Continuous Multiple Importance Sampling

REX WEST, The University of Tokyo, Japan
ILIYAN GEORGIEV, Autodesk, United Kingdom
ADRIEN GRUSON, McGill University, Canada
TOSHIYA HACHISUKA, The University of Tokyo, Japan

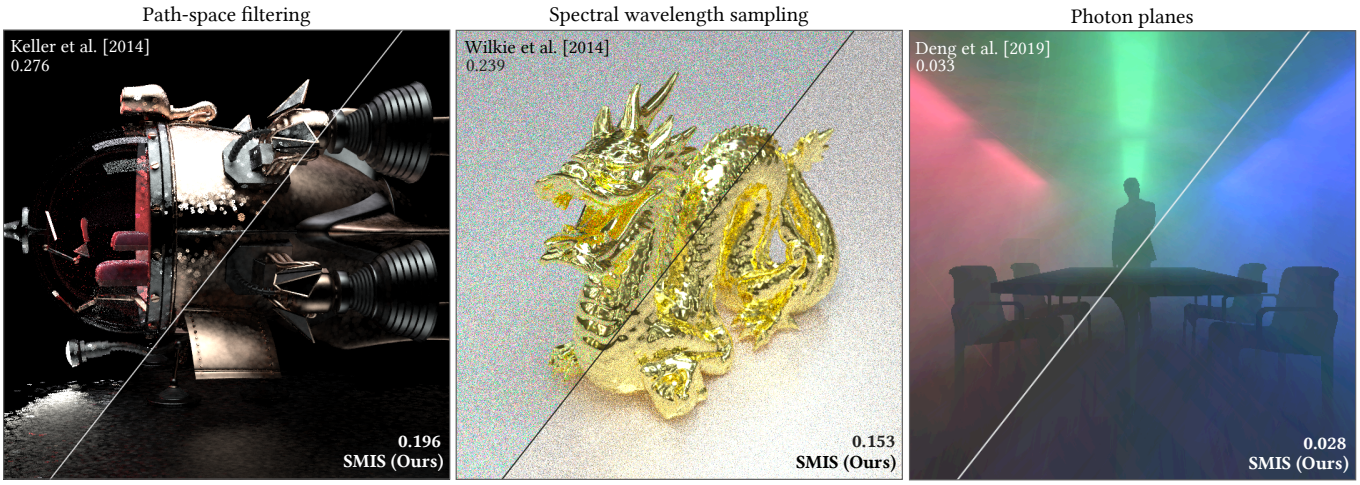

Fig. 1. Three applications of our framework to light transport simulation. We reformulate each application as a problem of combining a continuum of sampling techniques and leverage our continuous MIS (CMIS) formulation to derive an efficient weighting scheme. Based on this scheme, our practical stochastic MIS (SMIS) estimator outperforms existing state-of-the-art methods. For each image we report error in SMAPE units (see Section 4.4).

Multiple importance sampling (MIS) is a provably good way to combine a finite set of sampling techniques to reduce variance in Monte Carlo integral estimation. However, there exist integration problems for which a continuum of sampling techniques is available. To handle such cases we establish a continuous MIS (CMIS) formulation as a generalization of MIS to uncountably infinite sets of techniques. Our formulation is equipped with a base estimator that is coupled with a provably optimal balance heuristic and a practical stochastic MIS (SMIS) estimator that makes CMIS accessible to a broad range of problems. To illustrate the effectiveness and utility of our framework, we apply it to three different light transport applications, showing improved performance over the prior state-of-the-art techniques.

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

Additional Key Words and Phrases: multiple importance sampling, light transport, spectral rendering, path reuse, volume rendering

Authors' addresses: Rex West, The University of Tokyo, Japan; Iliyan Georgiev, Autodesk, United Kingdom; Adrien Gruson, McGill University, Canada; Toshiya Hachisuka, The University of Tokyo, Japan.

ACM Reference Format:
Rex West, Iliyan Georgiev, Adrien Gruson, and Toshiya Hachisuka. 2020. Continuous Multiple Importance Sampling. *ACM Trans. Graph.* 39, 4, Article 136 (July 2020), 12 pages. https://doi.org/10.1145/3386569.3392436

1 INTRODUCTION

Multiple importance sampling [Veach and Guibas 1995] (MIS) provides a framework for combining a set of sampling techniques in Monte Carlo integration. This combination is done by weighting the contribution of each sample produced by each sampling technique according to some heuristic.

MIS can be directly applied to problems where the set of sampling techniques is countable. However, there are certain problems where *an uncountably infinite number* (, a *continuum*) of techniques arises naturally. A generalization of MIS is needed for these problems as the classical formulation does not consider such cases.

We formally establish this continuous generalization of MIS, which we call *continuous MIS* (CMIS). Based on our formulation, we devise a CMIS estimator that combines a continuum of sampling techniques using a provably optimal balance heuristic. Since this estimator is not always practical, we propose an approximation to it—our *stochastic MIS* (SMIS) estimator—which is unbiased and extends classical MIS to stochastic technique selection.

To demonstrate the utility of our framework, we apply it to three applications in light transport simulation shown in Fig. 1. In path space filtering [Keller et al. 2014], CMIS allows us to reformulate

the problem as an MIS problem and to improve its robustness. In spectral rendering [Wilkie et al. 2014], CMIS helps reduce color noise by combining the contributions of a set of importance-sampled wavelengths. In volume rendering with photon planes [Deng et al. 2019], CMIS refines the weighting functions by forgoing analytical integration.

In summary, our main contributions are:

- an extension of MIS to a continuum of sampling techniques, equipped with a an optimal balance heuristic;
- a practical estimator approximating that optimal combination;
- three applications in light transport simulation where our estimators outperforms existing state-of-the-art techniques.

2 BACKGROUND AND RELATED WORK

Consider the definite integral I of a function $f\colon \mathcal{X} \to \mathbb{R}$ over some domain $\mathcal{X}$ and the n-sample Monte Carlo (MC) estimator $\langle I\rangle_n$ for it:

$$I = \int_{\mathcal{X}} f(x)\,\mathrm{d}x, \qquad \langle I\rangle_n = \frac{1}{n}\sum_{i=1}^{n}\frac{f(x_i)}{p(x_i)}, \tag{1}$$

where $\mathrm{d}x$ denotes an appropriate differential measure on $\mathcal{X}$. The variance of $\langle I\rangle_n$ generally becomes small when the sampling probability density function (PDF) $p(x)$ is approximately proportional to the integrand $f(x)$. Finding a single PDF that closely approximates f, however, is often difficult in practice.

Veach and Guibas [1995] proposed MIS as a means to combine several estimators with different PDFs p_t (for $t = 1, \ldots, T$), where each p_t potentially approximates a different feature of the integrand f. We refer to this method as *discrete MIS* (DMIS) to emphasize that it considers a *countable* set of techniques.

The two DMIS estimators proposed by Veach and Guibas can be derived by introducing a set of T functions $w_t(x)$ satisfying $\sum_{t=1}^{T} w_t(x) = 1$ and partitioning the integral I into a sum:

$$I = \int_{\mathcal{X}} \underbrace{\sum_{t=1}^{T} w_t(x)}_{=1} f(x)\,\mathrm{d}x = \sum_{t=1}^{T} \underbrace{\int_{\mathcal{X}} w_t(x) f(x)\,\mathrm{d}x}_{I_t} = \sum_{t=1}^{T} I_t. \tag{2}$$

The *one-sample DMIS estimator* then estimates this sum by choosing one integral I_t, with probability P_t, and estimating it by drawing a single sample x from the PDF p_t:

$$\langle I\rangle_{\mathrm{DMIS}} = \frac{\langle I_t\rangle_1}{\mathrm{P}_t} = \frac{w_t(x) f(x)}{\mathrm{P}_t p_t(x)}. \tag{3}$$

The *multi-sample DMIS estimator* takes n samples in total, explicitly estimating each integral I_t using $n_t = \mathrm{P}_t n$ samples $x_{t,i}$ from p_t:

$$\langle I\rangle_{\mathrm{MDMIS}} = \sum_{t=1}^{T}\langle I_t\rangle_{n_t} = \sum_{t=1}^{T}\frac{1}{n_t}\sum_{i=1}^{n_t}\frac{w_t(x_{t,i}) f(x_{t,i})}{p_t(x_{t,i})}. \tag{4}$$

The weighting functions w_t can be arbitrarily chosen as long as they satisfy $\sum_{t=1}^{T} w_t(x) = 1$ whenever $f(x) \neq 0$ and $w_t(x) = 0$ whenever $p_t(x) = 0$ [Veach 1997]. The *DMIS balance heuristic*

$$\hat{w}_t(x) = \frac{\mathrm{P}_t p_t(x)}{\sum_{t'=1}^{T} \mathrm{P}_{t'} p_{t'}(x)} \tag{5}$$

is a provably good choice, and Veach and Guibas [1995] showed that it minimizes the variance of the one-sample estimator (3).

DMIS provides a general framework for combining estimators. However, its generality leaves room for improvement when applied to certain problems or when additional information about the integrand and estimators is known. Considerable research has focused on optimizing the sample allocation among techniques [Pajot et al. 2011; He and Tang 2014; Havran and Sbert 2014; Sbert et al. 2016; Sbert and Havran 2017; Sbert et al. 2018]. Effort has also gone into improving the weighting heuristics; Georgiev et al. [2012], Popov et al. [2015], and Sbert et al. [2018] proposed schemes that take advantage of domain-specific auxiliary information. The method of Grittmann et al. [2019] augments the balance heuristic with variance estimates, and Karlík et al. [2019] instead proposed to optimize the sampling densities for balance-heuristic combination. Kondapaneni et al. [2019] derived the truly optimal weighting function for the multi-sample estimator (4) by allowing weights to be negative. Elvira et al. [2015] investigated variants of DMIS with stochastic technique selection, similarly to our stochastic MIS estimator (Section 3.2). Built upon DMIS, all these prior works are limited to handling countable sets of techniques.

More closely related to our work, Deng et al. [2019] have considered a continuum of photon-plane orientations for computing single scattering in participating media. They proposed an analytical weighting scheme over this continuum to reduce variance and ameliorate the effect of singularities. We formalize this basic idea by introducing a general framework for combining uncountable sets of sampling techniques. We also show how other rendering problems can benefit from being interpreted as such combinations.

3 CONTINUOUS MULTIPLE IMPORTANCE SAMPLING

To handle uncountable sets of sampling techniques, we devise an extension of DMIS that we call *continuous MIS*. We first establish our formulation by extending the dimensionality of the integration problem. We then derive a CMIS estimator as an ordinary MC estimator for that extended integral, along with a continuous balance heuristic with provable optimality similar to that of DMIS. We also show how DMIS can be derived from our continuous formulation. Lastly, we propose a practical unbiased estimator, based on stochastic technique selection, that makes CMIS accessible to a broad range of integral estimation problems.

