

Voronoi integration of the rendering equation

Franck Vandewière¹, Nicolas Chenavier², Samuel Delepoulle¹, and
Christophe Renaud¹

¹LISIC, Université du Littoral Côte d'Opale, Calais, France

²LMPA, Université du Littoral Côte d'Opale, Calais, France

*{franck.vandewiele, nicolas.chenavier, samuel.delepoulle,
christophe.renaud}@univ-littoral.fr*

Abstract

In photorealistic image rendering, Monte Carlo methods are the foundation for integration of the rendering equation in modern approaches. However, despite their effectiveness, traditional Monte Carlo methods often face challenges in controlling variance, resulting in noisy visual artifacts in some difficult to render regions of the image.

In this work, we propose a new approach to the integration of the rendering equation, introducing a Voronoi tessellation reweighting backed by a Poisson point process sampling strategy to address some of the limitations of standard Monte Carlo methods.

From a theoretical point of view, we show that the variance based on a Poisson-Voronoi tessellation is smaller than that induced by the Monte Carlo method when the intensity of the underlying process is arbitrarily large and the function to be integrated satisfies a Hölder condition.

1 Monte Carlo methods

1.1 Background on Monte Carlo Integration

Monte Carlo methods are a class of computational algorithms that use random sampling to estimate solutions of integral problems. They do not require prior knowledge of the exact solution and make no demanding assumptions about the functions to integrate. They can be applied to a wide category of problems and are commonly used in various fields such as artificial intelligence, physics, finance, and risk assessment.

Monte Carlo integration has a relatively slow convergence rate, with the error decreasing as $1/\sqrt{n}$, where n is the sample size. This means that to halve the error, it is necessary to quadruple the number of samples, which can be computationally expensive for high-precision results [10].

Nevertheless, it is still a very interesting approach for two main reasons:

- the decrease of the error does not depend on the dimension, which makes the method particularly interesting for high-dimensional problems;
- the method makes no assumptions about the function other than its integrability.

Let $f : \mathbf{R}^d \rightarrow \mathbf{R}$ a function with compact support integrable over \mathbf{R}^d , equipped with the Euclidean norm $|\cdot|$.

To integrate f , Monte Carlo methods assign the same weight to samples $f(x_i)$ taken at locations $(x_i)_{i \leq n}$:

$$\int_{\mathbf{R}^d} f(x) dx \simeq \frac{1}{n} \sum_{i=1}^n f(x_i).$$

This approach, by simply averaging the function values, ignores the spatial distribution and density of the samples. In particular, sampling clumps and clusters may be over-represented in the estimation of the integral.

1.2 Voronoi reweighting

Voronoi tessellations [5] provide an effective representation of the spatial distribution of samples. By defining zones of influence and proximity relations between sites, Voronoi cells offer a convenient and natural means to characterize the samples' geometric features. In particular, reweighting each sample's contribution according to the volume of its Voronoi cell is a straightforward and intuitive method for correcting Monte Carlo estimates, as it naturally mitigates the influence of sample clusters.

Voechovský et al. [15] apply a Voronoi reweighting scheme to identify under-sampled regions of the integration domain. However, the authors note that simply clipping Voronoi cells at the domain boundaries results in a biased estimation. To ensure an unbiased estimate of the integral, they propose to arrange samples in a periodic tiling, which entails a high computational cost. Guo &

Eisemann [2], on the other hand, propose to refine the natural Voronoi reweighting based on clipped cells by applying a correction factor to restore unbiasedness. While this approach is agnostic to the sampling distribution, it is not supported by a theoretical variance upper bound.

In the following, let S be a locally finite subset of \mathbf{R}^d , that is, a set such that its restriction to any compact subset of \mathbf{R}^d is finite. Denote by

$$C_S(x) = \{y \in \mathbf{R}^d : |y - x| \leq |y' - x| \text{ for all } y' \in S \setminus \{x\}\}$$

the Voronoi cell with nucleus x and by $v_S(x)$ its volume.

We consider this Voronoi weighting:

$$\sum_{x \in S} f(x)v_S(x),$$

where f is some real function with compact support in \mathbf{R}^d .

Our goal is to analyze the consistency of this estimator and derive a theoretical upper bound on its variance. We demonstrate its effectiveness through experiments on analytical benchmarks as well as in the challenging context of the rendering equation [4].

2 Approximation of an integral using a Poisson-Voronoi tessellation

For theoretical purposes, we consider below a function $f : \mathbf{R}^d \rightarrow \mathbf{R}$ with compact support W satisfying a Hölder condition, namely

$$|f(x) - f(y)| \leq k|x - y|^\alpha,$$

for some $k \geq 0$ and $\alpha \in (0, 1]$. Let η_n be a Poisson point process with intensity n in \mathbf{R}^d , i.e. a point process in \mathbf{R}^d satisfying the two following properties:

- (i) for any bounded Borel subset B in \mathbf{R}^d , the random variable $\#(\eta_n \cap B)$ has a Poisson distribution with parameter $n \text{vol}_d(B)$;
- (ii) for any pairwise disjoint Borel subsets B_1, \dots, B_k , $k \geq 1$, in \mathbf{R}^d , the random variables $\#(\eta_n \cap B_1), \dots, \#(\eta_n \cap B_k)$ are independent.

From a theoretical point of view, such a process is natural because it allows expectations of sums to be computed using Slivnyak-Mecke formula (see Appendix A). In practice, to simulate it in a bounded window W , one simply generates a random number N following a Poisson distribution with parameter $n \text{vol}_d(W)$ and then, conditionally on N , throws N points independently of each other and uniformly at random in W . Using notations similar to those in Section 1.2, we denote by

$$C_{\eta_n}(x) = \{y \in \mathbf{R}^d : |y - x| \leq |y' - x| \text{ for all } y' \in \eta_n \setminus \{x\}\}$$

the Voronoi cell with nucleus $x \in \eta_n$ and by $v_{\eta_n}(x)$ its volume.

Proposition 1. *With the above notations, we have*

$$(i) \quad \mathbb{E} \left[\sum_{x \in \eta_n} f(x) v_{\eta_n}(x) \right] = \int_{\mathbf{R}^d} f(x) dx;$$

$$(ii) \quad \mathbb{V} \left[\sum_{x \in \eta_n} f(x) v_{\eta_n}(x) \right] \leq cn^{-1-\frac{2\alpha}{d}}, \text{ for some } c > 0 \text{ depending only on } d \text{ and } f.$$

Approximating the integral of a given function (especially when it is Hölder continuous) using the values it takes on a set of points, weighted by the volume of Voronoi cells, is classical and naturally arises in quantization problems (see [6]). However, the underlying methods consider a deterministic set of points rather than a (random) point process and are computationally expensive.

Using a Voronoi tessellation based on a Poisson point process on \mathbf{R}^d (and not only on the support of f) is important in order to obtain assertion (i), because the underlying process is stationary. Conversely, a Voronoi tessellation based on a point process defined only on the support of f would yield a biased estimator of the integral, since it would not be stationary.

