ReSTIR PG: Path Guiding with Spatiotemporally Resampled Paths

ZHENG ZENG, NVIDIA, USA and University of California, Santa Barbara, USA MARKUS KETTUNEN, NVIDIA, Finland CHRIS WYMAN, NVIDIA, USA

LIFAN WU, NVIDIA, USA

RAVI RAMAMOORTHI, NVIDIA, USA and University of California, San Diego, USA

LING-QI YAN, Mohamed bin Zayed University of Artificial Intelligence, UAE

DAQI LIN, NVIDIA, USA

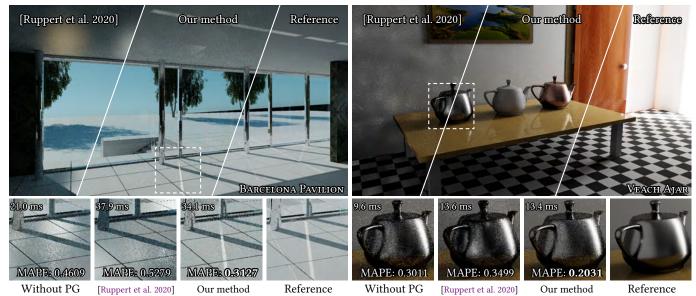


Fig. 1. Our ReSTIR Path Guiding (ReSTIR PG) framework integrates path guiding into ReSTIR Path Tracing (ReSTIR PT) by fitting guiding distributions from ReSTIR's spatiotemporally resampled paths and using them to guide the generation of initial candidates. Compared to integrating prior real-time path guiding methods [Derevyannykh 2022; Ruppert et al. 2020], which rely on raw path-traced samples, and the original ReSTIR PT without guiding (Without PG), our method produces cleaner results with lower noise and significantly reduces correlation artifacts. All methods are evaluated with ReSTIR PT enabled and rendered at 1 sample per pixel (1spp). Our method consistently achieves the best results and the lowest Mean Absolute Percentage Error (MAPE) values while preserving real-time performance.

We present ReSTIR Path Guiding (ReSTIR-PG), a real-time method that extracts guiding distributions from resampled paths produced by ReSTIR and uses them to generate improved initial candidates for the next frame. While ReSTIR significantly reduces variance through spatiotemporal resampling, its effectiveness is ultimately limited by the quality of the initial candidates, which are often poorly distributed and introduce correlation artifacts. Our key observation is that ReSTIR's accepted paths already approximate the

Authors' Contact Information: Zheng Zeng, zhengzeng@ucsb.edu, NVIDIA, USA and University of California, Santa Barbara, USA; Markus Kettunen, mkettunen@nvidia.com, NVIDIA, Finland; Chris Wyman, cwyman@nvidia.com, NVIDIA, USA; Lifan Wu, lifanw@nvidia.com, NVIDIA, USA; Ravi Ramamoorthi, ravir@cs.ucsd.edu, NVIDIA, USA and University of California, San Diego, USA; Ling-Qi Yan, lingqi.yan@mbzuai. ac.ae, Mohamed bin Zayed University of Artificial Intelligence, UAE; Daqi Lin, daqil@nvidia.com, NVIDIA, USA.

SA Conference Papers '25, Hong Kong, Hong Kong

© 2025 Copyright held by the owner/author(s).

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25), December 15–18, 2025, Hong Kong, Hong Kong, https://doi.org/10.1145/3757377.3763813.

target path contribution density, and that their bounce directions follow the ideal distribution for local path guiding – the product of incident radiance and the cosine-weighted BSDF. We exploit this structure to fit lightweight guiding distributions using each frame's resampled paths by density estimation. Compared to conventional guiding based on raw path-traced samples, ReSTIR-PG closes the loop between guiding and resampling. Our method achieves lower variance, faster response time to scene change, reduced correlation artifacts, all while preserving real-time performance.

CCS Concepts: \bullet Computing methodologies \rightarrow Rendering.

Additional Key Words and Phrases: Realistic Rendering

ACM Reference Format:

Zheng Zeng, Markus Kettunen, Chris Wyman, Lifan Wu, Ravi Ramamoorthi, Ling-Qi Yan, and Daqi Lin. 2025. ReSTIR PG: Path Guiding with Spatiotemporally Resampled Paths. In *SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25)*, December 15–18, 2025, Hong Kong, Hong Kong. ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3757377.3763813

1 Introduction

ReSTIR is widely adopted in real-time path tracing for its ability to reduce variance by resampling paths across space and time [Bitterli et al. 2020; Lin et al. 2022; Ouyang et al. 2021]. However, it often suffers from correlation artifacts – such as low-frequency boiling or blotchy noise – in challenging scenes (see Fig. 1). These arise when isolated fireflies propagate through reuse, a phenomenon known as sample impoverishment [Sawhney et al. 2024]. Consequently, final quality depends heavily on the quality of the initial path samples, which are often poorly distributed in complex cases.

A natural solution is path guiding. Prior work typically fits guiding distributions from raw path-traced samples, requiring dense data and often involving separate training and rendering phases – conditions feasible offline but problematic in real-time. In dynamic scenes, 1-spp samples are often sparse and noisy, and accumulating them over time introduces lag in the distribution.

We introduce *ReSTIR Path Guiding (ReSTIR PG)*, which instead fits guiding distributions with the resampled outputs of ReSTIR – specifically, the selected path at each pixel after spatiotemporal resampling. Our key observation is that ReSTIR samples already approximate the desired path contribution distribution and provide high-quality input for guiding. By sharing high-contribution paths across space, ReSTIR increases effective sampling density. We find that even a single frame of such samples is sufficient, avoiding temporal accumulation and lag.

Our approach constructs local guiding distributions through unweighted density estimation over the local incident directions from the accepted paths, stored in a world-space hash grid. We show this is theoretically equivalent to an ideal guiding distribution under certain assumptions. A similar result was derived in prior work [Schüßler et al. 2022] using a 9D GMM over vertex triplets. In contrast, we generalize the formulation to per-pixel path populations and design a lightweight algorithm that avoids high-dimensional fitting and conditional evaluation at runtime, making it practical for real-time.

This creates a positive feedback loop between guiding and resampling: ReSTIR provides high-quality input to guiding, and guiding improves the next frame's samples for ReSTIR. Unlike prior approaches that train guiding on raw pre-reuse paths, our method incorporates the benefits of reuse directly into guiding, improving quality without increasing temporal correlation.

Overall, ReSTIR PG improves both the initial sampling and the effectiveness of resampling. Compared to guiding based on 1spp raw path-traced data using state-of-the-art path guiding techniques, our method reduces noise, suppresses correlation artifacts, and adapts faster to scene changes. While incurring additional overhead, it still preserves real-time performance when integrated into the ReSTIR Path Tracing (ReSTIR PT) framework.

Overall, our contributions are:

 We provide a theoretical insight that ReSTIR-selected paths approximate ideal local guiding distributions (Section 3.3) and that these distributions can be efficiently constructed via unweighted density estimation over the local incident directions (Section 4.2).

- We introduce the ReSTIR PG framework, integrating path guiding into ReSTIR PT. Unlike prior methods, it learns from ReSTIR's resampled outputs instead of raw path-traced samples, preserving the quality benefits of resampling (Section 4).
- We demonstrate that integrating our path guiding into Re-STIR PT reduces variance and suppresses correlation artifacts, outperforming prior guiding methods in both image quality and responsiveness (Section 5).

2 Related work

2.1 Path Guiding

Path guiding leverages prior sample knowledge to guide the construction of new paths. Most techniques fall under *local path guiding*, which steers individual bounces using 5D or 7D spatio-directional distributions of incident radiance. Global path guiding that fits higher dimensional distributions has also been explored [Reibold et al. 2018]. Methods are typically categorized into offline and real-time variants based on performance constraints.