3.1 CMIS formulation

We begin by denoting $\mathcal{T}$ a space (of arbitrary dimension) that permits sampling according to a certain PDF. We call $\mathcal{T}$ the *technique space* and its elements t each identify a sampling technique. We also introduce the notion of a continuous weighting function $w\colon \mathcal{T} \times \mathcal{X} \to \mathbb{R}$ with the property $\int_{\mathcal{T}} w(t,x)\,\mathrm{d}t = 1$, where $\mathrm{d}t$ is an appropriate differential measure on $\mathcal{T}$. With this definition, we extend the dimension of the integral I in a way similar to Eq. (2), using Fubini's theorem:

$$I = \int_{\mathcal{X}} \underbrace{\int_{\mathcal{T}} w(t,x)\,\mathrm{d}t}_{=1} f(x)\,\mathrm{d}x = \int_{\mathcal{T}}\int_{\mathcal{X}} w(t,x) f(x)\,\mathrm{d}x\mathrm{d}t. \tag{6}$$

Note that in this formulation the weight function *integrates* to one, as opposed to *summing* to one in the DMIS formulation (2).

CMIS estimator. Our *continuous MIS* (CMIS) estimator arises as an ordinary one-sample estimator for the extended integral (6):

$$\langle I\rangle_{\text{CMIS}} = \frac{w(t,x)f(x)}{p(t,x)} = \frac{w(t,x)f(x)}{p(t)p(x|t)}. \tag{7}$$

Here, t and x are continuous random variables distributed according to a joint PDF $p(t,x)$ which factorizes into $p(t)p(x|t)$ when the sampling of the integration variable x depends on t. The estimator (7) bears similarity to the one-sample DMIS estimator (3). The term $p(t)$ is now the density (rather than unitless probability) for choosing technique t, and $p(x|t)$ is the technique's (conditional) PDF. For the estimator to be unbiased, the weighting function w must satisfy

$$(C_1) \quad \int_{\mathcal{T}} w(t,x)\,\mathrm{d}t = 1 \text{ whenever } f(x) \neq 0, \text{ and} \tag{8a}$$

$$(C_2) \quad w(t,x) = 0 \text{ whenever } p(t,x) = 0. \tag{8b}$$

Analogously to DMIS, these two conditions imply that at any point x where $f(x) \neq 0$, there is a t for which $p(t,x) = p(t)p(x|t) > 0$.

CMIS balance heuristic. Subject to the conditions in Eq. (8), the weighing function w can be arbitrarily chosen in our CMIS estimator (7). Different choices will lead to different variance of the estimator. A simple uniform weight $w_{\text{u}}(t,x) = 1/\int_{\mathcal{T}} \mathrm{d}t$, however, is unlikely to result in a good technique combination. We ideally want to choose a w that minimizes the variance of the CMIS estimator (7). Appendix A derives this variance-optimal weighting function:

$$\bar{w}(t,x) = \frac{p(t)p(x|t)}{\int_{\mathcal{T}} p(t')p(x|t')\,\mathrm{d}t'} = \frac{p(t,x)}{\int_{\mathcal{T}} p(t',x)\,\mathrm{d}t'} = \frac{p(t,x)}{p(x)}. \tag{9}$$

This weighting function can be naturally interpreted as the continuous counterpart to the classical DMIS balance heuristic (5); that function is also normalized over all available techniques and minimizes the variance of the one-sample DMIS estimator (3).

On the right-hand side of Eq. (9), $p(x) = \int_{\mathcal{T}} p(t',x)\,\mathrm{d}t'$ is the marginal PDF of x. Using the optimal $\bar{w}$ in Eq. (7) thus corresponds to an ordinary MC estimator for Eq. (1) that samples from this PDF:

$$\langle I\rangle_{\text{CMIS}} = \frac{\bar{w}(t,x)f(x)}{p(t,x)} = \frac{\cancel{p(t,x)}f(x)}{p(x)\cancel{p(t,x)}} = \frac{f(x)}{p(x)}. \tag{10}$$

Using the DMIS balance heuristic (5) similarly corresponds to an ordinary MC estimator for Eq. (1), which samples from a weighted sum of distributions [Veach 1997]. In the continuous case of CMIS, the summation becomes integration, and the resulting distribution becomes a marginal rather than a weighted sum.

Discussion. Note that simply letting the number of techniques T in DMIS grow infinitely does not yield CMIS in the limit. Deng et al. [2019] took this approach in an attempt to support MIS over a continuous space of photon-plane orientations. However, even when $T \to \infty$, the number of sampling techniques remains *countably* infinite with a cardinality of $\aleph_0$ [Enderton 1977]; this, in fact, is readily supported by DMIS via Eq. (2). Our CMIS formulation instead explicitly considers a continuous space $\mathcal{T}$ representing an *uncountably* infinite set of techniques with a cardinality of $\aleph_1 > \aleph_0$. Such a continuous formulation requires fundamental changes in the definition of the weighting function and the use of the probability *density* of choosing techniques. Additionally, our formulation allows $\mathcal{T}$ to have arbitrary dimension, which we will leverage in Sections 4 to 6.

Going in the opposite direction (i.e., reducing CMIS to a countable set of techniques) is comparatively straightforward. One way is to directly discretize the space $\mathcal{T}$ by selecting a set of techniques, turning the CMIS integral (6) into the DMIS sum (2) which can then be estimated via Eqs. (3) and (4). Another option is to make the joint PDF $p(t,x)$ piecewise constant w.r.t. the identifier t. Estimation can then be done via the CMIS estimator (7), or by again reducing Eq. (6) to Eq. (2) by leveraging the effective partitioning of $\mathcal{T}$ into subspaces $\mathcal{T}_i$, each with an associated technique t_i with PDF $p_i(x) = p(x|t_i)$. The integral over $\mathcal{T}$ in Eq. (6) can then be broken up into a sum of integrals over $\mathcal{T}_i$, and denoting $w_i(x) = \int_{\mathcal{T}_i} w(t,x)\,\mathrm{d}t$ yields Eq. (2).

Finally, a direct multi-sample counterpart to Eq. (4) for CMIS would not be practical as it would require drawing an uncountably infinite total number of samples, even with one sample per technique. As an alternative, the CMIS estimator (7) could be generalized to the form in Eq. (1), averaging over $n(t)$ samples for each technique t. For clarity, we restrict our discussion to the 1-sample case.

3.2 SMIS formulation

Evaluating the balance-heuristic CMIS estimator (10) requires evaluating the marginal PDF integral $p(x)$ in its denominator. While this integral is sometimes available in closed form, often it is not. One approach is to construct an unbiased estimator for $1/p(x)$ [Booth 2007] which appears as a multiplicative term in the estimator. However, this method requires special care when $p(x) > 1$, which is generally the case as $p(x)$ is a probability density. We take a simpler approach to approximating the CMIS estimator.

SMIS estimator. We achieve unbiased estimation by starting from a balance-heuristic CMIS estimator (10) that averages over n independent sample pairs $(t_1,x_1),\ldots,(t_n,x_n)$. We then *reuse* the samples $t_i \sim p(t_i)$ to estimate the marginal $p(x)$ in the balance heuristic $\bar{w}$ (9):

$$\underbrace{\frac{1}{n}\sum_{i=1}^{n} \frac{p(t_i,x_i)}{\int_{\mathcal{T}} p(t,x_i)\mathrm{d}t} \cdot \frac{f(x_i)}{p(t_i,x_i)}}_{n\text{-sample CMIS}} \approx \underbrace{\frac{1}{n}\sum_{i=1}^{n} \frac{p(t_i,x_i)}{\frac{1}{n}\sum_{j=1}^{n} \frac{p(t_j,x_i)}{p(t_j)}} \cdot \frac{f(x_i)}{p(t_i,x_i)}}_{n\text{-sample SMIS approximation}}. \tag{11}$$

Using $p(t_i,x_i) = p(t_i)p(x_i|t_i)$, the approximation on the right-hand side simplifies to

$$\langle I\rangle_{\text{SMIS}} = \sum_{i=1}^{n} \frac{\dot{w}(t_i,x_i)f(x_i)}{p(x_i|t_i)} = \sum_{i=1}^{n} \frac{f(x_i)}{\sum_{j=1}^{n} p(x_i|t_j)}, \tag{12}$$

where $\dot{w}(t_i,x) = \frac{p(x|t_i)}{\sum_{j=1}^{n} p(x|t_j)}$. Note that the approximation in Eq. (11) yields a *biased* estimator for the CMIS estimator on the left, due to Jensen's [1906] inequality. However, it does still give an *unbiased* estimator for the sought integral I, as we show in Appendix B. We call the result in Eq. (12) our *stochastic MIS (SMIS)* estimator.

Relation to DMIS. Our SMIS estimator closely resembles a DMIS estimator (4) with $T = n$ techniques and one sample for each (i.e., $n_t = 1, \forall t$). The key difference is that in SMIS each technique is chosen stochastically from a set $\mathcal{T}$ for every independent realization (i.e., evaluation) of the estimator. SMIS can be interpreted as estimating a random DMIS-discretization of the CMIS integral (6) (see discussion in Section 3.1). DMIS corresponds to the special case where all techniques are fixed for all realizations. In fact, both DMIS

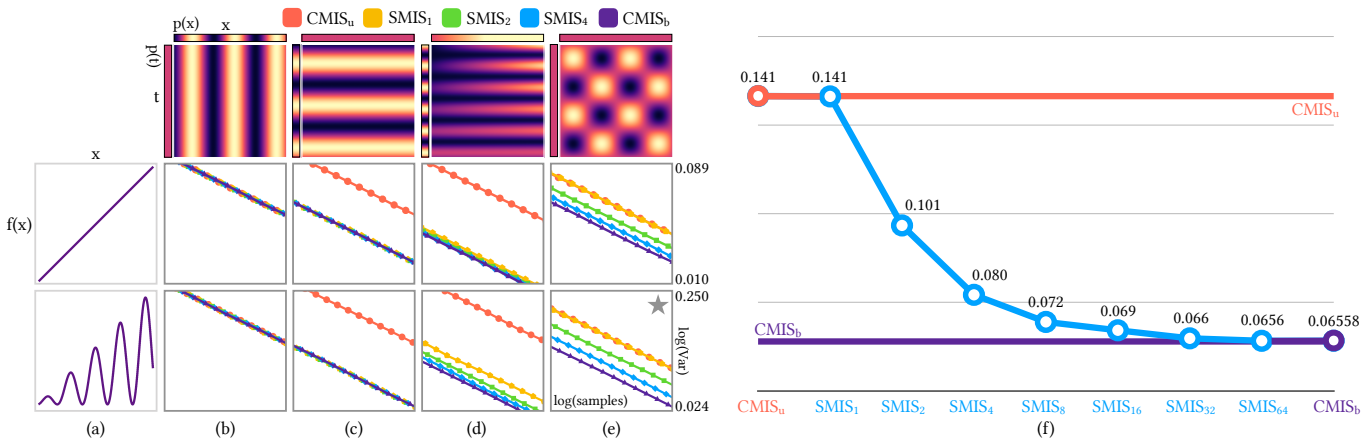

Fig. 2. Experiment comparing the variance of our CMIS and SMIS estimators on two 1D functions (a) integrated using four 2D joint PDFs (top row), i.e., a 1D technique space. For CMIS, we consider uniform (CMIS_u) and balance-heuristic weighting (CMIS_b), as well as several n-sample SMIS variants (SMIS_n). In every plot, at every measurement point we equalize the total number of samples among estimators by adjusting their number of realizations. As expected, CMIS_u performs worst. The SMIS_n variants have consistently lower variance, approaching that of the optimal CMIS_b as n increases. In (f) we plot this convergence for the ⋆-labeled integrand-PDF configuration. Increasing the number of combined samples is an effective way to reduce the variance of SMIS.

estimators (3) and (4) can be made stochastic without voiding their unbiasedness by sampling the techniques prior to estimation, as we discuss in Appendix B. SMIS thus inherits the flexibility of DMIS in its choice of weighting functions and sample allocation, as well as its unbiasedness conditions. Specifically, the weights must sum up to one over all *selected* techniques, whereas in CMIS the normalization is over all *available* techniques. We focus on the specific SMIS form in Eq. (12) to emphasize its interpretation as a CMIS approximation.

