

Correlation-Aware Multiple Importance Sampling for Bidirectional Rendering Algorithms

Pascal Grittmann¹ Iliyan Georgiev² Philipp Slusallek^{1,3}

¹Saarland University, Germany ²Autodesk, United Kingdom ³DFKI, Germany

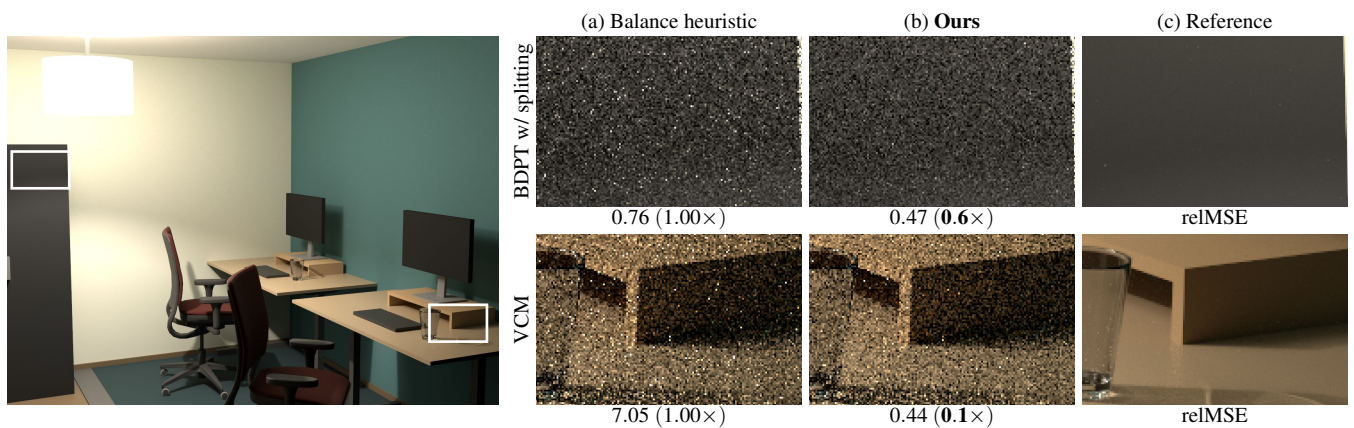


Figure 1: A scene featuring complex indirect illumination (lamp shade) and caustics (glass) – a prime use-case of bidirectional algorithms. We show two such methods: bidirectional path tracing with splitting (top row), and vertex connection and merging (bottom row). (a) Both exhibit problems with MIS in this scenario, due to correlation by shared path prefixes. (b) Our simple heuristic solves these problems.

Abstract

Combining diverse sampling techniques via multiple importance sampling (MIS) is key to achieving robustness in modern Monte Carlo light transport simulation. Many such methods additionally employ correlated path sampling to boost efficiency. Photon mapping, bidirectional path tracing, and path-reuse algorithms construct sets of paths that share a common prefix. This correlation is ignored by classical MIS heuristics, which can result in poor technique combination and noisy images. We propose a practical and robust solution to that problem. Our idea is to incorporate correlation knowledge into the balance heuristic, based on known path densities that are already required for MIS. This correlation-aware heuristic can achieve considerably lower error than the balance heuristic, while avoiding computational and memory overhead.

CCS Concepts

• **Computing methodologies** → **Rendering; Ray tracing;**

Keywords: light transport, bidirectional path tracing, VCM, multiple importance sampling, path correlation

1. Introduction

Monte Carlo integration is the standard approach to physically based rendering. The global illumination in a scene can be estimated by tracing randomly sampled light transport paths. Paths can be generated in a variety of ways: unidirectionally from the camera [Kaj86], via guiding [Jen95; VKŠ*14], or from both the camera and the emitters [VG95a; HPJ12; GKDS12]. Different path sampling techniques can be combined via multiple importance sampling (MIS) [VG95b], which is a key ingredient in achieving robustness under varying illumination and scene configurations.

Some rendering algorithms achieve efficiency by generating correlated paths. They reduce the sampling cost by constructing paths that share a common prefix. This is the overarching idea behind photon mapping [Jen96], path splitting [AK90] or distribution ray tracing [CPC84], and path reuse methods [WGGH20; KDB14]. Many of these approaches also make use of MIS to increase robustness, e.g. vertex connection and merging (VCM) [GKDS12; HPJ12], or the method of Popov et al. [PRDD15] which utilizes splitting in bidirectional path tracing (BDPT) [VG95b] by tracing multiple shadow rays from the same point. However, MIS and the

popular balance heuristic operate under the assumption of independent sampling, i.e. no correlation between individual paths. Lacking a practical alternative, the aforementioned algorithms rely on the balance heuristic. This is problematic as illustrated in Fig. 1: the heuristic produces poor combination weights, creating an overly dark image with strong outliers. Given a splitting technique that concatenates a single prefix with n suffix paths, the balance heuristic treats these as n mutually independent full paths. It is oblivious to the increased variance resulting from using a shared prefix.

The impact of correlation on MIS has received little attention in prior work. Even the recently derived optimal weights [KVG*19] assume mutually independent paths and samples. For the two examples in Fig. 1, ad-hoc solutions have been proposed [PRDD15; JG18; Jen19]. Besides being specific to one of the two cases, these approaches can produce unsatisfactory results or rely on unintuitive parameters. A more general and effective solution would be to incorporate variance estimates [GGSK19], but as we will discuss later, that approach is neither robust nor efficient enough to be practical in the two cases in Fig. 1.

In this paper, we investigate when and why correlated paths are problematic for MIS and propose a simple and effective heuristic as a remedy that is based solely on sampling densities. Even without proof of optimality, the empirical evaluation of our heuristic shows consistent improvement over the balance heuristic and prior work across all our test scenes. Implementing our heuristic is straightforward as it relies on the same quantities as the balance heuristic. It also adds no noteworthy computation or memory overhead.

