

Title:

Efficient Caustic Rendering with Lightweight Photon Mapping

Authors:

Pascal Grittmann, Arsène Pérard-Gayot, Philipp Slusallek, Jaroslav Křivánek

This is the peer reviewed version of the following article:

Grittmann, P., Pérard-Gayot, A., Slusallek, P. and Křivánek, J. (2018), Efficient Caustic Rendering with Lightweight Photon Mapping. *Computer Graphics Forum*, 37: 133-142. doi:10.1111/cgf.13481,

which has been published in final form at <https://doi.org/10.1111/cgf.13481>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self- Archived Versions.



Efficient Caustic Rendering with Lightweight Photon Mapping

Pascal Grittmann^{1,3} Arsène Pérard-Gayot¹ Philipp Slusallek^{1,2} Jaroslav Krivánek^{3,4}

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

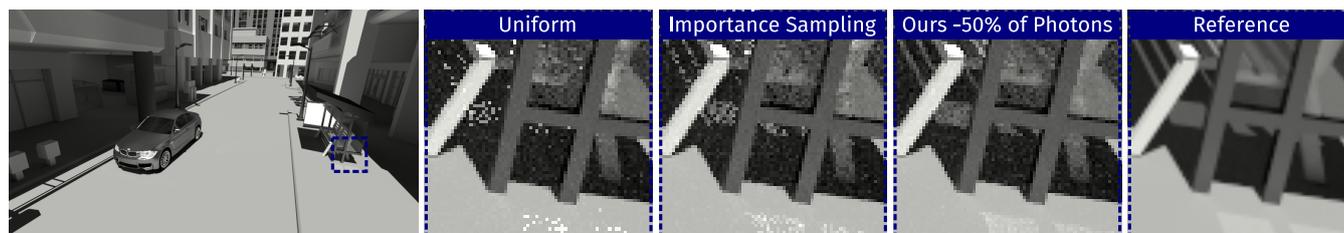


Figure 1: Different photon emission strategies in the CAR scene. We achieve a significantly better photon density inside caustics with fewer photons. The emission directions and the number of light paths (and thus photons) are optimized automatically by our method.

Abstract

Robust and efficient rendering of complex lighting effects, such as caustics, remains a challenging task. While algorithms like vertex connection and merging can render such effects robustly, their significant overhead over a simple path tracer is not always justified and – as we show in this paper – also not necessary. In current rendering solutions, caustics often require the user to enable a specialized algorithm, usually a photon mapper, and hand-tune its parameters. But even with carefully chosen parameters, photon mapping may still trace many photons that the path tracer could sample well enough, or, even worse, that are not visible at all.

Our goal is robust, yet lightweight, caustics rendering. To that end, we propose a technique to identify and focus computation on the photon paths that offer significant variance reduction over samples from a path tracer. We apply this technique in a rendering solution combining path tracing and photon mapping. The photon emission is automatically guided towards regions where the photons are useful, i.e., provide substantial variance reduction for the currently rendered image. Our method achieves better photon densities with fewer light paths (and thus photons) than emission guiding approaches based on visual importance. In addition, we automatically determine an appropriate number of photons for a given scene, and the algorithm gracefully degenerates to pure path tracing for scenes that do not benefit from photon mapping.

CCS Concepts

•Computing methodologies → Ray tracing;

1. Introduction

Lighting in real-world scenes produces a broad range of visual effects. Direct and smooth indirect illumination, for instance, are well resolved by the path tracing algorithm, following light transport paths from the camera. Strong, concentrated indirect illumination, and especially caustics, can be rendered more efficiently by tracing paths from the light sources. In other words, for every lighting effect, there is a path sampling technique that can sample it significantly better – with lower variance – than other techniques.

Different scenes feature different effects, so no single sampling technique is always better than others. Therefore, bidirectional algorithms, combining multiple sampling techniques via multiple importance sampling (MIS) [VG95b], have been introduced [Vea97, GKDS12, HJ12, KGH*14]. Such algorithms are *robust* in the sense that their performance does not drastically drop

due to the presence of a specific lighting effect, but this robustness comes at the cost of generally unsatisfactory performance: Most scenes do not feature all possible lighting effects everywhere in the rendered image, and the additional overhead of combining all the sampling techniques – that are included ‘just in case’ – is often high enough to eliminate any advantage such algorithms may have over a much simpler and faster path tracer. Therefore, path tracing is currently the default method of choice in most production renderers. Bidirectional methods are usually enabled manually by the user when they are required. When enabled, they typically apply across the entire image and cannot focus only on the areas where they make a difference.

Ideally, a renderer should only use the most efficient sampling technique(s) for any part of the image. With MIS, the contributions of each technique are weighted (i.e., reduced) in a way that tries to

minimize variance – but only after the necessary data has already been computed. Having many samples with a low MIS weight reduces the overall efficiency of the combined estimator, because the cost (number of samples) is much higher than it needs to be.

We propose a method to improve efficiency in one common setting: Scenes that are for the most part handled well by a path tracer but require light tracing or photon mapping to capture small but important caustic effects. A typical example are exteriors such as the one shown in Figure 1. A path tracer can render most of that image efficiently, except for the caustics caused by the car, the windows of the buildings, and the bus stop, which are better handled by a photon mapper. Photon mapping in such an exterior scene, however, can be extremely costly since many photons may need to be traced and stored to achieve a sufficient photon density in the caustics.

To address this problem, Jensen [Jen96] proposed a heuristic classification of photons into *global* and *caustic* photon maps and used so-called projection maps to guide caustics photons toward objects that are heuristically deemed to be ‘specular enough’ to produce caustic effects. With the right parameter settings, the algorithm can produce efficient renders, but this comes at the cost of unreasonable expectations on the user’s technical background. Furthermore, heuristic identification of caustic casters as specular objects does not apply to complex material models, not to mention the fact that other factors, such as light source size, are no less important in determining whether an object produces a visible caustic.

Our goal is robust, yet lightweight caustics rendering accessible to users with little technical background. To that end, we propose a technique to identify photon paths that offer a significant variance improvement over samples from a path tracer. The technique is based on path probabilities and avoids the inherently ambiguous heuristic classification of objects as potential caustic casters. We apply this technique in a rendering solution combining path tracing and photon mapping. Photon emission is automatically guided towards regions where the photons can provide substantial variance reduction over the path tracer. We show that this strategy significantly outperforms guiding based solely on visual importance [VKŠ*14]. In addition, we automatically determine an appropriate number of photons for a given scene, and the algorithm gracefully degenerates to pure path tracing for simple scenes where photon mapping makes no difference.

