Image-Space Russian Roulette for Monte Carlo Rendering

GEONU NOH, Gwangju Institute of Science and Technology, South Korea
CHANU YANG, Gwangju Institute of Science and Technology, South Korea
BOCHANG MOON*, Gwangju Institute of Science and Technology, South Korea

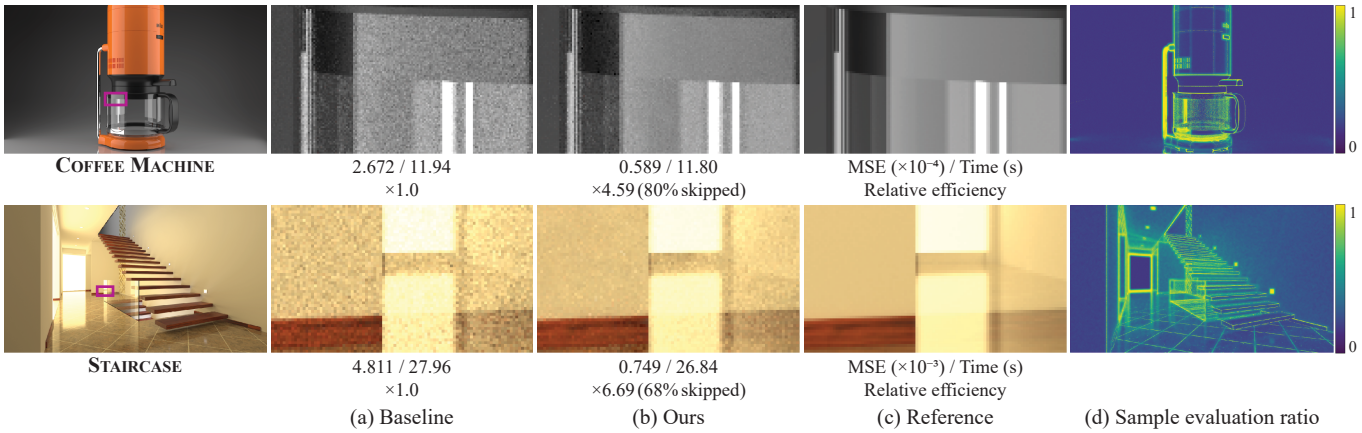

Fig. 1. Equal-time comparison between baseline path tracing with vertex-level Russian roulette (RR) (a) and our method (b). By integrating our unbiased image-space RR on top of baseline path tracing without modifying its underlying vertex-level RR, our approach skips a subset of sample evaluations prior to camera-ray construction while approximating their contributions, achieving lower error under the same rendering time. We measure efficiency as $1/(\text{MSE} \times \text{time})$ and report relative efficiency gains with respect to the equal-time baseline. For the two scenes, our method skips 80% and 68% of sample evaluations, respectively, reducing computation while mitigating RR-induced noise amplification. The per-pixel fraction of evaluated samples (d) visualizes how sample evaluations are spatially reduced. 3D model courtesy of cekuhnen (COFFEE MACHINE) and NewSee2l035 (STAIRCASE).

Russian roulette (RR) is a widely used technique for reducing computational cost in Monte Carlo (MC) rendering, which often requires a large number of samples to achieve low-noise results. In standard practice, RR probabilistically terminates suffix paths at path vertices when their expected contribution is low, returning a predefined termination value (e.g., zero) in place of the skipped contribution. In this paper, we introduce an image-space RR framework that operates directly in image space, enabling early termination of entire light paths before ray traversal. We generalize RR by introducing sample-dependent termination values, allowing skipped contributions to be approximated rather than discarded. To make this formulation practical, we exploit spatial correlation through a seed-shifting strategy that reuses primary samples across neighboring pixels. This enables efficient construction of both termination values and termination probabilities using correlated evaluations. We demonstrate that our method achieves lower error than baseline path tracing under equal-time conditions across a variety of scenes while remaining unbiased and compatible with existing techniques such as vertex-level RR.

*Corresponding author.

Authors' Contact Information: Geonu Noh, skarudxor2@gm.gist.ac.kr, Gwangju Institute of Science and Technology, Gwangju, South Korea; Chanu Yang, yangchanu@gm.gist.ac.kr, Gwangju Institute of Science and Technology, Gwangju, South Korea; Bochang Moon, bmoon@gist.ac.kr, Gwangju Institute of Science and Technology, Gwangju, South Korea.

This work is licensed under a Creative Commons Attribution 4.0 International License.
SA Conference Papers '26, Kuala Lumpur, Malaysia
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2842-6/2026/12
https://doi.org/10.1145/3829340.3842204

CCS Concepts: • **Computing methodologies → Ray tracing**.

Additional Key Words and Phrases: Monte Carlo rendering, global illumination, variance reduction, Russian roulette, image-space techniques

ACM Reference Format:
Geonu Noh, Chanu Yang, and Bochang Moon. 2026. Image-Space Russian Roulette for Monte Carlo Rendering. In *SIGGRAPH Asia 2026 Conference Papers (SA Conference Papers '26), December 01–04, 2026, Kuala Lumpur, Malaysia.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3829340.3842204

1 Introduction

Monte Carlo (MC) rendering approximates the light transport equation by stochastically simulating light paths formed through recursive ray generation. While conceptually simple, achieving low-noise results requires a large number of samples, leading to substantial computational cost. Improving efficiency, often measured as the inverse of the product of MSE and rendering time, remains a central challenge in rendering.

A widely used technique for improving efficiency is Russian roulette (RR), which stochastically terminates recursive ray generation in an unbiased manner. In practice, RR is typically applied at path vertices, where additional rays are terminated when the expected contribution of suffix paths is low. This process involves two components: a termination probability, which determines whether to continue tracing, and a termination value, which replaces the skipped contribution. Existing approaches commonly use non-adaptive termination values (e.g., zero [Rath et al. 2022;

Vorba and Křivánek 2016] or a globally estimated incoming radiance [Szécsi et al. 2003; Szirmay-Kalos et al. 2005]) and focus on designing effective termination probabilities using local or estimated path contributions.

While vertex-level RR has been extensively studied and proven effective, it is limited to decisions made after path generation begins. This raises a natural question: can RR be applied earlier in image space to skip entire light paths before traversal? Such an image-space extension offers the potential for additional efficiency gains, since entire paths can be avoided before any ray traversal. At the same time, it remains compatible with vertex-level RR for surviving paths.

However, designing a practical image-space RR is challenging. Unlike vertex-level RR, image-space decisions must be made with limited information. In addition, the contribution of a light path can be arbitrarily large, making such early termination risky. As a result, naïvely applying RR at this level leads to severe variance amplification, since surviving samples must be reweighted by the inverse survival probability to maintain unbiasedness.

These challenges motivate us to revisit the design of RR from a different perspective. Prior work has primarily focused on optimizing the termination probability while relying on non-adaptive termination values. In contrast, we generalize RR by introducing sample-dependent termination values that approximate skipped path contributions, unlike conventional RR, which discards them using fixed termination values (e.g., zero).

To make this formulation practical, we exploit spatial coherence in MC rendering through a seed-shifting strategy that reuses primary samples across neighboring pixels. This produces correlated path evaluations for constructing adaptive termination values and probabilities. Our contributions are as follows:

- We propose a generalized image-space RR formulation with sample-dependent termination values, enabling unbiased stochastic skipping of light-path evaluations before any ray traversal.
- We introduce a seed-shifting strategy with shared primary samples that exploits spatial correlation to construct adaptive termination values and probabilities.
- We further refine the estimator using a control-variate formulation that reduces variance while retaining the computational savings of image-space RR.
- We demonstrate that our unbiased image-space RR improves rendering efficiency and is complementary to existing techniques such as vertex-level RR and path guiding.

2 Related Work

RR has become a fundamental component of efficient MC rendering algorithms, including light transport simulation [Tokuyoshi and Harada 2019; Veach 1997], volumetric rendering [Herholz et al. 2019; Kettunen et al. 2021], and advanced sampling strategies [Bauszat et al. 2017; Meyer et al. 2024]. In this section, we focus on vertex-level RR methods that are most closely related to our approach.

Vertex-level RR methods determine whether to continue path traversal at each vertex so that unimportant suffix paths can be skipped in an unbiased manner to balance variance and computational cost [Pharr et al. 2016]. A central problem in these methods is the design of the termination probability, which should be high for low-contribution paths. This probability is typically adjusted using local information at the vertex, such as albedo [Arvo and Kirk 1990] or path throughput [Pharr et al. 2016; Szirmay-Kalos et al. 2005]. More advanced approaches estimate the expected incoming radiance at a vertex and optimize the probability or splitting factor for improved efficiency [Rath et al. 2022; Vorba and Křivánek 2016].

Although extensive work has focused on designing effective termination probabilities, comparatively less attention has been paid to the design of another parameter: the termination value. A common choice is to use a global value, often set to zero, implicitly assuming that the remaining contribution of a terminated path is negligible [Pharr et al. 2016; Rath et al. 2022; Vorba and Křivánek 2016]. Alternatively, prior work has proposed using a global approximation of incoming radiance as the termination value [Szécsi et al. 2003; Szirmay-Kalos et al. 2005].