2. Background and problem statement

We begin with a brief overview of path-space Monte Carlo light transport integration, and discuss when and why correlated paths are problematic for multiple importance sampling.

2.1. Monte Carlo path integration

The value of each pixel in a rendered image can be expressed as an integral over the space of all light-carrying paths \bar{x} [Vea97]:

$$F = \int_{\mathcal{P}} f(\bar{x}) d\mu(\bar{x}), \quad (1)$$

where \mathcal{P} includes all paths $\bar{x} = \mathbf{x}_0\mathbf{x}_1 \dots \mathbf{x}_k$ of any length k and $\mu(\bar{x})$ is the product area measure. An ordinary Monte Carlo estimator for this integral samples n paths \bar{x}_i with densities $p(\bar{x}_i)$ and has the form:

$$\langle F \rangle_n = \sum_{i=1}^n \frac{f(\bar{x}_i)}{np(\bar{x}_i)}. \quad (2)$$

The error of this estimator is characterized by its variance. If the n paths are *mutually independent*, i.e., uncorrelated, the variance reads:

$$\mathbb{V}[\langle F \rangle_n] = \int_{\mathcal{P}} \frac{f^2(\bar{x})}{np(\bar{x})} d\mu(\bar{x}) - \frac{1}{n} F^2. \quad (3)$$

That is, the variance depends on the choice of probability density function (PDF) p , and it decreases inversely proportionally to the number n of uncorrelated samples (i.e., paths).

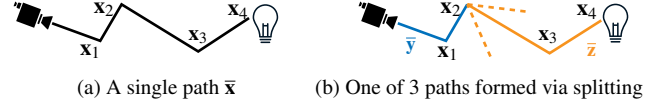


Figure 2: Without splitting (a), we trace independent full paths. Splitting (b) generates multiple full paths $\bar{x}_i = \bar{y}\bar{z}_i$ that share a prefix \bar{y} (in blue) and have mutually independent suffixes \bar{z}_i (in orange).

2.2. Path correlation through splitting

Some rendering methods increase efficiency by sampling correlated paths. One such example is path splitting [AK90; VK16] which constructs n paths $\bar{x}_i = \bar{y}\bar{z}_i$ that all share a prefix \bar{y} , each having an independently sampled suffix \bar{z}_i (see Fig. 2). Splitting estimators take the form

$$\langle F \rangle_{\text{split}} = \sum_{i=1}^n \frac{f(\bar{y}\bar{z}_i)}{np(\bar{y}\bar{z}_i)}. \quad (4)$$

Due to the shared prefix \bar{y} , the paths \bar{x}_i are no longer mutually independent – they are correlated.

The variance $\mathbb{V}[\langle F \rangle_{\text{split}}]$ of the splitting estimator can be expressed as a weighted sum of the variance $V_{\bar{y}}$ due to the prefix and the variance $V_{\bar{z}}$ due to the suffix [BM97]. The prefix variance,

$$V_{\bar{y}} := \mathbb{V} \left[\frac{F_{\bar{z}}(\bar{y})}{p(\bar{y})} \right] = \int_{\mathcal{P}} \frac{F_{\bar{z}}^2(\bar{y})}{p(\bar{y})} d\mu(\bar{y}) - F^2, \quad (5)$$

is the variance of a hypothetical estimator that knows the exact integral $F_{\bar{z}}(\bar{y}) = \int_{\mathcal{P}} f(\bar{y}\bar{z}) d\mu(\bar{z})$ over all suffix paths. The suffix variance is the expectation of the variance of a primary estimator $\langle F \rangle_1$ over all possible prefixes \bar{y} :

$$V_{\bar{z}} := \mathbb{E}[\mathbb{V}[\langle F \rangle_1 | \bar{y}]] = \int_{\mathcal{P}} \frac{f^2(\bar{x})}{p(\bar{x})} d\mu(\bar{x}) - \int_{\mathcal{P}} \frac{F_{\bar{z}}^2(\bar{y})}{p(\bar{y})} d\mu(\bar{y}). \quad (6)$$

While with independent sampling the total variance is inversely proportional to the sample count, i.e., $\mathbb{V}[\langle F \rangle_n] = \frac{1}{n} (V_{\bar{y}} + V_{\bar{z}})$, splitting reduces only the suffix variance:

$$\mathbb{V}[\langle F \rangle_{\text{split}}] = V_{\bar{y}} + \frac{1}{n} V_{\bar{z}}. \quad (7)$$

Splitting can still be an efficient strategy, particularly when $V_{\bar{y}}$ is low, e.g., due to highly glossy bounces in the prefix \bar{y} , or if the cost of sampling the prefix is high.

Expanding the terms in Eq. (7), we can write the variance in a form similar to that of independent sampling (3):

$$\mathbb{V}[\langle F \rangle_{\text{split}}] = \int_{\mathcal{P}} \frac{f^2(\bar{x})}{np(\bar{x})} d\mu(\bar{x}) + \frac{n-1}{n} \int_{\mathcal{P}} \frac{F_{\bar{z}}^2(\bar{y})}{p(\bar{y})} d\mu(\bar{y}) - F^2. \quad (8)$$

Compared to the variance of independent sampling (3), the value is increased by the second integral term – the covariance. Intuitively, the impact of the correlation depends on the value of $p(\bar{y})$ compared to the full-path density $p(\bar{x})$. This insight is a key motivation behind our heuristic presented in Section 3.