2. Related Work

For many scenes, path tracing (PT) [Kaj86] is the most efficient choice. A path tracer offers many benefits, like adaptive sampling, controllable image plane stratification, and many possibilities to reduce noise by filtering [ZJL*15]. With advanced importance sampling methods [VKŠ*14, HEV*16, MGN17], it is even possible to render complex indirect illumination and some low-frequency caustics in acceptable time. Furthermore, a path tracer can take advantage of the information provided by a photon map [Jen95] to trace hard to find illumination paths. Kaplanyan et al. [KD13] propose a (biased) method to allow a path tracer to sample certain non-physical paths that are represented by delta distributions.

Photon mapping [Jen96] was initially proposed as a two-pass method, integrated into a distribution path tracer. As discussed

above, the algorithm’s efficiency relies on hand-tuned parameters and a heuristic classification of objects as caustic-casters. The original idea was improved in various ways: Photon mapping can be made consistent and use less memory [HOJ08] or render distributed ray tracing effects [HJ09]. Furthermore, importance Sampling methods for photon tracing [VKŠ*14, PP98] can be used to avoid tracing invisible or low-contribution photons. Photon differentials [SFES07] can be used to improve the quality of sharp caustics. Directly visible caustics can be rendered efficiently by tracing paths from the light and connecting them to the camera [DW95]. For physically correct scenes without delta distributions, the bias from photon mapping can be eliminated as well, by replacing the density estimation with connections [QSH*15].

Bidirectional path tracing (BPT) [VG95a, PLW98] combines paths traced from the camera and the light sources via MIS. Photon mapping can be combined with BPT, as done with vertex connection and merging (VCM) [GKDS12, HPJ12]. Many path sampling techniques are used, with the hope that at least one technique will perform well in every situation. In practice, most of the time, most of the path sampling techniques are not needed and cause a serious overhead over a simple path tracer. We propose a method to reduce the number of samples invested in techniques that are less important for the current scene.

Metropolis-based methods explore the path space using Markov chains. Metropolis light transport (MLT) [VG97] mutates paths sampled by BPT to achieve better importance sampling. The multiplexed Metropolis light transport (MMLT) algorithm [HKD14] adds an additional parameter using tempering to choose the best sampling techniques based on the MIS weights. Difficult specular transport can be rendered with Markov chain approaches using regularization [KD13] or manifold exploration [JM12]. The Metropolis algorithm can also be used in conjunction with Photon mapping [ŠOHK16, GRŠ*17, HJ11], so as to distribute photons according to visibility, image contribution, or in a way that ensures a uniform error distribution. However, all these methods suffer from the drawbacks common to MCMC approaches. For instance, it is difficult to assess the convergence of the image, since the path space is not explored uniformly.

Metropolised VCM [ŠOHK16] distributes photons according to their MIS weighted image contribution. That improves efficiency by tracing fewer photons with low MIS weights. However, their method requires storing all camera vertices to evaluate the target function. Also, the MIS weights for VCM tend to assign an overly high weight to the photon mapper.

3. Background

This section reviews the basic concepts behind our derivation: The path integral formulation of light transport and multiple importance sampling.

3.1. Path Integral

The original rendering equation [Kaj86] can be rewritten as an integral over all light-transporting paths [Vea97]:

$$I = \int_{\Omega} f(\bar{x}) d\mu(\bar{x}) \quad (1)$$

The value I of a pixel in the image is given by the integral over the space of all paths Ω . A path $\bar{x} = x_0, \dots, x_n$ is a sequence of mutually visible points (on surfaces), where x_0 is on a light source and x_n on the camera sensor. The measurement contribution function

$$f(\bar{x}) = L_e(x_0)G(x_0 \leftrightarrow x_1) \left[\prod_{i=1}^{n-1} \rho_s(x_i)G(x_i \leftrightarrow x_{i+1}) \right] W_e(x_n) \quad (2)$$

determines the image contribution of the radiance travelling along \bar{x} . The emitted radiance L_e and the pixel sensitivity W_e are modified by the throughput of the path: The product of all geometry terms G and scattering distributions ρ_s [Vea97].

The differential product area measure $d\mu(\bar{x})$ is defined as the product of the differential surface areas at every point x_i along the path. When sampling a path \bar{x} for a Monte Carlo estimator of Equation (1), the corresponding pdf is therefore:

$$p(\bar{x}) = \prod_{i=0}^n p(x_i) \quad (3)$$

where $p(x_i)$ is the (surface area) pdf of sampling the vertex x_i . The resulting Monte Carlo estimator is then:

$$I \approx \frac{f(\bar{x})}{p(\bar{x})} \quad (4)$$

The variance of this estimator will be lower the closer $p(\bar{x})$ is to proportionality with $f(\bar{x})$ – the motivating idea behind path guiding approaches like [VKŠ*14].

3.2. Multiple Importance Sampling

Multiple importance sampling (MIS) [VG95b] weights the samples from a set of techniques to achieve the lowest possible variance. Weights are determined by a heuristic, measuring how efficient a technique is for the given sample compared to all other techniques. A common MIS heuristic is the balance heuristic. For the case of two sampling techniques A and B , this weight is:

$$w_A(x) = \frac{N_{APA}(x)}{N_{APA}(x) + N_{BPB}(x)} \quad (5)$$

The weight $w_A(x)$ for a technique A and sample x is computed based on the pdf $p(x)$ and the number of samples N of each technique.

The combined Monte Carlo estimator of Equation (1) for two techniques PM (photon mapping) and PT (path tracer) is then:

$$I \approx \frac{1}{N_{PT}} \sum_i^{N_{PT}} w_{PT}(\bar{x}_i) \frac{f(\bar{x}_i)}{p_{PT}(\bar{x}_i)} + \frac{1}{N_{PM}} \sum_i^{N_{PM}} w_{PM}(\bar{x}_i) \frac{f(\bar{x}_i)}{p_{PM}(\bar{x}_i)} \quad (6)$$

This combined estimator will converge to the correct result if for all \bar{x} where $f(\bar{x}) \neq 0$:

1. $w_{PT}(\bar{x}) + w_{PM}(\bar{x}) = 1$
2. either $p_{PT}(\bar{x}) \neq 0$ or $p_{PM}(\bar{x}) \neq 0$

A consequence of the first condition is that the samples are always weighted down by the MIS weights. If many samples from one technique are weighted down significantly, this can result in an overall less efficient estimator than disabling that technique entirely.

A consequence of the second condition is that the sampling distribution of the photon mapper can be modified freely – provided the path tracer remains unbiased. While existing photon guiding methods like [VKŠ*14] employ regularization to keep both the path tracer and the photon mapper correct estimators on their own, we take a different path: We deliberately set $p_{PM}(\bar{x}) = 0$ for all \bar{x} that the path tracer can sample efficiently. This approach improves the efficiency of the combined estimator by reducing the number of samples that will be weighted down significantly by MIS anyway. It also minimizes the number of photons that are used, thereby improving the efficiency even further.

4. Theory

As observed before, efficiency can be improved by restricting costly estimators like photon mapping to the subset of the domain that is most challenging for a simple path tracer. In this section, we propose an intuitive heuristic to determine whether a given photon should be sampled. We refer to such photons as being *useful*.

