Geometric Sample Reweighting for Monte Carlo Integration

JERRY JINFENG GUO and ELMAR EISEMANN, Delft University of Technology, the Netherlands

We present a general sample reweighting scheme and its underlying theory for the integration of an unknown function with low dimensionality. Our method produces better results than standard weighting schemes for common sampling strategies, while avoiding bias. Our main insight is to link the weight derivation to the function reconstruction process during integration. The implementation of our solution is simple and results in an improved convergence behavior. We illustrate its benefit by applying our method to multiple Monte Carlo rendering problems.

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

Additional Key Words and Phrases: Sampling and Reconstruction, Monte Carlo Integration, Sample Reweighting, Rendering

1 INTRODUCTION

Monte Carlo (MC) techniques form the foundation of realistic image synthesis for decades [Cook et al. 1984]. The principle is simple: a function is sampled and the samples are combined to approximate its integral. Standard MC is often referred to as *brute-force*, as its implementation is simple but the variance of the estimation can be high and convergence slow. One method to usually improve the integral approximation is to reconstruct the underlying function from the samples and much previous work devoted its attention to particular cases (e.g., shadows [Egan et al. 2009] or depth of field [Soler et al. 2009]). In this work, we revisit the reconstruction process. We derive an easy-to-implement algorithm to compute sample weights that generally improves the approximation when compared to standard weights for general MC integration.

Our observation is that standard sample weights are often less accurate for lower sampling rates because they do not properly reflect the integration domain nor the local sample density. Our weighting scheme considers all samples of a given set and defines weights based on a geometric partitioning of a low-dimensional integration domain. It results in a consistent estimator that outperforms standard weighting schemes. A major contribution of our work is the derivation of an unbiased estimator. It builds upon this partitioning and applies to sets of independent and identically distributed uniform random samples or stratified samples. Specifically, we propose a novel weighting scheme that is easy to implement and builds upon a sound theoretical derivation. It integrates well into existing rendering pipelines, can be parallelized in conjunction with the unbiased estimator, and we demonstrate its benefit over existing schemes via several rendering problems.

We will first cover prior work and MC integration. We then give the motivation behind our approach (Sec. 3) and present the core of our solution (Sec. 4). Numerical performance and applications to rendering are presented in Sec. 5.

2 BACKGROUND

MC methods. Since the 80s [Cook et al. 1984], MC integration plays a major role in rendering complex effects, such as motion blur, depth of field, and soft shadows. The complete light transport is described by the *rendering equation* [Kajiya 1986], which can be solved using *path tracing* as an associated MC solution. Nevertheless, not all samples taken during the evaluation of an integral contribute strongly to the result. One strategy to modify subsequent sample choices is to rely on previous samples, i.e., a Markov process. Metropolis sampling [Veach and Guibas 1997] can handle complex light path configurations by extensively exploring contributing paths once they are discovered. Multidimensional k-d trees [Guo et al. 2018; Hachisuka et al. 2008] can store samples in a global structure, which can then be used as a means to control future sample placement.

While standard Monte Carlo (MC) methods solve a definite integration $I = \int_\Omega f(x)\mathrm{d}x$ of a function f over a finite support $\Omega \subset \mathbb{R}^d$ by using a random sample set ($\{x_i \in \Omega\}$) with the resulting estimator being $\hat{I}_{MC} = \frac{1}{N}\sum_{i=1}^{N} f(x_i)$, *importance sampling* influences the sampling process via a probability distribution function (pdf) $p : \Omega \to \mathbb{R}$ [Veach and Guibas 1995]. The resulting unbiased estimator is:

$$\hat{I}_p = \frac{1}{N}\sum_{i=1}^{N}\frac{f(x_i)}{p(x_i)}, \tag{1}$$

which effectively weighs samples differently. Importance sampling is interesting when having knowledge about the scene. For instance, importance sampling the light source works better in scenes with small or point light sources [Debevec 2008; Dutre et al. 2006]. Sampling according to the BSDF works better with glossy to highly glossy surfaces[Lafortune et al. 1997; Shirley 1991; Ward 1992]. Multiple importance sampling (MIS) combines different such sampling strategies [Veach and Guibas 1995].

Reweighting. Our solution focuses on the weighting of samples interpreted as an improved function reconstruction. Different weight definitions have been shown to be beneficial for rendering, e.g., derived in Sobolev spaces [Marques et al. 2018]. However, these previous solutions target hemispherical illumination integrals and are not generally applicable to other problems. A reweighting scheme was also proposed for addressing firefly artifacts [Zirr et al. 2018] but the solution is biased and limited to narrow application scenarios.

Other specialized reconstruction techniques exist, including solutions for soft shadows [Egan et al. 2009], defocus blur [Soler et al. 2009], and motion blur [Egan et al. 2011], which lead to significant

Authors' address: Jerry Jinfeng Guo, j.guo-3@tudelft.nl; Elmar Eisemann, e.eisemann@tudelft.nl, Computer Graphics and Visualization, Department of Intelligent Systems, Faculty of EEMCS, Delft University of Technology, the Netherlands.

XXXX-XXXX/2019/7-ART
https://doi.org/10.1145/nnnnnnn.nnnnnnn

improvements. More complex reconstructions for light fields [Lehtinen et al. 2011] have proven very successful but are biased (though consistent).