Offline path guiding. Offline approaches build spatio-directional distributions from large sample sets, collected either during preprocessing or iteratively during rendering. A wide range of directional structures have been explored, including histograms [Hey and Purgathofer 2002; Jensen 1995], directional trees [Lafortune and Willems 1995; Müller et al. 2017; Rath et al. 2020], cosine lobes [Bashford-Rogers et al. 2012], Gaussian [Dodik et al. 2022; Herholz et al. 2016; Huang et al. 2024; Vorba et al. 2014], von Mises-Fisher mixtures [Ruppert et al. 2020], and neural models [Bako et al. 2019; Dong et al. 2023; Huang et al. 2024; Müller et al. 2019; Zhu et al. 2021]. These directional models are typically stored in spatial structures (e.g., grids or trees) or implicitly encoded by neural networks.

Somewhat different from the above methods, Schüßler et al. [2022] propose to fit a 9D Gaussian Mixture Model (GMM) over vertex triplets to directly approximate the full contribution function. During path construction, they condition this model on the last two vertices of the path to obtain a 3D conditional distribution for sampling the next vertex. They show that this conditional distribution is theoretically proportional to the ideal local guiding distribution under their assumptions. We extend this idea by deriving a more general formulation that accounts for per-pixel path distributions, allowing us to reuse ReSTIR samples that are already adapted to each pixel's target function, rather than generating training paths separately for path guiding. Furthermore, we design a practical algorithm that operates efficiently without incurring the overhead of high-dimensional model fitting or conditional evaluations.

Despite their effectiveness in offline rendering, these methods require far more samples than typically available in real-time settings. We therefore focus on real-time variants.

Real-time path guiding. Real-time methods simplify directional structures and focus on guiding direct or one-bounce indirect lighting. Dittebrandt et al. [2020] compress directional quadtrees [Müller et al. 2017] using bit fields and rely on previous-frame radiance estimates. Pantaleoni [2020] explore histograms, GMMs, and trees in the real-time context, with caching to stabilize radiance estimates.

Derevyannykh [2022] and Dittebrandt et al. [2023] use a single screen-space Gaussian or vMF lobe per pixel, with the latter incorporating MCMC for effective VMM construction. Lu et al. [2024] forgo spatial-directional structures entirely, tracing secondary rays to populate irradiance voxels and use them for directional sampling, resembling many-light strategies. This method avoids temporal lag.

These methods are designed to address the scarcity and high variance of path-traced path samples typically available in real-time rendering, employing techniques to reuse guiding information across space and time. Different from these approaches, we study how to build guiding distributions using ReSTIR path samples—paths that are already spatiotemporally resampled.

2.2 ReSTIR

Resampling is widely used in graphics to improve sample distributions by stochastically selecting from weighted candidates. Resampled Importance Sampling (RIS) [Talbot et al. 2005] formulates resampling for Monte Carlo integration. The introduction of reservoirbased spatiotemporal importance resampling (ReSTIR) [Bitterli et al. 2020] significantly advanced the efficiency of resampling. We recommend the gentle introduction by Wyman et al. [2023] for a comprehensive review of the theory and the algorithm.

ReSTIR has been extended in both rendering content and theory. It supports diffuse GI [Ouyang et al. 2021], general path tracing [Lin et al. 2022] (ReSTIR PT), volume rendering [Lin et al. 2021], and subsurface scattering [Werner et al. 2024]. In the theoretical direction, Generalized RIS [Lin et al. 2022] introduces shift mappings to allow reusing samples from different domains, Conditional RIS [Kettunen et al. 2023] enables subpath resampling in conditional spaces, and ReSTIR-MCMC [Sawhney et al. 2024] interleaves resampling with mutation to allow temporal decorrelation.

Beyond screen-space reservoirs, variants store samples in worldspace grids [Boissé 2021], voxels [Boksansky et al. 2021], surface caches [Weinrauch et al. 2023], or texture space [Li and Guo 2024]. Area ReSTIR [Zhang et al. 2024] resamples over pixel footprints for better reuse on high-frequency materials.

Other efforts improve ReSTIR's robustness and efficiency, including presampled light tiles for cache-efficient candidate generation [Wyman and Panteleev 2021], novel resampling MIS weights for better neighbor weighting [Pan et al. 2024; Tokuyoshi 2023], and cache-based path tracing cost reduction by irradiance probes [Majercik et al. 2021] and photons [Kern et al. 2024]. ReSTIR also helps differentiable rendering [Chang et al. 2023; Wang et al. 2023] by enabling sample reuse between gradient steps. Zhang et al. [2025] use ReSTIR to select the most contributive lights for shadow map generation. Alber et al. [2025] also employ resampling but adopt stochastic multiple importance sampling (SMIS) instead of RIS, and construct a Markov chain-based vMF mixture model, whereas we directly extract a vMF mixture model from ReSTIR samples.

To our knowledge, no prior work yet explores the connection between ReSTIR and path guiding. Our work is the first that uses path guiding to improve ReSTIR samples and the first to use ReSTIRgenerated samples to improve the quality of path guiding.

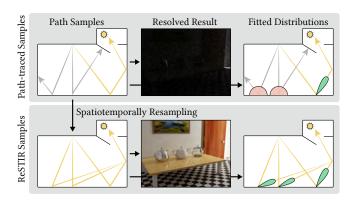


Fig. 2. Example in 2D of the PT samples and ReSTIR samples. The top row shows the 1spp PT samples - initial candidates generated by the path tracer - along with their resolved result and the corresponding fitted guiding distributions.

Problem analysis

In light transport, the color I_i of a pixel j can be computed by integrating the path contribution function $f(\bar{\mathbf{x}})$ over the pixel's path space Ω_i . Path tracers solve this by randomly sampling paths $\bar{\mathbf{x}} = [\mathbf{x}_0, \mathbf{x}_1, ..., \mathbf{x}_D]$ of different lengths D from the camera \mathbf{x}_0 to the lights. This forms the Monte Carlo estimator that unbiasedly estimates I_i :

$$I_{j} = \int_{\Omega_{j}} f(\bar{\mathbf{x}}) \, d\bar{\mathbf{x}} \approx \frac{1}{N} \sum_{i=1}^{N} \frac{f(\bar{\mathbf{x}})}{p(\bar{\mathbf{x}})} \,, \tag{1}$$

where $p(\bar{\mathbf{x}})$ is the sampling probability density of path $\bar{\mathbf{x}}$, and Nrepresents the number of paths traced through the pixel *j*.

To minimize the variance caused by a small number of samples, we aim to importance sample the integrand $f(\bar{\mathbf{x}})$ – ideally having the probability density $p(\bar{\mathbf{x}}) \propto f(\bar{\mathbf{x}})$, which requires only 1 sample to converge. However, the path contribution $f(\bar{\mathbf{x}})$ is generally unknown before sampling the path $\bar{\mathbf{x}}$, making it difficult to achieve.

Consider the expanded form of I_i (assume pinhole camera) which integrates over the surface area \mathcal{A} :

$$I_j = \int_{\mathscr{A}} W_e^{(j)}(\mathbf{x}_0 \to \mathbf{x}_1) L_i(\mathbf{x}_1 \to \mathbf{x}_0) G(\mathbf{x}_0 \leftrightarrow \mathbf{x}_1) d\mathbf{x}_1, \qquad (2)$$

where $W_e^{(j)} = h_j \cdot W_e$ is the per pixel sensor response function, G is the visibility-including geometry term between two points, and the incoming radiance L_i can be recursively integrated as

$$L_i(\mathbf{x}_k \to \mathbf{x}_{k-1}) = L_e(\mathbf{x}_k \to \mathbf{x}_{k-1}) + L_r(\mathbf{x}_k \to \mathbf{x}_{k-1})$$

where $L_{\mathbf{r}}(\mathbf{x}_k \to \mathbf{x}_{k-1}) =$

$$\int_{\mathcal{A}} \rho(\mathbf{x}_{k+1} \to \mathbf{x}_k \to \mathbf{x}_{k-1}) G(\mathbf{x}_k \leftrightarrow \mathbf{x}_{k+1}) L_i(\mathbf{x}_{k+1} \to \mathbf{x}_k) \, \mathrm{d}\mathbf{x}_{k+1} \quad (3)$$

for $k \ge 1$, and ρ is the BSDF. A path tracer usually recursively applies importance sampling on the integral $L_{\rm r}$ that computes the contribution from reflected radiance. This is usually written as a 1 spp recursive estimator in the solid angle measure where an incoming direction ω_i is sampled conditioned on the current vertex **x** and the outgoing direction ω_o to find the next vertex:

$$\langle L_{\mathbf{r}}(\mathbf{x}, \omega_o) \rangle = \frac{\rho(\mathbf{x}, \omega_o, \omega_i) L_i(\mathbf{x}, \omega_i) \cos \theta}{p(\omega_i | \mathbf{x}, \omega_o)}.$$
 (4)

Since the recursive term $L_i(\mathbf{x}, \omega_i)$ is unknown, common techniques consider local importance sampling of limited terms, like BSDF sampling or next-event estimation of direct illumination.