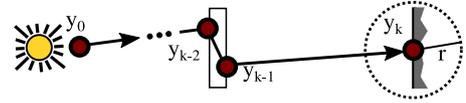


Figure 2: The photon y_k generated by the light path \bar{y} is useful, if reaching y_k from the light is more likely than reaching y_0 from anywhere within the photon mapping radius r around y_k .

Consider the example from Figure 2. Intuitively, the photon y_k is useful if the path tracer has a significantly lower probability to sample the path $\bar{y} = y_0 \dots y_k$ than the photon mapper. This yields the first version of our *usefulness measure* U' :

$$U'(\bar{y}) = \frac{p_{PM}(\bar{y})}{p_{PT}(\bar{y})} \quad (7)$$

With the pdfs of the path tracer and the photon mapper

$$p_{PT}(\bar{y}) = p_{PT}(y_k)p_{PT}(y_{k-1}|y_k)\dots p_{PT}(y_0|y_1) \quad (8)$$

$$p_{PM}(\bar{y}) = p_{PM}(y_0)p_{PM}(y_1|y_0)\dots p_{PM}(y_k|y_{k-1}) \quad (9)$$

The pdf of the path tracer $p_{PT}(\bar{y})$, in Equation (8), contains the pdf of sampling the point y_k from anywhere in the scene. While that could be approximated via density estimation over the camera paths, we use a simpler and more efficient approach: We assume that $p_{PT}(y_k)$ is uniform and non-zero only within the photon mapping radius around y_k . Therefore,

$$p_{PT}(y_k) = \frac{1}{\pi r^2} \quad (10)$$

where r is the photon mapping radius. Now, Equation (11) compares the probability of the photon mapper to sample the photon y_k with the probability of the path tracer to sample the same path starting anywhere within the photon mapping radius around y_k . As a consequence, a photon is more useful if the photon mapping radius is small compared to the size of the light source, which is part of the pdf $p_{PM}(y_0)$.

Theoretically, $U'(\bar{y})$ would now be greater than one whenever the photon mapper has a lower variance for the path \bar{y} than the path

tracer. Unfortunately, the efficiency of photon mapping is mainly due to path re-use, as was demonstrated in [GKDS12]. Hence, without accounting for the number of photons that are used, U' is not a reasonable measure of usefulness. To solve this problem, we propose the measure of usefulness $U(\bar{y})$ of a photon that was generated at the end of a light path \bar{y} :

$$U(\bar{y}) = \frac{N_{min} PPM(\bar{y})}{PPT(\bar{y})}, \quad (11)$$

where N_{min} is the minimum number of photons for which the path tracer would still be less efficient at sampling the path \bar{y} than the photon mapper. Based on this equation, we classify a photon as useful, if $U(\bar{y}) > 1$. This is somewhat similar to MIS with the maximum heuristic.

5. Implementation

Our method is based on a stripped-down version of VCM, which we refer to as vertex merging (VM), similar to Georgiev et al. [GKDS12]. The VM algorithm combines a path tracer with the most important techniques for caustic rendering: light tracing and photon mapping. Every photon is connected to the camera with a shadow ray. Density estimation is performed at every vertex of a camera path, except for the first one.

On every light source, we learn an emission distribution based on the MIS weighted image contribution of the photons emitted from that light. We refer to the resulting algorithm as VM with emission guiding (VM+EG). We combined this approach with the heuristic from Equation (11): Only the useful photons are considered when updating the distribution. We refer to this final version of the algorithm as VM+EG+U. Pseudocode is given in Algorithm 1.

Algorithm 1 Lightweight Photon Mapping: VM+EG(+U)

```

1: function RENDER_ITERATION
2:   ▷ Build the photon map
3:   photon_map = TRACE_PHOTONS( $N_{PM}$ )
4:   ▷ Splat all photons on the image
5:   for all photons  $y_k$  do
6:     ADD_CONTRIB_LT( $y_k$ )
7:      $y_k$ .contrib += contrib_lt
8:    $N_{PM} = 0$ 
9:   ▷ Trace one camera path per pixel
10:  for all pixel  $j$  in image do
11:     $\bar{z} \leftarrow$  TRACE_CAMERA_PATH( $j$ )
12:    for all vertices  $z_i$  along  $\bar{z}$  do
13:      ADD_CONTRIB_PT( $z_i$ )
14:      for all nearby photons  $y_k$  do
15:        ADD_CONTRIB_PM( $y_k$ )
16:         $y_k$ .contrib += contrib_pm
17:      if Equation (12) holds then
18:         $N_{PM} += 1$ 
19:   ▷ Update depending on whether / which guiding is used
20:  if algorithm == "VM+EG" then
21:    UPDATE_EMISSION_DISTRIB(all_photons)
22:  else if algorithm == "VM+EG+U" then
23:    UPDATE_EMISSION_DISTRIB(useful_photons)

```

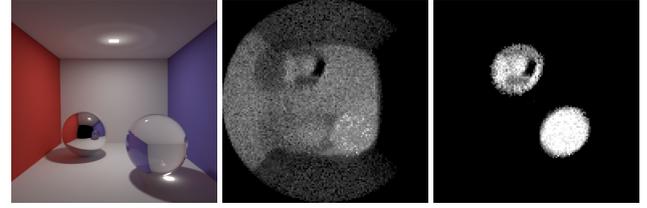


Figure 3: The emission histogram of the contribution (center) corresponds to the light source’s view of the scene. The histogram when only considering useful photons (right) results in emission only towards the specular spheres.

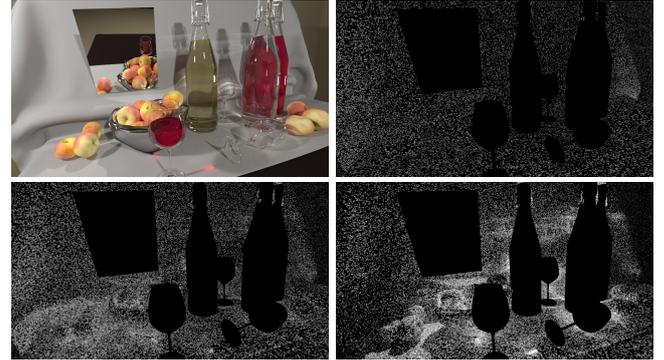


Figure 4: Photon densities in the STILL LIFE scene. Using the MIS weighted contribution of all photons (bottom left) is an improvement over uniform emission (top right). Our method (bottom right) improves the result by only tracing photons that represent caustics or indirect illumination due to caustics.

5.1. Learning

Whenever a photon is used to estimate a pixel value, the MIS weighted contribution is accumulated (c.f. lines 7 and 16 in Algorithm 1). At the end of every iteration, the accumulated contributions are used to update the pdfs used for emission sampling (lines 20-23). The number of pixels with more than a certain percentage of their luminance from photon mapping or light tracing determines the number of light paths to trace in the next iteration (lines 8, 17-18). Initially, the emission pdfs are all uniform.

