

Variance Analysis for Monte Carlo Integration

†Adrien Pilleboue^{1,*} †Gurprit Singh^{1,*} David Coeurjolly^{2,*} Michael Kazhdan^{3,*} Victor Ostromoukhov^{1,2,*}
†joint first authors ¹Université Lyon 1 ²CNRS/LIRIS UMR 5205 ³Johns Hopkins University

Abstract

We propose a new spectral analysis of the variance in Monte Carlo integration, expressed in terms of the power spectra of the sampling pattern and the integrand involved. We build our framework in the Euclidean space using Fourier tools and on the sphere using spherical harmonics. We further provide a theoretical background that explains how our spherical framework can be extended to the hemispherical domain. We use our framework to estimate the variance convergence rate of different state-of-the-art sampling patterns in both the Euclidean and spherical domains, as the number of samples increases. Furthermore, we formulate design principles for constructing sampling methods that can be tailored according to available resources. We validate our theoretical framework by performing numerical integration over several integrands sampled using different sampling patterns.

CR Categories: I.3.3 [Computer Graphics]: Three-Dimensional Graphics and Realism—Display Algorithms I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Monte Carlo Integration;

Keywords: Stochastic Sampling, Monte Carlo Integration, Fourier Analysis, Spherical Harmonics, Global Illumination

1 Introduction

Numerical integration schemes such as Monte Carlo methods are widely used in high quality production rendering (*e.g.*, to estimate per pixel radiance). This estimation is error-prone, as Monte Carlo methods provide an approximation of the true integral. Specifically, the quality of integration depends strongly on the sampling pattern. For example, regular structures in sampling can result in strong aliasing (structured artifacts) in the final rendered image. However, by using appropriately placed stochastically-generated samples, this aliasing in rendered images can be transformed into less objectionable noise [Crow 1977; Dippé and Wold 1985; Cook 1986; Mitchell 1987; Shirley 1991], which is tightly related to variance in Monte Carlo integration.

In this work, we study variance from the sampling perspective and propose a mathematical framework that predicts variance in Monte Carlo integration. Our framework shows that the variance in integration is directly related to the power spectra of the sampling pattern and the integrand under study. We build our mathematical framework in the Euclidean space and later extend it to the

(hemi-)spherical domain, which is crucial in light transport estimation. To our knowledge, this is the first theoretical and computational analysis of variance in Monte Carlo integration on the sphere and the hemisphere. To study variance in integration we perform spectral analysis of various sampling patterns in both the Euclidean and the (hemi-)spherical domains, using Fourier analysis for Euclidean space and spherical harmonics for the (hemi-)sphere.

1.1 Contributions

Our primary contribution is a mathematical framework that establishes a direct relation between the variance in Monte Carlo integration in the Euclidean and (hemi-)spherical domains with the power spectra of the sampling pattern and the integrand. We use our mathematical framework to study:

Spectral analysis of sampling patterns We perform spectral analysis of various state-of-the-art sampling patterns on Euclidean and spherical domains. We provide theoretical insights to perform similar analysis on hemispherical domain using spherical harmonics.

Monte Carlo convergence We extend our theoretical framework to analyze the best- and worst-case variance convergence rates—for a given class of functions—in Monte Carlo integration in the Euclidean and the (hemi-)spherical domains.

The remainder of this paper is split into two main parts. In the first part, we study the relation between the error in integration and the frequency content of the integrand and the sampling pattern. We define the notion of *homogeneous sampling* (Sec. 3.1) that allows expressing error only in terms of variance in MC integration. By restricting our analysis to homogeneous sampling, we derive the closed-form expression for variance in MC integration in terms of power spectra of the integrand and the sampling pattern. We show that this relation holds in the Euclidean (Sec. 4.2) and the spherical (Sec. 5.2) domains. All mathematical notations and symbols used in our mathematical formulation can be found in Table 1. In the second part, we directly use our variance formulation to study the MC integration convergence rates of various theoretical sampling patterns with well-defined spectral profiles (Sec. 7). We then use these theoretical convergence rates to build a taxonomy to classify existing sampling methods. We show empirical results (Sec. 8) to support our theoretical study and present future directions (Sec. 9) for further exploration.

