Specular Manifold Sampling for Rendering High-Frequency Caustics and Glints

TIZIAN ZELTNER, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
ILIYAN GEORGIEV, Autodesk, United Kingdom
WENZEL JAKOB, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

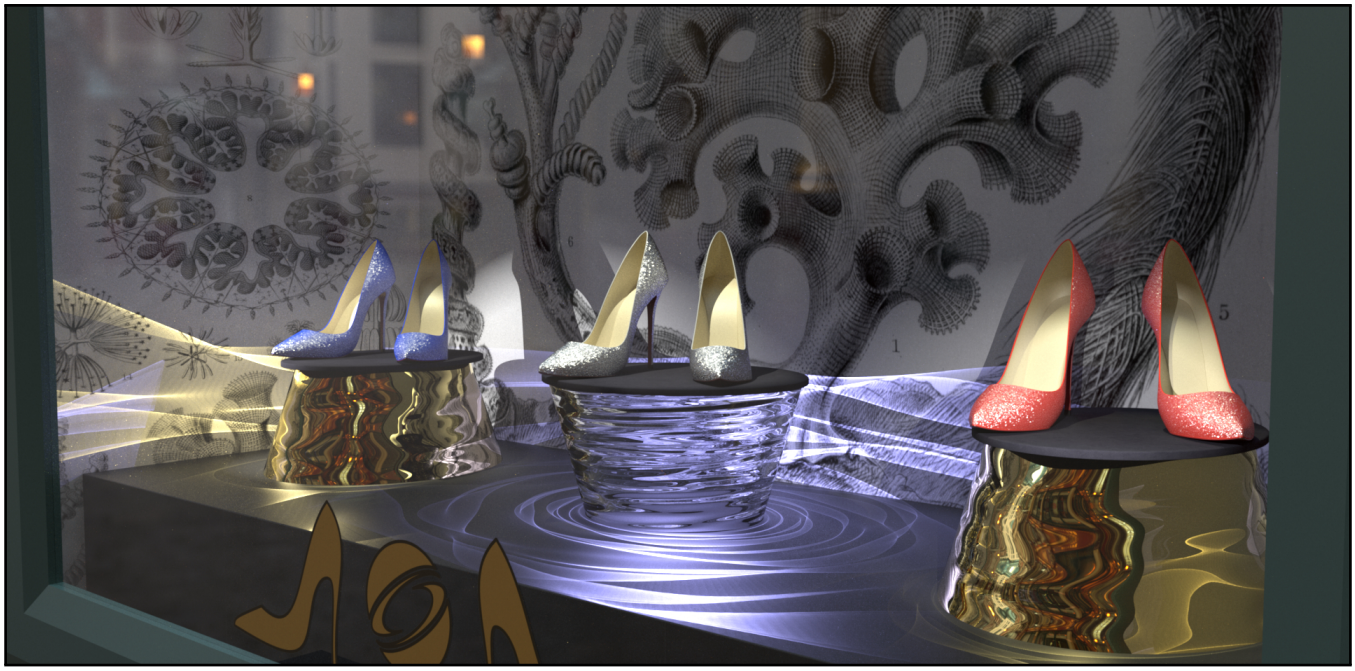

Fig. 1. Rendering of a shop window featuring a combination of challenging-to-sample light transport paths with specular-diffuse-specular ("SDS") interreflection: the two golden normal-mapped pedestals are illuminated by spot lights and project intricate caustic patterns following a single reflection from the metallic surface, while the transparent center pedestal generates caustics via double refraction. The glinty appearance of the shoes arises due to specular microgeometry encoded in a high-frequency normal map. This image was rendered by an ordinary unidirectional path tracer using our new specular manifold sampling strategy. The remaining noise is due to indirect lighting by caustics, which is not explicitly sampled by our technique. The background image is "Hexactinellae" from *Art Forms in Nature* by Ernst Haeckel.

Scattering from specular surfaces produces complex optical effects that are frequently encountered in realistic scenes: intricate caustics due to focused reflection, multiple refraction, and high-frequency glints from specular microstructure. Yet, despite their importance and considerable research to this end, sampling of light paths that cause these effects remains a formidable challenge.

In this article, we propose a surprisingly simple and general sampling strategy for specular light paths including the above examples, unifying the previously disjoint areas of caustic and glint rendering into a single framework. Given two path vertices, our algorithm stochastically finds a specular subpath connecting the endpoints. In contrast to prior work, our method supports high-frequency normal- or displacement-mapped geometry, samples specular-diffuse-specular ("SDS") paths, and is compatible with standard Monte Carlo methods including unidirectional path tracing. Both unbiased and biased variants of our approach can be constructed, the latter often significantly reducing variance, which may be appealing in applied settings (e.g. visual effects). We demonstrate our method on a range of challenging scenes and evaluate it against state-of-the-art methods for rendering caustics and glints.

CCS Concepts: • **Computing methodologies** → **Rendering**.

Additional Key Words and Phrases: Specular light paths, SDS paths, Caustics, Glints

ACM Reference Format:
Tizian Zeltner, Iliyan Georgiev, and Wenzel Jakob. 2020. Specular Manifold Sampling for Rendering High-Frequency Caustics and Glints. *ACM Trans. Graph.* 39, 4, Article 149 (July 2020), 15 pages. https://doi.org/10.1145/3386569.3392408

Authors' addresses: Tizian Zeltner, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland, tizian.zeltner@epfl.ch; Iliyan Georgiev, Autodesk, United Kingdom, iliyan.georgiev@autodesk.com; Wenzel Jakob, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland, wenzel.jakob@epfl.ch.

1 INTRODUCTION

Efficient simulation of the physical process of light transport and scattering constitutes one of the classic problems in the field of computer graphics. Over the past decades, research on this challenging problem has gravitated towards methods that perform Monte Carlo integration over high-dimensional sets of light transport paths, fueling the development of a toolbox of sophisticated path sampling strategies that are now in widespread use. While Monte Carlo rendering has been extremely successful in both industry and academia, there are many settings where these methods are surprisingly brittle and subject to catastrophically poor convergence.

One particularly problematic case are scenes containing *specular paths*, i.e., paths involving chains of interactions with smooth metallic or refractive surfaces. Their constrained nature makes such light paths difficult to find: at specular interfaces, light must satisfy the law of reflection or Snell's law, which drastically lowers the probability of sampling a valid configuration connecting the camera to a light source. Caustics or glittery surfaces that exhibit random patterns of highlights due to specular micro-geometry are visually striking examples of such specular paths, though they are often much more subtle while still having a negative effect on overall convergence. The large family of *specular-diffuse-specular* ("SDS") paths with only a single diffuse vertex surrounded by specular chains simply cannot be found at all and is normally absent in rendered images. The caustics seen through the shop window in Fig. 1 are an example of this configuration—they are only visible thanks to the proposed sampling strategy. The glints on the shoes are even more challenging and involve *only specular vertices*: a point spotlight, perfectly specular microstructure, the glass window, and a pinhole camera.

Faced with these challenges, research in the last two decades has produced a large set of tailored methods that are each able to render specific types of specular paths. However, no existing method tackles the general case, and every approach brings along an intrinsic set of limitations: for instance, bidirectional sampling [Lafortune and Willems 1993; Veach and Guibas 1995a] from cameras and light sources often improves convergence, but does not help with glints or SDS paths, and tends to be brittle when the two partial paths only rarely find each other[1]. Biased methods [Jensen 1996; Hachisuka and Jensen 2009] that relax the problem via spatial smoothing support SDS paths and greatly improve convergence but do not handle glints and introduce blur that is often undesirable. Specialized glint rendering methods for normal-mapped surfaces [Yan et al. 2014, 2016] are excellent at their task but do not help with other types of specular paths and require many gigabytes of memory to hold precomputed data structures. Finally, path sampling techniques based on specular manifolds [Jakob and Marschner 2012; Hanika et al. 2015a] are very powerful in theory but tend to be extremely fragile in practice, particularly when the scene geometry features high-frequency detail.

In this paper, we present a remarkably simple technique for sampling specular chains connecting two specified shading points—including glints and SDS paths—in an unbiased manner. Our approach builds on the theory of specular manifolds but significantly improves its practicality in formerly challenging cases, e.g. when working with high-frequency displaced or normal-mapped geometry. Concretely, our contributions are:

(1) A unified manifold sampling strategy for rendering reflective and refractive caustics.
(2) A specialized variant for rendering glints, which reduces memory usage hundred-fold compared to prior work.
(3) A biased variant of the method with reduced variance.
(4) A two-pass sampling strategy for normal-mapped surfaces.
(5) Changes to the specular manifold constraints of Jakob and Marschner [2012] that improve robustness and convergence.

The main limitation of our sampling technique is that variance increases significantly for longer specular chains, hence our experiments mainly focus on short chains with 1 or 2 vertices. That said, our approach is highly extensible, and we believe that future work could address this limitation.

We show how our method can be deployed in a unidirectional path tracer and compare its performance to prior work. We will release our implementation as open source to ensure reproducibility.

2 PRIOR WORK

Path tracing, proposed in Kajiya's seminal 1986 paper, is the foundation of an extensive body of work on unbiased Monte Carlo rendering including later bidirectional extensions [Lafortune and Willems 1993; Veach and Guibas 1995a] and specialized guiding strategies that exploit past observations to improve the quality of generated samples [Vorba et al. 2014; Müller et al. 2017; Simon et al. 2018]. However, even state-of-the-art methods with both bidirectional sample generation and guiding fail to find SDS paths and produce high variance when the scene contains objects with glinty microstructure.

Methods based on photon maps [Jensen 1996] are able to resolve the issues with SDS paths by introducing spatial blur that relaxes the original problem. These methods emit photons from light sources in a first pass that are stored on surfaces and queried in a secondary density estimation phase. Later work has shown how very large number of photons can be sampled progressively [Hachisuka and Jensen 2009], and how photon lookups can be incorporated in bidirectional rendering algorithms [Georgiev et al. 2012; Hachisuka et al. 2012] combining many different strategies using multiple importance sampling [Veach and Guibas 1995b]. While photon maps are an excellent choice for certain path classes, they can introduce objectionable blur, and they do not handle important cases including caustics on non-diffuse surfaces or glints. The path space can also be selectively mollified to introduce bias only for path types that were otherwise impossible to sample [Kaplanyan and Dachsbacher 2013]. In contrast, our work focuses on general path sampling in the original non-relaxed problem.

Fermat's principle states that specular paths are *extremal*, i.e., they locally maximize or minimize the time that light requires to travel from one end to the other. One way of generating such paths thus entails optimizing path length or solving an equivalent root finding problem. This idea was pioneered by Mitchell and Hanrahan [1992], who render caustics from implicitly defined curved reflectors using interval arithmetic and Newton-Raphson iteration to find 1-bounce specular reflections connecting a given pair of vertices. We also experimented with a similar approach in an early stage of this

[1] e.g., in the common setting when the camera is inside a building that is lit from outside.

work and found that the highly conservative nature of interval arithmetic often causes intervals to be too large to allow for systematic pruning of the solution space. Another approach involves precomputed hierarchical data structures that partition triangle meshes and bound the positions and normals of each subtree. Building on such a position-normal hierarchy, Walter et al. [2009] propose an efficient pruning strategy for specular chains with a single specular refraction (e.g. underwater caustics).

Specialized methods for glint rendering implement an effective BRDF that averages the behavior of a high-frequency specular surface over a given surface region. The methods of Yan et al. [2014, 2016] do so by constructing position-normal hierarchies in texture space to solve an approximate form of this problem for normal-mapped surfaces. These tree data structures tend to become extremely large, requiring tens of gigabytes of memory in our experiments. Procedural glints [Jakob et al. 2014; Kuznetsov et al. 2019; Wang et al. 2019] can be rendered with a much lower storage footprint, but are of course significantly less flexible. Recent work on glint rendering has focused on incorporating wave-optical effects [Yan et al. 2018], and directional bases for filtering product integrals involving glints [Gamboa et al. 2018].