Our method is independent of the application scenario and unbiased. It handles general functions and links the weights to *Voronoi* cell volumes. The latter has also been studied in the context of anti-aliasing problems [Mitchell 1990], for which the 2 dimensional voronoi cell volumes bounded with a pixel are directly used as sample weights and leads to improved anti-aliasing effects, but the theory has not been further developed for unbiased solutions, nor generalized to other contexts. Voronoi cell size has been used as weights for Monte Carlo integration in [Vorechovsk et al. 2016], where two ways of treating boundaries have been proposed. In this work, Voronoi cells of given set of samples within a domain are either bounded and clipped by the domain boundary, or extended by periodically adding auxiliary samples that extend the domain. Both approaches are shown to improve numerical performance of MC integrations. However, as we show in the Sec.4, directly using Voronoi cell size as weight results in an biased estimate. Our solution takes advantage of Voronoi tessellation and remains unbiased.

3 FORMULATION AND PROBLEM STATEMENT

Referring to Eq. 1, the estimator of importance sampling is a sum of function value $f(x_i)$ times a weight $\Delta(x_i)$, generally:

$$\Delta(x_i) = \frac{1}{Np(x_i)}. \tag{2}$$

Similarly, Riemann integration approximates an integral using function values $f(x_i)$ times a weight $w(x_i)$:

$$\hat{I} = \sum_{i=1}^{N} w(x_i) f(x_i). \tag{3}$$

The Riemann weights stem from a partitioning of the support Ω into *hypervolumes*. In 1D, these hypervolumes are intervals. Each hypervolume contains exactly one sample and its volume defines the sample's weight.

The weights $\Delta(x_i)$ are typically easy to compute but cannot be considered hypervolumes; they would overlap or introduce gaps and cannot easily be linked to a partitioning of Ω. Only with increasing number, due to the stochastic nature of the process, when the samples densely cover the support Ω, the difference in the weight definitions becomes negligible. See Fig. 1 (a), (b) and (c) for an illustration. In consequence, especially for low sample counts, the weights do not well reflect an approximation of the function.

4 GEOMETRIC SAMPLE REWEIGHTING

Our goal is to associate weights to samples that define an improved function reconstruction during the integration. We will first define a consistent solution, inspired by Riemann integration. This solution is independent of the sampling pattern and can be applied on any sample set as a post process to improve the approximation. This reweighting is consistent, but not unbiased for all sampling strategies. We then propose a modification to obtain an unbiased estimator for the cases of uniform random samples that are *independent and identically distributed* (i.i.d.), and samples with stratification. See Sec section 6 for the possibility to generalize our method for samples generated with an analytically known pdf.

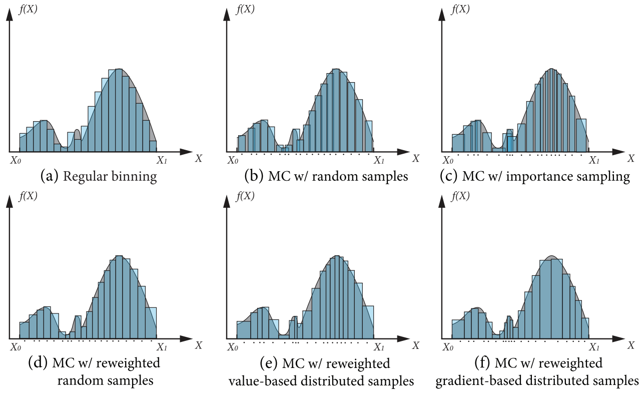

(a) Regular binning

(b) MC w/ random samples

(c) MC w/ importance sampling

(d) MC w/ reweighted random samples

(e) MC w/ reweighted value-based distributed samples

(f) MC w/ reweighted gradient-based distributed samples

Fig. 1. Top row: Three integration methods using the same amount of function evaluations (i.e., 20 samples): (a) Riemann sum through regular binning (according to right side value) (b) MC integration using uniform random samples; (c) MC integration using samples that are distributed according to a pdf w.r.t. function value. Notice that in (a) and (b), the associated bin widths are equal, i.e., $\frac{1}{20}$. Bin widths in (c) are adjusted according to its density determined during sample generation. Notice also the overlaps and gaps between sample bins as illustrated in (b) and (c). Bottom row: illustrations of our methods: (d) uniform random samples with our reweighting; (e) samples distributed according to function value with our reweighting; and (f) samples distributed according to function gradient with our reweighting. Notice the absence of gaps/overlaps and bin widths being adjusted according to sample positions.

4.1 Consistent Estimator

Riemann integration typically assumes a regular partitioning of the domain. Using a *Voronoi diagram* of the sample points, it is possible to partition the domain Ω to take sample density into account. A Voronoi diagram is a partition into regions such that the points in each region share the same closest sample location. It can be shown that the Voronoi cell corresponds to the intersection of half spaces defined by hyperplanes that are equidistant to two sample points. The theory of Voronoi diagrams is beyond the scope of this paper but more details can be found in [Aurenhammer 1991; De Berg et al. 1997].

In our case of a D dimensional problem setting, the diagram will be bounded by the hypercube $(0, 1)^D$, the domain from which samples are drawn. The volume of each Voronoi cell determines the corresponding sample weight and given that the cells are intersections of half-spaces, they are convex and their volume can be easily computed. The resulting estimator of our approach is $\hat{I}_{CON} = \sum_{i=1}^{N} w_{CON}(x_i) f(x_i)$, where $w_{CON}(x_i) = \frac{|V_i|}{|\Omega|}$.

Implicitly, this construction approximates the integrand via a piecewise-constant representation. Intuitively, to take the most benefit from this interpretation, samples should be chosen with respect to the gradient of the function. Fig. 1 shows an illustration of this strategy. In principle, even more advanced approximations could be used, yet it turns out that such weight definitions, while consistent, lead to a biased estimate. In the following, we will show the reasoning behind this and derive an unbiased estimator for i.i.d. uniform sampling and stratified sampling.