Local path guiding techniques usually build a 5D spatio-directional distribution that approximates the incoming radiance field at different regions, i.e. $p_{\rm guide}(\omega_i|\mathbf{x}) \stackrel{<}{\lesssim} L_i(\mathbf{x},\omega_i)$. This complements the BSDF sampling for sampling the whole product and is usually combined with the BSDF sampling technique via MIS. However, an ideal sampling would have PDF $p(\omega_i|\mathbf{x},\omega_o) \propto \rho(\mathbf{x},\omega_o,\omega_i)L_i(\mathbf{x},\omega_i)\cos\theta$. This adds the conditioning on ω_o and improves the sampling for glossy materials. A number of *product guiding* techniques [Diolatzis et al. 2020; Herholz et al. 2016; Huang et al. 2024; Müller et al. 2019] were proposed to approximate this sampling by fitting 7D distributions, but generally incur high computational overhead. They still need to do MIS with BSDF sampling for robustness.

Our method is based on ReSTIR. ReSTIR can resample paths according to a target function $\hat{p}(\bar{\mathbf{x}}) = f(\bar{\mathbf{x}})$ for each pixel. This creates a population of well-distributed paths which we can leverage to create local path guiding distributions. We introduce ReSTIR and its theory first, before building the link with path guiding.

3.1 GRIS and ReSTIR

Resampled Importance Sampling (RIS) [Talbot et al. 2005] can be used to resample from candidate samples to approximate the distribution proportional to a target function $\hat{p}(x)$, for which a closed-form sampling method usually does not exist. RIS works by probabilistically selecting a Y from i.i.d. candidate samples $\{X_i\}_{i=1}^M$, suboptimally distributed with source PDF p(x), with weights $w_i = \frac{1}{M}\frac{\hat{p}(X_i)}{p(X_i)}$. This can also be used to estimate an integral $I = \int_{\Omega} f(x) \ dx$ by computing $\langle I \rangle = f(Y)W_Y$ where

$$W_Y = \frac{1}{\hat{p}(Y)} \sum_{i=1}^{M} w_i$$
 (5)

is the *unbiased contribution weight* that estimates the actual unknown reciprocal PDF $1/p_Y(Y)$.

As candidate sample count $M \to \infty$, $Y \sim \hat{p}(Y)/\int_{\Omega} \hat{p}(x) \, dx$. Re-STIR [Bitterli et al. 2020] finds that the candidate sample generation cost can be amortized by reusing the selected samples between pixels and frames to recursively feed further RIS passes, dramatically improving the efficiency of RIS. The math is later formalized as Generalized RIS (GRIS) [Lin et al. 2022], where differently distributed, potentially correlated samples can be used as candidates. The candidates need not have known PDF; unbiased contribution weights like W_Y can replace the reciprocal source PDF in future resampling passes. Shift mappings $T_i:\Omega_i\to\Omega$ can be used to map X_i generated from different domains to the current domain, allowing reusing spatial and temporal samples without adding bias. GRIS uses a more general resampling weight $w_i=m_i(Y_i)\hat{p}(Y_i)W_{Y_i}$ where $Y_i=T_i(X_i)$ and $W_{Y_i}=W_{X_i}|T_i'|$. The resampling MIS weight $m_i(Y_i)$ dynamically weights each technique and guarantees unbiasedness.

ReSTIR builds a reservoir per pixel each frame, which is essentially a (X, W_X, c) tuple, where c is a *confidence weight* applied in m_i that controls the influence in weighting, favoring reservoirs with more effective samples. A merge operation reads samples from all input reservoirs and does GRIS using them as candidates, writing the result to the current reservoir and updating c to be the sum of all input confidence weights.

ReSTIR can be summarized in three passes.

- (1) Initialization: Resample candidates (e.g. path traced samples) with RIS. Store selected X with $(X, W_X, c = 1)$.
- (2) Temporal reuse: Merge a last frame's reservoir (fetched by the motion vector), but clamp its confidence weight c_{temp} to avoid unconstrained correlation and stale distribution in dynamic scenes.
- (3) Spatial reuse: Merge a couple of randomly selected neighboring pixels' reservoirs with the current reservoir. Compute the pixel color using the reservoir sample and keep the reservoir for the next frame.

Recursive spatial and temporal reuse ideally lead to a quick exponential growth of effective sample count, reducing noise to a low level in a few frames. Yet if the initial samples are poor, a subset of high-energy samples can dominate the resampling process, known as sample impoverishment [Sawhney et al. 2024]. With a capped $c_{\rm temp}$, excessive correlation usually exhibits as pulsating correlation blobs and streaks. Shifting reuse to suffix space [Kettunen et al. 2023] or adding Metropolis mutations [Sawhney et al. 2024] helps reduce correlation, but often requires heavy computation to obtain significant improvement.

We aim for a more direct route: improving the quality of initial candidate samples. While ReSTIR can keep improving the distribution of the final selected sample over frames, the initial candidate samples are always sampled from a suboptimal distribution – that offered by a path tracer. Our goal is to improve them by path guiding.

3.2 Observation

A straightforward approach is to build guiding distributions from initial candidates generated by the path tracer (*PT samples*). Many offline path guiding methods, can, in principle, be applied in this context. Yet real-time budgets allow only one sample per pixel, so PT samples are often too sparse and noisy for reliable fitting (Figure 2 top). The key question becomes: *how can we obtain dense*, *low-variance samples for effective guiding in real time?*

Previous real-time path guiding methods [Derevyannykh 2022; Dittebrandt et al. 2020, 2023; Lu et al. 2024; Pantaleoni 2020] addressed this challenge through various forms of reuse, such as spatiotemporal sample reuse, temporal reuse of sufficient statistics of model parameters, and caching radiance to reduce variance. Yet they often suffer from slow adaptation to scene changes and sometimes show high variance, largely due to an inherent trade-off between stability and timeliness of the distribution.

We notice another sample set already produced at no extra cost: the 1 spp outputs of ReSTIR's spatiotemporal resampling. These *ReSTIR samples* share high-energy paths between pixels, boosting density, exhibit lower variance (Figure 2 bottom), and stay current via unbiased temporal reuse [Lin et al. 2022]. They are therefore

ideal for guiding. To our knowledge, these samples have not been previously explored for constructing guiding distributions.

However, it is not immediately clear how to fit a path guiding distribution with ReSTIR samples. Fitting the path distribution directly would require global path guiding techniques, generally impractical in real-time. Applying local path guiding techniques is also nontrivial, as they normally rely on access to local incident radiance as weights in fitting - information not readily recoverable after spatiotemporal resampling¹.

We propose to project the global path information of ReSTIR samples onto local decisions, specifically, directional sampling at each bounce. In subsection 3.3, we show that this projection only requires knowing the local bounce directions and the resulting distribution allows us to perform the ideal local guiding.

3.3 Key insight

Our key insight is that if we have a set of path samples distributed according to the path contribution functions - such as the ReSTIR samples - the distribution of local bounce directions at each path vertex is proportional to the product of the incident radiance and cosine-weighted BSDF, which is exactly the distribution that we want in local path guiding. Our derivation is as follows.

Recall that our goal is to construct a directional distribution:

$$p(\omega_i \mid \mathbf{x}, \omega_o) \propto \rho(\mathbf{x}, \omega_o, \omega_i) L_i(\mathbf{x}, \omega_i) \cos \theta,$$
 (6)

where ρ is the BSDF, L_i is the incoming radiance, and θ is the angle between ω_i and the surface normal at **x**.