We revisit the design of the termination value and generalize it to a data-dependent formulation that varies across samples. By adaptively designing both the termination probability and termination value in image space, our method improves the efficiency of MC rendering. Our image-space RR is complementary to techniques applied at path vertices, including vertex-level RR, neural two-level Monte Carlo [Dereviannykh et al. 2025], and neural control variates [Müller et al. 2020], as these methods remain applicable to light paths that survive our RR decisions.

Path guiding. Path guiding is another widely used variance reduction technique that performs importance sampling at path vertices by adapting the sampling distribution of secondary rays to the incoming radiance. Various approaches approximate this radiance such as tree structures [Diolatzis et al. 2020; Müller et al. 2017], a voxel-based method [Lu et al. 2024], an adaptive octree for focal-light effects [Rath et al. 2023], mixture models [Dodik et al. 2022; Herholz et al. 2016; Ruppert et al. 2020], and neural networks [Dong et al. 2023; Huang et al. 2024; Rath et al. 2025; Zhu et al. 2021]. Similar to vertex-level RR, path guiding techniques can be applied to light paths that survive our RR decisions.

Image-space techniques in rendering. Exploiting spatial coherence among nearby pixels has been extensively studied in rendering. A common approach reduces MC noise by filtering pixel estimates using their spatial neighborhood through carefully designed filters [Zwicker et al. 2015], including classical MSE-based optimization [Li et al. 2012; Moon et al. 2014; Sakai et al. 2024, 2025] and learning-based methods [Bako et al. 2017; Balint et al. 2023; Choi et al. 2024; Wang et al. 2024; Yu et al. 2021].

A more aggressive approach explicitly enforces statistical correlation across nearby pixels through correlated sampling and constructs pixel estimators that combine neighboring samples, such as gradient-domain rendering [Kettunen et al. 2015; Lehtinen et al. 2013; Yan et al. 2025], spatiotemporal reservoir resampling [Bitterli et al. 2020; Lin et al. 2022; Zhang et al. 2024], pixel combiners [Back et al. 2020, 2023], and image-space control variates [Rousselle et al. 2016; Yang and Moon 2025]. We also leverage correlation among nearby pixels induced by shared random numbers. However, unlike

these methods, we incorporate this correlation into an RR estimator to skip expensive path evaluations.

Adaptive sampling. Image-space adaptive sampling allocates different numbers of samples across pixels according to estimated errors [Firmino et al. 2023; Moon et al. 2014; Rousselle et al. 2011; Sakai et al. 2025; Salehi et al. 2022]. Although adaptive sampling and our image-space RR can reduce computation in low-error image regions, they are complementary. Specifically, our approach keeps the number of RR trials per pixel uniform and adapts the computational cost within each trial: a trial either evaluates a complete light path or substitutes its contribution with a termination value. Thus, adaptive sampling determines where additional trials are allocated, whereas image-space RR determines which path evaluations are performed within those trials.

3 Background

In this section, we review the pixel estimator in the primary sample space (PSS) and define the problem setting. The rendering equation [Kajiya 1986] can be expressed in path-space form [Veach 1997] as

$$I_i = \int_{\mathcal{X}} f_i(x)\, d\mu(x), \tag{1}$$

where I_i denotes the value of pixel i, $x \in \mathcal{X}$ represents a light path, and f_i is the measurement contribution function.

Following Kelemen et al. [2002], this path integral can be rewritten over the primary sample space $\mathcal{U}$ as

$$I_i = \int_{\mathcal{U}} \frac{f_i(S(u))}{p(S(u))}\, du, \tag{2}$$

where $u \in \mathcal{U}$ is a vector of uniform random numbers, $S : \mathcal{U} \to \mathcal{X}$ maps primary samples to light paths, and $p(S(u))$ is the probability density of the path generated from u. A Monte Carlo (MC) estimator for I_i is given by

$$\hat{I}_i = \frac{1}{N} \sum_{n=1}^{N} g_i(u_i^n), \tag{3}$$

where $g_i(u) := \frac{f_i(S(u))}{p(S(u))}$ and u_i^n is the n-th primary sample for pixel i.

Problem specification. The estimation error $|\hat{I}_i - I_i|$ of the standard MC estimator $\hat{I}_i$ decreases at a rate of $O(1/\sqrt{N})$. Achieving acceptable noise levels therefore requires evaluating $g_i(u)$ for a large number of samples, which dominates the computational cost. Our goal is to reduce computational cost using an image-space RR estimator that skips the evaluation of $g_i(u)$ for a subset of samples in an unbiased manner before path traversal, while treating g_i as a black-box function. This formulation enables integration with existing optimizations, such as vertex-level RR, applied to the surviving light paths.

4 Image-Space Russian Roulette

In this section, we first examine a naïve extension of RR to image space and analyze its limitations in Sec. 4.1. We then introduce a generalized formulation with sample-dependent termination values in Sec. 4.2 and study how its variance depends on the termination value design in Sec. 4.3.

4.1 Naïve Image-Space Russian Roulette

A straightforward extension of RR to image space is to apply it directly to the pixel estimator $g_i(u)$ (Eq. 3) using a constant termination value C_i for the pixel i, as

$$\langle I_i^{RR;C_i} \rangle = \begin{cases} \frac{g_i(u) - q_i(u) C_i}{1 - q_i(u)} & \text{if } \zeta > q_i(u), \\ C_i & \text{otherwise,} \end{cases} \tag{4}$$

where $q_i(u)$ is the termination probability and ζ is an independent random variable uniformly distributed in $[0, 1)$.

To understand the behavior of the termination value C_i in this simple RR design, we examine two representative choices of C_i using a fixed termination probability on top of a path tracer, as shown in Fig. 2.

With the zero termination value ($C_i = 0$), which corresponds to the standard choice, image-space RR significantly amplifies variance. At the image-space level, the implicit assumption that the terminated sample has a negligible contribution is difficult to satisfy, since the skipped quantity corresponds to the contribution of an entire light path, unlike vertex-level RR, which considers only a suffix path. This estimator evaluates $g_i(u)$ for only a fraction $1 - q_i(u)$ of the samples on average. However, the contributions of the surviving samples are amplified by a factor of $1/(1 - q_i(u))$, leading to large variance.

An ideal choice for the per-pixel constant is to set $C_i = I_i$, where the ground truth pixel value I_i is assumed to be known for this experiment. As shown in Fig. 2, although this choice significantly reduces variance compared to the zero-termination case, it still does not provide a clear improvement over the baseline without image-space RR.

These observations indicate that the termination value plays a critical role in image-space RR, and constant-value approximations are insufficient for achieving meaningful efficiency gains. This motivates generalizing the termination value from a constant C_i to a sample-dependent function.

4.2 Generalized Image-Space Russian Roulette

We generalize the image-space RR estimator in Eq. 4 by replacing the constant termination value C_i with a sample-dependent function $\phi_i(u)$, resulting in

$$\langle I_i^{RR;\phi_i} \rangle = \begin{cases} \frac{g_i(u) - q_i(u) \phi_i(u)}{1 - q_i(u)} & \text{if } \zeta > q_i(u), \\ \phi_i(u) & \text{otherwise,} \end{cases} \tag{5}$$

where the constant C_i in Eq. 4 is replaced with a sample-dependent termination value $\phi_i(u)$. The simple RR estimator (Eq. 4) is a special case obtained by setting $\phi_i(u) = C_i$.

Let us rewrite the generalized RR estimator (Eq. 5) in an equivalent but more compact form using a Bernoulli random variable $B \in \{0, 1\}$ with success probability $1 - q_i(u)$ as

$$\langle I_i^{RR;\phi_i} \rangle = \phi_i(u) + \frac{B}{1 - q_i(u)} \big(g_i(u) - \phi_i(u)\big), \tag{6}$$

which provides a more convenient form for analysis and will be used throughout the following sections.

The key advantage of this generalized formulation (Eqs. 5 and 6) is that it enables reducing the per-sample residual $g_i(u) - \phi_i(u)$, which

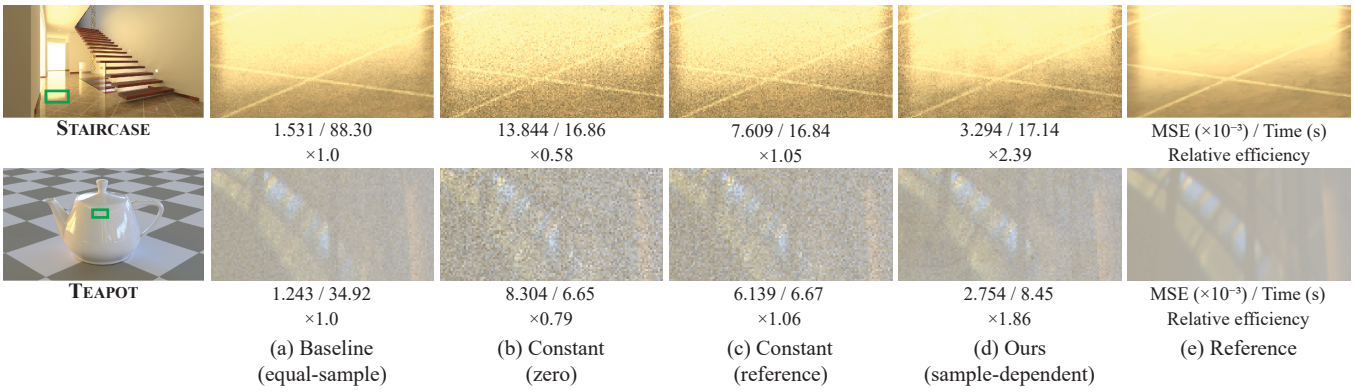

Fig. 2. Effect of termination value design in image-space RR under an equal-sample setting. All methods, including the baseline (a) and RR variants (b–d), use the same number of samples, while RR skips a subset of evaluations. This equal-sample comparison uses a fixed 80% skip rate for camera-ray evaluations to isolate the variance induced by the termination-value design; Sec. 5 introduces the adaptive termination probability used in our full method. For constant termination values, we consider (b) the standard zero value and (c) an ideal but impractical choice using the reference pixel value. Our sample-dependent termination value (d) mitigates the variance amplification of RR compared to constant designs, enabling practical efficiency gains. 3D model courtesy of NewSee2l035 (STAIRCASE) and Benedikt Bitterli (TEAPOT).

directly controls the noise amplification by the factor $1/(1-q_i(u))$. By allowing $\phi_i(u)$ to vary across samples, the termination value can be designed to be positively correlated with $g_i(u)$, thereby reducing the residual.