4.2 Deriving an Unbiased Estimator

i.i.d. Uniform Samples. The reason the direct use of the Voronoi cells' volume is biased is due to the samples whose cell shares a boundary with the domain boundary. To illustrate this situation, we will first consider the 1D case with a set X of N (with $N > 3$) i.i.d. **uniform** samples $X := \{x_i \in (0,1)\}$ before generalizing to D dimensions, and then to stratified sampling.

One Dimension. Let us assume that the one-dimensional sample set X is sorted from smallest to largest value. We are interested in the expected extent of each Voronoi cell, for which we need to derive the expected distance between two adjacent samples. For this reason, we first determine the expected position of sample x_i.

From *order statics* [David and Nagaraja 2004], we know that the distribution of the i-th i.i.d. sample follows the *beta* distribution, i.e.,

$$p_i(x) = \frac{x^{i-1}(1-x)^{N-i}}{\int_0^1 t^{i-1}(1-t)^{N-i}\mathrm{d}t}.$$

The expected position of the ordered i-th sample x_i is then:

$$\begin{aligned}
E[x_i] &= \int_0^1 x \cdot \frac{x^{i-1}(1-x)^{N-i}}{\int_0^1 t^{i-1}(1-t)^{N-i}\mathrm{d}t}\mathrm{d}x \\
&= N \cdot \binom{N-1}{i-1} \cdot \int_0^1 x \cdot x^{i-1}(1-x)^{N-i}\mathrm{d}x \\
&= N \cdot \binom{N-1}{i-1} \cdot \int_0^1 x^i \cdot \sum_{k=0}^{N-i} \binom{N-i}{k} 1^{N-i-k}(-x)^k \mathrm{d}x \\
&= N \cdot \binom{N-1}{i-1} \cdot \int_0^1 \sum_{k=0}^{N-i} \binom{N-i}{k} (-1)^k x^{k+i} \mathrm{d}x \\
&= N \cdot \binom{N-1}{i-1} \cdot \sum_{k=0}^{N-i} \binom{N-i}{k} \frac{(-1)^k}{i+k+1} \\
&= N \cdot \frac{(N-1)!}{(N-i)!(i-1)!} \cdot \frac{\Gamma(i+1)\Gamma(N-i+1)}{\Gamma(N+2)} \\
&= N \cdot \frac{(N-1)!}{(N-i)!(i-1)!} \cdot \frac{i!(N-i)!}{(N+1)!} = \frac{i}{N+1}.
\end{aligned}$$

Consequently, we have $E[|x_i - x_{i+1}|] = \frac{1}{N+1}$ for $i = 1$ to $N-1$ and a similar condition holds for $E[x_1 - 0] = E[1 - x_N] = \frac{1}{N+1}$. The expected weight is then $\frac{1}{N+1}$ for samples x_i with $i = 2$ to $N-1$ and $\frac{3}{2}\frac{1}{(N+1)}$ for samples x_1 and x_N. The latter weights are larger due to the intervals containing the two boundaries of the domain. Using these weights directly, leads to a consistent but biased estimator.

To render the estimator unbiased, we introduce a per-sample correction coefficient C:

$$\hat{I}_{GR} = \sum_{i=1}^{N} w_{GR}(x_i) f(x_i), \text{ where } w_{GR}(x_i) = \frac{|V_i|}{C(x_i)\,|\Omega|}. \tag{4}$$

These factors have to be carefully chosen — for instance, $C = 1$ would lead to the previously-derived consistent but biased result. The correction coefficient should indeed modify the expected contribution of a sample x_i to equal $\frac{1}{N}f(x_i)$. Following the weight derivation, an unbiased estimator in 1D, we would then define $C(x_1) = C(x_N) = \frac{2(N+1)}{3N}$ and $C(x_i) = \frac{N+1}{N}$ for all other samples. As most samples still share an identical correction factor, it keeps us close to the interpretation of the Voronoi cell volume. In higher dimensions, the definition is less straightforward.

D Dimensions. To derive the correction coefficient C from Eq.4 in D dimensions, we assume a set of N (with $N \geq 3^D$, i.e., intuitively, this results in at least one inner point and two boundary points along each dimension) samples in $\Omega = (0,1)^D$. We define the *boundary order* $b(x_i)$ of a sample as the amount of cell boundaries of its Voronoi cell that are part of the domain boundary. For instance, in the above one dimensional example, $b(x_1) = b(x_N) = 1$, and for all other sample points, we have $b(x_i) = 0$.

The cardinality of samples of order d is defined as: $|X_d| = \operatorname{card}\{b(x_i) = d, \forall i \in [1,N]\}$. For such a sample set of N samples, the expected cardinality of samples of order d is $E[|X_d|] = \binom{D}{d}(\sqrt[D]{N} - 2)^d 2^{D-d}$. This formula is the d-th term in the bionomial expansion of $[(\sqrt[D]{N} - 2) + 2]^D$. To understand this result, one should recall that the expected position of all samples forms a regular grid. Thus, this grid will have a resolution of $n = \sqrt[D]{N}$ along each axis. Starting with one axis, we would find n samples with two boundary samples of order one and all others samples are inner points of order zero. Repeating these samples n times along a new dimension will increment the order of the first repeated set of samples and the last, as these represent a new boundary along this dimension. For all other samples, their boundary order remains unchanged. This process can be done for all D dimensions, thus implying the binomial expansion.

