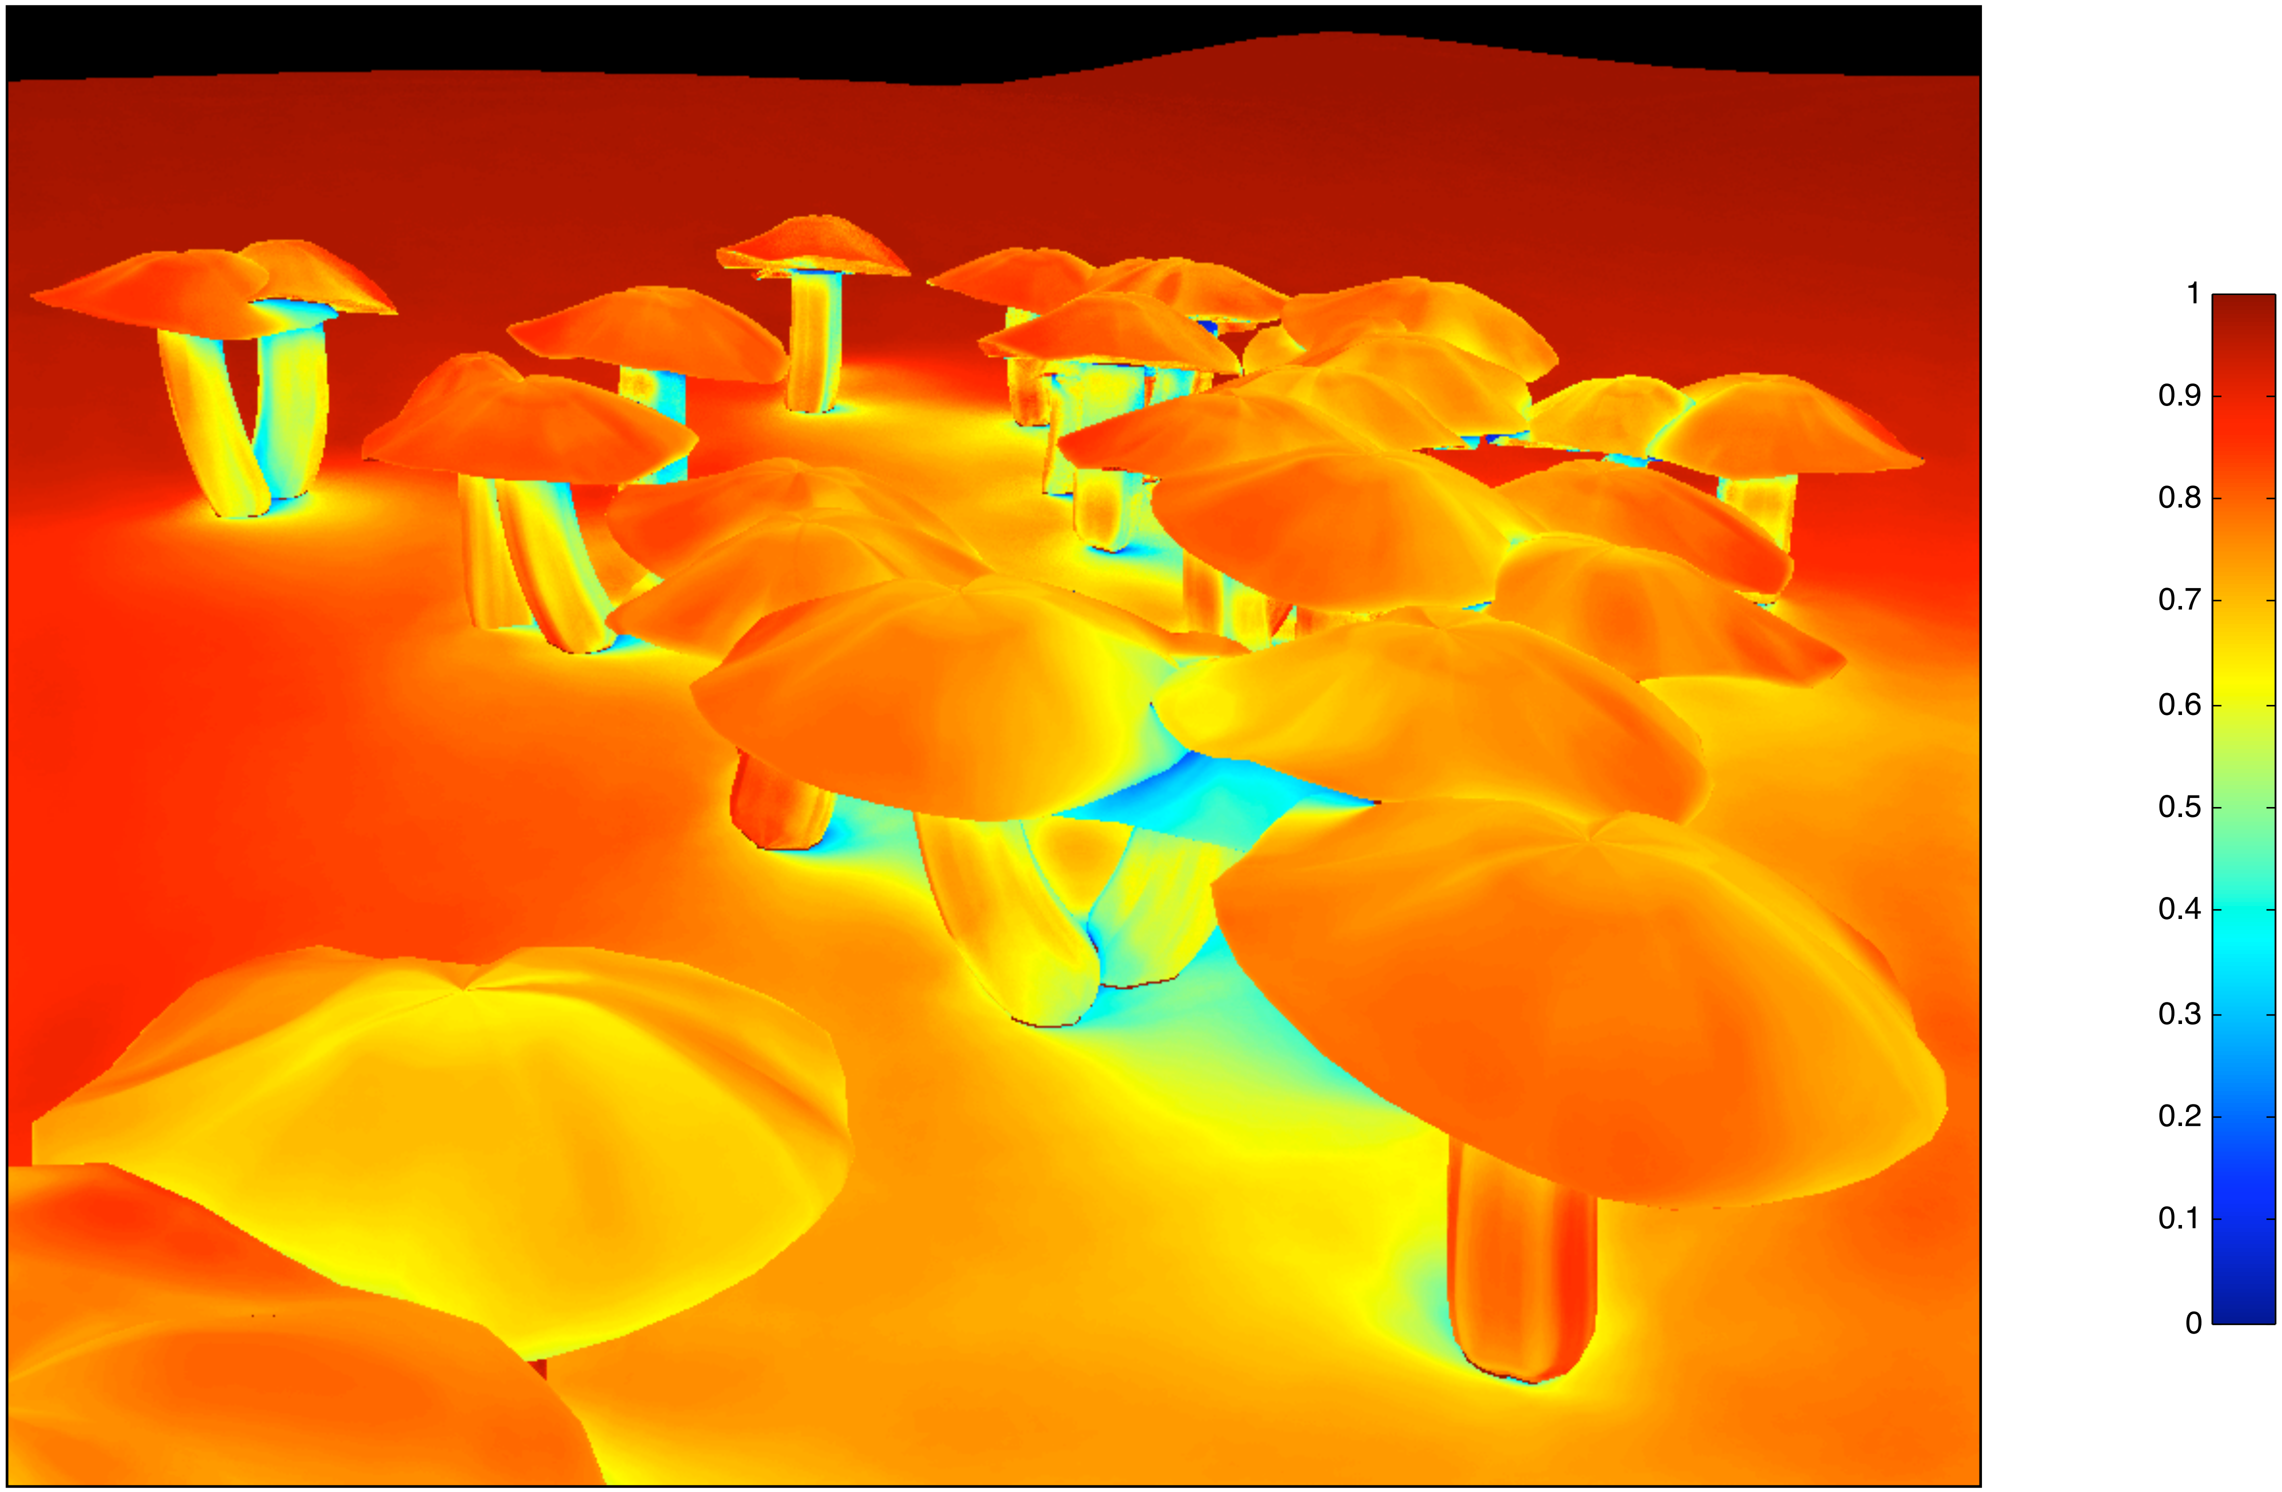

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Correlation

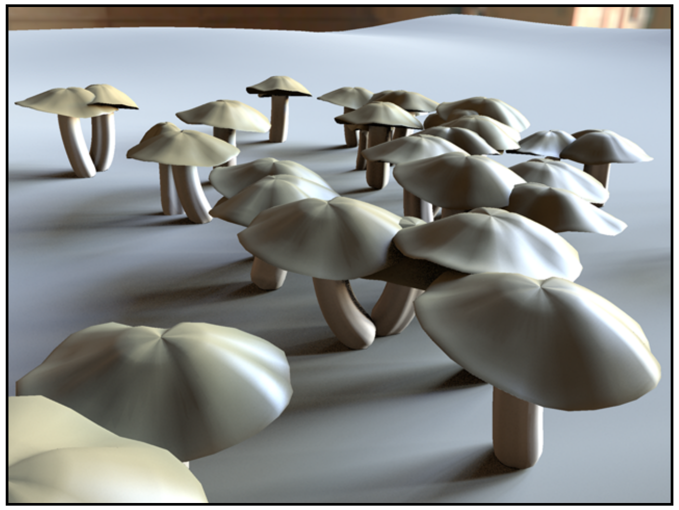

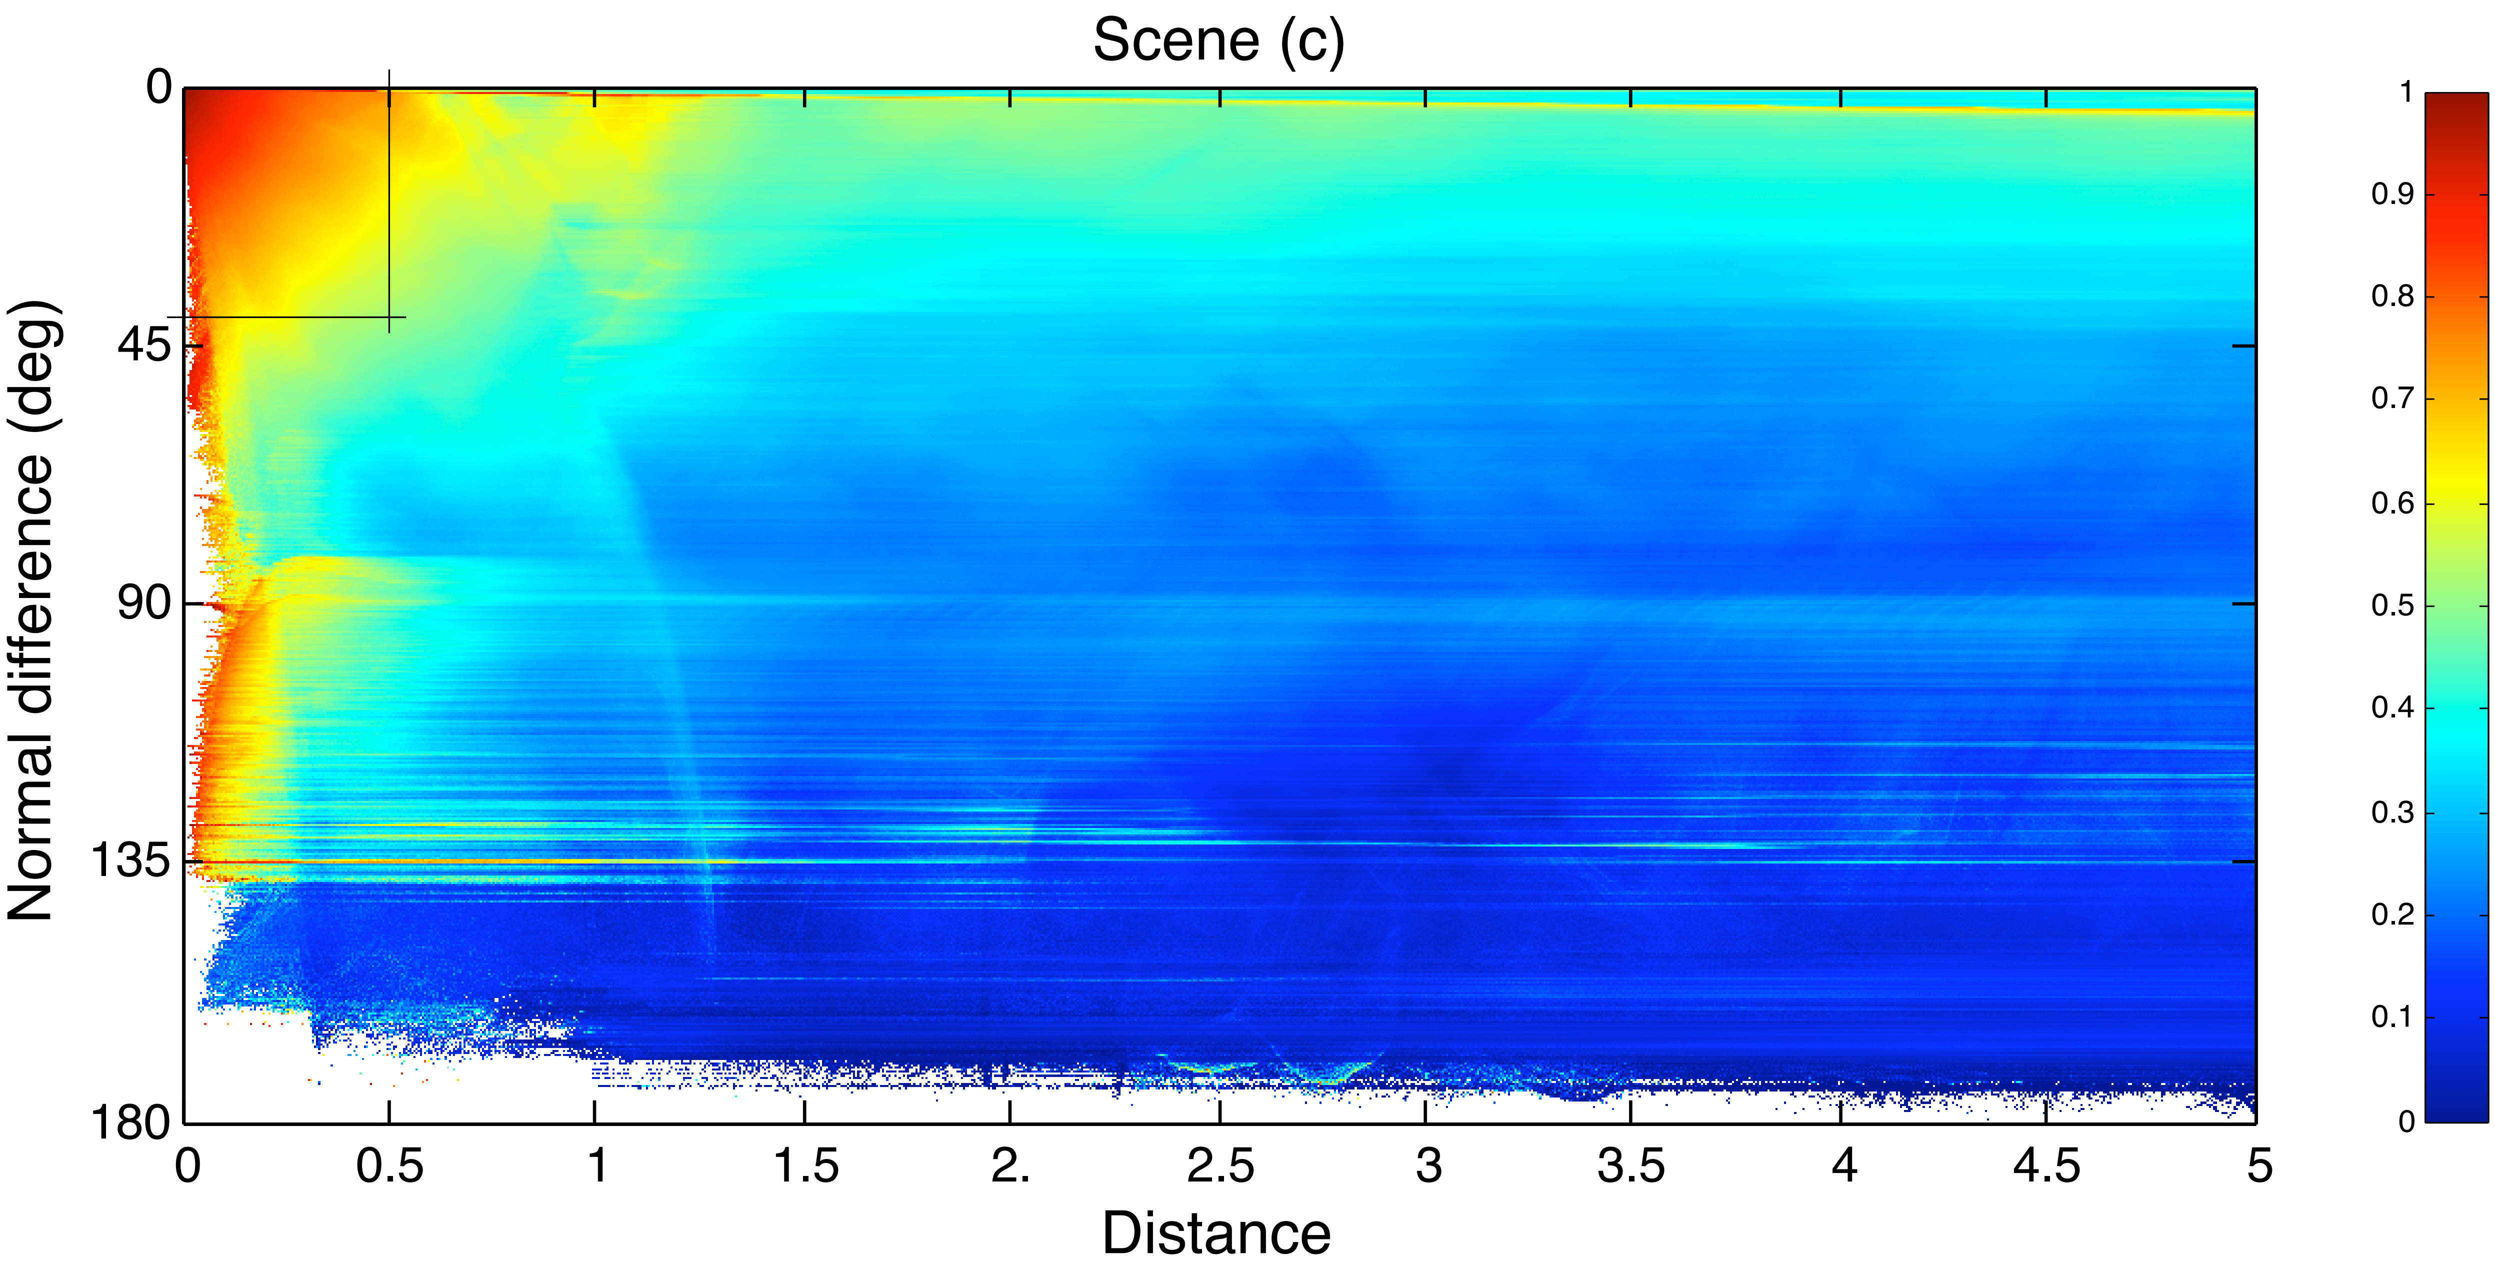

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Scene (d)