Discussion. The advantage of the balance-heuristic SMIS estimator (12) over its CMIS counterpart (10) is that the former provides a practical way of combining a continuum of techniques by only evaluating the conditional PDFs $p(x|t)$ of a finite subset rather than the marginal $p(x)$. Another interpretation is that SMIS provides an unbiased means to approximately sample from that marginal.

On the other hand, SMIS is an approximation to the CMIS estimator and the balance heuristic is thus no longer guaranteed to minimize the variance. We analyze that variance empirically in the following subsection. Future work can investigate the development of better combination heuristics, tailored to SMIS, that account for the stochastic technique selection.

Additionally, SMIS can be more costly than CMIS since its balance heuristic requires evaluating n conditional PDFs for each of n samples. In the case that these n^2 PDF evaluations are too costly, one may instead opt to seek for analytically integrable (albeit sub-optimal) weights with CMIS. We demonstrate the use of such weights in Section 6.

Finally, SMIS also allows the techniques to be chosen from a countable set, as we discuss in Appendix B. Elvira et al. [2015] explored similar estimators where n techniques are selected uniformly out of n (with and without replacement). In contrast, our SMIS allows selecting any subset of techniques, with non-uniform probability. When the number of available techniques—even if finite—is large, SMIS may provide an effective way of trading increased variance for a reduced estimator invocation cost.

3.3 Variance analysis

In the SMIS estimator (12), there are numerous ways to allocate a fixed budget of N technique-sample pairs; below we denote SMIS_n a configuration that takes n pairs in each of N/n (averaged) independent estimator realizations. We perform a numerical experiment, considering three such variants: SMIS_1, SMIS_2, and SMIS_4. We compare these to two CMIS estimators, based on uniform weighting $w_\text{u}(t,x) = 1/\int_\mathcal{T} \mathrm{d}t$ (CMIS_u) and the balance heuristic (9) (CMIS_b), respectively.

Experiment. In Fig. 2, we plot the variance of two 1D integrands $f(x)$ and four 2D joint PDFs $p(t,x)$ with a 1D technique space $\mathcal{T}$. Column (a) shows the two integrands. The top row shows the joint PDFs, along with the corresponding 1D technique-selection marginals $p(t)$ (vertically) and 1D sample-selection marginals $p(x)$ (horizontally). For the variance plots, at each measurement point we average increasing numbers of estimator realizations, adjusted so that each estimator consumes the same number of total samples N.

Intuitively, we expect the provably optimal CMIS_b estimator to perform best, and the "uninformed" CMIS_u to perform worst as it weighs all techniques equally. The SMIS variants, which aim to approximate CMIS_b, should fall somewhere in between.

In Fig. 2b, all sampling techniques are both equally likely, since $p(t) = \text{const}$, and have identical conditional PDFs $p(x|t)$. This results in there being effectively a single sampling distribution $p(x)$, and all estimators reduce to having the same expression.

In Fig. 2c, the joint PDF from (b) is transposed: all techniques t have constant conditionals $p(x|t)$ but $p(t)$ varies. In this case, CMIS_u is affected by the variation in $p(t)$ as it evaluates the full joint PDF $p(x,t) = p(x|t)p(t)$. All other estimators are agnostic to this variation as they only evaluate $p(x|t)$ or $p(x)$.

In Fig. 2d, the technique PDF $p(t)$ is a sinusoid and the $p(x|t)$ vary from uniform to linear. Each conditional PDF $p(x|t)$ has a different variance level, where techniques with a more linear density perform better than techniques with a more uniform density. Here CMIS_u

performs as badly as in (c). SMIS approaches the performance of CMIS_b as the number of samples per realization increases.

Lastly, in Fig. 2e we show a case where the technique and sample marginals are uniform, but the conditional densities $p(x|t)$ vary. When $p(t)$ is constant, the only source of estimation variance is the conditional $p(x|t)$. In this case, the CMIS_u and SMIS_1 estimators perform the same as they effectively both multiply the conditional density $p(x|t)$ by a constant. Once again, increasing the number of samples pairs n in SMIS_n improves its performance asymptotically toward the optimal CMIS_b. For the ⋆-labeled function-PDF configuration, we illustrate this variance convergence in Fig. 2f.

Discussion. In practice, when a continuous set of sampling techniques is available, or when a PDF is parameterized by a continuous parameter t, a reasonable approach is to construct an ordinary one-sample estimator based on a randomly chosen value for t. Such an estimator is identical to the one-sample balance-heuristic SMIS estimator $f(x)/p(x|t)$. With a larger N-sample budget, one would typically average N independent realizations of such estimators. Our results in Fig. 2 indicate that it is better to instead batch as many samples n as possible into N/n realizations of SMIS_n (12). For example, 1 realization of the SMIS_4 estimator shows a lower variance than the average of 4 realizations of SMIS_1, approaching that of the optimal CMIS_b estimator for increasing samples per realization.

Note, however, that SMIS_n requires n^2 PDF evaluations compared to the total of $n^2/2$ for $\text{SMIS}_{n/2}$ over two realizations. Depending on the relative cost of integrand and PDF evaluation, this implies that there exists a configuration that maximizes the efficiency of SMIS. Finding this optimal balance is left as future work.

4 APPLICATION: PATH REUSE

Our first practical application of the CMIS framework targets accelerating Monte Carlo path tracing via path reuse. Path sampling entails a high computational cost whose amortization is desirable, e.g., by reusing (sub)paths across multiple pixel estimations. In prior unbiased approaches, such reuse is limited by the discrete nature of DMIS. Path and (ir)radiance filtering methods [Ward and Heckbert 1992; Křivánek et al. 2005; Keller et al. 2014] are more flexible but add bias and are hampered by simplistic weighting heuristics. We show how CMIS can ameliorate these issues. As a proof of concept, we describe a practical path-space filtering method with a biased variant that achieves significant improvement over prior work.

4.1 Problem statement

The value I of a pixel in the rendered image can be expressed as an integral over the space $\mathcal{P}$ of all possible light transport paths $\overline{\mathbf{x}} = \mathbf{x}_1\mathbf{x}_2\ldots$ that connect the eye and the light sources:

$$I = \int_{\mathcal{P}} f(\overline{\mathbf{x}})\,\mathrm{d}\overline{\mathbf{x}} = \int_{\mathcal{P}} f_\text{e}(\overline{\mathbf{y}}) \underbrace{\int_{\mathcal{P}} f_\text{l}(\overline{\mathbf{y}},\overline{\mathbf{z}})\,\mathrm{d}\overline{\mathbf{z}}}_{I(\overline{\mathbf{y}})}\,\mathrm{d}\overline{\mathbf{y}} = \int_{\mathcal{P}} f_\text{e}(\overline{\mathbf{y}})\, I(\overline{\mathbf{y}})\,\mathrm{d}\overline{\mathbf{y}}, \quad (13)$$

which we decompose into integrals over *prefix subpaths* $\overline{\mathbf{y}}$ and *suffix subpaths* $\overline{\mathbf{z}}$, with $\overline{\mathbf{x}} = \overline{\mathbf{y}}\overline{\mathbf{z}}$.[1] We split the path energy contribution $f(\overline{\mathbf{x}}) = f_\text{e}(\overline{\mathbf{y}})f_\text{l}(\overline{\mathbf{y}},\overline{\mathbf{z}})$ into a contribution $f_\text{e}(\overline{\mathbf{y}})$ from the eye and a contribution $f_\text{l}(\overline{\mathbf{y}},\overline{\mathbf{z}})$ to the light. The latter includes a connection term accounting for the scattering distributions at the last vertex of $\overline{\mathbf{y}}$ and the first vertex of $\overline{\mathbf{z}}$, and their mutual orientation.

[1] Note that $\mathcal{P} = \mathcal{M}^\infty$ is infinite-dimensional, where $\mathcal{M}$ is the set of surface points in the scene. For our purpose its decomposition into $\mathcal{P}\times\mathcal{P} = \mathcal{M}^\infty$ is effectively the same space, since every light transport path has a finite length.

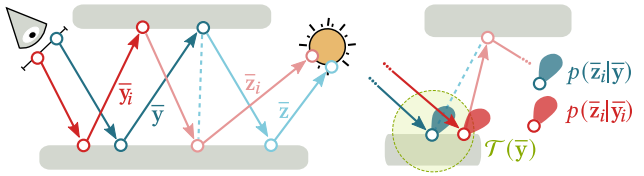

Fig. 3. (a): Two paths are each split into a prefix and a suffix: $\overline{\mathbf{x}} = \overline{\mathbf{y}}\overline{\mathbf{z}}$ and $\overline{\mathbf{x}}_i = \overline{\mathbf{y}}_i\overline{\mathbf{z}}_i$, respectively. The suffix $\overline{\mathbf{z}}_i$ is reused via a connection to prefix $\overline{\mathbf{y}}$. (b): Restricting connections to those suffixes $\overline{\mathbf{z}}_i$ whose prefix endpoints lie in the vicinity of $\overline{\mathbf{y}}$. Unlike previous works' ad-hoc weighting functions, we weight contributions using balance-heuristic SMIS which only requires the conditional densities $p(\overline{\mathbf{z}}_i|\overline{\mathbf{y}})$ and $p(\overline{\mathbf{z}}_i|\overline{\mathbf{y}}_i)$ to achieve an unbiased estimate.

