Efficient Caustic Rendering with Lightweight Photon Mapping

Pascal Grittmann[1,3] Arsène Pérard-Gayot[1] Philipp Slusallek[1,2] Jaroslav Křivánek[3,4]

[1]Saarland University [2]DFKI Saarbrücken [3]Charles University, Prague [4]Render Legion

The Idea Behind Guiding

- Importance sampling of the L_i/W_i term (path tracing / particle tracing)
- Combine with importance sampling of the BSDF
- Ideally results in perfect importance sampling of the entire Light Transport Equation (LTE)!

- **How to importance sample L_i?**
- Many approaches
- Usually store a representation of L_i at some point in the scene and interpolate them
- Methods differ in what representations they choose and how they obtain them

Reduces Variance

(plotted with low pass filter)

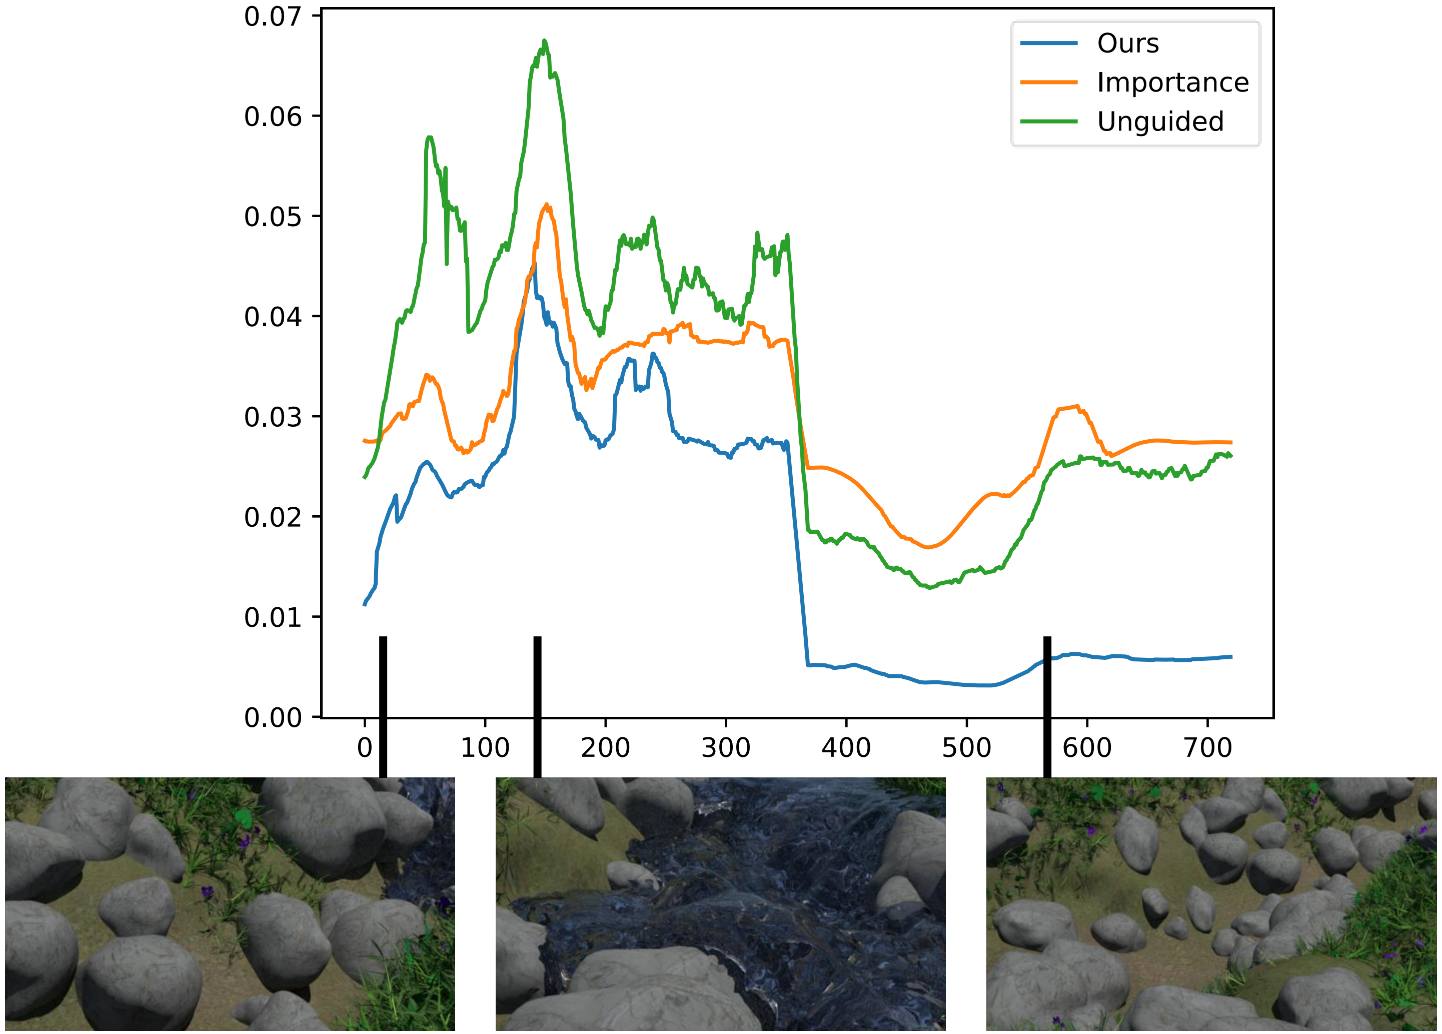

Photon Mapping Already Does Guiding [Jen96]

- Heuristic classification of materials as “glossy”
- Projection of caustic-casters
- “Caustic map”

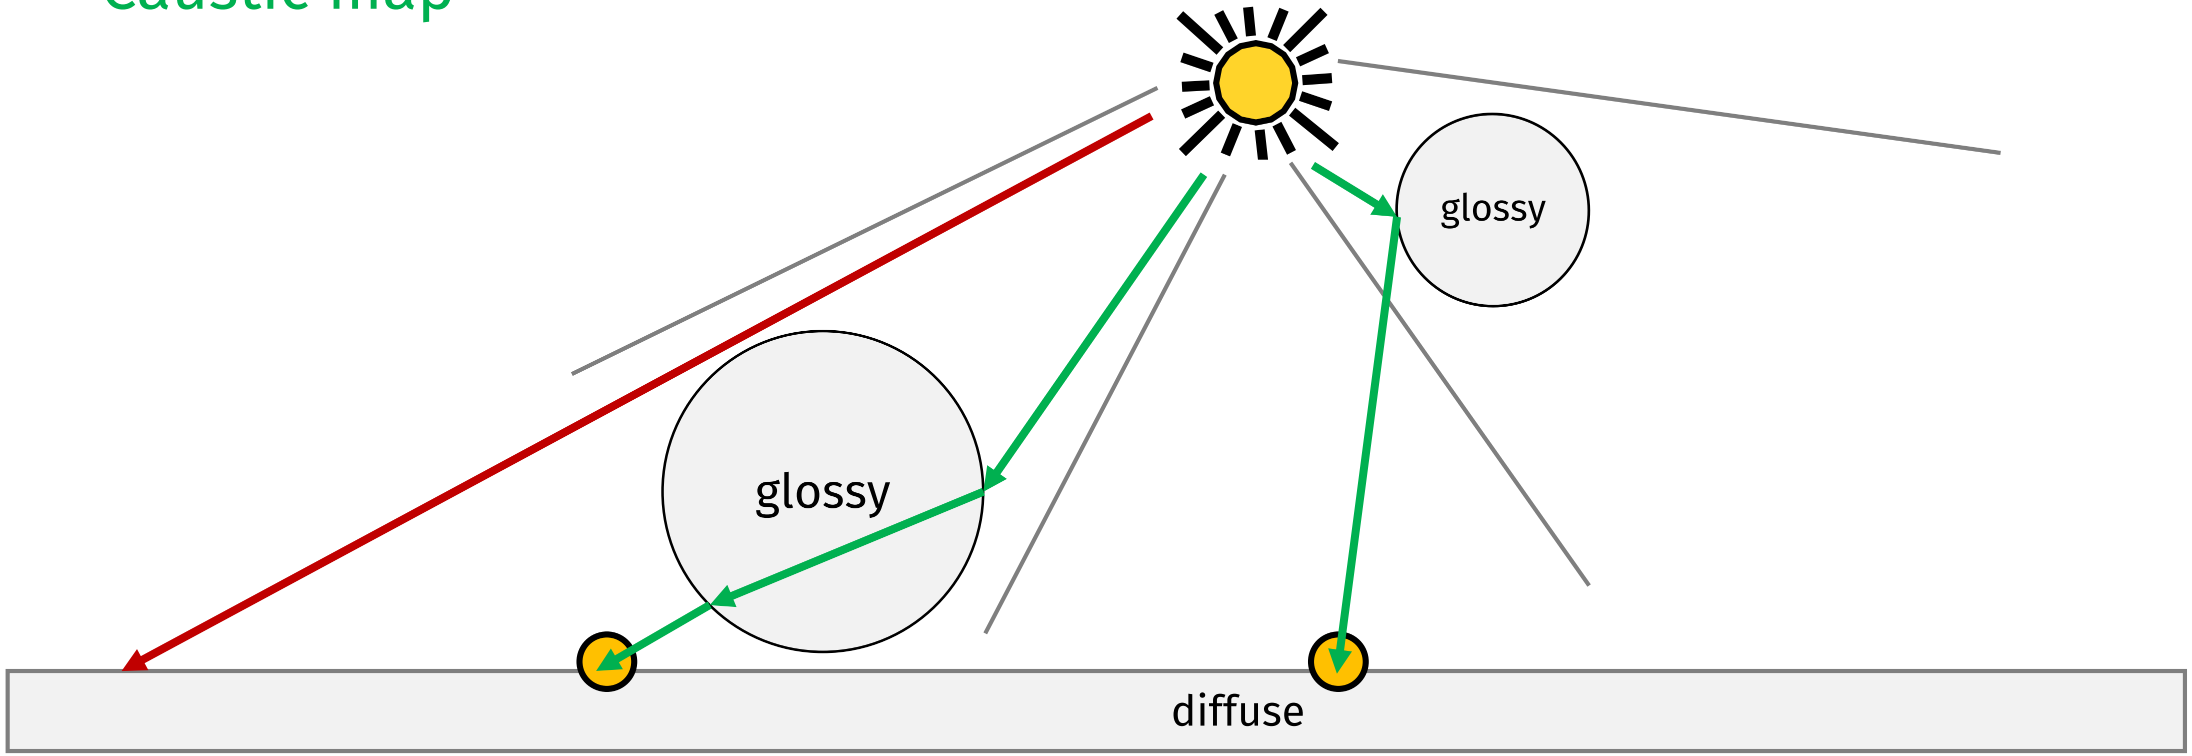

Path Guiding Using the Photon Map

- One of the first approaches to guide
- Uses nearby photons to construct a histogram of incident radiance
- Samples a cell of this histogram and a direction within the cell (uniformly)
- Histogram is a grid, each cell maps to a part of the hemisphere