2 Related work

Sampling in Euclidean domain To improve the quality of image synthesis, researchers have studied various Monte Carlo [Crow 1977; Dippé and Wold 1985; Cook 1986; Mitchell 1987] and Quasi-Monte Carlo based [Niederreiter 1992; Lemieux 2009; Keller et al. 2012] sampling patterns. Many algorithms have been proposed to improve the quality of sampling by studying their spectral properties. Ulichney [1987] was the first to provide qualitative characterization of a good sampling pattern, which is now commonly called *Blue Noise*. Mitchell [1991] has also pointed out that energy in low-frequency part of the Fourier spectrum of the sampling pattern should be avoided. Since then, different algorithms [Balzer et al. 2009; Schlömer et al. 2011; de Goes et al. 2012] and tile-based

* firstname.lastname@liris.cnrs.fr

* misha@cs.jhu.edu

Symbol	Definition
\mathbb{D}, N	sampling domain (toroidal unit square \mathcal{T}^d , unit sphere S^2 , unit hemisphere \mathcal{H}^2), with number of samples N .
$\mu(\mathbb{D})$	Lebesgue measure of domain \mathbb{D}
θ, ϕ	colatitude $\theta \in [0, \pi]$ and longitude $\phi \in [0, 2\pi]$, on the sphere or the hemisphere
l, m	degree and order of spherical harmonics, both values are integers with $l \geq 0$ and $m \in [-l, l]$.
$Y_l^m(\theta, \phi)$	Spherical harmonic function of degree l and order m whose inner product over the sphere is 4π
$P_l^m(\cos \theta)$	associated Legendre polynomial of degree l and order m .
Θ	the Fourier domain without the DC peak.
Φ	the Fourier domain with the DC peak.
$G(x)$	arbitrary integrable function on \mathbb{D} .
$\mathcal{P}_G(\cdot)$	power spectrum of G in the toroidal domain or angular power spectrum of G in the spherical domain.
$\check{\mathcal{P}}_G(\cdot)$	radial mean of the expected value of \mathcal{P}_G in the toroidal domain or expected value of \mathcal{P}_G in the spherical domain.
$\bar{g}, \ g\ $	conjugate and magnitude operators for a complex variable g , respectively.
$\langle \cdot, \cdot \rangle$	inner product operator
$\langle \cdot \rangle, \text{Var}(\cdot)$	expectation and variance operators
E-, S-	equation numbering notation used to refer to the Euclidean (E) and spherical (S) domains, respectively.
X	bold face letter represents a random variable.

Table 1: Notations and mathematical symbols used in this paper. Note that, G and g are dummy variables which get replaced by symbols applicable in respective domains based on the context in the paper.

methods [Cohen et al. 2003; Kopf et al. 2006; Ostromoukhov 2007] have been proposed to achieve the blue noise quality of stochastic sampling patterns. Recently, researchers [Zhou et al. 2012; Heck et al. 2013; Öztireli and Gross 2012] have proposed methods for generation of distributions with the desired power spectra in order to reduce the variance of the Monte Carlo integration. Wachtel and colleagues [2014] presented a first tile-based method that allows advanced spectral control. However, these works do not provide any closed-form relation between the variance and the sampling power spectra that could leverage spectral control in variance reduction. In this work, we derive an explicit expression relating the variance in MC integration with the power spectra of the sampling pattern and the integrand, showing how variance can be controlled or predicted directly from the power spectra.

Sampling in spherical domain Sphere sampling has been well studied in various domains [Górski et al. 2005; McEwen and Wiaux 2011; Choirat and Seri 2013; Cui and Freedman 1997]. In computer graphics, spherical sampling has been actively studied for rendering purposes [Arvo 1995; Arvo 2001]. Ureña and colleagues [2013] proposed an area preserving spherical rectangle parametrization to reduce variance in scattering computation from planar rectangular emitters. Marques and colleagues [2013] introduced spherical Fibonacci lattices to improve QMC sampling in the spherical domain. Li and colleagues [2010] performed sphere sampling using planar warpings while Bowers and colleagues [2010] directly generate samples on manifolds and studied their spectral properties. A relaxation based method was proposed by Xu and colleagues [2012] to generate blue noise samples directly on the surface of the sphere. Our focus in this work is to analyze existing stochastic sampling patterns and their spectral properties in the spherical domain and to better understand their consequences in light transport simulation.