5.2. Emission Guiding

We construct histograms in primary sample space (the uniform random numbers used to sample the emission), in the spirit of [Jen95]. We implemented our emission guiding for area light sources and directional light sources. This setup covers all three corner cases of light source descriptions: delta distributions, infinitely far away light sources, and lights with an actual surface. Extending our method to other types of light sources is straightforward. We also accumulate the total (useful) image contribution of each light to distribute the budget of light paths among the light sources. Figure 3 shows the resulting histograms after a single iteration, Figure 4 the resulting photon distribution.

The histogram resolution is set (and updated) proportionally to the number of rays emitted from the light source. Additionally, a Gaussian filter is used to reduce noise in the histograms. Most caus-

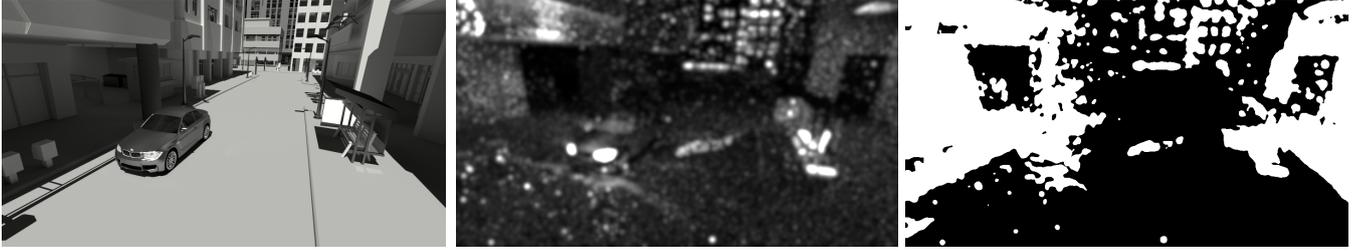


Figure 5: The MIS weighted contribution of the useful photons (center, exposure +4) is essentially an image of all caustics and indirect illumination due to the caustics. The image shown here is the blurred result after a single iteration. The image on the right shows which pixels our method classified as requiring photons (in white). Note that this does not only cover the caustics themselves but also pixels that receive indirect illumination due to caustics. For instance, the wall on the left is indirectly illuminated by the car.

tics are due to lights that are small in comparison to the surface they illuminate, therefore we neglect the spatial domain of area light sources. This results in additional blurring of the histogram.

In contrast to methods like [VKŠ*14, MGN17], we do not add a uniform density on top of our result. Therefore, the contribution of the photon mapper on its own will be biased. Fortunately, the MIS combination with the unbiased path tracer will eliminate that bias from the combined estimator.

5.3. Number of Photons

A common method to determine the number of photons per iteration is to use the same number of light paths as camera paths (or a constant multiple thereof), in order to balance the efforts invested into the two techniques. The number of photons is then the number of vertices along these light paths. If only a small fraction of the image contains caustics, then the photon density in those regions would be unnecessarily high when using one light path for every camera path. Therefore, we instead trace one light path for every pixel that receives a significant MIS weighted contribution from techniques involving useful photons. In order to do that, we use the following equation:

$$I_{PM,j} + I_{LT,j} > \alpha I_j \quad (12)$$

Here, $I_{PM,j}$ and $I_{LT,j}$ refer to the MIS weighted contribution of photon mapping and light tracing, respectively, to the pixel j . The estimated value of the pixel j is denoted as I_j , and multiplied with a threshold α , which we chose to be 1%. With $\alpha = 1\%$, the inequality is true whenever photon mapping and light tracing account for more than 1% of the estimated pixel value. We trace one light path for every pixel for which this inequality holds (see Algorithm 1). The result for the CAR scene is shown in Figure 5.

6. Evaluation and Results

We compared the performance of four different approaches: A simple path tracer (as a baseline reference), the VM algorithm with uniform photon emission, VM with emission guiding based on the visual importance of all photons (VM+EG), and VM with emission guiding based on only the useful photons (VM+EG+U). All algorithms were implemented in the same (in-house) renderer, share the majority of their code, and were tested on the same hardware: a desktop PC with 32GB RAM and an Intel i7-4790 CPU.

Our method makes it possible to render scenes that uniform photon emission cannot handle efficiently. A simple example of such a scene is given in Figure 6. Due to a large base plane, rendering the TORUS scene with unguided photon emission requires millions of light paths in every iteration, most of which will not be visible. Otherwise, as shown in the image, the photon mapper produces images with high variance – and bias – which will converge slowly.

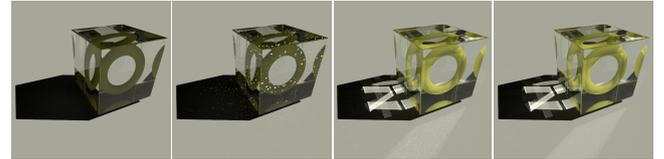


Figure 6: A torus inside a glass cube on a large plane, illuminated by a directional light source, after one minute of rendering. From left to right: path tracing (cannot sample caustics due to the delta distribution), uniform emission (would need many more photons), ours (same number of photons), and the reference.

The CAR scene in Figure 8 is a more realistic scene with similar issues as the TORUS. The majority of the visible part of the scene can be sampled efficiently by the path tracer. Every photon emitted towards these regions would be wasted. Indeed, our algorithm only emits photons towards the windows, the car, and the bus stop, resulting in significantly lower levels of noise with only a quarter of the photons (and thus memory) and half the number of light paths (and thus computational time) than the contribution-guided emission. This is also visible in the emission histograms (Figure 7).

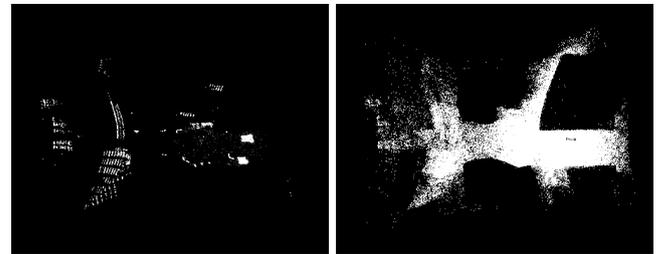


Figure 7: The emission histogram constructed by our method (on the left) for the CAR scene results in far fewer and far more important emission directions than the histogram based on all photons (on the right).