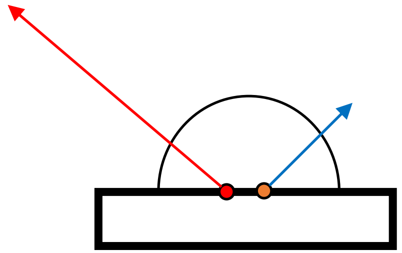

0	0	0	0
0	2	0	0
0	0	4	0
0	0	0	0

The red photon has a luminance of 2
The blue one a luminance of 4

Gaussian Mixture Models – Vorba et al. 2014

- Fits mixtures of gaussians to the incident radiance/importance at a set of points in the scene
- Project hemisphere onto plane, incident directions as bivariate Gaussians over that plane
- Gaussians are easy to sample and easy to update
- Long training pass before actual rendering (~15-30 min)

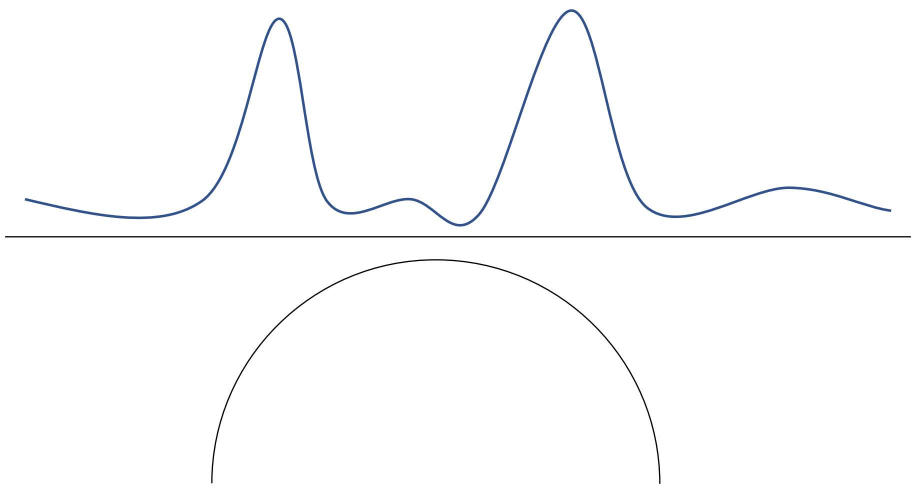

Vorba GMM – Training Phase

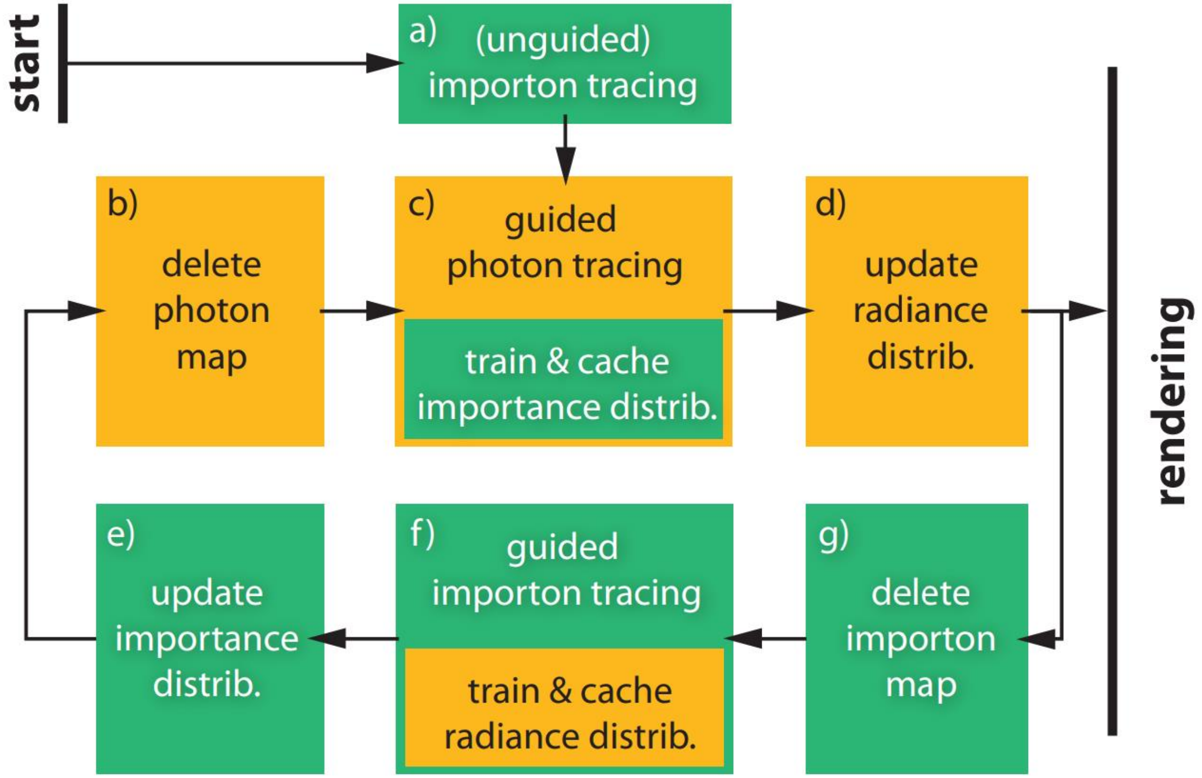

Guide Photons According to Visual Importance

- [PP98] [VKS*14] [SOHK16]
- Using importance sampling or MCMC

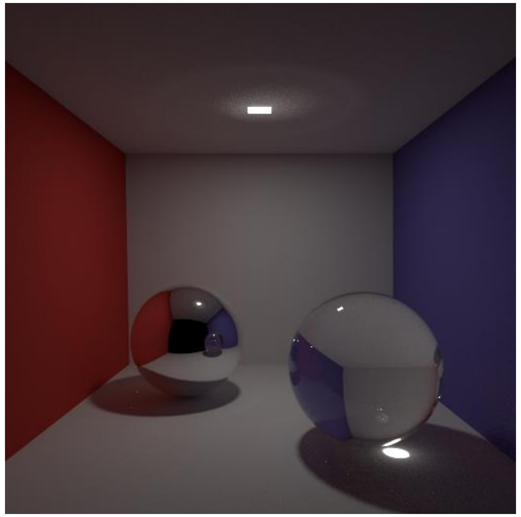

Example Scene

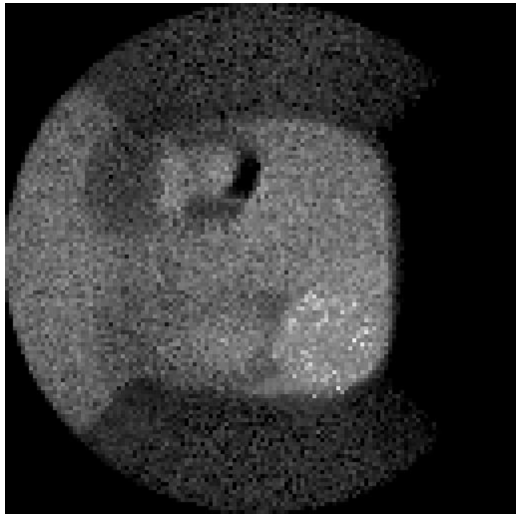

Visual importance

Our Method

- Guide emission based on visual importance
- Limit to paths with high variance form the path tracer

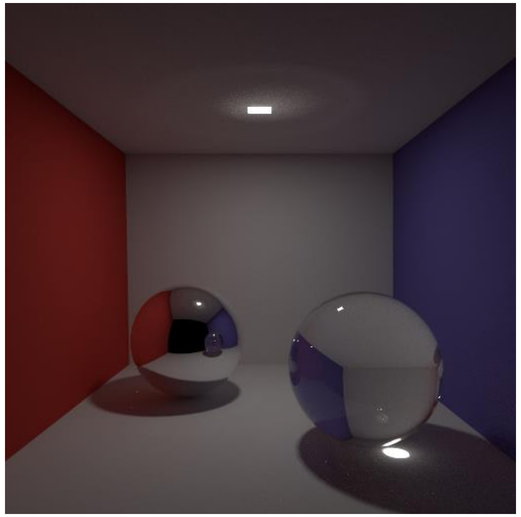

Example Scene

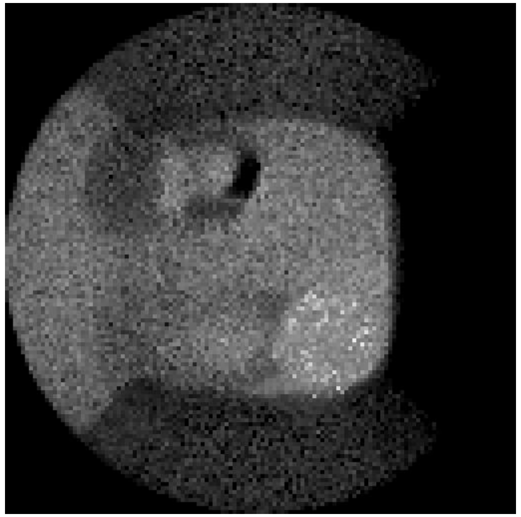

Visual importance of all photons

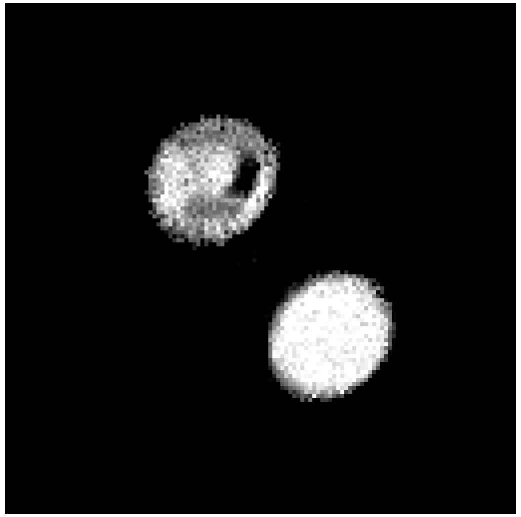

Our Method: only “useful” photons

Our Method Relies Only on Path Probabilities

- No (implicit) material classification
- Accounts for the (relative) size of the light source