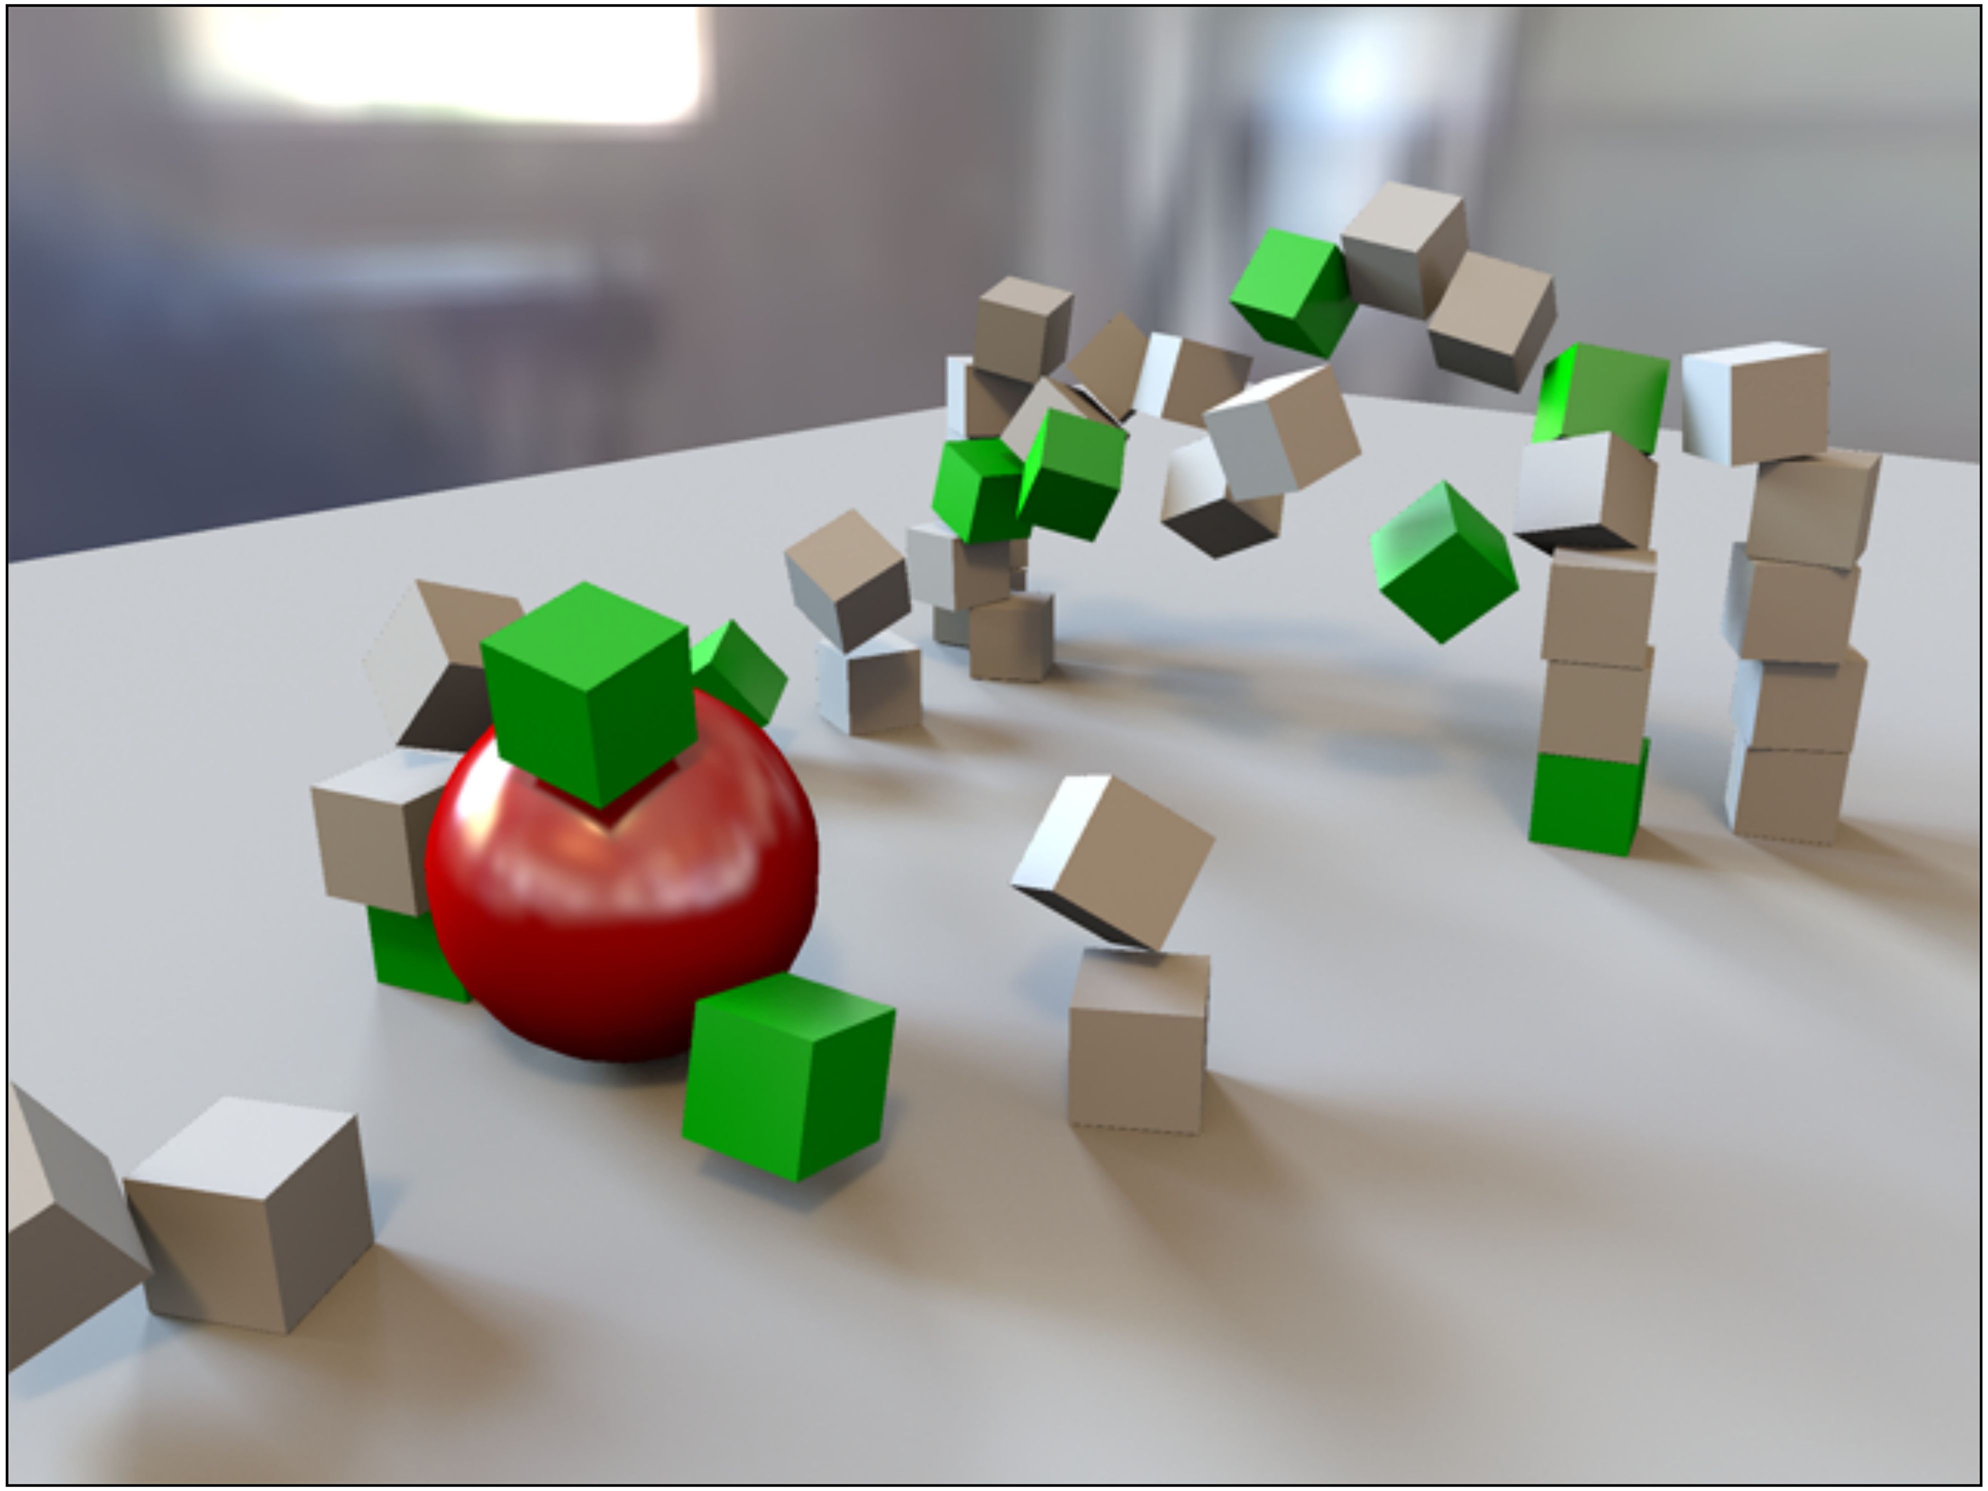

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Correlation

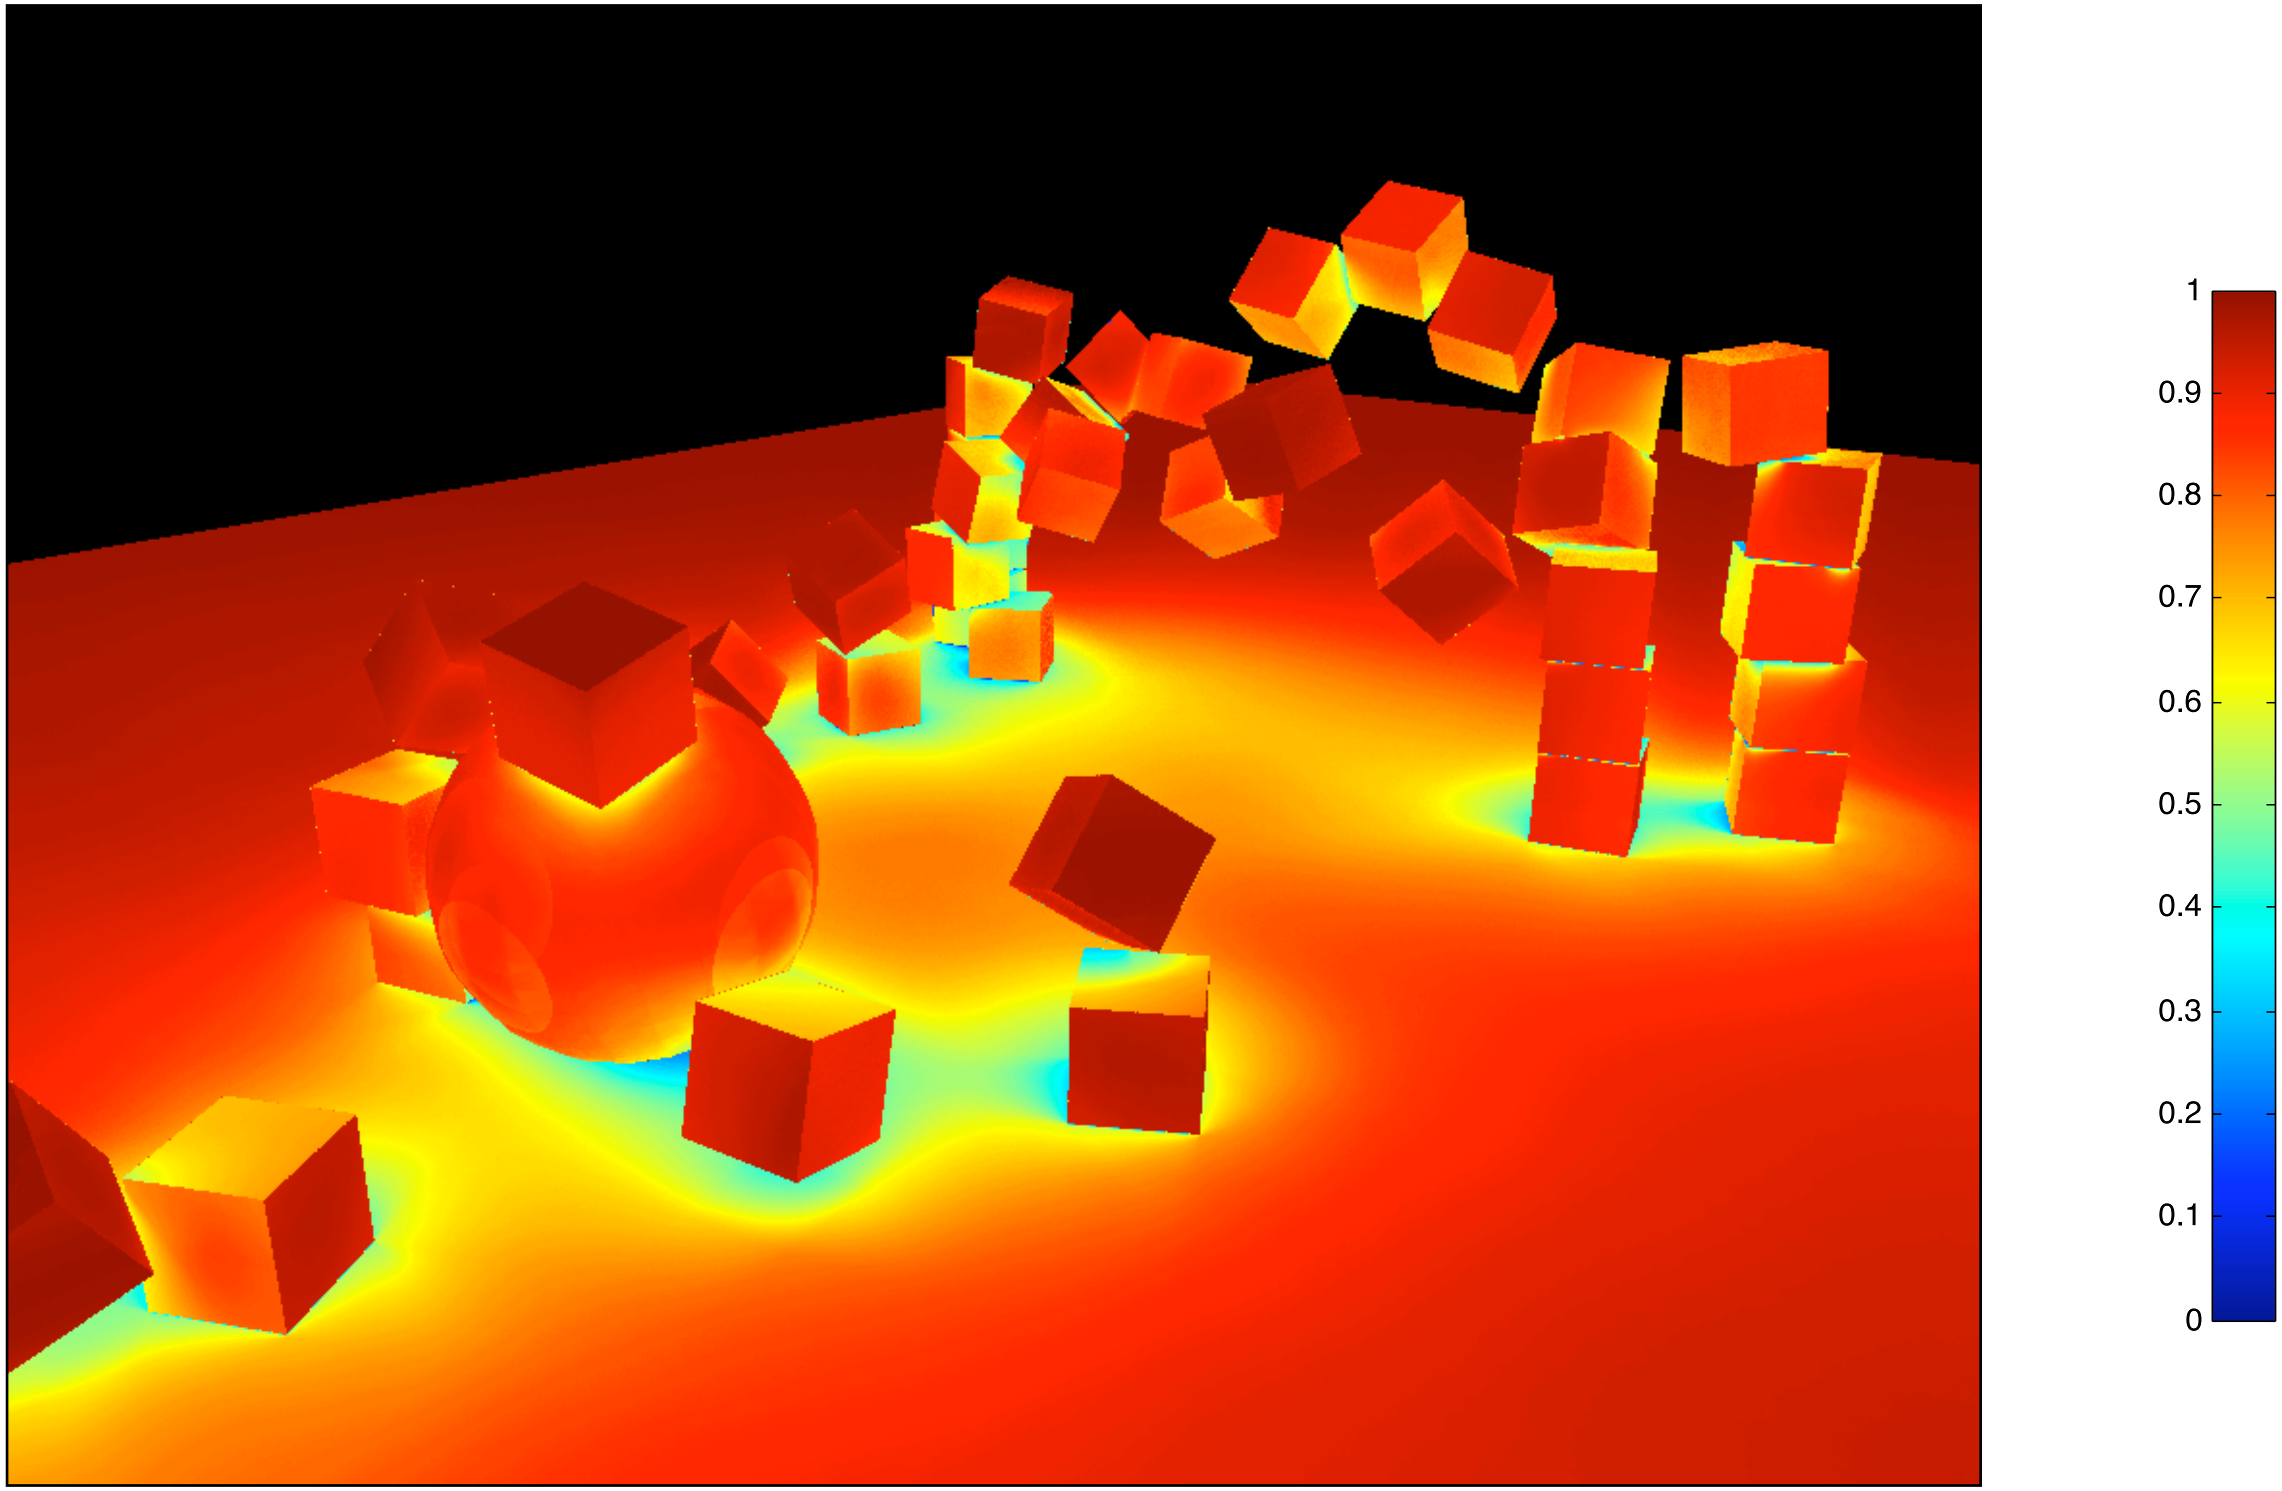

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Correlation

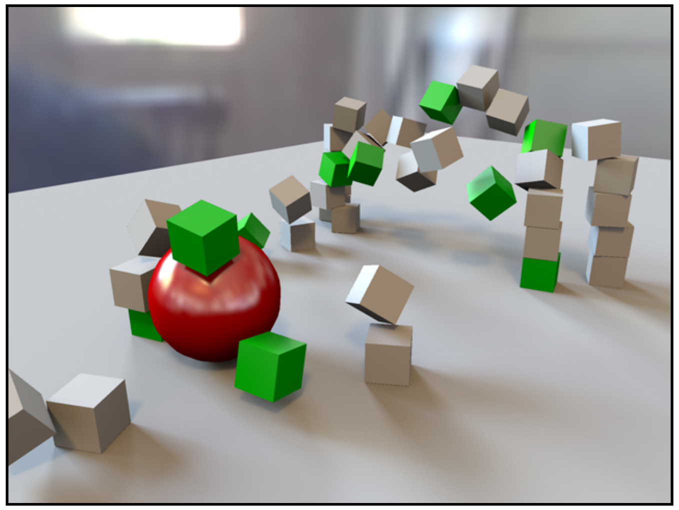

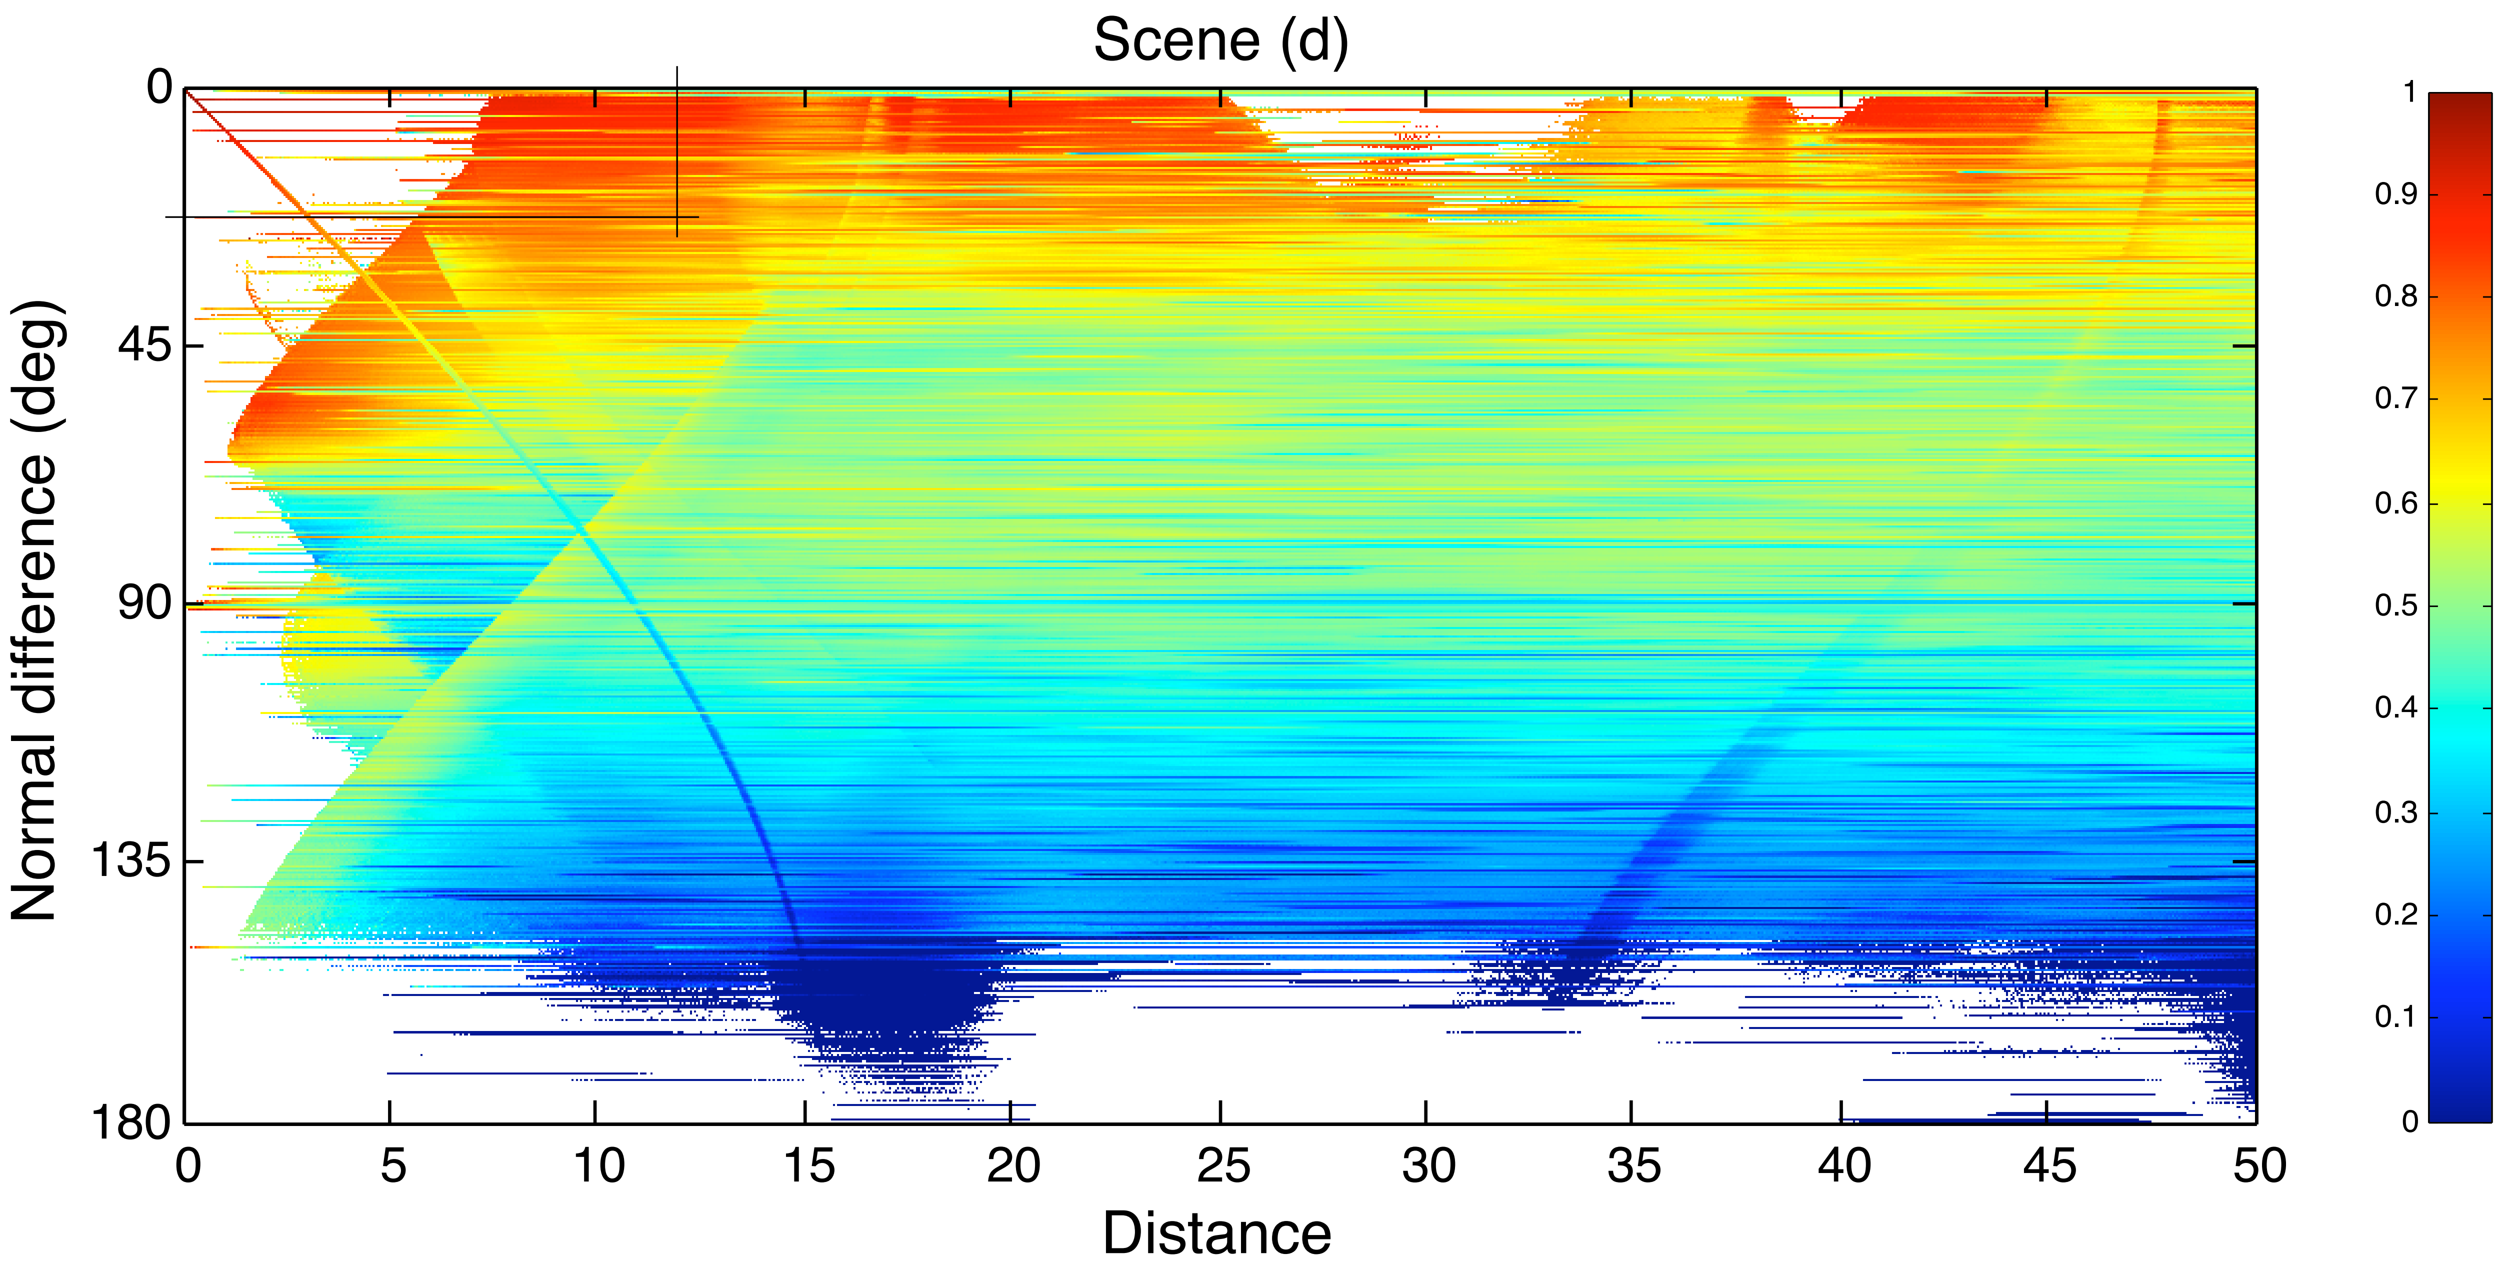

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

2. The Visibility Cache

During rendering, insert cache records:

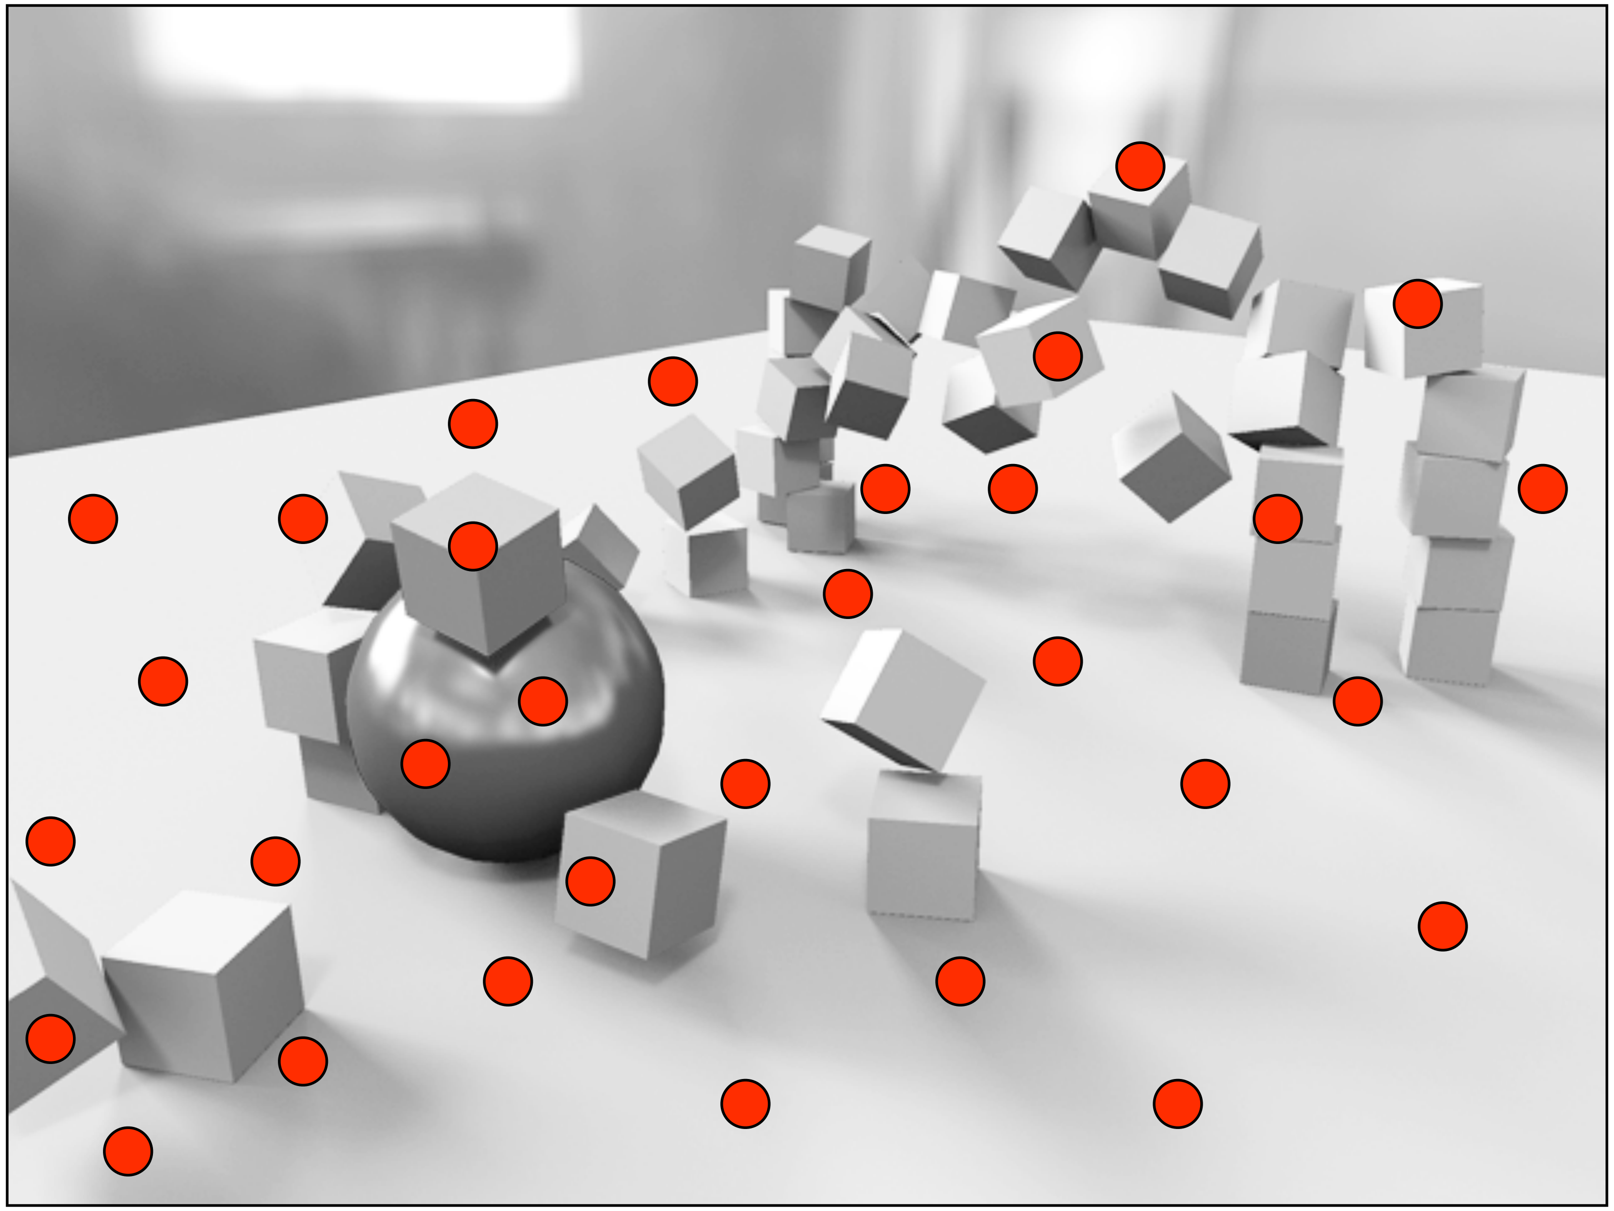

Each record is a low-res visibility map (e.g. 32x32)

Typically scene: 10k - 20k records

⇒ **5-10 M rays**

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Where to place records?

Several heuristics based on measurements

- Geometric
- Average visibility difference
- Max number of wrong visibility queries