We analyze this estimator in Sec. 4.3 and present a practical design of $\phi_i(u)$ together with $q_i(u)$ in Sec. 5.

4.3 Analysis

We analyze the unbiasedness and variance of the generalized image-space RR estimator, revealing how the termination value influences the additional variance introduced by RR.

The conditional expectation of the generalized RR estimator in Eq. 6 is given by

$$\mathbb{E}[\langle I_i^{RR;\phi_i}\rangle \,|\, u] = \phi_i(u) + \big(g_i(u) - \phi_i(u)\big) = g_i(u), \tag{7}$$

leading to $\mathbb{E}[\langle I_i^{RR;\phi_i}\rangle] = \mathbb{E}[g_i(u)] = I_i$. Thus, the estimator remains unbiased for arbitrary choices of $\phi_i(u)$ and $q_i(u)$.

Also, its conditional variance is given by

$$\mathbb{V}[\langle I_i^{RR;\phi_i}\rangle \mid u] = \frac{q_i(u)}{1-q_i(u)}(g_i(u)-\phi_i(u))^2. \tag{8}$$

Using the law of total variance together with Eqs. 7 and 8, the variance of the estimator (Eq. 6) can be written as

$$\begin{aligned}\mathbb{V}[\langle I_i^{RR;\phi_i}\rangle] &= \mathbb{V}[\mathbb{E}[\langle I_i^{RR;\phi_i}\rangle \,|\, u]] + \mathbb{E}[\mathbb{V}[\langle I_i^{RR;\phi_i}\rangle \mid u]] \\ &= \mathbb{V}[g_i(u)] + \mathbb{E}\left[\frac{q_i(u)}{1-q_i(u)}(g_i(u)-\phi_i(u))^2\right],\end{aligned} \tag{9}$$

which consists of the variance of the original estimator and the additional variance introduced by RR. This decomposition shows that the additional variance is proportional to the squared residual $(g_i(u)-\phi_i(u))^2$.

To gain further insight into the termination value, suppose the termination probability is constant, i.e., $q_i(u) = q$. For a constant termination value design, i.e., $\phi_i(u) = C_i$, the additional variance term becomes

$$\frac{q}{1-q}\mathbb{E}\left[(g_i(u)-C_i)^2\right] = \frac{q}{1-q}\left(\mathbb{V}[g_i(u)] + (\mathbb{E}[g_i(u)]-C_i)^2\right). \tag{10}$$

The optimal constant termination value that minimizes this term is given by $C_i^* = \mathbb{E}[g_i(u)] = I_i$, leading to

$$\mathbb{V}[\langle I_i^{RR;C_i^*}\rangle] = \frac{1}{1-q}\mathbb{V}[g_i(u)]. \tag{11}$$

Hence, even with the optimal constant termination value, the ordinary MC variance is amplified by the factor of $1/(1-q)$. This makes constant-value RR particularly ineffective in image space, in contrast to vertex-level RR, where suffix path contributions, which are often smaller in magnitude, typically exhibit lower variance than the full path contribution $g_i(u)$.

In contrast, our generalized design introduces an additional degree of freedom through $\phi_i(u)$, enabling reduction of the residual term and lowering RR-induced variance even under a constant termination probability. In particular, when $q_i(u) = q$, comparing Eq. 9 and Eq. 11 shows that variance can be reduced when the residual magnitude $|g_i(u) - \phi_i(u)|$ is small. We construct such a sample-dependent termination value in the next section.

5 Practical Image-Space Russian Roulette

In this section, we describe a practical image-space RR estimator. Based on the analysis in Sec. 4.3, an effective design should (i) construct termination values that are positively correlated with $g_i(u)$, and (ii) adapt termination probabilities based on the residual $g_i(u) - \phi_i(u)$.

Fig. 3 illustrates our two-stage pipeline, which is performed at every rendering iteration. At each iteration, we independently generate one primary sample per pixel and evaluate the corresponding 1-spp image together with four diagonally shifted evaluations. All five path evaluations contribute to the final estimator, and their costs are included in all reported rendering times; the pre-pass is

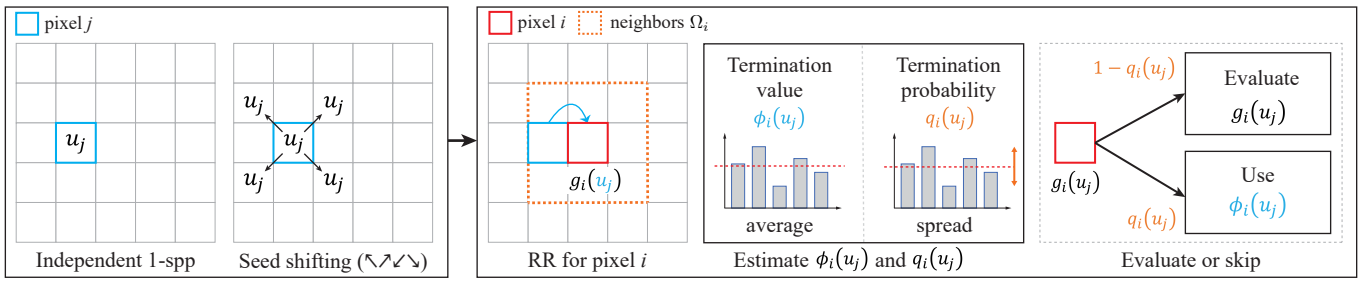

Fig. 3. Overview of our image-space RR framework. The pipeline consists of two main stages. (a) Pre-pass: we render a 1-spp image using independent sampling, assigning each pixel j a primary sample u_j, which is then reused at spatially shifted pixel locations (seed shifting) to generate correlated evaluations. (b) Image-space RR: for each pixel i, primary samples from neighboring pixels are reused to construct termination values $\phi_i(u_j)$ and probabilities $q_i(u_j)$, which guide per-sample decisions to evaluate $g_i(u_j)$ or use $\phi_i(u_j)$. The final pixel estimate is obtained by averaging contributions from both stages.

therefore not discarded overhead. We then apply image-space RR using the resulting correlated samples. The primary samples, seed shifts, and RR decisions are regenerated independently at every iteration, and the final image is obtained by averaging the per-iteration estimates.

5.1 Seed Shifting

We employ a simple seed-shifting scheme, where the same primary sample is reused at spatially shifted pixel locations. Let $D = \{(-\Delta, -\Delta), (+\Delta, -\Delta), (-\Delta, +\Delta), (+\Delta, +\Delta)\}$ denote a fixed set of diagonal offsets (see Fig. 3). The seed-shift size is set to $\Delta = 8$.

We first assign a unique random seed to each pixel and render a 1-spp image, where each pixel j is associated with a primary sample u_j. We then reuse the same primary sample u_j (i.e., identical random numbers) at shifted pixel locations by sharing the same random numbers according to the offsets in D. This produces multiple correlated evaluations of the same sample across nearby pixels.

For each pixel i, we construct an image-space RR estimator that operates by reusing each primary sample u_j of neighboring pixel $j \in \Omega_i$, where Ω_i denotes a local neighborhood centered at pixel i. For Ω_i, we randomly select 50% of pixels within a 15×15 window centered at pixel i, unless otherwise specified.

Specifically, the RR estimator is defined as

$$\hat{I}_i^{RR} = \frac{1}{|\Omega_i|} \sum_{j \in \Omega_i} \left[\phi_i(u_j) + \frac{B_j}{1 - q_i(u_j)} \left(g_i(u_j) - \phi_i(u_j) \right) \right]. \quad (12)$$

For each sample u_j, the contribution $g_i(u_j)$ is evaluated only when $B_j = 1$; otherwise, it is approximated by $\phi_i(u_j)$. This formulation corresponds to a multi-sample estimator of the generalized RR estimator (Eq. 6), where each primary sample u_j is reused from a neighboring pixel j.

We describe how to construct the termination value $\phi_i(u_j)$ and termination probability $q_i(u_j)$ in the subsequent sections.

5.2 Adaptive Termination Value via Seed Shifting

A key component of the estimator (Eq. 12) is the termination value $\phi_i(u_j)$, which approximates the contribution $g_i(u_j)$ when the sample is skipped.