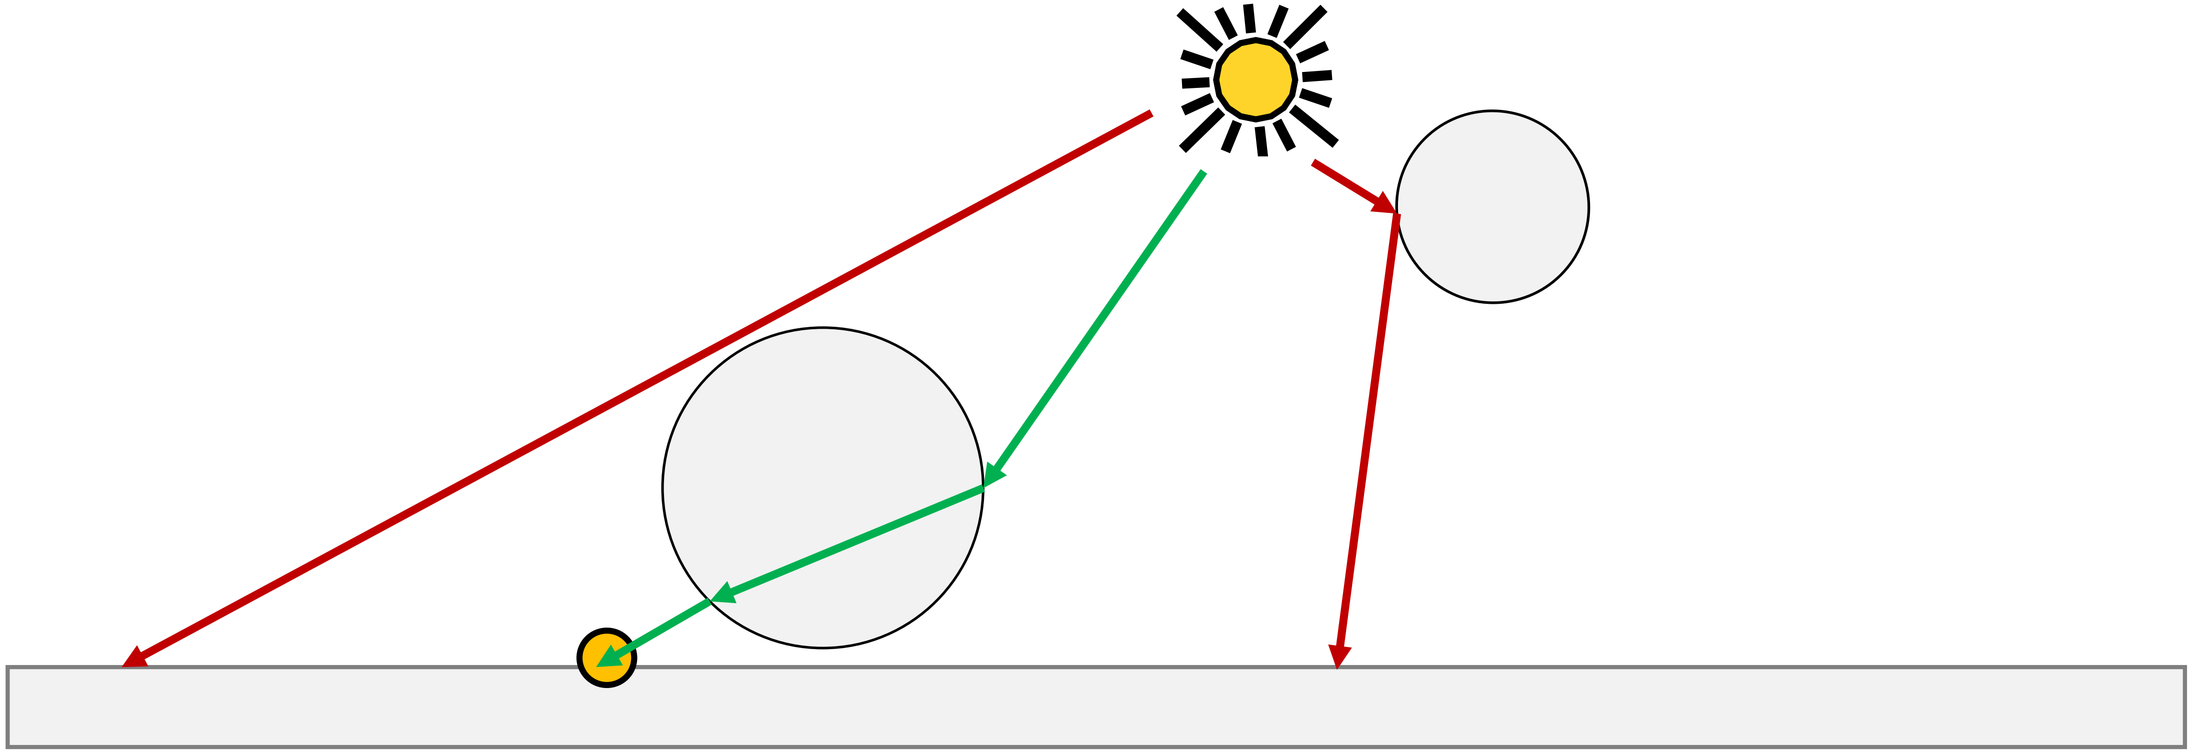

The Lightweight Photon Mapping Algorithm

- Based on VCM / UPS – [GKDS12] [HPJ12]
- Goal: More efficient solution for large scenes with a few small caustics
- MIS Combination of
  - Light Tracer
  - Photon Mapper
  - Path Tracer

Motivation / Idea

- Existing methods: Try to be unbiased for all estimators
  - Looses main advantage of MIS!
- Why not ignore estimators that we know will contribute little?
  - A la maximum heuristics or alpha-max heuristics – but only where necessary

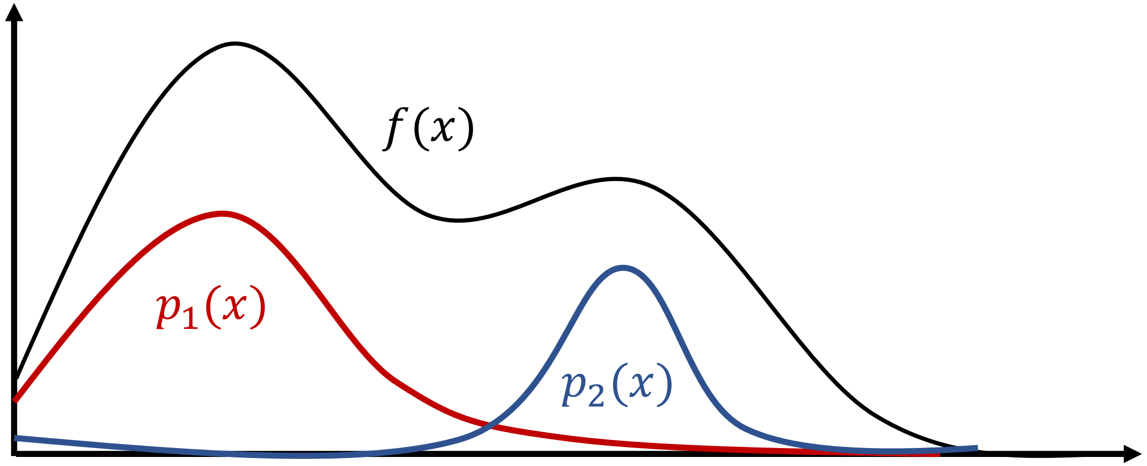

- Can restricting costly estimators to regions of high variance result in more efficient combined algorithms?

The Notion of “Useful” Photons

$$\frac{N_{min}\ p_{PM}(y)\ \pi r^2}{p_{PT}(y|y_k)} > 1$$

“The **photon mapper** can reach a point within r with **higher probability than the path tracer**, using only N_{min} light paths”

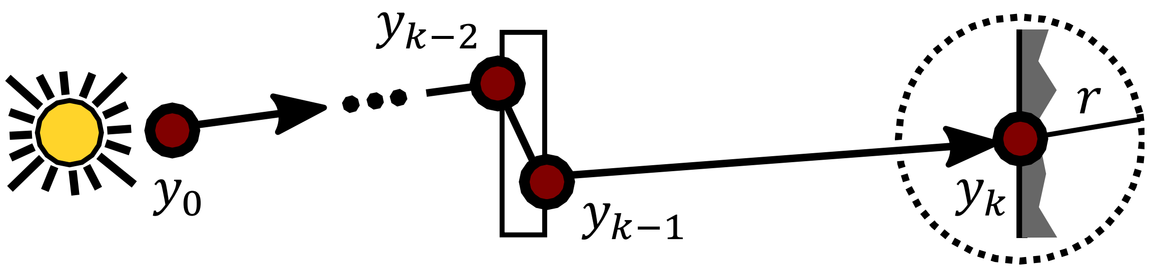

How Many Photons Should We Trace?
- One Per Pixel Influenced by Caustics

- VCM: One light path per pixel
- With guiding: Fewer light paths are needed!

$$I = I_{PM} + I_{LT} + I_{PT}$$

Rendered Image

$$I_{PM} + I_{LT}$$

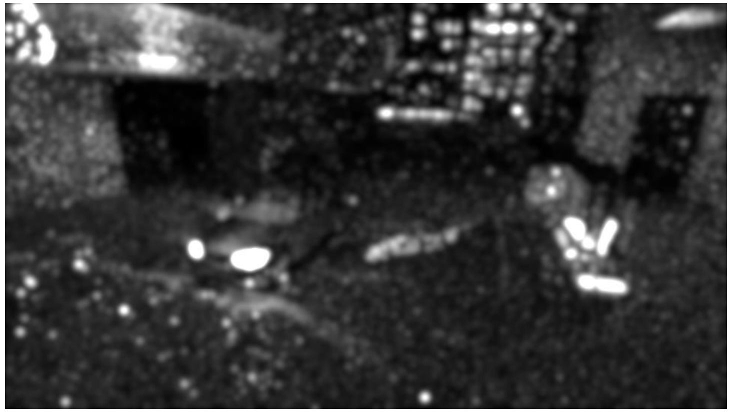

PM / LT Contribution
(exposure +5)

$$\frac{I_{PM} + I_{LT}}{I_{PM} + I_{LT} + I_{PT}} > 1\%$$

Pixel Classification

Is that Number of Light Paths Optimal?