Now, suppose each pixel *j* has reached its target distribution, i.e. $\bar{\mathbf{x}} \sim f(\bar{\mathbf{x}})$ for each pixel. Because different pixels see a different part² of \mathcal{A} for its possible primary hits \mathbf{x}_1 , each has a different normalization factor C_i defined as

$$C_j = \frac{1}{\int_{\Omega_j} f(\bar{\mathbf{x}}) \, \mathrm{d}\bar{\mathbf{x}}},\tag{7}$$

integrating over that pixel's paths with non-zero pixel response.

Consider a global path population formed by all *n* pixels (stratified samples of the whole path space). This implies

$$p(\bar{\mathbf{x}}) = \frac{1}{n} \sum_{j=1}^{n} C_j f(\bar{\mathbf{x}}) [h_j(\mathbf{x}_0 \to \mathbf{x}_1) > 0]$$
 (8)

for all paths $\bar{\mathbf{x}}$ of all lengths in the space $\Omega = \bigcup_{D=2}^{\infty} \mathscr{A}^D$ (excluding the directly visible emissive contribution).

We begin with some term definitions. The contribution function f is defined separately for each path of length D as

$$f([\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_D])$$

$$= W_e^G(\mathbf{x}_0, \mathbf{x}_1) \left(\prod_{u=1}^{D-1} \rho(\mathbf{x}_{u+1} \to \mathbf{x}_u \to \mathbf{x}_{u-1}) G(\mathbf{x}_{u+1} \leftrightarrow \mathbf{x}_u) \right) L_e(\mathbf{x}_D),$$

where ρ is the BSDF, G is the geometry term, and $W_e^G(\mathbf{x}_0, \mathbf{x}_1) =$ $W_{\rm e}(\mathbf{x}_0 \to \mathbf{x}_1)G(\mathbf{x}_0 \leftrightarrow \mathbf{x}_1)$. For brevity, we define

$$T(\mathbf{x}_{u}) = \begin{cases} \rho \left(\mathbf{x}_{u+1} \to \mathbf{x}_{u} \to \mathbf{x}_{u-1}\right) G\left(\mathbf{x}_{u+1} \leftrightarrow \mathbf{x}_{u}\right) & (u > 1) \\ W_{e}^{G}(\mathbf{x}_{0}, \mathbf{x}_{1}) \rho \left(\mathbf{x}_{2} \to \mathbf{x}_{1} \to \mathbf{x}_{0}\right) G\left(\mathbf{x}_{2} \leftrightarrow \mathbf{x}_{1}\right) & (u = 1) \end{cases}$$
(9)

so that the contribution function becomes

$$f([\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_k]) = \left(\prod_{u=1}^{k-1} T(\mathbf{x}_u)\right) L_e(\mathbf{x}_k). \tag{10}$$

Next, we express the conditional distribution of the next vertex \mathbf{x}_{u+1} of a path, where $u \geq 1$. By defining the singular case $\int_{\mathcal{A}^0} f(\bar{\mathbf{x}}) \, \mathrm{d}\mathbf{x}_{u+2} \cdots \mathrm{d}\mathbf{x}_{u+1} \coloneqq f([\mathbf{x}_0, \dots, \mathbf{x}_{u+1}]) \text{ for brevity, we have}$

$$p(\mathbf{x}_{u+1} \mid \mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_u) = \frac{p(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_u, \mathbf{x}_{u+1})}{p(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_u)}$$
(11)

$$= \frac{\sum_{D=u+1}^{\infty} \int_{\mathscr{A}^{D-u-1}} p(\bar{\mathbf{x}}) \, d\mathbf{x}_{u+2} \cdots d\mathbf{x}_{D}}{\sum_{D=u+1}^{\infty} \int_{\mathscr{A}^{D-u}} p(\bar{\mathbf{x}}) \, d\mathbf{x}'_{u+1} \, d\mathbf{x}_{u+2} \cdots d\mathbf{x}_{D}}.$$
 (12)

We substitute Eq. (8) and Eq. (10):

$$= \frac{\sum_{D=u+1}^{\infty} \int_{\mathcal{A}^{D-u-1}} \left(\frac{\frac{1}{n} \sum_{j=1}^{n} C_{j} [h_{j}(\mathbf{x}_{0} \rightarrow \mathbf{x}_{1}) > 0]}{\sum_{D=u+1}^{\infty} \int_{\mathcal{A}^{D-u}} \left(\frac{\frac{1}{n} \sum_{j=1}^{n} C_{j} [h_{j}(\mathbf{x}_{0} \rightarrow \mathbf{x}_{1}) > 0]}{\sum_{l=1}^{\infty} \int_{\mathcal{A}^{D-u-1}} \left(\prod_{l=1}^{l-1} T(\mathbf{x}_{l})\right) L_{e}(\mathbf{x}_{D}) d\mathbf{x}_{u+2} \cdots d\mathbf{x}_{D}}\right)}$$

$$= \frac{\sum_{D=u+1}^{\infty} \int_{\mathcal{A}^{D-u}} \left(\prod_{l=1}^{l-1} T(\mathbf{x}_{l})\right) L_{e}(\mathbf{x}_{D}) d\mathbf{x}_{u+2} \cdots d\mathbf{x}_{D}}{\sum_{D=u+1}^{\infty} \int_{\mathcal{A}^{D-u}} \left(\prod_{l=1}^{l-1} T(\mathbf{x}_{l})\right) L_{e}(\mathbf{x}_{D}) d\mathbf{x}_{u+1}' d\mathbf{x}_{u+2} \cdots d\mathbf{x}_{D}}\right)$$

$$(13)$$

Factor $\prod_{t=1}^{u-1} T(\mathbf{x}_t)$ cancels:

$$= \frac{T(\mathbf{x}_u) \sum_{D=u+1}^{\infty} \int_{\mathcal{A}^{D-u-1}} \left(\prod_{t=u+1}^{D-1} T(\mathbf{x}_t) \right) L_e(\mathbf{x}_D) \, d\mathbf{x}_{u+2} \cdots d\mathbf{x}_D}{\sum_{D=u+1}^{\infty} \int_{\mathcal{A}^{D-u}} \left(\prod_{t=u}^{D-1} T(\mathbf{x}_t) \right) L_e(\mathbf{x}_D) \, d\mathbf{x}'_{u+1} \, d\mathbf{x}_{u+2} \cdots d\mathbf{x}_D}$$
(14)

The sum multiplying $T(\mathbf{x}_u)$ is now exactly $L_i(\mathbf{x}_{u+1} \to \mathbf{x}_u)$:

$$= \frac{T(\mathbf{x}_u)L_i(\mathbf{x}_{u+1} \to \mathbf{x}_u)}{\sum_{D=u+1}^{\infty} \int_{\mathcal{A}^{D-u}} \left(\prod_{t=u}^{D-1} T(\mathbf{x}_t)\right) L_e(\mathbf{x}_D) d\mathbf{x}'_{u+1} d\mathbf{x}_{u+2} \cdots d\mathbf{x}_D}.$$
 (15)

From the expression above, we observe that the earlier vertices $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{u-2}$ have no influence on this conditional distribution. Hence, we can simplify the conditional to:

 $^{^1\}text{To}$ compute an estimate of $L_i(\mathbf{x}_k,\omega_i)$ on a path vertex \mathbf{x}_k requires subpath PDFs. But only the unbiased contribution weight of the whole path is known after resampling. 2 We assume box filters ($h_{j}=1$ within the pixel). But our derivation can be easily generalized to non-box filters.

$$p\left(\mathbf{x}_{u+1} \mid \mathbf{x}_{u-1}, \mathbf{x}_{u}\right) \tag{16}$$

$$= \frac{T(\mathbf{x}_u)L_i(\mathbf{x}_{u+1} \to \mathbf{x}_u)}{\sum_{D=u+1}^{\infty} \int_{\mathscr{A}^{D-u}} \left(\prod_{t=u}^{D-1} T(\mathbf{x}_t)\right) L_e(\mathbf{x}_D) \, \mathrm{d}\mathbf{x}'_{u+1} \, \mathrm{d}\mathbf{x}_{u+2} \cdots \, \mathrm{d}\mathbf{x}_D}$$
(17)

$$\propto \rho \left(\mathbf{x}_{u+1} \to \mathbf{x}_u \to \mathbf{x}_{u-1} \right) G \left(\mathbf{x}_{u+1} \longleftrightarrow \mathbf{x}_u \right) L_i(\mathbf{x}_{u+1} \to \mathbf{x}_u).$$
 (18)

Note RHS is a result of observing that the denominator of LHS is a constant and substituting Equation 9.