Concerning (ii), we notice that the upper bound, namely $n^{-1-\frac{2\alpha}{d}}$, is the term of a convergent series. In particular, thanks to the the Tchebychev's inequality and the Borel-Cantelli lemma, it implies

$$\sum_{x \in \eta_n} f(x) v_{\eta_n}(x) \xrightarrow[n \rightarrow \infty]{} \int_{\mathbf{R}^d} f(x) dx \quad \text{a.s..}$$

It is noteworthy that the variance converges to zero at a faster rate than that of the Monte Carlo method, for which the rate is $O(n^{-1})$. To the best of our knowledge, such a result is new.

An interesting point is that, when f does not satisfy a Hölder condition but is of the form $f = \mathbb{1}_W$, where W is some admissible set of \mathbf{R}^d , then

$$\mathbb{V} \left[\sum_{x \in \eta_n} f(x) v_{\eta_n}(x) \right] \leq cn^{-1-\frac{2}{d}} = O(n^{-1-\frac{1}{d}});$$

see [3, 12, 13]. See also [9] for a recent work on Poisson-Delaunay approximation. Here again, our rate of convergence decreases faster.

See appendix A for a proof of this result.

3 Application to integration problems

3.1 Leveraging Poisson-Voronoi tessellations for numerical estimations

Proposition 1 offers an appealing reduction technique for Monte Carlo estimations of Hölder continuous functions. In the following, we will consider the

numerical computation of the integral of real functions defined in \mathbf{R}^2 with compact support $W = [-\frac{1}{2}, \frac{1}{2}]^2$. To apply this result to this kind of problem, we need to sample a real function with a Poisson point process in \mathbf{R}^2 . To restrict a Poisson process η_n to a finite sampling of a compact window W , we first draw an integer N from a Poisson distribution with parameter $|W|n = n$. Then, given N , x_1, \dots, x_N are sampled uniformly from W .

Even if this construction allows us to reduce it to a sampling equivalent to a Poisson process, our estimation still faces a difficulty: a finite set of nuclei x_1, \dots, x_N gives rise to unbounded Voronoi cells.

For this reason, we consider a *Stretched window Poisson Point Process* (SPPP) sampling strategy: instead of considering a Voronoi tessellation from nuclei drawn in $W = [-\frac{1}{2}, \frac{1}{2}]^2$, we draw random nuclei from a stretched window $W' = [-\frac{1}{2} - \varepsilon, \frac{1}{2} + \varepsilon]^2$, where $\varepsilon > 0$ is chosen to ensure that Voronoi cells of the nuclei contained in W are bounded, with a certain level of confidence.

Note that the extra nuclei drawn from outside W are used to constrain the Voronoi tessellation, and not for sampling function f . By choosing a suitable confidence level, the acceptance-rejection method yield a valid Voronoi partition for estimating the integral of f over W via the result of Proposition 1 with minimum overhead.

To choose ε , we rely on the following argument.

Let us assume that n points are distributed within a window $W = [-\frac{1}{2}, \frac{1}{2}]^2$. Let $W' = [-\frac{1}{2} - \varepsilon, \frac{1}{2} + \varepsilon]^2$ the corresponding stretched window for some ε . We proceed *as if* dealing with a Poisson point process of intensity n in the plane. Let $X_n = \{x_1, \dots, x_n\}$ denote n points sampled independently and uniformly at random in W , and $X'_n = \{x'_1, \dots, x'_{n'}\}$ denote the n' points (also thrown independently and uniformly at random) in the strip $W' \setminus W$, distributed independently of the points in X . Since the area of the strip $W' \setminus W$ is $4\varepsilon + 4\varepsilon^2$, if we want to preserve the density, the number of points n' distributed in the strip is given by $n' = \lfloor n(4\varepsilon + 4\varepsilon^2) \rfloor$. Let $\tilde{X}_n = X_n \cup X'_n$ and observe that $\tilde{X}_n \cap W = X_n$. Recall that η_n denote a Poisson point process of intensity n . We have :

$$\begin{aligned} \mathbb{P}(\text{reject the simulation}) &= \mathbb{P}(\exists x \in \tilde{X}_n \cap W : C_{\tilde{X}_n}(x) \text{ is unbounded}) \\ &\leq \mathbb{P}(\exists x \in \tilde{X}_n \cap W : R(C_{\tilde{X}_n}(x)) > \varepsilon) \\ &\underset{n \rightarrow \infty}{\sim} \mathbb{P}(\exists x \in \eta_n \cap W : R(C_{\eta_n}(x)) > \varepsilon) \end{aligned}$$

where, for any locally finite subset $\chi \subset \mathbf{R}^2$, the quantity $R_\chi(x)$ denotes the circumscribed radius associated with $C_\chi(x)$, i.e.

$$R_\chi(x) = \inf\{r > 0 : B(x, r) \supset C_\chi(x)\}.$$

Moreover

$$\begin{aligned}
\mathbb{P}(\exists x \in \eta_n \cap W : R(C_{\eta_n}(x)) > \varepsilon) &\leq \mathbb{E} \left[\sum_{x \in \eta_n \cap W} \mathbb{1}_{R(C_{\eta_n}(x)) > \varepsilon} \right] \\
&= n \int_W \mathbb{P}(R(C_{\eta_n \cup \{x\}}(x)) > \varepsilon) dx \\
&= n \mathbb{P}(R(C_{\eta_n \cup \{0\}}(0)) > \varepsilon) \\
&= n \mathbb{P}(R(C_{\eta_n \cup \{0\}}(0)) > \varepsilon n^{1/2}).
\end{aligned}$$

In the above expression, the first equality follows from the Slivnyak-Mecke formula (see e.g. Corollary 3.2.3 in [11]), the second equality follows from the facts that the area of W equals 1 and that η_n is stationary, and the third one follows from the relation $\eta_n \stackrel{D}{=} n^{-1/2} \eta_1$. Thanks to [1], Eq. (6) this gives

$$\mathbb{P}(\exists x \in \eta_n \cap W : R(C_{\eta_n}(x)) > \varepsilon) \leq 4\pi n^2 \varepsilon^2 e^{-\pi n \varepsilon^2}.$$

Therefore, as n is large, the probability to reject the simulation is lower than $4\pi n^2 \varepsilon^2 e^{-\pi n \varepsilon^2}$.

In our experimental setup, we use a target level of confidence of $\frac{1}{1,000}$. With $n = 10,000$, for instance, this gives $\varepsilon \simeq 0.0255$.

3.2 Poisson-Voronoi estimators

We recall the notation introduced previously. $X_n = \{x_1, \dots, x_n\}$ represents a set of n points sampled independently and uniformly in the compact window $W = [-\frac{1}{2}, \frac{1}{2}]^2$; W' denotes the stretched window $[-\frac{1}{2} - \varepsilon, \frac{1}{2} + \varepsilon]^2$ for some $\varepsilon > 0$. Conversely, the set of supplementary points $X'_n = \{x'_1, \dots, x'_{n'}\}$ is sampled independently and uniformly in the strip $W' \setminus W$, and $\tilde{X}_n = X_n \cup X'_n$ is the augmented set of initial points.

Let $f : \mathbf{R}^2 \rightarrow \mathbf{R}$ be a function with compact support W . From proposition 1, we derive three estimators for integration problems. Figure 1 summarizes the main characteristics of each estimator.