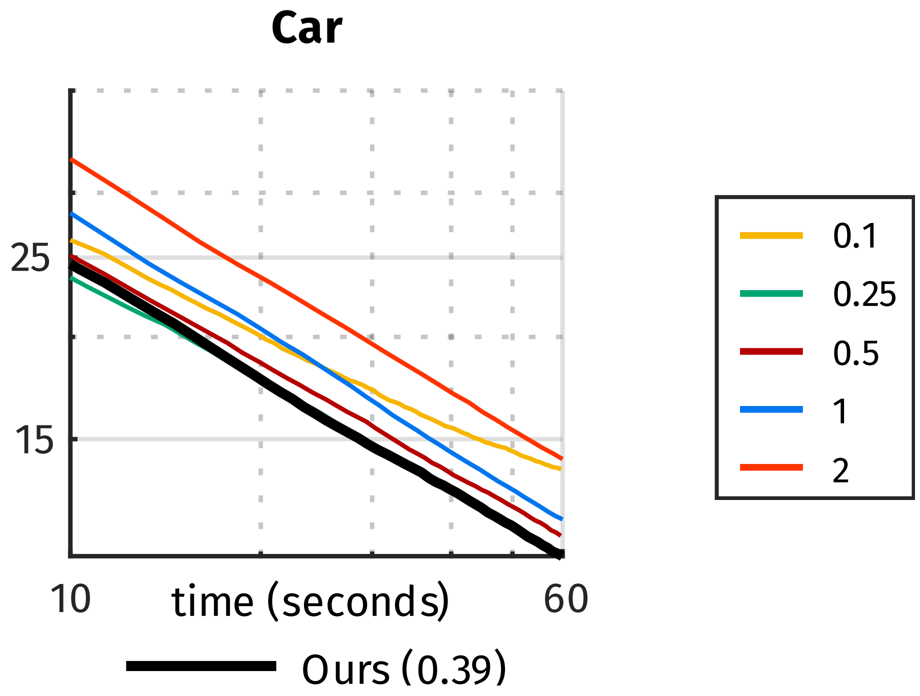

→ Optimal for large scenes with small Caustics

Is that Number of Light Paths Optimal?

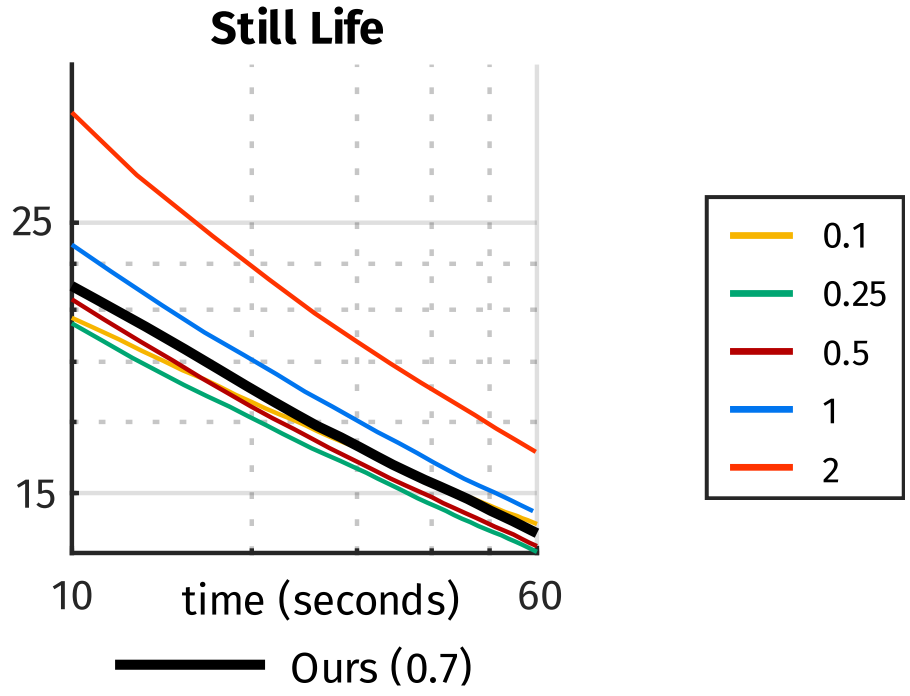

→ Complex SDS paths require more samples from the path tracer

Is that Number of Light Paths Optimal?

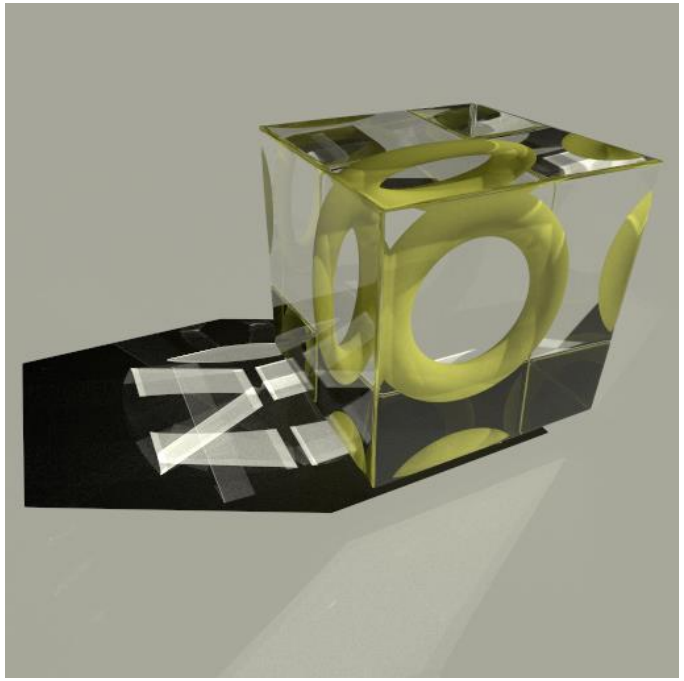

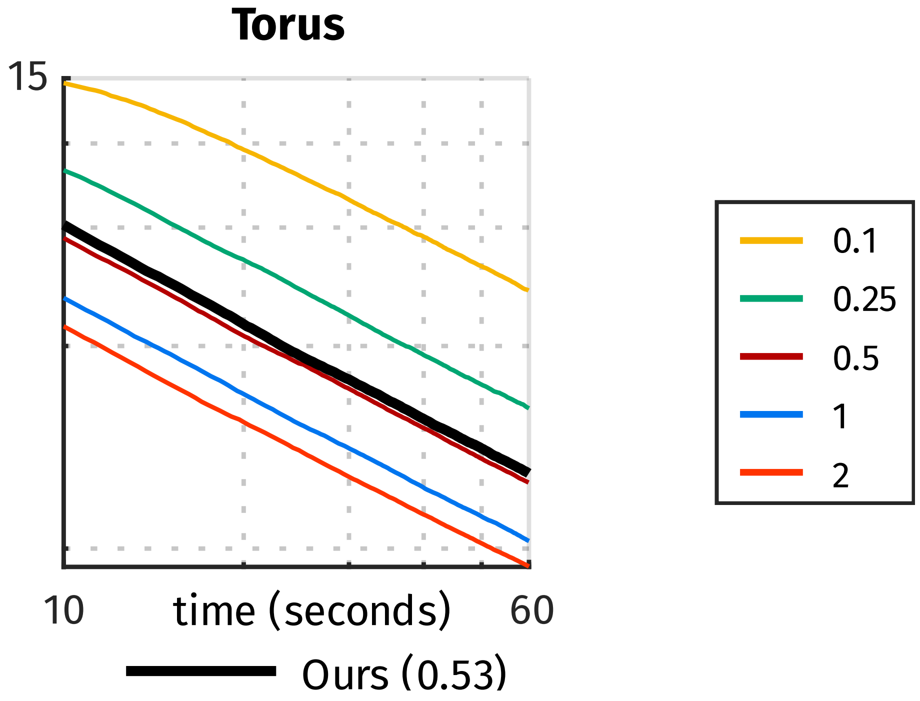

→ For scenes that are trivial except for the caustics, a higher number would be more efficient