To construct $\phi_i(u_j)$, we exploit the seed-shifting scheme. Since the same primary sample u_j is reused at multiple shifted pixel locations, we obtain several correlated evaluations of this sample across nearby pixels. We define

$$\phi_i(u_j) = \frac{1}{W_i^j} \sum_{k \in K_j} w_{i,k} g_k(u_j), \quad (13)$$

where the normalization factor $W_i^j = \sum_{k \in K_j} w_{i,k}$, and K_j denotes the set of shifted pixels associated with sample u_j (see Fig. 3 (a)).

To compute the weight $w_{i,k}$, we employ joint bilateral weighting:

$$w_{i,k} = \exp\left(-\frac{\|n_i - n_k\|^2}{b_n^2} - \frac{\|\rho_i - \rho_k\|^2}{b_\rho^2} \right), \quad (14)$$

where n_i and ρ_i denote the normal and albedo at pixel i, respectively. We set the bandwidths to $b_n = b_\rho = 0.1$. We compute the G-buffers for the weights once prior to the rendering stage.

When the shifted pixels associated with u_j have substantially different G-buffer features from pixel i (e.g., when $W_i^j < 0.1$), the estimate $\phi_i(u_j)$ becomes unreliable. We then fall back to standard MC by evaluating $g_i(u)$ using an independently drawn primary sample u. Reusing u_j would also remain unbiased, but we use this fresh sample to avoid retaining unnecessary spatial correlation. The fallback occurs infrequently (e.g., approximately 4.6% of reused samples on average for the two scenes in Fig. 2).

Since $\phi_i(u_j)$ is constructed from the same primary sample u_j as the target contribution $g_i(u_j)$, it is therefore expected to be positively correlated with $g_i(u_j)$, reducing the residual $g_i(u_j) - \phi_i(u_j)$. As shown in Fig. 2, this sample-dependent construction effectively mitigates noise amplification compared to constant termination values. However, its effectiveness decreases in regions with low inter-pixel correlation (e.g., near edges), which motivates adapting the termination probability in the next section.

5.3 Adaptive Termination Probability via Seed Shifting

The termination probability $q_i(u_j)$ (in Eq. 12) controls how often the true contribution $g_i(u_j)$ is evaluated or approximated by $\phi_i(u_j)$.

From the variance analysis (Eq. 9), the additional variance introduced by RR increases with the squared residual $(g_i(u_j) - \phi_i(u_j))^2$.

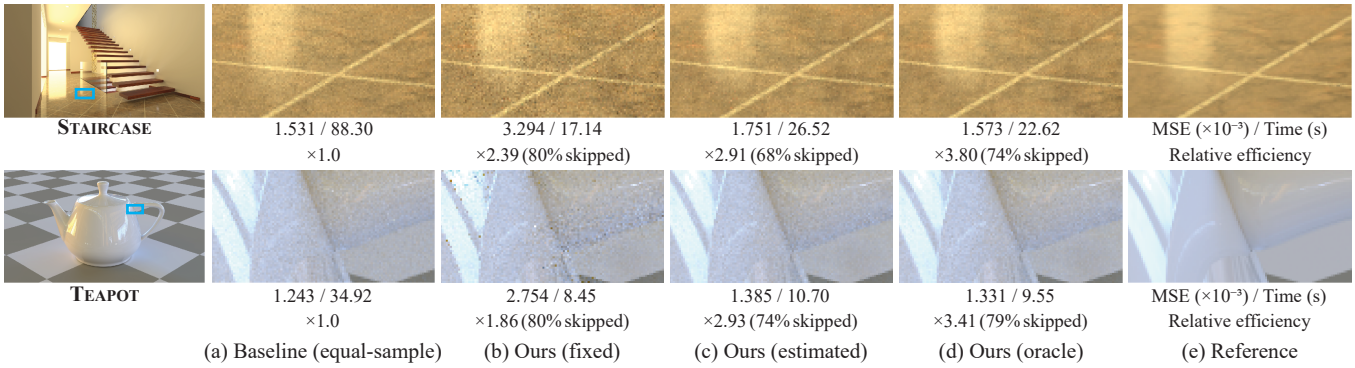

Fig. 4. Equal-sample comparison of image-space RR under different termination probability designs. The baseline path tracing result is shown in (a), and RR variants (b–d) use our sample-dependent termination values. We compare (b) a fixed termination probability, (c) our adaptive probability estimated via seed shifting, and (d) an oracle based on the true residual, i.e., the squared difference between the true contribution and the termination value. The oracle assumes access to the true contribution at no computational cost and is therefore not achievable in practice. While a gap remains between our method and the oracle, our adaptive design effectively reduces noise amplification compared to the fixed-probability case. 3D model courtesy of NewSee2l035 (STAIRCASE) and Benedikt Bitterli (TEAPOT).

This suggests that samples with larger residuals should be evaluated more frequently, whereas samples whose termination values closely approximate their contributions can be skipped more aggressively.

To approximate this residual magnitude, we again use correlated evaluations obtained via seed shifting. Specifically, we approximate the unknown residual using neighboring samples by replacing $g_i(u_j)$ with $g_k(u_j)$:

$$(g_i(u_j) - \phi_i(u_j))^2 \approx \xi_i(u_j) = \frac{1}{|K_j|} \sum_{k \in K_j} \left(g_k(u_j) - \phi_i(u_j)\right)^2. \quad (15)$$

We then set the survival probability as

$$1 - q_i(u_j) = \text{clamp}\left(\sqrt{\xi_i(u_j)},\ 0.1,\ 1.0\right). \quad (16)$$

Since $\xi_i(u_j)$ approximates the squared residual, we use $\sqrt{\xi_i(u_j)}$ as a proxy for its magnitude when setting the survival probability. We also enforce $1 - q_i(u_j) \geq 0.1$ to avoid excessive variance amplification due to the factor $q_i(u_j)/(1-q_i(u_j))$ (Eq. 9), and $1-q_i(u_j) \leq 1.0$ to ensure valid probabilities.

Because $\sqrt{\xi_i(u_j)}$ scales with scene radiance, the heuristic in Eq. 16 is not invariant to global changes in light intensity, which can alter the survival probability before clamping. Deriving a more principled termination-probability criterion that maximizes rendering efficiency while accounting for radiance scale is left for future work.

Fig. 4 compares three choices of termination probability: a fixed value, our estimated probability, and an oracle with the true residual $(g_i(u_j) - \phi_i(u_j))^2$. While the MSE gap between our method and the oracle is noticeable, our simple heuristic via seed shifting effectively reduces noise amplification compared to the fixed-probability case.

5.4 Connection to Control Variates and Refinement

To further reduce variance, we refine the estimator (Eq. 12) by modifying the first occurrence of the termination value. The second occurrence of $\phi_i(u_j)$ must depend on the same primary sample u_j to preserve correlation with $g_i(u_j)$. In contrast, the first term does not affect the residual and can therefore be replaced with a less noisy estimate.

We thus replace it with a less noisy estimate $\bar{\phi}_i(u_j)$, yielding

$$\hat{I}_i^{RR} = \frac{1}{|\Omega_i|} \sum_{j \in \Omega_i} \left[\bar{\phi}_i(u_j) + \frac{B_j}{1 - q_i(u_j)} \left(g_i(u_j) - \phi_i(u_j)\right)\right]. \quad (17)$$

We define $\bar{\phi}_i(u_j)$ as

$$\bar{\phi}_i(u_j) := \frac{1}{W_i^j} \sum_{k \in K_j} w_{i,k} \bar{g}_k, \quad (18)$$

where $\bar{g}_k$ is the average MC estimate at pixel k obtained from the pre-pass stage. Since both $\bar{g}_k$ and the estimator $g_k(u_j)$ in Eq. 13 are unbiased estimates of the ground truth I_k for pixel k and $w_{i,k}$ (Eq. 14) is determined only from G-buffers (computed before the sampling stage), $\mathbb{E}[\bar{\phi}_i(u_j)] = \mathbb{E}[\phi_i(u_j)]$.

Analysis. Since the refinement does not modify the residual term multiplied by the Bernoulli variable, the RR-induced conditional variance remains identical to Eq. 8. In contrast, the conditional expectation (Eq. 7) becomes $\bar{\phi}_i(u_j) + g_i(u_j) - \phi_i(u_j)$. Nonetheless, the modified estimator remains unbiased because its unconditional expectation is unchanged due to $\mathbb{E}[\bar{\phi}_i(u_j)] = \mathbb{E}[\phi_i(u_j)]$.

Applying the law of total variance, the per-sample variance of the refined estimator is given by

$$\mathbb{V}\left[\bar{\phi}_i(u_j) + g_i(u_j) - \phi_i(u_j)\right] + \mathbb{E}\left[\frac{q_i(u_j)}{1 - q_i(u_j)} (g_i(u_j) - \phi_i(u_j))^2\right], \quad (19)$$

where the second term corresponds to the RR-induced variance analyzed in Sec. 4.3 and remains unchanged (see Eq. 9), while the first term corresponds to the variance of the refined base estimator. This variance term can be reduced when the residual $g_i(u_j) - \phi_i(u_j)$ is small, unlike the original variance term $\mathbb{V}[g_i(u_j)]$ in Eq. 9.