For every pixel, we seek to estimate Eq. (13) by sampling a prefix subpath $\overline{\mathbf{y}}$ from the eye, followed by n suffix subpaths $\overline{\mathbf{z}}_i$:

$$\langle I\rangle = \frac{f_\text{e}(\overline{\mathbf{y}})}{p(\overline{\mathbf{y}})}\langle I(\overline{\mathbf{y}})\rangle_n = \frac{f_\text{e}(\overline{\mathbf{y}})}{p(\overline{\mathbf{y}})}\sum_{i=1}^{n} w(\overline{\mathbf{y}},\overline{\mathbf{z}}_i)\frac{f_\text{l}(\overline{\mathbf{y}},\overline{\mathbf{z}}_i)}{p(\overline{\mathbf{z}}_i)}, \quad (14)$$

where w is some weighting function. Distribution ray tracing [Cook 1986] corresponds to splitting at $\overline{\mathbf{y}}$, with known $p(\overline{\mathbf{z}}_i) = p(\overline{\mathbf{z}}_i|\overline{\mathbf{y}})$ and $w = 1/n$, which is costly as the suffixes $\overline{\mathbf{z}}_i$ are specific to $\overline{\mathbf{y}}$. Instead, we want to utilize non-splitting path tracing [Kajiya 1986] and reconnect the prefix $\overline{\mathbf{y}}$ of each path $\overline{\mathbf{x}}$ to the suffixes $\overline{\mathbf{z}}_i$ of n other paths $\overline{\mathbf{x}}_i = \overline{\mathbf{y}}_i\overline{\mathbf{z}}_i$ (Fig. 3a). Such reuse amortizes the cost of sampling each suffix over multiple pixel estimates. Unfortunately, we cannot directly use the $\overline{\mathbf{z}}_i$ in the estimator (14), for the following reasons:

(1) each suffix $\overline{\mathbf{z}}_i$ is sampled according to a different distribution given its own prefix: $\overline{\mathbf{z}}_i \sim p(\overline{\mathbf{z}}_i|\overline{\mathbf{y}}_i)$;
(2) the possible prefixes $\overline{\mathbf{y}}_i$ form an uncountably infinite set.

Existing path-reuse methods sidestep these issues by restricting the set of prefixes $\overline{\mathbf{y}}_i$ to those generated from a deterministically chosen n pixels [Bekaert et al. 2002; Bauszat et al. 2017], combining the n estimates via balance-heuristic DMIS: $w = \hat{w}$. Filtering based methods instead consider all prefixes $\overline{\mathbf{y}}_i$ with endpoints in the vicinity of the endpoint of $\overline{\mathbf{y}}$, ignoring the above issues and resorting to heuristic weighting w to construct a biased version of Eq. (14) [Keller et al. 2014].

4.2 CMIS formulation

Seeking to lift the limitations of prior work, our key observation in tackling general path reuse is to interpret it as a problem of combining sampling techniques from a continuous space. Indeed, the above two properties are a natural fit for our CMIS framework, by considering the prefixes $\overline{\mathbf{y}}_i$ as *technique identifiers* and the suffixes $\overline{\mathbf{z}}_i$ as the *samples* from these techniques. The technique space $\mathcal{T}$ is then the set $\mathcal{P}$ of all possible prefixes.

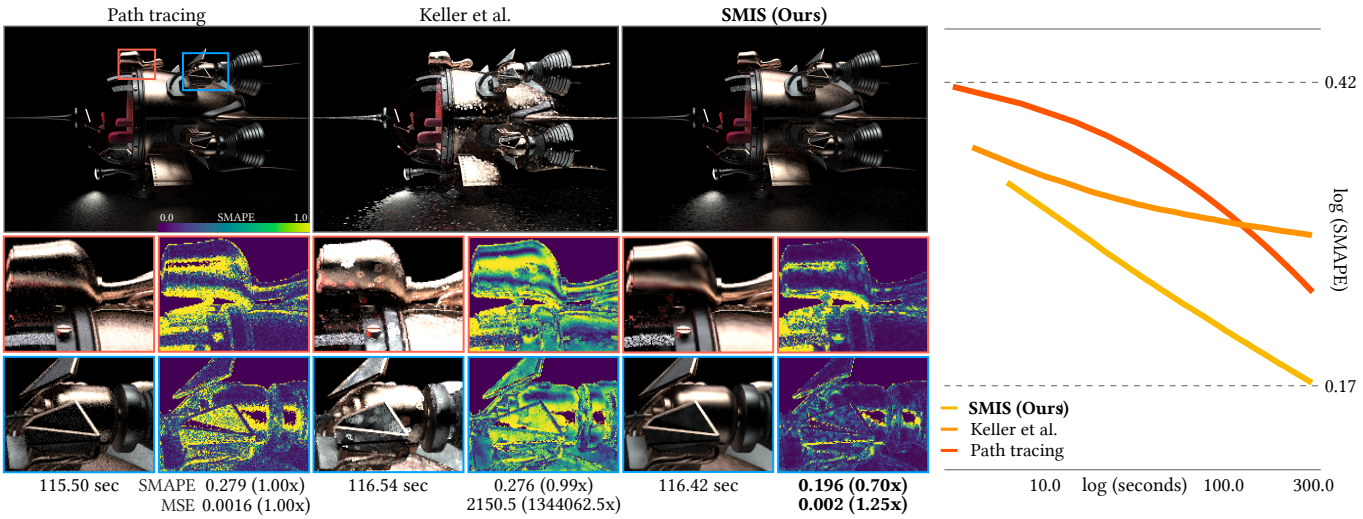

Fig. 4. Equal-time comparison between path tracing, Keller et al.'s [2014] path filtering method, and our SMIS path filtering estimator. At low sample counts, the latter two outperform path tracing by efficiently amortizing path sampling data. Thanks to its robust sample combination, our method outperforms path tracing also in the long run. Notice its accuracy in resolving illumination on small-scale glossy objects—a notoriously difficult task for filtering methods.

Using these definitions of techniques and samples, we extend the suffix integral $I(\overline{\mathbf{y}})$ from Eq. (13) by following the steps in Eq. (6):

$$I(\overline{\mathbf{y}}) = \int_{\mathcal{P}} \underbrace{\int_{\mathcal{T}} w(\overline{\mathbf{y}}', \overline{\mathbf{z}})\, \mathrm{d}\overline{\mathbf{y}}'}_{=1} f_{\mathrm{l}}(\overline{\mathbf{y}}, \overline{\mathbf{z}})\, \mathrm{d}\overline{\mathbf{z}} = \int_{\mathcal{T}} \int_{\mathcal{P}} w(\overline{\mathbf{y}}', \overline{\mathbf{z}}) f_{\mathrm{l}}(\overline{\mathbf{y}}, \overline{\mathbf{z}})\, \mathrm{d}\overline{\mathbf{z}} \mathrm{d}\overline{\mathbf{y}}', \quad (15)$$

where w is normalized over $\mathcal{T}$ for every suffix subpath $\overline{\mathbf{z}}$. Note that $\mathcal{T}$ can depend on $\overline{\mathbf{y}}$ and will typically be $\mathcal{T}(\overline{\mathbf{y}}) \subseteq \mathcal{P}$.

CMIS estimator. We can now construct an estimator for Eq. (15) of the form in Eq. (7) but averaging over n realizations $(\overline{\mathbf{y}}_i, \overline{\mathbf{z}}_i)$:

$$\langle I(\overline{\mathbf{y}}) \rangle_{\mathrm{CMIS}} = \frac{1}{n} \sum_{i=1}^{n} \frac{w(\overline{\mathbf{y}}_i, \overline{\mathbf{z}}_i) f_{\mathrm{l}}(\overline{\mathbf{y}}, \overline{\mathbf{z}}_i)}{p(\overline{\mathbf{y}}_i) p(\overline{\mathbf{z}}_i | \overline{\mathbf{y}}_i)}. \quad (16)$$

This estimator constructs n complete paths by connecting a given prefix $\overline{\mathbf{y}}$ to the suffixes $\overline{\mathbf{z}}_i$ of other, independently sampled paths $\overline{\mathbf{x}}_i = \overline{\mathbf{y}}_i \overline{\mathbf{z}}_i$. Plugging $\langle I(\overline{\mathbf{y}}) \rangle_{\mathrm{CMIS}}$ into Eq. (14) in place of $\langle I(\overline{\mathbf{y}}) \rangle$ yields an unbiased estimator for the pixel value I, as desired.

SMIS estimator. As shown in Section 3.1, the optimal choice of weighting function w in CMIS estimators is the balance heuristic $\bar{w}$ (9). In our setting, evaluating $\bar{w}$ involves computing the normalization factor $p(\mathbf{z}_i) = \int_{\mathcal{T}} p(\overline{\mathbf{y}}') p(\overline{\mathbf{z}}_i | \overline{\mathbf{y}}')\, \mathrm{d}\overline{\mathbf{y}}'$ for each $\overline{\mathbf{z}}_i$, which is the unconditional density for sampling $\overline{\mathbf{z}}_i$ as a continuation of *any possible* prefix subpath $\overline{\mathbf{y}}'$. This integral is generally not available in closed form. We can resort to SMIS which alleviates its evaluation:

$$\langle I(\overline{\mathbf{y}}) \rangle_{\mathrm{SMIS}} = \sum_{i=1}^{n} \frac{\dot{w}(\overline{\mathbf{y}}_i, \overline{\mathbf{z}}_i) f_{\mathrm{l}}(\overline{\mathbf{y}}, \overline{\mathbf{z}}_i)}{p(\overline{\mathbf{z}}_i | \overline{\mathbf{y}}_i)} = \sum_{i=1}^{n} \frac{f_{\mathrm{l}}(\overline{\mathbf{y}}, \overline{\mathbf{z}}_i)}{\sum_{j=1}^{n} p(\overline{\mathbf{z}}_i | \overline{\mathbf{y}}_j)}. \quad (17)$$

Plugged into Eq. (14), $\langle I(\overline{\mathbf{y}}) \rangle_{\mathrm{SMIS}}$ yields an unbiased path-reuse pixel estimator, all of whose terms are evaluable. Notably, the normalization factor in Eq. (17) is the sum of densities for generating the suffix $\overline{\mathbf{z}}_i$ from the n sampled prefixes $\overline{\mathbf{y}}_j$. Compared to a splitting estimator of the form in Eq. (14) [Cook et al. 1984], our SMIS (17) amortizes the cost of sampling each suffix across n pixel estimates.

4.3 Practical path-space filtering algorithm

The SMIS estimator (17) provides a framework for unbiased subpath reuse in path tracing. Practical implementations, however, need to be mindful about setting the parameter n which controls the amount of reuse. Each connection to a suffix $\overline{\mathbf{z}}_i$ entails evaluating the n densities $p(\overline{\mathbf{z}}_i | \overline{\mathbf{y}}_j)$ for sampling its first point from the endpoints of every prefix $\overline{\mathbf{y}}_j$. Each estimator realization thus requires casting n^2 visibility rays. This cost can be managed by shrinking the technique space $\mathcal{T}$, for example restricting it to the vicinity of the endpoint of $\overline{\mathbf{y}}$ (Fig. 3b). Even so, the cost of accurate PDF evaluation, which maintains unbiasedness by ensuring zero-PDF samples are assigned zero weight, remains high and can prevent a net improvement.

We borrow approximations from path-space filtering [Keller et al. 2014] to make our method practical. Specifically, we assume reused paths are always visible when evaluating connections $f_{\mathrm{l}}(\overline{\mathbf{y}}, \overline{\mathbf{z}}_i)$ and conditional densities $p(\overline{\mathbf{z}}_i | \overline{\mathbf{y}}_j)$. This is a reasonable assumption since each $\overline{\mathbf{z}}_i$ is by construction visible from a *nearby* $\overline{\mathbf{y}}_i$. We also reuse the scattering distribution evaluation at the first vertex of $\overline{\mathbf{z}}_i$ rather than re-evaluating it for each connection. These approximations greatly reduce the computational cost by introducing some bias.