We note that the previously discussed caustic and glint rendering techniques of Mitchell and Hanrahan [1992], Walter et al. [2009], and Yan et al. [2014, 2016] find all possible solutions within their supported path classes and are thus fully deterministic. This yields pixel-perfect renderings when the scene contains no other sources of variance, but this is rarely the case in practice. In contrast, our method stochastically samples individual solutions; the added variance due to this random decision can then be reduced along with complementary sources of variance (environment mapping, indirect illumination, and so on).

Our algorithm is closely related to Jakob and Marschner's *manifold exploration* (ME) method [2012], which is a Markov chain Monte Carlo (MCMC) perturbation strategy that analyzes the differential geometry of the manifold of valid specular paths to make proposals in the framework of Metropolis Light Transport [Veach 1998]. The natural constraint representation [Kaplanyan et al. 2014; Hanika et al. 2015b] significantly expands on this idea in the more general case of multiple glossy interactions. One downside of these perturbation strategies is that they can only make very small changes to a specular path; exploration of the larger space relies on many repetitions of this basic operation. When the geometry is characterized by high-frequency detail, the manifold of valid specular paths tends to become very complex, and simple local steps are insufficient for global exploration.

The *manifold next-event estimation* (MNEE) technique of Hanika et al. [2015a] applies the ME equation-solving iteration in a pure Monte Carlo context. Starting with an incorrect initial specular path, MNEE iteratively attempts to walk towards a valid solution via projection and tangential steps on the specular manifold. Due to the simple deterministic initialization, it can find at most a single solution, which works well on smooth geometry (e.g. spheres) but breaks down in more challenging cases. MNEE was later adapted as a connection strategy for bidirectional path tracing [Speierer et al. 2018]. Other recent work has explored the potential of next-event estimation strategies that sample multiple vertices in the context of volume rendering [Koerner et al. 2016; Weber et al. 2017].

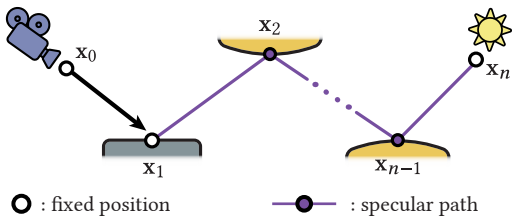

Fig. 2. Example specular path between the sensor and an emitter position.

The proposed method, named *specular manifold sampling* (SMS), is a generalization of the MNEE approach: using a stochastic initialization and an unbiased sampling weight estimator, we are able to find solutions on complex geometry where manifold-based techniques were previously inapplicable. We also propose a simple 2-stage stochastic initialization for normal-mapped surfaces and demonstrate that SMS straightforwardly generalizes to the related problem of glint rendering.

We provide a brief review of path-space manifolds [Jakob 2013] that are the foundation of our work before delving into specifics.

3 BACKGROUND

Consider a light transport path[2] $\bar{\mathbf{x}} = \mathbf{x}_0, \ldots, \mathbf{x}_n$ containing a chain of specular vertices between two non-specular endpoints $\mathbf{x}_1$ and $\mathbf{x}_n$. For example, $\mathbf{x}_1$ could be a shading point on a surface, and $\mathbf{x}_n$ a position on a light source (Fig. 2). Although this path is an element of a large and high-dimensional path space, it effectively lies on a much lower-dimensional subspace, since each specular vertex imposes physical constraints that collapse some dimensions of the ambient space: when refraction takes place at a vertex, it must for example satisfy Snell's law, removing all of its continuous degrees of freedom.

Jakob and Marschner [2012] characterize these constraints via a function $\mathbf{c}_i$ associated with each vertex $\mathbf{x}_i$. It projects Walter et al.'s [2007] generalized half-vector $\mathbf{h}(\mathbf{x}_i, \boldsymbol{\omega}_\mathrm{i}, \boldsymbol{\omega}_\mathrm{o})$ onto the local tangent space:

$$\mathbf{c}_i(\mathbf{x}_{i-1}, \mathbf{x}_i, \mathbf{x}_{i+1}) = \mathbf{T}(\mathbf{x}_i)^T \, \mathbf{h}(\mathbf{x}_i, \overrightarrow{\mathbf{x}_i\mathbf{x}_{i-1}}, \overrightarrow{\mathbf{x}_i\mathbf{x}_{i+1}}). \tag{1}$$

Here, $\mathbf{T}(\mathbf{x}_i)$ is a 3×2 matrix of tangent vectors, and $\overrightarrow{\mathbf{ab}}$ denotes a normalized vector pointing from $\mathbf{a}$ to $\mathbf{b}$. The constraint function takes on the value zero when the vertices $\mathbf{x}_{i-1}$, $\mathbf{x}_i$, and $\mathbf{x}_{i+1}$ are correctly positioned so that the relevant physical laws hold, subsuming both cases of specular reflection and refraction. A valid specular path can then be summarized as a root of the combined function $\mathbf{C}(\bar{\mathbf{x}}) = [\mathbf{c}_2, \ldots, \mathbf{c}_{n-1}]^T$ that stacks all specular path constraints. The function is normally parameterized via the UV coordinates $\mathbf{uv}_i$ of specular vertices, in which case its inputs and outputs have a matching number of dimensions.

The main application of this formulation is that it enables the use of numerical root finding techniques to solve for paths that lie

[2]In this article, overlined quantities like $\bar{\mathbf{x}}$ characterize an entire light path, while non-overlined quantities refer to other entities (e.g., individual vertices).

on the specular manifold. In particular, solutions to the equation $\mathbf{C}(\overline{\mathbf{uv}}) = 0$ can be found using multivariate Newton iteration:

$$\overline{\mathbf{uv}}_{i+1} = \overline{\mathbf{uv}}_i - (\nabla \mathbf{C}(\overline{\mathbf{uv}}_i))^{-1} \cdot \mathbf{C}(\overline{\mathbf{uv}}_i). \tag{2}$$

Here, $\overline{\mathbf{uv}}_i$ refers to the UV coordinates of all specular vertices at iteration i. The Jacobian $\nabla \mathbf{C}$ in (2) is a block tridiagonal matrix consisting of 2×2 blocks and can be inverted cheaply.

In practice, consistent UV parameterizations of scene geometry are often not available, and Jakob and Marschner instead propose to rely on local tangential parameterizations (e.g. of individual triangles). In this case, the Newton step (2) will generally produce vertices that are detached from the scene geometry, and an additional ray tracing-based reprojection step is required after each iteration. This predictor-corrector scheme, named *manifold walk*, has quadratic convergence when $\overline{\mathbf{uv}}_0$ is sufficiently close to a solution.[3]

We briefly note that another important role of the Jacobian $\nabla \mathbf{C}$ is the computation of the *generalized geometry term* $G(\mathbf{x}_1 \leftrightarrow \cdots \leftrightarrow \mathbf{x}_n)$ relating differential solid angles at one end of the chain ($\mathbf{x}_1$) to differential surface area at the other end ($\mathbf{x}_n$). This term takes the role of the traditional geometry term [Veach 1998] when a path contains specular vertices.

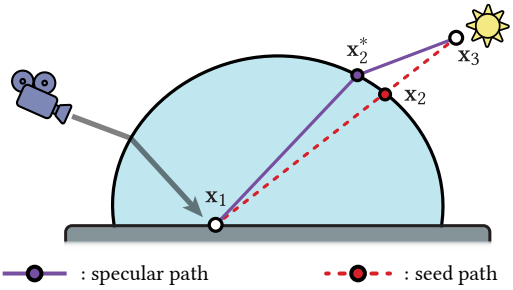

Fig. 3. Finding a specular refraction through a spherical interface using manifold next-event estimation. Note how the seed path is usually very close to the valid specular path in this scenario, which means that the manifold walk has a high success probability of converging to the desired solution.

3.1 Manifold next-event estimation

Manifold walks can also be applied in regular Monte Carlo methods, as illustrated by Hanika et al.'s [2015a] *manifold next-event estimation (MNEE)* method. Their approach is specific to refractive chains and initializes the Newton iteration with a tentative path $\bar{\mathbf{x}}_0$ corresponding to a straight-line connection (Fig. 3 and Alg. 1). Repeated evaluation of Eq. (2) along with reprojection then leads to two possible outcomes: either the manifold walk diverges, or it converges to a valid specular path. The probability of finding a valid path in this way is generally much higher than in an ordinary path tracer, reducing the variance of rendered images.

Prior to MNEE, the fundamental problem impeding the integration of manifold walks into pure Monte Carlo methods had been the determination of the probability $p(\bar{\mathbf{x}}^*)$ of finding a solution $\bar{\mathbf{x}}^*$ given an initial state $\bar{\mathbf{x}}_0$ chosen with density $p(\bar{\mathbf{x}}_0)$. This probability

[3]This is clearly the case in the original application of Jakob and Marschner [2012] whose MCMC perturbation strategy takes small steps on the specular manifold.

ALGORITHM 1: Manifold next-event estimation [Hanika et al. 2015a]

Input: Shading point $\mathbf{x}_1$ and emitter position $\mathbf{x}_n$ with density $p(\mathbf{x}_n)$
Output: Estimate of radiance traveling from $\mathbf{x}_n$ to $\mathbf{x}_1$

$\mathbf{s} = (\mathbf{x}_2, \ldots, \mathbf{x}_{n-1}) \leftarrow$ collect specular vertices along ray $\mathbf{x}_1 \rightarrow \mathbf{x}_n$
$\mathbf{s}^* = (\mathbf{x}_2^*, \ldots, \mathbf{x}_{n-1}^*) \leftarrow$ `manifold_walk`$(\mathbf{x}_1, \mathbf{s}, \mathbf{x}_n)$
return $f_s(\mathbf{s}^*) \cdot G(\mathbf{x}_1^* \leftrightarrow \cdots \leftrightarrow \mathbf{x}_n^*) \cdot L_e(\mathbf{x}_n) \,/\, p(\mathbf{x}_n)$

is given by a marginalization over the path space $\mathcal{P}$ [Hanika et al. 2015a]:

$$p(\bar{\mathbf{x}}^*) = \int_{\mathcal{P}} p(\bar{\mathbf{x}}^* \mid \bar{\mathbf{x}}_0)\, p(\bar{\mathbf{x}}_0)\, \mathrm{d}\bar{\mathbf{x}}_0. \tag{3}$$

The computation of this expression is mathematically daunting because $p(\bar{\mathbf{x}}^* \mid \bar{\mathbf{x}}_0)$ encapsulates the behavior of an unpredictable equation-solving iteration. MNEE sidesteps this problem using two clever tricks: First, it always uses the same initialization, turning $p(\bar{\mathbf{x}}_0)$ into a Dirac delta function. Second, it relies on Multiple Importance Sampling (MIS) [Veach and Guibas 1995a] to fall back to standard path tracing when the manifold walk fails. In this case, the remaining term $p(\bar{\mathbf{x}}^* \mid \bar{\mathbf{x}}_0)$ cancels out in the estimator.

However, MNEE also suffers from several fundamental limitations that we intend to address in this article: Because the method can find at most a single specular path, it does little to reduce variance when there are multiple solutions. Most of the examples in the original paper [Hanika et al. 2015a] are extremely simple shapes such as spheres[4] or cylinders, where a single solution indeed suffices to render most specular paths. However, this is clearly no longer the case when the specular geometry is more complex. Also, while the straight-line initialization $\bar{\mathbf{x}}_0$ proposed by Hanika et al. works well for refractive caustics, it is unclear how this approach could be generalized to the reflection case. Finally, MNEE does not support paths that give rise to specular glints.

4 METHOD

We now turn to the proposed method named *specular manifold sampling* (SMS), initially focusing on a simple unbiased algorithm that generalizes to cases where multiple specular paths connect two given endpoints. We then propose several extensions: the first reduces variance at the cost of nonzero bias, the second replaces Jakob and Marschner's specular manifold constraints with improved variants, the third samples paths in two stages to improve performance on normal-mapped surfaces, and the last streamlines the algorithm for glint rendering. Many of these extensions are modular and can be combined as desired. Figure 6 shows a preview of several combinations applied to the problem of rendering refractive caustics, comparing our results to MNEE and a brute-force reference.