Voronoi estimator is the immediate adaptation of proposition 1 for a SPPP (Figure 1a):

$$E_V(f) = \sum_{x \in X_n} f(x) v_{\tilde{X}_n}(x).$$

Filtered Voronoi estimator is another take on proposition 1. While SPPP guarantees that the Voronoi cells are bounded, in practice, their diameter might be very large. For this reason, this estimator filters out Voronoi cells that might extend beyond W' . Let $X_n^- \subseteq X_n$ denote the subset of nuclei for which the associated Voronoi cell, computed based on the

augmented configuration \tilde{X}_n , lies entirely within the domain W' (Figure 1b). We define the filtered Voronoi estimator the following way:

$$E_F(f) = \frac{1}{\sum_{x \in X_n^-} v_{\tilde{X}_n}(x)} \sum_{x \in X_n^-} f(x) v_{\tilde{X}_n}(x).$$

Clipped Voronoi estimator discards the stretching strategy entirely and clips Voronoi cells on the borders of W (Figure 1c). This estimator is expected to be biased:

$$E_C(f) = \sum_{x \in X_n} f(x) |C_{X_n}(x) \cap W|.$$

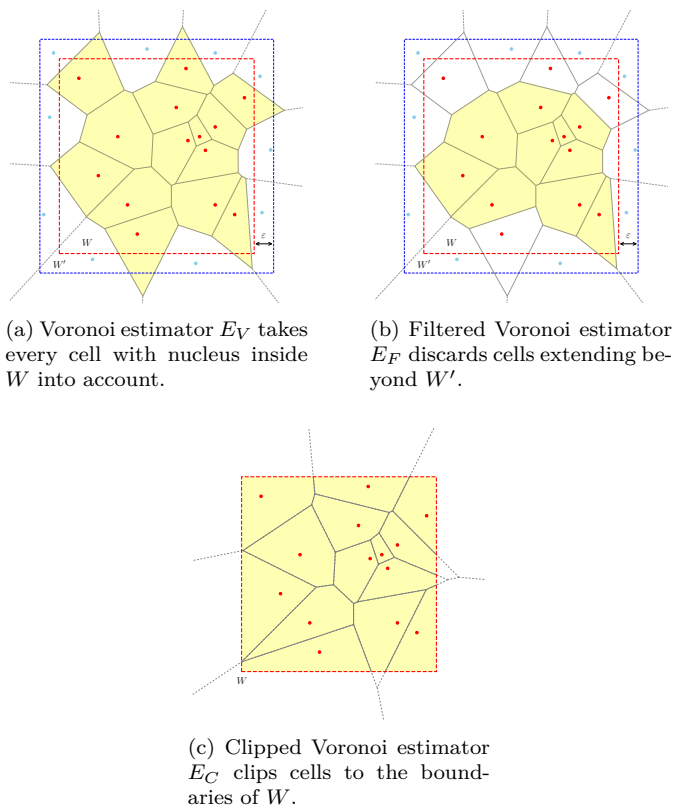


Figure 1: Three Voronoi integration estimators over $W = [-\frac{1}{2}, \frac{1}{2}]^2$. Voronoi cells nuclei inside W are depicted in red, and auxiliary nuclei in blue. Cells used by the estimators are highlighted in yellow.

4 Experiment

4.1 Experimental goals and hypothesis

The primary goal is to observe that the Voronoi weighting approach, which uses the area of Voronoi cells to weight sample contributions, provides a more efficient and stable estimate of an integral compared to the unweighted average used in Monte Carlo. This efficiency and stability are quantified by comparing the standard deviation of the estimates for different integration methods.

A list of functions is used to test the convergence of the methods (see table 1 for their definition).

Each of these functions exhibit some form of irregularity at the origin:

- `holder_α` is Hölder continuous with exponent $\alpha \in (0, 1]$; it is not differentiable at the origin if $\alpha < 1$;
- `not_holder` is continuous on W but not Hölder and therefore not differentiable (at the origin);
- `discontinuity` is undefined at the origin, thus not Hölder continuous on W , yet \mathcal{C}^∞ on $W \setminus \{(0, 0)\}$.

Table 1: Definition of the used functions

name	definition
<code>holder_α</code>	$ xy ^\alpha$
<code>not_holder</code>	$xy \sin\left(\frac{1}{x}\right) \sin\left(\frac{1}{y}\right)$
<code>discontinuity</code>	$(x^2 + y^2)^{-\frac{1}{2}}$

Table 2: Comparison of standard deviations and average execution durations (ms) over 10,000 runs at 4,096 samples per pixel (spp).

Fonction	Standard Deviation			Time		
	mc	vor	fvor	mc	vor	fvor
<code>holder_1</code>	0.00086	0.000171	0.000167	2.280	201.970	149.870
<code>holder_0.5</code>	0.001786	0.000291	0.000292	4.040	204.820	150.300
<code>holder_0.1</code>	0.001461	0.000299	0.000297	3.780	216.570	156.930
<code>holder_0.01</code>	0.000244	$< 10^{-6}$	$< 10^{-6}$	4.070	229.090	151.750
<code>not_holder</code>	0.000585	0.000179	0.000178	5.450	228.610	159.330
<code>discontinuity</code>	0.10716	0.079052	0.10411	2.530	231.970	152.400

4.2 Results

4.2.1 Experimental labeling

In the following experiments, the standard Monte Carlo estimation algorithm is denoted `mc`. Voronoi estimators E_V , E_F and E_C are denoted `vor`, `fvor` and `cvor` respectively.

4.2.2 Analysis

Figure 2 shows that all three estimators achieve significant variance reduction for function `not_c1`. The clipped estimator E_C , however, exhibits significant estimation bias. Consequently, being of limited practical interest, this estimator will not be systematically evaluated in subsequent experiments. Table 2 shows that E_V and E_F achieve significant variance reduction for the Hölder function cases considered. Though Proposition 1 does not hold for this case, a comparable variance reduction is also observed for the function `not_holder`. However, Voronoi estimators fail to outperform classical Monte Carlo for function `discontinuity`, which has a singularity at the origin.

During our experiments, we found that in practice, the SPPP sampling strategy leads to a much lower rejection rate than implied by the target confidence level.

In addition to generating the random sample distribution, standard Monte Carlo requires only repeated evaluations of the integrand. For explicitly defined functions, this is a computationally trivial task. In contrast, our method requires computing a Voronoi tessellation, a costly operation that dominates the total runtime. Table 2 and Figure 3 show that execution times disqualify Voronoi reweighting methods for the integration of explicitly defined functions. However, for functions where evaluation is costly, and specifically of an order of magnitude comparable to the Voronoi tessellation construction, the rapid convergence of our method might mitigate the computational overhead. For this reason, the next section is dedicated to applying this approach to the rendering equation.

5 Results with MC integration for Global illumination

5.1 The Rendering Equation and its Monte Carlo Estimation

Path tracing is a Monte Carlo estimator for the rendering equation: pixel values are expressed by an integral which is numerically solved by averaging random samples of light transport paths.