Details in the paper

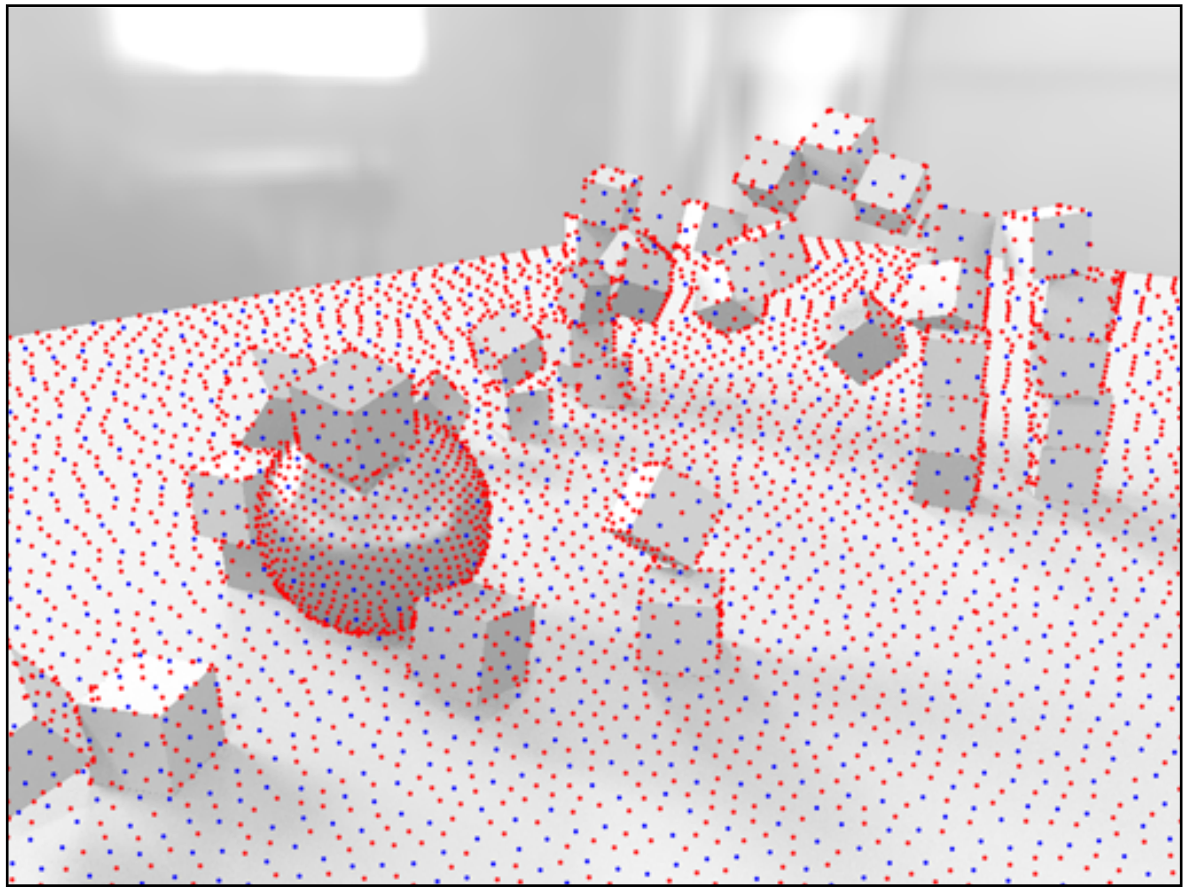

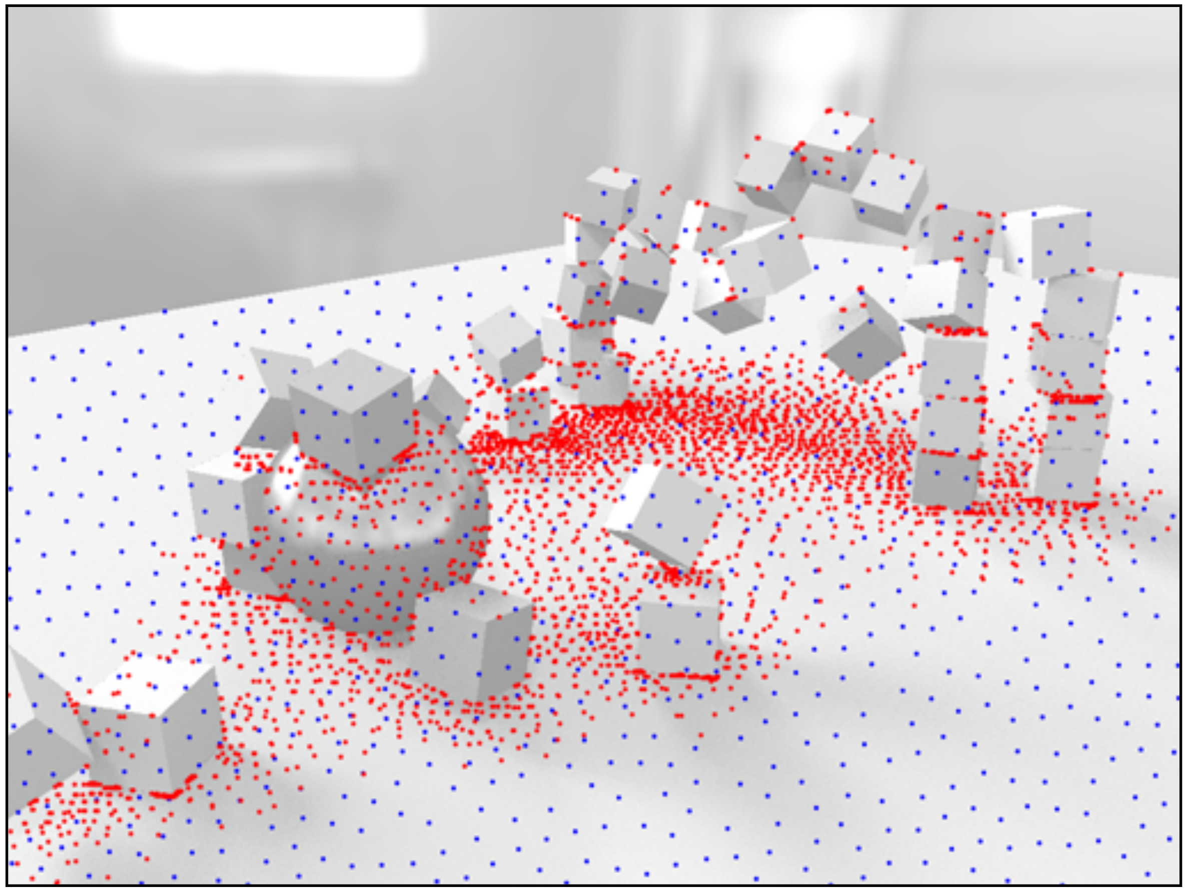

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Example: Static Scene

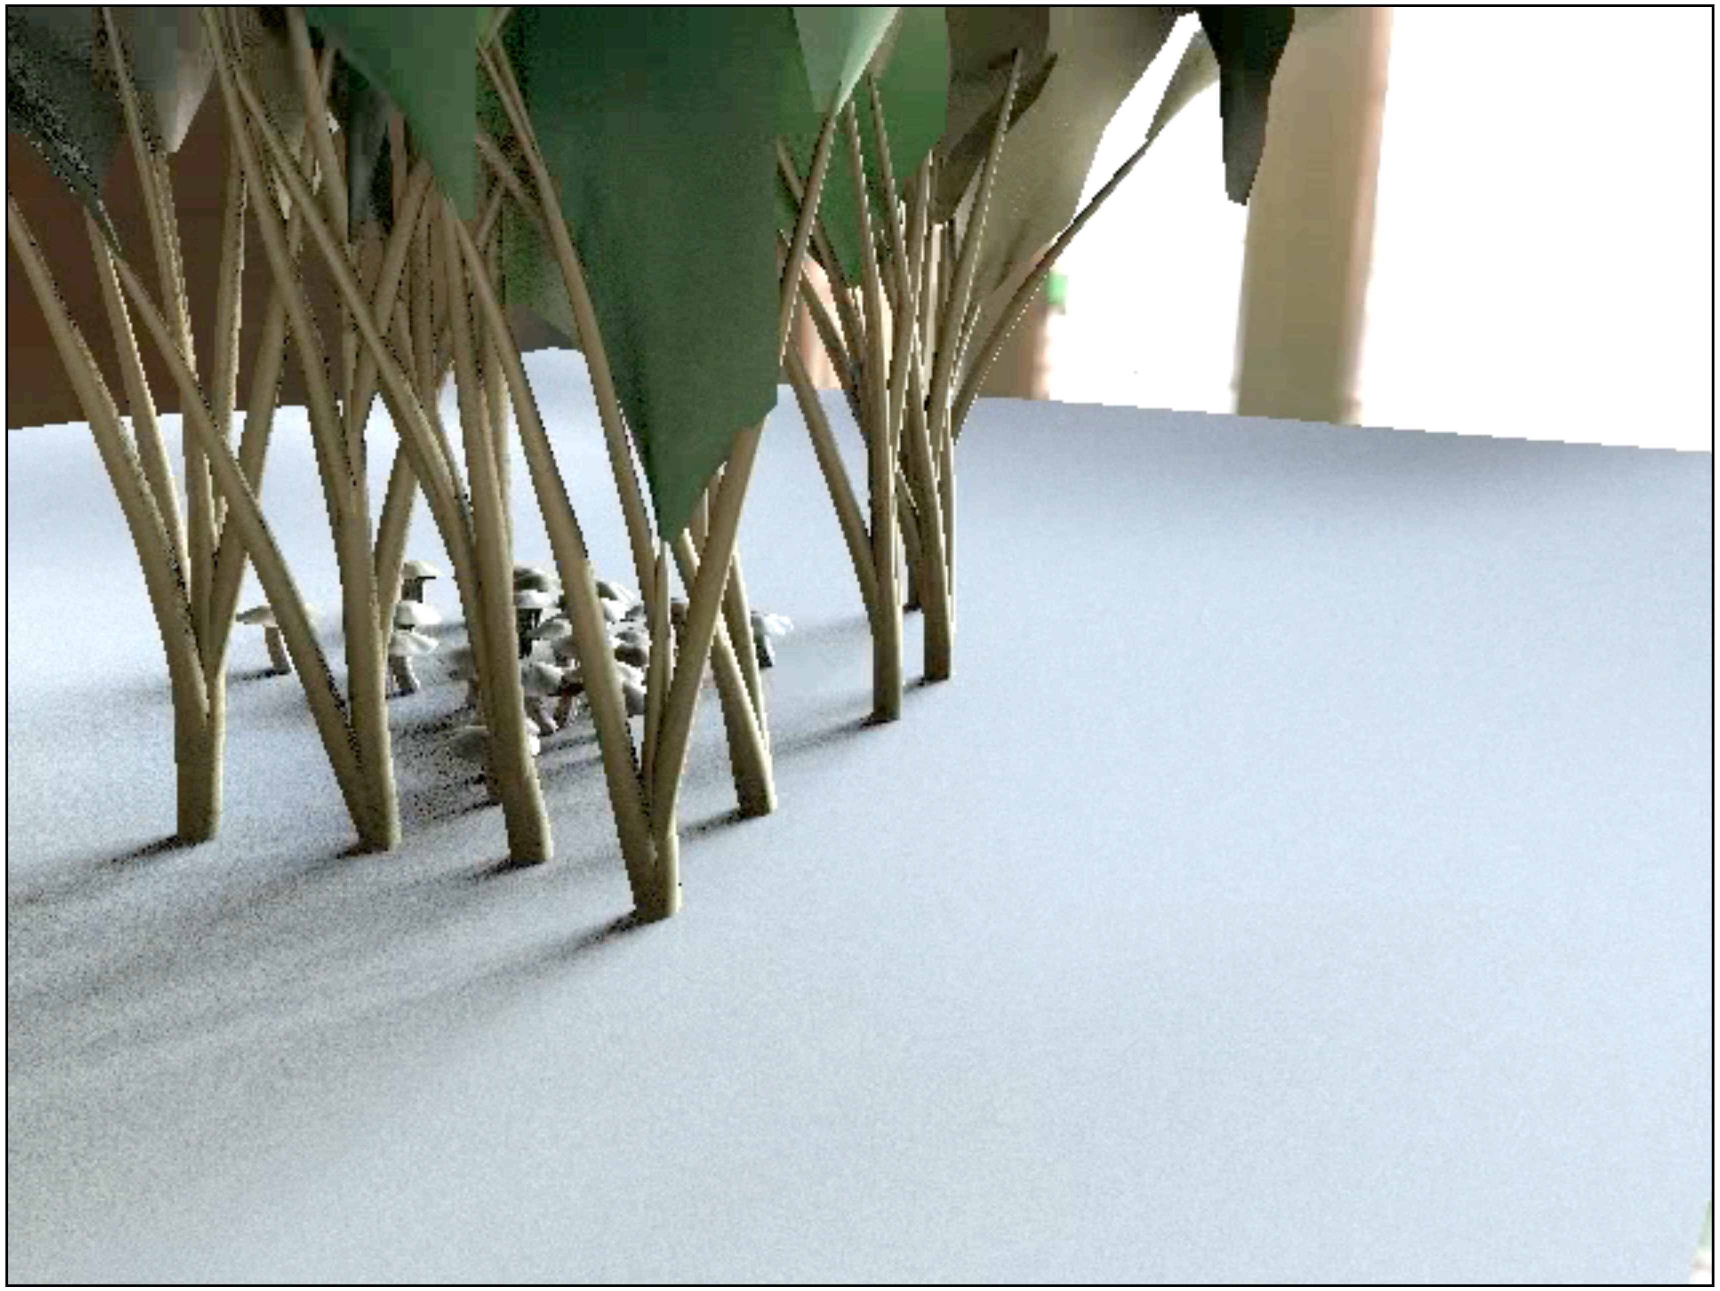

Video

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Example: Static Scene

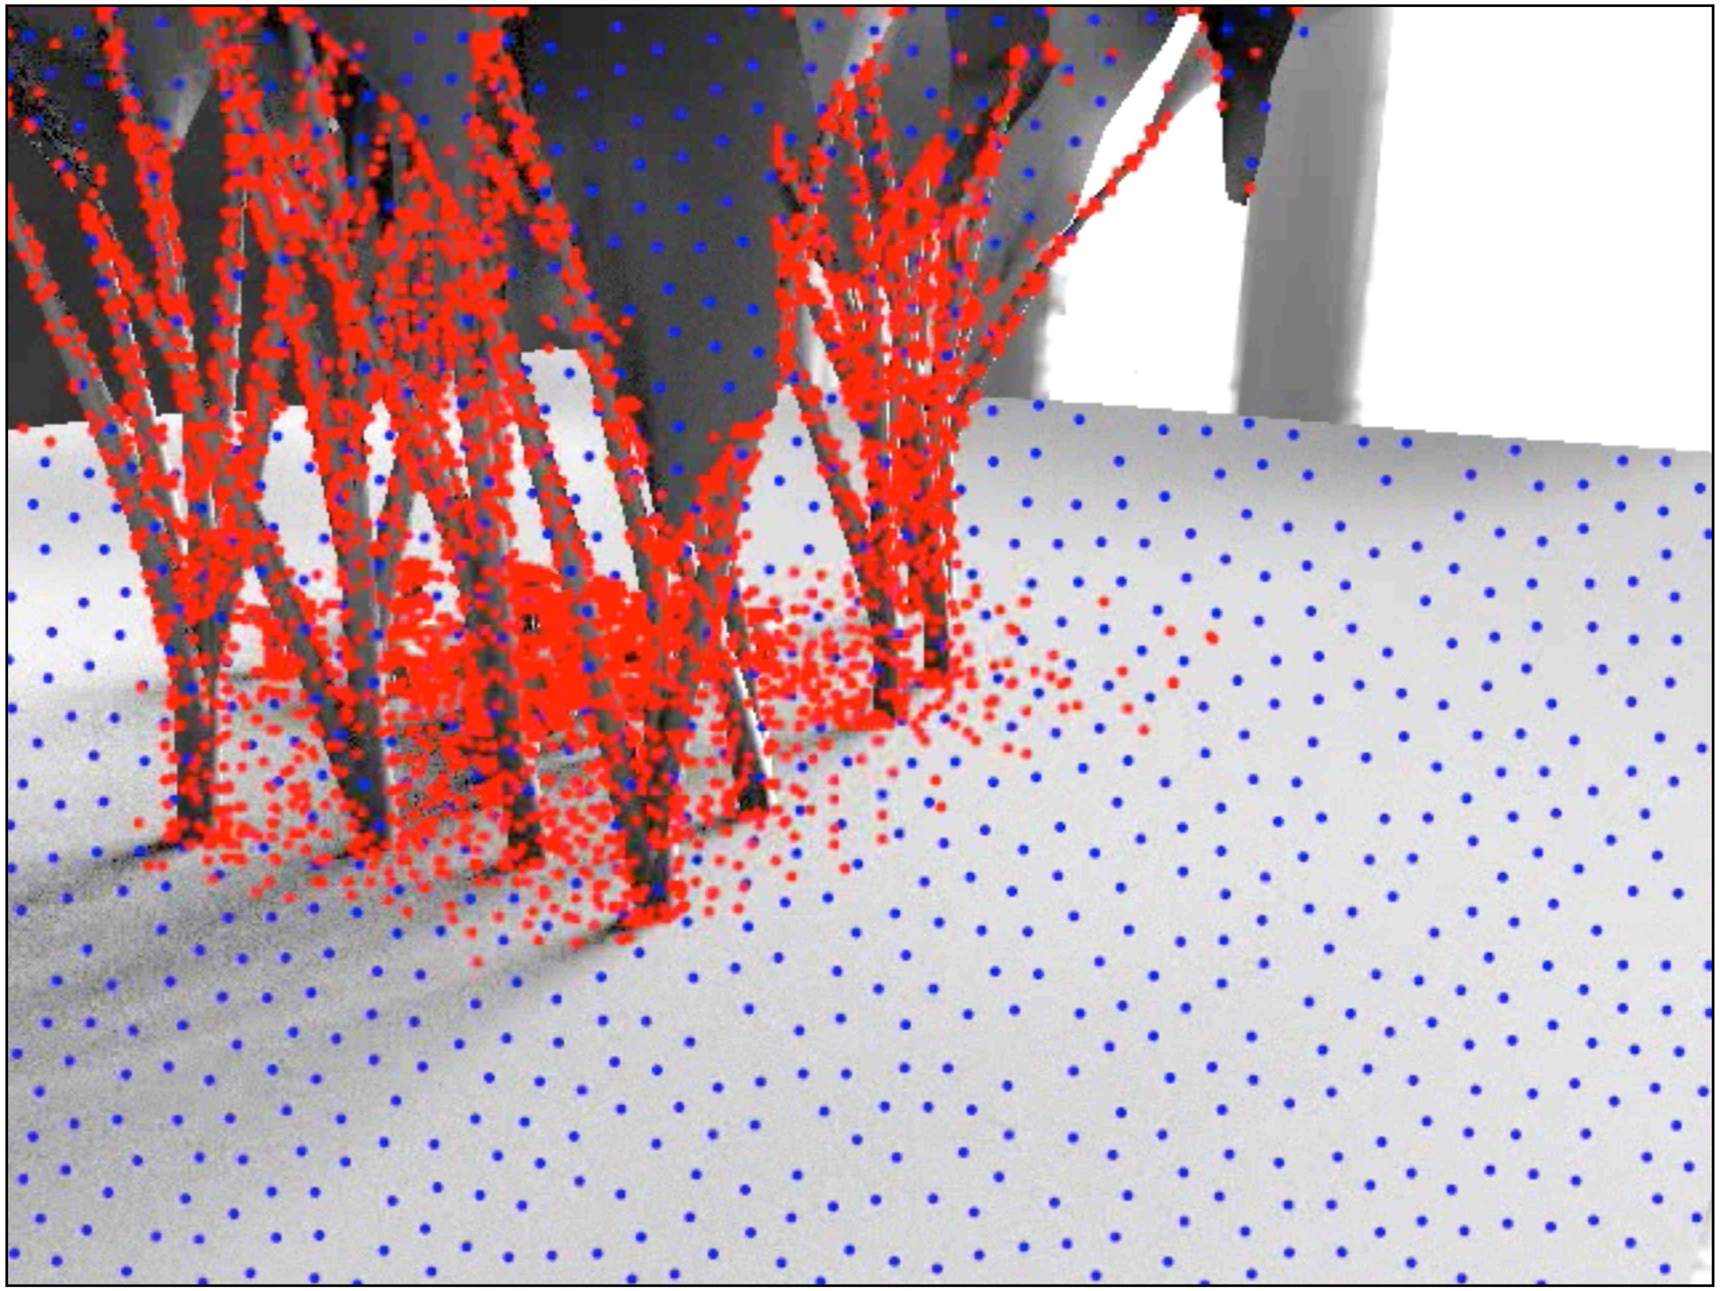

Video

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Example: Dynamic Scene

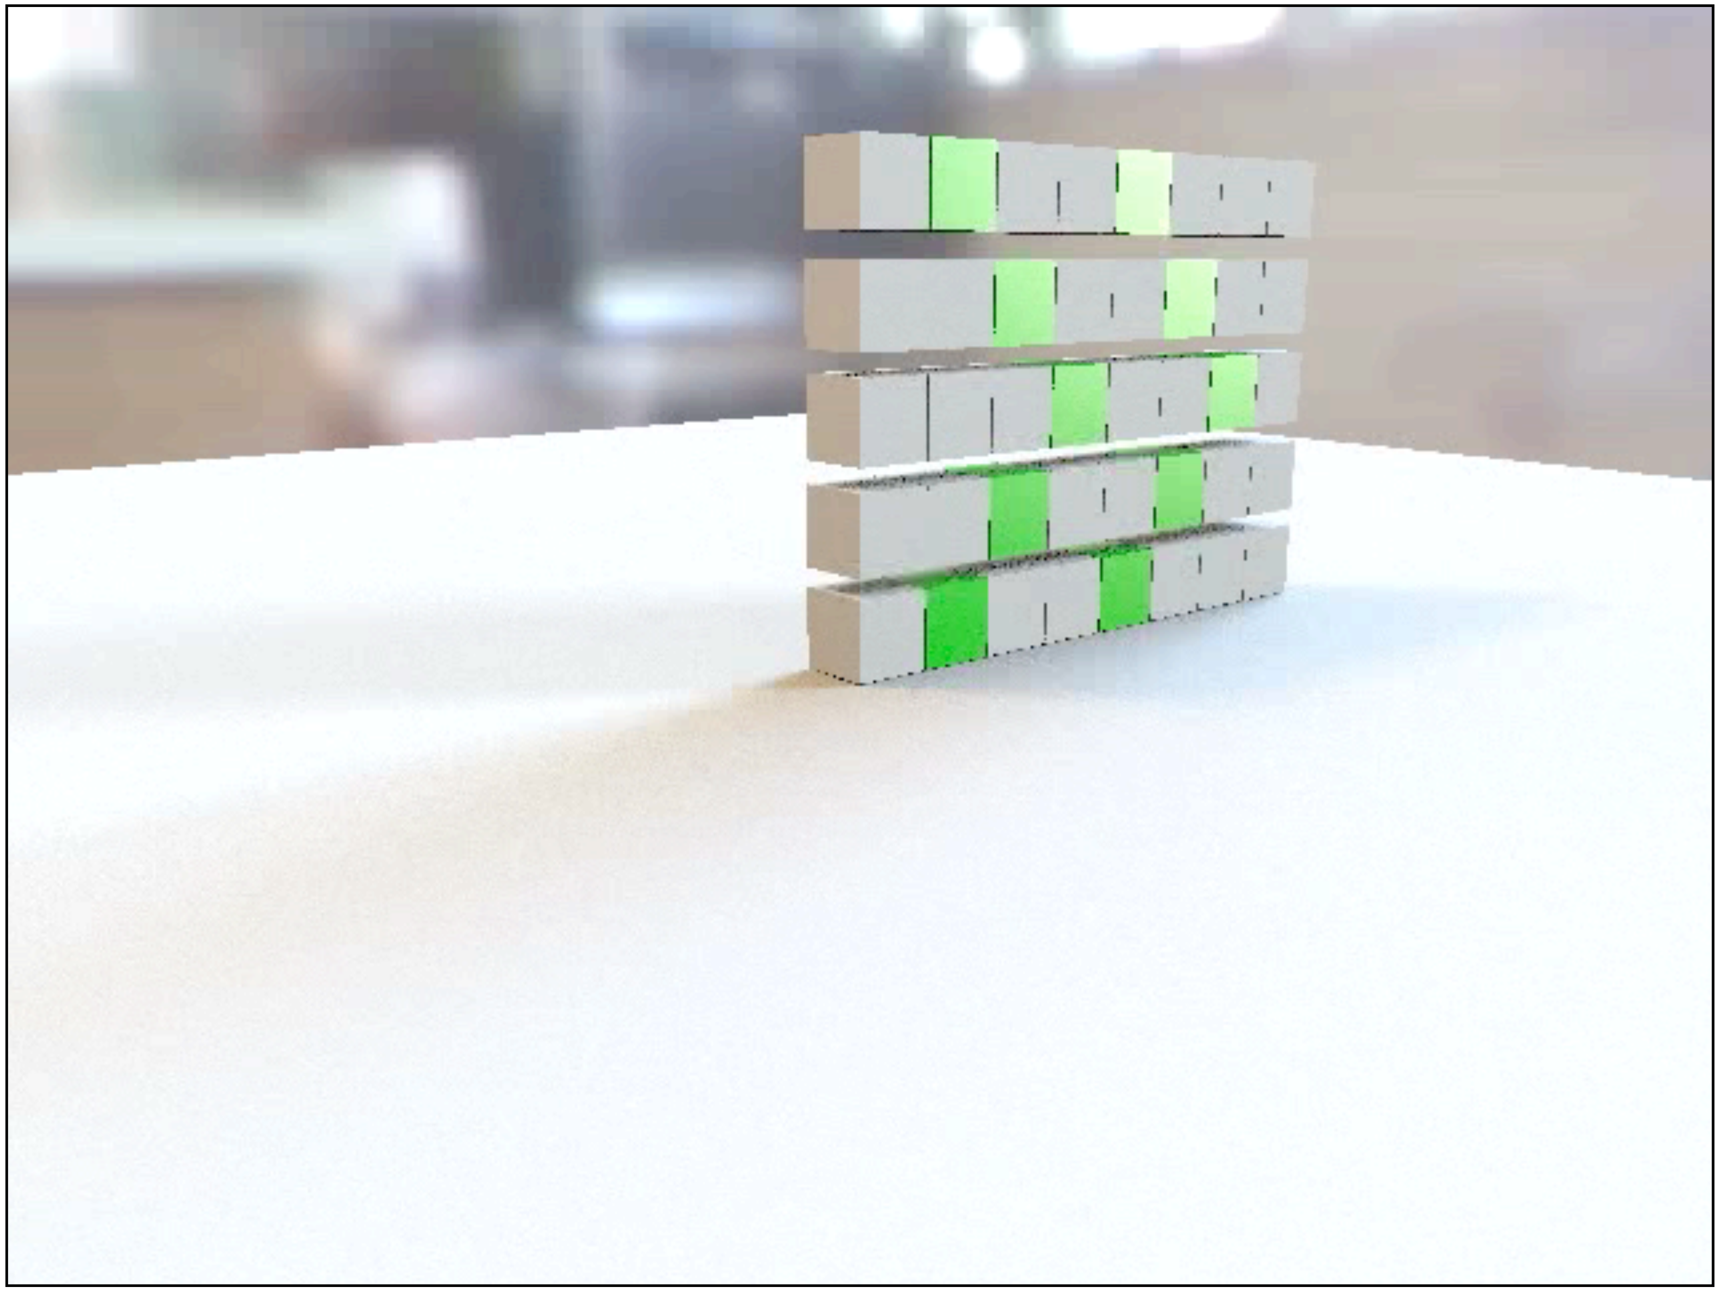

Video

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Example: Dynamic Scene

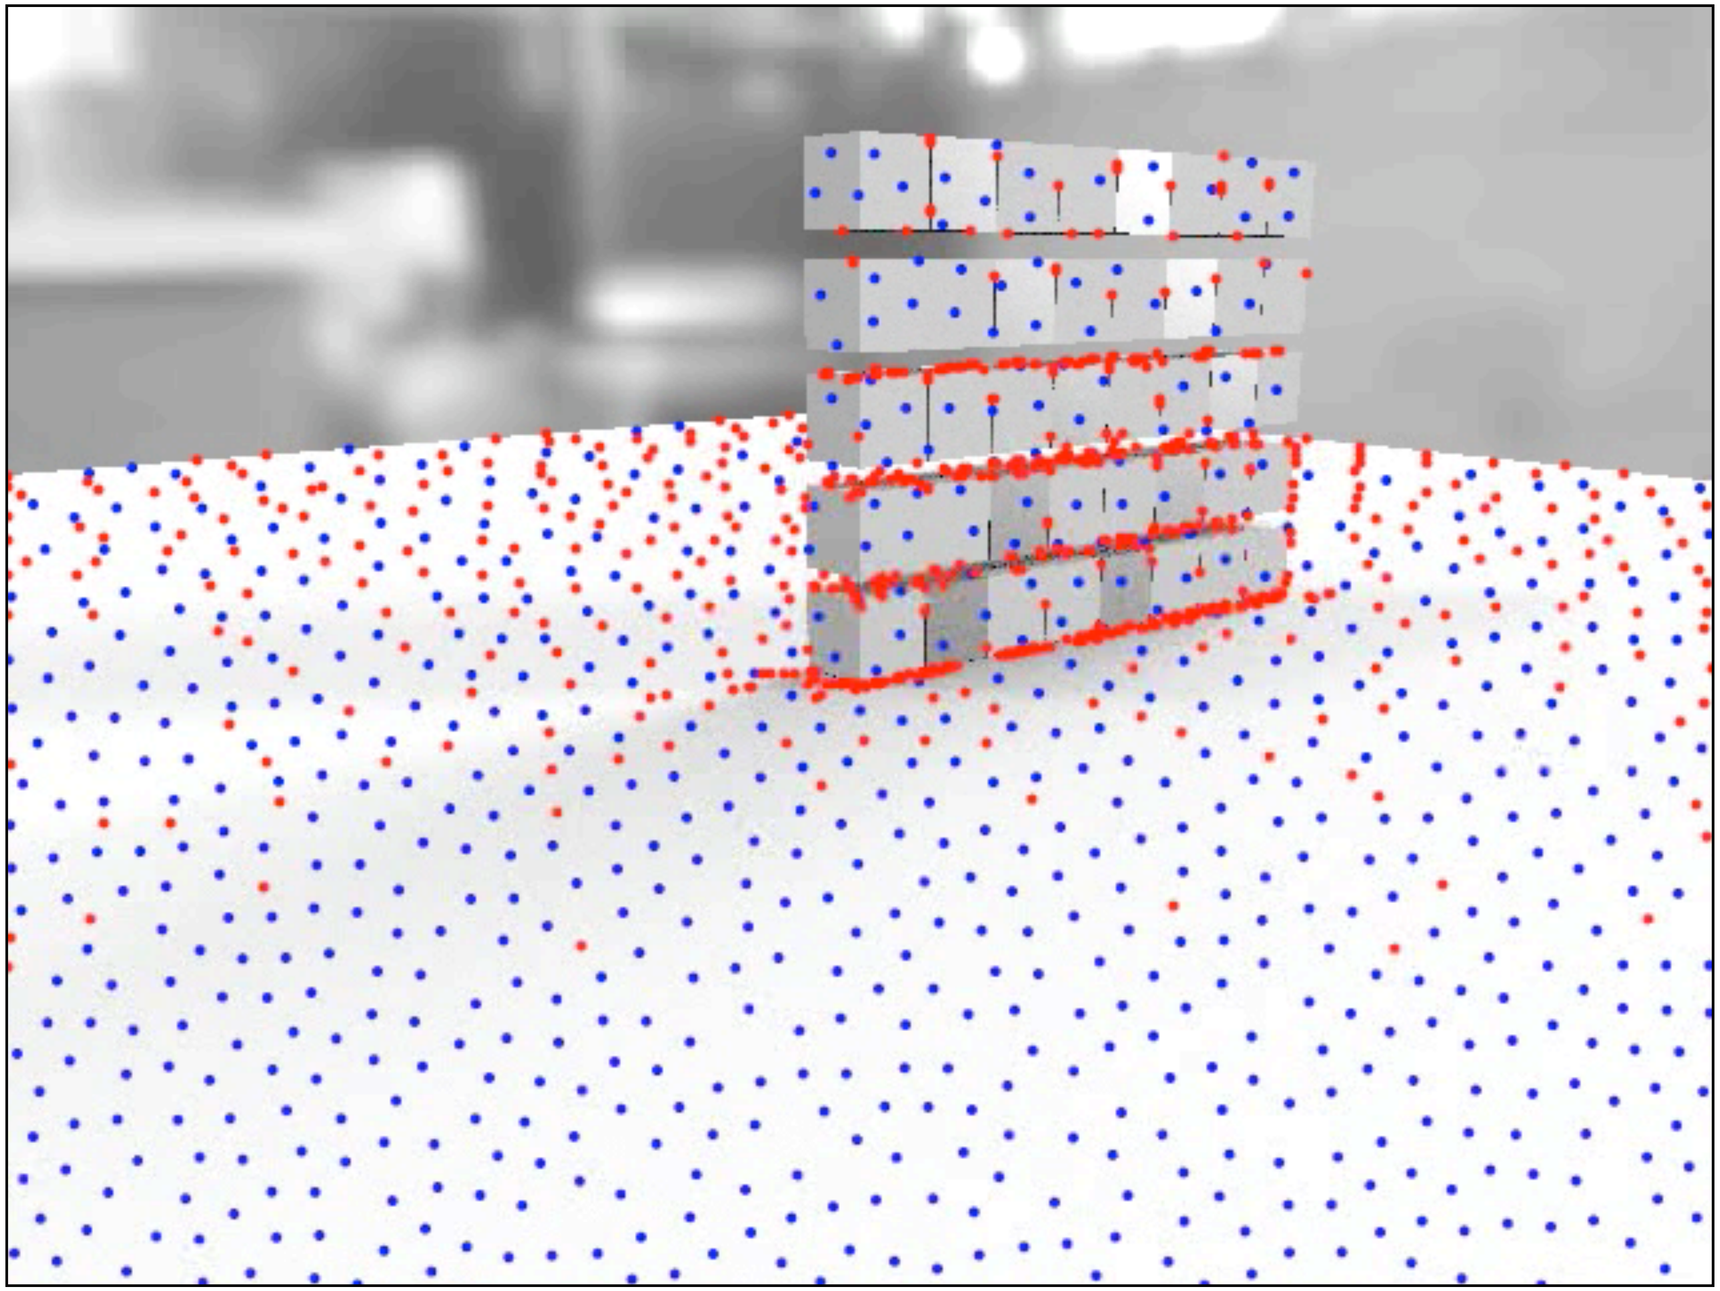

Video

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Implementation

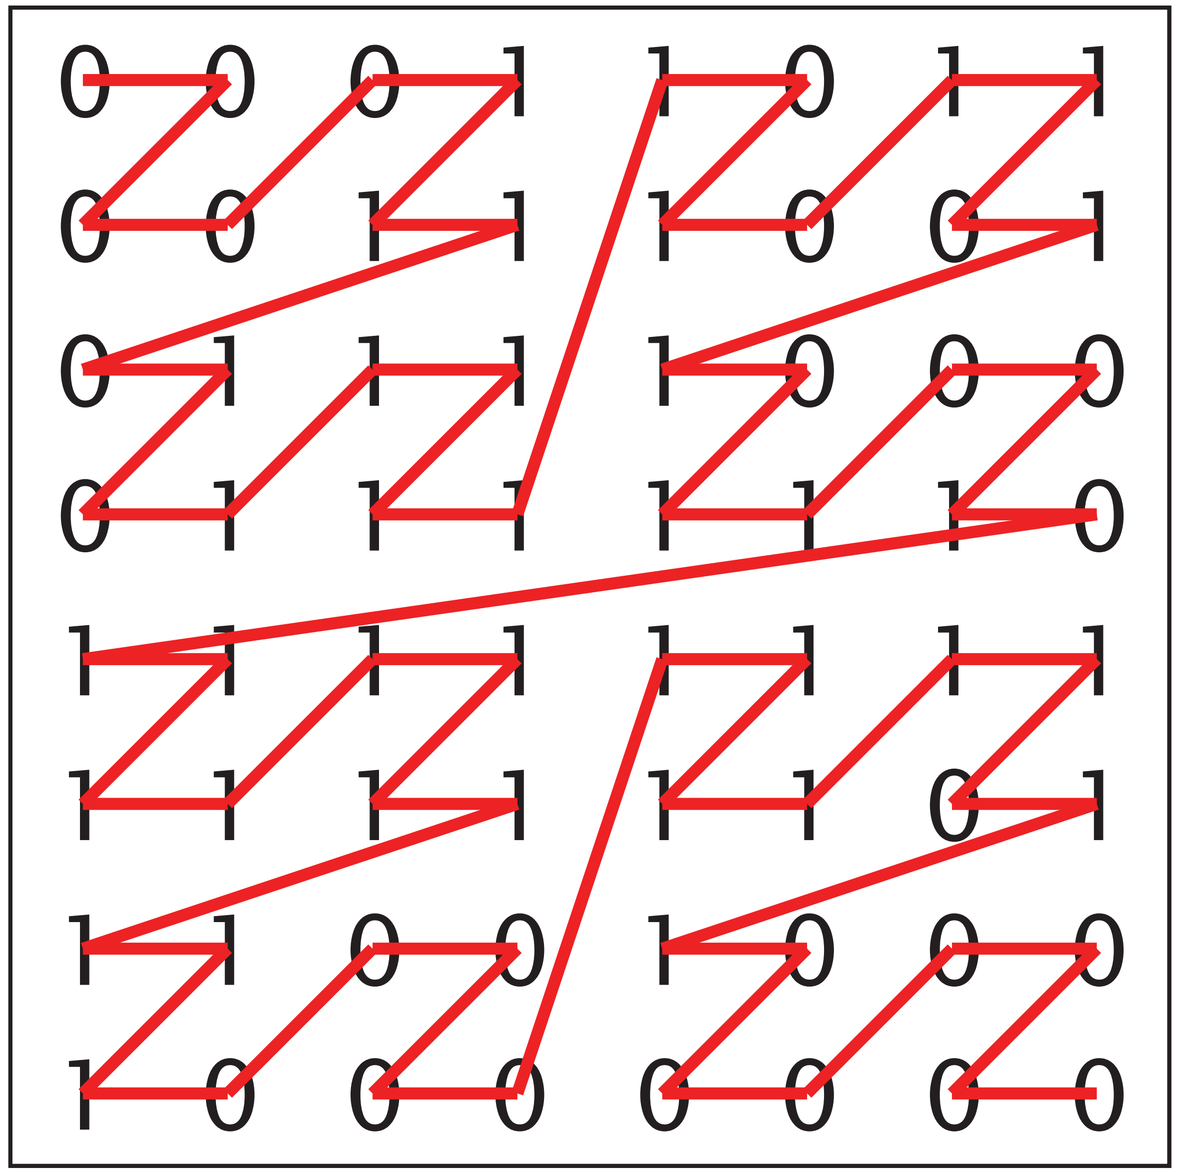

00000111010111111010110110...10110000000

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Implementation

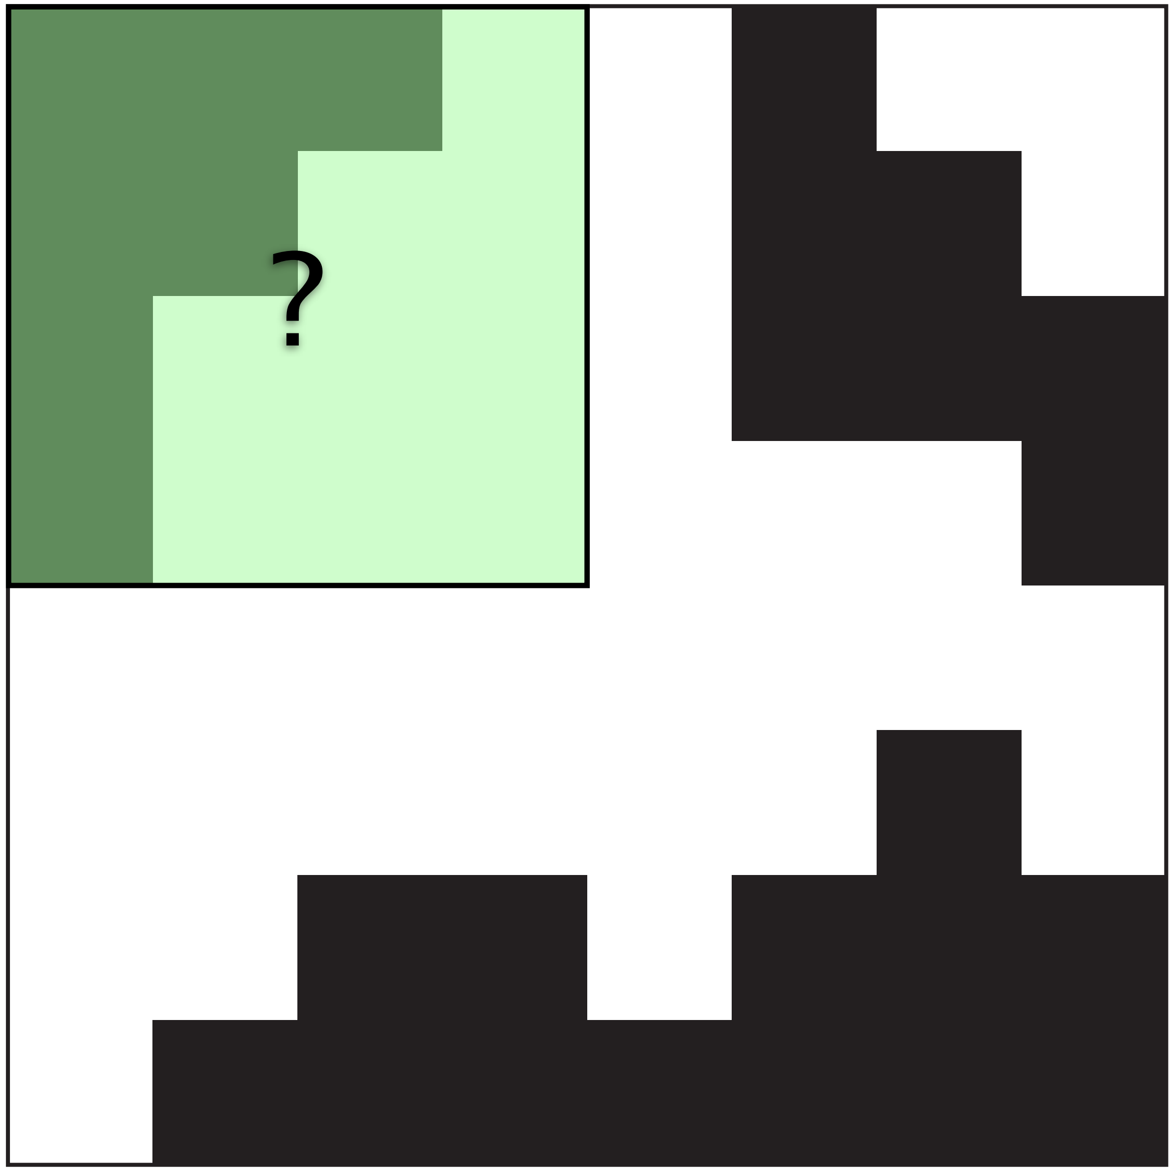

0000011101011111 1010110110 . . . 1011000000

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Implementation

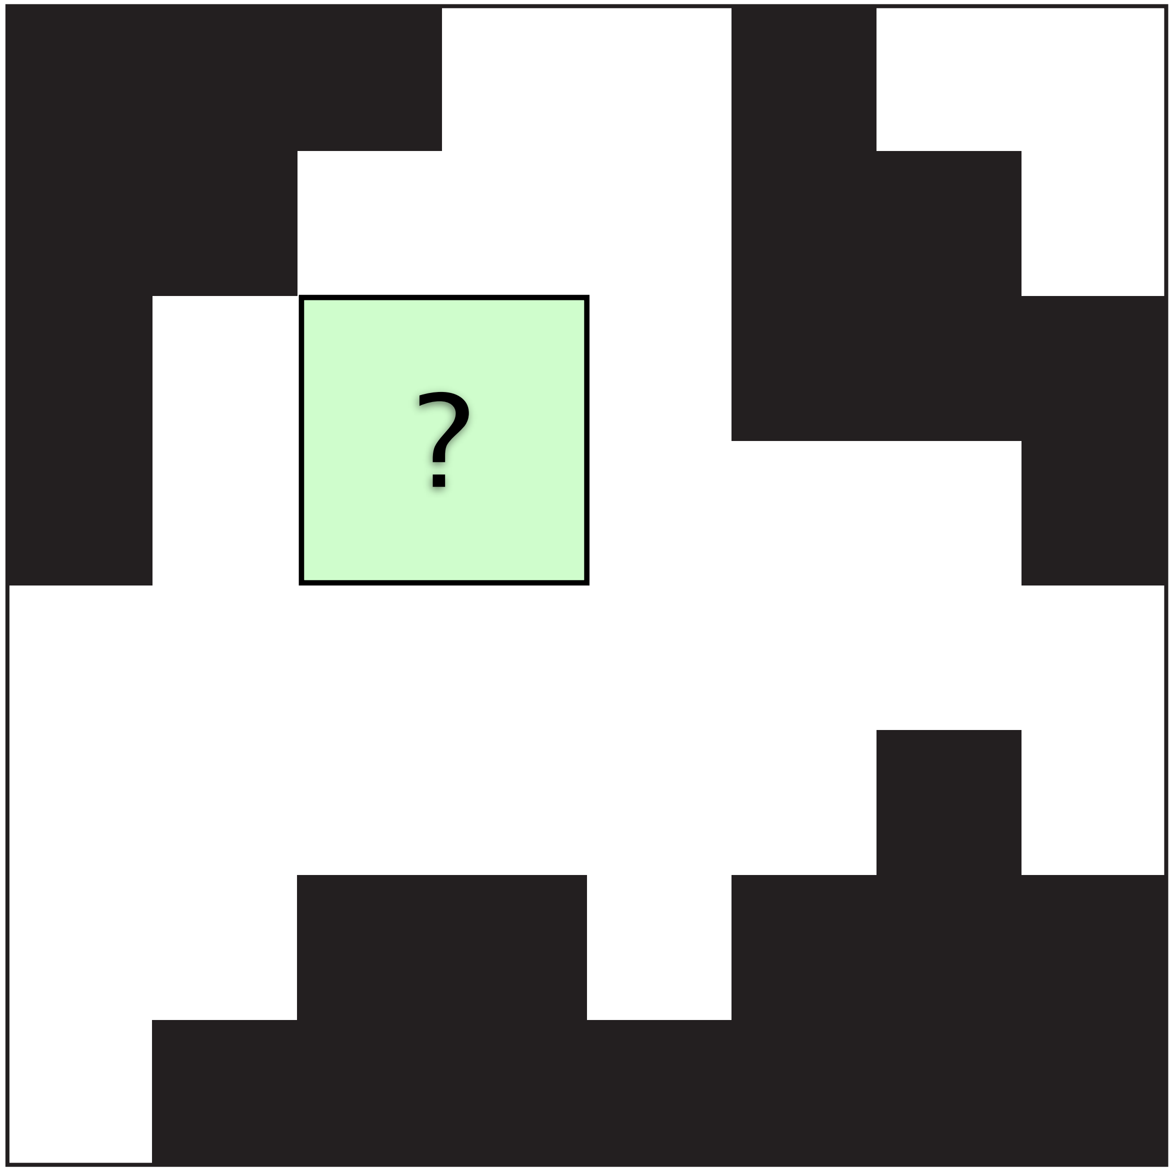

0000011101011111010110110 . . . 10110000000

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Visibility Difference

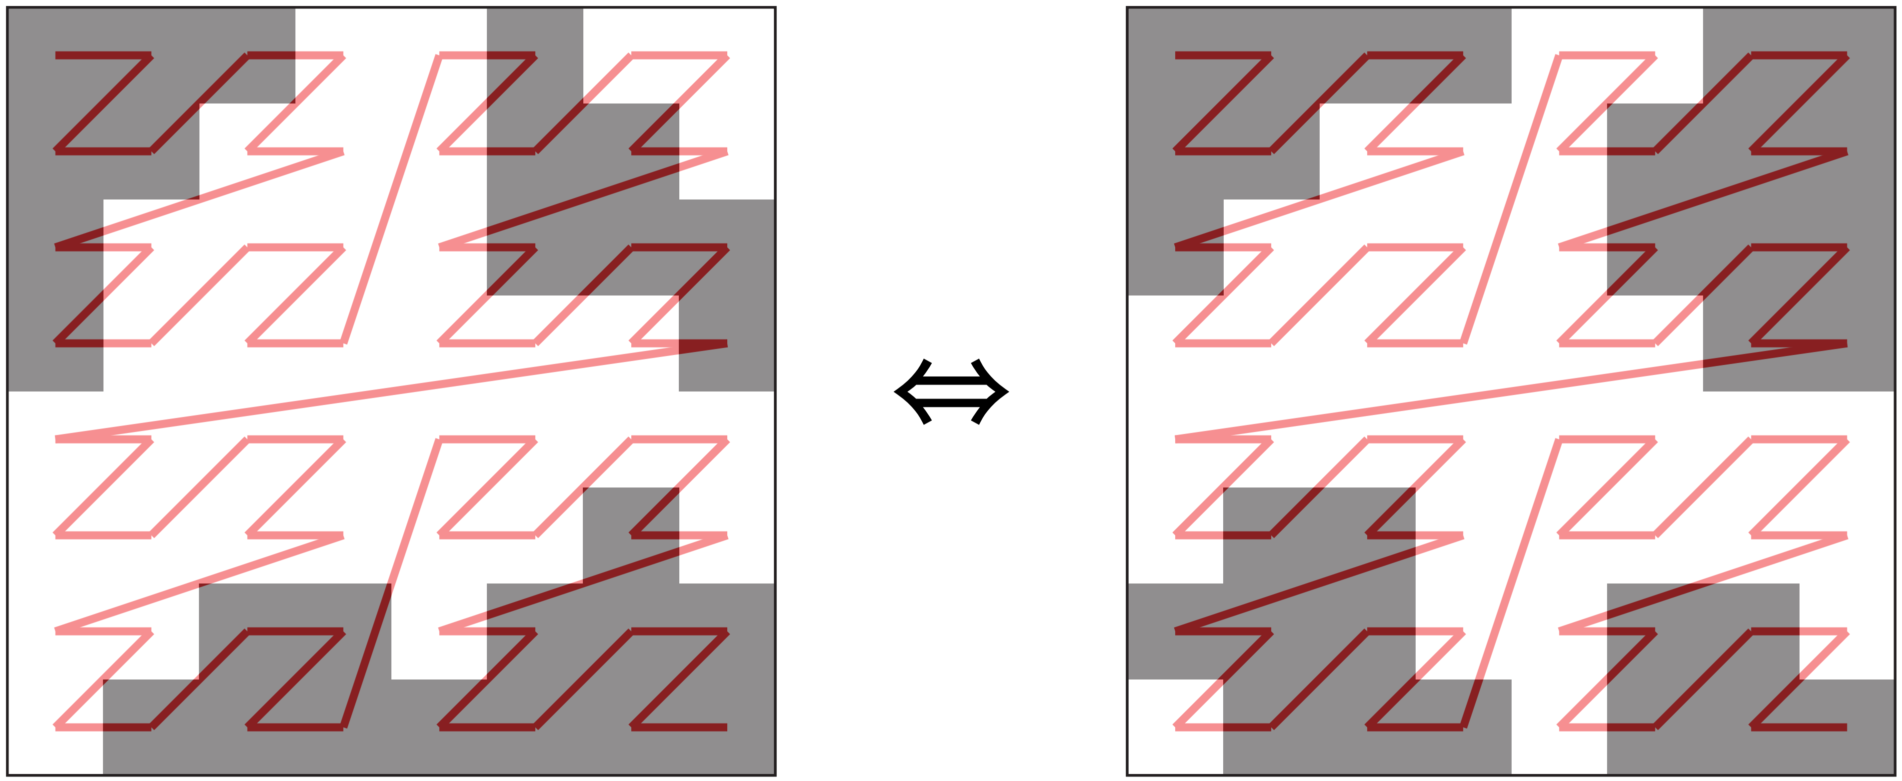

0000011101011111101011011011001011111111111000001111110110000000
0000001101111111111000001011000011101101001001001111111110100100

0000010000100000010011010000001000010010110001000000001000100100

1) XOR

SSE: 128 bits at the time

2) Count bits ⇒ 15 ⇒ **23.4%** difference

also using SSE

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Memory Use

Example

Visibility map size: 32x32 pixels

1 bit/pixel ⇒ **128 bytes + 44 bytes** book-keeping

Number of cache records: 11,900

Total memory use: **2 MB** (fits in CPU's L2 cache)

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

App #1 Monte Carlo

Use <u>Control Variates</u> to reduce noise:

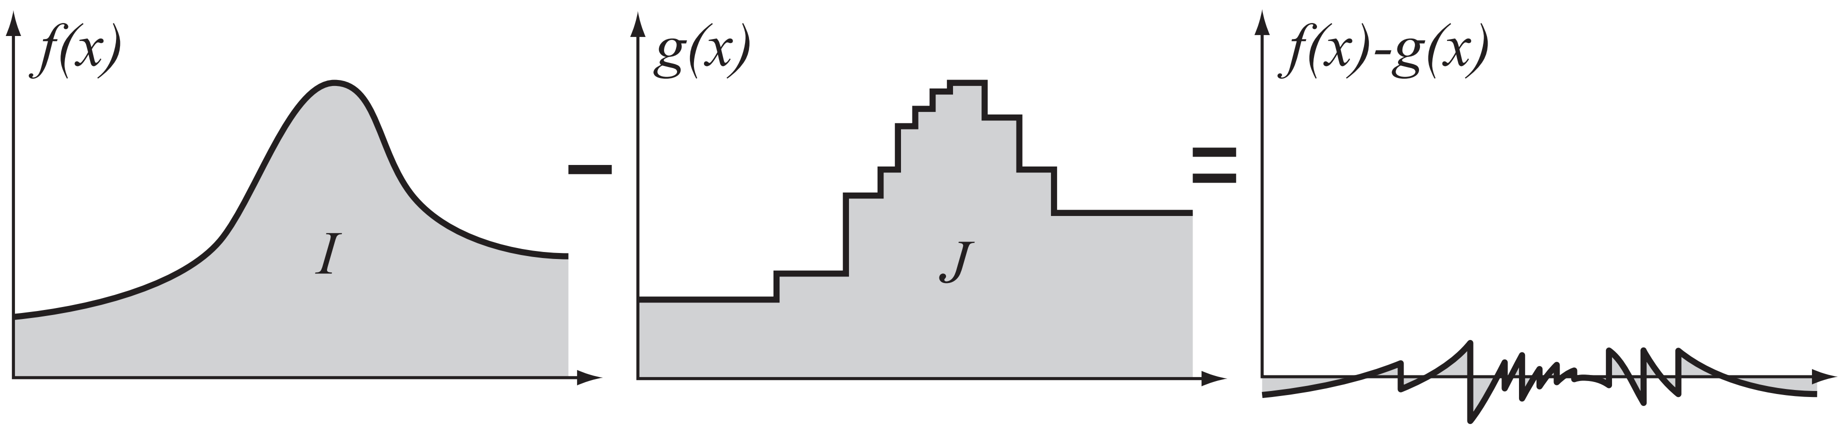

Original function | Correlated approximation | Difference has lower variance

$$L_{\text{out}} = \int LBV\, d\omega = \int LBV - L\tilde{B}\tilde{V}\, d\omega + \int L\tilde{B}\tilde{V}\, d\omega$$

Petrik Clarberg and Tomas Akenine-Möller, “*Exploiting Visibility Correlation in Direct Illumination*”, Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

App #1 Monte Carlo

Use Control Variates to reduce noise:

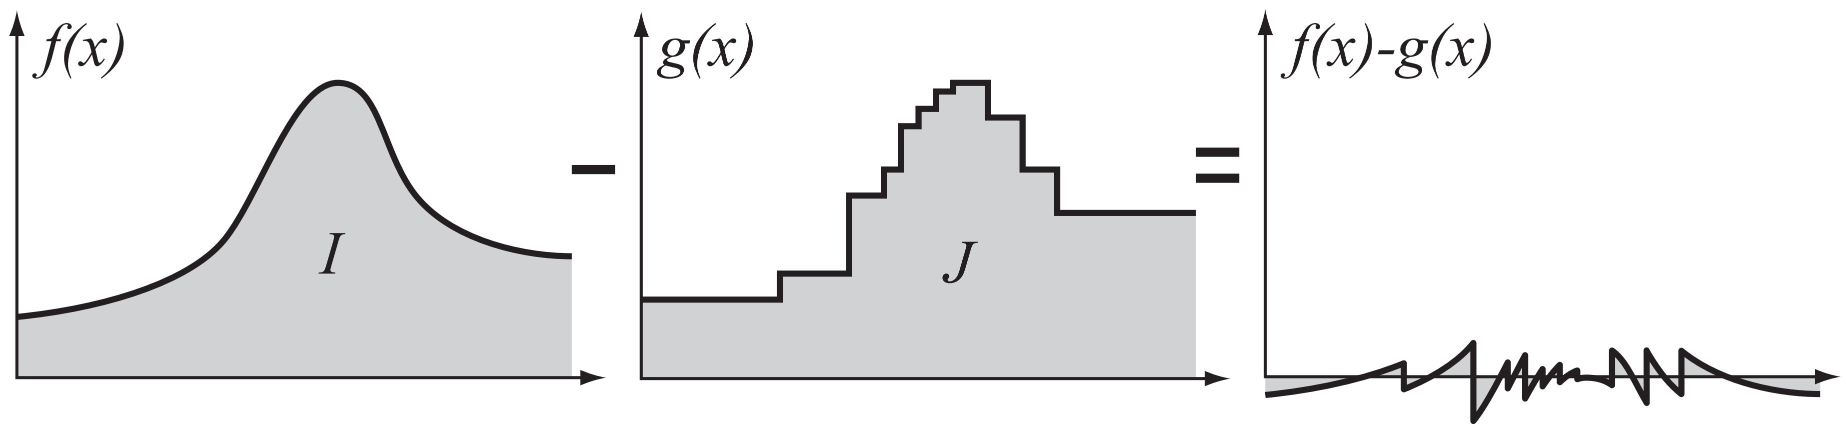

Original function

Correlated approximation

Difference has lower variance

$$L_{\text{out}} = \int LBV\,d\omega = \int LBV - L\tilde{B}\tilde{V}\,d\omega + \int L\tilde{B}\tilde{V}\,d\omega$$

Given by Visibility Cache

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

App #1 Monte Carlo

Use Control Variates to reduce noise:

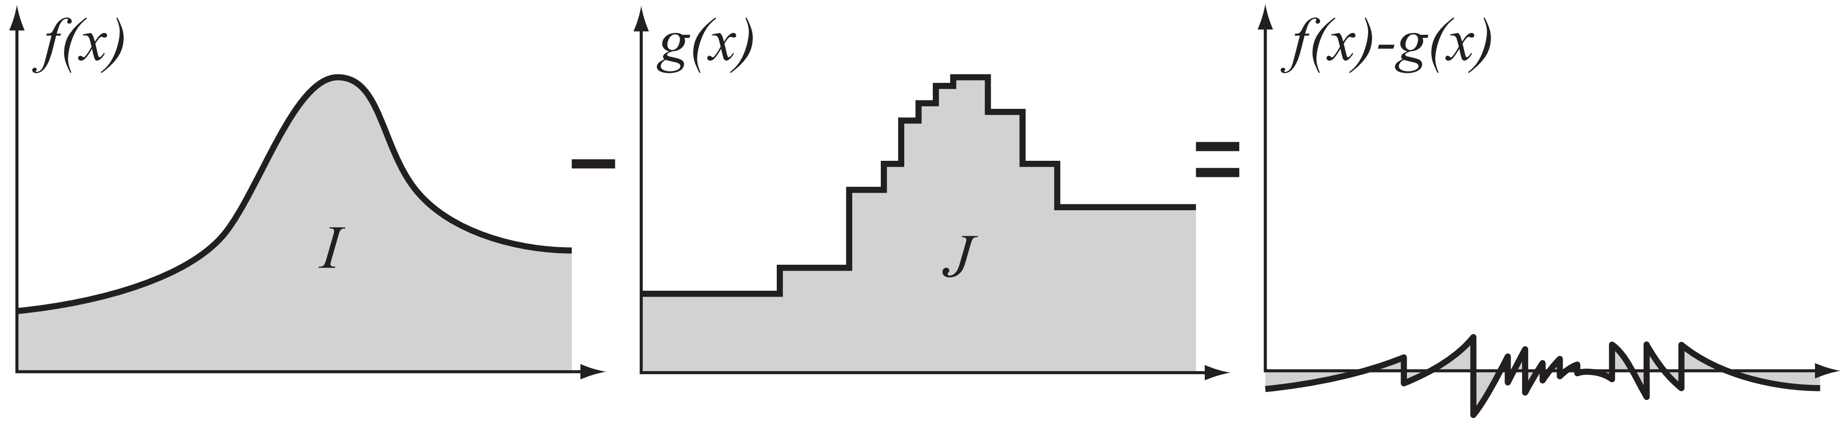

Original function

Correlated approximation

Difference has lower variance

$$L_{\text{out}} = \int LBV\,d\omega = \int LBV - L\tilde{B}\tilde{V}\,d\omega + \int L\tilde{B}\tilde{V}\,d\omega$$

Given by existing product sampling algorithms [Cline et al. 2006, Clarberg et al. 2008]

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

App #1 Monte Carlo

Use Control Variates to reduce noise:

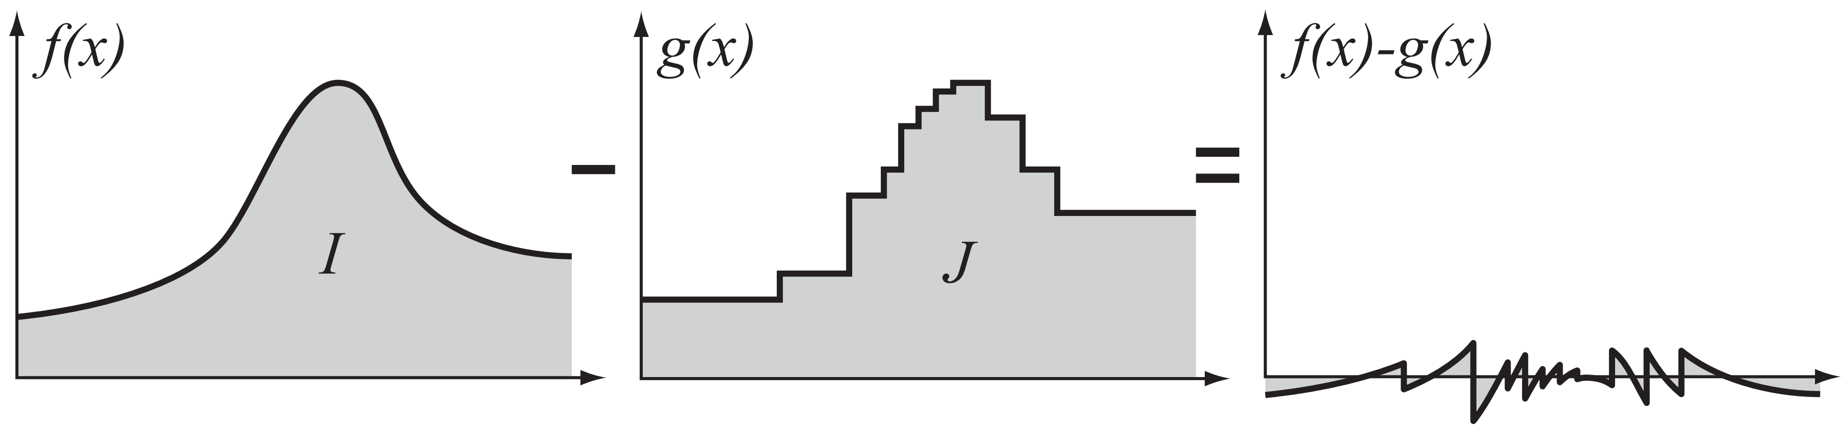

Original function | Correlated approximation | Difference has lower variance

Unbiased !

$$L_{\text{out}} = \int LBV\,d\omega = \int LBV - L\tilde{B}\tilde{V}\,d\omega + \int L\tilde{B}\tilde{V}\,d\omega$$

Evaluated using importance sampling | Given by quick quadtree product

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Equal-Time Results

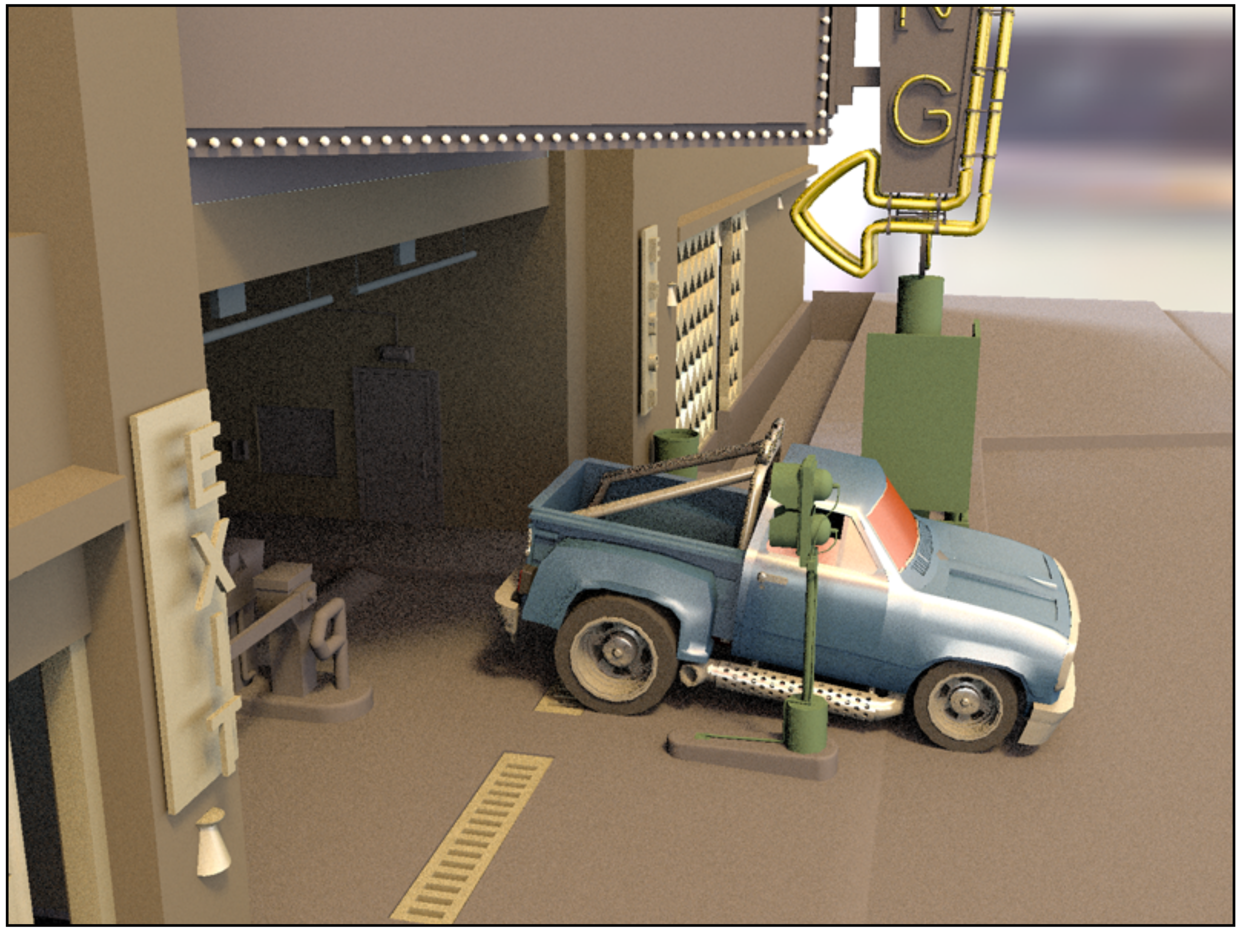

State-of-the-art
product sampling
(39 samples/pixel)

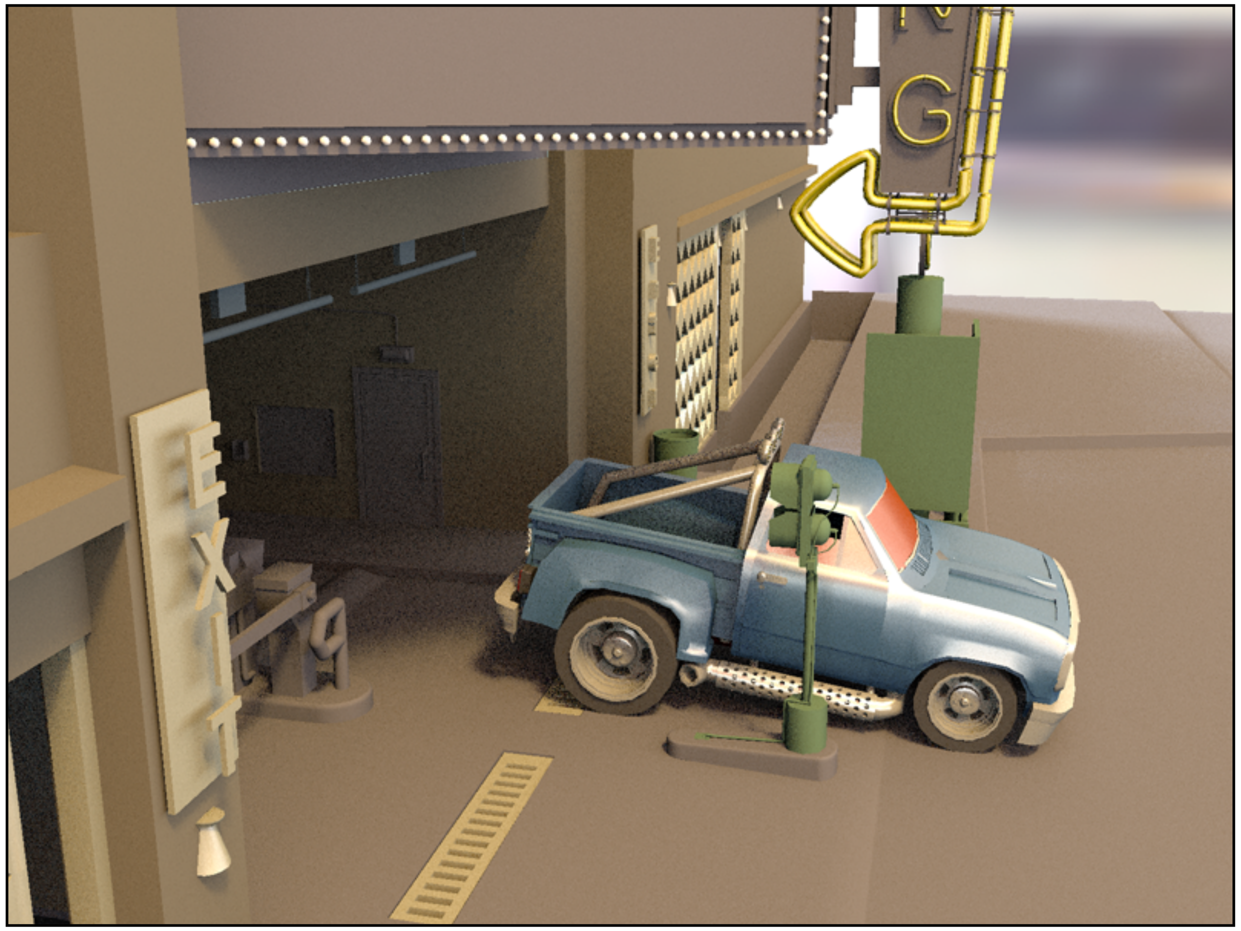

Our algorithm
(30 samples/pixel)