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	-

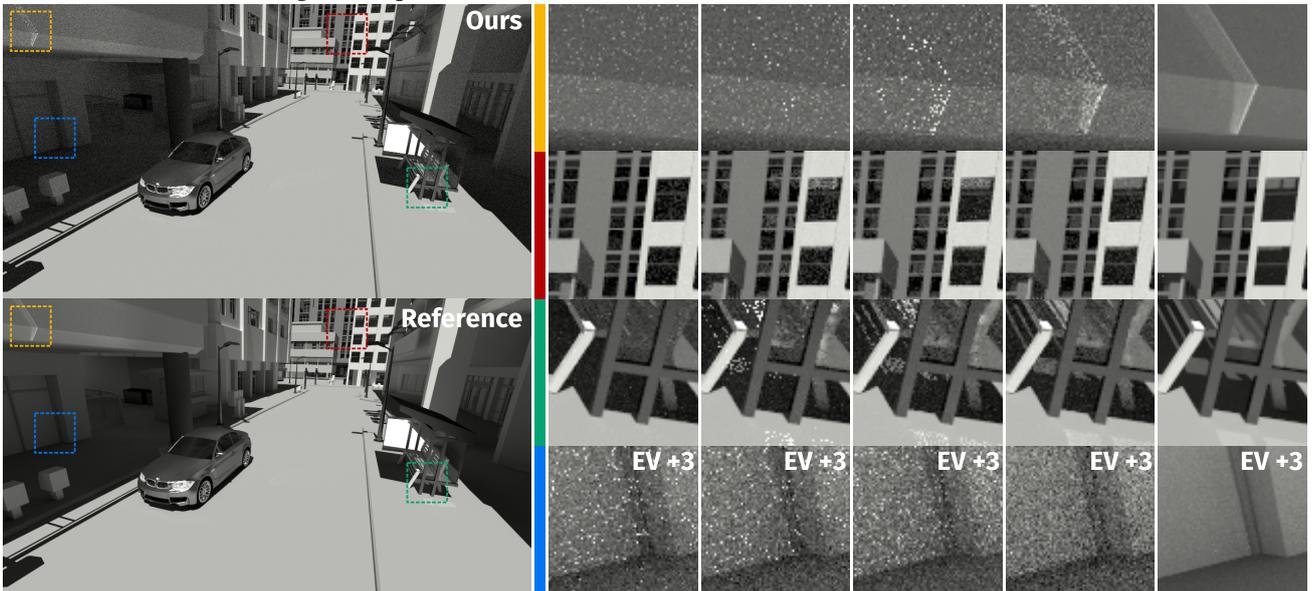


Figure 8: Equal-time comparison after one minute of rendering for the CAR scene. In this scene, the path tracer cannot sample some of the caustics at all (directional light source and perfect specularity). Our method results in half the number of light paths getting traced and therefore also a significantly 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. The reference was rendered with VM+EG within two hours.

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	-

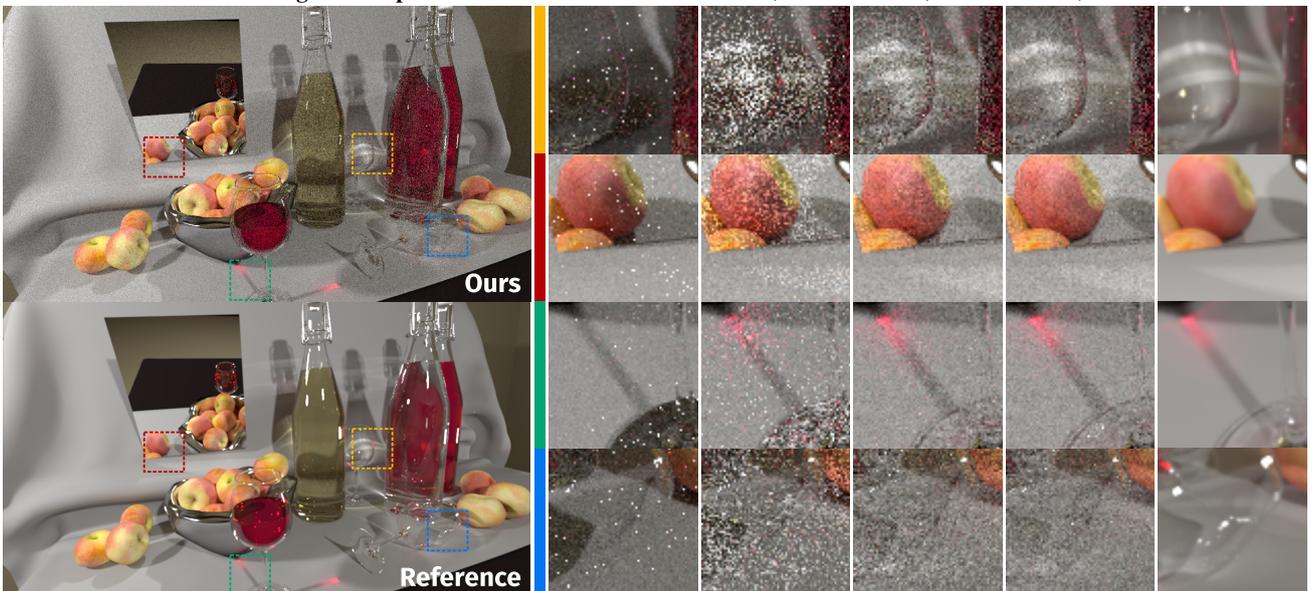


Figure 9: Equal-time comparison after one minute of rendering for the STILL LIFE scene. Here, the difference between our method and contribution-based guiding (VM+EG) is less visible than in the CAR scene (Fig. 8), because the majority of the image is influenced by caustics. The reference was rendered with VM+EG (60 hours).

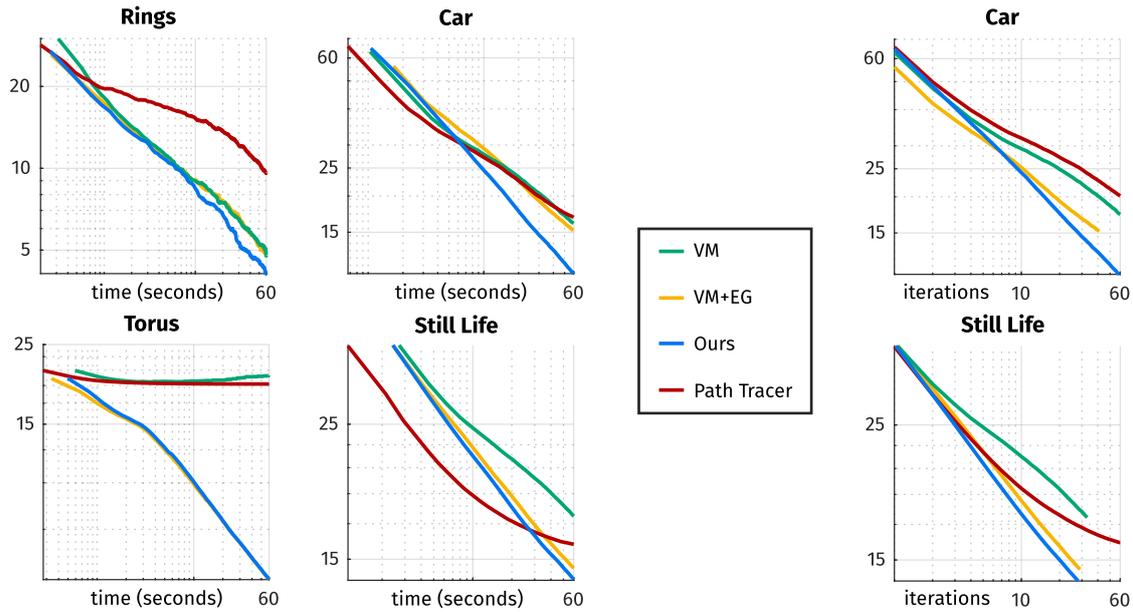


Figure 10: Equal-time (one minute) and equal-iteration-count convergence rate (log-RMSE) for some of our test scenes. Our method has either better or identical convergence rates in all our test scenes, even on a per-iteration level.