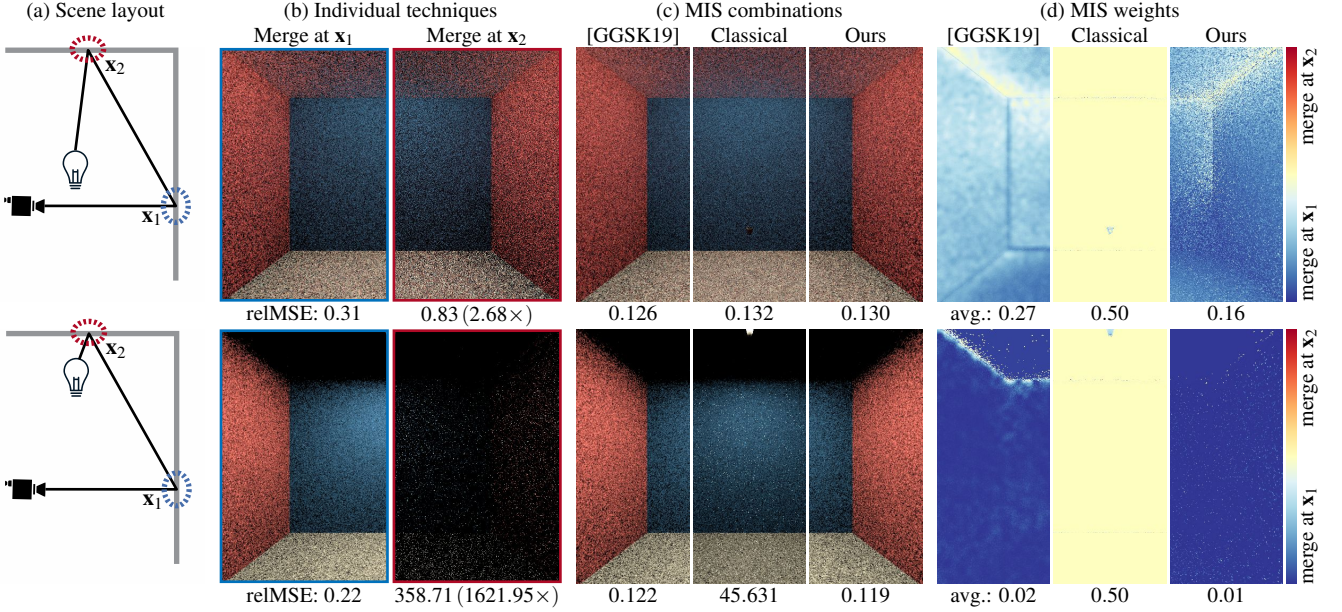


Figure 3: A box with diffuse materials and a small area light (shining upwards) positioned far from the ceiling (top) and close to it (bottom). We render length-3 paths and consider vertex merging techniques only. The variance of merging at \mathbf{x}_2 increases as the light moves closer to the ceiling, resulting in a dark image with all energy focused in a few outliers. The classical balance heuristic does not capture that increased variance and assigns equal weights to the two merging techniques in both scene configurations, producing a poor combination in the latter.

2.3. Multiple importance sampling

Importance sampling can reduce the variance of a Monte Carlo estimator by using a sampling technique whose density p is (at least roughly) proportional to the integrand. While in practice a single technique cannot achieve perfect proportionality, several approximate techniques can be combined together via multiple importance sampling (MIS). The multi-sample MIS estimator [VG95b]

$$\langle F \rangle_{\text{MIS}} = \sum_t \sum_{i=1}^{n_t} w_t(\bar{\mathbf{x}}_{t,i}) \frac{f(\bar{\mathbf{x}}_{t,i})}{n_t p_t(\bar{\mathbf{x}}_{t,i})} \quad (9)$$

evaluates n_t samples with density p_t from each technique t . The balance heuristic [VG95b],

$$w_t(\bar{\mathbf{x}}) = \frac{n_t p_t(\bar{\mathbf{x}})}{\sum_k n_k p_k(\bar{\mathbf{x}})}, \quad (10)$$

is a provably good choice for the weighting function, provided that the n_t samples from each technique are mutually independent [Vea97]. Unfortunately, when the samples (for some techniques) are correlated, the balance heuristic drifts farther from optimality. The reason is that it minimizes only the second moment in the variance of $\langle F \rangle_{\text{MIS}}$, i.e., the first integral in Eqs. (3) and (8). While this is reasonable with independent samples, ignoring the second integral in (8) – the covariance – can yield a poor technique combination. The balance heuristic effectively assumes that both $V_{\bar{\mathbf{y}}}$ and $V_{\bar{\mathbf{z}}}$ decrease with increasing n , while in reality only $V_{\bar{\mathbf{z}}}$ does. Nevertheless, lacking an alternative, the balance heuristic is used even in cases involving correlated techniques [GKDS12; HPIJ12; PRDD15; NID20].

2.4. Case study: Vertex connection and merging

Photon mapping [Jen96] is a popular method to render caustics [ŠK19; EK20; GPSK18]. It is a splitting estimator: each camera prefix subpath branches out into multiple (photon) suffix subpaths. The difference to the classical splitting estimator (4) is that the suffixes are sampled from the emitters. Vertex connection and merging (VCM) [GKDS12; HPIJ12] constructs one such “merging” estimator at each vertex on a camera subpath, combining all estimates using the balance heuristic. While in path tracing the splitting factor n is typically small [VK16], in VCM it is in the order of millions (i.e., the number of subpaths started from the emitters).

It only takes a simple scene to demonstrate the catastrophic failure of MIS with correlated samples. Figure 3 shows a diffuse box illuminated by a small area light source at two different positions: near the floor (top row) and near the ceiling (bottom row). We consider length-3 paths $\bar{\mathbf{x}} = \mathbf{x}_0 \mathbf{x}_1 \mathbf{x}_2 \mathbf{x}_3$ (i.e., one-bounce indirect illumination) and vertex merging techniques only. For each path there are thus two possible techniques, merging at \mathbf{x}_1 and \mathbf{x}_2 respectively. We compare three variants of the balance heuristic for combining these two techniques: classical (10), variance-aware [GGSK19], and our proposed (introduced in Section 3). As the light source moves closer to the ceiling, the variance of merging at \mathbf{x}_2 explodes as the camera subpath is less likely to find the shrinking, brightly illuminated spot on the ceiling. However, this is not reflected in the weights of the classical balance heuristic which ends up producing an extremely noisy image.