4.1 Finding all solutions

We initially restrict ourselves to specular chains with a single smooth reflection or refraction. Our technique however extends to chains with surface roughness and multiple interactions analogously to Hanika et al. [2015a]; see Section 5.2. In this restricted setting, there

[4]Rendering of caustics in 3D models of human eyes constitutes an important use case of MNEE in the entertainment industry [Chiang and Burley 2018].

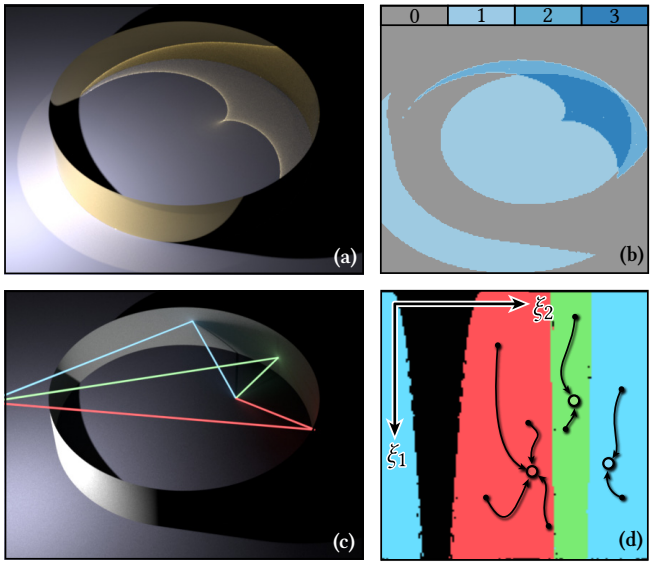

Fig. 4. **(a)**: Multiple solutions of the specular path constraint form a superposition of caustics in the Ring scene. **(b)**: Color map showing the number of solutions at each shading point. **(c)**: The three solution paths at a particular point. **(d)**: The basins of convergence (in primary sample space) corresponding to those solution paths. All manifold walks started at a point (black dot) inside a region converge to the associated solution (colored dot).

is a discrete and finite set[5] of solutions connecting two given endpoints. Our approach to finding them is simple: whereas MNEE performs manifold walks using a fixed initialization, SMS randomly samples the initial guess from a probability distribution $p(\bar{\mathbf{x}}_0)$. Newton's method exhibits quadratic convergence when the starting point is sufficiently close to a root, hence all solutions will be found with a nonzero probability—however, the probability of successful convergence is unknown.

There is a great degree of latitude in the choice of an initial guess: we could uniformly generate positions on specular surfaces, or sample the BSDF of the preceding vertex $\mathbf{x}_1$ and find starting points via ray tracing. Regardless of what approach is used, we assume that the implementation of this sampling strategy takes two uniformly distributed random numbers $\boldsymbol{\xi} = (\xi_1, \xi_2) \in [0,1)^2 =: \mathcal{U}$ as input and warps them to the desired distribution. Our objective now is to generate an initial sample $\boldsymbol{\xi}$ that is close enough to the solution, so that the Newton iteration (2) will take us there.

When inspecting the convergence behavior of these manifold walks on the *primary sample space* $\mathcal{U}$, we generally observe multiple *basins of convergence* $\mathcal{B}_k \subseteq \mathcal{U}$, each containing a point $\boldsymbol{\xi}^{(k)}$ identified with a corresponding solution vertex $\mathbf{x}_2{}^{(k)}$. Figure 4 illustrates the situation on a simple scene with a cardioid caustic. Newton's method is known to produce convergence basins that

[5]We note that cases with a continuous 1D subspace of solutions can be constructed, for instance when $\mathbf{x}_1$ and $\mathbf{x}_3$ lie at the center-line of a perfect cylindrical mirror. This is of no relevance for rendering natural scenes, since an arbitrarily small perturbation of the surface geometry would break the symmetries that are needed to create a 1D solution subspace. We ignore this corner case similarly to prior work [Walter et al. 2009].

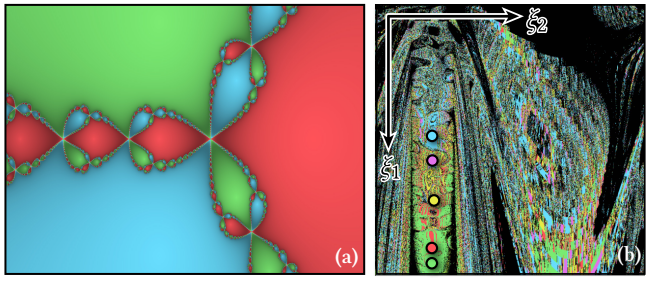

Fig. 5. **(a)**: Newton fractal $z^3 - 1 = 0$. **(b)**: Complex convergence basins with 5 unique solutions encountered in one of our scenes.

potentially have an extremely complex geometric structure [Hubbard et al. 2001] and can even be fractal. For example, Fig. 5a shows convergence towards three different solutions of a simple polynomial equation on the set of complex numbers and Fig. 5b shows a particularly challenging convergence plot that we encountered after applying a normal map to the Ring scene shown in Fig. 4. Recall that our goal is to use these solutions in a Monte Carlo estimator, where an algorithm that merely finds solutions is insufficient—we must also know the discrete probability p_k of finding a particular path vertex $\mathbf{x}_2{}^{(k)}$. An unbiased estimator is then given by a standard MC ratio $f(\bar{\mathbf{x}})/p(\bar{\mathbf{x}})$, where $p(\bar{\mathbf{x}})$ in the denominator contains p_k.

Since the samples $\boldsymbol{\xi}$ are uniformly distributed, this probability is simply the area of the associated convergence basin on $\mathcal{U}$:

$$p_k = \int_{\mathcal{U}} \mathbb{1}_{\mathcal{B}_k}(\boldsymbol{\xi}) \, \mathrm{d}\boldsymbol{\xi}. \tag{4}$$

However, exact evaluation or precomputation of this integral is clearly infeasible: as discussed, $\mathcal{B}_k$ can have an extremely complex shape which also depends on the position of the path endpoints. On the other hand, a simple unbiased estimator is given by

$$\langle p_k \rangle = \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}_{\mathcal{B}_k}(\boldsymbol{\xi}_i), \tag{5}$$

where $\boldsymbol{\xi}_i \in \mathcal{U}$ is a sequence of i.i.d. uniform variates. Unfortunately, this approach is flawed: usage of p_k occurs in the *denominator* of the path throughput weight. Since $\mathbb{E}[1/X] \neq 1/\mathbb{E}[X]$, using this estimator could introduce significant bias: for example, $\langle p_k \rangle$ can equal zero if all N tries fail to converge to the basin $\mathcal{B}_k$, in which case the estimated path throughput weight would be infinite! Fortunately, an unbiased estimator for the inverse $\langle 1/p_k \rangle$ can be created using a simple iterative approach.

4.2 Unbiased SMS

The problem of computing an unbiased MC estimate of the reciprocal of an integral was studied by Booth [2007]. Recently, Qin et al. [2015] built on this idea to create an unbiased photon gathering strategy. The key idea underlying Booth's approach is surprisingly simple: turning the inverse into a geometric series moves the problematic integral from the denominator to the numerator:

$$\frac{1}{p_k} = \frac{1}{\int_{\mathcal{U}} \mathbf{1}_{\mathcal{B}_k}(\boldsymbol{\xi}) \, \mathrm{d}\boldsymbol{\xi}} = \frac{1}{1-a} = \sum_{i=0}^{\infty} a^i, \tag{6}$$

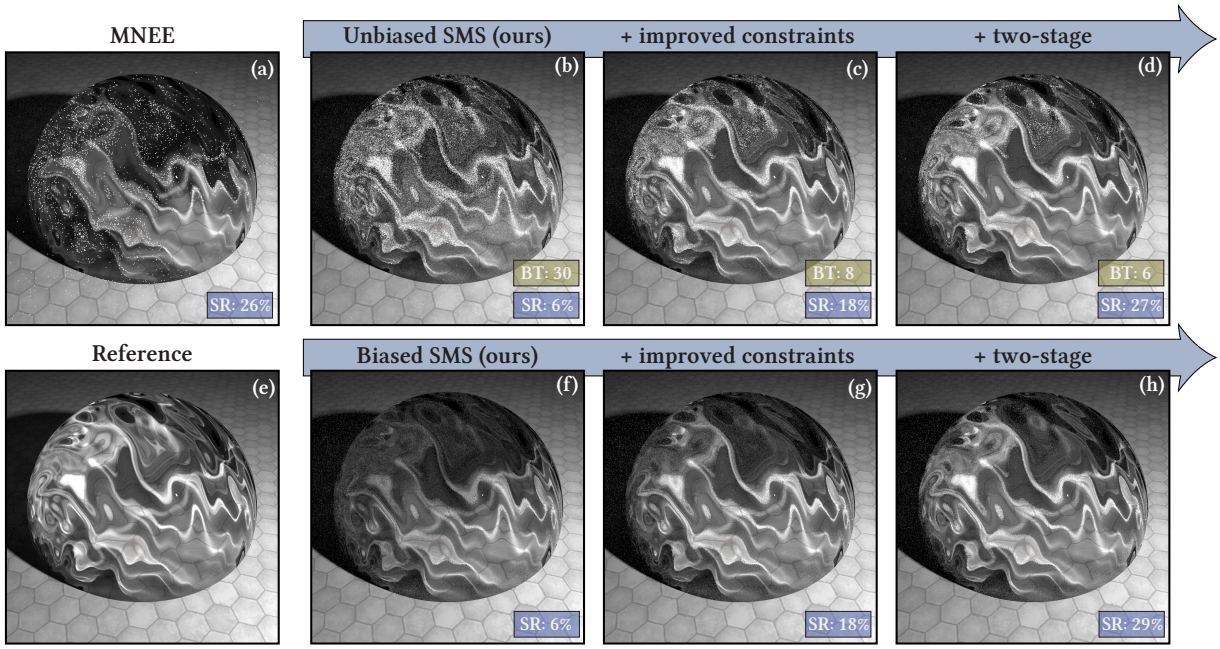

Fig. 6. Our two methods and their extensions illustrated on a normal-mapped dielectric sphere illuminated by a small area light; equal-time renderings (1 minute). Small insets summarize manifold walk success rate (SR) and average number of Bernoulli trials (BT), where applicable. **(a)**: Previous work (MNEE) fails to capture the full complexity of the caustic as it only finds at most one refractive path per shading point. For the remaining energy it falls back to path tracing with high variance. **(b)**: Our unbiased SMS method on its own. **(c)** - **(d)**: Adding the constraint and two-stage manifold walk improvements which increase the success rate while reducing the iteration count. **(f)** - **(h)**: Same sequence but using our biased SMS variant which suppresses the noise of the reciprocal probability estimation. As the success rate increases from left to right, bias goes down. The remaining bias mainly manifests itself in the regions where the unbiased counterpart remains noisy. The biased version uses a sample set size of $M = 16$. **(e)**: The path-traced reference was rendered for 5 hours.

where $a = 1 - \int_{\mathcal{U}} \mathbf{1}_{\mathcal{B}_k}(\xi)\,\mathrm{d}\xi$. This expansion is legal as long as $|a| < 1$, which in our case—integrating an indicator function over the unit square—is clearly satisfied. Unbiased estimation of the reciprocal then entails repeated manifold walks with i.i.d. initial points, denoted as $\langle a \rangle_j$ below:

$$\langle 1/p_k \rangle = 1 + \sum_{i=1}^{\infty} \prod_{j=1}^{i} \langle a \rangle_j. \tag{7}$$

Here, $\langle a \rangle_j = 0$ when manifold walk j has converged to root $\xi^{(k)}$ and $\langle a \rangle_j = 1$ if it has found another root or diverged. The above expression can thus be understood as a simple counting process: we run repeated manifold walks *until* $\xi^{(k)}$ is found, and the number of trials then provides an unbiased estimate of $1/p_k$. This result can also be understood in terms of the geometric distribution, which models the number of Bernoulli trials needed until a certain event with probability p_k takes place. Here again, the expected number of attempts is $1/p_k$. Simulating a geometric distribution therefore provides an unbiased estimator of the sought reciprocal and constitutes the base ingredient of our unbiased SMS scheme, which we lay out in Alg. 2.