Expressed in terms of the solid angle form, this becomes:

$$p(\omega_i \mid \mathbf{x}_u, \omega_o) \propto \rho(\mathbf{x}_u, \omega_o, \omega_i) L_i(\mathbf{x}_u, \omega_i) \cos \theta.$$
 (19)

This result implies that, conditioned on the vertex \mathbf{x}_u and the outgoing direction ω_o , the distribution of the incident direction ω_i is proportional to the product of the incident radiance, the cosine term, and the BSDF, precisely the ideal distribution for local guiding.

Consequently, given ReSTIR samples - which are approximately distributed according to the path contribution function³ – we can directly extract the desired local directional distributions at each vertex without the need to explicitly evaluate incident radiance.

Note that Schüßler et al. [2022]'s derivation is a special case of our derivation (n = 1 pixel). They assume paths distributed globally with $p(\bar{\mathbf{x}}) \sim f(\bar{\mathbf{x}})$. Although generalizing their derivation to ours is conceptually direct (cancellation of the C_i sums), the conclusions affect the design choices - in their case they perform resampling in grid cells from global path population according to $f(\bar{\mathbf{x}})$, whereas we can directly use the ReSTIR sampled paths that are (approximately) distributed to Equation 8 to populate a guiding distribution.

ReSTIR Path Guiding

In this section, we present our ReSTIR Path Guiding (ReSTIR PG) framework (Figure 3), which integrates path guiding into ReSTIR PT [Lin et al. 2022]. In ReSTIR PG, guiding distributions over incident directions are constructed from ReSTIR samples and used to guide the generation of initial candidates in the subsequent frame.

4.1 Overview

Our goal is to construct a local guiding distribution that is proportional to the product of the incident radiance, the cosine term, and the BSDF. As shown in Section 3.3 and Eq. (19), the conditional distribution $p(\omega_i \mid \mathbf{x}, \omega_o)$, estimated directly from the ReSTIR samples, can serve as an effective approximation to this ideal target.

However, similar to the 9D GMM over vertex triplets used in prior work [Schüßler et al. 2022], explicitly representing this conditional distribution requires fitting a 7D high-dimensional model and querying conditionals at runtime. This is computationally expensive and impractical for real-time rendering.

To address this, we propose to reduce the dimensionality, resulting in a 2D distribution $p(\omega_i)$. We show that this marginal, when estimated from ReSTIR samples, still approximates the ideal target under reasonable assumptions (Section 4.2).

Based on this insight, our method proceeds as follows: first, we partition the scene into spatial cells. Next, after the spatiotemporal reuse stage of ReSTIR PT, we collect ReSTIR samples and splat their local incident directions into the corresponding cells (Section 4.3). At the end of each frame, we fit a guiding distribution in each cell to represent $p(\omega_i)$ (Section 4.4). Finally, in the subsequent frame, we use the fitted distributions to guide the generation of new initial candidate samples (Section 4.5).

4.2 Dimensionality reduction

We aim to reduce the dimensionality of the conditional distribution $p(\omega_i \mid \mathbf{x}, \omega_0)$, which is estimated directly from ReSTIR samples and is proportional to the ideal importance sampling distribution:

$$p(\omega_i \mid \mathbf{x}, \omega_o) \propto \rho(\mathbf{x}, \omega_o, \omega_i) L_i(\mathbf{x}, \omega_i) \cos \theta.$$
 (20)

First, we reduce the dimensionality to 4D by omitting the position x. Following the practices of prior path guiding work [Müller et al. 2017; Pantaleoni 2020; Ruppert et al. 2020], we partition the scene into spatial cells and fit a directional distribution within each cell. This removes the explicit dependence on \mathbf{x} at the cost of averaging the distributions. Within each spatial cell, the distribution becomes

$$p(\omega_i \mid \omega_o) \propto \rho(\omega_o, \omega_i), L_i(\omega_i) \cos \theta.$$
 (21)

Still, representing the 4D conditional distribution $p(\omega_i \mid \omega_o)$ using standard directional models remains impractical under real-time constraints. We further reduce the dimensionality to 2D and completely avoid explicit and costly conditional modeling by leveraging the relationship between marginal and conditional distributions:

$$p(\omega_i) = \int_{\mathbb{H}^2} p(\omega_i \mid \omega_o) p(\omega_o) \, d\omega_o. \tag{22}$$
 Substituting Eq. (21) into Eq. (22) yields

$$p(\omega_i) \propto L_i(\omega_i) \cos(\theta) \int_{\mathbb{H}^2} \rho(\omega_o, \omega_i) p(\omega_o) d\omega_o,$$
 (23)

which shows that the distribution of the incident direction ω_i is proportional to the product of the incident radiance, the cosine term, and a BSDF-weighted integral over the outgoing direction distribution $p(\omega_0)$. Unlike the 4D conditional distribution $p(\omega_i \mid \omega_0)$, this 2D marginal distribution $p(\omega_i)$ can be efficiently constructed using standard directional representations.

However, an important question arises: is this marginal distribution $p(\omega_i)$ – that is, the right-hand side of Eq. (23) – still practically useful for guiding?

On diffuse surfaces, where the BSDF ρ is constant with respect to direction, we show that the marginal distribution $p(\omega_i)$ is also proportional to the ideal target. In this case, the BSDF term can be factored out of the integral in Eq. (23), and the integral over $p(\omega_o)$ evaluates to one due to normalization. Hence, Eq. (23) reduces to

$$p(\omega_i) \propto \rho L_i(\omega_i) \cos(\theta),$$
 (24)

which matches the ideal importance sampling distribution for diffuse surfaces. We visualize the distributions of the LHS that estimated from ReSTIR samples and RHS at a path vertex in a purely diffuse VEACH AJAR scene in Fig. 4. The visualization shows that the distribution of incident directions (LHS) closely matches the distribution of the product of BSDF, incident radiance, and cosine term (RHS).

³This assumption breaks down when ReSTIR fails completely – for example, when no effective reuse is possible. Typically, even with 1spp input, ReSTIR produces samples close to the ideal distribution, though with strong correlation due to poor initial samples.

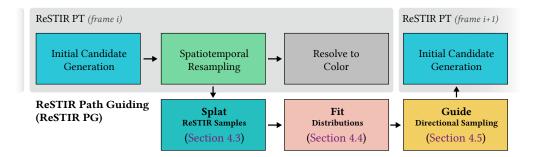


Fig. 3. Overview of our ReSTIR Path Guiding (ReSTIR PG) framework. Compared to the original design of ReSTIR PT, we introduce three additional stages: first, at frame i, we splat the bounce directions of ReSTIR samples into a spatial data structure; then, we fit guiding distributions in each spatial cell at the end of each frame i; and finally, we guide the initial candidate generation in the subsequent frame i + 1.

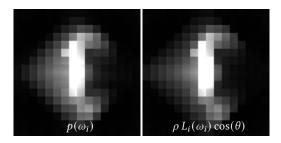


Fig. 4. Visualization of the distribution of incident directions $p(\omega_i)$ estimated from ReSTIR samples (left) and the product of BSDF, incident radiance, and cosine term $\rho L_i(\omega_i) \cos(\theta)$ (right) at a path vertex in a purely diffuse Veach Ajar scene.

On glossy surfaces, where the BSDF $\rho(\omega_i, \omega_o)$ varies with direction, we can no longer simplify Eq. (23) by factoring $\rho(\omega_i, \omega_o)$ out of the integral. However, we find this integral $\int \rho(\omega_i, \omega_o) p(\omega_o) d\omega_o$ can still serve a useful role: it effectively performs a BSDF-weighted averaging of the outgoing distribution $p(\omega_0)$ for each incident direction ω_i . This naturally embeds the directional preferences on incident directions of glossy reflections into the marginal distribution. In practice, we observe that this property holds as long as the outgoing directions ω_o used during evaluation and sampling are consistent with the distribution $p(\omega_o)$ of ReSTIR samples during construction. In our experiments, we did not observe it performing worse than a distribution based solely on incident radiance.