Results

Impact of the Full Method with Our Test Scenes

Photon Densities in the Cornell Box Variations

Reference

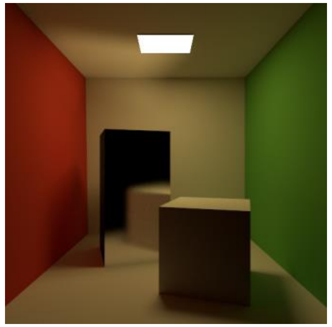

Photon density – Guiding with all Photons

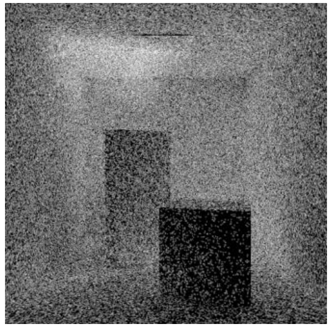

Photon density – Our

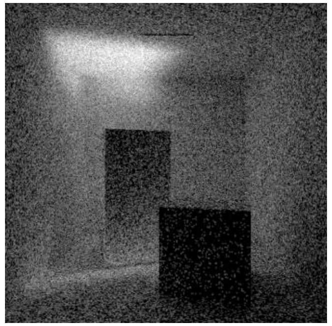

Photon Densities in the Cornell Box Variations

Reference

Photon density – Guiding with all Photons

Photon density – Our

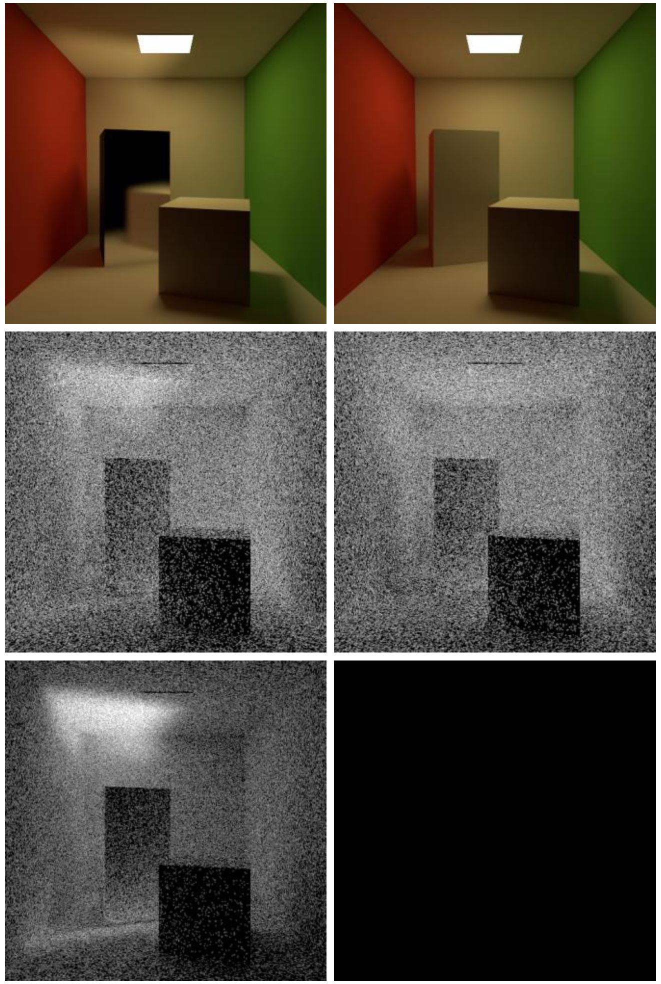

Photon Densities in the Cornell Box Variations

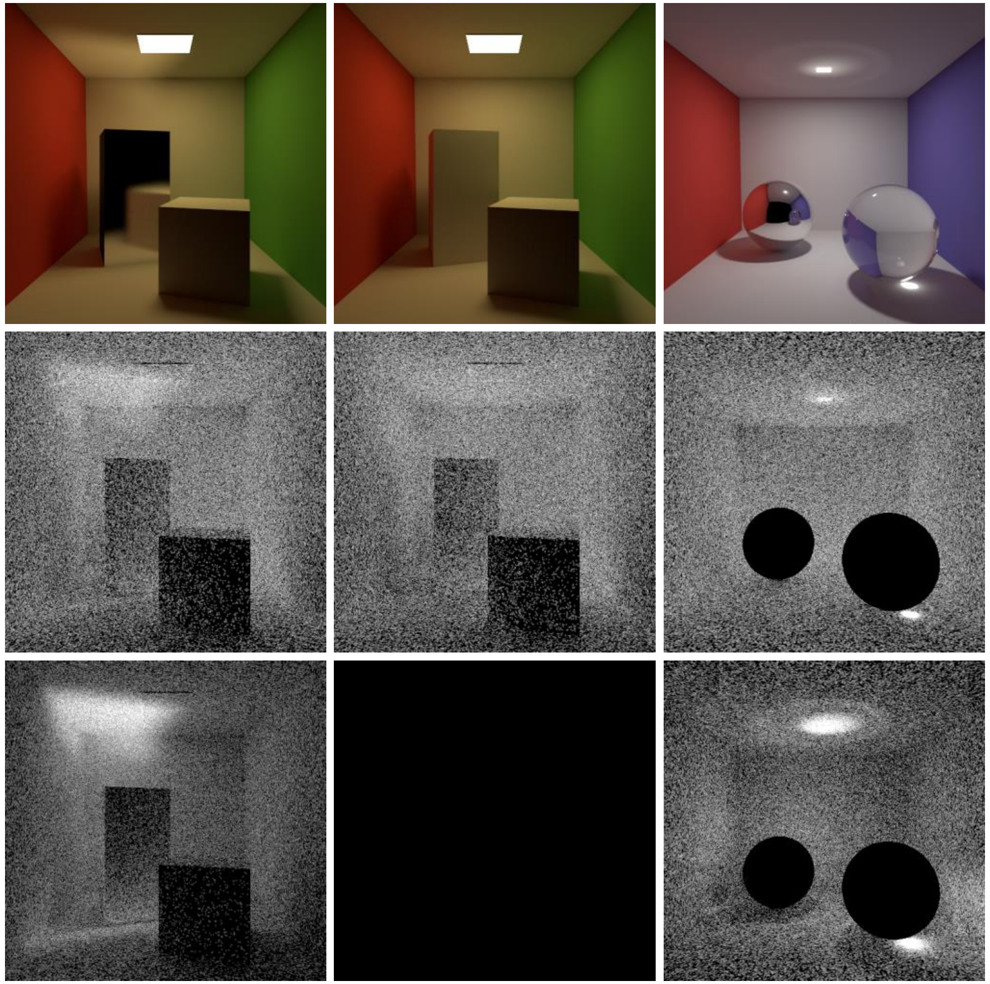

Photon Densities in the Cornell Box Variations

Reference

Photon density – Guiding with all Photons

Photon density – Our

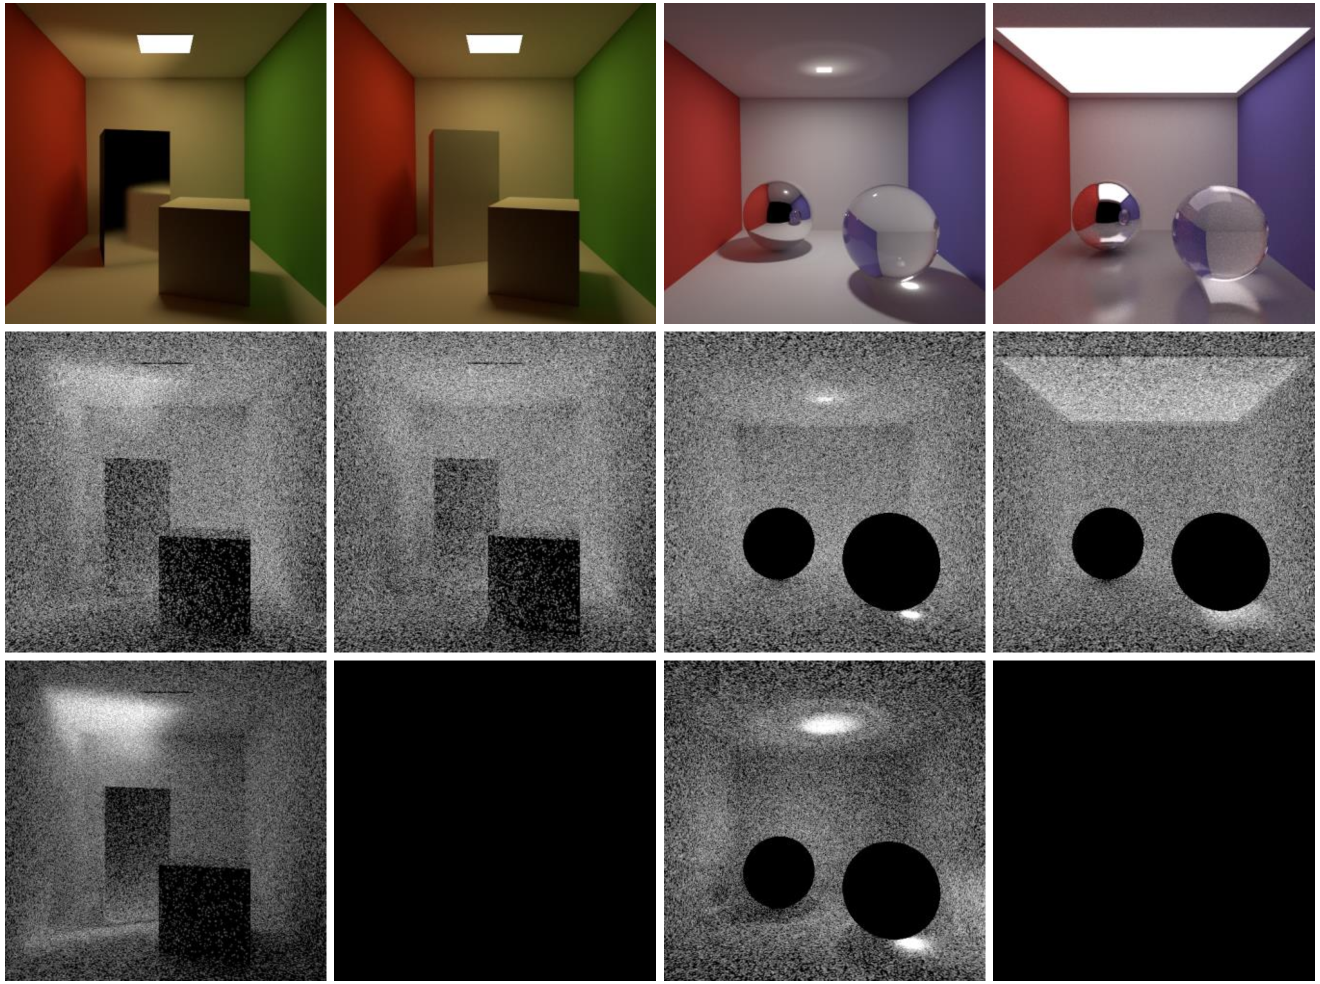

The Torus – Simple Example, Directional Light

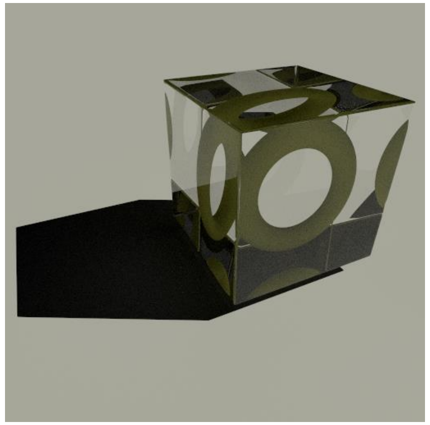

Path tracer
(delta light)

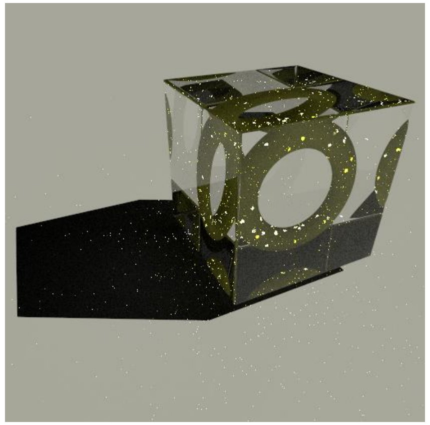

Unguided

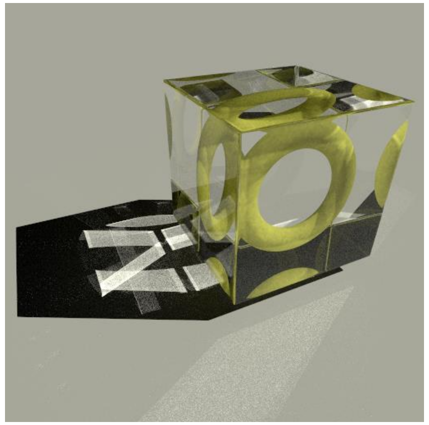

Our

Result identical to existing guiding approaches.