Petrik Clarberg and Tomas Akenine-Möller, “*Exploiting Visibility Correlation in Direct Illumination*”, Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Product Sampling

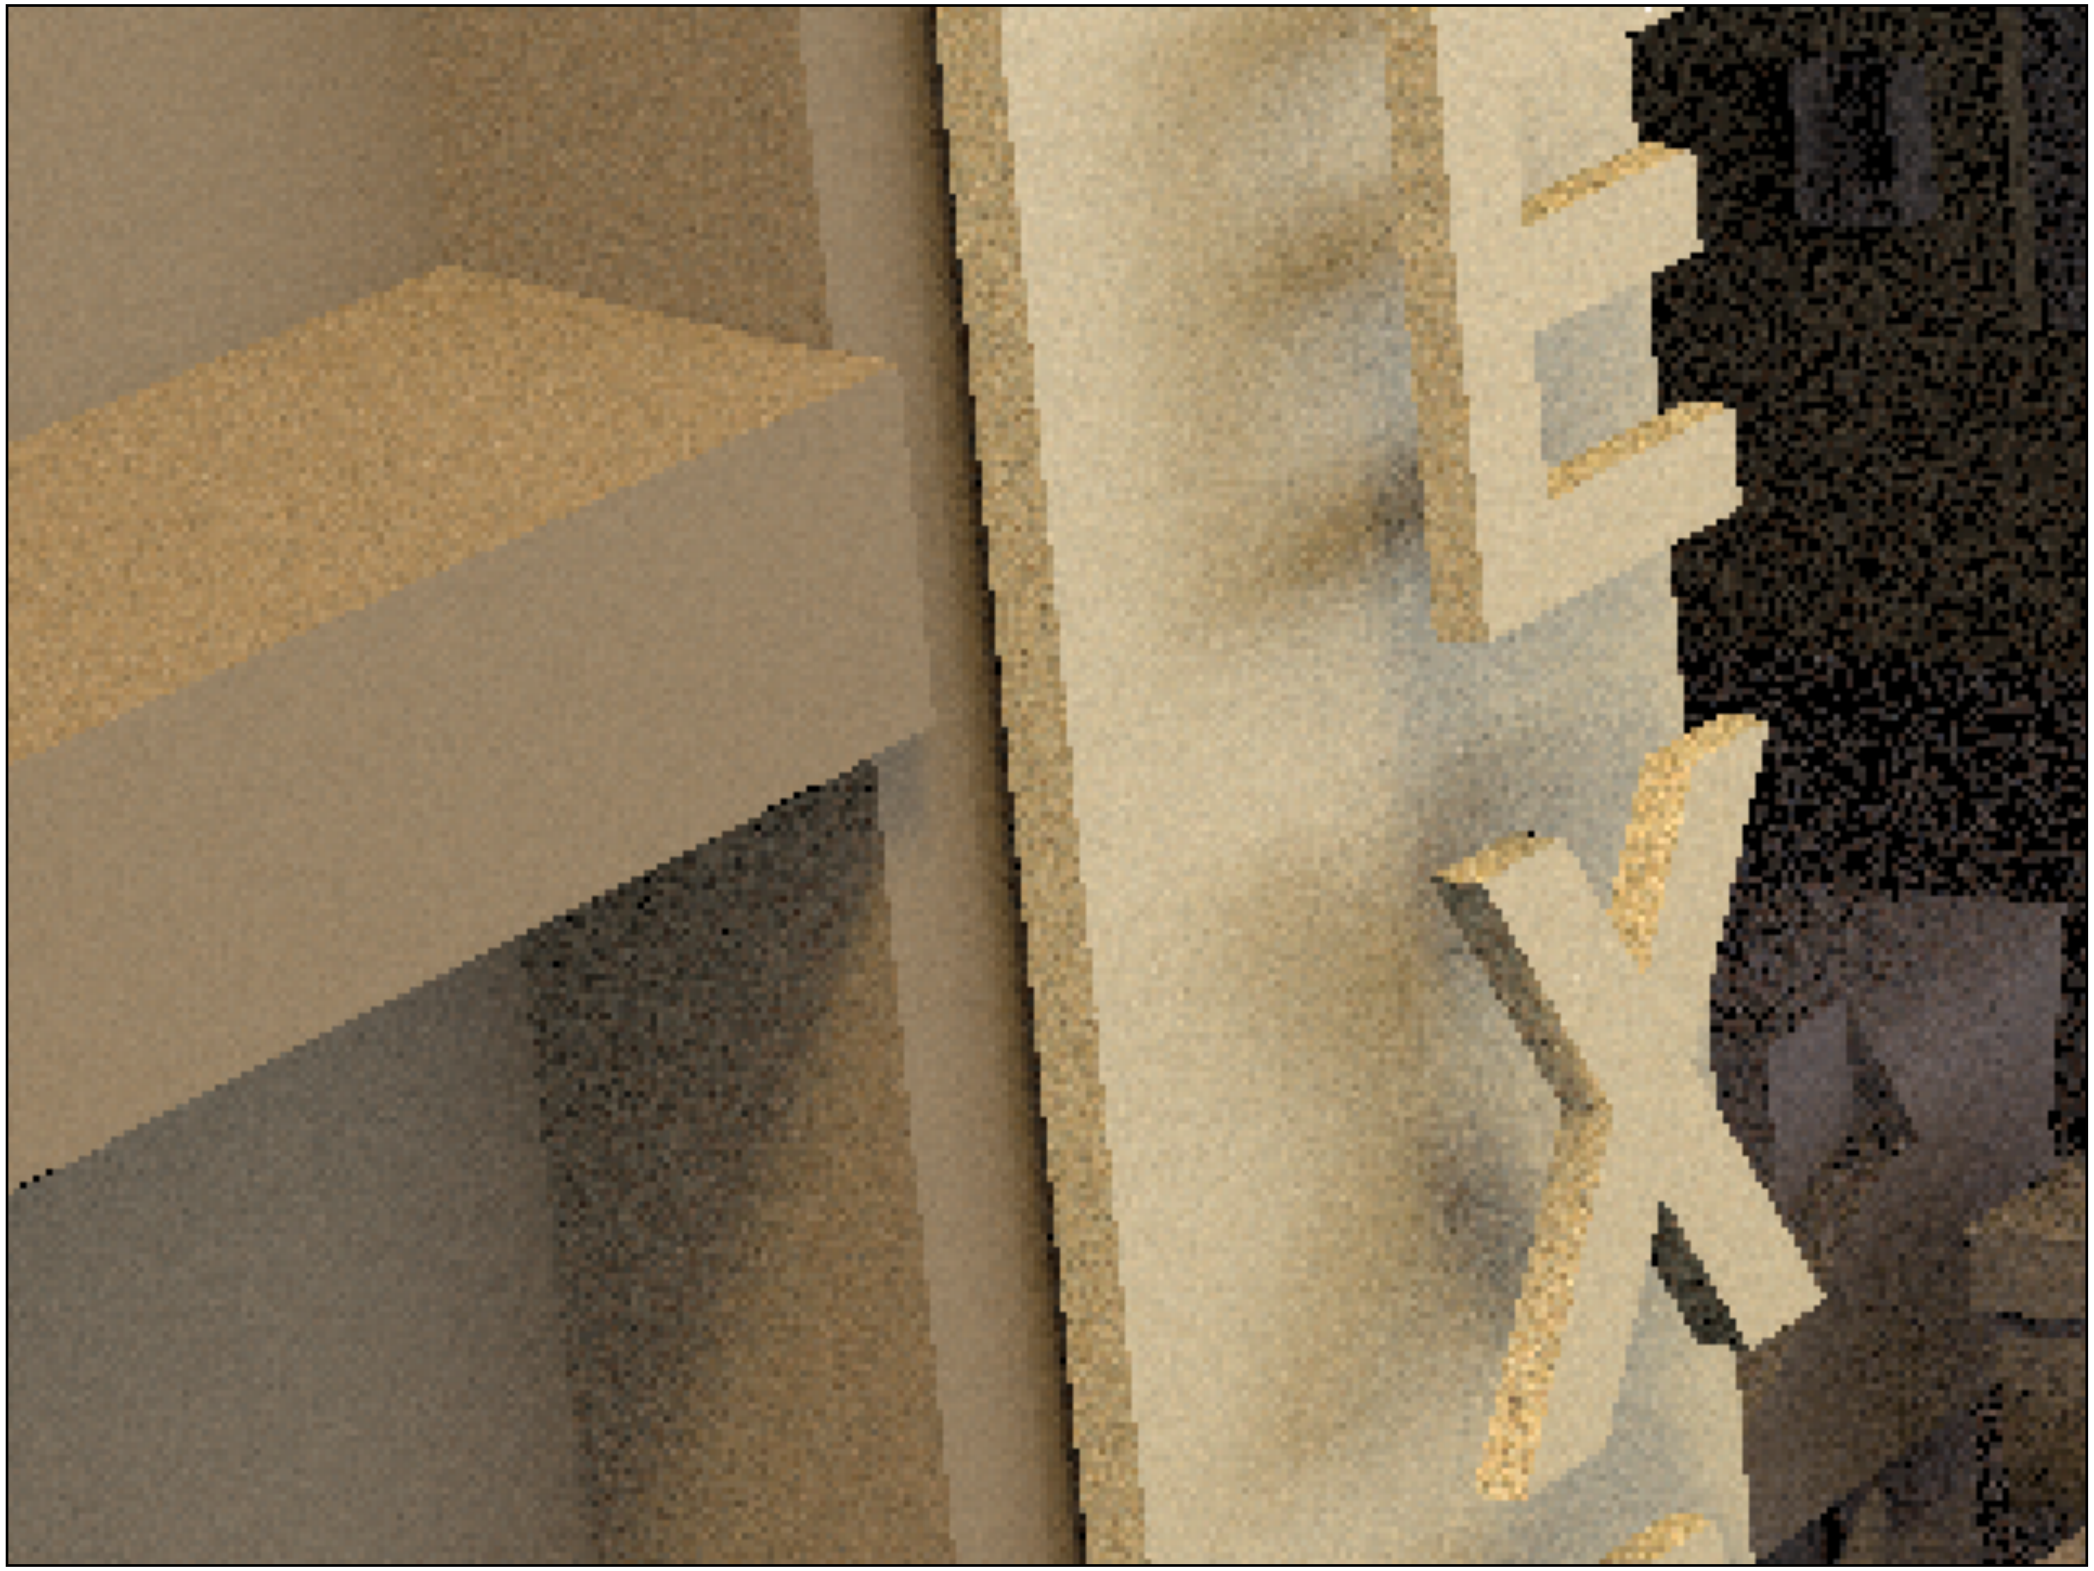

Petrik Clarberg and Tomas Akenine-Möller, “*Exploiting Visibility Correlation in Direct Illumination*”, Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Our Algorithm