To accommodate these differing characteristics of directional distributions across BSDF types and ensure that cells with the same BSDF are grouped together, we adopt the non-uniform spatial hash grid from prior work [Pantaleoni 2020], whose resolution adapts with viewing distance, and extend its hash function to explicitly encode the BSDF ID.

We adopt von Mises-Fisher (vMF) mixture models to represent the distribution $p(\omega_i)$: $\mathcal{V}(\omega;\Theta) = \sum_{k=1}^K \pi_k v(\omega;\Theta_k)$, with parameters $\Theta = \{\pi_k, \Theta_k\}_{k=1}^K$, where each $\pi_k \ge 0$, $\sum_k \pi_k = 1$, and $\Theta_k =$ $\{\mu_k, \kappa_k\}$ represents a vMF component with unit-length mean direction $\mu_k \in \mathbb{S}^{d-1}$ and concentration $\kappa_k \geq 0$. Unlike prior work that relies on a weighted EM algorithm [Vorba et al. 2014], we apply the standard EM algorithm, as our collected incident directions already

match the target distribution and do not require radiance-based weighting.

4.3 Splat ReSTIR samples

After the spatiotemporal reuse stage of ReSTIR PT, we collect the selected path samples at each pixel. Each sample contains the vertex position, surface normal, BSDF ID, and incident direction of every vertex along the path. We then splat all incident directions into a spatial hash grid according to these properties.

In each spatial cell, we perform the expectation step (E-step) of the EM algorithm. For each incident direction ω_n , we compute its responsibility $\gamma_k(\omega_n)$ with respect to each vMF component k based on the current mixture parameters. We then accumulate the sufficient statistics across N samples: the weighted sum of directions r_k and the total responsibility mass w_k for each component k. Full details of the E-step procedure are provided in the supplemental document. These statistics are then used in the maximization step (*M-step*) to update the vMF mixture model parameters in each cell.

4.4 Fit guiding distributions

After accumulating the sufficient statistics during the E-step, we proceed to fit the vMF mixture model in each spatial cell by performing the M-step of the EM algorithm. Our implementation uses four vMF components per spatial cell, balancing expressiveness and real-time performance, since more components increase runtime cost and fitting complexity.

For each component k, we estimate the new mean direction μ_k , concentration κ_k , and mixture weight π_k , using the sufficient statistics r_k , w_k from the E-step. Similar to prior work [Ruppert et al. 2020], these parameters are computed with appropriate regularization and clamping to ensure robustness. The resulting set $\Theta_k = \mu_k, \kappa_k$ and π_k define the directional guiding distribution for the cell, and are stored for use in the subsequent frame for initial candidate generation. Please refer to the supplemental document for the full update formulas.

4.5 Guide initial candidate generation

Finally, we use the fitted vMF mixture models from the previous frame to guide the generation of initial candidate samples in the current frame. No additional warm-up phase is required.

Following prior path guiding practices, we augment BSDF sampling with guiding from our fitted vMF mixture models. During path construction, we randomly choose between the two strategies with a fixed probability $\alpha=0.5$. To ensure unbiasedness and robustness when the guiding distribution is poor, we combine the two using one-sample MIS [Veach and Guibas 1995].

5 Results

We implement our method using the Falcor framework [Kallweit et al. 2022], building upon the original implementation of ReSTIR PT [Lin et al. 2022] and a path tracer with next-event estimation (NEE) enabled. We focus only on sampling of indirect lighting; for direct lighting, we employ ReSTIR Direct Illumination (ReSTIR DI) [Bitterli et al. 2020]. Results include both components to demonstrate the method's effectiveness in practice. We disable Russian Roulette and instead cap the bounce count. The maximum number of bounces is selected per scene to balance performance and quality: 9 for Barcelona Pavilion, 4 for Veach Ajar, 8 for Veach Bidir Room, 4 for Living Room, 6 for Bistro, and 3 for Sun Temple. All tests are performed at a resolution of 1920 × 1080 on a machine equipped with an NVIDIA RTX 6000 Ada GPU, Intel i9-10900K CPU, and 64GB of RAM. We report error values using the mean absolute percentage error (MAPE)⁴.

Runtime performance and memory overhead. We report runtime performance of all methods on the result image slices. Our method introduces non-negligible overhead but still achieves real-time performance: $\sim 25\%$ from collecting and splatting samples, $\sim 25\%$ from fitting, and $\sim 50\%$ from sampling. Runtime increases with bounce count and the number of valid cells. Each pixel stores ReSTIR samples recording incident directions, positions, normals, and material IDs of all path vertices, which at 1920×1080 amounts to ~ 102.8 MB per bounce. Directional distributions are stored in the hash grid, with each valid cell holding four vMF lobes and their sufficient statistics (152 bytes). The total memory of valid cells varies by scene, e.g., ~ 5 MB for Veach AJAR and ~ 76 MB for Bistro.

Improved initial candidates. Our method improves the quality of ReSTIR's initial candidates by integrating path guiding. As shown in Fig. 6 and Table 1, compared to the original candidates (PT samples) in ReSTIR PT, our approach produces a noticeably higher density of high-contribution samples, both visually and as evidenced by lower MAPE values. In particular, for the BARCELONA PAVILION scene, our initial candidates better capture the shadow beneath the glass windows.

Reduced variance and correlation artifacts. Our method fits guiding distributions using high-quality ReSTIR samples. For comparison, we adapt two prior path guiding methods [Derevyannykh 2022; Ruppert et al. 2020], which rely on raw path-traced samples to guide initial candidate generation. As shown in Fig. 5 and Fig. 8, our approach produces lower variance and significantly reduces correlation artifacts while maintaining a real-time performance. All methods are evaluated with ReSTIR PT enabled and rendered at 1spp. The result images in Fig. 5 and Fig. 8 are captured in a static

Table 1. Average number of contributing ReSTIR samples and raw path-traced (PT) samples per frame across spatial cells in each scene. ReSTIR consistently provides significantly more effective samples than raw PT.

Scene	ReSTIR Samples	PT Samples
Barcelona Pavilion	1,140	535
VEACH AJAR	3,055	164
Veach Bidir Room	3,326	1,372
LIVING ROOM	2,282	95
Bistro	241	35
Sun Temple	1,140	342

scene with a static camera. Note that even without path guiding, ReSTIR PT leverages spatiotemporal reuse and hybrid shifts to propagate occasionally discovered caustic paths, making the caustics visible (see Fig. 5). However, when too few pixels initially capture such paths and there are not enough neighboring samples to shift from, the caustics become barely perceptible (see Fig. 1). By contrast, our method introduces more caustic paths into the initial candidates through path guiding. We further evaluate the methods under dynamic lighting in Fig. 7. The baselines exhibit severe correlation artifacts (e.g., boiling blobs), whereas our method produces noticeably cleaner results. Please refer to the supplementary video for additional comparisons under motion.

6 Conclusion

We introduced *ReSTIR PG*, a novel framework that integrates path guiding into ReSTIR PT. Prior path guiding methods rely on 1spp raw path-traced samples; in real-time rendering, such samples are often noisy and sparse in space. Observing that ReSTIR's resampled samples already approximate the ideal local guiding distribution, our method directly extracts directional distributions from these high-quality samples. This enables efficient and effective guiding of the next batch of initial candidates, reducing noise and suppressing correlation artifacts in ReSTIR's results, and effectively closing the loop between guiding and reuse.

Since we only fit a marginal PDF of next-bounce directions, guiding quality can be suboptimal on glossy materials. Naively fitting a full 4D distribution is often slow and inaccurate. Future work could explore more expressive and efficient fitting strategies—such as neural networks—to better extract information from ReSTIR samples and improve guiding. Our method could also benefit from enhancements to ReSTIR itself, including the use of sample mutations or improved temporal shift mappings.

Acknowledgments

The authors thank the anonymous reviewers for valuable feedback, Songyin Wu for helping create the test scenes, and Aaron Lefohn for supporting this work.