The resulting SMIS estimator still has the form of Eq. (17) but does not entail any ray casting apart from sampling the initial full paths. Notably, the weighting function is based on a balance-heuristic combination of techniques, and it is normalized over the *sampled techniques*. This is the crucial difference to the otherwise very similar path-space filtering method of Keller et al. [2014]. Theirs utilizes conservative binary weighting normalized over the *samples*, culling them based on surface similarity at the prefix endpoints. They motivate this choice by empirical findings that directly weighting the samples by such similarity norms can increase variance. In contrast, our balance-heuristic weighting is solely based on sampling densities. Additionally, having derived our estimator from an unbiased one, all sources of bias are known and introduced systematically.

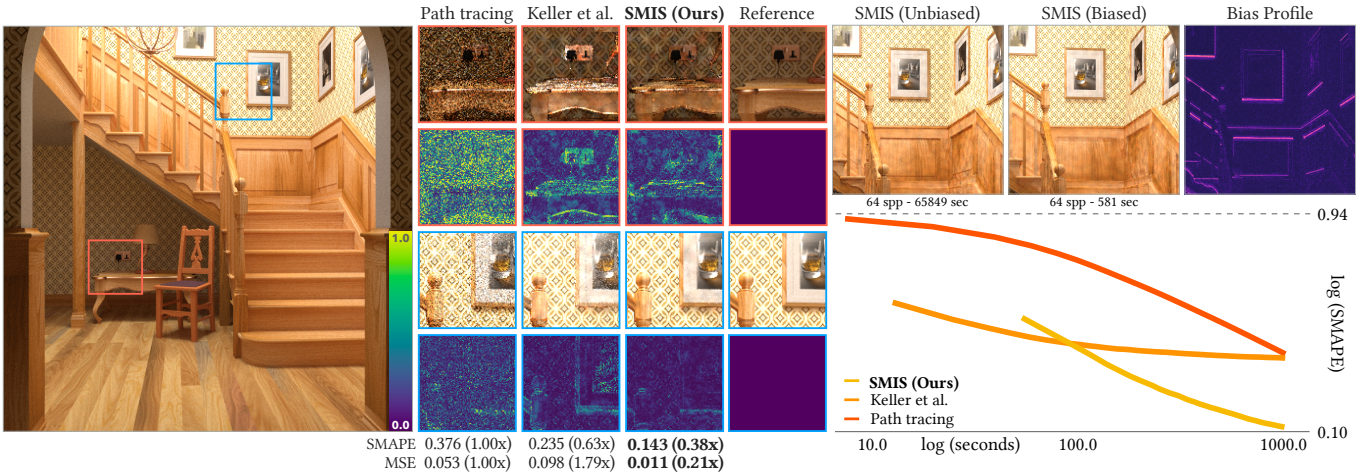

Fig. 5. Left: Equal-time comparison (350 sec) between path tracing, Keller et al.'s [2014] path filtering method, and our SMIS path filtering estimator. The latter two efficiently reduce the noise on diffuse surfaces, however ours also handles glossy materials without introducing significant bias. Right: Equal-sample comparison (64 samples/pixel) demonstrating our that our formulation can completely eliminate the bias in path filtering, though at a significant cost.

4.4 Results

We implemented our method as a two-stage algorithm. In the first stage, we sample all paths, one per pixel, storing their vertices in a hashed spatial grid. In the second stage, for the prefix $\overline{\mathbf{y}}$ of each path, we retrieve all nearby vertices of other paths and invoke our SMIS estimator for $\overline{\mathbf{y}}$ with the corresponding suffixes $\overline{\mathbf{z}}_i$. We set the range-search radius to 6 pixels in screen space, and reduce it at each iteration of progressive rendering.

We compare our method against Keller et al.'s [2014] and baseline path tracing (PT) using the same set of paths. We measure *symmetric mean absolute percentage error* (SMAPE): $E = \frac{1}{P}\sum_{i=1}^{P}\frac{|r_i - e_i|}{|r_i|+|e_i|}$, where r_i and e_i are the reference and estimated values for the i-th pixel, respectively. This metric is robust to outliers and does not rely on an arbitrary "epsilon" regularizer in the denominator, as do other relative metrics such as the relative mean squared error.

The scene in Fig. 4 contains many glossy materials lit by two area lights from the top and two on the space ship. The inability to handle such scenes is a known weakness of traditional filtering approaches. PT remains noisy, even if it can take many more samples than the other two methods in the allotted time. Keller et al.'s filtering replaces this noise with even more conspicuous artifacts. By treating the prefixes $\overline{\mathbf{y}}_i$ as sampling techniques and the suffixes $\overline{\mathbf{z}}_i$ as samples from these techniques, our SMIS weighting achieves a much superior combination. The zoom-ins highlight its ability to resolve detail on small-scale glossy geometry—a notoriously difficult proposition for filtering approaches.

Figure 5 shows a mostly diffuse scene. Here, both filtering methods show noticeable noise reduction compared to PT, particularly on the walls. However, Keller et al.'s method suffers from substantial bias on the glossy wall socket and table. Our method handles these cases well and halves the overall error. In the top right of the figure we compare our practical (yet biased) method to the unbiased variant that traces all rays necessary for accurate PDF and contribution evaluation. In this equal-sample comparison we do not shrink the radius over time. As expected, the additional visibility checks make the unbiased variant significantly more costly, and the biased one slightly blurs shadowed regions.

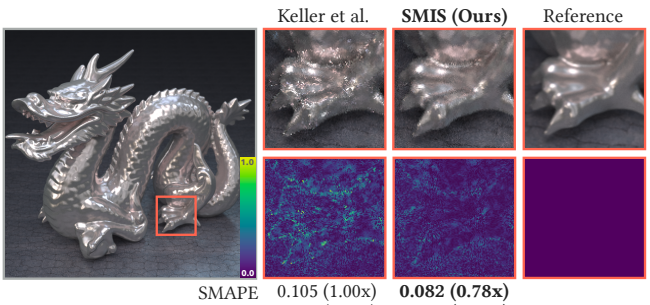

Fig. 6. Equal-sample comparison (8 paths/pixel) of Keller et al.'s [2014] weighting function and our balance-heuristic SMIS, all other things being equal. Keller et al.'s weighting culls most nearby samples in regions of high geometric variation and shows high variance. Our SMIS weighting can robustly handle such samples even on complex glossy geometry.

Figure 6 compares our method to a variant of Keller et al.'s [2014]'s method, where the only difference between the two is in the sample weighting. Keller et al.'s weighting function is based on a series of heuristics chosen to minimize the filtering bias. As a result, it culls most samples in regions of geometric complexity, reducing the filtering effectiveness. In contrast, our SMIS weighting is able to robustly handle such configurations, reducing variance and visual artifacts. Based solely on sampling densities, it manages to preserve the strengths of each technique, as seen in more traditional applications of MIS.

It is worth pointing out that, even when our method achieves lower error, uncorrelated path tracing images can look subjectively better in some cases. This is due to the positive pixel-error correlation inherent to the type of path reuse employed by path filtering. Breaking this correlation, or even turning it negative [Georgiev and Fajardo 2016; Heitz and Belcour 2019] via correlation-aware path reuse, is an interesting direction for future investigation.

5 APPLICATION: SPECTRAL RENDERING

Spectral rendering overcomes the limitations of tri-stimulus (RGB) rendering by extending the path integral (13) over the wavelength domain Λ (Fig. 7a):

$$I = \int_{\mathcal{P}} \int_{\Lambda} f(\overline{\mathbf{x}}, \lambda)\, \mathrm{d}\lambda \mathrm{d}\overline{\mathbf{x}} = \int_{\mathcal{S}} f(\overline{\mathbf{s}})\, \mathrm{d}\overline{\mathbf{s}}, \tag{18}$$

where $\mathcal{S} = \mathcal{P} \times \Lambda$ is the spectral path space and $\overline{\mathbf{s}} = \{\overline{\mathbf{x}}, \lambda\} \in \mathcal{S}$.

Estimating the above integral requires additionally sampling a wavelength λ for each path $\overline{\mathbf{x}}$, resulting in increased variance in the form of color noise (see Figs. 8 and 9). To reduce this noise, the method of Wilkie et al. [2014] evaluates multiple wavelengths for each path. Given a sampled wavelength λ_h (called the *hero wavelength*), which conditions the sampling of $\overline{\mathbf{x}}$, their method evaluates n spectral paths $\{\overline{\mathbf{x}}, \lambda_1\}, \ldots, \{\overline{\mathbf{x}}, \lambda_h\}, \ldots, \{\overline{\mathbf{x}}, \lambda_n\}$. The wavelength samples are uniformly spaced (Fig. 7b), thus *every* sample λ_i in that set could be the hero wavelength. This corresponds to having n sampling techniques for each wavelength, which are combined via balance-heuristic DMIS:

$$\langle I \rangle_{\text{HeroMIS}} = \sum_{i=1}^{n} \frac{f(\overline{\mathbf{x}}, \lambda_i)}{\sum_{j=1}^{n} p(\lambda_j) p(\overline{\mathbf{x}}|\lambda_j)}. \tag{19}$$

Such a multi-wavelength evaluation allows this estimator to achieve tangible reduction in color noise. However, its uniform wavelength spacing is sub-optimal for spectral power distributions (SPDs) that concentrate energy at a few peaks (e.g., fluorescent lights). We show how our CMIS framework can be leveraged for further noise reduction via more flexible and effective wavelength importance sampling.

5.1 CMIS formulation

We begin by extending the dimensionality of the spectral path integral (18) via a weighting function w over some technique space $\mathcal{T}$, and then construct a CMIS estimator, following Section 3.1:

$$I = \int_{\mathcal{S}} \underbrace{\int_{\mathcal{T}} w(\xi, \overline{\mathbf{s}})\, \mathrm{d}\xi}_{=1} f(\overline{\mathbf{s}})\, \mathrm{d}\overline{\mathbf{s}}, \qquad \langle I \rangle_{\text{CMIS}} = \frac{w(\xi, \overline{\mathbf{s}}) f(\overline{\mathbf{s}})}{p(\xi, \overline{\mathbf{s}})}. \tag{20}$$

Here, ξ denotes a technique that conditions the sampling of $\overline{\mathbf{s}}$, i.e., $p(\xi, \overline{\mathbf{s}}) = p(\xi) p(\overline{\mathbf{s}}|\xi)$. The optimal choice for w is the CMIS balance heuristic (9), however the marginal of $p(\xi, \overline{\mathbf{s}})$ over ξ may not be available in closed form. We can use SMIS instead:

$$\langle I \rangle_{\text{SMIS}} = \sum_{i=1}^{n} \frac{\dot{w}(\overline{\mathbf{s}}_i, \xi_i) f(\overline{\mathbf{s}}_i)}{p(\overline{\mathbf{s}}_i|\xi_i)} = \sum_{i=1}^{n} \frac{f(\overline{\mathbf{x}}, \lambda_i)}{p(\lambda_i|\xi_i) \sum_{j=1}^{n} p(\overline{\mathbf{x}}|\lambda_j, \xi_j)}, \tag{21}$$

where $\dot{w}(\overline{\mathbf{s}}_i, \xi_i) = \frac{p(\overline{\mathbf{x}}|\lambda_i, \xi_i)}{\sum_{j=1}^{n} p(\overline{\mathbf{x}}|\lambda_j, \xi_j)}$ and $p(\overline{\mathbf{s}}|\xi) = p(\lambda|\xi) p(\overline{\mathbf{x}}|\lambda, \xi)$. On the right-hand side, we restrict the estimator to reuse the same base path $\overline{\mathbf{x}}$ for every spectral path, i.e., $\overline{\mathbf{s}}_i = \{\overline{\mathbf{x}}, \lambda_i\}$, similarly to Wilkie et al. [2014]. The difference to their estimator is that our SMIS formulation provides more freedom in the distribution of wavelengths λ_i (Fig. 7c). Specifically, they can be sampled according to an arbitrary distribution and independently from each other. The sampling of $\overline{\mathbf{x}}$ is conditioned on one of the wavelengths λ_i, chosen uniformly.