The rendering equation [4] is the fundamental mathematical model that describes how light interacts with surfaces in a scene. For any non trivial scene, this equation has no analytic solution. To estimate its solution, the Path Tracing algorithm relies on a stochastic simulation of the interaction of light with the

Dispersion of Voronoi and MC Estimators

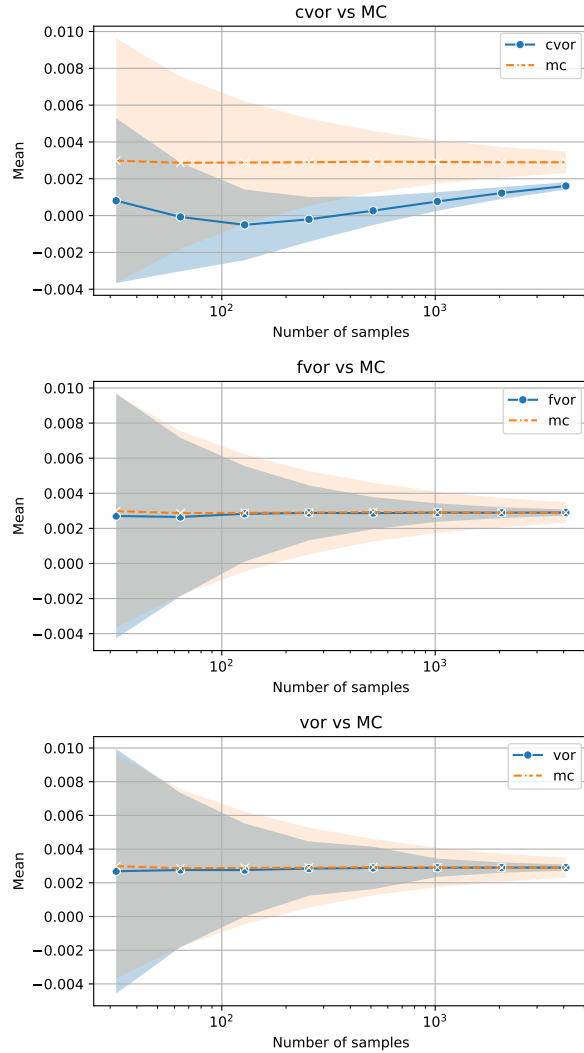


Figure 2: Performance of integrators for function `not_holder`. Each plot compares a reweighting integration method against standard Monte Carlo (labeled `mc`). Solid lines represent the average estimate over 10,000 runs, with sample count n ranging from 2^5 to 2^{14} . Shaded regions indicate the dispersion of these estimates within ± 1 standard deviation. The estimators E_V , E_C and E_F (labeled `vor`, `cvor` and `fvor` respectively) demonstrate significant variance reduction. As expected, the clipped Voronoi reweighting method `cvor` exhibits estimation bias, while still reducing the variance.

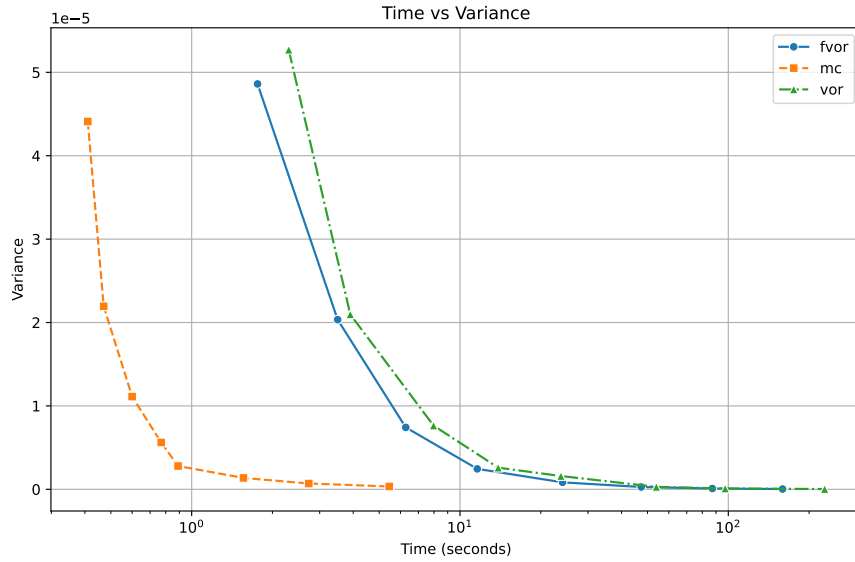


Figure 3: Comparative efficiency of three algorithms for evaluating the integral of the function `not_holder`.

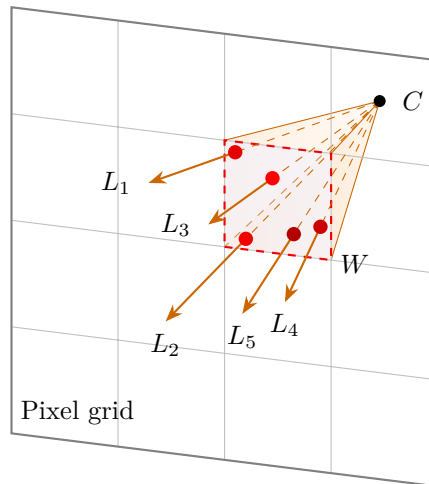


Figure 4: Simulating numerous potential light paths connecting a virtual camera C to light sources leads to the accumulation of light contributions (L_i) over a pixel area W , defining its final color. In a standard Monte Carlo framework, these contributions are averaged without accounting for the geometry of their distribution.

geometry and the materials of the scene. Being a stochastic algorithm, Path Tracing depends on variance reduction techniques for achieving an accurate, noise-free visual result. Figure 4 illustrates how the color of each individual pixel to display on the screen is computed: by simulating numerous light paths traversing it [7].

5.2 Experimental renderer

YAPT stands for Yet Another Path Tracer [14]. It is an experimental rendering engine developed to investigate and test novel sampling and integration techniques in the field of physically-based rendering. As a path tracer, YAPT simulates the way light interacts with objects in a virtual environment to produce realistic images (Figure 5).

The core architectural principle is the separation between the two fundamental processes of Monte Carlo integration:

Sampler responsible for generating the sample distribution used to estimate the light transport integral. This component determines the spatial distribution of light samples (e.g., uniform, stratified, or SPPP)

Integrator responsible for calculating the light transport (radiance) along the paths defined by the sampler. This component implements the specific rendering algorithm, such as classic Path Tracing or techniques incorporating explicit light sampling like Next Event Estimation (NEE), which is used in these experiments.

5.3 Convergence and Efficiency Analysis

The core purpose of YAPT is evaluating the performance and convergence of new integration methods compared to standard Monte Carlo techniques. This engine was used to evaluate the image rendering performance of estimators E_V and E_F , for which experimental results in section 4 demonstrate a variance reduction for Hölder functions integrated over a compact window. In the context of the rendering equation, such an assumption is in fact not very restrictive, since pixels are squares, hence bounded sets, and the radiance can be modeled by a sufficiently regular function. Experiments were performed using the classic Cornell Box scene, a standard benchmark in computer graphics. Figure 5 visually illustrates the convergence process inherent to Monte Carlo rendering. By progressively increasing the sample count per pixel, the approximated image approaches the true solution, resulting in a reduction of variance, perceived as visual noise. During our experiments, the following rendering settings were chosen : sample size ranging from 64 to 4,096, maximum path depth 8, and Next Event Estimation.