The two techniques differ only in the direction in which the edge $\mathbf{x}_1 \mathbf{x}_2$ is sampled. And the corresponding path densities are equal

when the vertices \mathbf{x}_1 and \mathbf{x}_2 are both diffuse, since in that case $p(\mathbf{x}_1 \rightarrow \mathbf{x}_2) = p(\mathbf{x}_2 \rightarrow \mathbf{x}_1) = G(\mathbf{x}_1 \leftrightarrow \mathbf{x}_2)/\pi$, where G is the geometry term (which is symmetric w.r.t. \mathbf{x}_1 and \mathbf{x}_2). The classical balance heuristic thus assigns equal weights to the techniques, regardless of the geometry of the actual path.

While the variance-aware weights [GGSK19] can solve this problem in this simple case, they are not practical. They require computing variance estimates of every single merging technique, for every path length. Crucially, these estimates need to be accurate, which is expensive for techniques that produce nothing but outliers. (We have used 128 samples/pixel for the experiment in Fig. 3.) Hence, this method is not efficient in this setting.

The approach of Popov et al. [PRDD15], which sets $n = 1$ in the classical heuristic (10), would not work. For one, n cancels out when only merging techniques are being combined. Worse, when vertex connection techniques are included, setting $n = 1$ effectively disables all merging techniques, reverting VCM to bidirectional path tracing. The heuristics of Jendersie et al. [JG18; Jen19] address this specific failure case but use unintuitive parameters and can sometimes perform worse than the classical balance heuristic.

3. Correlation-aware balance heuristic

Our goal is to incorporate correlation into MIS in a practical, lightweight, and robust manner. The optimal way of doing so would involve computing a large number of complex integrals, which has been shown to be impractical [PRDD15]. As an alternative, we propose a heuristic that, like the balance heuristic, relies only on readily available path densities. It utilizes the prefix and suffix subpath densities in addition to the full-path density. The key to making this approach work is to convert these densities into unitless and hence comparable quantities.

Our idea is to apply a correction to the balance heuristic (10), in a similar fashion to Grittmann et al. [GGSK19]:

$$w_t(\bar{\mathbf{x}}) = \frac{c_t(\bar{\mathbf{x}})n_t p_t(\bar{\mathbf{x}})}{\sum_k c_k(\bar{\mathbf{x}})n_k p_k(\bar{\mathbf{x}})}, \quad (11)$$

where the correction factor $c_t(\bar{\mathbf{x}})$ has the bounds

$$\frac{1}{n_t} \leq c_t(\bar{\mathbf{x}}) \leq 1. \quad (12)$$

A group of correlated paths cannot yield higher variance than a single path, hence the lower bound of $1/n_t$. They also cannot reduce the variance more than independent samples do, hence the upper bound of one. In the following, we motivate our heuristic to find a suitable $c_t(\bar{\mathbf{x}})$ and describe how to evaluate it. In essence, we keep c_t close to its lower bound of $1/n_t$ unless we have evidence of low correlation. This follows the reasoning of Popov et al. [PRDD15].

3.1. Density-based correction factor

The classical MIS heuristics, including the balance heuristic, are based on the assumption that high sampling density implies low variance [Vea97]. We expand on this idea by inspecting the subpath densities to measure how much splitting increases the estimator variance over independent sampling. For paths with high prefix density (e.g. with specular interactions), the impact on the variance




Figure 4: Given a path $\bar{\mathbf{x}} = \bar{\mathbf{y}}\bar{\mathbf{z}}$ (in black), we want to compute the probability of sampling a similar path, i.e., one falling within the shaded region. We compute the probability that each edge would produce a vertex \mathbf{x}'_i that falls within a disc around the actual vertex \mathbf{x}_i . The blue and orange lines visualize how these are computed for the camera prefix and the light suffix, respectively.

is expected to be low. Conversely, low prefix density (e.g. due to a diffuse bounce) indicates high prefix variance $V_{\bar{\mathbf{y}}}$, which reduces the impact of the $1/n$ factor in Eq. (7). The correction factor $c_t(\bar{\mathbf{x}})$ should thus be roughly proportional to the density of the prefix $\bar{\mathbf{y}}$.

Let $P(\bar{\mathbf{x}}) \in [0, 1]$ be a measure of path density which we will define below. For a given path $\bar{\mathbf{x}}$, we set the correction factor $c_t(\bar{\mathbf{x}})$ to the ratio of the prefix density to the full-path density:

$$c_t(\bar{\mathbf{x}}) = \max\left(\frac{P(\bar{\mathbf{y}})}{P(\bar{\mathbf{x}})}, \frac{1}{n_t}\right), \quad (13)$$

clipping it to the known lower bound of $1/n_t$. We require $P(\bar{\mathbf{y}}) \leq P(\bar{\mathbf{x}})$, which retains the balance heuristic (i.e., $c_t(\bar{\mathbf{x}}) = 1$) when the prefix $\bar{\mathbf{y}}$ has high relative density. This is the case when, e.g., the prefix is a glossy subpath and the suffix is diffuse. If the prefix density is low, e.g., due to diffuse interactions (as in Fig. 3), the MIS weight (11) is reduced – unless the suffix undergoes even more diffuse interactions.