Variance analysis of integration Recent papers have conducted a comprehensive analysis of variance in Monte Carlo integration. In particular, Durand [2011] investigated error in integration and related error to the spectral properties of the associated sampling patterns. Ramamoorthi and colleagues [2012] focused specifically on visibility to assess error due to sampling patterns in soft shadow rendering. Subr and Kautz [2013] proposed a mathematical formulation that relates the variance in Monte Carlo integration directly to the variance of sampling Fourier coefficients taken over multiple

realizations. In this work, we go one step further, and relate the variance in MC integration directly to the power spectrum of the sampling pattern. Using this relation, we can predict the variance of integration and tailor new sampling patterns according to an individual application’s requirements. We further extend our analysis to the (hemi-)spherical domain.

Convergence analysis Convergence rate of various sampling methods in numerical integration has been extensively studied on the unit sphere [Hesse et al. 2010]. A rich literature on QMC integration [Brauchart et al. 2014], can also be found for the worst-case error in the spherical domain. In image synthesis, Marques and colleagues [2013] studied the worst-case error for spherical Fibonacci point sets for illumination integrals. Ramamoorthi and colleagues [2012] analytically studied convergence rate of jittered sampling, while Mitchell [1996] obtained similar convergence rates empirically. Recently, Subr and colleagues [2014] studied convergence of various importance based samplers. In this work, we perform our variance convergence analysis on state-of-the-art stochastic (non-adaptive) samplers in both the Euclidean and spherical domains.

Spherical harmonics in graphics Spherical harmonics (SH) has been widely used in computer graphics, for inverse rendering [Ramamoorthi and Hanrahan 2001], precomputed radiance transfer [Sloan et al. 2002], and shape analysis [Kazhdan 2007]. SH are also used for importance sampling of spherical functions [Jarosz et al. 2009] and for performing arbitrary BRDF shading for low-frequency lighting [Kautz et al. 2002]. In our work, we use spherical harmonics to derive our theoretical formulations and to study the spectral properties of different sampling patterns on the sphere and the hemisphere.

3 Monte Carlo estimator

Monte Carlo (MC) integration is a numerical method for estimating the integral, I of a function F . The MC estimator, I_N , averages over N stochastic samples taken over the integrand in a given sampling domain \mathbb{D} . We consider samples $\{\mathbf{s}_1, \dots, \mathbf{s}_N\}$ as a set of equi-weighted random variables (Fig. 1). The resulting MC estimator is

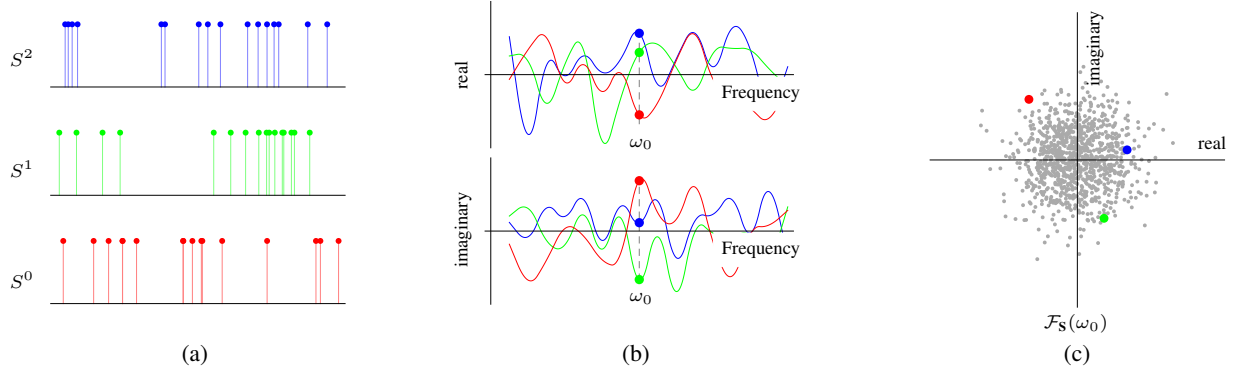


Figure 1: Illustration of the random nature of the Fourier coefficients of a 1D white noise sampling pattern. (a) Three realizations, S^0 , S^1 and S^2 of the sampling pattern \mathbf{S} . (b) The real and imaginary parts of the Fourier transform of these realizations. Points on the dashed line correspond to the three possible values of the Fourier transform for a given frequency ω_0 . (c) Distribution of values of $\mathcal{F}_{\mathbf{S}}(\omega_0)$ in the complex plane for 1024 realizations. The first three realizations are shown in their respective colors. Note that homogeneous sampling patterns have random Fourier coefficients that are uniformly distributed on each concentric circle in the complex plane.