Figure 5: Illustration of Monte Carlo rendering convergence with YAPT. The estimate gradually approaches the exact scene radiance as the sample count increases, resulting in a progressively more accurate estimation of the scene’s true radiance.

5.4 Quantifying Estimator Performance

Figure 6 provides a quantitative analysis of different estimators by plotting the mean estimated radiance. Shaded regions represent the ± 1 standard deviation interval. These evaluations were performed for a specific, difficult-to-sample pixel located near a light source boundary. Accurately evaluating the color of such a pixel is indeed difficult. During the sampling process, contributions collected from direct lighting in the source surface will have high amplitude, whereas off-source contributions, which are indirectly lit, will be evaluated with significant variance, due to the diffuse nature of the material of the walls of the scene.

Figure 6 demonstrates that Voronoi integration methods achieve lower variance than standard Monte Carlo approach for identical sample sizes when target pixel is located at the boundary of a light source.

Figure 7 summarizes the key findings regarding the Voronoi estimators. The Voronoi estimators achieve superior performance by effectively reducing variance in regions exhibiting significant spatial radiance heterogeneity, specifically areas with steep radiance gradients. This improvement is particularly noticeable along light source boundaries, where the adaptive nature of the Voronoi weighting enables a more robust approximation of the radiance integral across sharp discontinuities like the edge of a projected light source. Additionally, this method also leads to faster convergence in areas involving specular light interactions on glossy materials. However, the benefits of the method are not uniform across the scene, especially where the radiance field is smooth or highly stochastic. Specifically, this method offers no substantial gain for surfaces with diffuse (Lambertian) reflection or in regions exhibiting low or weak spatial correlation. This limitation is inherent to adaptive sampling techniques, which offer limited advantage over uniform sampling when the underlying light function is smooth or lacks structure [16]. The geometry of the Voronoi tessellation for a direct



(a) Voronoi estimation versus uniform Monte Carlo (b) Filtered Voronoi estimation vs Monte Carlo

Figure 6: Evolution of the mean radiance estimate and ± 1 standard deviation interval as a function of the sample size. Results obtained with the Voronoi estimator E_V `spp/vor` and the filtered Voronoi estimator E_F `spp/fvor` are compared to those from random sampling using Monte Carlo estimation `rnd/mc` in each graph. For this pixel, located along a source boundary, both Voronoi estimators exhibit a faster convergence rate than standard MC estimation.

light, a specular reflection or a diffuse wall is shown in Figure 8.


Finally, the Voronoi-weighted integration method achieves significant reduction in variance in regions exhibiting steep radiance gradients, at the cost of some computational overhead. This overhead (see figure 9a) stems from constructing the Voronoi tessellation and computing the corresponding cell areas. Figure 9b, however, shows that this overhead is largely offset by the variance reduction achieved by the methods E_V and E_F , making them more efficient than a standard Monte Carlo approach.

A more global performance indicator demonstrate an overall improvement of the numerical precision of the images computed by Voronoi estimators (Figure 7.) For a Cornell Box scene, Standard Monte Carlo estimation achieves an MSE value of 0.030, while a Voronoi estimation achieves a lower error (MSE=0.015).


6 Discussion perspectives

The integration techniques presented in this paper converge faster than a standard Monte Carlo approach for Hölder continuous functions. In practice, these methods prove more efficient than standard Monte Carlo when the integrand is computationally expensive.


Proposition 1 is a new result regarding Poisson-Voronoi tessellations. In future work, this result could be extended to point processes that are not necessarily Poissonian and that include determinantal processes, in particular the Ginibre process. Using the notations of Proposition 1, a natural question that




(a) regular Monte Carlo with sample size 256



(b) difference map with reference image (MSE = 0.030)



(c) Voronoi ponderation vor with sample size 256



(d) difference map with reference image (MSE = 0.015)

Figure 7: This figure compares two integration methods in a Path Tracer for a fixed number of samples per pixel of $n = 256$. Top row: (a) Rendering of the Cornell Box scene using random Monte Carlo (MC) integration for $n = 256$ samples per pixel (spp). (b) Error map (difference) between the MC rendering (a) and a high-quality reference (reference calculated with $n = 4096$ spp and stratified sampling). Bottom row: (c) Rendering of the same scene using Voronoi-weighted integration for $n = 256$ spp. (d) Map of absolute errors (difference) between the Voronoi rendering (c) and the same reference image. A comparison of error maps (b) and (d) demonstrates that Voronoi-weighted integration **vor** (c) yields lower variance than standard MC integration at equal sample counts, most notably at radiance discontinuities, such as shadows and light source edges.

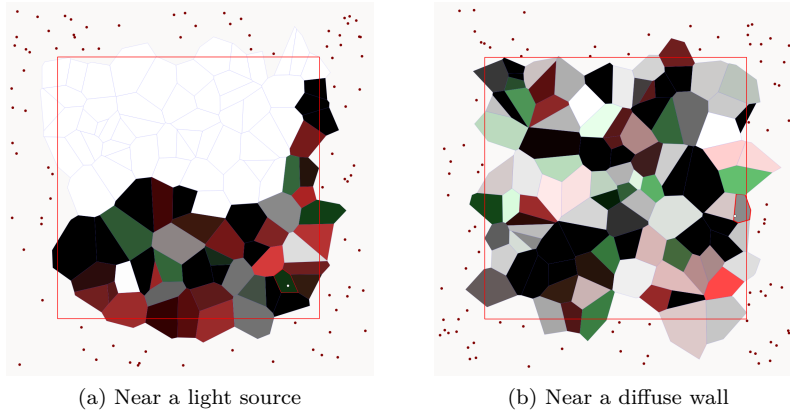




Figure 8: Pixels rendered with the E_V estimators in a Cornell Box scene. Voronoi cells are colored according to the radiance of the sample computed by path tracing. Pixel boundaries are outlined in red. Voronoi weighting methods prove highly effective in scenarios where the radiance field exhibits significant spatial heterogeneity. They exhibit superior performance at major discontinuities, such as light source edges (figure a). However, they offer negligible improvement for diffuse materials (figure b), a limitation they share with adaptive sampling techniques.



(a) CPU time as a function of sample size



(b) variance as a function of CPU time

Figure 9: Estimators performance for a pixel located at the boundary of a light source. (a) Voronoi-based estimations lead to a nearly identical increase in the global computation time for both `vor` and `fvor` methods. (b) However, given a fixed CPU time budget, both methods prove to be more efficient.

arises is whether $\sum_{x \in \eta_n} f(x)v_{\eta_n}(x)$ satisfies a central limit theorem, i.e. if

$$\frac{\sum_{x \in \eta_n} f(x)v_{\eta_n}(x) - \mathbb{E}[\sum_{x \in \eta_n} f(x)v_{\eta_n}(x)]}{\sqrt{\mathbb{V}[\sum_{x \in \eta_n} f(x)v_{\eta_n}(x)]}} \xrightarrow[n \rightarrow \infty]{} \mathcal{N}(0, 1)$$

in distribution.