Unbiased SMS is trivially added to any existing implementation of MNEE. Its runtime cost is directly linked to the "complexity" of specular paths in the scene: when the geometry is relatively smooth, $\mathcal{U}$ contains a small number of solutions that are surrounded by large convergence basins. Lines 1-2 then rapidly converge to a particular solution, and only a few iterations are required to find that same solution once more in lines 4-9. However, when the geometry is complex, many solutions may exist, and their convergence basins also occupy smaller area in primary sample space. The required number of trial iterations in lines 4-9 is a potential cause for concern in this case. Furthermore, the variance of such an estimator for a specific solution p_k based on a geometric distribution is equal to $(1 - p_k) \,/\, p_k^2$, which can become very large when $p_k \approx 0$.

ALGORITHM 2: Unbiased specular manifold sampling

Input: Shading point $\mathbf{x}_1$ and emitter position $\mathbf{x}_3$ with density $p(\mathbf{x}_3)$
Output: Estimate of radiance traveling from $\mathbf{x}_3$ to $\mathbf{x}_1$

1 $\mathbf{x}_2 \leftarrow$ sample a specular vertex as initial position
2 $\mathbf{x}_2^* \leftarrow$ `manifold_walk`$(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$
3 $\langle 1/p_k \rangle \leftarrow 1$ ▷ Estimate inverse probability of sampling $\mathbf{x}_2^*$
4 **while** *true* **do**
5 $\quad \mathbf{x}_2 \leftarrow$ sample specular vertex as above
6 $\quad \mathbf{x}_2' \leftarrow$ `manifold_walk`$(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$
7 $\quad$ **if** $\|\mathbf{x}_2' - \mathbf{x}_2^*\| < \varepsilon$ **then**
8 $\quad\quad$ **break**
9 $\quad \langle 1/p_k \rangle \leftarrow \langle 1/p_k \rangle + 1$
10 **return** $f_s(\mathbf{x}_2^*) \cdot G(\mathbf{x}_1 \leftrightarrow \mathbf{x}_2 \leftrightarrow \mathbf{x}_3) \cdot \langle 1/p_k \rangle \cdot L_e(\mathbf{x}_3) \,/\, p(\mathbf{x}_3)$

Another fundamental issue with unbiased SMS is that a considerable amount of computation is in some sense not used optimally: repeated Bernoulli trials potentially find many additional solutions, yet Alg. 2 only cares about one bit of information: whether or not these match the solution that was previously found in line 2. It would be desirable that the method leverages this additional information to improve the estimates, but it is challenging to do without introducing bias. That said, the presence of some bias is often acceptable if this improves other aspects, such as running time or variance, without introducing undesirable visual artifacts in renderings. Motivated by this, we devise a biased variant of SMS that addresses the discussed concerns.

4.3 Biased SMS

The biased variant of our method (Alg. 3) replaces the unbounded number of trial iterations by a fixed budget of M samples. Furthermore, instead of sampling one path, and then performing Bernoulli trials to estimate its reciprocal probability, we simply cluster the M samples into a set of unique solutions $\mathbf{x}_2{}^{(l)}$ $(l = 1, \dots, L)$. A biased estimate of the reciprocal is then given by the relative number of occurrences n_l of each solution, which also avoids the potential issues with a division by zero discussed earlier. The (biased) estimator for the total throughput at shading point $\mathbf{x}_1$ then becomes:

$$\frac{1}{M}\sum_{l=1}^{L} n_l \frac{f(\mathbf{x}_2{}^{(l)})}{p(\mathbf{x}_2{}^{(l)})} \approx \frac{1}{M}\sum_{l=1}^{L} n_l f(\mathbf{x}_2{}^{(l)}) \frac{M}{n_l} = \sum_{l=1}^{L} f(\mathbf{x}_2{}^{(l)}). \quad (8)$$

Compared to the original unbiased approach, this variant has a fixed iteration count, and it exploits the information provided by all samples. Note that it is *consistent* and will converge to the true solution as $M \to \infty$. In our experiments, we observe that the resulting bias is manifested as energy loss: regions of complex caustics that are only rarely found by the Newton iteration appear darker, while the unbiased variant produces the correct intensity at the cost of substantially increased variance and runtime. We believe that trading variance for energy loss in this way could be preferable in applied contexts (e.g. visual effects). Section 5 and the supplemental video both present a number of results contrasting the two methods.

When setting $M = 1$, biased SMS is closely related to a biased version of MNEE where no MIS is applied to account for paths that cannot be sampled using the manifold walk. The two algorithms are however not equivalent. Consider a situation where only one valid solution exists but the straight-line initialization doesn't converge to it. In this case, MNEE can produce arbitrarily high variance—or in case of point lights will miss the contribution entirely. SMS however will still find the solution with non-zero probability.

Another interesting aspect of the biased variant is that it generates many samples at once. Coherence in this computation could be amenable to vectorized execution using modern SIMD instruction sets, such as AVX512.

ALGORITHM 3: Biased specular manifold sampling

Input: Shading point $\mathbf{x}_1$, emitter position $\mathbf{x}_3$ with density $p(\mathbf{x}_3)$, and trial set size M
Output: Estimate of radiance traveling from $\mathbf{x}_3$ to $\mathbf{x}_1$

1 $S = \{\}$ ▷ Set of **unique** solutions
2 **for** $i = 1, \dots, M$ **do**
3 $\quad \mathbf{x}_2 \leftarrow$ sample a specular vertex as initial position
4 $\quad \mathbf{x}_2^* \leftarrow$ `manifold_walk`$(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$
5 $\quad S = S \cup \{\mathbf{x}_2^*\}$
6 `result` = 0 ▷ Accumulate contribution
7 **for** $k = 1, \dots,$ `size`(S) **do**
8 $\quad$ `result +=` $f_s(\mathbf{x}_2{}^{(k)}) \cdot G(\mathbf{x}_1 \leftrightarrow \mathbf{x}_2{}^{(k)} \leftrightarrow \mathbf{x}_3) \cdot L_e(\mathbf{x}_3) \,/\, p(\mathbf{x}_3)$
9 **return** `result`

4.4 Improved specular manifold constraints

One significant difference of our method compared to all previous applications of manifold walks is that we require the Newton solver to take very large steps starting from an invalid state. In contrast, MNEE renders refractive caustics with a straight-line initialization that is generally already very close to the final solution, and applications of manifold walks to MCMC rendering [Jakob and Marschner 2012; Lehtinen et al. 2013] only make small perturbations to existing valid paths, in which case the Newton iteration converges rapidly. During the development of our technique, we found that manifold walks would often converge surprisingly poorly when initialized randomly, which led to serious convergence issues even in the case of simple and smooth geometry (Fig. 8). We realized that these two aspects are related: when taking large steps, Newton iterations based on the original specular manifold constraints often produce invalid back-facing solutions that impede convergence.

The main issue here is how the specular manifold constraint in Eq. (1) encodes specular configurations via half-vector projections. While this term conveniently subsumes both reflective and refractive cases with one equation, the formulation does not distinguish between front- and back-facing solutions (Fig. 7).

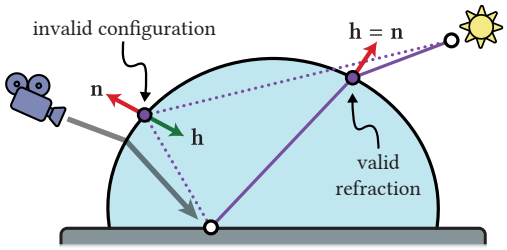

Fig. 7. The half-vector constraint of Jakob and Marschner [2012] often produces invalid back-facing solutions.

We propose a new constraint function that removes this ambiguity. After choosing between reflection and transmission, note that the law of reflection and Snell's law, respectively, fully determine the scattered direction $\mathrm{S}(\boldsymbol{\omega}, \mathbf{n}, \eta)$ given the incident direction $\boldsymbol{\omega}$, surface normal $\mathbf{n}$, and relative index of refraction η:

$$\mathrm{S}(\boldsymbol{\omega}, \mathbf{n}, \eta) = \begin{cases} \mathrm{S}^{\mathrm{r}}(\boldsymbol{\omega}, \mathbf{n}), & \text{if reflection,} \\ \mathrm{S}^{\mathrm{t}}(\boldsymbol{\omega}, \mathbf{n}, \eta), & \text{if transmission,} \end{cases} \quad \text{where} \quad (9)$$

$$\mathrm{S}^{\mathrm{r}}(\boldsymbol{\omega}, \mathbf{n}) = 2\langle \boldsymbol{\omega}, \mathbf{n} \rangle \mathbf{n} - \boldsymbol{\omega}, \quad (10)$$

$$\mathrm{S}^{\mathrm{t}}(\boldsymbol{\omega}, \mathbf{n}, \eta) = -\eta\,(\boldsymbol{\omega} - \langle \boldsymbol{\omega}, \mathbf{n} \rangle \mathbf{n}) - \mathbf{n}\sqrt{1 - \eta^2(1 - \langle \boldsymbol{\omega}, \mathbf{n} \rangle^2)}. \quad (11)$$

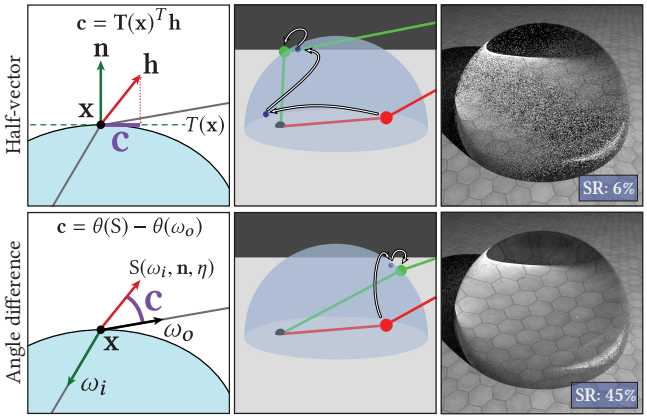

Fig. 8. Comparison between the previous half-vector based constraint function (**top**) and our new one based on angle differences (**bottom**). **Left**: 2D illustrations of the constraints. **Center**: For a random initial point (red) on the sphere, the half-vector constraint converges to an invalid configuration after three iterations. Our improved formulation results in better global convergence of the Newton solver. **Right**: Corresponding equal-time renderings on a dielectric sphere lit by a small area light source. Manifold walks are started at uniformly sampled positions over the glass surface. The walk success rate (SR) in the insets is measured over the entire image, where pixels outside the caustic count as sampling failures.