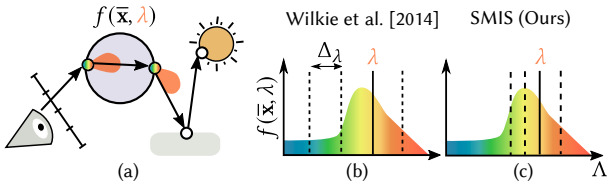

Fig. 7. (a): The contribution of path $\overline{\mathbf{x}}$ depends on the wavelength λ. Sampling and evaluating a single wavelength per path produces color noise in the rendered image. (b): Wilkie et al.'s [2014] method reduces color noise by applying DMIS on a set of equi-spaced wavelengths. (c): Our SMIS estimator is more flexible and allows importance sampling the wavelengths on the continuous spectral domain, resulting in further color noise reduction.

5.2 Implementation and results

For our experiments, we used the CIE 1931 standard observer chromatic response curves, and converted the images into the sRGB space. We used the technique of Jakob and Hanika [2019] to convert input RGB reflectance values into smooth spectra. All compared images use an equal number of samples.

In our implementation of the SMIS estimator (21), every technique ξ gives the same distribution, i.e., $p(\overline{\mathbf{s}}|\xi) = p(\overline{\mathbf{x}}, \lambda) = p(\lambda) p(\overline{\mathbf{x}}|\lambda)$. Each wavelength λ_i is sampled according to a PDF $p(\lambda)$ that is proportional to the product of the observer response and a mixture of the scenes' light-source SPDs. We use a stratified sample pattern that is warped according to that PDF; to demonstrate the flexibility of our formulation, we also consider a variant that performs unstratified, independent sampling (referred to as SMIS_i below).

We compare our two SMIS variants to four other estimators: ordinary Uniform and SpectralIS which sample one wavelength per path, HeroMIS (19), and an unpractical brute-force estimator for Eq. (18) averaging the contributions of 512 wavelength samples for each path. Each realization of HeroMIS and our two SMIS variants evaluates $n = 4$ wavelengths λ_i. All estimators use the aforementioned importance PDF $p(\lambda)$, except for Uniform which samples wavelengths uniformly. Note that HeroMIS can use this PDF only for the hero wavelength as the rest are equally spaced (see Fig. 7b). Our SMIS can importance sample each wavelength individually.

The scene in Fig. 8 is lit by a narrow-band F10 illuminant. The product of its SPD with the observer spectra is a spiky distribution that necessitates importance sampling, as evident when comparing Uniform to SpectralIS. HeroMIS brings some improvement but the uniform wavelength spacing prevents it from accurately importance sampling the paths' spectral contributions (see Fig. 7c). Our SMIS alleviates this restriction, and both variants achieve noticeably lower color noise, with SMIS approaching the unpractical brute-force estimator. On the top right of Fig. 8 we compare the effect of different wavelength importance PDFs on the color noise.

Figure 9 shows a scene containing smooth dielectric objects with wavelength-dependent refraction indices. For paths passing through these objects, only a single wavelength has non-zero contribution, thus neither our SMIS nor HeroMIS provides benefit. However, similarly to Fig. 8, for non-specular paths our SMIS estimator achieves error close to the unpractical brute-force estimator, thanks to more effective wavelength importance sampling.

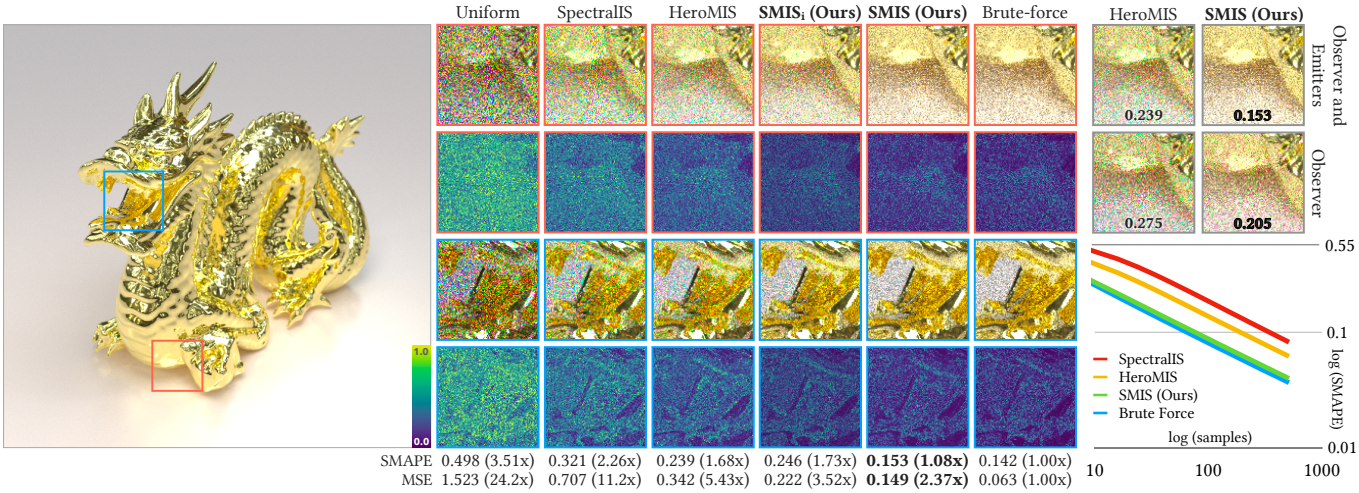

Fig. 8. Gold dragon with under an F10 illuminant with a spiky SPD. Wavelength importance sampling can significantly reduce color noise in this scene, as observed in SpectralIS compared to Uniform. By evaluating four wavelengths per path, HeroMIS and our SMIS achieve further improvement. Thanks to accurately placing each individual wavelength sample, SMIS is more effective, with results very close to brute-force 512-wavelength evaluation.

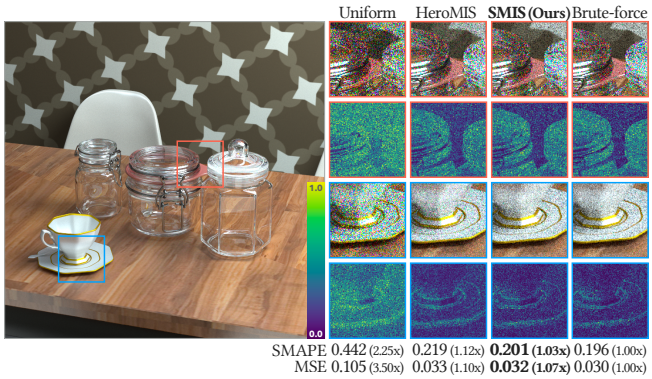

Fig. 9. Scene containing multiple, slightly dispersive glass objects under a D65 illuminant. Our SMIS achieves the lowest error by accurately importance sampling the wavelengths according to the product of the observer response and the illuminant SPD.

6 APPLICATION: VOLUME SINGLE SCATTERING

A continuum of sampling techniques also arises in some methods for volumetric light transport simulation. We show how the photon planes approach of Deng et al. [2019] can be interpreted as a path sampling technique with PDF dependent on a continuous parameter, i.e., technique identifier. Cast in our framework, their method can be seen as a CMIS estimator that uses a suboptimal closed-form weighting function. We formulate a balance-heuristic SMIS estimator that achieves lower variance through more accurate weighting.

6.1 Problem statement

Deng et al. [2019] introduced the concept of photon planes for computing single scattering from rectangular light sources in participating media. A photon plane is parameterized by a line segment s on the light and an outgoing direction ω_l (see Fig. 10a). An eye ray shot in direction ω_e gathers contribution from a photon plane if it intersects it.

We can interpret the above procedure as a sampling technique that uses a light-source frame $(\mathbf{u}_\alpha, \mathbf{v}_\alpha)$ rotated by an angle α (see Fig. 10b). Given α, we first sample an offset v along $\mathbf{v}_\alpha$, which forms the segment s. Given s, we sample an outgoing direction ω_l; the segment and direction together define the geometry of a (photon) plane with infinite length along ω_l. The technique succeeds to construct a path $\bar{\mathbf{x}}$ if a ray sampled independently from the camera intersects the plane. The last, light-source vertex on that path is determined by projecting the ray-plane intersection point onto the segment s. The resulting pixel estimator reads[2]

$$\langle I \rangle_s = \frac{f(\bar{\mathbf{x}})}{p(\bar{\mathbf{x}} \mid \alpha)}. \tag{22}$$

This estimator is unbiased for any $\alpha \in [0, \pi]$, each corresponding to a different rotation of the photon plane about the segment of the path $\bar{\mathbf{x}}$ that is parallel to ω_l. However, the path PDF $p(\bar{\mathbf{x}}|\alpha)$ contains a geometric singularity in configurations where the eye ray is parallel to the plane [Deng et al. 2019]. This singularity can be ameliorated by considering all plane orientations and weighting down the problematic configurations.

6.2 CMIS formulation

We can leverage our CMIS framework to pose a continuous technique-combination problem: α is a technique identifier and $\mathcal{T} = [0, \pi]$ is the technique space. Similarly to the previous two applications, we extend the dimensionality of the integration problem:

$$I = \int_{\mathcal{P}} f(\bar{\mathbf{x}}) \mathrm{d}\bar{\mathbf{x}} = \int_{\mathcal{T}} \int_{\mathcal{P}} w(\alpha, \bar{\mathbf{x}}) f(\bar{\mathbf{x}}) \, \mathrm{d}\bar{\mathbf{x}} \mathrm{d}\alpha, \tag{23}$$

[2]In practice, one photon-plane sample $(\alpha, v_\alpha, \omega_l)$ is reused for all pixels. This correlation has no effect on the interpretation of the method as a path sampling technique, as is the case with photon beams [Křivánek et al. 2014].