To achieve an unbiased estimator, we first compute the expected Voronoi volume $E[V_i]$ for a sample x_i. For D dimensions, we have $D+1$ boundary orders from 0 to D. As we are dealing with an i.i.d. uniform distribution, in each dimension, we have $n-1$ intervals between samples and two intervals with the boundary, leading to a total of $n+1$ intervals. Therefore, we have:

$$E[V_i] = \left(\frac{3}{2}\right)^{b(x_i)} \frac{|\Omega|}{n+1}$$

Again, for unbiasedness, we need $E[w_{GR}(x_i)] = 1/N$, thus each sample should expectedly contribute equally. The following definition of the correction coefficients fulfills this property:

$$C(x_i) = \left(\frac{3}{2}\right)^{b(x_i)} \frac{N}{n+1}, \tag{5}$$

because

$$\begin{aligned}
E[w_{GR}(x_i)] &= E\left[\frac{|V_i|}{C(x_i)\,|\Omega|}\right] = \frac{1}{|\Omega|}E\left[\frac{|V_i|}{C(x_i)}\right] \\
&= \frac{1}{|\Omega|}\frac{\left(\frac{3}{2}\right)^{b(x_i)}\frac{|\Omega|}{n+1}}{\left(\frac{3}{2}\right)^{b(x_i)}\frac{N}{n+1}} = \frac{1}{|\Omega|} \times \frac{|\Omega|}{N} = \frac{1}{N}. \blacksquare
\end{aligned}$$

Stratified Samples. The extension to stratified sampling is relatively straightforward, as each stratum is considered an independent unit. This means that the function is independently integrated in each stratum and its whole range is a composition of these units. In consequence, the boundary observation now applies to the boundary of each stratum. For a sample set X of size N generated with S strata, each stratum is expected to contain $\frac{N}{S}$ samples. Let $n = \sqrt[D]{N}$ and $s = \sqrt[D]{S}$, then the integration problem for a stratum with $\frac{N}{S}$ samples would imply the correction coefficients to be:

$$C(x_i) = \left(\frac{3}{2}\right)^{b(x_i)} \frac{N}{n+s}. \tag{6}$$

5 RESULTS

In this section, we first show the numerical performance (Sec. 5.1) of our scheme and show its application to a few rendering scenarios (Sec. 5.2). In all tests, we compare four different estimators:

(1) Standard MC (i.i.d. uniform sampling)
(2) Our weighted standard MC (i.i.d. uniform sampling)
(3) Stratified MC (stratified sampling)
(4) Our weighted stratified MC (stratified sampling)

5.1 Numerical Performance

The numerical performance is tested with two examples: one for a 1D MC integration and the other for a 2D MC integration, which are plotted in Figure 2. The 1D function is given as:

$$f(x) = 10 \times \begin{cases} \sqrt{-x^2 + 0.5x} & 0 < x < 0.25 \\ -\sqrt{-x^2 + x - 0.1875} + 0.25 & 0.25 < x < 0.5 \\ 20 \times (x - 0.5) & 0.5 < x < 0.55 \\ 1.0 & 0.55 < x < 0.65 \\ -20 \times (x - 0.7) & 0.65 < x < 0.7 \\ 0.1 \times \sin(10\pi \cdot (x - 0.7)) & 0.7 < x < 0.8 \\ 0.25 \times \sin(10\pi \cdot (x - 0.8)) & 0.8 < x < 0.9 \\ 0.5 \times \sin(10\pi \cdot (x - 0.9)) & 0.9 < x < 1.0 \end{cases}$$

For the 2D function, we take the *Lena* image [Munson 1996]. The functions were chosen to include discontinuities, large-scale variations and small scale changes and led to a representative behavior of several tests that we have performed. Generally, the MSE drops as more samples are added (Column 1). Our solutions outperform standard uniform sampling and even stratified sampling by several orders of magnitude and converges around 1000, 100 times faster respectively in 1D and 100, 10 times faster respectively in 2D.

For the case of stratified sampling, we illustrate different amounts of strata for the same sampling count (Column 2). Our weighting scheme makes this parameter less important, as it achieves a better function approximation.

We next investigate the impact of distributing samples into batches for which we estimate the function integration separately, before deriving the overall estimate by averaging, which would typically be the case for distributed computations. First, we fixate the amount of samples to 100K (Column 3). Notice that the performance of standard uniform sampling remains invariant with respect to the amount of samples per batch, as it is already an averaging process. Our solution results in a better approximation for more samples per batch, as it will approximate the function more faithfully, as expected. Similarly, stratified sampling also benefits from more samples per batch, but shows slower convergence.

We also investigate the effect of using different batch sizes for uniform (Column 4) and stratified sampling(Column 5). More batches thus means a higher overall sample count and all methods improve with the addition of batches. In all cases, the graphs stop after reaching 100K samples. Our solution performs best and the graphs also illustrate the convergence over several batches, due to its unbiasedness.

5.2 Application to Rendering

We implemented our method in Mitsuba [Jakob 2010], targeting one and two dimensional integration problems, namely motion blur (Sec. 5.2.1), dispersion (Sec. 5.2.2), depth of field (Sec. 5.2.3) and illumination integrals (Sec. 5.2.4). We evaluate MSE and visual appearance, as well as convergence behavior. For all implementations, our reweighting operates at a per-pixel level. We apply our method on the level of primary samples, thus all applicable local importance sampling techniques are utilized throughout the pipeline.

5.2.1 Motion Blur. To simulate motion blur, distribution rendering samples the time domain: For a pixel (i, j), the luminance $L_{(i,j)}$ is given by:

$$L_{(i,j)} = \int_{t_{\text{open}}}^{t_{\text{close}}} f_{(i,j)}(t)\mathrm{d}t,$$

with t_{open} and t_{close} being the shutter opening and closing time and f incorporating the shutter function. Since time is 1 dimensional, building a Voronoi partition means sorting and measuring the distance between samples. We tested our implementation in two scenes with animation (Fig. 3 and 4).