References

Lucas Domingo Alber, Johannes Hanika, and Carsten Dachsbacher. 2025. Real-Time Markov Chain Path Guiding for Global Illumination and Single Scattering. Proc. ACM Comput. Graph. Interact. Tech. 8, 1, Article 15 (May 2025), 18 pages. doi:10. 1145/3728296

⁴We compute MAPE as mean $\left(\frac{|I-I_{gt}|}{0.01 \cdot \text{mean}(\tilde{I}_{gt}) + \tilde{I}_{gt}}\right)$, where \tilde{I}_{gt} is a grayscale image.

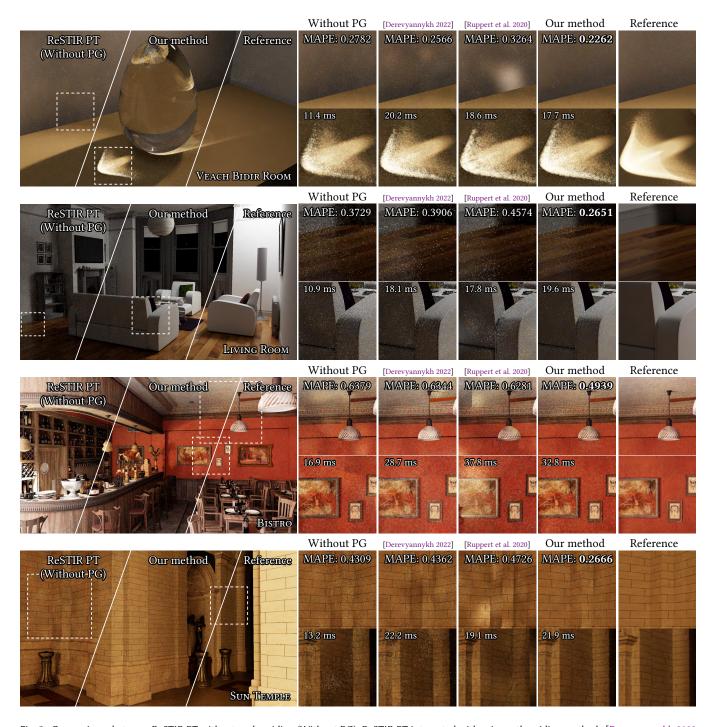


Fig. 5. Comparisons between ReSTIR PT without path guiding (Without PG), ReSTIR PT integrated with prior path guiding methods [Derevyannykh 2022; Ruppert et al. 2020], and our method. Note that all methods are evaluated with ReSTIR PT enabled and rendered with 1spp. Our method produces cleaner results with lower noise and fewer correlation artifacts. In all comparisons, our method achieves the lowest MAPE.

Steve Bako, Mark Meyer, Tony DeRose, and Pradeep Sen. 2019. Offline deep importance sampling for Monte Carlo path tracing. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 527–542. Thomas Bashford-Rogers, Kurt Debattista, and Alan Chalmers. 2012. A significance cache for accelerating global illumination. In Computer Graphics Forum, Vol. 31. Wiley Online Library, 1837-1851.

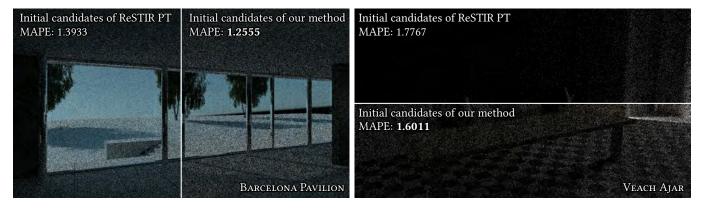


Fig. 6. Comparison of initial candidate samples from ReSTIR PT (1spp raw path-traced) and our ReSTIR-PG method (1spp path-guided). Our approach significantly improves the quality of initial candidates. As evidenced by the lower MAPE, our method produces cleaner candidates with fewer noisy and outlier samples. This directly contributes to reduced noise and correlation artifacts in the final rendered images after applying ReSTIR, as shown in Fig. 1. In the VEACH AJAR scene, we changed the right-most teapot from transparent to glossy to better showcase our method's performance on glossy materials.



Fig. 7. Comparisons of ReSTIR PT without path guiding (Without PG), ReSTIR PT integrated with prior path guiding methods [Derevyannykh 2022; Ruppert et al. 2020], and our method on the Veach Ajar Box Light scene, featuring a challenging moving box light. All methods are evaluated with ReSTIR PT enabled and rendered at 1 sample per pixel (1spp). Our method produces cleaner results with lower noise and fewer correlation artifacts, and achieves the lowest MAPE.



Fig. 8. Comparison of ReSTIR PT without path guiding (Without PG), ReSTIR PT with prior path guiding methods [Derevyannykh 2022; Ruppert et al. 2020], and our method on the Veach Ajar Glossy scene, a modified version of Veach Ajar with more glossy surfaces, showing glossy inter-reflection. All methods are evaluated with ReSTIR PT enabled and rendered at 1 sample per pixel (1 spp). Even on glossy surfaces, our method yields cleaner results with lower noise, fewer correlation artifacts, and the lowest MAPE.

Benedikt Bitterli, Chris Wyman, Matt Pharr, Peter Shirley, Aaron E. Lefohn, and Wojciech Jarosz. 2020. Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting. 39, 4, Article 148 (Jul 2020), 17 pages. doi:10.1145/3386569.3392481 Guillaume Boissé. 2021. World-Space Spatiotemporal Reservoir Reuse for Ray-Traced Global Illumination. In SIGGRAPH Asia 2021 Technical Communications (Tokyo, Japan) (SA '21 Technical Communications). Association for Computing Machinery, New York, NY, USA, Article 22, 4 pages. doi:10.1145/3478512.3488613