In the context of image rendering, our experiments highlighted that these techniques achieve rapid variance reduction in regions where the radiance field exhibits high-frequency structures. We demonstrated that a rejection scheme based on the dilation of the Voronoi seed space yields well-formed tessellations for this type of evaluation. Rather than relying on rejection, it would be of interest to investigate an adaptive approach; such a method would not reject an invalid tessellation but instead attempt to refine it by further expanding the integration window.

We restricted our analysis to window-like domains in dimension 2. Applying similar techniques beyond the primary ray would be a natural extension. Specifically, constructing a Voronoi tessellation within the directional hemisphere at each scene bounce may be worth considering. Furthermore, exploring Voronoi tessellations within the higher-dimensional space of paths represents another promising research perspective.

Contributions All authors contributed to the study conception and design. Material preparation, data collection and analysis were performed by FV and SD. The first draft of the manuscript was written by FV and all authors commented on previous versions of the manuscript. All authors read and approved the final manuscript.

References

- [1] P. Calka. The distributions of the smallest disks containing the Poisson-Voronoi typical cell and the Crofton cell in the plane. *Adv. in Appl. Probab.*, 34(4):702–717, 2002.
- [2] J. Guo and E. Eisemann. Geometric sample reweighting for monte carlo integration. *Computer Graphics Forum*, 40(7):109–119, October 2021.
- [3] M. Heveling and M. Reitzner. Poisson-Voronoi approximation. *Ann. Appl. Probab.*, 19(2):719–736, 2009.
- [4] James T. Kajiya. The rendering equation. *ACM SIGGRAPH Computer Graphics*, 20(4):143–150, August 1986.
- [5] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Wiley Series in Probability and Statistics. John Wiley & Sons Ltd., Chichester, second edition, 2000.

- [6] Gilles Pagès. A space quantization method for numerical integration. *Journal of Computational and Applied Mathematics*, 89(1):1–38, 1998.
- [7] Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically based rendering: From theory to implementation*. MIT Press, Cambridge, MA, Mar 2023.
- [8] M. Reitzner, E. Spodarev, and D. Zaporozhets. Set reconstruction by Voronoi cells. *Adv. in Appl. Probab.*, 44:938–953, 2012.
- [9] M. Reitzner and A. Strotmann. Poisson-delaunay approximation. Available at <https://arxiv.org/pdf/2410.23003>., 2024.
- [10] Christian Robert and George Casella. *Monte Carlo Statistical Methods*. Springer texts in statistics. Springer, New York, NY, November 2010.
- [11] R. Schneider and W. Weil. *Stochastic and Integral Geometry*. Probability and its Applications (New York). Springer-Verlag, Berlin, 2008.
- [12] M. Schulte. A central limit theorem for the Poisson-Voronoi approximation. *Adv. in Appl. Math.*, 49(3-5):285–306, 2012.
- [13] C. Thäle and J. E. Yukich. Asymptotic theory for statistics of the Poisson-Voronoi approximation. *Bernoulli*, 22(4):2372–2400, 2016.
- [14] Franck Vandewiele and Samuel Delepouille. prise-3d/yapt: 1.0.3, 2025.
- [15] Miroslav Vořechovský, Václav Sadílek, and Jan Eliáš. Voronoi weighting of samples in monte carlo integration. In *Proceedings of the 2nd International Conference on Uncertainty Quantification in Computational Sciences and Engineering (UNCECOMP 2017)*, pages 478–491, Rhodes Island, Greece, 2017. Institute of Structural Analysis and Antiseismic Research School of Civil Engineering National Technical University of Athens (NTUA) Greece.
- [16] M. Zwicker, W. Jarosz, J. Lehtinen, B. Moon, R. Ramamoorthi, F. Rouselle, P. Sen, C. Soler, and S.-E. Yoon. Recent advances in adaptive sampling and reconstruction for monte carlo rendering. *Comput. Graph. Forum*, 34(2):667–681, May 2015.

A Proof of proposition 1

Proof. Assertion (i) directly follows from the Slivnyak-Mecke formula (see e.g. Corollary 3.2.3 in [11]). Indeed, the latter implies

$$\begin{aligned} \mathbb{E} \left[\sum_{x \in \eta_n} f(x) v_{\eta_n}(x) \right] &= n \int_{\mathbf{R}^d} f(x) \mathbb{E} [v_{\eta_n \cup \{x\}}(x)] \, dx \\ &= \int_{\mathbf{R}^d} f(x) \, dx \end{aligned}$$

since $\mathbb{E} [v_{\eta_n \cup \{x\}}(x)] = \mathbb{E} [v_{\eta_n \cup \{0\}}(0)] = n^{-1}$.

To deal with (ii), we apply the Poincaré inequality (see e.g. Section 2 in [8]), which claims that

$$\mathbb{V} [F(\eta_n)] \leq n \mathbb{E} \left[\int_{\mathbf{R}^d} (F(\eta_n \cup \{x\}) - F(\eta_n))^2 dx \right]$$

for all measurable function $F : \mathbf{N} \rightarrow \mathbf{R}_+$ such that $\mathbb{E} [F^2(\eta_n)] < \infty$, where \mathbf{N} denotes the set of all locally finite counting measures in \mathbf{R}^d . This implies

$$\mathbb{V} \left[\sum_{x \in \eta_n} f(x) v_{\eta_n}(x) \right] \leq n \mathbb{E} \left[\int_{\mathbf{R}^d} \left(\sum_{y \in \eta_n \cup \{x\}} f(y) v_{\eta_n \cup \{x\}}(y) - \sum_{y \in \eta_n} f(y) v_{\eta_n}(y) \right)^2 dx \right]. \quad (1)$$

Now, let x be fixed and let $y \in \eta_n$. Denote by $\mathcal{F}_{\eta_n}(y)$ the Voronoi flower associated with y , i.e.

$$\mathcal{F}_{\eta_n}(y) = \bigcup_{z \in C_{\eta_n}(y)} B(z, |z - y|),$$

where $B(z, r)$ is the (closed) ball centered at z with radius r . Notice that, if $x \notin \mathcal{F}_{\eta_n}(y)$, then $v_{\eta_n \cup \{x\}}(y) = v_{\eta_n}(y)$. Therefore

$$\begin{aligned} & \sum_{y \in \eta_n \cup \{x\}} f(y) v_{\eta_n \cup \{x\}}(y) - \sum_{y \in \eta_n} f(y) v_{\eta_n}(y) \\ &= \sum_{y \in \eta_n} f(y) (v_{\eta_n \cup \{x\}}(y) - v_{\eta_n}(y)) + f(x) v_{\eta_n \cup \{x\}}(x) \\ &= \sum_{y \in \eta_n} f(y) (v_{\eta_n \cup \{x\}}(y) - v_{\eta_n}(y)) \mathbb{1}_{x \in \mathcal{F}_{\eta_n}(y)} + f(x) v_{\eta_n \cup \{x\}}(x) \\ &= \sum_{y \in \mathcal{N}_{\eta_n \cup \{x\}}(x)} f(y) (v_{\eta_n \cup \{x\}}(y) - v_{\eta_n}(y)) + f(x) v_{\eta_n \cup \{x\}}(x), \end{aligned}$$

where the last equality comes from the fact that $x \in \mathcal{F}_{\eta_n}(y)$ iff $x \in \mathcal{N}_{\eta_n \cup \{x\}}(y)$, i.e. iff $y \in \mathcal{N}_{\eta_n \cup \{x\}}(x)$, and where $\mathcal{N}_{\eta_n \cup \{x\}}(x)$ denotes the set of neighbors of x in the Voronoi tessellation associated with $\eta_n \cup \{x\}$. Furthermore, adapting an argument of Section 4 in [3], we have

$$v_{\eta_n \cup \{x\}}(x) = \sum_{y \in \mathcal{N}_{\eta_n \cup \{x\}}(x)} (v_{\eta_n}(y) - v_{\eta_n \cup \{x\}}(y)).$$