5.2.2 Spectral Rendering. Light dispersion can happen at reflective or refractive dielectric materials, leading to effects such as rainbows, resulting from different wavelengths travelling in different directions. Spectral sampling simulates multiple wavelengths in order to capture such effects. To reduce the complexity of the additional spectral dimension[Bergner et al. 2009], *hero wavelength spectral sampling*[Wilkie et al. 2014] can be used as an approximation:

$$\hat{I}_{(i,j)} = \frac{1}{N} \sum_{k=1}^{N} \frac{f_{(i,j)}(\lambda_k)}{p_{(i,j)}(\lambda_k)}$$

Our implementation of spectral sampling uses 15-bin wavelengths. Hero wavelength sampling is used with 3 shifted additional wavelength samples [Wilkie et al. 2014]. We tested our method with two scenes configured with dispersive di-electric materials (Fig. 5 and 6).

As shown in the results, our method brings down colour noise significantly and dispersive regions look much smoother at low sample rate.

5.2.3 Defocus Blur. A camera with aperture leads to defocus blur/depth of field effects. The aperture is usually modelled as a 2D shape, e.g., a square, a circle, or a star, which is sampled to determine the origin of each primary sample ray, which passes through the position on the focal plane corresponding to the current pixel. For lens aperture

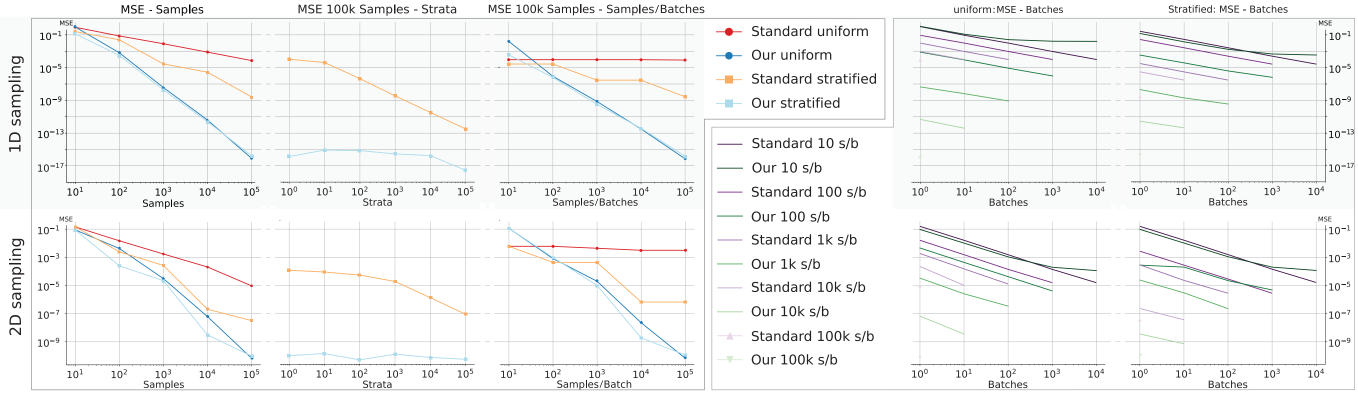

Fig. 2. We apply our geometric sample reweighting to one and two dimensional MC integration problems.

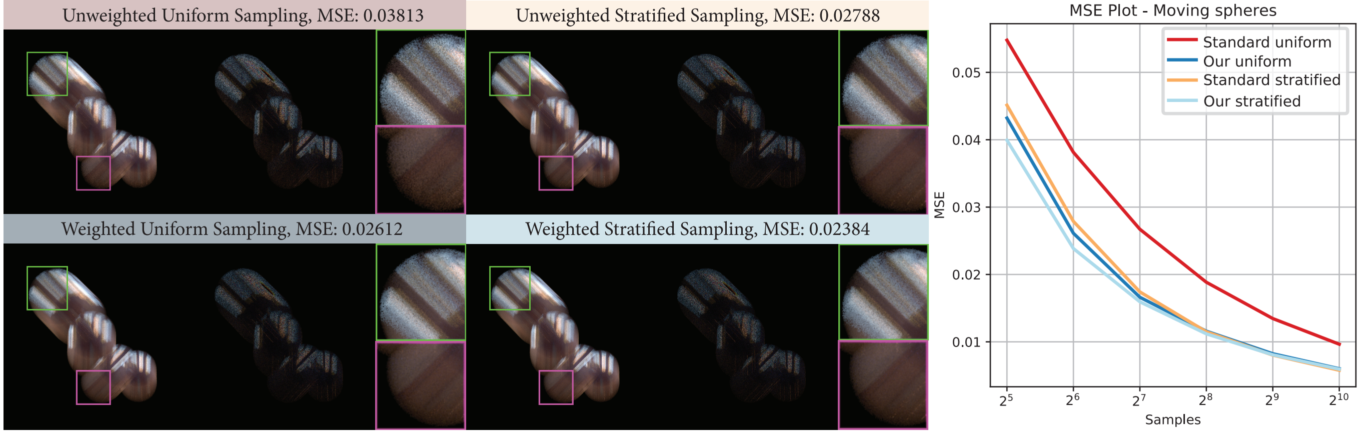
Unweighted Uniform Sampling, MSE: 0.03813
Unweighted Stratified Sampling, MSE: 0.02788
Weighted Uniform Sampling, MSE: 0.02612
Weighted Stratified Sampling, MSE: 0.02384
MSE Plot - Moving spheres
Standard uniform
Our uniform
Standard stratified
Our stratified
MSE
0.05
0.04
0.03
0.02
0.01
Samples
2^5
2^6
2^7
2^8
2^9
2^10

Fig. 3. Four highly glossy spheres moving in different directions with 64 samples per pixel. In each subfigure: corresponding render, difference with reference and highlighted regions.