3.2. Unitless path density measure

We cannot use the raw densities of $\bar{\mathbf{x}}$ and $\bar{\mathbf{y}}$ directly in Eq. (13), because they have different units (due to the different number of vertices) and are scale-dependent. We therefore consider the *unitless* probability to sample a similar (sub)path instead. That is, $P(\bar{\mathbf{x}})$ is the probability that each vertex \mathbf{x}'_i of another path $\bar{\mathbf{x}}'$ lies within distance r to the corresponding vertex \mathbf{x}_i (see Fig. 4). For the subpaths $\bar{\mathbf{y}}$ and $\bar{\mathbf{z}}$ these are respectively the products

$$P(\bar{\mathbf{y}}) = \prod P(\mathbf{y}_{i-1} \rightarrow \mathbf{y}_i) \text{ and } P(\bar{\mathbf{z}}) = P(\mathbf{z}_0) \prod P(\mathbf{z}_{i-1} \rightarrow \mathbf{z}_i), \quad (14)$$

where $P(\mathbf{x}_{i-1} \rightarrow \mathbf{x}_i)$ is the probability that a \mathbf{x}'_i lies within radius r around \mathbf{x}_i , when sampled from \mathbf{x}_{i-1} . The probability for the full path is given by basic probability laws:

$$P(\bar{\mathbf{x}}) = P(\bar{\mathbf{z}} \cap \bar{\mathbf{y}}) = P(\bar{\mathbf{z}}) + P(\bar{\mathbf{y}}) - P(\bar{\mathbf{z}})P(\bar{\mathbf{y}}). \quad (15)$$

Computing the probability for each vertex in Eq. (14) exactly would involve costly integration, but a cheap approximation is available [GKDS12]:

$$P(\mathbf{x}_i \rightarrow \mathbf{x}_j) = \int_{D_r} p(\mathbf{x}_i \rightarrow \mathbf{x}) d\mathbf{x} \approx \min(\pi r^2 p(\mathbf{x}_i \rightarrow \mathbf{x}_j), 1). \quad (16)$$

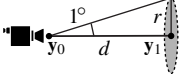
The approximation assumes that the sampling density is constant inside the r -neighborhood D_r of \mathbf{x}_i , which is a disk. Clamping the

result to one ensures we get a valid probability. Such clamping occurs when the probability density is focused in a region smaller than the disk. In practice, if the radius is well chosen this only happens for highly glossy scattering and small area lights.

3.3. Radius

What remains is setting the radius r in Eq. (16). The problem is similar to choosing a good radius for photon mapping: Finding the best possible value is difficult, but simple heuristics can go a long way. While in photon mapping the radius trades bias for variance, we have different considerations to make. If the radius is too large, the probabilities will all be close to one. This would make $P(\bar{x}) \approx 1$, and we revert to the balance heuristic. If, on the other hand, the radius is too small, we penalize glossy bounces too much.

We could set the radius as a fraction of the scene extent; however, this approach is not robust, e.g., when the camera sees only a small part of a large scene. An alternative is to use pixel footprints [ŠOHK16]. The resulting radius then depends on the image resolution and the camera field of view. But increasing the image resolution should not alter $P(\bar{x})$ as it has little effect on the path sampling variance. Therefore, we replace the pixel footprint by a related quantity independent of the resolution or field of view – the footprint of a one-degree viewing angle:

$$r := d \tan \frac{\pi}{180} \approx 0.0175d, \quad (17)$$


where $d = \|y_1 - y_0\|$ is the distance to the hit point along the camera ray, as sketched on the right above.

4. Evaluation

We have applied our correlation-aware balance heuristic (11) to two rendering algorithms that sample correlated paths: vertex connection and merging (VCM) and bidirectional path tracing (BDPT) with splitting. In both methods, evaluating the correction factor (13) is straightforward as it relies only on densities that the classical balance heuristic already requires. We implemented the methods using the SeeSharp rendering framework. The source code of our experiments is included in the supplemental and also available on GitHub.

Our heuristic does not add any measurable overhead over the balance heuristic. In the results presented below, we use the same set of path samples when comparing the various MIS combinations on each scene. Thus, any differences between the images are solely due to differences in the weighting.

4.1. Vertex connection and merging

Our VCM implementation follows the original formulation [GKDS12; HPJ12], with the exception that we forego merging at the second camera vertex, y_1 . This technique produces an image identical to light tracing but with added bias due to blurring. We compare our heuristic (11) to the classical balance heuristic (10), the variance-aware balance heuristic [GGSK19], as well as the approach of Jendersie [Jen19] (which supersedes that of Jendersie and Grosch [JG18]).

We show results on three scenes in Fig. 5. On ROUGH GLASSES and MINIMALIST ROOM, the balance heuristic (a) suffers from the same failure as in Fig. 3. On the HOME OFFICE scene, it works well despite the indirect illumination around the walls close to the ceiling: the light source is large, so the variance is high, the density is low, and merging on the ceiling is beneficial.

The weighting scheme of Jendersie [Jen19] (b) eliminates the outliers in the top two scenes. Unfortunately, on the ROUGH GLASSES scene, which features surfaces with varying roughness, it also deteriorates the quality of glossy reflections compared to the balance heuristic (bottom zoom-ins). This limits the utility of that weighting scheme, since glossy reflections of caustics are a key strength of VCM. The scheme also performs somewhat worse than the balance heuristic on the HOME OFFICE scene.

More consistent improvements over the balance heuristic can be achieved with variance-aware weighting (c) [GGSK19]. However, the variance estimates this method uses are inaccurate, so it struggles at removing outliers completely, as seen in the MINIMALIST ROOM scene. Additionally, the high computational overhead of variance estimation becomes a concern at larger path lengths. For the MINIMALIST ROOM and HOME OFFICE scenes we cap the path length to five, which is somewhat expensive but still manageable for the variance-aware scheme. However, highly glossy scenes require simulating much longer paths. With up to ten bounces on the ROUGH GLASSES scene, variance-aware weighting takes roughly $4.7 \times$ longer to render than the balance heuristic or our approach.

Our heuristic (d) successfully eliminates outliers and retains glossy reflections at all roughness levels. It never performs worse than the balance heuristic in any of our test scenes. The only case where it does not deliver the best variance reduction is in the region shown in the top zoom-ins of the MINIMALIST ROOM scene. There, the distances and angles between the wall, lamp shade, and light source are similar, and they are all diffuse. The path density does not fully capture the differences in variance.