Relation with Control Variates. This refinement (Eq. 17) can be interpreted as introducing a control variate into the base estimator

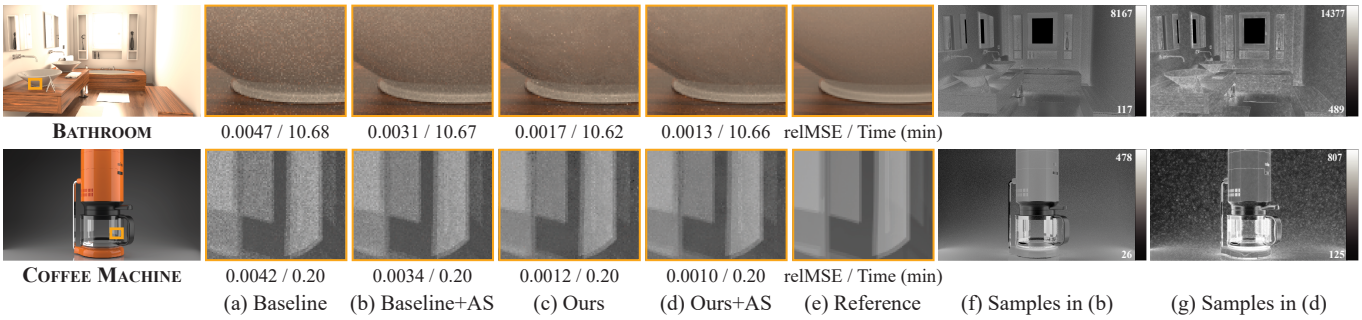

Fig. 5. Equal-time comparisons of baseline path tracing and our method with and without adaptive sampling (AS). The corresponding per-pixel sample-count maps for Baseline+AS and Ours+AS are shown in (f) and (g), respectively. 3D model courtesy of nacimus (Bathroom) and cekuhnen (Coffee Machine).

(Eq. 12). In particular, when no samples are skipped (i.e., $q_i(u_j) = 0$), the refined estimator reduces to

$$\hat{I}_i^{RR} = \frac{1}{|\Omega_i|} \sum_{j \in \Omega_i} \left[g_i(u_j) - \phi_i(u_j) + \bar{\phi}_i(u_j) \right], \tag{20}$$

where $\phi_i(u_j)$ acts as a control variate and $\bar{\phi}_i(u_j)$ serves as a less noisy unbiased estimate of $\mathbb{E}[\phi_i(u_j)]$. This refinement is related to image-space control variates [Yang and Moon 2025], which combine neighboring estimates using data-dependent coefficients and consequently introduce bias. In contrast, our method embeds an image-space control variate with a fixed coefficient of one into RR, enabling unbiased stochastic skipping of expensive path evaluations. As shown in Fig. 7, this combination improves efficiency by reducing variance through the control-variate refinement and computational cost through sample skipping.

6 Results and Discussion

In this section, we evaluate our method using the path tracer in Mitsuba [Jakob 2010], with next-event estimation and standard throughput-based RR. All experiments are conducted on an AMD Ryzen 9 9950X at full-HD resolution (1920×1080), with errors computed against reference images rendered with large sample counts.

6.1 Compatibility with the Baseline

Fig. 9 shows MSE convergence for baseline path tracing (PT) and our method across ten scenes. Although the efficiency gain varies across scenes, our method achieves lower MSE at equal times than baseline PT, except in the Glass of Water scene.

Fig. 8 presents qualitative comparisons under equal-time and equal-sample settings. Under equal-time conditions, our method produces more accurate images than the baseline due to improved efficiency (up to 14.13× for the Foggy Bathroom). Compared to the equal-sample baseline, our method evaluates only a subset of light paths through image-space RR, yet achieves comparable or even lower error in several scenes. This behavior reflects the combined effect of reduced per-sample cost from image-space RR and reduced variance from the control variate refinement (Sec. 5.4).

6.2 Compatibility with Optimized Vertex-Level Methods

We evaluate the compatibility of our method with two vertex-level optimizations: (1) practical path guiding (PPG [Müller et al. 2017]) and (2) efficiency-aware RR and splitting (EARS [Rath et al. 2022]). We observe that both methods tend to exhibit larger relative efficiency gains over baseline PT when efficiency is computed using relative MSE (relMSE) rather than MSE. Therefore, we report both metrics together with their corresponding relative efficiency gains.

As shown in Fig. 10, PPG improves rendering efficiency over baseline PT through optimized path sampling, whereas EARS does so through per-vertex RR and splitting. Our method further improves the efficiency of both methods by reducing per-sample cost through stochastic skipping of expensive light-path evaluations in image space, demonstrating its complementarity with existing vertex-level rendering optimizations.

6.3 Compatibility with Adaptive Sampling

Fig. 5 presents a proof-of-concept combination of conventional adaptive sampling (AS) with baseline PT and our method. In both cases, AS estimates per-pixel relative variance and iteratively allocates samples accordingly. For our method, it adjusts the number of RR trials, $|\Omega_i|$ in Eq. 17, instead of using a fixed neighborhood sampling ratio of 50%. AS lowers relMSE for both methods by concentrating samples in high-error pixels. Because the allocations are estimated from the same samples, both implementations are biased; unbiased integration of AS with our RR framework remains future work.

6.4 Parameter Analysis

Our method has two main parameters: the seed-shift size Δ and the neighborhood sampling ratio r within a fixed 15×15 neighborhood defining Ω_i used in Eq. 17. We observe that the relative efficiency gain over baseline PT remains stable across a broad range of parameter values. Across the five scenes shown in Fig. 8, the average efficiency gain varies by approximately 3% for $\Delta = 4$–12 and 13% for $r = 30\%$–70%. Accordingly, we use $\Delta = 8$ and $r = 50\%$ by default.

6.5 Limitations and Future Work

A primary limitation of our method is its reliance on inter-pixel correlation induced by shared primary samples. When such correlation is weak, the approximation quality of the termination value

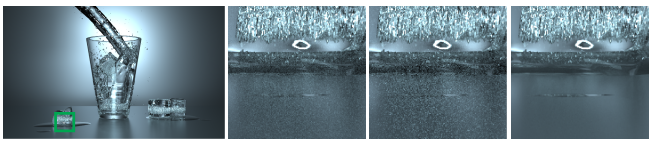

(a) Baseline (b) Ours (c) Reference

Fig. 6. A failure case for our method. In the insets, reflection from the floor followed by further reflections and refractions among specular objects weakens the inter-pixel correlation, making ours noisier than baseline PT. 3D model courtesy of BhaWin (Glass of Water).

degrades, reducing the effectiveness of our method, as illustrated in Fig. 6. Also, RR-induced correlated noise is governed primarily by the prediction error of the termination value. This error depends on the inter-pixel correlation induced by shared primary samples, which varies spatially, leading to uneven noise. Because the residual $g_i(u_j) - \phi_i(u_j)$ in Eq. 12 is signed, inverse-survival-probability weighting can produce both bright and dark outliers. As future work, extending the method with more expressive approximations of the contribution function could further improve robustness in scenes with more complex local variations.

For baseline PT and our method, MSE and relMSE primarily reflect variance and reference-scaled variance, respectively. Nevertheless, neither metric fully characterizes the spatial structure or perceptual visibility of noise. In particular, the spatially correlated residuals in our results can be more noticeable than independent noise at a similar MSE. Designing neighborhood sampling schemes that reduce inter-pixel correlation while retaining the efficiency benefits of image-space RR is an interesting direction for future work.

Another limitation is that our method can be incompatible with image denoisers that assume independence among pixel estimates. Although our estimator remains unbiased, this correlation may violate assumptions used by standard denoisers, which are often applied when a denoising bias is acceptable. To evaluate out-of-the-box denoising compatibility, we apply the publicly available pretrained OIDN model [Áfra 2026] to baseline PT and our results at the equal-time budget used in Fig. 8 and at approximately one quarter of that budget. Fig. 11 compares denoising results for the Bathroom and Lamp scenes; the remaining visual results are provided in the supplemental material. As can be seen in Table 1, Ours + OIDN achieves lower relMSE than PT + OIDN in four of the five scenes at the full-time budget and in two at the reduced-time budget. Thus, the improvement of our RR estimator over baseline PT is not fully transferred to the denoised results, since OIDN effectively suppresses independent noise but tends to leave some of our correlated noise intact. These results measure compatibility with a publicly available pretrained denoiser rather than the potential of a denoiser designed specifically for our noise distribution. We leave the development of such denoisers for future work.

Our current implementation primarily targets offline rendering, where the five pre-pass evaluations per iteration can be afforded. Applying image-space RR in real-time rendering, where typically only a few samples per pixel are available, would require a different estimator design. Extending the method to a spatio-temporal formulation that exploits correlation across both pixels and frames is a promising future direction.

Table 1. RelMSE of baseline PT and ours before and after OIDN, using the time budget of Fig. 8 (top) and approximately one quarter of that budget (bottom). For each budget, the lowest relMSE per scene is shown in bold.