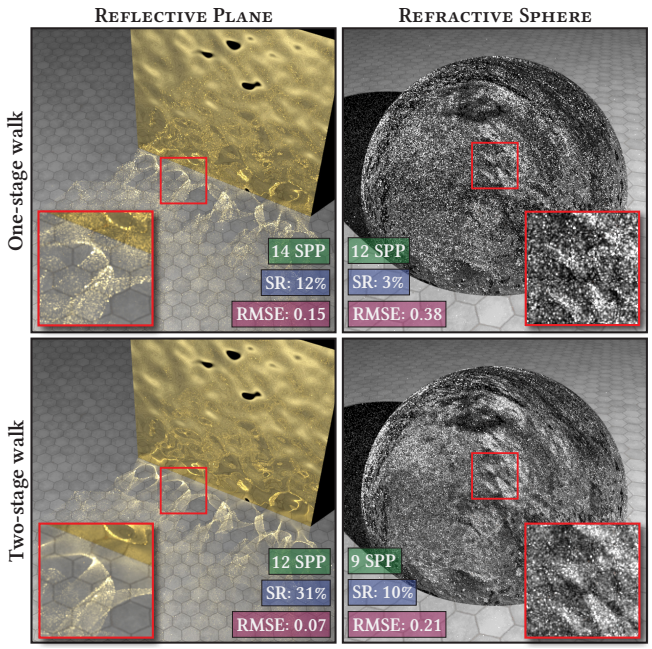

Fig. 9. Two scenes showing specular reflection and refraction on normal-mapped surfaces, rendered with unbiased SMS. We show equal-time comparison (1 minute) between starting points distributed uniformly over the shapes (**top**) and our two-pass manifold walks with improved starting points (**bottom**). Insets show computed number of samples per pixel (SPP), root mean square error (RMSE) compared to a converged reference, and manifold walk success rate (SR) measured over all pixels. Note how the cost of the additional manifold walk is amortized by a faster probability estimate resulting in a similar number of computed samples between the two methods.

When a vertex is in a valid physical configuration, its incident and outgoing directions satisfy $\boldsymbol{\omega}_\mathrm{o} = \mathrm{S}(\boldsymbol{\omega}_\mathrm{i}, \mathbf{n}, \eta)$. Another way of encoding specular manifold constraints thus entails measuring the difference between $\boldsymbol{\omega}_\mathrm{o}$ and $\mathrm{S}(\boldsymbol{\omega}_\mathrm{i}, \mathbf{n}, \eta)$. We experimented with different implementations of such a distance function, including simple 1D angle measurements $\angle(\boldsymbol{\omega}_\mathrm{o}, \mathrm{S}(\boldsymbol{\omega}_\mathrm{i}, \mathbf{n}, \eta))$ that produce a non-square Jacobian $\nabla \mathbf{C}$ requiring the use of a pseudoinverse in Newton iterations. Ultimately, we found that 2D constraints $\mathbf{c}_i : \mathbb{R}^2 \to \mathbb{R}^2$ exhibit better behavior, and we therefore measure a difference in the spherical coordinates of both vectors[6]:

$$\mathbf{c}_i = \begin{pmatrix} \theta(\mathrm{S}(\overrightarrow{\mathbf{x}_i\mathbf{x}_{i-1}}, \mathbf{n}_i, \eta_i)) - \theta(\overrightarrow{\mathbf{x}_i\mathbf{x}_{i+1}}) \\ \phi(\mathrm{S}(\overrightarrow{\mathbf{x}_i\mathbf{x}_{i-1}}, \mathbf{n}_i, \eta_i)) - \phi(\overrightarrow{\mathbf{x}_i\mathbf{x}_{i+1}}) \end{pmatrix} \tag{12}$$

where $\theta(\boldsymbol{\omega}) = \cos^{-1}(\boldsymbol{\omega}_z)$ and $\phi(\boldsymbol{\omega}) = \mathrm{atan2}(\boldsymbol{\omega}_y, \boldsymbol{\omega}_x)$ determine associated spherical coordinates. Figure 8 showcases the significantly improved convergence due to these constraints.

Note that the scattering operation $\mathrm{S}(\cdot)$ can fail in configurations with total internal reflection. In such cases we are still able to evaluate the constraint for the opposite light direction by taking the difference between $\mathrm{S}(\boldsymbol{\omega}_\mathrm{o}, \mathbf{n}, \eta^{-1})$ and $\boldsymbol{\omega}_\mathrm{i}$.

4.5 Two-stage manifold walks

Comparing the nature of specular paths in a simple cardioid caustic (Fig. 4) to a reflection from normal-mapped geometry (Fig. 9, left column), we observe that the more complex caustic is a superposition of many different solutions with a fairly localized effect. The space $\mathcal{U}$ is largely empty, containing only a few small convergence basins that are clustered together (Fig. 10a). The probability of finding a valid solution is therefore low. Furthermore, estimating the reciprocal probability depends on our ability to *find the same solution once more*, which is even less likely. Unbiased SMS estimates therefore tend be slow and noisy, while biased SMS loses a considerable amount of energy.

To address these problems, we will exploit prior knowledge to devise an improved strategy for generating initial guesses. To motivate our approach, consider a caustic generated by a smooth planar surface (e.g., a flat version of the metal surface shown in Fig. 9). Points on this caustic all correspond to a single solution with a large convergence basin in $\mathcal{U}$. If we begin to make the surface more complex, e.g., by perturbing positions or shading normals, this single solution splits into multiple nearby solutions with disproportionately smaller convergence basins. Stronger geometric perturbations tend to produce more solutions that are spread further apart. The aforementioned issues with SMS could be addressed if there was a way to predict the locations of these solutions and construct starting points in their proximity in a more targeted manner. We propose a simple technique to create such a targeted initial guess for the special case of normal-mapped surfaces.[7]

[6] The subtraction of azimuth angles requires special handling due to their periodicity: the result must be mapped onto $[-\pi, \pi]$ using a floating point modulo operation.

[7] The high-level idea is more general and can likely be extended to other approaches for introducing surface detail, such as displacement maps.

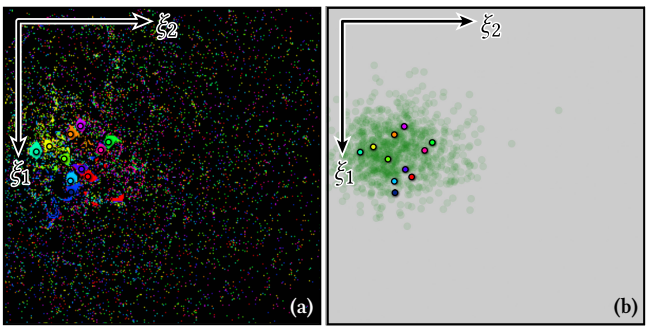

Fig. 10. **(a)**: Convergence basins of the Reflective Plane scene in Fig. 9a, for one shading point on the ground inside the caustic region. **(b)**: Sampling density (green) of intermediate points after the first manifold walk in our two-pass approach, which focuses on the cluster of solutions.

Our *two-stage* sampling approach performs two manifold walks: the first stage ignores the specified normal map and finds specular paths on the original smooth surface. Such a smooth manifold walk will converge to a point that roughly lies at the center of the cluster of solutions in Fig. 10a. However, instead of an ordinary manifold walk, our first stage relies on an *offset manifold walk* [Jakob and Marschner 2012], whose manifold constraint tilts the normal of the underlying surface. The normal perturbation is randomly chosen from the distribution of normals that are present in the normal map. Cheaper approximations are also usable—we use a Gaussian approximation of the entire normal map obtained from the lowest MIP level of a LEAN map [Olano and Baker 2010]. Importantly, this offset normal is chosen before each manifold walk and does not change during the iteration which converges rapidly and with high probability (i.e. the convergence basin is large). This initial randomized manifold walk brings us into the proximity of the various solutions (green density in Fig. 10b); a second manifold walk, on the bumpy surface, started from there takes us all the way to a solution.

Estimation of the reciprocal probability of this adapted sampling strategy is surprisingly easy: the only requirement is that two-stage sampling is consistently used in both Lines 1-2 and Lines 4-9 of Alg. 2. It would generally appear that the two-stage sampling should be more costly, but in our experiments (e.g. Fig. 9) we observed roughly equal performance. The reason for this is that the reciprocal probability estimator requires fewer iterations to rediscover a solution. At the same time, variance is reduced noticeably.

4.6 Glints

Our method generalizes straightforwardly to the problem of rendering *glints*, which are miniscule subpixel reflections of a light source with narrow angular support (e.g. the sun) on high-frequency specular microgeometry. Rendering glints using standard Monte Carlo techniques tends to be prohibitively expensive, since millions of samples per pixel may be needed to obtain an acceptable result.

Interestingly, the problem of finding glints is almost identical to the caustic case, the main exception being that the search is now constrained to small surface regions observed by individual pixels and their reconstruction filters. Previous general glint rendering techniques [Yan et al. 2014, 2016] evaluate an effective integrated BRDF over the entire pixel footprint and produce very high quality estimates at the cost of burdensome 4D spatio-directional hierarchies that are necessary to organize and query the distribution of normals. Our goal is to implement an unbiased estimate of such a query using manifold walks.

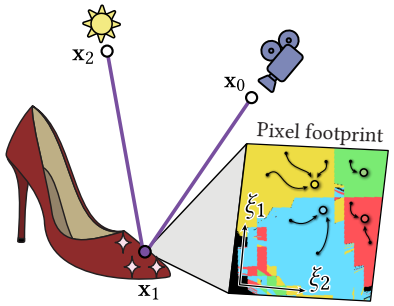

Fig. 11. Specular manifold sampling can also find connections within a pixel footprint to render glints that arise due to specular microstructure.

Our method generates random starting points within the pixel and then runs unbiased or biased SMS to find solutions (Fig. 11). In contrast to prior work, which introduced a small amount of *intrinsic roughness* to relax the problem, we solve the unmodified problem with a discrete set of solutions to find a path $\bar{\mathbf{x}} = \mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2$ connecting the camera to a sampled emitter position via a single[8] specular reflection. After tracing the initial camera ray, we approximate the projected pixel footprint with a small parallelogram based on ray differentials [Igehy 1999]. We perform local manifold walks in UV space to refine the initial guess. This leads to much higher performance compared to the caustic case as no costly ray tracing operation is needed to re-project onto the surface. We terminate paths that step outside the parallelogram since they are unlikely to find usable solutions. Like prior work, we assume that the endpoints $\mathbf{x}_0$ and $\mathbf{x}_2$ are distant, in which case changes in the half-vector across the parallelogram are minimal. It can thus be approximated by a constant, which simplifies the search for solutions: in particular, the specular manifold constraint can be simplified to a function that attempts to equate this fixed half-vector and the local shading normal, both expressed in 2D slope space. Like previous work, we use high-resolution normal maps to encode the subpixel surface details in our scenes. In principle, the method could be extended to other types of normal variation or actual geometric displacement.

To robustly apply our glint rendering technique in scenes with complex lighting (e.g. high-frequency environment maps), we further combine our SMS strategy with standard BSDF sampling in a multiple importance sampling (MIS) framework. We found it very effective to use approximate MIS weights based on the directional distributions from the light source and the effective BSDF of the pixel footprint provided by a LEAN map. As shown in Fig. 12 this approach successfully separates the regions of the integrands where one sampling strategy is preferable over the other. While precomputing a LEAN map adds some storage overhead, it is orders of magnitude lower than Yan et al.'s spatio-directional trees. We also

[8] Glints involving multiple specular reflections on displaced geometry can in principle be found using our method and could be an interesting topic for future work.

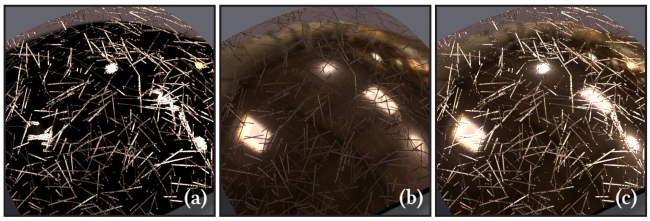

Fig. 12. Curved metal surface with small scratches lit by a high-frequency environment map containing multiple light sources. Weighting using multiple importance sampling (MIS) effectively uses each technique where it performs best: **(a)**: Our method (unbiased SMS) captures the direct illumination from the bright light sources inside the scratches. **(b)**: Standard BSDF sampling focuses on reflections from flat regions. **(c)**: Combined result.

note that our method always renders the non-smoothed input normal map without geometric approximations—the LEAN map is only used as a proxy to compute an effective BSDF sampling density over a pixel footprint, enabling unbiased combination of our sampling strategy with other standard techniques via MIS.