In Fig. 6 we explore different choices for the radius parameter in our heuristic. We have found our default choice (17) to be close to optimal on most scenes. In very uniformly lit scenes, e.g. TARGET PRACTICE, a larger radius produces slightly better results. There, diffuse bounces from the camera add little variance and should be penalized less. This is a limitation of our heuristic as in those cases a low prefix density does not indicate high variance.

4.2. Bidirectional path tracing

For our second application, we reproduce the correlation problem of BDPT with multiple connections per camera vertex [PRDD15] in a simple setting: we use multiple shadow rays for next-event estimation. In Fig. 7 we compare our heuristic against setting $n = 1$ in the balance heuristic [PRDD15] and the variance-aware weights [GGSK19]. Again, we feed the same set of path samples to all methods. To highlight the problem and make noise more visible, we use 100 shadow rays per camera vertex. Lower counts produce similar results, only less pronounced. Figure 1 shows an example with 10 shadow rays, and Fig. 8 compares different splitting factors. All full-size images are included in the supplemental material.

The balance heuristic produces outliers in the LIVING ROOM




Figure 5: Equal-sample comparison of different MIS heuristics for VCM (ours in bold). We provide error values (relMSE) per crop and over the entire image. The values in parentheses are relative to the balance heuristic (lower is better). In contrast to previous work, our heuristic is consistently better than the balance heuristic, providing significant error reduction in failure cases and slight improvement otherwise.

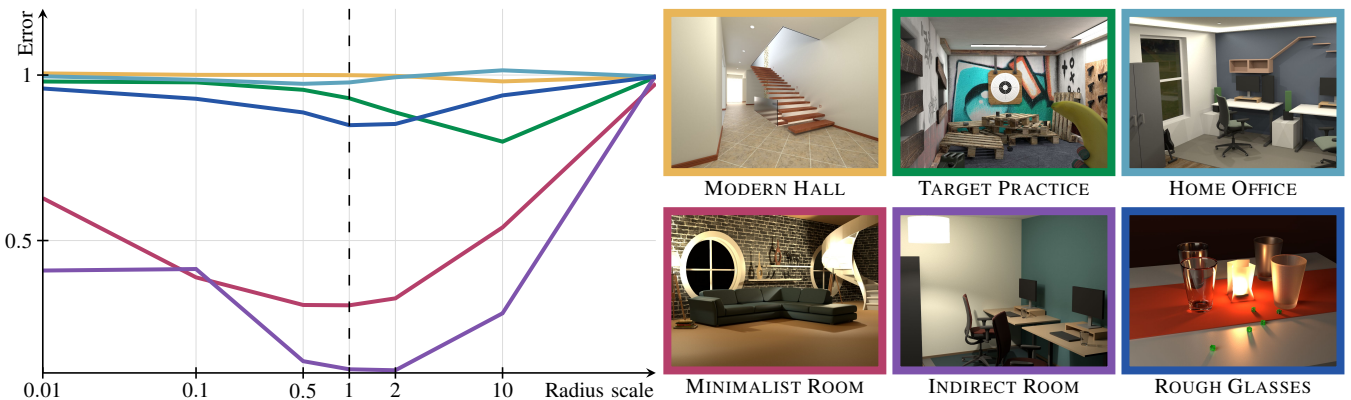


Figure 6: Error relative to the balance heuristic (lower is better) when using different radii for our heuristic. Each line corresponds to one scene. The values on the x-axis are scaling factors applied to the radius computed using Eq. (17). Our default choice (dashed line) is close to the optimal for all scenes except TARGET PRACTICE. That scene features uniform illumination where diffuse bounces from the camera add little variance. A larger radius penalizes such bounces less and improves that specific result, but performs notably worse on all other scenes.