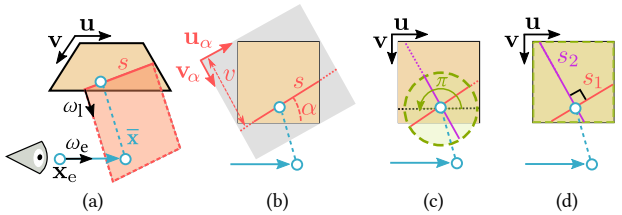

Fig. 10. (a) A path $\bar{\mathbf{x}}$ (in blue) is constructed when an eye ray with direction ω_e intersects a given photon plane (in red). The plane is parameterized by a line segment s on the light source and an outgoing direction ω_l. (b) The segment s is obtained by sampling an offset v in the rotated frame $(\mathbf{u}_\alpha, \mathbf{v}_\alpha)$. (c) Deng et al. [2019] ignore the plane width $\|s\|$ to enable analytic CMIS-style weighting of all plane orientations $\alpha \in [0, \pi]$ that yield the same path $\bar{\mathbf{x}}$. (d) Our SMIS formulation considers a finite set of orientations, accounting for the exact sampling density of each plane (including width).

where w is a valid CMIS weighting function (8). The corresponding balance-heuristic CMIS estimator reads

$$\langle I \rangle_{\text{CMIS}} = \frac{\bar{w}(\alpha, \bar{\mathbf{x}}) f(\bar{\mathbf{x}})}{p(\alpha, \bar{\mathbf{x}})}, \quad \text{with} \quad \bar{w}(\alpha, \bar{\mathbf{x}}) = \frac{p(\alpha, \bar{\mathbf{x}})}{\int_{\mathcal{T}} p(\alpha', \bar{\mathbf{x}})\, \mathrm{d}\alpha'}. \quad (24)$$

This estimator requires evaluating the marginal $\int_{\mathcal{T}} p(\alpha', \bar{\mathbf{x}})\, \mathrm{d}\alpha'$ which is the likelihood of sampling $\bar{\mathbf{x}}$ using any plane orientation. The joint PDF being integrated is

$$p(\alpha, \bar{\mathbf{x}}) = p(\alpha) p(\bar{\mathbf{x}}|\alpha) = p(\alpha) p(v|\alpha) p(\mathbf{x}_e, \omega_l, \omega_e) |(\vec{s} \times \omega_l) \cdot \omega_e|, \quad (25)$$

where $\vec{s}$ is the (normalized) segment direction, $p(v|\alpha) = \|s\|/A$ with $\|s\|$ and A respectively being the segment length and light-source area, and the cosine between the plane normal and ω_e is an area-measure conversion Jacobian. Unfortunately, $p(\alpha, \bar{\mathbf{x}})$ cannot be analytically integrated over α due to the dependence of $p(v|\alpha)$ on the segment length $\|s\|$ which varies with orientation (see Fig. 10b).

CMIS estimation. Deng et al. [2019] devised a specialized weighting scheme to account for an infinite number of photon-plane orientations. In our CMIS framework, their scheme can be interpreted as using a simpler, closed-form weighting function $\tilde{w}$ in place of the CMIS balance heuristic $\bar{w}$. That function is proportional to only the ray-plane cosine which is analytically integrable as it excludes the orientation-varying plane width (see Fig. 10c):

$$\tilde{w}(\alpha, \bar{\mathbf{x}}) = \frac{|(\vec{s} \times \omega_l) \cdot \omega_e|}{\int_{\mathcal{T}} |(\vec{s}' \times \omega_l) \cdot \omega_e| \mathrm{d}\alpha'} = \frac{|(\vec{s} \times \omega_l) \cdot \omega_e|}{\frac{2}{\pi} \sqrt{((\mathbf{u} \times \omega_l) \cdot \omega_e)^2 + ((\mathbf{v} \times \omega_l) \cdot \omega_e)^2}}, \quad (26)$$

where $\mathbf{u}$ and $\mathbf{v}$ are the normalized edges of the light source.

SMIS estimation. In addition to closed-form CMIS weighting, we can employ balance-heuristic SMIS and use exact sampling PDFs. The estimator is a straightforward adaptation of Eq. (12):

$$\langle I \rangle_{\text{SMIS}} = \sum_{i=1}^{n} \frac{\dot{w}(\alpha_i, \bar{\mathbf{x}}_i) f(\bar{\mathbf{x}}_i)}{p(\bar{\mathbf{x}}_i|\alpha_i)} = \sum_{i=1}^{n} \frac{f(\bar{\mathbf{x}}_i)}{\sum_{j=1}^{n} p(\bar{\mathbf{x}}_i|\alpha_j)}. \quad (27)$$

It considers a finite number of photon-plane orientations α_j and uses the exact conditional PDF $p(\bar{\mathbf{x}}_i|\alpha_j)$ from Eq. (25) for each.

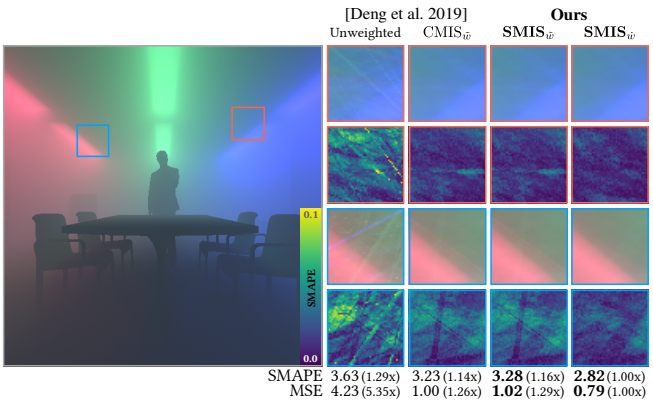

Fig. 11. Equal-sample comparison of previous estimators by Deng et al. [2019] (unweighted and weighted) and our SMIS estimators. The SMAPE and MSE metrics values are scaled by 10^2 and 10^4 respectively. Planes with disproportionately high contribution can occur in both CMIS$_{\tilde{w}}$ and SMIS$_{\tilde{w}}$ as their weighting scheme ignores the photon-plane width. SMIS$_{\dot{w}}$ includes this missing factor to further reduce the variance.

6.3 Implementation and results

We compare four estimators: unweighted (22), Deng et al.'s [2019] analytical approximate weighting (26) (CMIS$_{\tilde{w}}$), and two SMIS variants (27), one using the exact PDFs in the balance heuristic (SMIS$_{\dot{w}}$) and one using the cosine approximation (26) (SMIS$_{\tilde{w}}$). For SMIS, we use just two stratified α samples per estimator (27) realization, making the two planes exactly perpendicular (Fig. 10d). We generate one plane and only upon its intersection we evaluate the second one. This approach introduces a negligible overhead compared to the analytically weighted variants as the second plane involves only weight computation. We use the same set of 40,960 photon planes for all methods, and gather them using one sample per pixel.

Figure 11 shows a scene featuring homogeneous participating media with an isotropic phase function lit by six rectangular light sources. The unweighted estimator produces conspicuous visual artifacts due to the singularity in the photon-plane orientation. CMIS$_{\tilde{w}}$ and SMIS$_{\tilde{w}}$ avoid these artifacts effectively; however, the contributions of narrow, bright photon planes are still visible as thin beams in the rendered image. SMIS$_{\dot{w}}$ properly accounts for the smaller sampling density of such narrow planes and weights them down, yielding further variance reduction.

7 CONCLUSION

This paper formally establishes the first general-purpose framework for combining a continuum of sampling techniques in Monte Carlo integration. Our CMIS formulation is equipped with a provably optimal balance heuristic, and the unbiased SMIS approximation provides a technique combination framework with immediate practical utility. We demonstrate the versatility of our framework on a range of rendering applications, generalizing path reuse and multi-wavelength sampling by framing them as technique-combination problems, and improving photon-plane weighting in volumetric rendering. We see numerous opportunities for future investigation, both in improving the core components of the framework and in

expanding its utility beyond the few example applications in this paper.

Similarly to DMIS, both our CMIS and SMIS frameworks can potentially benefit from designing domain-specific weighting heuristics [Georgiev et al. 2012; Popov et al. 2015; Sbert et al. 2018]. Another interesting route to explore is an alternative to the SMIS balance heuristic with lower variance bounds that incorporates additional information, e.g., variance [Grittmann et al. 2019]. One important aspect of our SMIS is the performance trade-off it comes with: higher-n SMIS$_n$ estimators achieve better variance reduction but at the cost of squared number of PDF evaluations. A formal variance analysis of SMIS could reveal the optimal trade-off. It could also provide guidance for designing novel weighting heuristics.

There are many problems in light transport simulation where DMIS is applied to a continuous domain by enforcing strict limitations, e.g., path reuse [Bekaert et al. 2002; Bauszat et al. 2017] and gradient-domain rendering [Hua et al. 2019]. SMIS provides a simple and general tool for lifting such limitations as we have demonstrated in this paper, enabling the exploration of more efficient estimators. Last but not least, SMIS provides the means for achieving unbiased estimation in settings where unknown marginal densities are involved, e.g., bidirectional instant radiosity [Segovia et al. 2006].

ACKNOWLEDGMENTS

Some of the assets used in this paper come from third-party sources. The Breakfast Room and Staircase scenes in Figs. 5 and 9 are courtesy of Blend Swap user *Wig42*. The Spaceship scene in Figs. 1 and 4 is courtesy of Blend Swap user *thecali*. The Boardroom scene is courtesy of Anat Grynberg, Greg Ward, and Morgan McGuire [2017]. The Dragon model in Figs. 1, 6 and 8 is courtesy of the Stanford Computer Graphics Laboratory. The Glass container models in Fig. 9 are courtesy of Blend Swap user *ruwo*. The cracked asphalt texture in Fig. 6 is courtesy of Tomasz Grabowiecki.

This work has been partially funded by JSPS KAKENHI Grant Numbers JP18KK0309 and JP19H04123, and a grant from Autodesk.

REFERENCES

Anonymous. *"Two random variables were talking in a bar. They thought they were being discrete but I heard their chatter continuously".* Annals of Statistical Jokes.

Pablo Bauszat, Victor Petitjean, and Elmar Eisemann. 2017. Gradient-Domain Path Reusing. *ACM Trans. Graph.* 36, 6, Article Article 229 (Nov. 2017), 9 pages. https://doi.org/10.1145/3130800.3130886

Philippe Bekaert, Mateu Sbert, and John Halton. 2002. Accelerating Path Tracing by Re-Using Paths. In *Proceedings of the 13th Eurographics Workshop on Rendering*. Eurographics Association, 125–134.

T. E. Booth. 2007. Unbiased Monte Carlo Estimation of the Reciprocal of an Integral. 156, 3 (2007), 403–407. https://doi.org/10/gfzq76

Robert L. Cook. 1986. Stochastic Sampling in Computer Graphics. 5, 1 (Jan. 1986), 51–72. https://doi.org/10/cqwhcc

Robert L. Cook, Thomas Porter, and Loren Carpenter. 1984. Distributed Ray Tracing. 18, 3 (July 1984), 137–145. https://doi.org/10/c9thc3

Xi Deng, Shaojie Jiao, Benedikt Bitterli, and Wojciech Jarosz. 2019. Photon surfaces for robust, unbiased volumetric density estimation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 38, 4 (jul 2019). https://doi.org/10.1145/3306346.3323041

Victor Elvira, Luca Martino, David Luengo, and Mónica Bugallo. 2015. Generalized Multiple Importance Sampling. *Statist. Sci.* 34 (11 2015). https://doi.org/10.1214/18-STS668

H.B. Enderton. *Elements of Set Theory.* Elsevier Science.