- Jakub Boksansky, Paula Jukarainen, and Chris Wyman. 2021. Rendering Many Lights with Grid-Based Reservoirs. In Ray Tracing Gems II. Springer, 351-365
- Wesley Chang, Venkataram Sivaram, Derek Nowrouzezahrai, Toshiya Hachisuka, Ravi Ramamoorthi, and Tzu-Mao Li. 2023. Parameter-space ReSTIR for differentiable and inverse rendering. In ACM SIGGRAPH 2023 Conference Proceedings. 1-10.
- Mikhail Derevyannykh. 2022. Real-Time Path-Guiding Based on Parametric Mixture Models. In Eurographics 2022 - Short Papers, Nuria Pelechano and David Vanderhaeghe (Eds.). The Eurographics Association. doi:10.2312/egs.20221024
- Stavros Diolatzis, Adrien Gruson, Wenzel Jakob, Derek Nowrouzezahrai, and George Drettakis. 2020. Practical product path guiding using linearly transformed cosines. In Computer graphics forum, Vol. 39. Wiley Online Library, 23-33.
- Addis Dittebrandt, Johannes Hanika, and Carsten Dachsbacher. 2020. Temporal sample reuse for next event estimation and path guiding for real-time path tracing. (2020).
- Addis Dittebrandt, Vincent Schüßler, Johannes Hanika, Sebastian Herholz, and Carsten Dachsbacher. 2023. Markov Chain Mixture Models for Real-Time Direct Illumination. In Computer Graphics Forum, Vol. 42. Wiley Online Library, e14881.
- Ana Dodik, Marios Papas, Cengiz Öztireli, and Thomas Müller. 2022. Path Guiding Using Spatio-Directional Mixture Models. In Computer Graphics Forum, Vol. 41. Wiley Online Library, 172-189.
- Honghao Dong, Guoping Wang, and Sheng Li. 2023. Neural parametric mixtures for path guiding. In ACM SIGGRAPH 2023 Conference Proceedings. 1-10.
- Sebastian Herholz, Oskar Elek, Jiří Vorba, Hendrik Lensch, and Jaroslav Křivánek. 2016. Product importance sampling for light transport path guiding. In Computer Graphics Forum, Vol. 35. Wiley Online Library, 67-77.
- Heinrich Hey and Werner Purgathofer. 2002. Importance sampling with hemispherical particle footprints. In Proceedings of the 18th spring conference on Computer graphics. 107-114.
- Iiawei Huang, Akito Iizuka, Hajime Tanaka, Taku Komura, and Yoshifumi Kitamura, 2024. Online Neural Path Guiding with Normalized Anisotropic Spherical Gaussians. ACM Transactions on Graphics 43, 3 (2024), 1-18.
- Henrik Wann Jensen. 1995. Importance driven path tracing using the photon map. In Rendering Techniques' 95: Proceedings of the Eurographics Workshop in Dublin, Ireland, June 12–14, 1995 6. Springer, 326–335.
- Simon Kallweit, Petrik Clarberg, Craig Kolb, Tom'aš Davidovič, Kai-Hwa Yao, Theresa Foley, Yong He, Lifan Wu, Lucy Chen, Tomas Akenine-Möller, Chris Wyman, Cyril Crassin, and Nir Benty. 2022. The Falcor Rendering Framework. https://github. $com/NVIDIAGameWorks/Falcor\ https://github.com/NVIDIAGameWorks/Falcor.$
- René Kern, Felix Brüll, and Thorsten Grosch. 2024. ReSTIR FG: Real-Time Reservoir Resampled Photon Final Gathering. In Eurographics Symposium on Rendering, Eric Haines and Elena Garces (Eds.). The Eurographics Association. doi:10.2312/sr. 20241155
- Markus Kettunen, Daqi Lin, Ravi Ramamoorthi, Thomas Bashford-Rogers, and Chris Wyman. 2023. Conditional Resampled Importance Sampling and ReSTIR. In SIG-GRAPH Asia 2023 Conference Papers. 1-11.
- Eric P Lafortune and Yves D Willems. 1995. A 5D tree to reduce the variance of Monte Carlo ray tracing. In Rendering Techniques' 95: Proceedings of the Eurographics Workshop in Dublin, Ireland, June 12-14, 1995 6. Springer, 11-20.
- Tianyu Li and Xiaoxin Guo. 2024. Real-time Seamless Object Space Shading. (2024).
- Daqi Lin, Markus Kettunen, Benedikt Bitterli, Jacopo Pantaleoni, Cem Yuksel, and Chris Wyman. 2022. Generalized Resampled Importance Sampling: Foundations of ReSTIR. 41, 4 (2022), 75:1-75:23. doi:10.1145/3528223.3530158
- Daqi Lin, Chris Wyman, and Cem Yuksel. 2021. Fast Volume Rendering with Spatiotemporal Reservoir Resampling. 40, 6 (2021), 279:1-279:16. doi:10.1145/3478513.3480499
- Haolin Lu, Wesley Chang, Trevor Hedstrom, and Tzu-Mao Li. 2024. Real-Time Path Guiding Using Bounding Voxel Sampling. ACM Transactions on Graphics (TOG) 43, 4 (2024), 1-14
- Zander Majercik, Thomas Müller, Alexander Keller, Derek Nowrouzezahrai, and Morgan McGuire. 2021. Dynamic diffuse global illumination resampling. In ACM SIGGRAPH
- Thomas Müller, Markus Gross, and Jan Novák. 2017. Practical path guiding for efficient light-transport simulation. In Computer Graphics Forum, Vol. 36. Wiley Online Library, 91-100.
- Thomas Müller, Brian McWilliams, Fabrice Rousselle, Markus Gross, and Jan Novák. 2019. Neural importance sampling. ACM Transactions on Graphics (ToG) 38, 5 (2019),
- Yaobin Ouyang, Shiqiu Liu, Markus Kettunen, Matt Pharr, and Jacopo Pantaleoni. 2021. ReSTIR GI: Path Resampling for Real-Time Path Tracing. Computer Graphics Forum 40, 8 (2021), 17-29. doi:10.1111/cgf.14378
- Xingyue Pan, Jiaxuan Zhang, Jiancong Huang, and Ligang Liu. 2024. Enhancing Spatiotemporal Resampling with a Novel MIS Weight. In Computer Graphics Forum, Vol. 43. Wiley Online Library, e15049.
- Jacopo Pantaleoni. 2020. Online path sampling control with progressive spatio-temporal filtering. SN Computer Science 1, 5 (2020), 279.
- Alexander Rath, Pascal Grittmann, Sebastian Herholz, Petr Vévoda, Philipp Slusallek, and Jaroslav Křivánek. 2020. Variance-aware path guiding. ACM Transactions on Graphics (TOG) 39, 4 (2020), 151-1.

- Florian Reibold, Johannes Hanika, Alisa Jung, and Carsten Dachsbacher. 2018. Selective guided sampling with complete light transport paths. ACM Transactions on Graphics (TOG) 37, 6 (2018), 1-14.
- Lukas Ruppert, Sebastian Herholz, and Hendrik PA Lensch. 2020. Robust fitting of parallax-aware mixtures for path guiding. ACM Transactions on Graphics (TOG) 39,
- Rohan Sawhney, Dagi Lin, Markus Kettunen, Benedikt Bitterli, Ravi Ramamoorthi, Chris Wyman, and Matt Pharr. 2024. Decorrelating restir samplers via mcmc mutations. ACM Transactions on Graphics 43, 1 (2024), 1-15.
- Vincent Schüßler, Johannes Hanika, Alisa Jung, and Carsten Dachsbacher. 2022. Path guiding with vertex triplet distributions. In Computer Graphics Forum, Vol. 41. Wiley Online Library, 1-15.
- Justin Talbot, David Cline, and Parris Egbert. 2005. Importance Resampling for Global Illumination. In Eurographics Symposium on Rendering. 139-146. doi:10.2312/EGWR/
- Yusuke Tokuyoshi. 2023. Efficient spatial resampling using the pdf similarity. Proceedings of the ACM on Computer Graphics and Interactive Techniques 6, 1 (2023),
- Eric Veach and Leonidas J Guibas. 1995. Optimally combining sampling techniques for Monte Carlo rendering. In Proceedings of the 22nd annual conference on Computer graphics and interactive techniques. 419-428.
- Iiří Vorba, Ondřei Karlík, Martin Šik, Tobias Ritschel, and Jaroslav Křivánek, 2014. On-line learning of parametric mixture models for light transport simulation. ACM Transactions on Graphics (TOG) 33, 4 (2014), 1-11.
- Yu-Chen Wang, Chris Wyman, Lifan Wu, and Shuang Zhao. 2023. Amortizing samples in physics-based inverse rendering using ReSTIR. ACM Transactions on Graphics (TOG) 42, 6 (2023), 1-17,
- Alexander Weinrauch, Wolfgang Tatzgern, Pascal Stadlbauer, Alexis Crickx, Jozef Hladky, Arno Coomans, Martin Winter, Joerg H. Mueller, and Markus Steinberger. 2023. Effect-based Multi-viewer Caching for Cloud-native Rendering. ACM Trans. Graph. 42, 4, Article 87 (July 2023), 16 pages. doi:10.1145/3592431
- Mirco Werner, Vincent Schüßler, and Carsten Dachsbacher. 2024. ReSTIR Subsurface $Scattering \ for \ Real-Time \ Path \ Tracing. \ Proceedings \ of \ the \ ACM \ on \ Computer \ Graphics$ and Interactive Techniques 7, 3 (2024), 1-19.
- Chris Wyman, Markus Kettunen, Daqi Lin, Benedikt Bitterli, Cem Yuksel, Wojciech Jarosz, Pawel Kozlowski, and Giovanni De Francesco. 2023. A Gentle Introduction to ReSTIR: Path Reuse in Real-time. In ACM SIGGRAPH 2023 Courses (Los Angeles, California). doi:10.1145/3587423.3595511
- Chris Wyman and Alexey Panteleev. 2021. Rearchitecting Spatiotemporal Resampling for Production. In ACM/EG Symposium on High Perfrormance Graphics. 23-41. doi:10. 2312/hpg.20211281
- 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. doi:10.1145/3658210
- Song Zhang, Daqi Lin, Chris Wyman, and Cem Yuksel. 2025. Many-Light Rendering Using ReSTIR-Sampled Shadow Maps. Computer Graphics Forum 44, 2 (2025), e70059. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.70059 doi:10.1111/
- Shilin Zhu, Zexiang Xu, Tiancheng Sun, Alexandr Kuznetsov, Mark Meyer, Henrik Wann Jensen, Hao Su, and Ravi Ramamoorthi. 2021. Photon-driven neural reconstruction for path guiding. ACM Transactions on Graphics (TOG) 41, 1 (2021), 1-15.