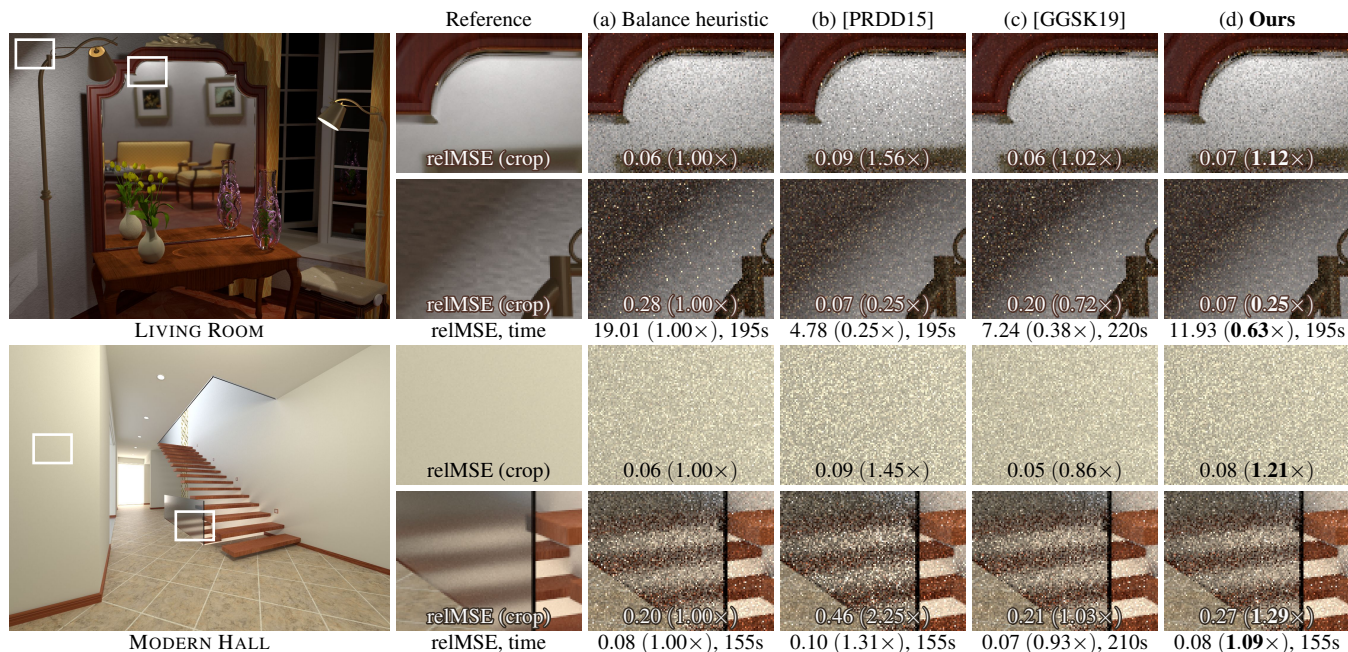


Figure 7: Equal-sample comparison of different MIS heuristics for BDPT with 100 shadow rays for next-event estimation. Although the improvement is smaller in this case compared to VCM in Fig. 5, our heuristic still performs best on average.




Figure 8: Zoom-ins of the INDIRECT ROOM scene from Fig. 1 rendered with different shadow-ray counts. With the balance heuristic, variance increases when more rays are used.

scene due to paths splitting on the lamp shade after bouncing off the diffuse wall (bottom zoom-ins). To eliminate these outliers, Popov et al. [PRDD15] set $n = 1$ in the balance heuristic. Unfortunately, doing so introduces new outliers in the mirror reflection (top zoom-ins). Camera prefixes that bounce off the highly glossy mirror do not increase variance, hence the weight of such paths should not be reduced. The presence of outliers makes the variance estimates in the variance-aware weighting unreliable, preventing it from improving noticeably over the balance heuristic. Our method consistently improves on the balance heuristic over the entire image.

The MODERN HALL scene poses an interesting challenge. Illuminated by numerous light sources from various directions, the variance in the camera prefix is low despite the diffuse interactions. Our heuristic captures most, though not all of that effect. This is the only scene where our heuristic is outperformed by the balance heuristic. Our result is still closer to the balance heuristic than the overly conservative approach of Popov et al. In this scene, the only

method that manages to fully retain the balance-heuristic performance are the more costly variance-aware weights.

5. Limitations and future work

The BDPT application shows the limitations of our simple heuristic. If a scene is dominated by uniform and diffuse illumination, as is the case in the MODERN HALL, the sampling density is a poor indicator of the variance, as zero variance would be achieved with a rather uniform density. That is, a low density can sometimes still yield low variance. By design, our heuristic does not reflect that. Nevertheless, it performs consistently better than the more aggressive solution of Popov et al., while remaining cheap to evaluate.

We have restricted our discussion to surface scattering from finite light sources; however, extending our heuristic to infinite lights or volumetric scattering [KGH*14] should be straightforward. All that is needed is an analogy of the disc approximation used to make the sampling densities unitless and comparable. Applying our ideas to integration problems beyond rendering could also be interesting. However, due to the complex ways correlation affects variance, it is questionable whether a general *and* practical solution is possible.

We do not have hard proofs about the optimality of our heuristic. The result is bounded by the upper bound of Popov et al. [PRDD15], and we only deviate from that if we have reason to believe the correlation does not increase variance. While our empirical results might be sufficient evidence to warrant the use in practice, it is still worthwhile to look for additional guarantees, or alternative heuristics based on a more rigorous mathematical derivation.

Multiple importance sampling in the context of VCM also ig-

nores another important aspect: the bias due to merging. Accounting for this bias in the MIS weights could help reduce artifacts. In that direction, it would be interesting to augment the variance-aware weights by bias estimates [HJJ10], or to make our simpler heuristic bias-aware.

Lastly, having more robust MIS weights opens up possibilities for more creative correlated rendering algorithms. Some prior work has avoided correlation, to also avoid issues with MIS [NID20]. Allowing correlation and utilizing our heuristic could provide additional improvements to such methods. Furthermore, there has been some work on guiding photons based on MIS-weighted contributions [GPSK18; ŠK19]. The efficiency and/or robustness of such applications could also benefit from our heuristic.

6. Conclusion

We propose a simple heuristic to account for correlation due to path splitting in multiple importance sampling for light transport simulation. Our heuristic is efficient and relies on quantities that are already required by the balance heuristic. In contrast to prior work, our approach has no overly harmful effects on the technique combination in cases where the correlation has little or no impact on the variance. Implementing our heuristic takes little effort and introduces negligible overhead.

Acknowledgments

We thank Sebastian Herholz for proof-reading and many writing suggestions. All figures except Figs. 2 and 4 were made with the help of Mira Niemann's [figure generator](#). The MINIMALIST ROOM and MODERN HALL scenes were shared by [Wig42](#) on Blend Swap, and TARGET PRACTICE was modelled by Vladislav Hnatovskiy and Mira Niemann. Above all, we are grateful to the late Jaroslav Krivánek, without whom this project would not have happened.