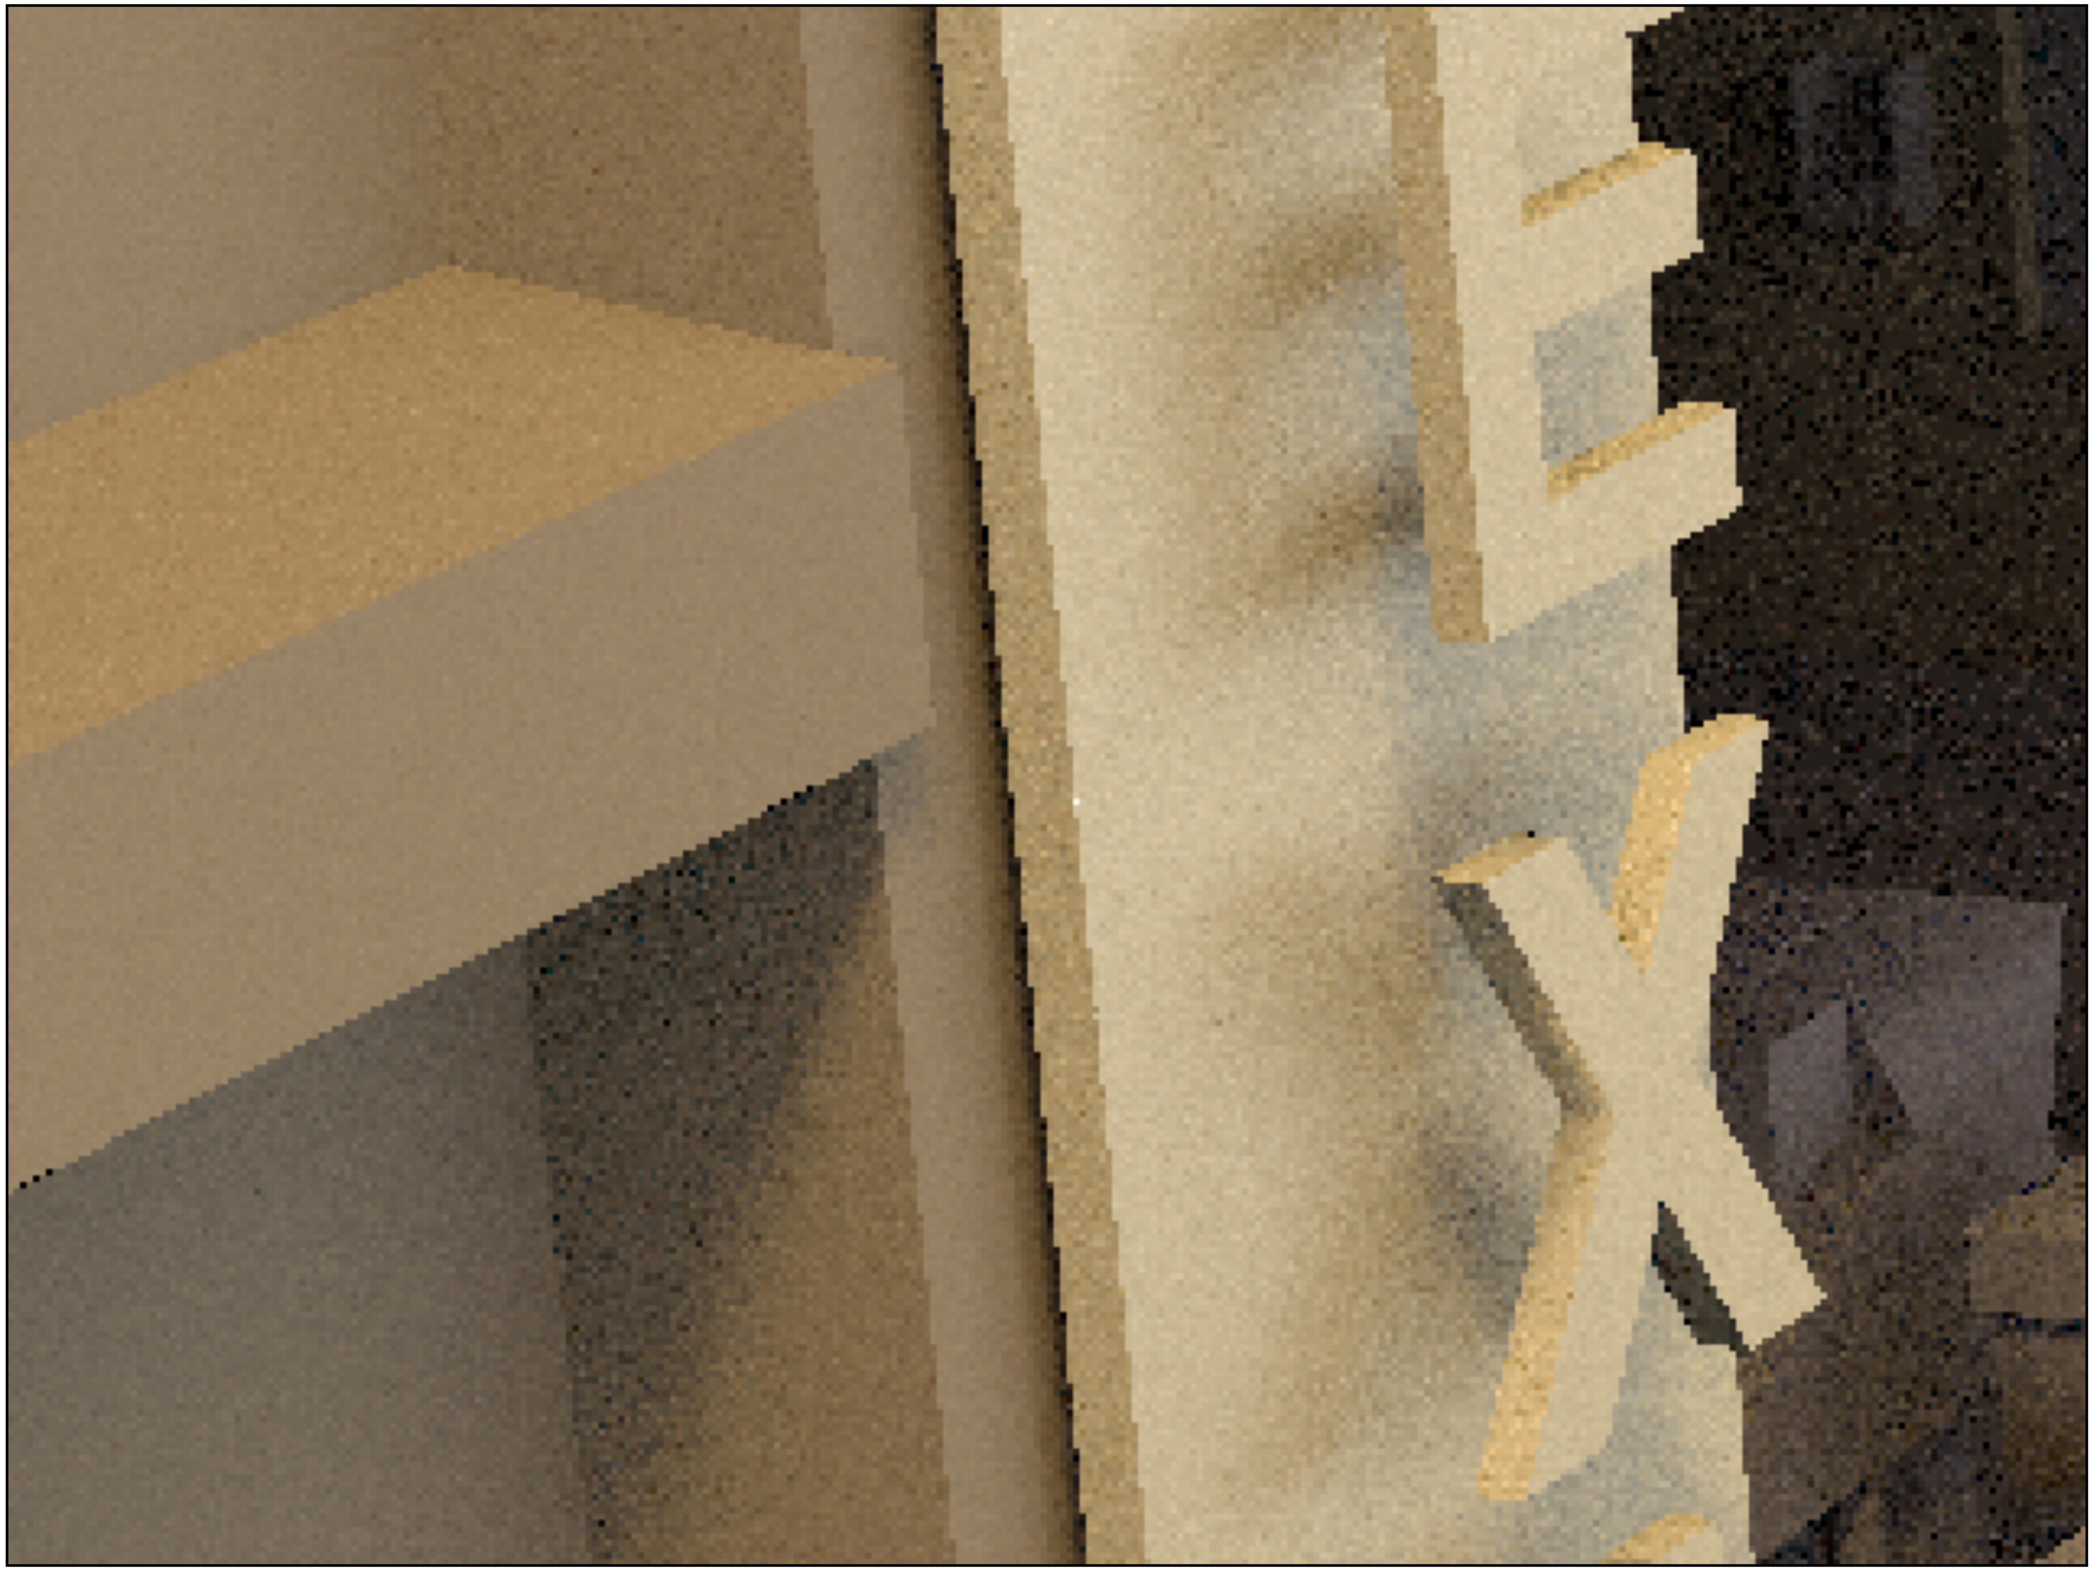

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Product Sampling

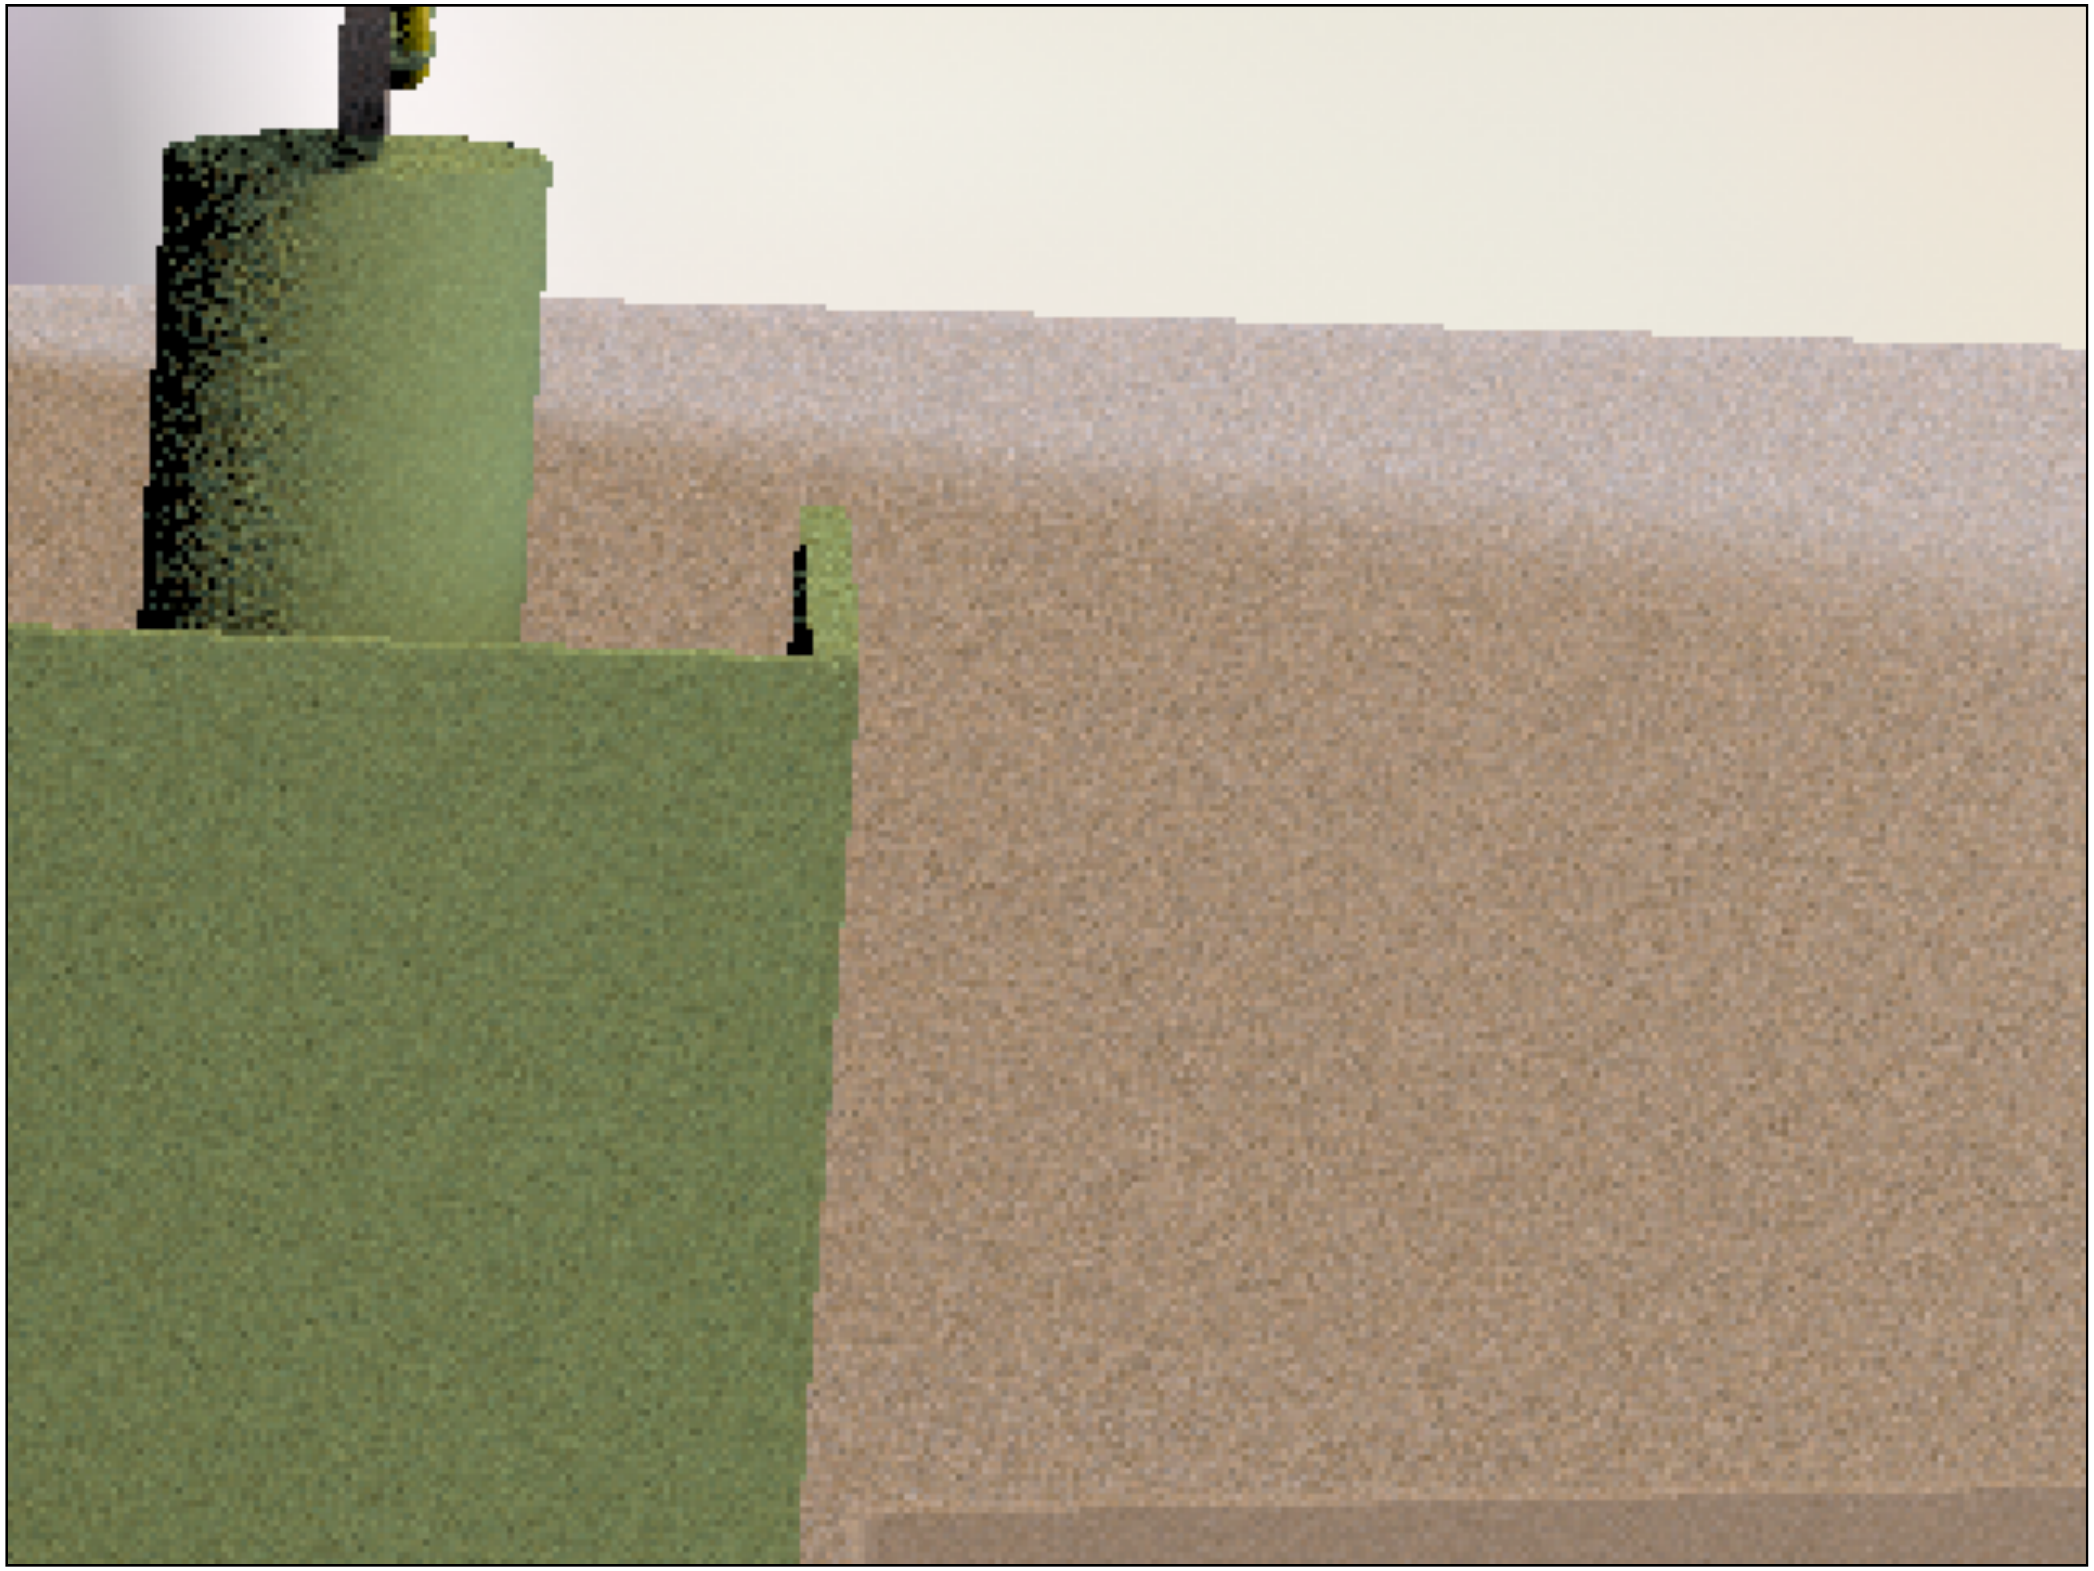

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Our Algorithm

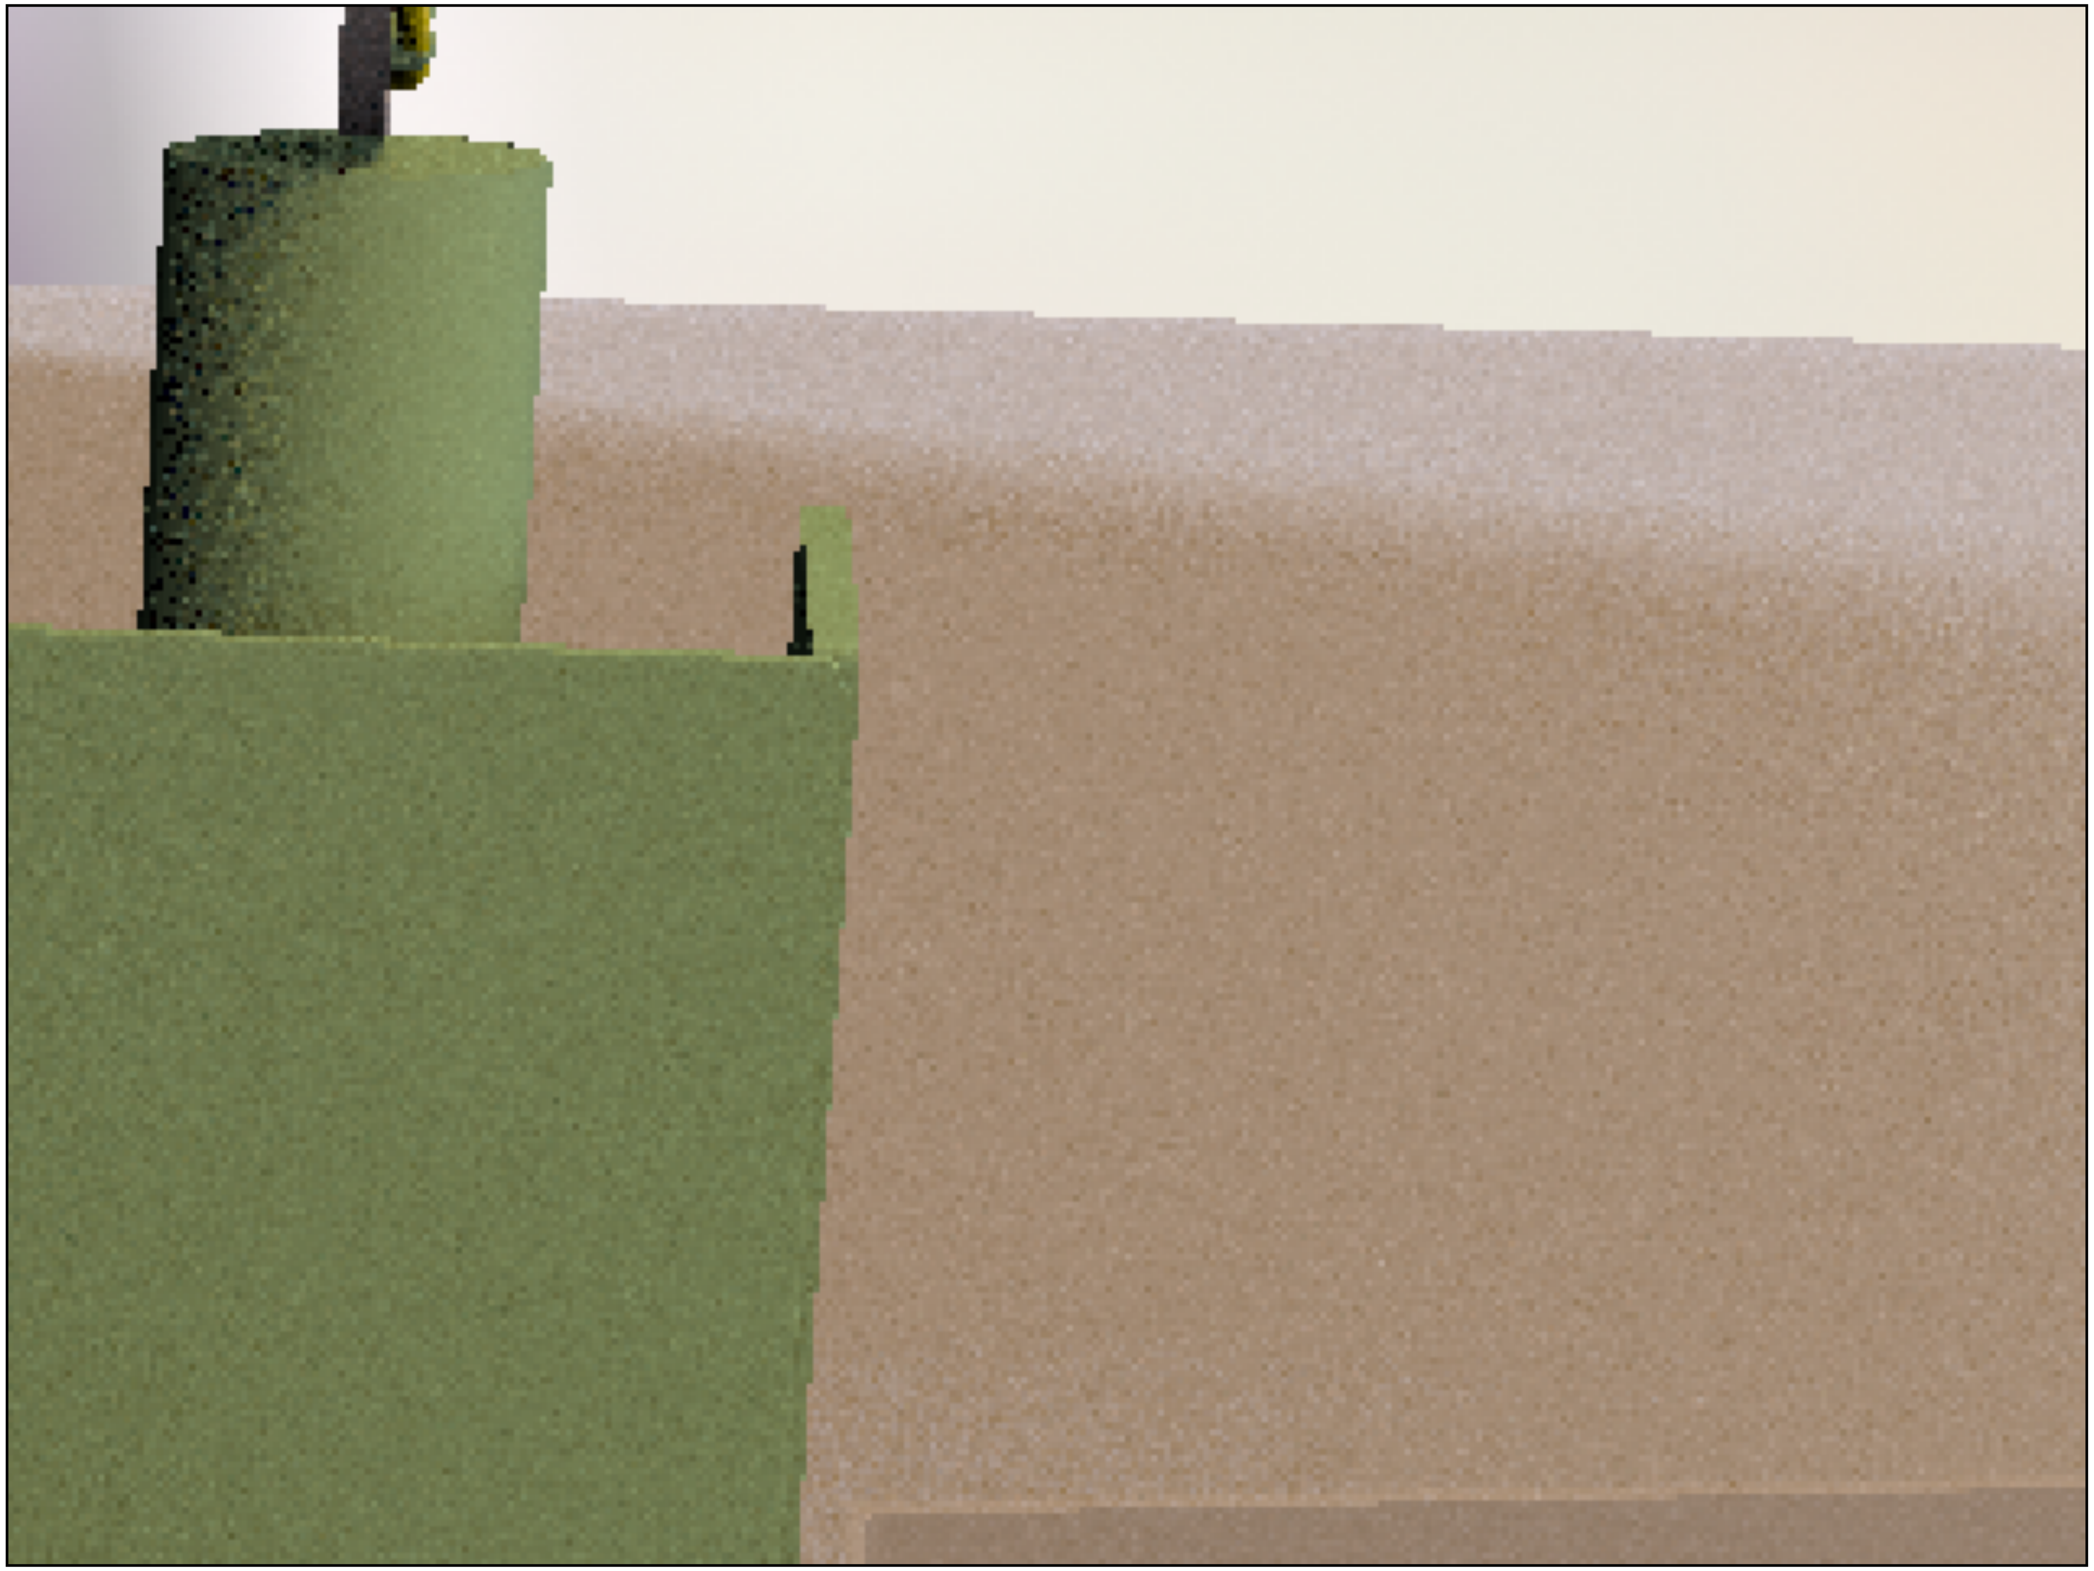

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Product Sampling

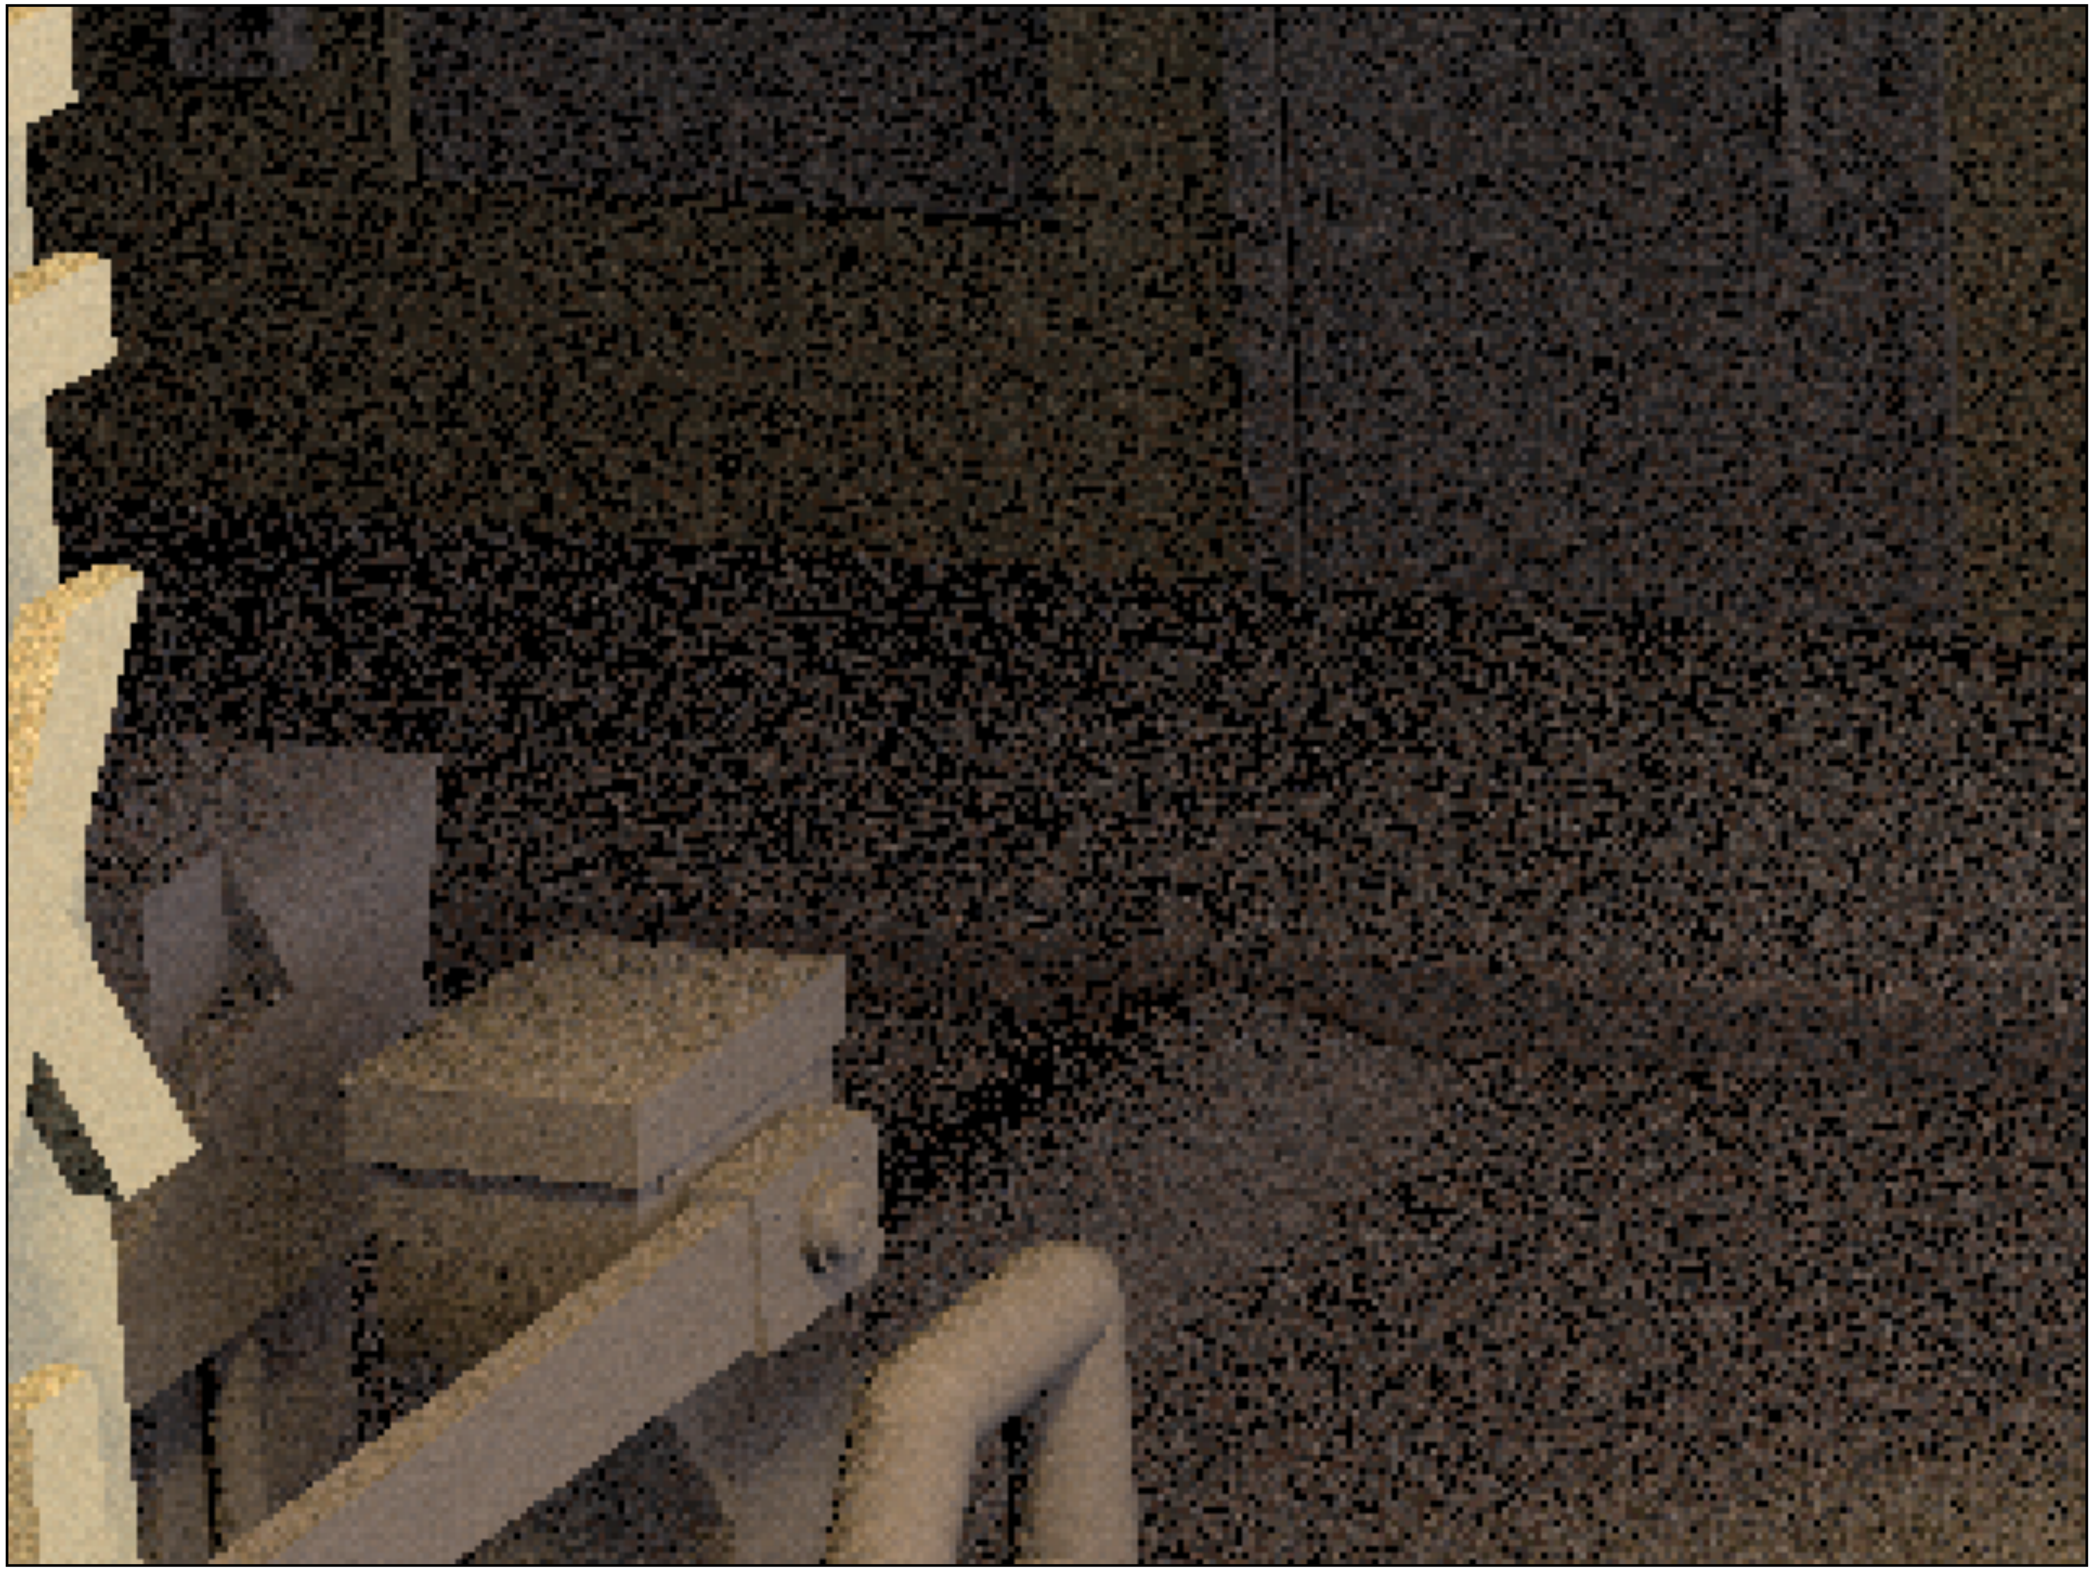

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Our Algorithm

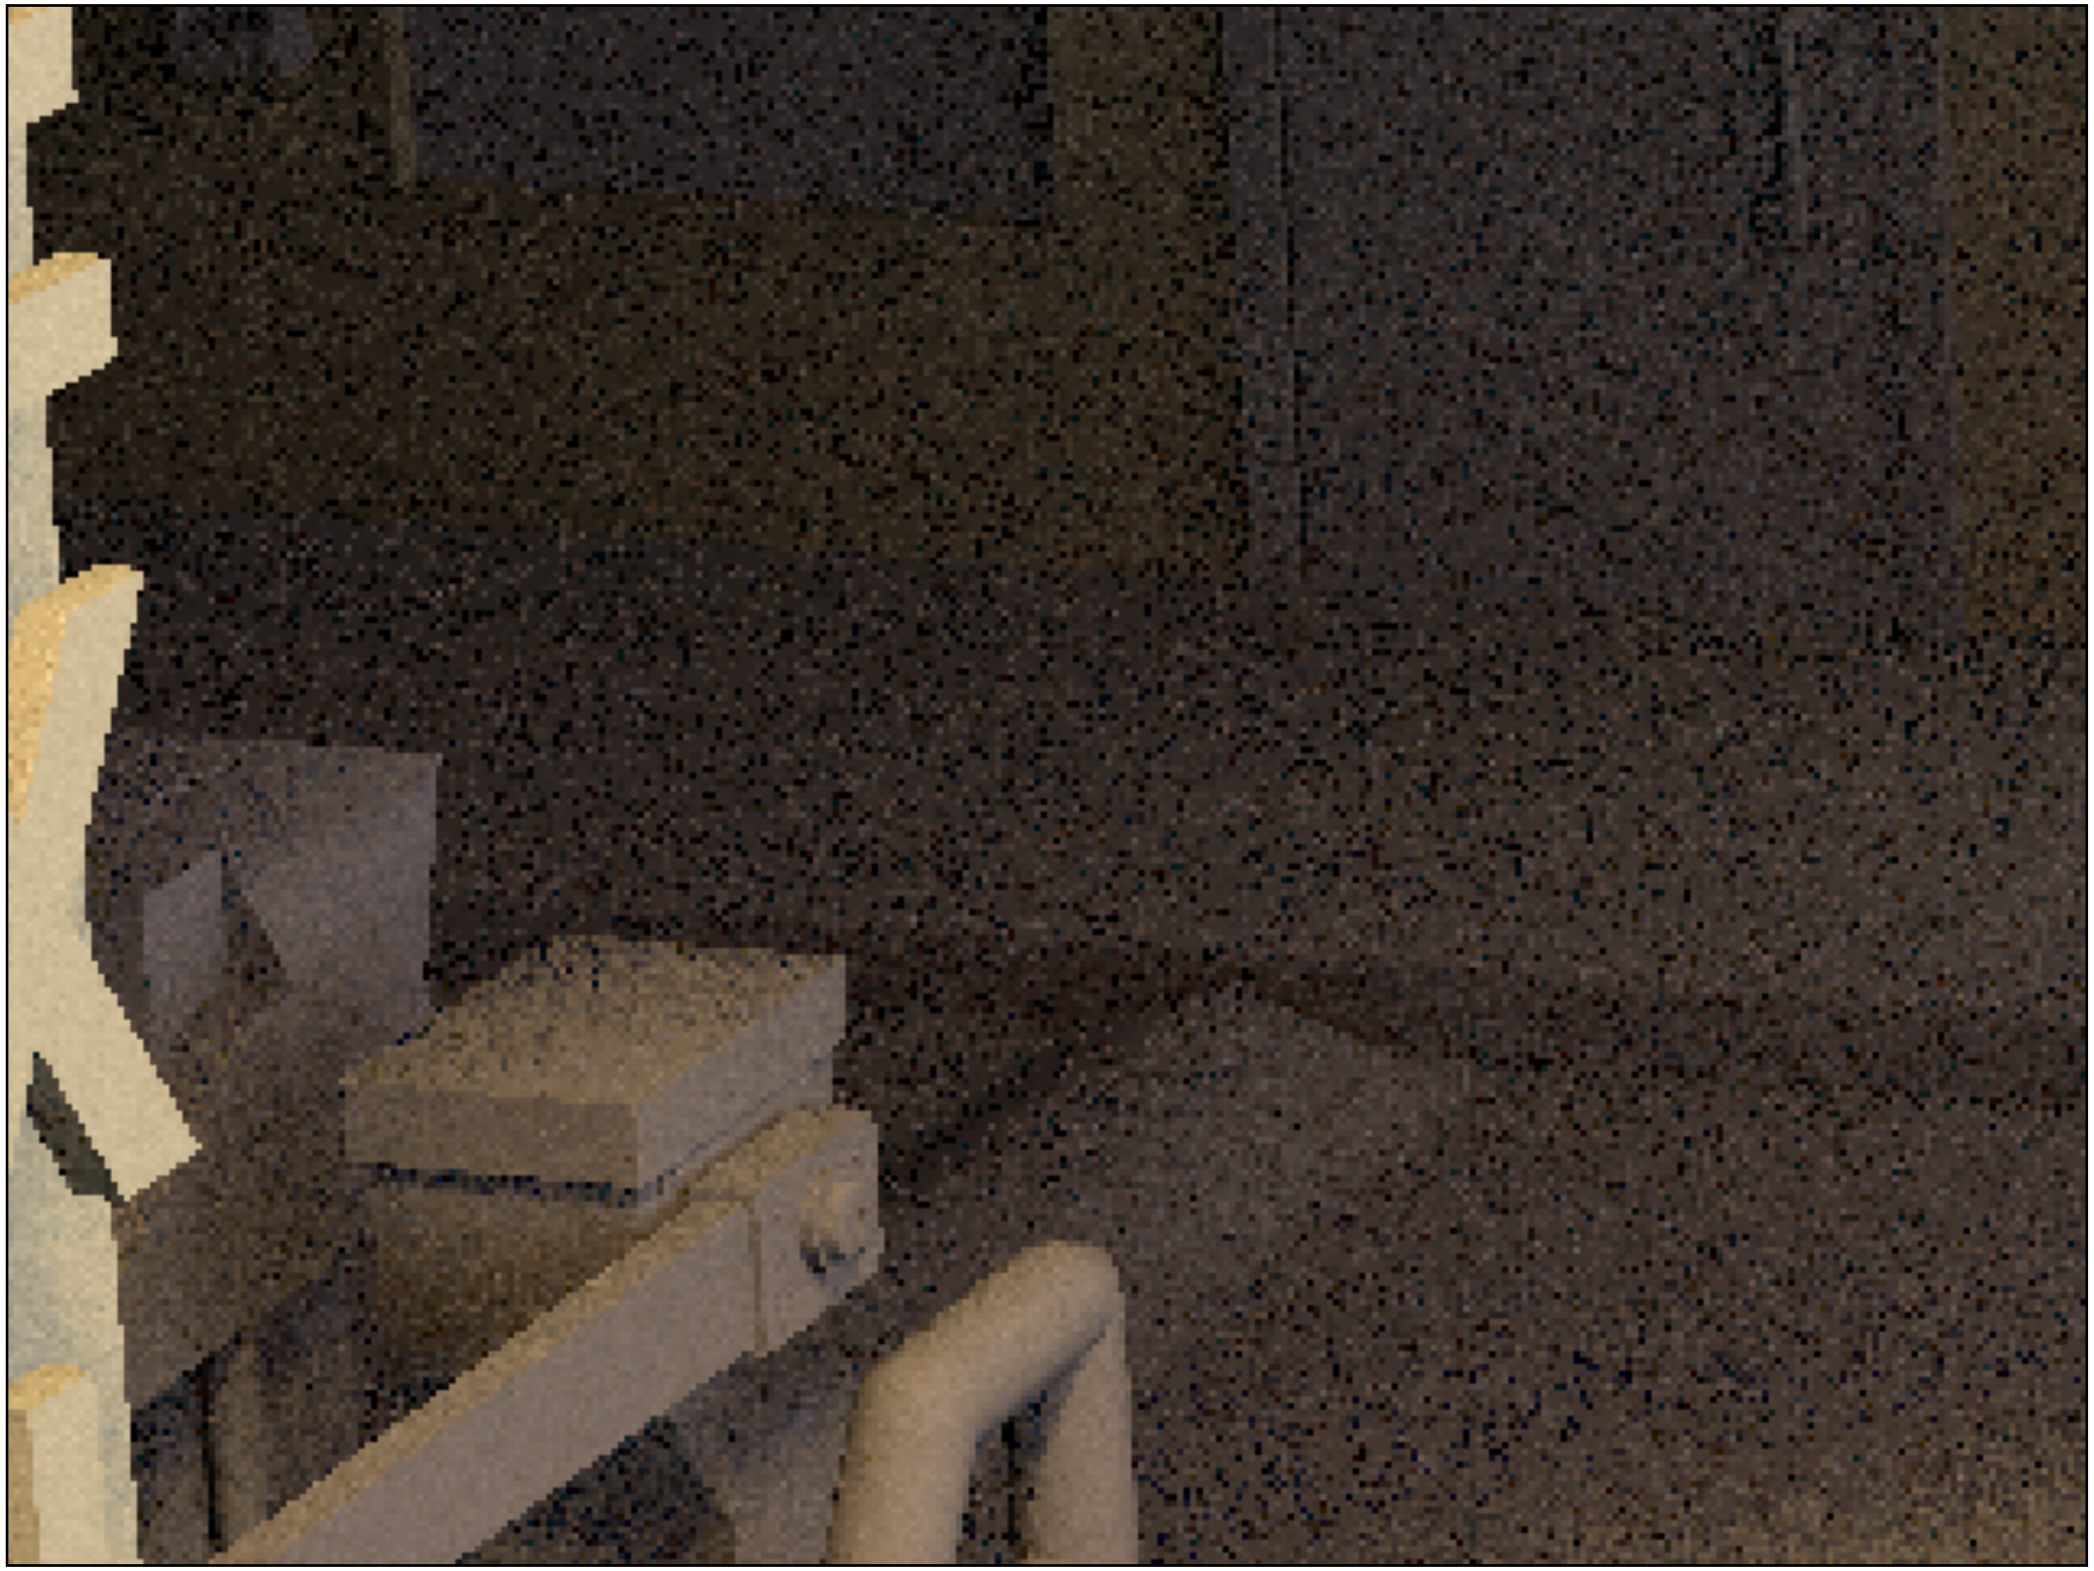

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

App #1 Summary

Good use of control variates (finally)

Unbiased even if visibility approximation is off

Animations: Re-use over time <u>without</u> lagging

Up to 5x variance reduction

Depends on scene complexity

- Some scenes have too low visibility correlation
- Largest improvement is in shadow regions

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

App #2 Ambient Occlusion

Pre-integrate records: $A_j = \int_\Omega V_j(\quad)(n \cdot \quad)d$

Approximate as weighted sum: $A(x) = \sum_i w_i A_i$

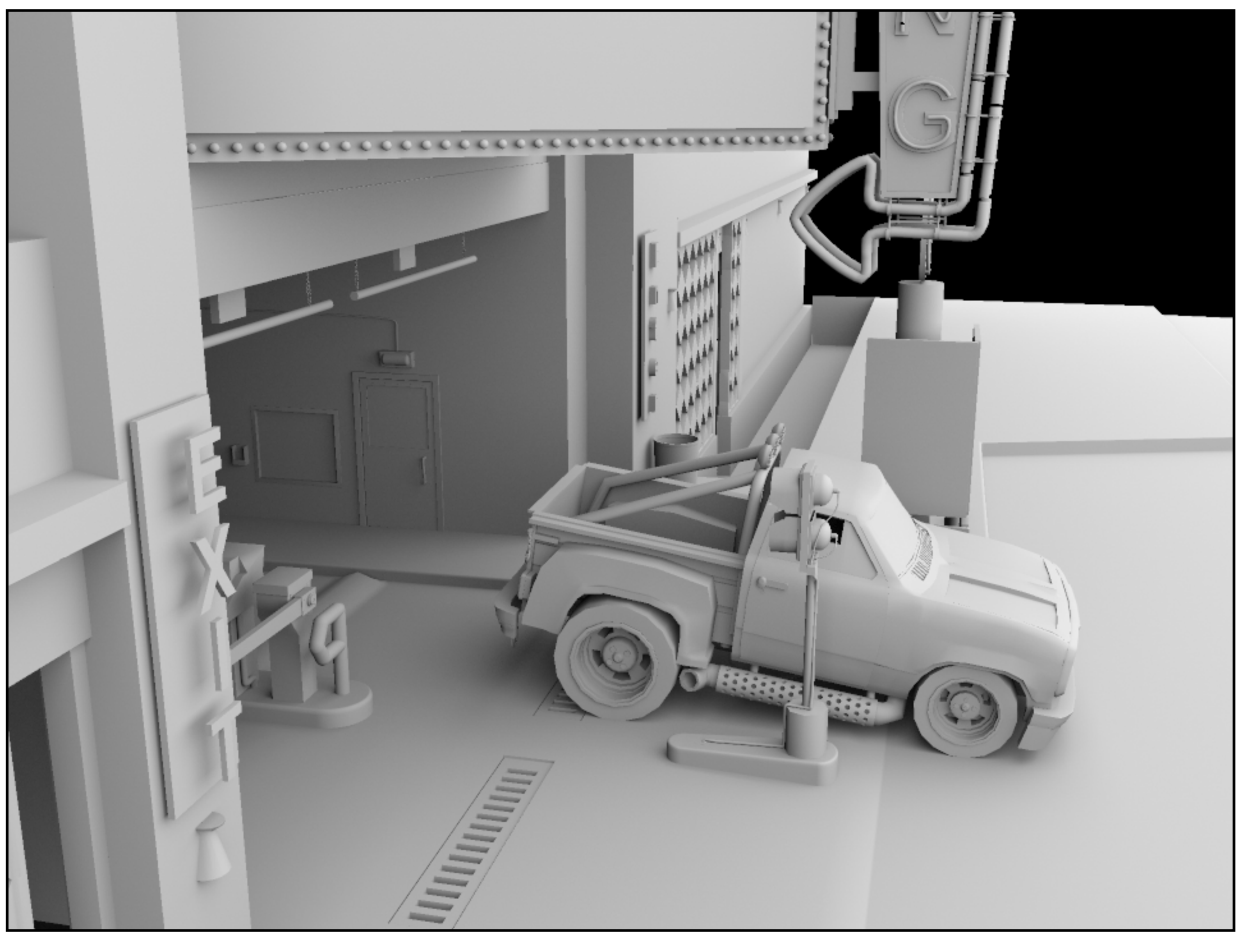

Reference (1024 rays)

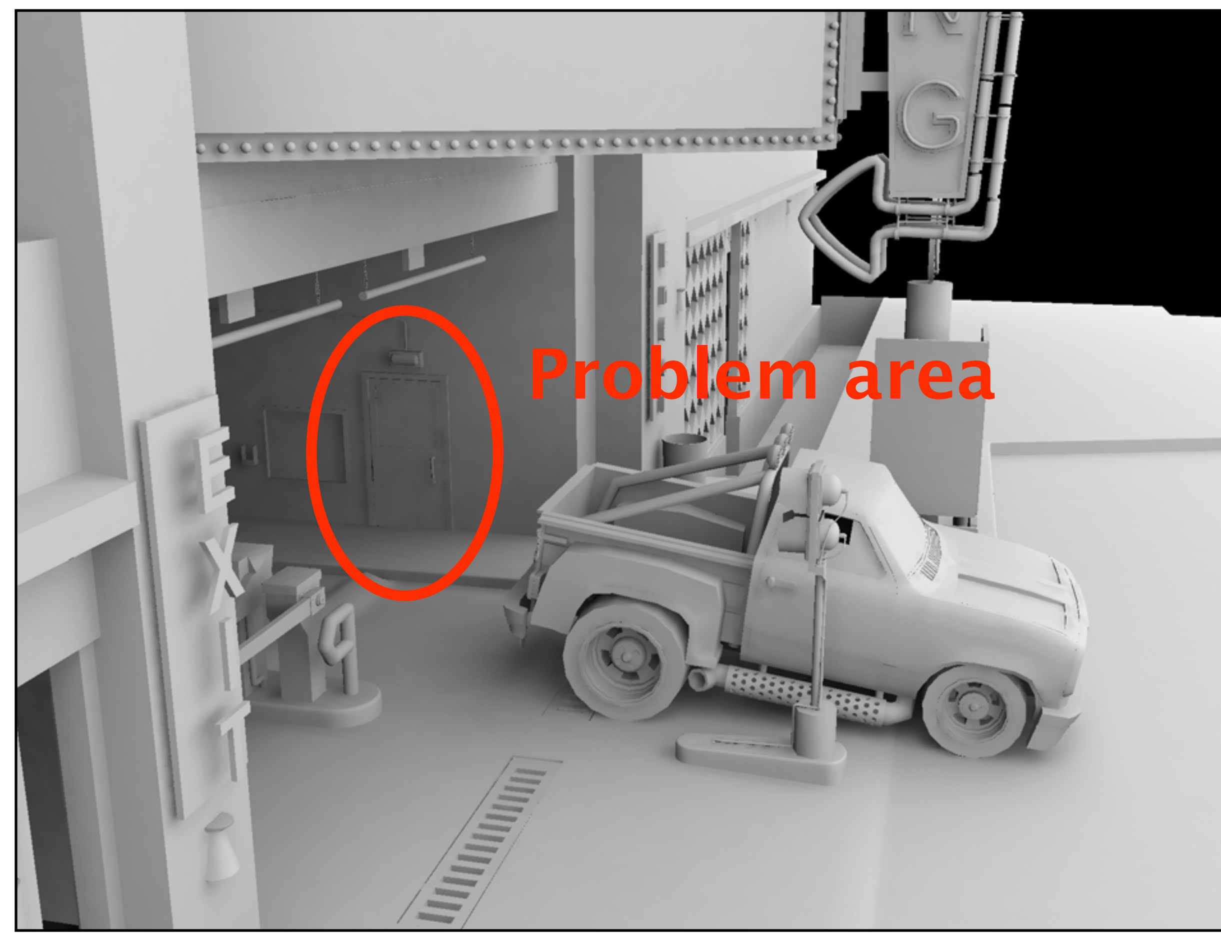

Using cache (average 17.7 rays)

> 60x speedup

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Issues

Amortized cost per pixel is very low, but the quality is not (yet) perfect

Hard to detect small geometric features:

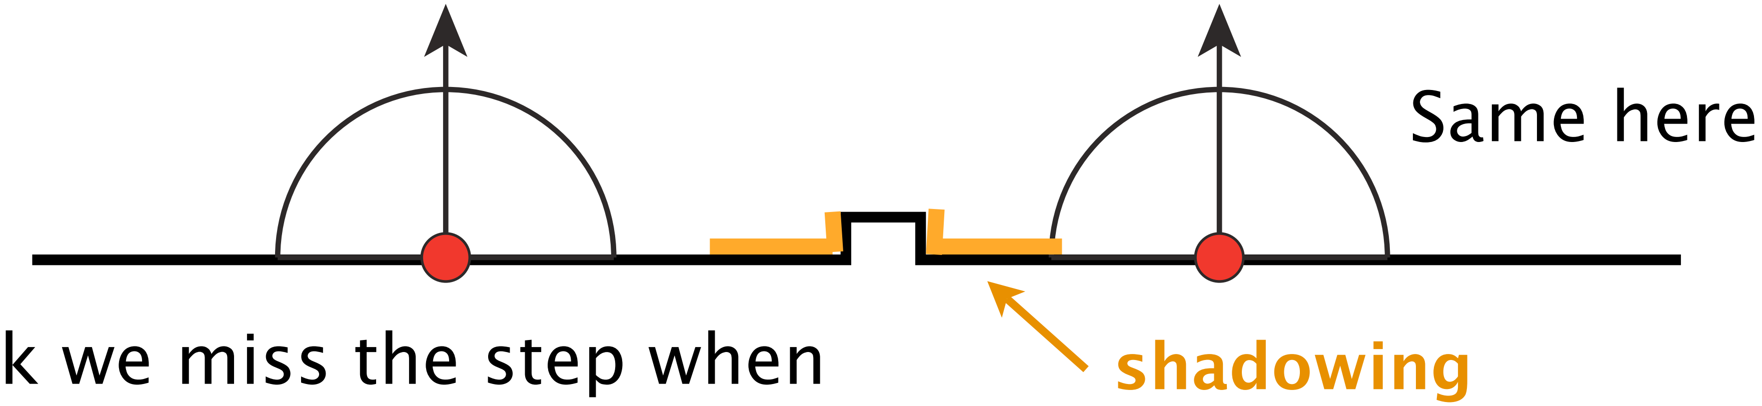