defined as:

$$\mathbf{I}_N := \frac{\mu(\mathbb{D})}{N} \sum_{k=1}^N F(\mathbf{s}_k), \quad (1)$$

where $\mu(\mathbb{D})$ represents the Lebesgue measure of \mathbb{D} . Here, a sampling pattern can be written as a random signal \mathbf{S} made of N Dirac functions located at sample positions— $\mathbf{S}(x) = \sum \delta(x - \mathbf{s}_k)$ for $k = 1 : N$ —which can be represented in continuous form [Durand 2011] as:

$$\mathbf{I}_N = \frac{\mu(\mathbb{D})}{N} \int_{\mathbb{D}} \mathbf{S}(x) F(x) dx. \quad (2)$$

To build our analysis we consider sampling patterns which are homogeneous.

3.1 Homogeneous sampling

A *homogeneous sampling pattern* has statistical properties invariant to translation over the domain \mathbb{D} . The notion of homogeneous distribution of points is commonly used in physics [Hansen and McDonald 1990] for compound gases. In the computer graphics community, this notion is similar to the “widesense stationary process” described by Dippé and Wold [1985] which is discussed only for *Poisson sampling*. However, in our mathematical formulation, we extend the domain of application of this *widesense stationary* notion by making any kind of sampling pattern translation invariant, that is, homogeneous. In MC integration, the consequence is that a homogeneous sampling pattern produces no bias, so that error only arises from the variance. For more details, please refer to Appendix A. Given this property, we can restrict our error analysis to variance only.

Many state-of-the-art sampling methods such as White noise, Poisson Disk, and all optimization-based methods starting from a white noise distribution (e.g., the methods of Schlömer et al. [2011] and de Goes et al. [2012]) are homogeneous. Surprisingly, jittered sampling and Latin hypercube sampling are non-homogeneous. This can be explained by the fact that both methods rely on a subdivision of the sampling domain that is fixed over all realizations. However, any sampling method can be transformed into a homogeneous sampling by uniformly and randomly translating each realization of the sampling pattern. Analogously, sampling patterns on the unit sphere (S^2) can be made homogeneous by uniformly and randomly rotating the distribution.

4 MC variance analysis for Euclidean space

In this section, we perform variance analysis in (d -dimensional) Euclidean space. We represent variance in terms of the power spectra of both \mathbf{S} and F . For this purpose, we first represent the MC estimator in the Fourier domain. Then, we derive the variance expression in the Fourier domain using the homogeneity of the sampling patterns. We perform a similar analysis on the sphere (S^2) and hemisphere (\mathcal{H}^2) domains in Sec. 5. Without loss of generality, we consider our sampling domain to be a toroidal unit square $[0, 1]^d$, denoted \mathcal{T}^d .

4.1 Monte Carlo estimator in Fourier domain

The MC estimator \mathbf{I}_N (Eq. (2)) can be characterized in the Fourier domain (Φ) using the fact that dot-product of functions (the integral of the product) is equal to the dot-product of their Fourier coefficients. This makes the MC estimator, (Eq. (2)):

$$\mathbf{I}_N = \frac{\mu(\mathcal{T}^d)}{N} \int_{\Phi} \mathcal{F}_{\mathbf{S}}(\omega) \overline{\mathcal{F}_F(\omega)} d\omega, \quad (\text{E-3})$$

where $\mathcal{F}_{\mathbf{S}}$ and \mathcal{F}_F represent the Fourier transform of \mathbf{S} and F , respectively.