References

- [AK90] ARVO, JAMES and KIRK, DAVID. "Particle transport and image synthesis". *Proceedings of the 17th annual conference on Computer graphics and interactive techniques*. 1990, 63–66 1, 2.
- [BM97] BOLIN, MARK R and MEYER, GARY W. "An error metric for Monte Carlo ray tracing". *Eurographics Workshop on Rendering Techniques*. Springer. 1997, 57–68 2.
- [CPC84] COOK, ROBERT L, PORTER, THOMAS, and CARPENTER, LOREN. "Distributed ray tracing". *Proceedings of the 11th annual conference on Computer graphics and interactive techniques*. 1984, 137–145 1.
- [EK20] ESTEVEZ, ALEJANDRO CONTY and KULLA, CHRISTOPHER. "Practical Caustics Rendering with Adaptive Photon Guiding". *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks*. 2020, 1–2 3.
- [GGSK19] GRITTMANN, PASCAL, GEORGIEV, ILIYAN, SLUSALLEK, PHILIPP, and KRIVÁNEK, JAROSLAV. "Variance-Aware Multiple Importance Sampling". *ACM Trans. Graph. (SIGGRAPH Asia 2019)* 38.6 (Nov. 17–20, 2019), 152:1–152:9 2–5.
- [GKDS12] GEORGIEV, ILIYAN, KRIVÁNEK, JAROSLAV, DAVIDOVIČ, TOMÁŠ, and SLUSALLEK, PHILIPP. "Light transport simulation with vertex connection and merging." *ACM Trans. Graph.* 31.6 (2012), 192–1 1, 3–5.
- [GPSK18] GRITTMANN, PASCAL, PÉRARD-GAYOT, ARSÈNE, SLUSALLEK, PHILIPP, and KRIVÁNEK, JAROSLAV. "Efficient Caustic Rendering with Lightweight Photon Mapping". *Comput. Graph. Forum (EGSR '18)*. Vol. 37. 4. Wiley Online Library. 2018, 133–142 3, 8.
- [HJJ10] HACHISUKA, TOSHIYA, JAROSZ, WOJCIECH, and JENSEN, HENRIK WANN. "A progressive error estimation framework for photon density estimation". *ACM Transactions on Graphics (TOG)* 29.6 (2010), 1–12 8.
- [HPJ12] HACHISUKA, TOSHIYA, PANTALEONI, JACOPO, and JENSEN, HENRIK WANN. "A path space extension for robust light transport simulation". *ACM Trans. Graph. (TOG)* 31.6 (2012), 191 1, 3, 5.
- [Jen19] JENDERSIE, JOHANNES. "Variance Reduction via Footprint Estimation in the Presence of Path Reuse". *Ray Tracing Gems*. Springer, 2019, 557–569 2, 4, 5.
- [Jen95] JENSEN, HENRIK WANN. "Importance Driven Path Tracing using the Photon Map". *Rendering Techniques*. 1995 1.
- [Jen96] JENSEN, HENRIK WANN. "Global illumination using photon maps". *Rendering Techniques '96*. Springer, 1996, 21–30 1, 3.
- [JG18] JENDERSIE, JOHANNES and GROSCH, THORSTEN. "An Improved Multiple Importance Sampling Heuristic for Density Estimates in Light Transport Simulations." *EGSR (EI&I)*. 2018, 65–72 2, 4, 5.
- [Kaj86] KAJIYA, JAMES T. "The Rendering Equation". *SIGGRAPH Comput. Graph.* 20.4 (Aug. 1986), 143–150. ISSN: 0097-8930 1.
- [KDB14] KELLER, ALEXANDER, DAHM, KEN, and BINDER, NIKOLAUS. "Path space filtering". *ACM SIGGRAPH 2014 Talks*. 2014 1.
- [KGH*14] KRIVÁNEK, JAROSLAV, GEORGIEV, ILIYAN, HACHISUKA, TOSHIYA, et al. "Unifying Points, Beams, and Paths in Volumetric Light Transport Simulation". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 33.4 (July 2014) 7.
- [KVG*19] KONDAPANENI, IVO, VÉVODA, PETR, GRITTMANN, PASCAL, et al. "Optimal Multiple Importance Sampling". *ACM Trans. Graph. (SIGGRAPH 2019)* 38.4 (July 28–Aug. 1, 2019), 37:1–37:14 2.
- [NID20] NABATA, KOSUKE, IWASAKI, KEI, and DOBASHI, YOSHINORI. "Resampling-aware Weighting Functions for Bidirectional Path Tracing Using Multiple Light Sub-Paths". *ACM Trans. Graph. (SIGGRAPH 2020)* 39.2 (2020), 1–11 3, 8.
- [PRDD15] POPOV, STEFAN, RAMAMOORTHY, RAVI, DURAND, FREDO, and DRETTAKIS, GEORGE. "Probabilistic connections for bidirectional path tracing". *Computer Graphics Forum*. Vol. 34. 4. Wiley Online Library. 2015, 75–86 1–5, 7.
- [ŠK19] ŠIK, MARTIN and KRIVÁNEK, JAROSLAV. "Implementing One-Click Caustics in Corona Renderer". (2019) 3, 8.
- [ŠOHK16] ŠIK, MARTIN, OTSU, HISANARI, HACHISUKA, TOSHIYA, and KRIVÁNEK, JAROSLAV. "Robust light transport simulation via metropolised bidirectional estimators". *ACM Trans. Graph. (TOG)* 35.6 (2016), 245 5.
- [Vea97] VEACH, ERIC. *Robust Monte Carlo methods for light transport simulation*. Stanford University PhD thesis, 1997 2–4.
- [VG95a] VEACH, ERIC and GUIBAS, LEONIDAS. "Bidirectional estimators for light transport". *Photorealistic Rendering Techniques*. Springer, 1995, 145–167 1.
- [VG95b] VEACH, ERIC and GUIBAS, LEONIDAS J. "Optimally Combining Sampling Techniques for Monte Carlo Rendering". *SIGGRAPH '95*. ACM. 1995, 419–428 1, 3.
- [VK16] VORBA, JIŘÍ and KRIVÁNEK, JAROSLAV. "Adjoint-driven Russian Roulette and Splitting in Light Transport Simulation". *ACM Trans. Graph.* 35.4 (July 2016), 42:1–42:11. ISSN: 0730-0301 2, 3.
- [VKŠ*14] VORBA, JIŘÍ, KARLÍK, ONDŘEJ, ŠIK, MARTIN, et al. "On-line Learning of Parametric Mixture Models for Light Transport Simulation". *ACM Trans. Graph. (Proceedings of SIGGRAPH 2014)* 33.4 (2014) 1.
- [WGGH20] WEST, REX, GEORGIEV, ILIYAN, GRUSON, ADRIEN, and HACHISUKA, TOSHIYA. "Continuous multiple importance sampling". *ACM Transactions on Graphics (TOG)* 39.4 (2020), 136–1 1.