4.7 Integration into rendering algorithms

Specular manifold sampling is a general building block in the design of rendering algorithms, and there is considerable flexibility in its usage. We experimented with different ways of incorporating it into a unidirectional path tracer and ultimately opted for uniform seed-point generation on specially marked "caustic caster" shapes. For two-bounce caustics, shown in Section 5, the seed point on the second interface is found by tracing a ray through the first vertex and performing either reflection or refraction depending on the type of surface. SMS is very effective for sampling high-frequency caustic paths from (near) specular surfaces and often greatly outperforms standard emitter or BSDF sampling strategies. However, in the case of significant surface roughness or smooth illumination, standard strategies remain superior. Ideally, all of these strategies would be combined via MIS to improve their robustness, as we do in Section 4.6 for the special case of glints. However, determination of suitable probabilities in the more general caustic setting remains an open problem. SMS iterations are also relatively expensive, and we use them even at shading points where no possible connections exist. Conservative criteria that specify when SMS should be used are another interesting avenue for future work.

5 RESULTS

We now present several results rendered with our method as well as comparisons to relevant prior work. We based our implementation on the Mitsuba 2 renderer [Nimier-David et al. 2019]. All time measurements in this article were recorded on a compute node with 2 Xeon 6132 processors, each with 14 2.6 GHz cores. Many of the rendered results contain small insets with the manifold walk success rate (SR) for the corresponding image. These numbers may seem low at first, but are averaged over the full image—including pixels that do not contain any caustics for SMS to find.

Many of the results below are additionally shown in animated form as part of the supplemental video:

- Comparison of SMS against prior work on three scenes (Fig. 14): Swimming Pool, Reflective Plane, and Refractive Sphere.
- Temporal coherency and the effect of the trial set size M in biased SMS (Fig. 15).
- Two-bounce refractive caustics in the Double-Refractive Slab scene (Fig. 18).
- Temporal coherency of SMS on glints in the Shoes and Kettle scenes (Fig. 19).

5.1 Shop window scene

Figure 1 showcases our method in a example involving glints and caustics on complex geometry. It features several different caustics rendered using unbiased SMS: reflective caustics caused by a spotlight shining on the gold pedestals, and a two-bounce refractive caustic from a point light hidden inside the middle pedestal. The shoes have a glittery appearance that is rendered using biased SMS for glints. We disable caustics from higher order scattering (e.g. two or more bounces between the pedestals). Many of the caustics in this scene are SDS paths that a regular path tracer cannot find. We can still compare to a path tracer by adding some surface roughness (Beckmann NDF with $\alpha = 0.005$) to all specular shapes. Figure 13 shows such a comparison, where we only visualize the contributions sampled by SMS. Given equal render time, our method outperforms regular path tracing by a significant margin.

The scene also illustrates a current limitation of our method: even though SMS can effectively sample light connections through specular interfaces, there is still significant variance in the scene coming from other challenging paths, mainly the indirect lighting caused by the caustics; see Fig. 13b and the remaining noise in Fig. 1. Our method cannot sample these paths explicitly; the issue could be addressed by path guiding methods.

Fig. 13. Modified version of the Shop Window scene from Fig. 1 where a small amount of surface roughness is added to the previously purely specular surfaces, enabling a standard path tracer **(a)** to find the same light transport paths as our proposed method **(b)**. Both renderings are computed in equal time (20 minutes), showing only contributions due to caustics.

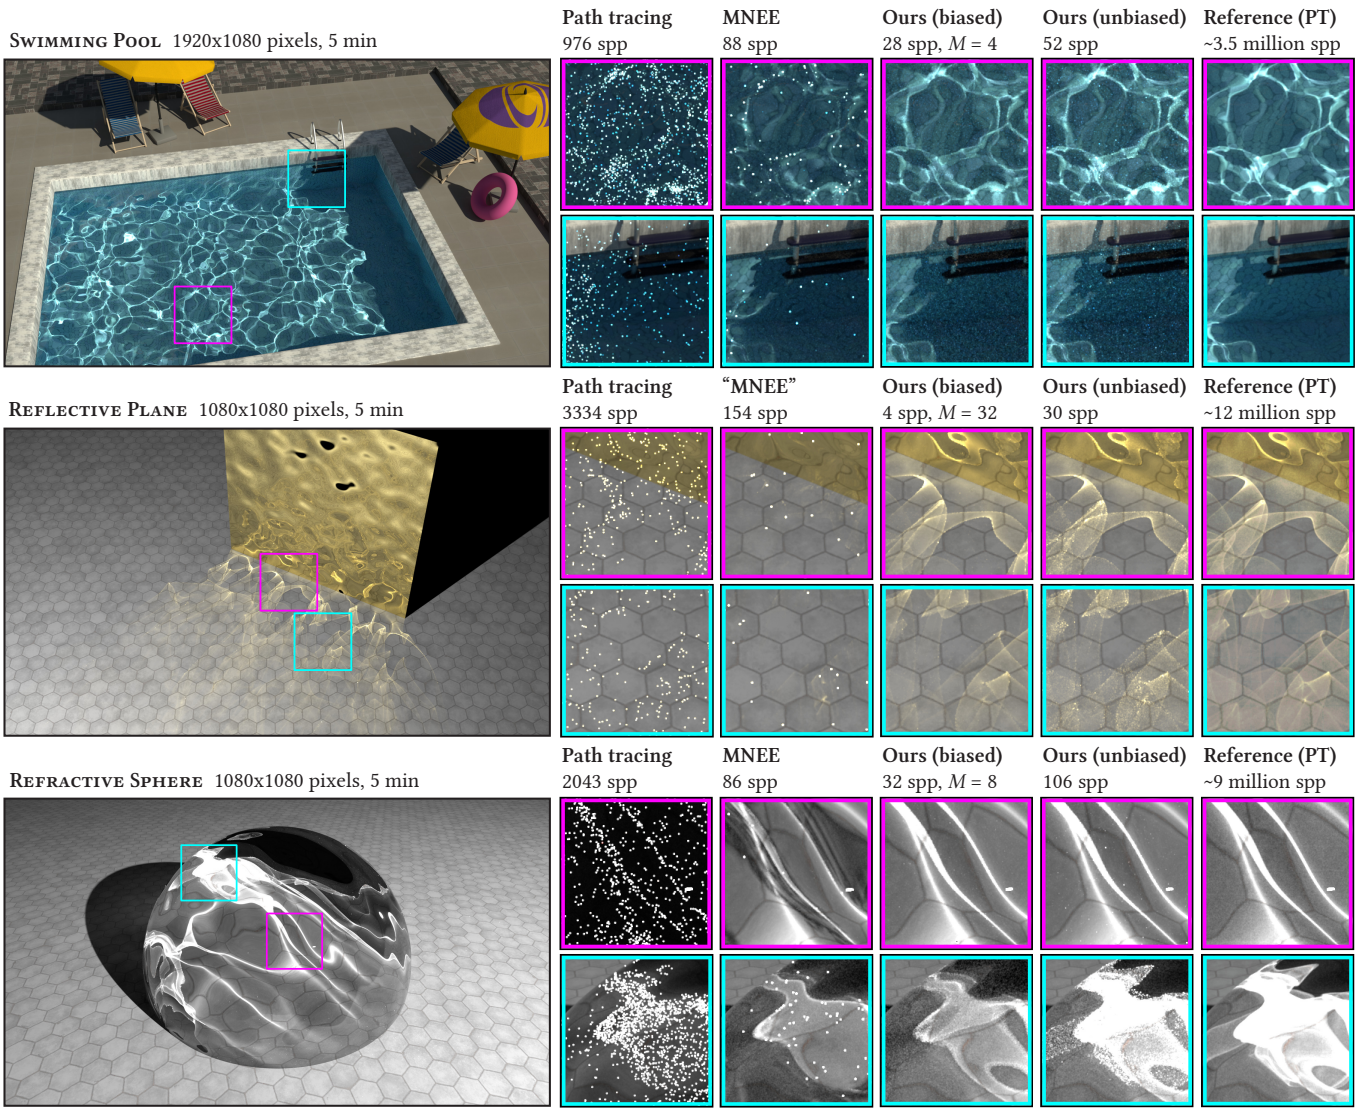

Fig. 14. Equal-time comparison (5 minutes) between path tracing, prior-work MNEE [Hanika et al. 2015a], and both unbiased and biased versions of our proposed SMS method. Previous work can only produce (at most) one connection to the light source via the specular interface, whereas our techniques can sample the full range of light paths either in an unbiased or biased way. The latter removes some of the high-frequency noise introduced by unbiased probability estimation and trades it for energy loss with the trial-set size parameter M. We report samples per pixel (spp) computed by each method, as well as the chosen M in the biased case.

5.2 Specular manifold sampling

In Fig. 14 we compare the effectiveness of SMS to brute-force path tracing and the previous state-of-the-art method MNEE of Hanika et al. [2015a]. We examine three scenes with challenging caustics due to normal-mapped surfaces. The Swimming Pool and Refractive caustic scenes are a famous examples where both uni- and bidirectional path tracing techniques fail to discover the prominent SDS paths that generate intricate patterns on the ground plane. MNEE improves on this but misses all but one light connection. For the remaining paths it has to either fall back to the brute-force strategy that has significant variance still or suffers from severe energy loss if those light connections are omitted. Our method finds all paths and significantly outperforms the others in equal time.

The Reflective Plane scene is an example where MNEE was previously not applicable. For a clearer comparison, we added a variation (called "MNEE" in quotes) that constructs a deterministic seed path by tracing a ray from the shading point towards the center of the object's bounding box. Since the caustic is the superposition of many individual solutions, this is clearly not sufficient and "MNEE" ends up finding only a very small part of the caustic. Biased versions of SMS can optionally be applied and reduce high-frequency noise caused by the unbiased probability estimate. This comes at a higher

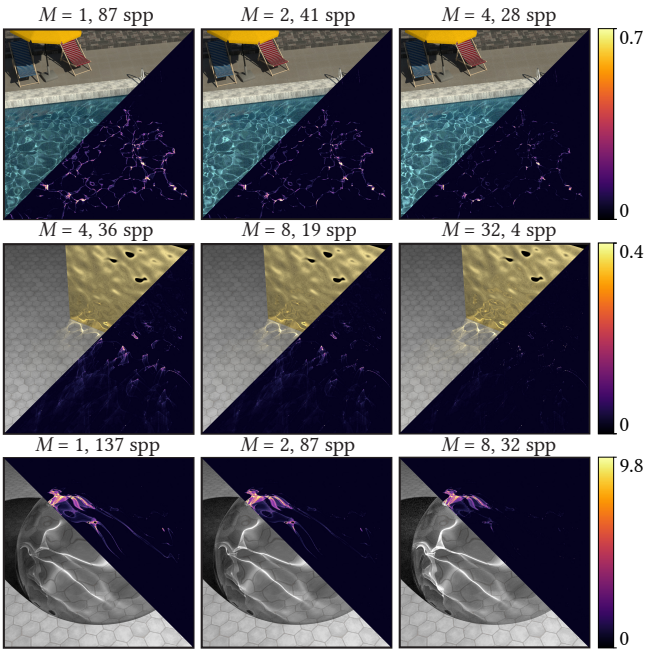

Fig. 15. Comparing our biased method with varying trial set sizes M, at equal render time of 5 minutes. As M increases, each individual radiance estimate becomes more expensive but less energy of the caustics is lost. Overlays show pixel-wise squared error compared to unbiased SMS.

cost per individual sample and some energy loss, the extent of which can be controlled by the trial set size M. We want to highlight that the biased method is free of artifacts and generates temporally coherent results despite its limited exploration of a random subset of the path space. Animated results of this comparison for different values of M are shown in the supplemental video.

5.3 Biased SMS

Our biased SMS approach involves an interesting tradeoff between sample variance and energy loss. The trial set size parameter M plays an important role here: it directly relates to how much effort the sampler spends per radiance estimate from caustics, and indirectly controls how much of the caustic is found. We explore this effect in Fig. 15 on the same scenes as used above. Choosing the optimal M is currently a user choice and could be investigated more in the future.