Car Scene – Large Exterior Scene, Small Caustics

Equal-time comparison (60 seconds)

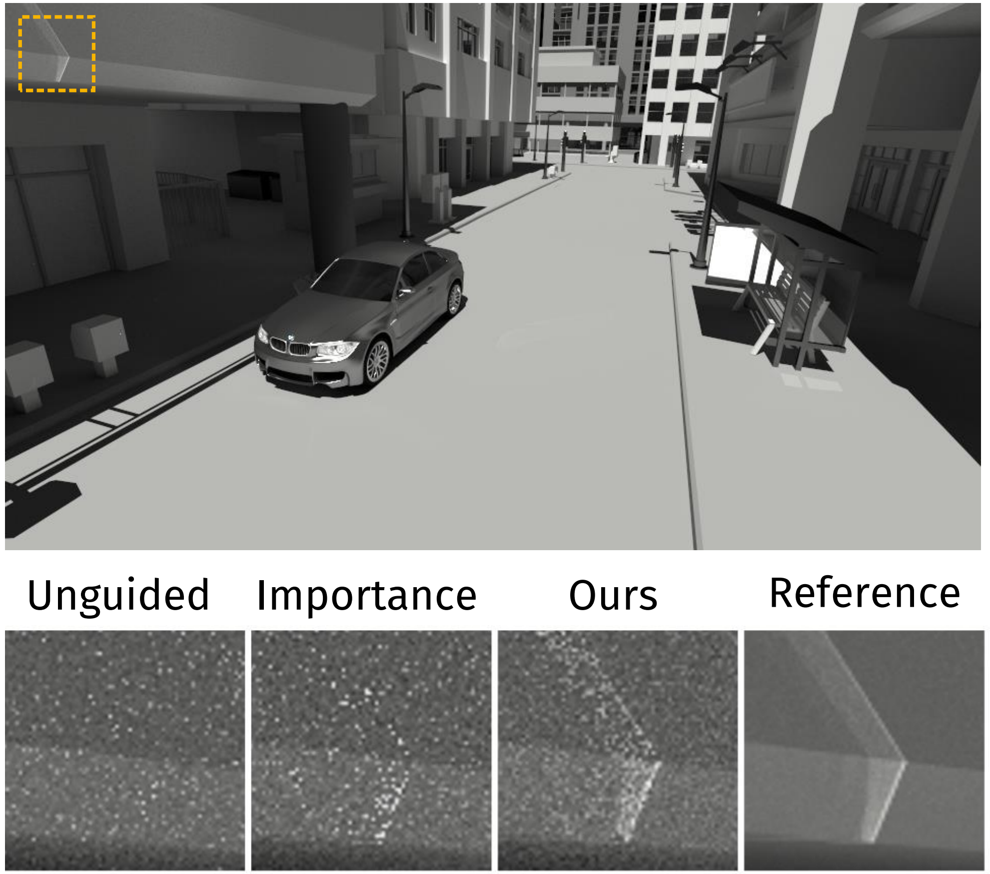

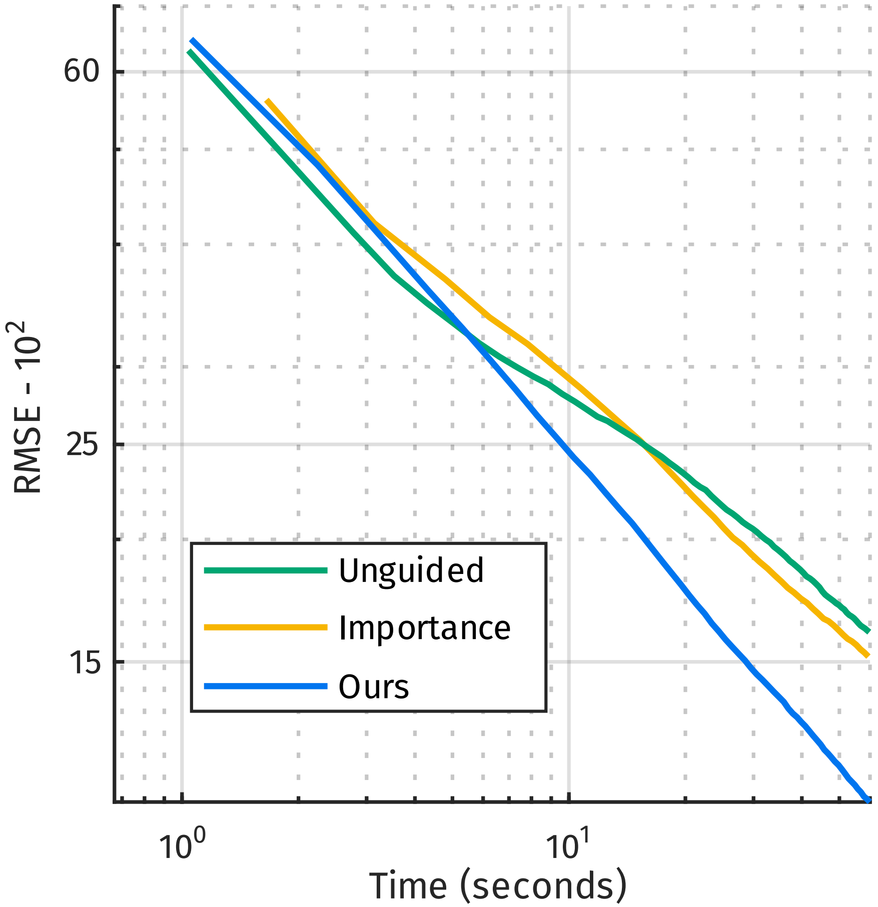

Car Scene – Large Exterior Scene, Small Caustics

Equal-time comparison (60 seconds)

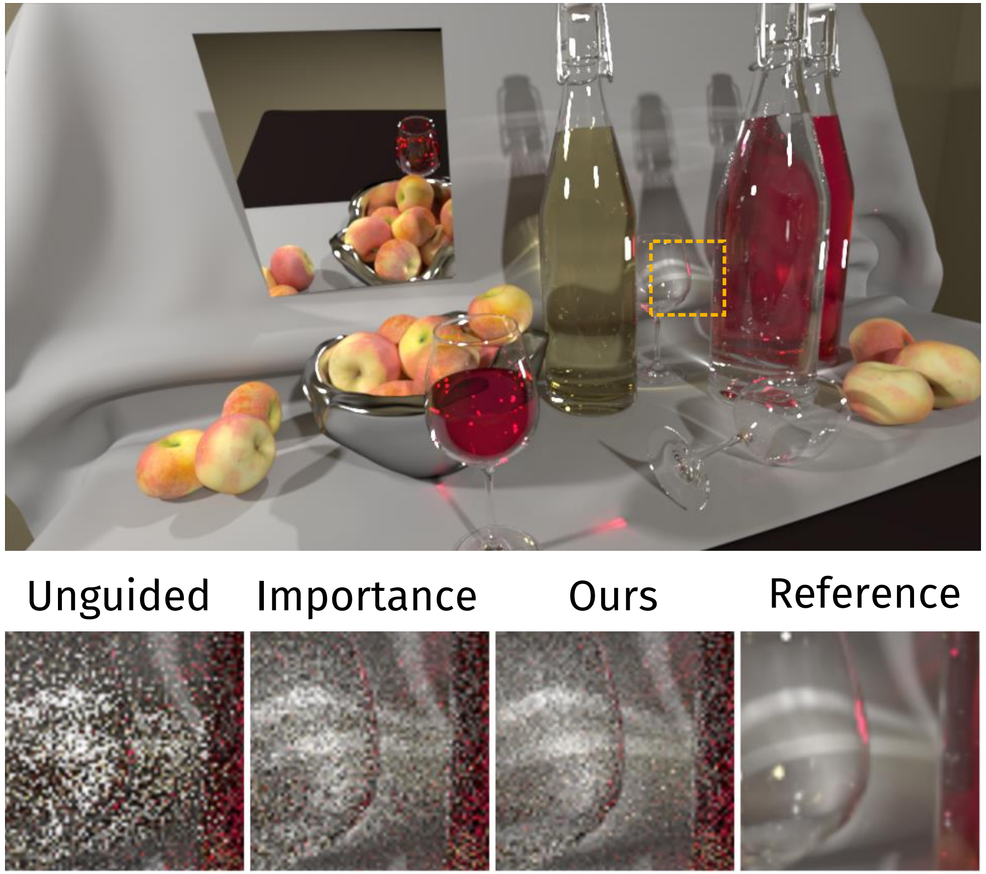

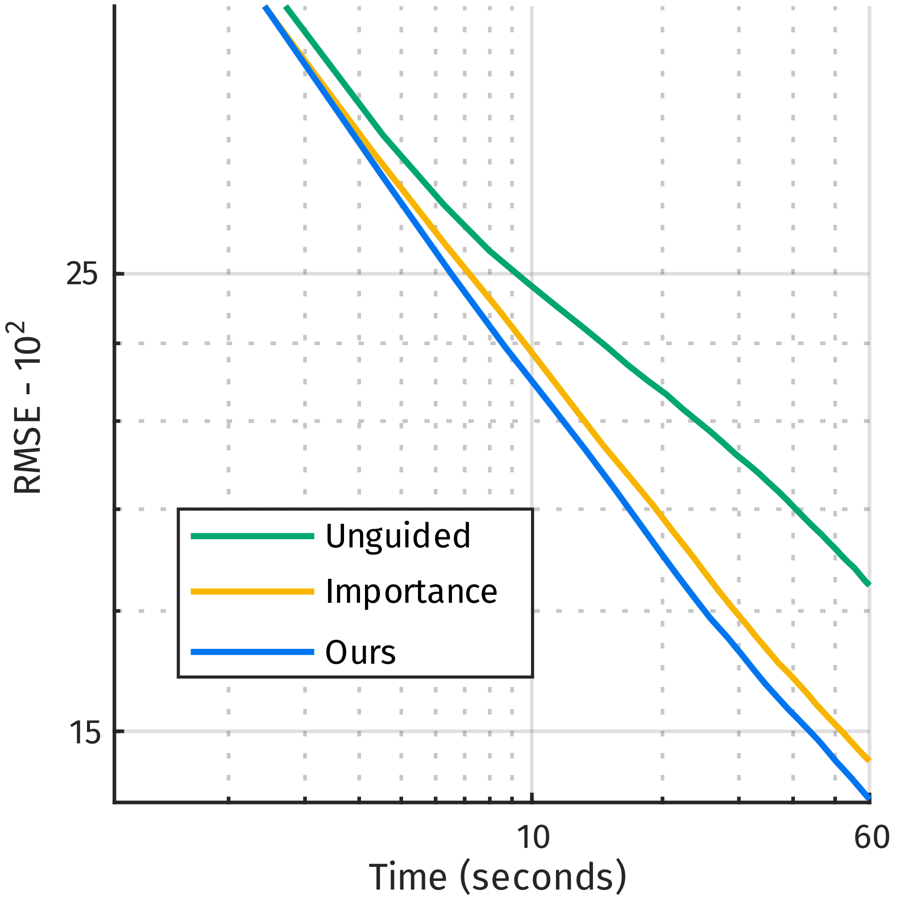

Algorithm:	Path Tracer	VM	VM+EG	Ours	Reference
RMSE:	1688.18	1608.44	1518.65	1079.11	-
Photons per Iteration (Average):	-	491,133	1,917,898	932,028	-
Light Paths per Iteration:	-	518,400	515,763	203,976	-