Method	Bathroom	Bedroom	Coffee Machine	Foggy Bathroom	Lamp
PT	0.004741	0.005302	0.004781	0.003665	0.002940
Ours	0.001743	0.002774	**0.001212**	0.000585	0.000366
PT + OIDN	0.001296	**0.000771**	0.001396	0.000341	0.000370
Ours + OIDN	**0.001065**	0.000863	0.001301	**0.000330**	**0.000230**
PT	0.018561	0.021015	0.016986	0.014568	0.010212
Ours	0.006816	0.010766	0.004859	0.002194	0.001334
PT + OIDN	**0.001951**	**0.001317**	0.002013	**0.000509**	0.000733
Ours + OIDN	0.002340	0.002047	**0.001988**	0.000772	**0.000631**

7 Conclusion

We presented an image-space RR formulation that allows skipped path contributions to be approximated rather than discarded. This perspective extends RR beyond its traditional role at path vertices, enabling computation to be reduced via early termination of camera rays before any ray traversal. Our approach demonstrates that spatial correlation across pixels can be effectively leveraged to guide stochastic termination decisions, yielding unbiased estimators that reduce per-sample cost while mitigating variance amplification. By integrating these ideas into a simple and practical framework, we show that image-space decisions can complement existing rendering techniques and provide a new avenue for improving the error-time tradeoff in MC rendering.

Acknowledgments

We thank the anonymous reviewers for their constructive comments and suggestions. This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) and the National Research Foundation of Korea (NRF), funded by the Korean government (MSIT) under grants 2019-0-01842 and RS-2026-25596155.

References

Attila T. Áfra. 2026. Intel® Open Image Denoise. https://www.openimagedenoise.org

James Arvo and David Kirk. 1990. Particle transport and image synthesis. In *Proceedings of the 17th Annual Conference on Computer Graphics and Interactive Techniques* (Dallas, TX, USA) *(SIGGRAPH '90)*. ACM, New York, NY, USA, 63–66.

Jonghee Back, Binh-Son Hua, Toshiya Hachisuka, and Bochang Moon. 2020. Deep combiner for independent and correlated pixel estimates. *ACM Trans. Graph.* 39, 6, Article 242 (Nov. 2020), 12 pages.

Jonghee Back, Binh-Son Hua, Toshiya Hachisuka, and Bochang Moon. 2023. Input-dependent uncorrelated weighting for Monte Carlo denoising. In *SIGGRAPH Asia 2023 Conference Papers* (Sydney, NSW, Australia) *(SA '23)*. ACM, New York, NY, USA, Article 9, 10 pages.

Steve Bako, Thijs Vogels, Brian Mcwilliams, Mark Meyer, Jan Novák, Alex Harvill, Pradeep Sen, Tony Derose, and Fabrice Rousselle. 2017. Kernel-predicting convolutional networks for denoising Monte Carlo renderings. *ACM Trans. Graph.* 36, 4, Article 97 (July 2017), 14 pages.

Martin Balint, Krzysztof Wolski, Karol Myszkowski, Hans-Peter Seidel, and Rafał Mantiuk. 2023. Neural partitioning pyramids for denoising Monte Carlo renderings. In

ACM SIGGRAPH 2023 Conference Proceedings (Los Angeles, CA, USA) *(SIGGRAPH '23)*. ACM, New York, NY, USA, Article 60, 11 pages.

Pablo Bauszat, Victor Petitjean, and Elmar Eisemann. 2017. Gradient-domain path reusing. *ACM Trans. Graph.* 36, 6, Article 229 (Nov. 2017), 9 pages.

Benedikt Bitterli, Chris Wyman, Matt Pharr, Peter Shirley, Aaron Lefohn, and Wojciech Jarosz. 2020. Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting. *ACM Trans. Graph.* 39, 4, Article 148 (Aug. 2020), 17 pages.

Hajin Choi, Seokpyo Hong, Inwoo Ha, Nahyup Kang, and Bochang Moon. 2024. Online neural denoising with cross-regression for interactive rendering. *ACM Trans. Graph.* 43, 6, Article 221 (Nov. 2024), 12 pages.

Mikhail Dereviannykh, Dmitrii Klepikov, Johannes Hanika, and Carsten Dachsbacher. 2025. Neural two-level Monte Carlo real-time rendering. *Computer Graphics Forum* 44, 2 (2025), e70050.

Stavros Diolatzis, Adrien Gruson, Wenzel Jakob, Derek Nowrouzezahrai, and George Drettakis. 2020. Practical product path guiding using linearly transformed cosines. *Computer Graphics Forum* 39, 4 (July 2020), 23–33.

Ana Dodik, Marios Papas, Cengiz Öztireli, and Thomas Müller. 2022. Path guiding using spatio-directional mixture models. *Computer Graphics Forum* 41, 1 (2022), 172–189.

Honghao Dong, Guoping Wang, and Sheng Li. 2023. Neural parametric mixtures for path guiding. In *ACM SIGGRAPH 2023 Conference Proceedings* (Los Angeles, CA, USA) *(SIGGRAPH '23)*. ACM, New York, NY, USA, Article 29, 10 pages.

Arthur Firmino, Jeppe Revall Frisvad, and Henrik Wann Jensen. 2023. Denoising-aware adaptive sampling for Monte Carlo ray tracing. In *ACM SIGGRAPH 2023 Conference Proceedings* (Los Angeles, CA, USA) *(SIGGRAPH '23)*. ACM, New York, NY, USA, Article 32, 11 pages.

Sebastian Herholz, Oskar Elek, Jiří Vorba, Hendrik Lensch, and Jaroslav Křivánek. 2016. Product importance sampling for light transport path guiding. *Computer Graphics Forum* 35, 4 (2016), 67–77.

Sebastian Herholz, Yangyang Zhao, Oskar Elek, Derek Nowrouzezahrai, Hendrik P. A. Lensch, and Jaroslav Křivánek. 2019. Volume path guiding based on zero-variance random walk theory. *ACM Trans. Graph.* 38, 3, Article 25 (June 2019), 19 pages.

Jiawei Huang, Akito Iizuka, Hajime Tanaka, Taku Komura, and Yoshifumi Kitamura. 2024. Online neural path guiding with normalized anisotropic spherical Gaussians. *ACM Trans. Graph.* 43, 3, Article 26 (April 2024), 18 pages.

Wenzel Jakob. 2010. Mitsuba renderer. http://www.mitsuba-renderer.org

James T Kajiya. 1986. The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*. 143–150.

Csaba Kelemen, László Szirmay-Kalos, György Antal, and Ferenc Csonka. 2002. A simple and robust mutation strategy for the Metropolis light transport algorithm. *Computer Graphics Forum* 21, 3 (2002), 531–540.

Markus Kettunen, Eugene D'Eon, Jacopo Pantaleoni, and Jan Novák. 2021. An unbiased ray-marching transmittance estimator. *ACM Trans. Graph.* 40, 4, Article 137 (July 2021), 20 pages.

Markus Kettunen, Marco Manzi, Miika Aittala, Jaakko Lehtinen, Frédo Durand, and Matthias Zwicker. 2015. Gradient-domain path tracing. *ACM Trans. Graph.* 34, 4, Article 123 (July 2015), 13 pages.

Jaakko Lehtinen, Tero Karras, Samuli Laine, Miika Aittala, Frédo Durand, and Timo Aila. 2013. Gradient-domain Metropolis light transport. *ACM Trans. Graph.* 32, 4, Article 95 (July 2013), 12 pages.

Tzu-Mao Li, Yu-Ting Wu, and Yung-Yu Chuang. 2012. SURE-based optimization for adaptive sampling and reconstruction. *ACM Trans. Graph.* 31, 6, Article 194 (Nov. 2012), 9 pages.

Daqi Lin, Markus Kettunen, Benedikt Bitterli, Jacopo Pantaleoni, Cem Yuksel, and Chris Wyman. 2022. Generalized resampled importance sampling: foundations of ReSTIR. *ACM Trans. Graph.* 41, 4, Article 75 (July 2022), 23 pages.

Haolin Lu, Wesley Chang, Trevor Hedstrom, and Tzu-Mao Li. 2024. Real-time path guiding using bounding voxel sampling. *ACM Trans. Graph.* 43, 4, Article 125 (July 2024), 14 pages.

Joshua Meyer, Alexander Rath, Ömercan Yazici, and Philipp Slusallek. 2024. MARS: multi-sample allocation through Russian roulette and splitting. In *SIGGRAPH Asia 2024 Conference Papers* (Tokyo, Japan) *(SA '24)*. ACM, New York, NY, USA, Article 20, 10 pages.

Bochang Moon, Nathan Carr, and Sung-Eui Yoon. 2014. Adaptive rendering based on weighted local regression. *ACM Trans. Graph.* 33, 5, Article 170 (Sept. 2014), 14 pages.

Thomas Müller, Markus Gross, and Jan Novák. 2017. Practical path guiding for efficient light-transport simulation. *Computer Graphics Forum* 36, 4 (June 2017), 91–100.

Thomas Müller, Fabrice Rousselle, Alexander Keller, and Jan Novák. 2020. Neural control variates. *ACM Trans. Graph.* 39, 6, Article 243 (Nov. 2020), 19 pages.

Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically based rendering: from theory to implementation* (3rd ed.). Morgan Kaufmann Publishers Inc.

Alexander Rath, Pascal Grittmann, Sebastian Herholz, Philippe Weier, and Philipp Slusallek. 2022. EARS: efficiency-aware Russian roulette and splitting. *ACM Trans. Graph.* 41, 4, Article 81 (July 2022), 14 pages.