5.4 Geometric displacement

Although many of the shown results involve caustics from specular surfaces with normal maps, our method is more general, and we also found it to be effective on surfaces with true geometric detail (e.g. from a displacement map), which we illustrate in Fig. 16. When not using our two-stage sampling method, which is currently limited to normal maps, manifold sampling performs equally well on both types of surfaces as confirmed by the similar success rates. This seems counter-intuitive at first, since smoother geometry should also result in a specular manifold that is easier to navigate. However,

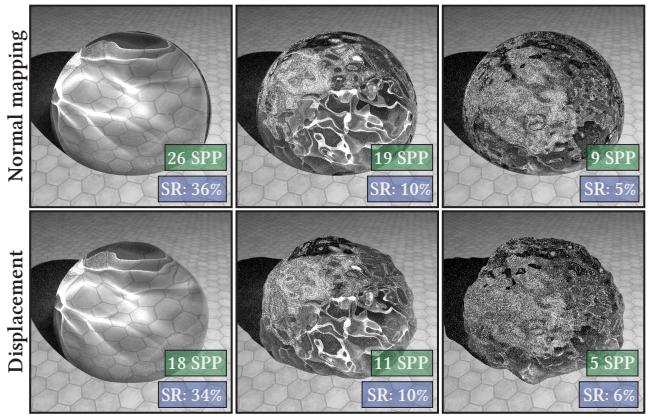

Fig. 16. **Top**: Sequence of refractive spheres with increasingly complex normal maps. **Bottom**: The same setup, but this time with actual displaced geometry. Insets show computed number of samples per pixel (SPP) and the manifold walk success rate (SR) measured over all pixels in the image. The manifold walks perform about equally on both geometry types.

the disagreement between actual geometry and the "fake" surface variation from normal maps can also limit the Newton solver's efficiency. The ray-tracing-heavy nature of the manifold walks and the underlying re-projection steps lead to slightly reduced performance in scenes with dense geometric tessellation.

5.5 Surface roughness

Like prior work on specular manifolds, our method generalizes to near-specular or rough surfaces by performing offset manifold walks previously described in Section 4.5. Here, we sample an *offset normal* from the material's microfacet distribution before the SMS step. The manifold walk then searches for specular connections involving that offset normal instead of the true shading normal. Only a finite number of solutions exists for this specific offset normal, and the probability estimate of therefore remains unchanged. Averaging over paths sampled in this manner converges to the correct solution [Hanika et al. 2015a].

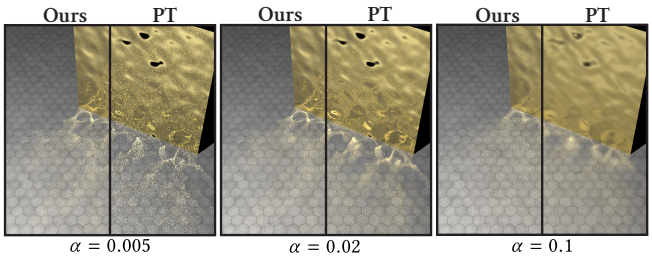

Fig. 17. Rough reflective plane with a rough metallic surface with Beckmann normal distribution. As the roughness increases from left to right, the path tracer (PT) becomes more capable of performing the connection to the light source on specular plane directly, whereas our proposed sampling strategy loses its efficiency. Both methods use equal render time of 5 minutes.

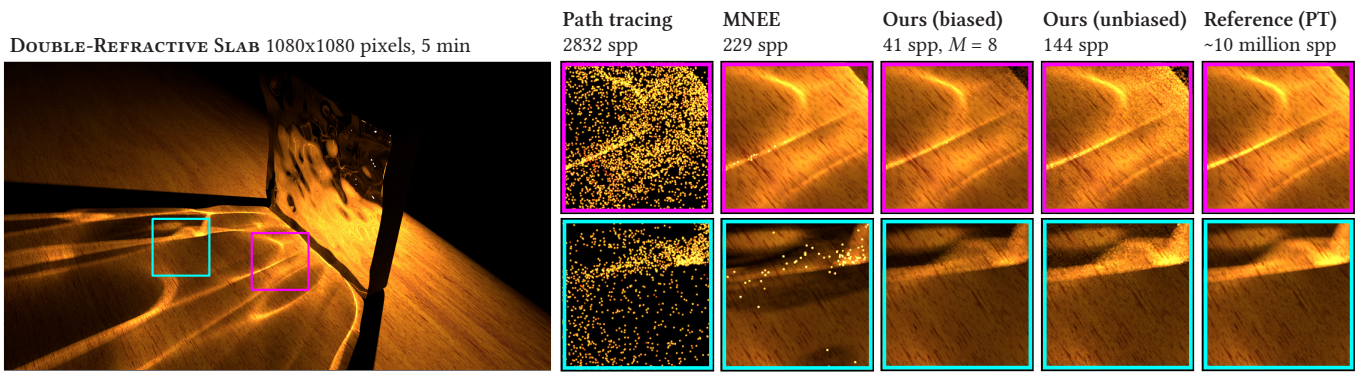

Fig. 18. Equal-time comparison (5 minutes) of our method on a challenging scene where SMS samples a light connection involving two consecutive specular refractions. The two-sided solid piece of glass is modeled with geometric displacement.

Figure 17 shows a sequence of increasingly rough reflective surfaces. Note how the caustic is at first sharp and full of high-frequency details, and becomes progressively blurry from left to right. This blur also enables a unidirectional path tracer to find valid light connections more often: in the limit case, every point on the surface is contributing to the shading point. As our method handles roughness by integrating over many perfectly specular light paths with randomized offset normals, the opposite is true for our method and its variance increases. We only recommend using our technique for specular or near-specular surfaces and switching over to conventional path sampling techniques in other cases. In the future, it would be interesting to incorporate a form of multiple importance sampling to robustly handle both extremes in addition to intermediate cases. Note that the same argument applies also to scenes with low-frequency lighting, e.g. largely constant environment maps.

5.6 Multiple specular interactions

Like MNEE, the principle behind our method generalizes to longer chains with multiple specular interactions. In Fig. 18, we show an intricate caustic pattern caused by double refraction through a solid displaced piece of glass. We again compare our unbiased and biased SMS variants to a standard path tracer and MNEE. Multiple interactions increase the dimension of the space of initial configurations that SMS must generate: we could generate initial rays to start the manifold walks in the same way as before, but at each interaction we additionally decide between reflection or refraction (if applicable), and whether or not to terminate the chain and attempt to connect to a light position. To keep variance manageable we found it best to limit SMS to a single family of light paths in this setting (e.g. paths of fixed length with only refractive events). In cases where this increased variance is not acceptable, the biased technique can still be applied—but for the same reasons there will be a potentially significant energy loss. As shown in the Double-Refractive Slab scene, SMS can still produce good results, but better strategies for sampling initial configurations will be required to turn SMS into a fully general sampling strategy that can efficiently find *all possible* chains of specular interactions.

5.7 Glints

Figure 19 examines the performance of our glint rendering technique on two scenes with complex microstructure specified using normal maps. We compare our SMS to the state-of-the-art method of Yan et al. [2016]. Both methods make use of MIS in this comparison. The Shoes scene features a glittery pair of shoes with procedural normal displacement by a Gaussian height field and is lit by a sky with an almost purely directional sun. The Kettle scene involves brushed metal with very strong anisotropy. It is illuminated by the Grace Cathedral environment map, which includes several bright and narrow light sources. Stochastic sampling of glint solutions is beneficial in these cases, since several complementary sources of variance can be reduced at the same time.

We found the biased SMS variant to generally be more practical for glint rendering compared to its unbiased counterpart.[9] The potentially unbounded number of iterations in the recursive unbiased probability estimator, combined with extremely high-frequency normal map detail, occasionally produces acute outliers in the pixel estimate that lead to poor convergence, as seen in the plots.

Note that all methods in Fig. 19 converge to slightly different results, but they all find the same individual glints and have very similar appearance overall. Our method uses 100–300× less memory compared to previous work (e.g. 110 MiB vs. 11 GiB). At the same time, it converges in an equal or shorter amount of time and still generates temporally coherent animations. Please see the supplemental video for an animated demonstration.

6 CONCLUSION

We introduced a simple and powerful specular path sampling technique that combines deterministic root finding with stochastic sampling in a pure Monte Carlo setting. The basic method can be used in a variety of different ways, and we demonstrated example applications in the context of efficient path tracing of glints and caustics. Our approach is not restricted to unidirectional path tracing, and we contemplate its utility in bidirectional and even MCMC methods, where manifold walks were originally proposed.

[9] The bias here only involves the probability estimate. In practice, even the "unbiased" version will not match a brute-force result perfectly due to the far-field approximation.

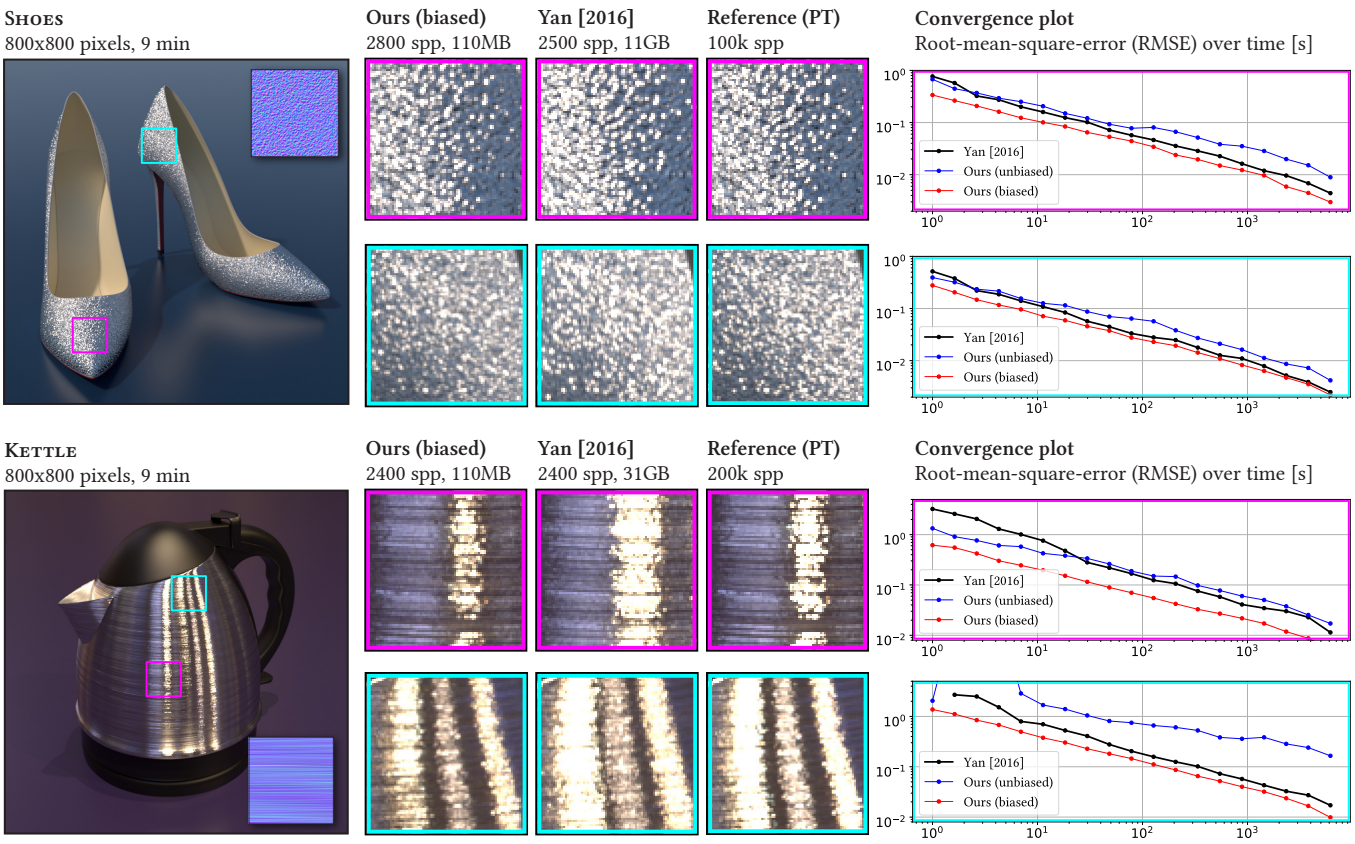

Fig. 19. Equal-time comparison (9 minutes) of our glint rendering method to prior work specific to this problem [Yan et al. 2016]. Shoes scene: highly directional illumination from the sun, Kettle scene: grace cathedral environment map where integration over multiple sources of variance (i.e. all the lights) is critical. Our method yields comparable results in the first scenario and is superior in the second. At the same time, our method requires 100–300× less memory. The insets and corresponding convergence plots focus on different parts of the glinty appearance.