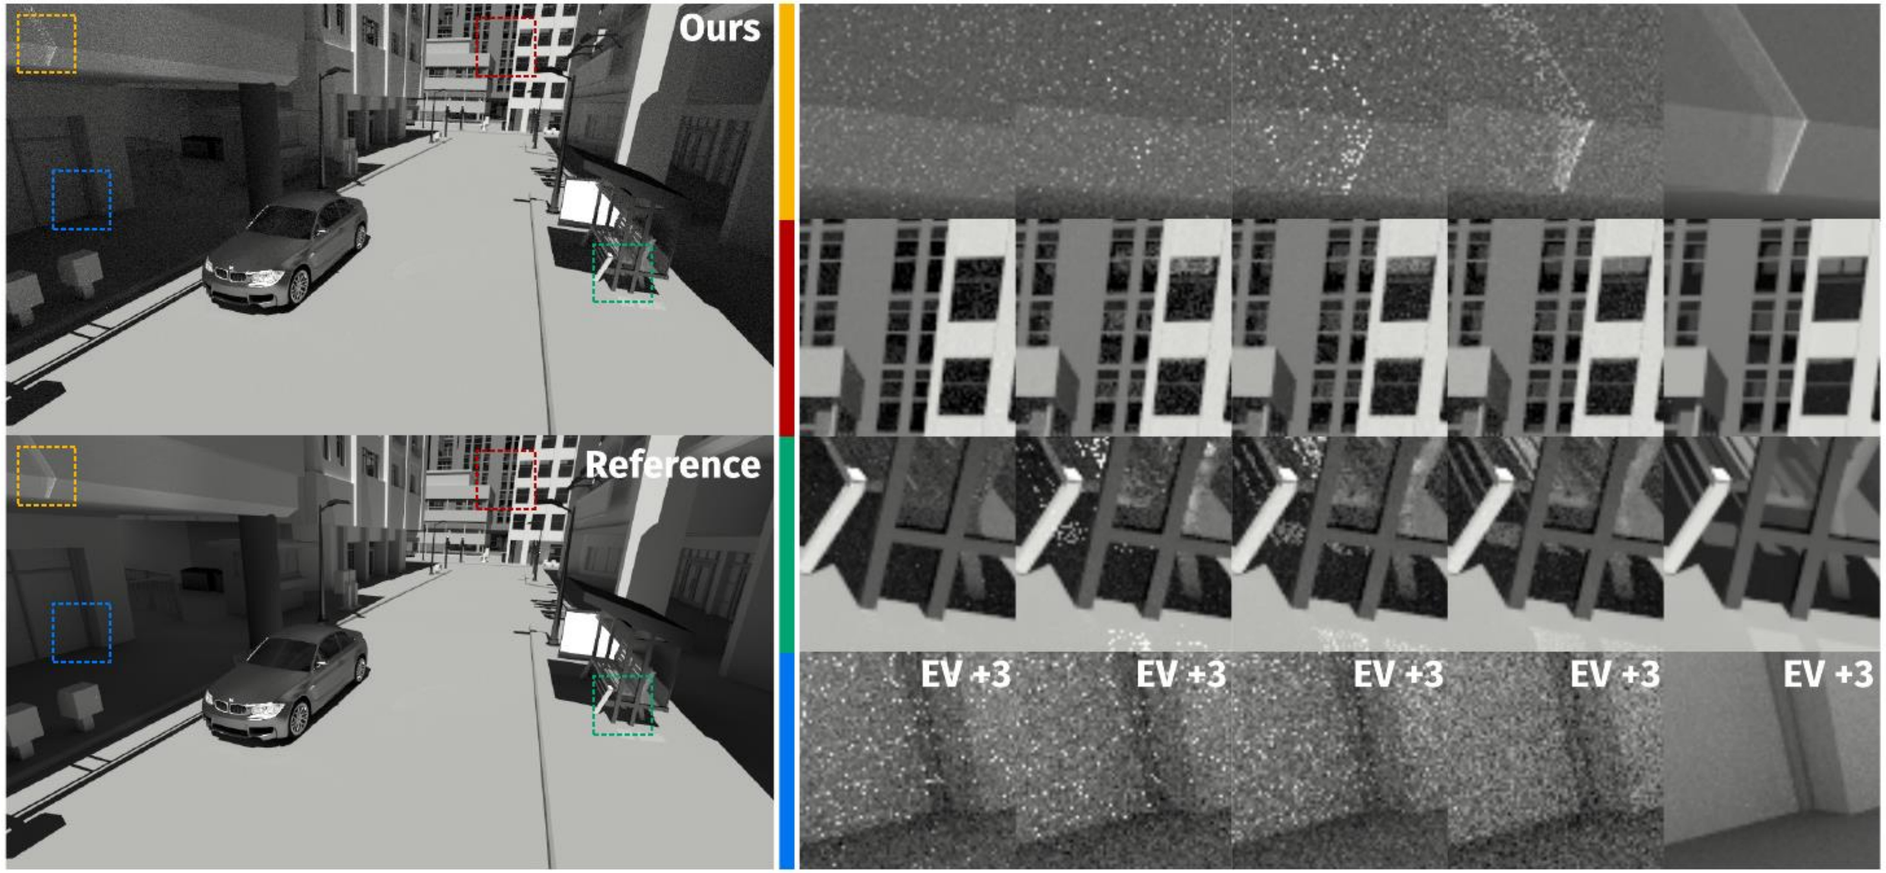

Figure 8: *Equal-time comparison after one minute of rendering for the* CAR *scene. In this scene, the path tracer cannot sample some of the caustics at all (directional light source and perfect specularity). Our method results in half the number of light paths getting traced and therefore also a significantly lower number of photons. The contribution-based guiding (VM+EG) results in even more photons than the uniform emission, because fewer light paths miss the scene entirely.*

Algorithm:	Path Tracer	VM	VM+EG	Ours	Reference
RMSE:	5397.38	6219.95	4765.8	4488.86	-
Photons per Iteration (Average):	-	2,408,366	2,418,036	1,337,252	-
Light Paths per Iteration:	-	518,400	498,460	362,475	-