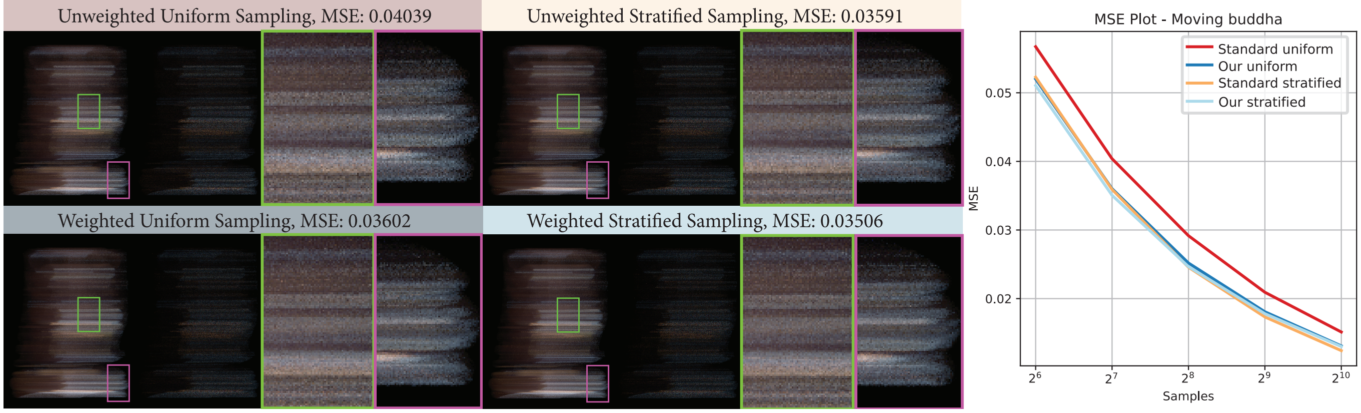
Unweighted Uniform Sampling, MSE: 0.04039
Unweighted Stratified Sampling, MSE: 0.03591
Weighted Uniform Sampling, MSE: 0.03602
Weighted Stratified Sampling, MSE: 0.03506
MSE Plot - Moving buddha
Standard uniform
Our uniform
Standard stratified
Our stratified
MSE
0.05
0.04
0.03
0.02
Samples
2^6
2^7
2^8
2^9
2^10

Fig. 4. Highly glossy Buddha moving horizontally with 128 samples per pixel. In each subfigure: corresponding render, difference with reference and highlighted regions.

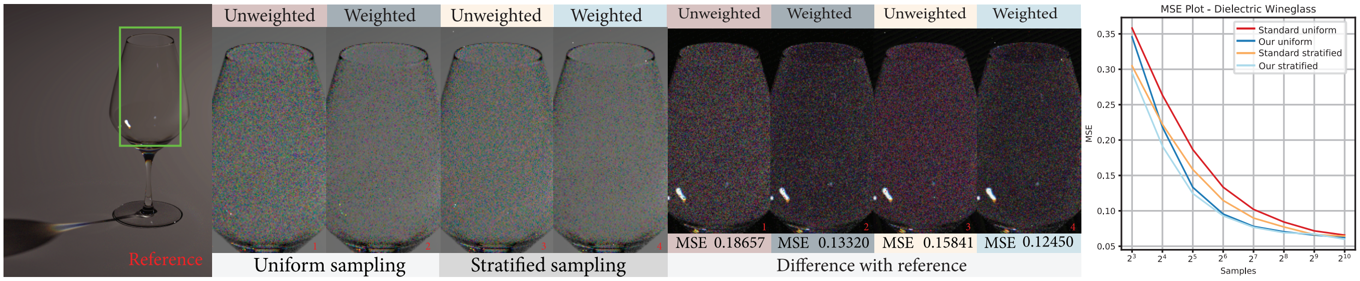

Fig. 5. Wineglass with dispersive dielectric materials with 64 and 32 samples per pixel. MSE for highlighted four plots are: 0.18657, 0.13320, 0.15841 and 0.12450 respectively.

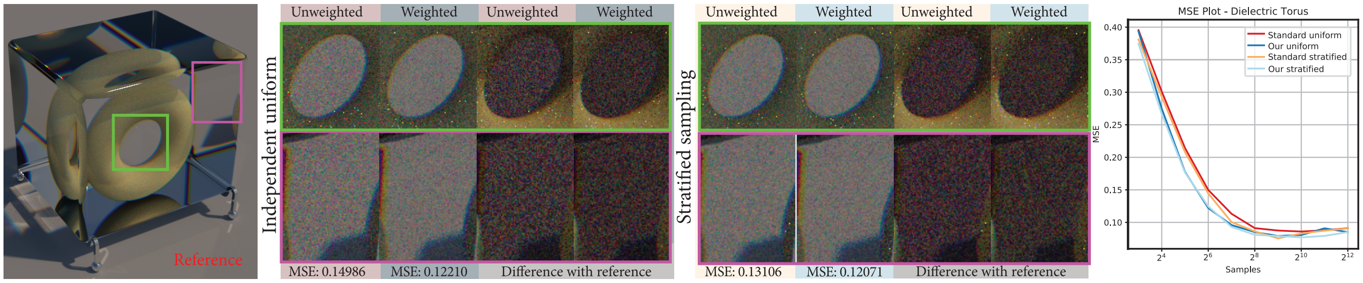

Fig. 6. Torus with dielectric materials with 32 samples per pixel.

$\mathbb{A} \subset \mathbb{R}^2$, we obtain:

$$L_{(i,j)} = \int_{\mathbb{A}} f_{(i,j)}(s)\mathrm{d}s.$$

To determine our weights, we use a 2D Voronoi diagram based on the aperture samples. We tested a simple glossy sphere illuminated using an environment map (Fig. 7).