Thus

$$\sum_{y \in \eta_n \cup \{x\}} f(y) v_{\eta_n \cup \{x\}}(y) - \sum_{y \in \eta_n} f(y) v_{\eta_n}(y) = \sum_{y \in \mathcal{N}_{\eta_n \cup \{x\}}(x)} (f(y) - f(x)) (v_{\eta_n \cup \{x\}}(y) - v_{\eta_n}(y)).$$

This together with (1) gives

$$\begin{aligned} \mathbb{V} \left[\sum_{x \in \eta_n} f(x) v_{\eta_n}(x) \right] &\leq n \int_{\mathbf{R}^d} \mathbb{E} \left[\left(\sum_{y \in \mathcal{N}_{\eta_n \cup \{x\}}(x)} |f(y) - f(x)| |v_{\eta_n}(y) - v_{\eta_n \cup \{x\}}(y)| \right)^2 \right] \\ &=: r_1(n) + r_2(n) \end{aligned} \quad (2)$$

where

$$r_1(n) := n \int_{\mathbf{R}^d} \mathbb{E} \left[\sum_{y \in \mathcal{N}_{\eta_n \cup \{x\}}(x)} |f(y) - f(x)|^2 |v_{\eta_n}(y) - v_{\eta_n \cup \{x\}}(y)|^2 \right] dx$$

and

$$\begin{aligned} r_2(n) &:= n \int_{\mathbf{R}^d} \mathbb{E} \left[\sum_{(y_1, y_2) \neq \in \mathcal{N}_{\eta_n \cup \{x\}}(x)^2} |f(y_1) - f(x)| |f(y_2) - f(x)| \right. \\ &\quad \left. \times |v_{\eta_n}(y_1) - v_{\eta_n \cup \{x\}}(y_1)| |v_{\eta_n}(y_2) - v_{\eta_n \cup \{x\}}(y_2)| \right] dx. \end{aligned}$$

We provide below a bound for $r_1(n)$. According to the Slivnyak-Mecke formula, we have

$$r_1(n) = n^2 \int_{(\mathbf{R}^d)^2} |f(y) - f(x)|^2 \mathbb{E} \left[|v_{\eta_n \cup \{y\}}(y) - v_{\eta_n \cup \{x, y\}}(y)|^2 \mathbb{1}_{y \in \mathcal{N}_{\eta_n \cup \{x, y\}}(x)} \right] dy dx$$

Since $v_{\eta_n \cup \{x, y\}}(y) \leq v_{\eta_n}(y)$, we get

$$\begin{aligned} r_1(n) &\leq 4n^2 \int_{(\mathbf{R}^d)^2} |f(y) - f(x)|^2 \mathbb{E} \left[v_{\eta_n \cup \{y\}}^2(y) \mathbb{1}_{y \in \mathcal{N}_{\eta_n \cup \{x, y\}}(x)} \right] dy dx \\ &= 4 \int_{(\mathbf{R}^d)^2} |f(y) - f(x)|^2 \mathbb{E} \left[v_{\eta \cup \{n^{1/d}y\}}^2(n^{1/d}y) \mathbb{1}_{n^{1/d}y \in \mathcal{N}_{\eta \cup \{n^{1/d}x, n^{1/d}y\}}(n^{1/d}x)} \right] dy dx, \end{aligned}$$

where the last equality comes from the fact that $\eta_n \stackrel{\mathcal{D}}{=} n^{-1/d} \eta$, where $\eta := \eta_1$ is a PPP with intensity 1. Taking the change of variables $x' = n^{1/d}x$ and $y' = n^{1/d}y$ and applying the Hölder's inequality, we obtain

$$r_1(n) \leq 4n^{-2} \int_{(\mathbf{R}^d)^2} |f(n^{-1/d}y') - f(n^{-1/d}x')|^2 \mathbb{E} \left[v_{\eta \cup \{y'\}}^2(y') \mathbb{1}_{y' \in \mathcal{N}_{\eta \cup \{x', y'\}}(x')} \right] dy' dx',$$

and therefore

$$\begin{aligned} r_1(n) &\leq 4n^{-2} \int_{(\mathbf{R}^d)^2} |f(n^{-1/d}y') - f(n^{-1/d}x')|^2 \mathbb{E} \left[v_{\eta \cup \{y'\}}^4(y') \right]^{1/2} \\ &\quad \times \mathbb{P} \left(y' \in \mathcal{N}_{\eta \cup \{x', y'\}}(x') \right)^{1/2} dy' dx' \end{aligned}$$

Since η is stationary, we have $\mathbb{E} \left[v_{\eta \cup \{y'\}}^4(y') \right] = \mathbb{E} [v^4(\mathcal{C})] < \infty$ for any $y' \in \mathbf{R}^d$, where $\mathcal{C} \stackrel{\mathcal{D}}{=} C_{\eta \cup \{0\}}(0)$ and $v(\mathcal{C})$ denote the typical cell associated with η and the volume of \mathcal{C} , respectively. Writing $W_n := n^{1/d}W$, we get

$$\begin{aligned} r_1(n) &\leq cn^{-2} \int_{(\mathbf{R}^d)^2} |f(n^{-1/d}y') - f(n^{-1/d}x')|^2 \mathbb{P} (y' \in \mathcal{N}_{\eta \cup \{x', y'\}}(x'))^{1/2} dy' dx' \\ &= 2cn^{-2} \int_{W_n} \left(\int_{\mathbf{R}^d} |f(n^{-1/d}y') - f(n^{-1/d}x')|^2 \mathbb{P} (y' \in \mathcal{N}_{\eta \cup \{x', y'\}}(x'))^{1/2} dy' \right) dx', \end{aligned}$$

where, in the last line, we used the fact that $f(n^{-1/d}y') = f(n^{-1/d}x')$ if $(x', y') \in (W_n^c)^2$ since W is the support of f . The factor 2 comes from the fact that the integrand is symmetric w.r.t. x' and y' (since $y' \in \mathcal{N}_{\eta \cup \{x', y'\}}(x')$ iff $x' \in \mathcal{N}_{\eta \cup \{x', y'\}}(y')$). Furthermore, for any locally finite subset $\chi \subset \mathbf{R}^d$ in general position and for any $x \in \chi$, let $R_\chi(x)$ be the circumscribed radius associated with $C_\chi(x)$, i.e.