For scenes dominated by caustics, our method is at least as good as purely contribution-based emission guiding, which can be seen in the STILL LIFE scene in Figure 9. The STILL LIFE features a set of perfectly specular objects on a table in the middle of a large empty room illuminated by three area light sources. The visible regions that do not require photons are fairly small. Still, our method is slightly more efficient than contribution-based emission guiding.

Figure 10 compares the convergence rates of the different approaches. Throughout all our scenes, we achieved at least equal convergence rates (e.g., the TORUS), though most of the time, our method was significantly faster. Our method is often also better on a per-iteration basis. Although every iteration of our method uses significantly fewer photons, these are concentrated in the most important regions, i.e., the caustics.

6.1. Tracing an Adequate Number of Photons

Our method to determine an efficient number of light paths significantly improves efficiency in scenes with only a few small caustics. It can even disable light tracing altogether. This is illustrated with the simple variations of the classic Cornell Box in Figure 11.

However, we do not claim that our heuristic is optimal. We compared the convergence rate using our heuristic with multiple arbitrarily chosen constant numbers (Figure 12). Our heuristic is close to optimal for scenes that feature only small caustics and require many samples from the path tracer. If the path tracer does not require many samples, as is the case with the TORUS scene, the efficiency could be improved by using even more light paths. On the other hand, scenes that require the path tracer to resolve some complicated reflections of caustics, like the STILL LIFE, might be rendered more efficiently with fewer photons. However, our method adapts to the scene and is a better choice (on average) than any arbitrary constant.

6.2. Choosing the Parameters

Our approach still requires two parameters: N_{min} (Section 4) and the pixel threshold (Section 5.3). The latter has a fairly intuitive meaning: a high percentage only uses photons for the brightest caustics. The parameter N_{min} , on the other hand, is less intuitive. Lower values result in fewer photons considered useful. Because our measure of usefulness only relies on the path probabilities, the optimal choice depends on implementation details of the renderer, like the photon mapping radius, and not on the scene. However, the relationship between these characteristics and the N_{min} parameter is not very intuitive. We tested different values across all our test scenes and found that, for our implementation, $N_{min} = 5000$ works well. We also verified this choice in a simple benchmarking scene, shown in Figure 13.

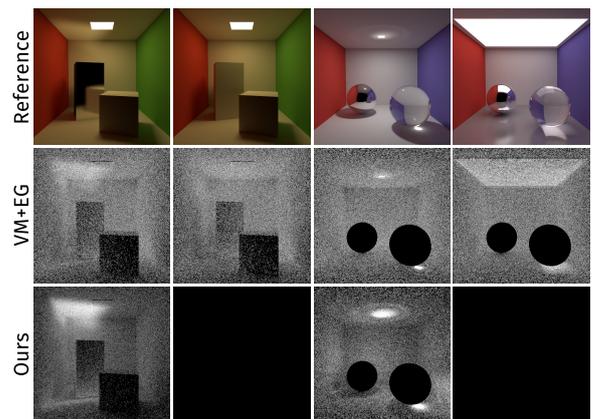


Figure 11: Variations of the Cornell Box scene and the resulting photon densities. Our usefulness heuristic focuses photons on caustics and indirect illumination due to caustics. Our pixel classification automatically disables photon mapping for the diffuse case and the large light source (because there are no useful photons).

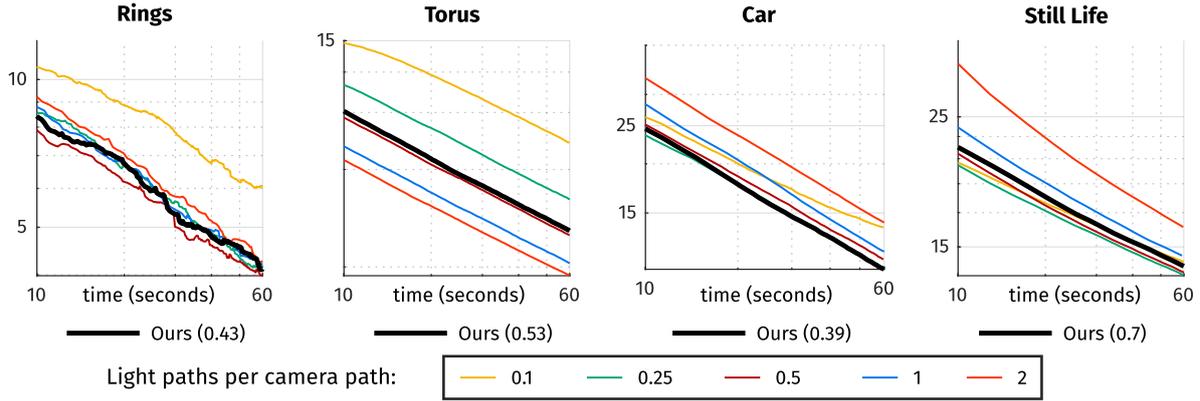


Figure 12: Equal-time (one minute, log-RMSE) comparison of our heuristic with arbitrarily chosen lower and higher numbers of light paths per camera path. The CAR scene demonstrates that our heuristic is close to optimal for cases where caustics only make up a small fraction of the image. While our heuristic is not always optimal, it performs better across different scenes than any fixed ratio.