5.2.4 Direct Illumination. Leaving out irrelevant terms, the luminance $L_{\mathbf{x}}$ at scattering point $\mathbf{x}$ with one bounce is given by:

$$\begin{aligned} L_{\mathbf{x}} &= L_e(\mathbf{x}) + L_{\text{direct}} + L_{\text{indirect}} \\ &= L_e(\mathbf{x}) + \int_{\mathbb{L}} f_s(\mathbf{x}) L_e(l \rightarrow \mathrm{x})\mathrm{d}l + \int_{\Omega} f_s(\mathbf{x}, \omega) L_i(\omega)\mathrm{d}\omega, \end{aligned}$$

where L_e denotes light emission and $l \in \mathbb{L}$ denotes all light sources. In this application, we use light sampling instead of random rays to ensure that the light source is always sampled. Our unbiased reweighting achieves the best convergence and, as shown in the insets, also the smoothest results (Fig. 8).

Observation. In all cases, our solution leads to smoother visual result and less black holes in the falloff regions. From the MSE plots, we can see that standard MC with uniform sampling has the worst performance, while our weighted stratified sampling generally has the best one. Our method improves both uniform sampling and stratified sampling. We can also see that even with uniform sampling as input, our weighted uniform sampling not only improves over the unweighted version, but also has a performance that is as good as our weighted stratified sampling. Precompute sample weights enables a negligible computation overhead.

6 CONCLUSION

The reweighting scheme in this paper enables a better approximation than standard MC weights. Our solution is general and does not require any prior knowledge about the integrating function. Implicitly, our method approximates this function via a reconstruction from the samples, but does not introduce a bias in the resulting estimator. We showed its practical benefit for various rendering problems.

While we focus on primary samples that are either i.i.d. uniform or stratified in this work, our method can also handle non-uniform sample sets following a distribution of $p(x)$. The expected position of the i-th sample x_i is then $N \cdot \binom{N-1}{i-1} \cdot \int_0^1 P^{-1}(x) \cdot x^{i-1}(1-x)^{N-i}\mathrm{d}x$, where $P(u) = \int p(u)\mathrm{d}u$. Unfortunately, it is necessary to integrate the distribution function. Approximate schemes remain an area of future work. Similarly, using the method in higher dimensions requires the computation of cell volumes in high-dimensional Voronoi diagrames, which can be costly. One could precompute these weights but we left such accelerations as future work. Finally, it is an exciting opportunity to exploit the generality of our solution to improve other integration problems.

REFERENCES

Franz Aurenhammer. 1991. Voronoi diagrams—a survey of a fundamental geometric data structure. *ACM Computing Surveys (CSUR)* 23, 3 (1991), 345–405. https://doi.org/10.1145/116873.116880

Steven Bergner, Mark S Drew, and Torsten Möller. 2009. A tool to create illuminant and reflectance spectra for light-driven graphics and visualization. *ACM Transactions on Graphics (TOG)* 28, 1 (2009), 5. https://doi.org/10.1145/1477926.1477931

Robert L Cook, Thomas Porter, and Loren Carpenter. 1984. Distributed ray tracing. In *ACM SIGGRAPH Computer Graphics*, Vol. 18. ACM, 137–145. https://doi.org/10.1145/800031.808590

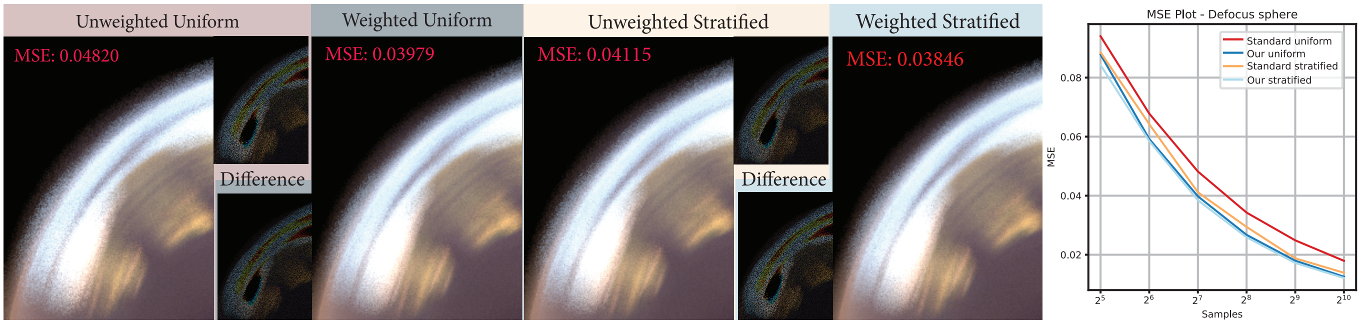

Fig. 7. Glossy sphere crossing the focal plane of the camera with 128 samples per pixel.

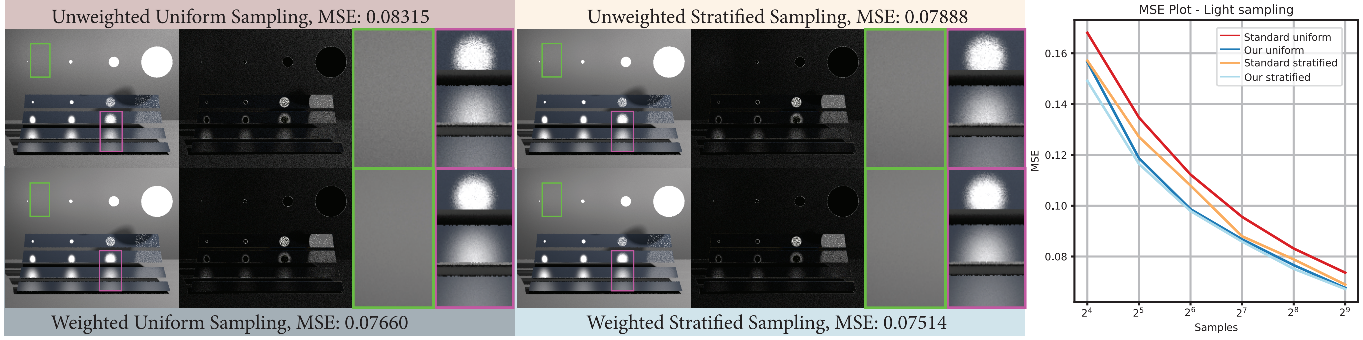

Fig. 8. Each image uses 4 primary rays per pixel and at each scattering event 256 light samples. Our solution is applied to the 256 samples.