$$R_\chi(x) = \inf\{r > 0 : B(x, r) \supset C_\chi(x)\}.$$

Since $\mathcal{N}_\chi(x) \subset B(x, 2R_\chi(x))$, we get

$$\begin{aligned} r_1(n) &\leq cn^{-2} \int_{W_n} \left(\int_{\mathbf{R}^d} |f(n^{-1/d}y') - f(n^{-1/d}x')|^2 \mathbb{P} (y' \in B(x', 2R_{\eta \cup \{x', y'\}}(x')))^{1/2} dy' \right) dx' \\ &\leq cn^{-2-\frac{2\alpha}{d}} \int_{W_n} \left(\int_{\mathbf{R}^d} |y' - x'|^{2\alpha} \mathbb{P} (y' \in B(x', 2R_{\eta \cup \{x'\}}(x')))^{1/2} dy' \right) dx', \end{aligned}$$

where, in the last line, we used the facts that f satisfies a Hölder condition with exponent $\alpha \in (0, 1]$ and that $R_{\eta \cup \{x', y'\}}(x') \leq R_{\eta \cup \{x'\}}(x')$. Thanks to the stationarity of η , for any $x' \in \mathbf{R}^d$, we have $R_{\eta \cup \{x'\}}(x') \stackrel{\mathcal{D}}{=} R(\mathcal{C})$, where $R(\mathcal{C})$ denotes the circumscribed radius of the typical cell. Therefore

$$\int_{\mathbf{R}^d} |y' - x'|^{2\alpha} \mathbb{P} (y' \in B(x', 2R_{\eta \cup \{x'\}}(x')))^{1/2} dy' = \int_{\mathbf{R}^d} |z|^{2\alpha} \mathbb{P} \left(R(\mathcal{C}) \geq \frac{|z|}{2} \right)^{1/2} dz.$$

It is known that the tail of $R(\mathcal{C})$ decreases exponentially fast to 0, see e.g. [1], which proves that the last term is finite. Integrating over $x' \in W_n$, we get

$$r_1(n) \leq cn^{-1-\frac{2\alpha}{d}}.$$

We provide below an upper bound for $r_2(n)$. This will be sketched because we will proceed in the same spirit as we did for $r_1(n)$. First, according to the Slivnyak-Mecke formula, we write

$$\begin{aligned} r_2(n) &= n^3 \int_{(\mathbf{R}^d)^3} |f(y_1) - f(x)| |f(y_2) - f(x)| \mathbb{E} \left[\left| v_{\eta_n \cup \{x, y_1, y_2\}}(y_1) - v_{\eta_n \cup \{y_1, y_2\}}(y_1) \right| \right. \\ &\quad \left. \times \left| v_{\eta_n \cup \{x, y_1, y_2\}}(y_2) - v_{\eta_n \cup \{y_1, y_2\}}(y_2) \right| \mathbb{1}_{(y_1, y_2) \in \mathcal{N}_{\eta_n \cup \{x, y_1, y_2\}}^2(x)} \right] dy_2 dy_1 dx. \end{aligned}$$

Taking the change of variables $x' = n^{1/d}x$, $y'_1 = n^{1/d}y_1$, $y'_2 = n^{1/d}y_2$ and using the facts that $\eta_n \stackrel{D}{=} n^{-1/d}\eta$, $v_{\eta \cup \{x', y'_1, y'_2\}}(y'_i) \leq v_{\eta \cup \{y'_1, y'_2\}}(y'_i)$, with $i \in \{1, 2\}$, we get

$$\begin{aligned} r_2(n) &\leq 4cn^3 n^{-2} n^{-3} \int_{W_n} \left(\int_{(\mathbf{R}^d)^2} |f(n^{-1/d}y'_1) - f(n^{-1/d}x')| |f(n^{-1/d}y'_2) - f(n^{-1/d}x')| \right. \\ &\quad \left. \times \mathbb{E} \left[v_{\eta \cup \{y'_1, y'_2\}}(y'_1) v_{\eta \cup \{y'_1, y'_2\}}(y'_2) \mathbb{1}_{(y'_1, y'_2) \in \mathcal{N}_{\eta \cup \{x', y'_1, y'_2\}}^2(x')} \right] dy'_2 dy'_1 \right) dx'. \end{aligned}$$

Notice that, without restriction, we have integrated over x by using symmetry arguments and the fact that f has compact support W . Furthermore, according to the Hölder's inequality, we have for any $(x, y'_1, y'_2) \in W_n \times \mathbf{R}^d \times \mathbf{R}^d$,

$$\begin{aligned} &\mathbb{E} \left[v_{\eta \cup \{y'_1, y'_2\}}(y'_1) v_{\eta \cup \{y'_1, y'_2\}}(y'_2) \mathbb{1}_{(y'_1, y'_2) \in \mathcal{N}_{\eta \cup \{x', y'_1, y'_2\}}^2(x')} \right] \\ &\leq \left(\mathbb{E} \left[v_{\eta \cup \{y'_1, y'_2\}}^4(y'_1) \right] \mathbb{E} \left[v_{\eta \cup \{y'_1, y'_2\}}^4(y'_2) \right] \mathbb{P} \left((y'_1, y'_2) \in \mathcal{N}_{\eta \cup \{x', y'_1, y'_2\}}^2(x') \right) \right)^{1/4} \\ &\leq c \left(\mathbb{P} \left((y'_1, y'_2) \in \mathcal{N}_{\eta \cup \{x', y'_1, y'_2\}}^2(x') \right) \right)^{1/4}, \end{aligned}$$

where, in the last line, we used the facts that $v_{\eta \cup \{y'_1, y'_2\}}(y'_i) \leq v_{\eta \cup \{y'_i\}}(y'_i) \stackrel{D}{=} v(\mathcal{C})$ and that $\mathbb{E} [v^4(\mathcal{C})] < \infty$. This, together with the facts that f is Hölder and $\mathcal{N}_{\eta \cup \{x', y'_1, y'_2\}}^2(x') \subset B(x', 2R_{\eta \cup \{x'\}}(x'))$ implies

$$\begin{aligned} r_2(n) &\leq cn^{-2-\frac{2\alpha}{d}} \int_{W_n} \left(\int_{(\mathbf{R}^d)^2} |x' - y'_1|^\alpha |x' - y'_2|^\alpha \right. \\ &\quad \left. \times \mathbb{P} \left(y'_1 \in B(x', 2R_{\eta \cup \{x'\}}(x')) \right)^{1/4} \mathbb{P} \left(y'_2 \in B(x', 2R_{\eta \cup \{x'\}}(x')) \right)^{1/4} dy'_1 dy'_2 \right) dx'. \end{aligned}$$

Following the same approach as for $r_1(n)$, we get

$$\begin{aligned} r_2(n) &\leq cn^{-2-\frac{2\alpha}{d}} \int_{W_n} \left(\int_{\mathbf{R}^d} |z|^\alpha \mathbb{P} \left(R(\mathcal{C}) \geq \frac{|z|}{2} \right)^{1/4} \right)^2 dx' \\ &\leq cn^{-1-\frac{2\alpha}{d}}. \end{aligned}$$

This concludes the proof of Proposition 1. \square