Alexander Rath, Marco Manzi, Sebastian Weiss, Tiziano Portenier, Farnood Salehi, Saeed Hadadan, and Marios Papas. 2025. Neural resampling with optimized candidate allocation. In *Eurographics Symposium on Rendering*, Beibei Wang and Alexander Wilkie (Eds.). The Eurographics Association.

Alexander Rath, Ömercan Yazici, and Philipp Slusallek. 2023. Focal path guiding for light transport simulation. In *ACM SIGGRAPH 2023 Conference Proceedings* (Los Angeles, CA, USA) *(SIGGRAPH '23)*. ACM, New York, NY, USA, Article 30, 10 pages.

Fabrice Rousselle, Wojciech Jarosz, and Jan Novák. 2016. Image-space control variates for rendering. *ACM Trans. Graph.* 35, 6, Article 169 (Dec. 2016), 12 pages.

Fabrice Rousselle, Claude Knaus, and Matthias Zwicker. 2011. Adaptive sampling and reconstruction using greedy error minimization. *ACM Trans. Graph.* 30, 6 (Dec. 2011), 1–12.

Lukas Ruppert, Sebastian Herholz, and Hendrik P. A. Lensch. 2020. Robust fitting of parallax-aware mixtures for path guiding. *ACM Trans. Graph.* 39, 4, Article 147 (Aug. 2020), 15 pages.

Hiroyuki Sakai, Christian Freude, Thomas Auzinger, David Hahn, and Michael Wimmer. 2024. A statistical approach to Monte Carlo denoising. In *SIGGRAPH Asia 2024 Conference Papers* (Tokyo, Japan) *(SA '24)*. ACM, New York, NY, USA, Article 68, 11 pages.

Hiroyuki Sakai, Christian Freude, Michael Wimmer, and David Hahn. 2025. Statistical error reduction for Monte Carlo rendering. In *Proceedings of the SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25)*. ACM, New York, NY, USA, Article 50, 12 pages.

Farnood Salehi, Marco Manzi, Gerhard Roethlin, Romann Weber, Christopher Schroers, and Marios Papas. 2022. Deep adaptive sampling and reconstruction using analytic distributions. *ACM Trans. Graph.* 41, 6, Article 259 (Nov. 2022), 16 pages.

László Szécsi, László Szirmay-Kalos, and Csaba Kelemen. 2003. Variance reduction for Russian-roulette. In *Proceedings of WSCG 2003*.

László Szirmay-Kalos, György Antal, and Mateu Sbert. 2005. Go with the winners strategy in path tracing. *J. WSCG* 13, 2 (2005), 49–56.

Yusuke Tokuyoshi and Takahiro Harada. 2019. Hierarchical Russian roulette for vertex connections. *ACM Trans. Graph.* 38, 4, Article 36 (July 2019), 12 pages.

Eric Veach. 1997. *Robust Monte Carlo methods for light transport simulation*. Ph. D. Dissertation. Stanford University, USA.

Jiří Vorba and Jaroslav Křivánek. 2016. Adjoint-driven Russian roulette and splitting in light transport simulation. *ACM Trans. Graph.* 35, 4, Article 42 (July 2016), 11 pages.

Qi Wang, Pengju Qiao, Yuchi Huo, Shiji Zhai, Zixuan Xie, Wei Hua, Hujun Bao, and Tao Liu. 2024. Neural kernel regression for consistent Monte Carlo denoising. *ACM Trans. Graph.* 43, 6, Article 222 (Nov. 2024), 14 pages.

Difei Yan, Zengyu Li, Lifan Wu, and Kun Xu. 2025. Generalized unbiased reconstruction for gradient-domain rendering. *ACM Trans. Graph.* 44, 6, Article 211 (Dec. 2025), 14 pages.

Chanu Yang and Bochang Moon. 2025. Imperfect image-space control variates for Monte Carlo rendering. *ACM Trans. Graph.* 44, 6, Article 205 (Dec. 2025), 11 pages.

Jiaqi Yu, Yongwei Nie, Chengjiang Long, Wenju Xu, Qing Zhang, and Guiqing Li. 2021. Monte Carlo denoising via auxiliary feature guided self-attention. *ACM Trans. Graph.* 40, 6, Article 273 (Dec. 2021), 13 pages.

Song Zhang, Daqi Lin, Markus Kettunen, Cem Yuksel, and Chris Wyman. 2024. Area ReSTIR: resampling for real-time defocus and antialiasing. *ACM Trans. Graph.* 43, 4, Article 98 (July 2024), 13 pages.

Shilin Zhu, Zexiang Xu, Tiancheng Sun, Alexandr Kuznetsov, Mark Meyer, Henrik Wann Jensen, Hao Su, and Ravi Ramamoorthi. 2021. Hierarchical neural reconstruction for path guiding using hybrid path and photon samples. *ACM Trans. Graph.* 40, 4 (July 2021).

Matthias Zwicker, Wojciech Jarosz, Jaakko Lehtinen, Bochang Moon, Ravi Ramamoorthi, Fabrice Rousselle, Pradeep Sen, Cyril Soler, and Sung-Eui Yoon. 2015. Recent advances in adaptive sampling and reconstruction for Monte Carlo rendering. *Computer Graphics Forum* 34, 2 (2015), 667–681.

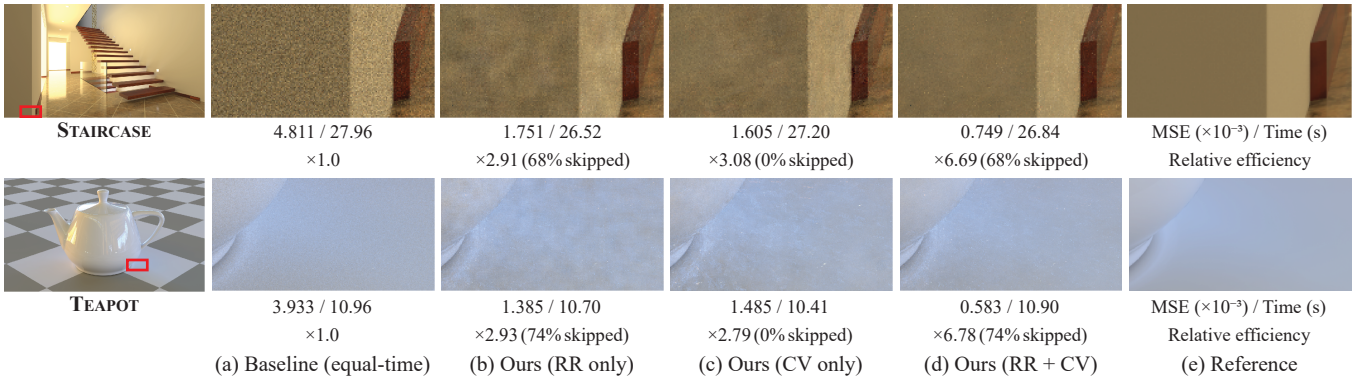

Fig. 7. Ablation of image-space Russian roulette (RR) and the control variate (CV) refinement (Sec. 5.4) under equal-time comparisons. We compare (a) baseline path tracing (PT), (b) our method using RR only, (c) using the CV refinement only (i.e., no sample skipping), and (d) the full method combining both RR and CV. RR alone (b) reduces computation but introduces additional variance, while the CV refinement (c) reduces the variance of the base estimator without sample skipping. Combining both (d) yields the best performance, showing that RR and CV address complementary aspects of efficiency—cost reduction and variance reduction, respectively. 3D model courtesy of NewSee2l035 (Staircase) and Benedikt Bitterli (Teapot).