Building on a simple unbiased algorithm, we presented several complementary extensions and improvements. For example, better strategies for sampling seeds paths can further improve convergence, and such heuristics are easy to integrate into our method without introducing bias. Improved manifold constraints expand the size of the convergence basins in primary sample space. A further change yields an intentionally biased estimator with desirable properties for production usage. Far-field approximations in the context of specular glints lead to a particularly simple iterative algorithm, whose steps no longer require the use of ray tracing operations. In the future, we would like to explore further acceleration of this variant leveraging vectorized execution.

Determining *when* to use our method is another important aspect for future investigation. Attempting many connections that are ultimately unsuccessful can consume a large amount of computation. While glints would benefit from simple culling heuristics, e.g. based on cones bounding the normal variation inside the pixel, the general case of caustics from arbitrary specular geometry is significantly more challenging. Combining our techniques with others via multiple importance sampling in this general setting is another pertinent problem.

Our article focuses mainly on the generation of subpaths with a single specular vertex. While our method in principle also generalizes to more complex path classes with multiple specular reflection, performance using our current strategy for choosing starting points remains suboptimal and could be an interesting topic for future work. We wish to pursue these and related improvements, and envision a unified path sampling strategy that elevates stochastic manifold walks to a standard building block in the design of Monte Carlo rendering methods.

ACKNOWLEDGMENTS

We thank Guillaume Loubet for many insightful discussions about glints and caustics, Olesya Jakob for designing the shop window scene in Fig. 1, Ling-Qi Yan and Miloš Hašan for sharing their implementation of prior work [Yan et al. 2016], and Héloïse de Dinechin for fixing a bug in the implementation of the specular constraints.

Our test scenes use textures from CC0 Textures and cgbookcase, and are lit by environment maps courtesy of HDRI Haven and Paul Debevec. The kettle model in Fig. 19 has been created by Blend Swap user *PrinterKiller*.

This work was funded by a grant from Autodesk.

REFERENCES

Thomas E. Booth. 2007. Unbiased Monte Carlo Estimation of the Reciprocal of an Integral. *Nuclear Science and Engineering* 156, 3 (2007), 403–407. https://doi.org/10.13182/NSE07-A2707 arXiv:https://doi.org/10.13182/NSE07-A2707

Matt Jen-Yuan Chiang and Brent Burley. 2018. Plausible Iris Caustics and Limbal Arc Rendering. In *ACM SIGGRAPH 2018 Talks* (Vancouver, British Columbia, Canada) *(SIGGRAPH '18)*. ACM, New York, NY, USA, Article 15, 2 pages. https://doi.org/10.1145/3214745.3214751

Luis E. Gamboa, Jean-Philippe Guertin, and Derek Nowrouzezahrai. 2018. Scalable Appearance Filtering for Complex Lighting Effects. *ACM Trans. Graph.* 37, 6, Article Article 277 (Dec. 2018), 13 pages. https://doi.org/10.1145/3272127.3275058

Iliyan Georgiev, Jaroslav Křivánek, Tomáš Davidovič, and Philipp Slusallek. 2012. Light transport simulation with vertex connection and merging. *ACM Trans. Graph.* 31, 6, Article 192 (Nov. 2012), 10 pages. https://doi.org/10.1145/2366145.2366211

Toshiya Hachisuka and Henrik Wann Jensen. 2009. Stochastic Progressive Photon Mapping. *ACM Trans. Graph.* 28, 5 (Dec. 2009), 1–8. https://doi.org/10.1145/1618452.1618487

Toshiya Hachisuka, Jacopo Pantaleoni, and Henrik Wann Jensen. 2012. A path space extension for robust light transport simulation. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 191.

Johannes Hanika, Marc Droske, and Luca Fascione. 2015a. Manifold Next Event Estimation. *Computer Graphics Forum (Proceedings of Eurographics Symposium on Rendering)* 34, 4 (June 2015), 87–97.

Johannes Hanika, Anton Kaplanyan, and Carsten Dachsbacher. 2015b. Improved Half Vector Space Light Transport. *Comput. Graph. Forum* 34, 4 (July 2015), 65–74. http://dl.acm.org/citation.cfm?id=2858834.2858842

John Hubbard, Dierk Schleicher, and Scott Sutherland. 2001. How to Find All Roots of Complex Polynomials by Newton's Method. *Inventiones mathematicae* 146 (10 2001), 1–33. https://doi.org/10.1007/s002220100149

Homan Igehy. 1999. Tracing Ray Differentials. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 179–186. https://doi.org/10.1145/311535.311555

Wenzel Jakob. 2013. *Light transport on path-space manifolds*. Ph.D. Dissertation. Cornell University.

Wenzel Jakob, Miloš Hašan, Ling-Qi Yan, Ravi Ramamoorthi, and Steve Marschner. 2014. Discrete Stochastic Microfacet Models. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 33, 4 (July 2014), 115:1–115:10. https://doi.org/10.1145/2601097.2601186

Wenzel Jakob and Steve Marschner. 2012. Manifold Exploration: A Markov Chain Monte Carlo Technique for Rendering Scenes with Difficult Specular Transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 31, 4 (July 2012), 58:1–58:13. https://doi.org/10.1145/2185520.2185554

Henrik Wann Jensen. 1996. Global illumination using photon maps. In *Rendering Techniques' 96*. Springer, 21–30.

James T. Kajiya. 1986. The Rendering Equation. In *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '86)*. Association for Computing Machinery, New York, NY, USA, 143–150. https://doi.org/10.1145/15922.15902

Anton Kaplanyan and Carsten Dachsbacher. 2013. Path Space Regularization for Holistic and Robust Light Transport. *Computer Graphics Forum* 32 (05 2013). https://doi.org/10.1111/cgf.12026

Anton S. Kaplanyan, Johannes Hanika, and Carsten Dachsbacher. 2014. The Natural-constraint Representation of the Path Space for Efficient Light Transport Simulation. *ACM Trans. Graph.* 33, 4, Article 102 (July 2014), 13 pages. https://doi.org/10.1145/2601097.2601108

David Koerner, Jan Novák, Peter Kutz, Ralf Habel, and Wojciech Jarosz. 2016. Subdivision Next-Event Estimation for Path-Traced Subsurface Scattering. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations* (Dublin, Ireland) *(EGSR '16)*. Eurographics Association, Goslar, DEU, 91–96.

Alexandr Kuznetsov, Miloš Hašan, Zexiang Xu, Ling-Qi Yan, Bruce Walter, Nima Khademi Kalantari, Steve Marschner, and Ravi Ramamoorthi. 2019. Learning Generative Models for Rendering Specular Microgeometry. *ACM Trans. Graph.* 38, 6, Article Article 225 (Nov. 2019), 14 pages. https://doi.org/10.1145/3355089.3356525

Eric Lafortune and Yves Willems. 1993. Bi-Directional Path Tracing. *Proceedings of Third International Conference on Computational Graphics and Visualization Techniques (Compugraphics'* 93 (01 1993).

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 Article 95 (July 2013), 12 pages. https://doi.org/10.1145/2461912.2461943

Don Mitchell and Pat Hanrahan. 1992. Illumination from Curved Reflectors. *SIGGRAPH Comput. Graph.* 26, 2 (July 1992), 283–291. https://doi.org/10.1145/142920.134082

Thomas Müller, Markus Gross, and Jan Novák. 2017. Practical Path Guiding for Efficient Light-Transport Simulation. *Comput. Graph. Forum* 36, 4 (July 2017), 91–100. https://doi.org/10.1111/cgf.13227

Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. 2019. Mitsuba 2: A Retargetable Forward and Inverse Renderer. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 38, 6 (Dec. 2019). https://doi.org/10.1145/3355089.3356498

Marc Olano and Dan Baker. 2010. LEAN Mapping. In *Proceedings of the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (Washington, D.C.) *(I3D '10)*. Association for Computing Machinery, New York, NY, USA, 181–188. https://doi.org/10.1145/1730804.1730834

Hao Qin, Xin Sun, Qiming Hou, Baining Guo, and Kun Zhou. 2015. Unbiased Photon Gathering for Light Transport Simulation. *ACM Trans. Graph.* 34, 6, Article Article 208 (Oct. 2015), 14 pages. https://doi.org/10.1145/2816795.2818119

Florian Simon, Johannes Hanika, Alisa Jung, and Carsten Dachsbacher. 2018. Selective guided sampling with complete light transport paths. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 37, 6 (Dec. 2018). https://doi.org/10.1145/3272127.3275030

Sebastien Speierer, Christophe Hery, Ryusuke Villemin, and Wenzel Jakob. 2018. Caustic Connection Strategies for Bidirectional Path Tracing. *Pixar Technical Memo #18-01*.

Eric Veach. 1998. *Robust Monte Carlo Methods for Light Transport Simulation*. Ph.D. Dissertation. Stanford, CA, USA. Advisor(s) Guibas, Leonidas J. AAI9837162.

Eric Veach and Leonidas Guibas. 1995a. Bidirectional estimators for light transport. In *Photorealistic Rendering Techniques*. Springer, 145–167.

Eric Veach and Leonidas J Guibas. 1995b. Optimally combining sampling techniques for Monte Carlo rendering. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. ACM, 419–428.

Jiří Vorba, Ondřej Karlík, Martin Šik, Tobias Ritschel, and Jaroslav Křivánek. 2014. On-Line Learning of Parametric Mixture Models for Light Transport Simulation. *ACM Trans. Graph.* 33, 4, Article Article 101 (July 2014), 11 pages. https://doi.org/10.1145/2601097.2601203

Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet Models for Refraction Through Rough Surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques* (Grenoble, France) *(EGSR'07)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 195–206. https://doi.org/10.2312/EGWR/EGSR07/195-206

Bruce Walter, Shuang Zhao, Nicolas Holzschuch, and Kavita Bala. 2009. Single Scattering in Refractive Media with Triangle Mesh Boundaries. *ACM Transactions on Graphics* 28, 3 (Aug. 2009), 92:1–8. https://doi.org/10.1145/1531326.1531398

Beibei Wang, Miloš Hašan, Nicolas Holzschuch, and Ling-Qi Yan. 2019. *Example-Based Microstructure Rendering with Constant Storage*. Technical Report 2019-08. UC Santa Barbara. https://cs.ucsb.edu/research/tech-reports/2019-08

Pascal Weber, Johannes Hanika, and Carsten Dachsbacher. 2017. Multiple Vertex Next Event Estimation for Lighting in Dense, Forward-Scattering Media. *Comput. Graph. Forum* 36, 2 (May 2017), 21–30. https://doi.org/10.1111/cgf.13103

Ling-Qi Yan, Miloš Hašan, Steve Marschner, and Ravi Ramamoorthi. 2016. Position-Normal Distributions for Efficient Rendering of Specular Microstructure. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2016)* 35, 4 (2016).

Ling-Qi Yan, Miloš Hašan, Wenzel Jakob, Jason Lawrence, Steve Marschner, and Ravi Ramamoorthi. 2014. Rendering Glints on High-Resolution Normal-Mapped Specular Surfaces. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2014)* 33, 4 (2014).

Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering Specular Microgeometry with Wave Optics. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2018)* 37, 4 (2018).