Big risk we miss the step when sampling the hemisphere

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Issues

Amortized cost per pixel is very low, but the quality is not (yet) perfect

Hard to detect small geometric features:

- We use minimum hit distance [Tabellion & Lamorlette 2004]

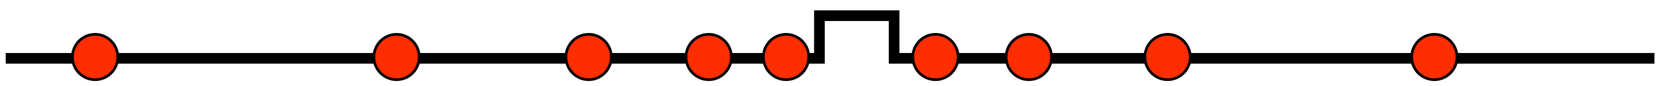

Exaggerated blurring due to filtering, some blotchiness

- We expect "visibility gradients" (~ irradiance gradients) to help

App #3 Fast Lighting Preview

Similar to Precomputed Radiance Transfer
Cost: One sparse quadtree traversal / pixel

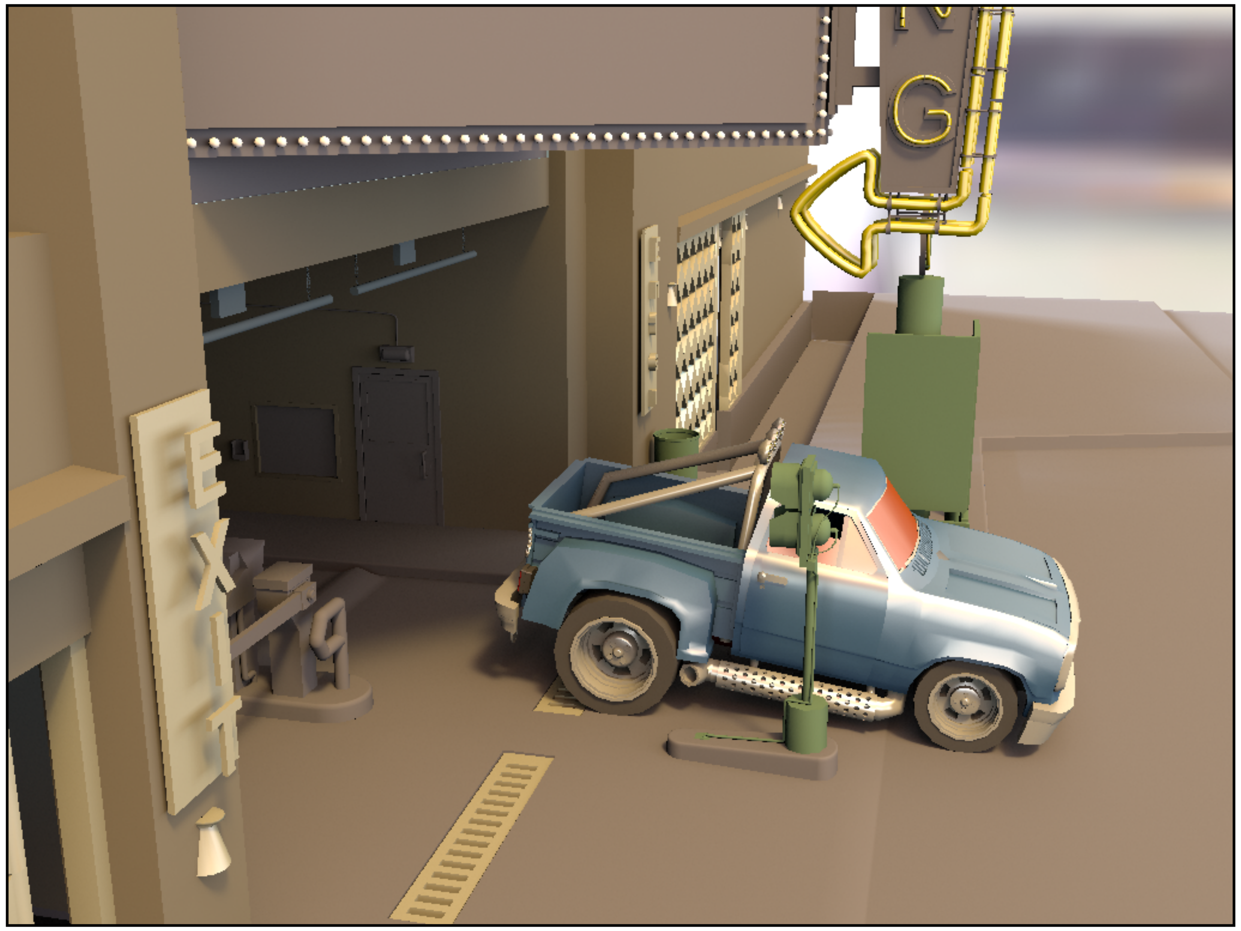

Reference

Preview using visibility cache

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Conclusion

Visibility is often strongly correlated

But, some scenes are just too complex (e.g. grass)

Visibility Caching

Store compact visibility maps sparsely in the scene

Interpolate and use for:

- Monte Carlo rendering using control variates
- Ambient occlusion
- Lighting design
- ...

Petrik Clarberg and Tomas Akenine-Möller, "*Exploiting Visibility Correlation in Direct Illumination*", Computer Graphics Forum (Proc. of EGSR 2008), vol. 27(4), pp. 1125-1136.

Thank You

Funding

Swedish Foundation for Strategic Research

Intel Corporation

Acknowledgements

3D Models:

- Christophe Desse, Matthew Thain, and Jacob Munkberg

Lund University Graphics Group (LUGG)

Anonymous reviewers!

Questions?

Please send to:

- Petrik Clarberg <petrik@cs.lth.se>