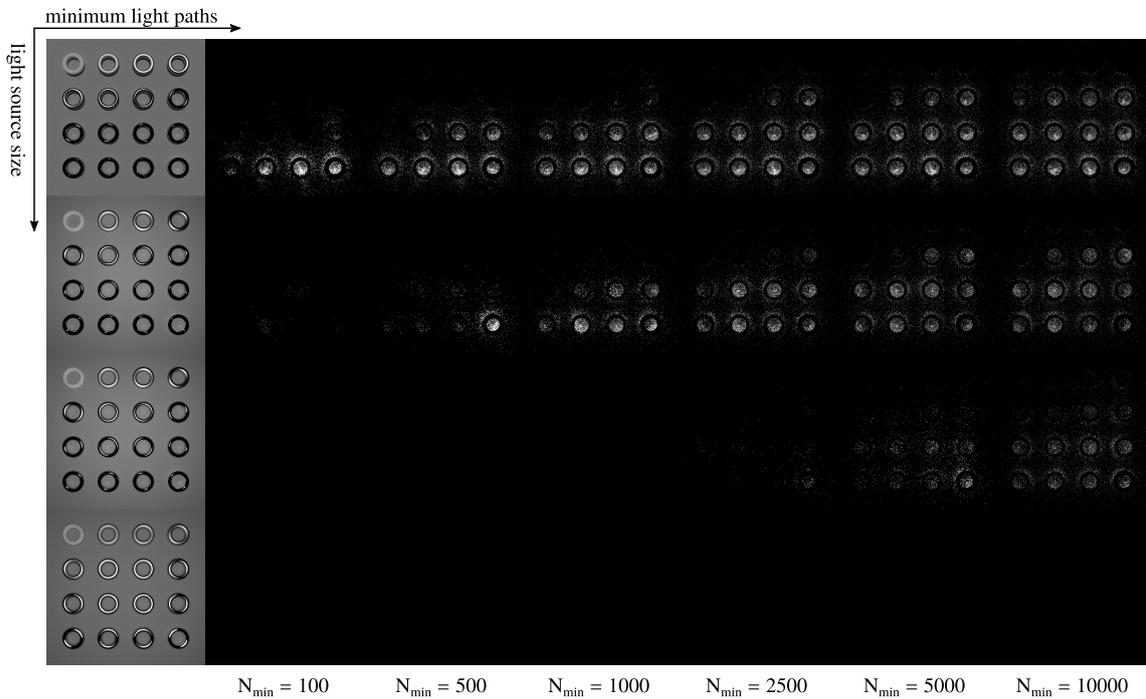


Figure 13: We placed a series of rings with decreasing roughness on a plane. We tested our algorithm with different light source sizes: directional light, small area light, medium-sized area light, large area light. The images from left to right show the photon densities when increasing the N_{min} parameter. Our parameter choice (5000) resulted in a reasonable photon distribution and the lowest level of noise.

7. Limitations and Future Work

The method we propose is quite simple. Yet, our results show that it can already have a significant impact. Various parts of both our theory and its practical application could be improved further. Overcoming the limitations outlined in the next sections could result in an even more efficient and more robust rendering method.

7.1. Limitations of the Usefulness Heuristic

The heuristic proposed in Equation (11) has three noteworthy weak-points that we had to compensate for in our application.

First, the parameter N_{min} has to be chosen empirically (c.f. Section 6.2). Also, it is influenced by the photon mapping radius. If the radius is not similar across scenes, then the choice of N_{min} may also depend on the scene. With a radius based on the pixel footprint, we found that a constant value for N_{min} works equally well across all our test scenes.

Secondly, the BRDF of the surface around the photon is ignored. Therefore, photons on glossy surfaces are considered to be as useful as photons on diffuse surfaces. For our application, we compensated for this by using the proposed heuristic only as a binary decision on top of the MIS weighted contribution. Because the MIS

weights for merging on a glossy surface are typically very low, our algorithm can still robustly handle scenes with many highly glossy surfaces.

Thirdly, our usefulness heuristic is quite similar to the MIS weights and shares a problem with those: The weights for certain effects involving strong indirect illumination are not optimal. Therefore, our approach will fall-back to the path tracer for almost all kinds of non-caustic indirect illumination. Also, caustics due to indirect illumination often receive a low weight as well.

7.2. Limitations of the Implementation

The histograms we used for guiding worked well in our scenes. However, they are not very flexible and can exhibit the teapot in a stadium problem. More adaptive representations like mixture models [VKŠ*14] or tree structures [MGN17] could be used to reduce the memory footprint and the required number of samples when building the distributions.

Caustics due to strong indirect illumination are an issue that simple emission guiding cannot completely solve. Minimizing the number of photons in that case would require guiding full paths, in the spirit of [VKŠ*14]. This could be achieved by propagating the usefulness of generated photons back to not only the light sources but also to all intermediate vertices. In this setting, it could also be interesting to use our usefulness measure with a Metropolis approach. For instance, the binary visibility and contribution target functions used by Šik et al. [ŠOHK16] could additionally apply our usefulness threshold, to focus on useful photons. Support for volumetric caustics should also be possible with these approaches.

Our method only rates the usefulness of a photon based on whether the path tracer could sample the same path. While this traces only photons that provide an actual benefit over path tracing, the photon distribution itself might not be optimal. For instance, in the CAR scene, the windows in the background receive a similar photon density to the caustics closer to the camera. This issue is somewhat alleviated by the fact that we also guide based on the image contribution. However, bright, yet far away, caustics might still receive too many photons. One way to improve our method further could therefore be to combine it with previous work trying to distribute photons in a way that achieves a more uniform distribution of error in the image, like [GRŠ*17].

7.3. Future Work

Discarding photons. The memory footprint of photon mapping could be reduced further by discarding photons that are classified as not useful by Equation (11). While our method prevents the emission of paths that will never cause useful photons, there might still be some photons generated at intermediate bounces, or due to regularization, that are not useful. These do not have to be stored or considered during the nearest neighbor search. Unfortunately, the memory savings from discarding these photons come at the cost of more involved MIS weight computations. In particular, the efficient method proposed by [Geo12], which we are also using, does not work anymore as the merging probabilities now depend on both the camera and the light sub-path.

Indirect illumination. Apart from caustics, bidirectional methods can also be useful for other effects, like indirect illumination. Rating the usefulness of a photon as a virtual point light, for instance, requires changes to Equation (11), although the basic ideas should still hold. Combining our method with a *guided path tracer*, like the ones proposed by [MGN17, VKŠ*14, BRDC12, HEV*16, Jen95], could improve the results even further. A guided path tracer can handle indirect illumination and many low-frequency caustics reasonably well, a fact that will be reflected in the pdf values of such a path, and therefore also in the usefulness heuristic.

Adaptive sampling and distributed rendering. Our method allows to detect which photons are useful for any given part of the image. For adaptive sampling, we could update the emission distribution once a part of the image has finished rendering, to only emit photons required for the remaining subset of the image. This could also improve the performance of photon mapping on a cluster.

8. Conclusion

We proposed a method to measure the usefulness of a given photon with respect to a path tracer. This usefulness is computed based only on path probabilities, without any scene-dependent heuristics or parameters. It can easily be used on top of traditional importance-driven emission guiding.

We tested our method for the most promising application: Guiding emission from the light sources towards those objects that cause visible caustics. Our experiments show that we can achieve significantly better equal-time and even equal-sample convergence rates than existing methods that guide the emission based solely on the visual importance. The method can be used with bidirectional path tracing to efficiently render glossy-diffuse and specular-diffuse paths, or in combination with vertex connection and merging to also capture specular-diffuse-specular and glossy-diffuse-glossy paths.