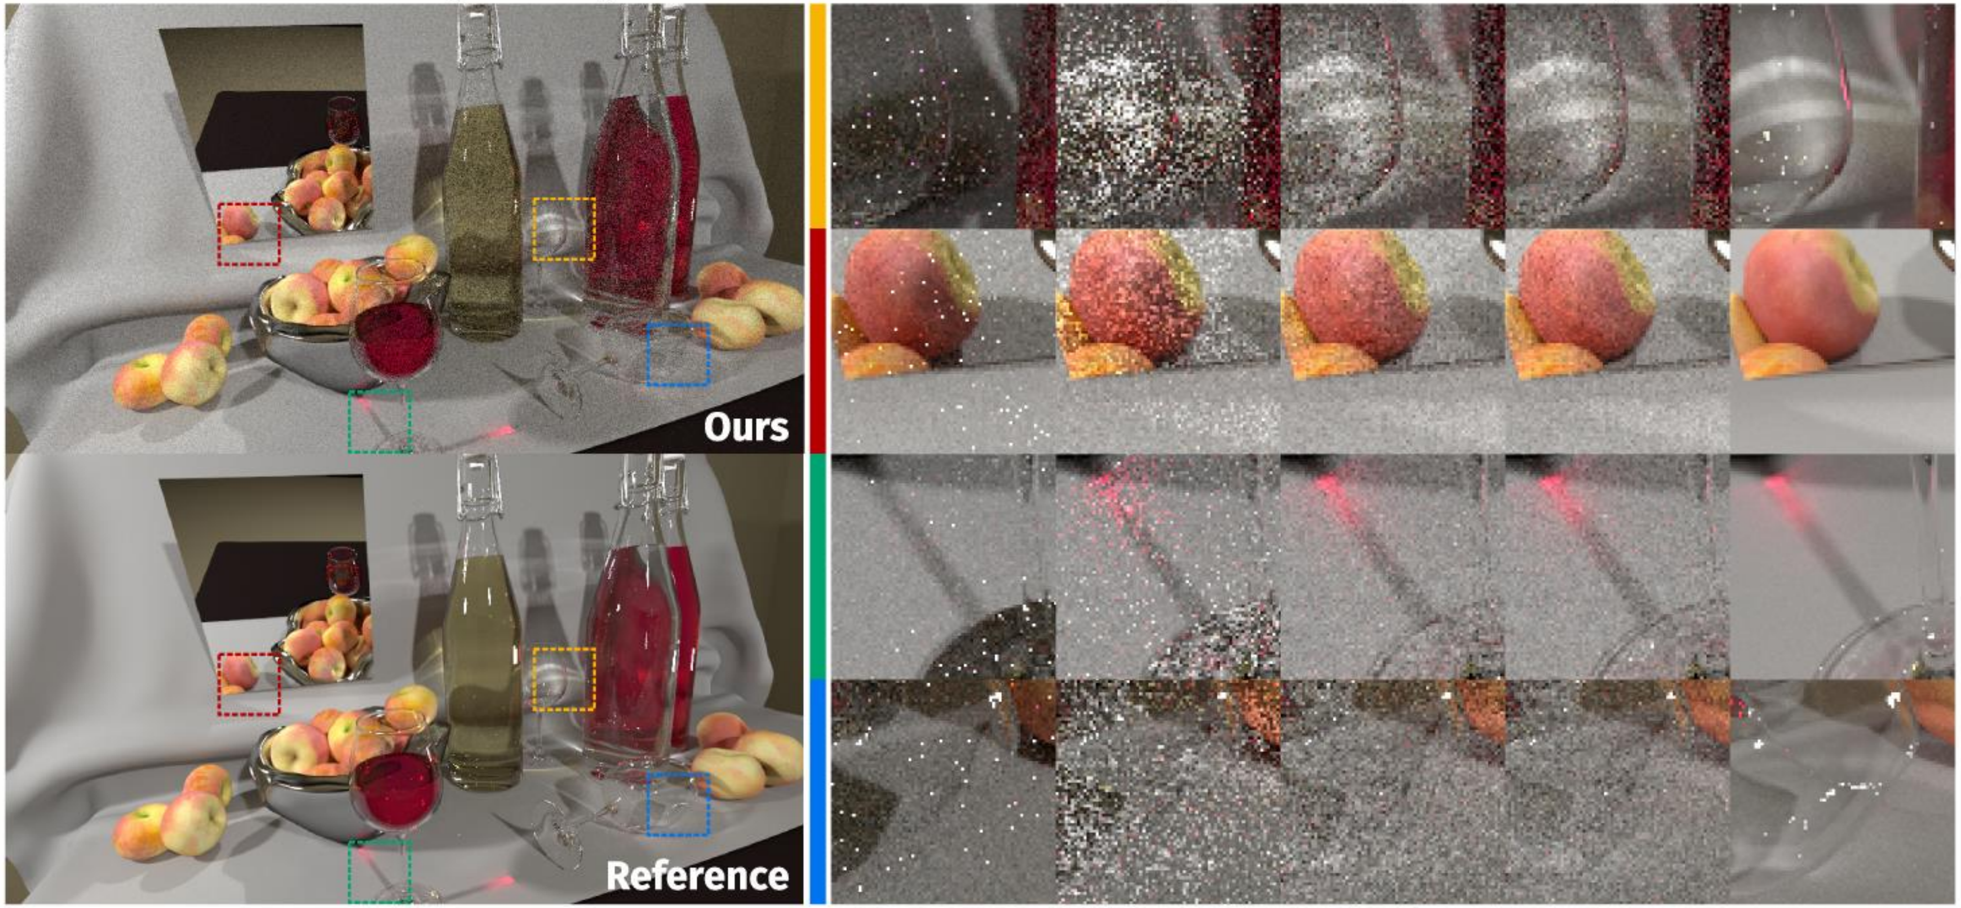

Figure 10: *Equal-time comparison after one minute of rendering for the* STILL LIFE *scene. Here, the difference between our method and contribution-based guiding (VM+EG) is slightly less visible than in the* CAR *scene (Fig. 8). This is because the majority of the image is influenced by caustics.*

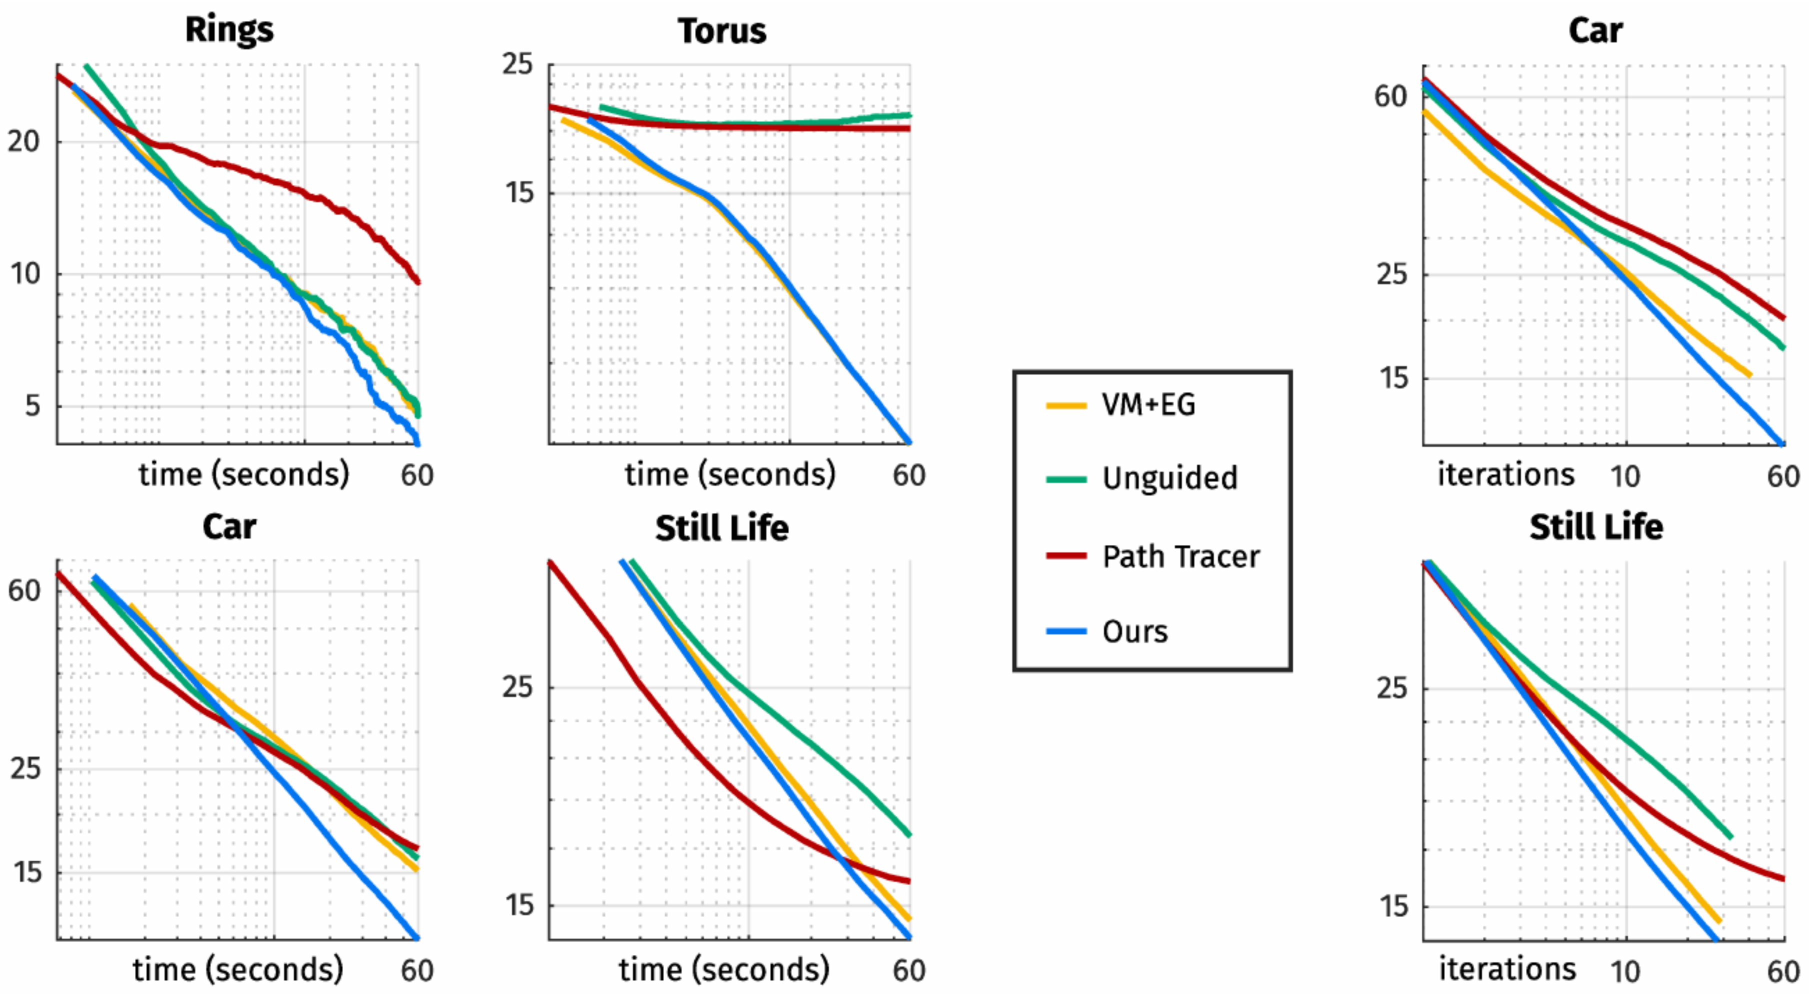

Figure 11: *Equal-time (one minute) and equal-iteration-count convergence rate (log-RMSE) for some of our test scenes. Our method has either better or identical convergence rates in all our test scenes, even on a per-iteration level.*

Limitations

- Only for caustic-casters directly in front of the light source
- Resorts to path tracing for (diffuse) indirect illumination

Efficient Caustic Rendering with Lightweight Photon Mapping

Pascal Grittmann   Arsène Pérard-Gayot   Philipp Slusallek   Jaroslav Křivánek

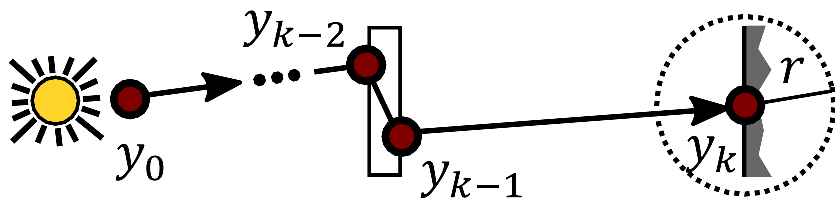

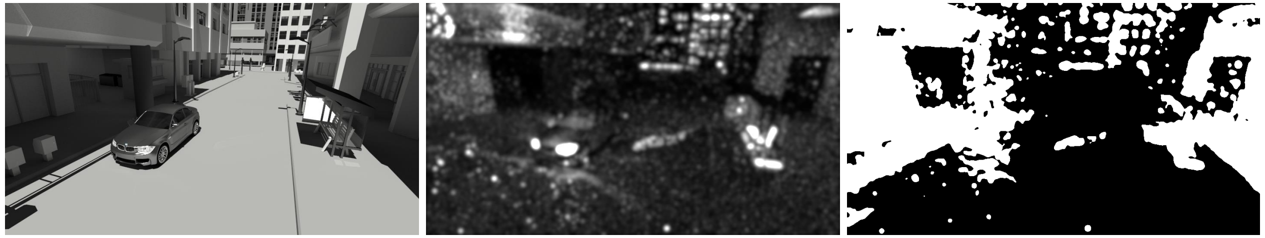

Reference   PM / LT contribution   Our pixel classification

Restrict costly estimators to a subset of the domain

→ More efficient MIS combination

Unguided   Importance   Ours   Reference

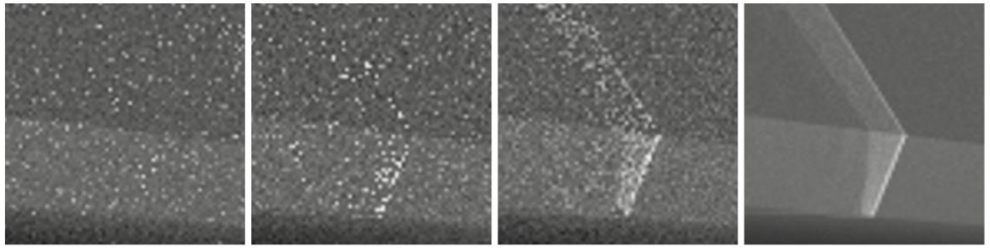

Reference

Importance driven

Our