Herbert Aron David and Haikady Navada Nagaraja. 2004. *Order statistics* (3rd ed.). Wiley Online Library. https://doi.org/10.1002/0471722162

Mark De Berg, Marc Van Kreveld, Mark Overmars, and Otfried Schwarzkopf. 1997. Computational geometry. In *Computational geometry*. Springer, 1–17.

Paul Debevec. 2008. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *ACM SIGGRAPH 2008 classes*. ACM, 32. https://doi.org/10.1145/1401132.1401175

Philip Dutre, Philippe Bekaert, and Kavita Bala. 2006. *Advanced global illumination*. AK Peters/CRC Press.

Kevin Egan, Florian Hecht, Frédo Durand, and Ravi Ramamoorthi. 2011. Frequency analysis and sheared filtering for shadow light fields of complex occluders. *ACM Transactions on Graphics (TOG)* 30, 2 (2011), 9. https://doi.org/10.1145/1944846.1944849

Kevin Egan, Yu-Ting Tseng, Nicolas Holzschuch, Frédo Durand, and Ravi Ramamoorthi. 2009. Frequency analysis and sheared reconstruction for rendering motion blur. In *ACM Transactions on Graphics (TOG)*, Vol. 28. ACM, 93. https://doi.org/10.1145/1540000

Jerry Jinfeng Guo, Pablo Bauszat, Jacco Bikker, and Elmar Eisemann. 2018. Primary Sample Space Path Guiding. In *Eurographics Symposium on Rendering - Experimental Ideas and Implementations*, Wenzel Jakob and Toshiya Hachisuka (Eds.). The Eurographics Association. https://doi.org/10.2312/sre.20181174

Toshiya Hachisuka, Wojciech Jarosz, Richard Peter Weistroffer, Kevin Dale, Greg Humphreys, Matthias Zwicker, and Henrik Wann Jensen. 2008. Multidimensional adaptive sampling and reconstruction for ray tracing. In *ACM Transactions on Graphics (TOG)*, Vol. 27. ACM, 33. https://doi.org/10.1145/1399504.1360632

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

James T Kajiya. 1986. The rendering equation. In *ACM Siggraph Computer Graphics*, Vol. 20. ACM, 143–150. https://doi.org/10.1145/280811.280987

Eric PF Lafortune, Sing-Choong Foo, Kenneth E Torrance, and Donald P Greenberg. 1997. Non-linear approximation of reflectance functions. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., 117–126. https://doi.org/10.1145/258734.258801

Jaakko Lehtinen, Timo Aila, Jiawen Chen, Samuli Laine, and Frédo Durand. 2011. Temporal light field reconstruction for rendering distribution effects. In *ACM Transactions on Graphics (TOG)*, Vol. 30. ACM, 55. https://doi.org/10.1145/2010324.1964950

Ricardo Marques, Christian Bouville, and Kadi Bouatouch. 2018. Optimal Sample Weights for Hemispherical Integral Quadratures. In *Computer Graphics Forum*. Wiley Online Library. https://doi.org/10.1111/cgf.13392

Don P Mitchell. 1990. The antialiasing problem in ray tracing. *Advanced Topics in Ray Tracing, SIGGRAPH 1990 Course Notes* (1990).

David C Munson. 1996. A note on Lena. *IEEE Transactions on Image Processing* 5, 1 (1996), 3–3.

Peter S Shirley. 1991. *Physically based lighting calculations for computer graphics*. Ph.D. Dissertation. University of Illinois at Urbana-Champaign.

Cyril Soler, Kartic Subr, Frédo Durand, Nicolas Holzschuch, and François Sillion. 2009. Fourier depth of field. *ACM Transactions on Graphics (TOG)* 28, 2 (2009), 18. https://doi.org/10.1145/1516522.1516529

Eric Veach and Leonidas J Guibas. 1995. Optimally combining sampling techniques for Monte Carlo rendering. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. ACM, 419–428. https://doi.org/10.1145/218380.218498

Eric Veach and Leonidas J Guibas. 1997. Metropolis light transport. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., 65–76. https://doi.org/10.1145/258734.258775

M. Vorechovsk, V. Sadilek, and J. Elias. 2016. Application of Voronoi Weights in Monte Carlo Integration with a Given Sampling Plan. In *International Journal of Reliability and Safety, Computing with Polymorphic Uncertain, Proceedings of REC2016*. Inderscience Publishers, 441–452.

Gregory J Ward. 1992. Measuring and modeling anisotropic reflection. In *ACM SIGGRAPH Computer Graphics*, Vol. 26. ACM, 265–272. https://doi.org/10.1145/142920.134078

Alexander Wilkie, Sehera Nawaz, Marc Droske, Andrea Weidlich, and Johannes Hanika. 2014. Hero wavelength spectral sampling. In *Computer Graphics Forum*, Vol. 33. Wiley Online Library, 123–131. https://doi.org/10.1111/cgf.12419

Tobias Zirr, Johannes Hanika, and Carsten Dachsbacher. 2018. Re-Weighting Firefly Samples for Improved Finite-Sample Monte Carlo Estimates. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 410–421. https://doi.org/10.1111/cgf.13335