Our method only traces photons that a path tracer could not sample efficiently. For large scenes with only some small caustics, we significantly reduced the number of photons per iteration. Thereby, our method now enables us to use photon mapping even for large exterior scenes that would otherwise have required millions of photons per iteration – and likely still provided unsatisfactory results.

Our results demonstrate that the efficiency of Monte Carlo estimators combined via Multiple Importance Sampling can be increased significantly by restricting the more expensive estimators to the subsets of the domain where they offer a significant improvement. We believe that this general idea could also benefit many other applications.

9. Acknowledgements

We thank the anonymous reviewers for their valuable comments and suggestions. This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 642841 (DISTRO). It was further supported by the Czech Science Foundation grant 16-18964S.

References

- [BRDC12] BASHFORD-ROGERS T., DEBATTISTA K., CHALMERS A.: A significance cache for accelerating global illumination. In *Computer Graphics Forum* (2012), vol. 31, pp. 1837–1851. 9
- [DW95] DUTRÉ P., WILLEMS Y. D.: Importance-driven Monte Carlo light tracing. In *5th Eurographics Workshop on Rendering*. Springer Berlin Heidelberg, Darmstadt, Germany, 1995, pp. 185–197. 2
- [Geo12] GEORGIEV I.: *Implementing vertex connection and merging*. Tech. rep., Saarland University, 2012. 9
- [GKDS12] GEORGIEV I., KŘIVÁNEK J., DAVIDOVIĆ T., SLUSALLEK P.: Light transport simulation with vertex connection and merging. *ACM Trans. Graph.* 31, 6 (Nov. 2012), 192:1–192:10. 1, 2, 4
- [GRŠ*17] GRUSON A., RIBARDIÈRE M., ŠIK M., VORBA J., COZOT R., BOUATOUCH K., KŘIVÁNEK J.: A spatial target function for Metropolis photon tracing. *ACM Transactions on Graphics (TOG)* 36, 1 (2017), 4. 2, 9
- [HEV*16] HERHOLZ S., ELEK O., VORBA J., LENSCH H., KŘIVÁNEK J.: Product importance sampling for light transport path guiding. In *Computer Graphics Forum* (2016), vol. 35, pp. 67–77. 2, 9
- [HJ09] HACHISUKA T., JENSEN H. W.: Stochastic progressive photon mapping. *ACM Transactions on Graphics (TOG)* 28, 5 (2009), 141. 2
- [HJ11] HACHISUKA T., JENSEN H. W.: Robust adaptive photon tracing using photon path visibility. *ACM Transactions on Graphics (TOG)* 30, 5 (2011), 114. 2
- [HKD14] HACHISUKA T., KAPLANYAN A. S., DACHSBACHER C.: Multiplexed Metropolis light transport. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 100. 2
- [HOJ08] HACHISUKA T., OGAKI S., JENSEN H. W.: Progressive photon mapping. *ACM Trans. Graph.* 27, 5 (Dec. 2008), 130:1–130:8. 2
- [HPJ12] HACHISUKA T., PANTALEONI J., JENSEN H. W.: A path space extension for robust light transport simulation. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 191. 1, 2
- [Jen95] JENSEN H. W.: Importance driven path tracing using the photon map. In *Rendering Techniques 95*. Springer, 1995, pp. 326–335. 2, 4, 9
- [Jen96] JENSEN H. W.: Global illumination using photon maps. In *Rendering Techniques 96*. Springer, 1996, pp. 21–30. 2
- [JM12] JAKOB W., MARSCHNER S.: Manifold exploration: A Markov chain Monte Carlo technique for rendering scenes with difficult specular transport. *ACM Trans. Graph.* 31, 4 (July 2012), 58:1–58:13. 2
- [Kaj86] KAJIYA J. T.: The rendering equation. In *ACM Siggraph Computer Graphics* (1986), vol. 20, ACM, pp. 143–150. 2
- [KD13] KAPLANYAN A. S., DACHSBACHER C.: Path space regularization for holistic and robust light transport. In *Computer Graphics Forum* (2013), vol. 32, pp. 63–72. 2
- [KGH*14] KŘIVÁNEK J., GEORGIEV I., HACHISUKA T., VÉVODA P., ŠIK M., NOWROUZEZAHRAI D., JAROSZ W.: Unifying points, beams, and paths in volumetric light transport simulation. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 103. 1
- [MGN17] MÜLLER T., GROSS M., NOVÁK J.: Practical path guiding for efficient light-transport simulation. 91–100. 2, 5, 9
- [PLW98] P. LAFORTUNE E., WILLEMS Y.: Bi-directional path tracing. In *Proceedings of Third International Conference on Computational Graphics and Visualization Techniques (Compugraphics' 01 1998)*, vol. 93. 2
- [PP98] PETER I., PIETREK G.: Importance driven construction of photon maps. In *Rendering Techniques 98*. Springer, 1998, pp. 269–280. 2
- [QSH*15] QIN H., SUN X., HOU Q., GUO B., ZHOU K.: Unbiased photon gathering for light transport simulation. *ACM Trans. Graph.* 34, 6 (Oct. 2015), 208:1–208:14. 2
- [SFES07] SCHJØTH L., FRISVAD J. R., ERLEBEN K., SPORRING J.: Photon differentials. In *Proceedings of the 5th international conference on Computer graphics and interactive techniques in Australia and South-east Asia* (2007), ACM, pp. 179–186. 2
- [ŠOHK16] ŠIK M., OTSU H., HACHISUKA T., KŘIVÁNEK J.: Robust light transport simulation via metropolis bidirectional estimators. *ACM Transactions on Graphics (TOG)* 35, 6 (2016), 245. 2, 9
- [Vea97] VEACH E.: *Robust Monte Carlo methods for light transport simulation*. No. 1610. Stanford University PhD thesis, 1997. 1, 2, 3
- [VG95a] VEACH E., GUIBAS L.: Bidirectional estimators for light transport. In *Photorealistic Rendering Techniques*. Springer, 1995, pp. 145–167. 2
- [VG95b] VEACH E., GUIBAS L. J.: Optimally combining sampling techniques for Monte Carlo rendering. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques* (1995), ACM, pp. 419–428. 1, 3
- [VG97] VEACH E., GUIBAS L. J.: Metropolis light transport. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (1997), ACM Press/Addison-Wesley Publishing Co., pp. 65–76. 2
- [VKŠ*14] VORBA J., KARLÍK O., ŠIK M., RITSCHER T., KŘIVÁNEK J.: On-line learning of parametric mixture models for light transport simulation. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 101. 2, 3, 5, 9
- [ZJL*15] ZWICKER M., JAROSZ W., LEHTINEN J., MOON B., RAMAMOORTHY R., ROUSSELLE F., SEN P., SOLER C., YOON S.-E.: Recent advances in adaptive sampling and reconstruction for Monte Carlo rendering. 667–681. 2