Iliyan Georgiev and Marcos Fajardo. 2016. Blue-Noise Dithered Sampling. 35:1–35:1. https://doi.org/10/gfznbx

Iliyan Georgiev, Jaroslav Křivánek, Stefan Popov, and Philipp Slusallek. 2012. Importance Caching for Complex Illumination. 31, 2 (June 2012), 701–710. https://doi.org/10/gbbdcc

Pascal Grittmann, Iliyan Georgiev, Philipp Slusallek, and Jaroslav Křivánek. 2019. Variance-Aware Multiple Importance Sampling. *ACM Trans. Graph. (SIGGRAPH Asia 2019)* 38, 6 (2019), 9. https://doi.org/10.1145/3355089.3356515

Vlastimil Havran and Mateu Sbert. 2014. Optimal combination of techniques in multiple importance sampling. 141–150. https://doi.org/10.1145/2670473.2670496

Yuanzhen He and Boxin Tang. 2014. A Characterization of Strong Orthogonal Arrays of Strength Three. 42, 4 (2014), 1347–1360. https://doi.org/10/gfznb4

Eric Heitz and Laurent Belcour. 2019. Distributing Monte Carlo Errors as a Blue Noise in Screen Space by Permuting Pixel Seeds Between Frames. *Computer Graphics Forum* (2019). https://doi.org/10.1111/cgf.13778

Binh-Son Hua, Adrien Gruson, Victor Petitjean, Matthias Zwicker, Derek Nowrouzezahrai, Elmar Eisemann, and Toshiya Hachisuka. 2019. A Survey on Gradient-Domain Rendering. *Computer Graphics Forum* 38 (05 2019), 455–472. https://doi.org/10.1111/cgf.13652

Wenzel Jakob and Johannes Hanika. 2019. A Low-Dimensional Function Space for Efficient Spectral Upsampling. *Computer Graphics Forum (Proceedings of Eurographics)* 38, 2 (March 2019).

J. L. W. V. Jensen. 1906. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Math.* 30 (1906), 175–193. https://doi.org/10.1007/BF02418571

James T. Kajiya. 1986. The Rendering Equation. 20, 4 (Aug. 1986), 143–150. https://doi.org/10/cvf53j

Ondřej Karlík, Martin Šik, Petr Vévoda, Tomáš Skřivan, and Jaroslav Křivánek. 2019. MIS Compensation: Optimizing Sampling Techniques in Multiple Importance Sampling. *ACM Trans. Graph. (SIGGRAPH Asia 2019)* 38, 6 (2019), 12. https://doi.org/10.1145/3355089.3356565

Alexander Keller, Ken Dahm, and Nikolaus Binder. 2014. Path Space Filtering *(SIGGRAPH '14)*. ACM, 68:1–68:1. https://doi.org/10/gfz6mr

Ivo Kondapaneni, Petr Vévoda, Pascal Grittmann, Tomas Skrivan, Philipp Slusallek, and Jaroslav Krivanek. 2019. Optimal Multiple Importance Sampling. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2019)* 38, 4 (July 2019), 37:1–37:14. https://doi.org/10.1145/3306346.3323009

Jaroslav Křivánek, Pascal Gautron, Sumanta Pattanaik, and Kadi Bouatouch. 2005. Radiance Caching for Efficient Global Illumination Computation. 11, 5 (2005), 550–561. https://doi.org/10/csf2sw

Jaroslav Křivánek, Iliyan Georgiev, Toshiya Hachisuka, Petr Vévoda, Martin Šik, Derek Nowrouzezahrai, and Wojciech Jarosz. 2014. Unifying Points, Beams, and Paths in Volumetric Light Transport Simulation. 33, 4 (July 2014), 103:1–103:13. https://doi.org/10/f6cz72

Morgan McGuire. 2017. *Computer Graphics Archive.* https://casual-effects.com/data

Anthony Pajot, Loic Barthe, Mathias Paulin, and Pierre Poulin. 2011. Representativity for Robust and Adaptive Multiple Importance Sampling. 17, 8 (2011), 1108–1121. https://doi.org/10/dpg3dn

Stefan Popov, Ravi Ramamoorthi, Fredo Durand, and George Drettakis. 2015. Probabilistic Connections for Bidirectional Path Tracing. *Comput. Graph. Forum* 34, 4 (July 2015), 75–86.

Mateu Sbert and Vlastimil Havran. 2017. Adaptive multiple importance sampling for general functions. *The Visual Computer* 33, 6 (01 Jun 2017), 845–855. https://doi.org/10.1007/s00371-017-1398-1

Mateu Sbert, Vlastimil Havran, and László Szirmay-Kalos. 2016. Variance Analysis of Multi-sample and One-sample Multiple Importance Sampling. *Computer Graphics Forum* 35 (10 2016), 451–460. https://doi.org/10.1111/cgf.13042

Mateu Sbert, Vlastimil Havran, and Laszlo Szirmay-Kalos. 2018. Multiple Importance Sampling Revisited: Breaking the Bounds. *EURASIP Journal on Advances in Signal Processing* 2018, 1 (Feb. 2018), 15. https://doi.org/10.1186/s13634-018-0531-2

Benjamin Segovia, Jean Claude Iehl, Richard Mitanchey, and Bernard Péroche. 2006. Bidirectional Instant Radiosity. 389–397.

Eric Veach. 1997. *Robust Monte Carlo Methods for Light Transport Simulation.* Ph.D. Thesis. Stanford University, United States – California.

Eric Veach and Leonidas J. Guibas. 1995. Optimally Combining Sampling Techniques for Monte Carlo Rendering, Vol. 29. 419–428. https://doi.org/10/d7b6n4

Gregory J. Ward and Paul S. Heckbert. 1992. Irradiance Gradients. In *CE_EGWR93*, Alan Chalmers, Derek Paddon, and François X. Sillion (Eds.). Consolidation Express Bristol, 85–98.

Alexander Wilkie, Sehera Nawaz, Marc Droske, Andrea Weidlich, and Johannes Hanika. 2014. Hero Wavelength Spectral Sampling. 33, 4 (June 2014), 123–131. https://doi.org/10/f6fgb4

A DERIVATION OF THE CMIS BALANCE HEURISTIC

Here we derive the CMIS balance heuristic (9). We do this by minimizing the variance of the CMIS estimator (7) which reads

$$\mathrm{V}[\langle I\rangle_{\mathrm{CMIS}}] = \mathrm{E}\left[\langle I\rangle_{\mathrm{CMIS}}^2\right] - \mathrm{E}[\langle I\rangle_{\mathrm{CMIS}}]^2 \tag{28a}$$

$$= \int_{\mathcal{X}}\int_{\mathcal{T}} \frac{w^2(t,x)f^2(x)}{p(t,x)}\,\mathrm{d}t\mathrm{d}x - I^2. \tag{28b}$$

Our objective is to find the function $\bar{w}$ that minimizes the variance expression (28b) under the constraints in Eq. (8). Since these constraints are enforced for each x separately, the problem of finding the optimal $\bar{w}$ reduces to minimizing just the inner integral, over $\mathcal{T}$:

$$\bar{w} = \operatorname*{argmin}_w \left[\int_{\mathcal{T}} \frac{w^2(t,x)f^2(x)}{p(t,x)}\mathrm{d}t\right] = \operatorname*{argmin}_w \left[\int_{\mathcal{T}} \frac{w^2(t,x)}{p(t,x)}\mathrm{d}t\right], \tag{29}$$

where the equality on the right holds because the term $f^2(x)$ is constant over $\mathcal{T}$. This problem can be solved using the Euler-Lagrange equation from the calculus of variations as follows. We first define

$$g(t,w) = \frac{w^2(t,x)}{p(t,x)} + \lambda w(t,x), \tag{30}$$

where λ is the (yet unknown) Lagrange multiplier. We then write the Euler-Lagrange equation for g, which we solve for w:

$$\frac{\mathrm{d}g}{\mathrm{d}w} - \frac{\mathrm{d}}{\mathrm{d}t}\underbrace{\left(\frac{\mathrm{d}g}{\mathrm{d}w'}\right)}_{=0} = 0 \quad\Leftrightarrow\quad 2\frac{w(t,x)}{p(t,x)} + \lambda = 0 \tag{31}$$

$$\Leftrightarrow\quad w(t,x) = -2\lambda p(t,x). \tag{32}$$

We can find λ by applying the constraint C_1 (8a), ., by integrating Eq. (32) over $\mathcal{T}$ and equating the result to 1. This yields $\lambda = -1/\left(2\int_{\mathcal{T}} p(t,x)\,\mathrm{d}t\right)$; substituting it back into Eq. (32) yields Eq. (9).

B UNBIASEDNESS OF THE SMIS ESTIMATOR

Here we show the unbiasedness of our SMIS estimator (12) by writing out its expected value over all random variables, i.e., the n independently sampled pairs $(t_i, x_i) \sim p(t_i, x_i) = p(t_i)p(x_i|t_i)$:

$$\mathrm{E}[\langle I\rangle_{\mathrm{SMIS}}] = \mathrm{E}\left[\langle I\rangle_{\mathrm{SMIS}}\right]_{(t_1,x_1),\ldots,(t_n,x_n)} \tag{33a}$$

$$= \mathrm{E}\left[\mathrm{E}[\langle I\rangle_{\mathrm{SMIS}}]_{x_1,\ldots,x_n}\right]_{t_1,\ldots,t_n} \tag{33b}$$

$$= \mathrm{E}\left[\mathrm{E}\left[\sum_{i=1}^{n} \dot{w}(x_i,t_i)\frac{f(x_i)}{p(x_i|t_i)}\right]_{x_1,\ldots,x_n}\right]_{t_1,\ldots,t_n} \tag{33c}$$

$$= \mathrm{E}\left[\underbrace{\sum_{i=1}^{n}\int_{\mathcal{X}} \dot{w}(x,t_i)\frac{f(x)}{\cancel{p(x|t_i)}}\cancel{p(x|t_i)}\,\mathrm{d}x}_{=\sum_{i=1}^{n} I_{t_i} = I;\ \text{see Eq. (2)}}\right]_{t_1,\ldots,t_n} = I. \tag{33d}$$

Note in Eqs. (33c) and (33d) that already the inner expectation over the x variables is equal to I. The outer expectation over the techniques t does not change this result which, in fact, also holds for any DMIS estimator, with any valid DMIS weighting heuristic. Furthermore, note that the result also does not change whether t are continuous or discrete random variables; the set of techniques can thus be countable or uncountable. DMIS corresponds to the special case where t_i are fixed and countable, e.g., finitely many.