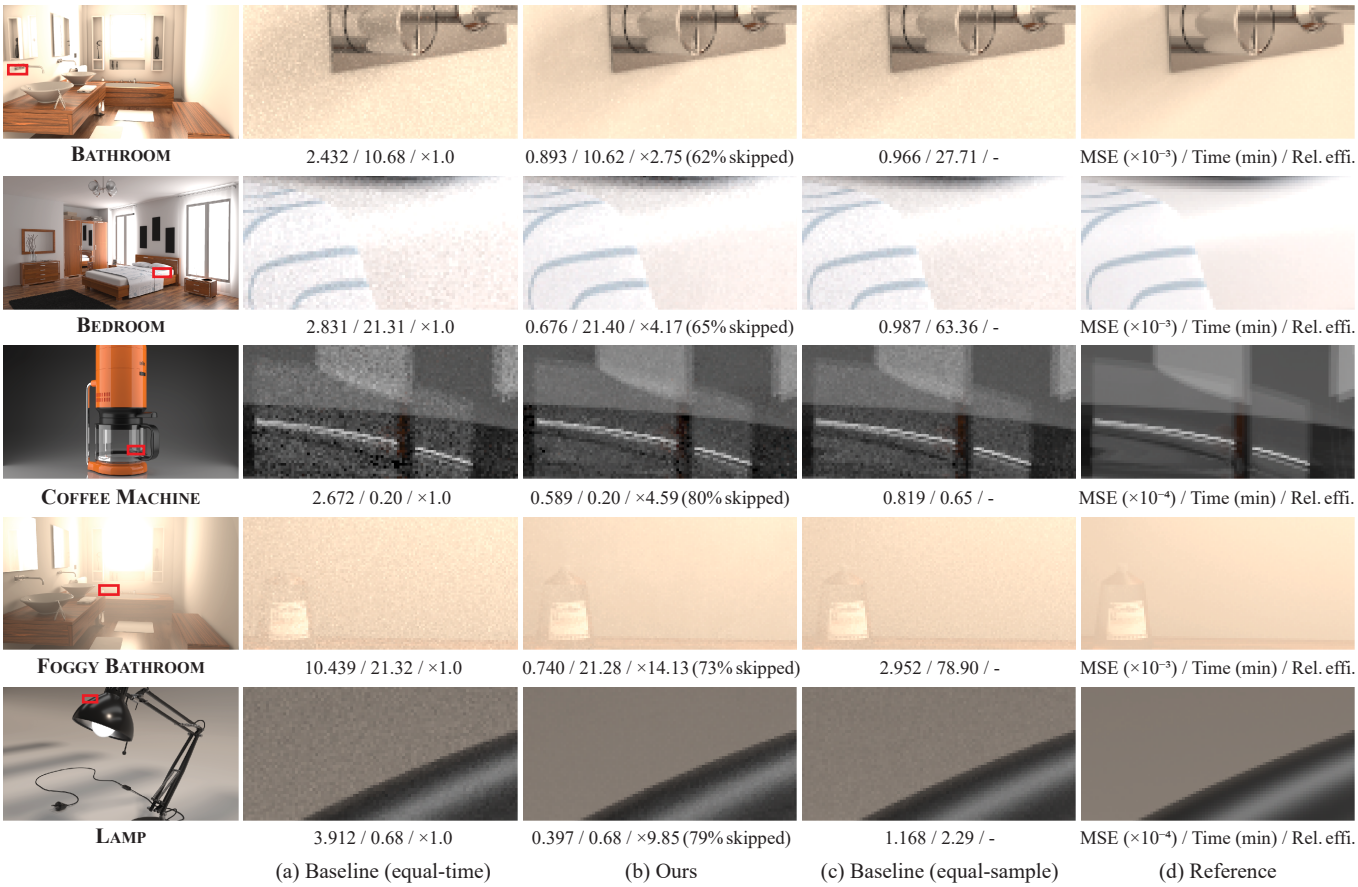

Fig. 8. Comparison of our method with baseline PT under equal-time and equal-sample settings. Under equal-time conditions, our method achieves lower error than the baseline due to improved efficiency. Notably, despite evaluating fewer light paths than the equal-sample baseline, our method achieves comparable or even lower error for the scenes above. This improvement is enabled by the CV refinement (Sec. 5.4), which reduces variance while preserving the computational benefits of image-space RR. 3D model courtesy of nacimus (Bathroom and Foggy Bathroom), SlykDrako (Bedroom), cekuhnen (Coffee Machine), and UP3D (Lamp).

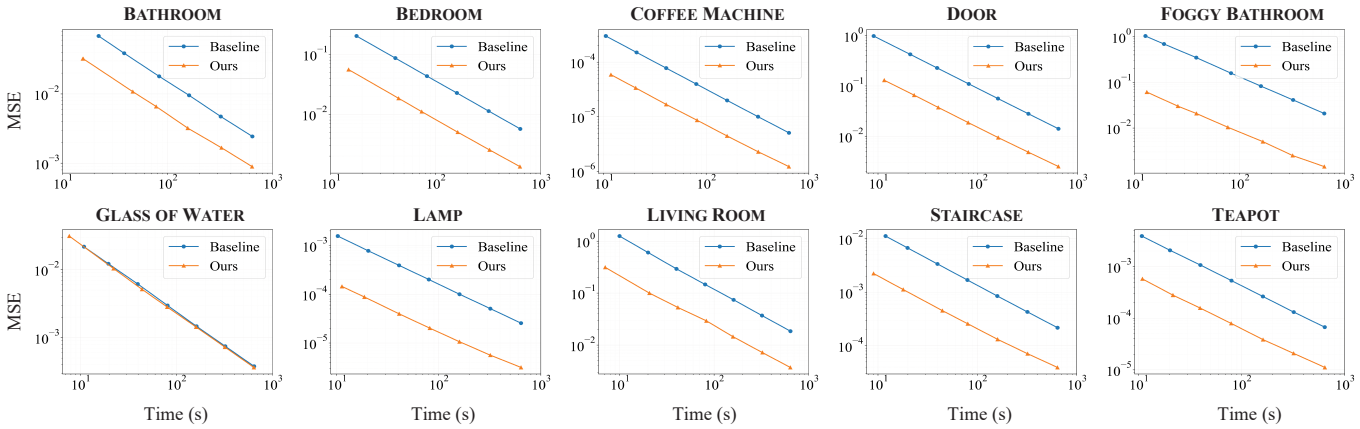

Fig. 9. MSE convergence of baseline PT and our method. Our method exhibits error reduction behavior similar to that of the baseline (i.e., similar slopes in the log-log plots). At the same time, it generally achieves lower error across scenes with varying degrees of inter-pixel correlation, while remaining unbiased.

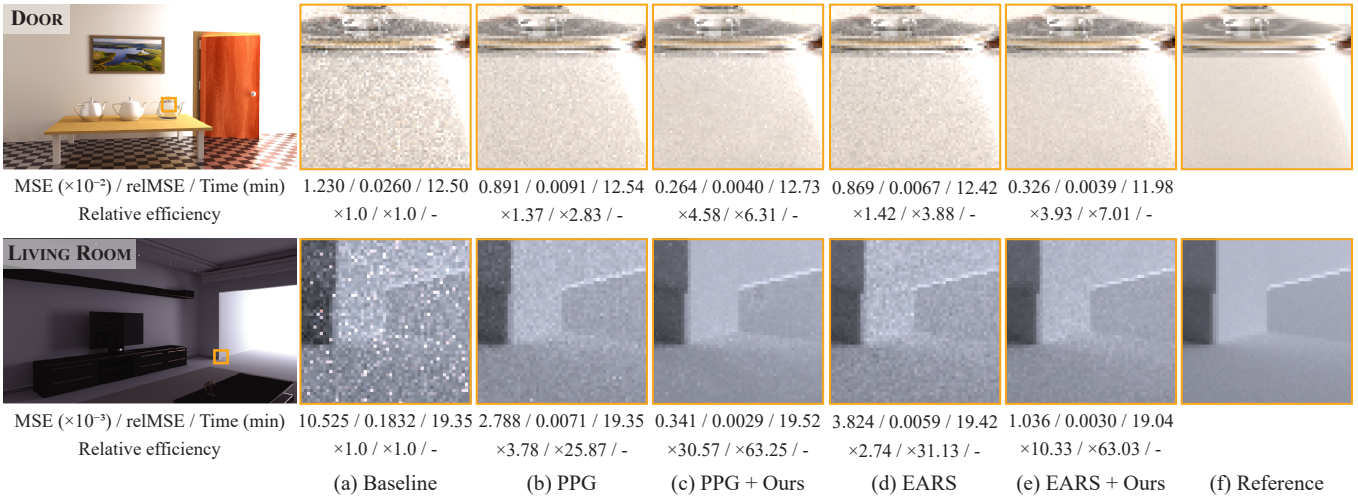

Fig. 10. Compatibility with practical path guiding (PPG) and efficiency-aware RR and splitting (EARS). While both PPG and EARS outperform baseline PT, our method further improves their efficiency. Specifically, we skip 50% of light-path evaluations for PPG + Ours and 40% for EARS + Ours in the Door scene, and 74% for PPG + Ours and 62% for EARS + Ours in the Living Room scene. 3D model courtesy of Benedikt Bitterli (Door) and Davide Tirindelli (Living Room).

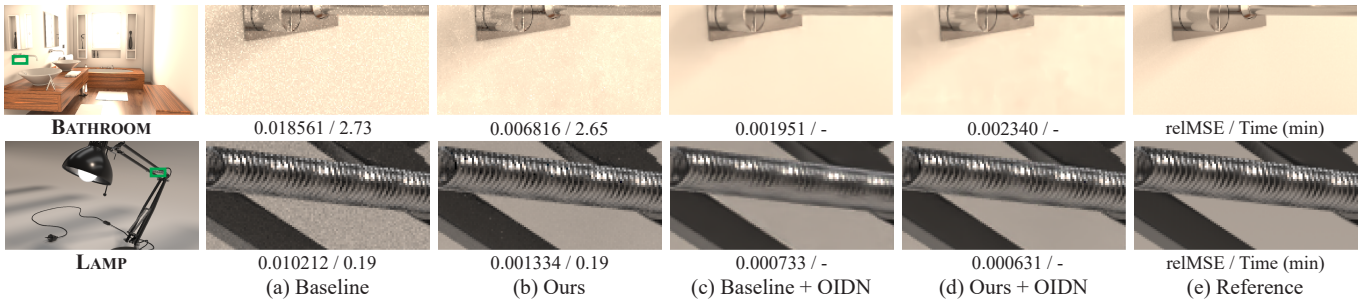

Fig. 11. Comparisons of baseline PT and our method before and after applying pretrained OIDN. Each scene uses approximately one quarter of its time budget in Fig. 8. In the unbiased raw results, ours (b) substantially reduces relMSE relative to baseline PT (a), but this gain is not fully preserved after denoising because OIDN suppresses the independent noise of baseline PT more effectively than the correlated noise produced by our method. 3D model courtesy of nacimus (Bathroom) and UP3D (Lamp).