4.2 Variance in spectral form in Euclidean domain

To perform variance analysis, we follow a similar approach to Durand [2011] and, Subr and Kautz [2013]. We use the spectral form of the MC estimator (Eq. (E-3)) in the definition of variance, $\text{Var}(\mathbf{I}_N) := \langle \|\mathbf{I}_N\|^2 \rangle - \|\langle \mathbf{I}_N \rangle\|^2$, giving:

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{T}^d)^2}{N^2} \int_{\Phi \times \Phi} \langle \mathcal{F}_{\mathbf{S},F}(\omega, \omega') \rangle d\omega d\omega' - \|\langle \mathbf{I}_N \rangle\|^2, \quad (\text{E-4})$$

with:

$$\mathcal{F}_{\mathbf{S},F}(\omega, \omega') := \mathcal{F}_{\mathbf{S}}(\omega) \cdot \overline{\mathcal{F}_F(\omega)} \cdot \overline{\mathcal{F}_{\mathbf{S}}(\omega')} \cdot \mathcal{F}_F(\omega'), \quad (\text{E-5})$$

which gives us a closed-form expression for variance. However, unlike Subr and Kautz [2013], who relate the variance in integration with the variance of the underlying sampler, we provide an explicit closed-form relation of variance with the power spectra of both the sampling pattern and the integrand. To derive this relation, we first simplify Eq. (E-4) by restricting our analysis to *homogeneous* sampling patterns which are unbiased. As a result, the second

term in Eq. (E-4)—that is, $\|\langle I_N \rangle\|^2$ —becomes equal to the square-norm of the integrand. Given that the integral of the integrand is equal to $\mathcal{F}_F(0)$ and the DC peak of the weighted (by N) sampling pattern is equal to 1, the second term ($\|\langle I_N \rangle\|^2$) gets cancelled out with the $\mathcal{F}_{S,F}(0,0)$ term. This restricts the integral over the Fourier domain without the DC peak frequency, denoted by Θ . Since the homogeneous sampling patterns, mentioned in Sec. 3.1, have statistical properties that are invariant to translation, it is equivalent to study the variance generated by the translated version of each realization, with the average taken over all translations. Formally, we can treat the torus as the group of translations, so that $\tau(\mathbf{S})$ denotes the translation of \mathbf{S} by an element $\tau \in \mathcal{T}^d$. Then, averaging equation Eq. (E-4) over all translations of \mathbf{S} , we get:

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{T}^d)}{N^2} \int_{\mathcal{T}^d \times \Theta \times \Theta} \langle \mathcal{F}_{\tau(\mathbf{S}),F}(\omega, \omega') \rangle d\omega d\omega' d\tau \quad (\text{E-6})$$

$$= \frac{\mu(\mathcal{T}^d)}{N^2} \int_{\mathcal{T}^d \times \Theta \times \Theta} \langle \mathcal{F}_{\mathbf{S},F}(\omega, \omega') \rangle e^{i2\pi\tau \cdot (\omega' - \omega)} d\omega d\omega' d\tau, \quad (\text{E-7})$$

where the exponential derives from the translation of the sampling pattern by a vector τ in the Fourier domain. When $\omega \neq \omega'$, the integral of the exponential part equals zero, so that only the case $\omega = \omega'$ contributes to the variance. Thus, we can remove one integral over Θ :

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{T}^d)}{N^2} \int_{\Theta} \langle \mathcal{F}_{\mathbf{S},F}(\omega, \omega) \rangle \int_{\mathcal{T}^d} d\tau d\omega \quad (\text{E-8})$$

$$= \frac{\mu(\mathcal{T}^d)^2}{N^2} \int_{\Theta} \langle \|\mathcal{F}_{\mathbf{S},F}(\omega, \omega)\|^2 \rangle d\omega. \quad (\text{E-9})$$

Finally, denoting by \mathcal{P}_F to the power spectrum of the integrand and by \mathcal{P}_S to the power spectrum of the sampling pattern normalized by N , and leveraging the fact that $\|\mathcal{F}_{\mathbf{S},F}(\omega, \omega)\|^2 = \|\mathcal{F}_S(\omega)\|^2 \cdot \|\mathcal{F}_F(\omega)\|^2$, we get:

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{T}^d)^2}{N} \int_{\Theta} \langle \mathcal{P}_S(\omega) \rangle \mathcal{P}_F(\omega) d\omega. \quad (\text{E-10})$$

This gives us the expression for the variance of the MC integration in terms of the power spectra of both the sampling pattern and the integrand in the toroidal domain.

Isotropic sampling If homogeneous sampling is isotropic (*i.e.*, the power spectrum is radially symmetric), then the variance can be computed from the radial mean power spectrum of the integrand and the sampling pattern. For this, we first rewrite the integral in Eq. (E-10) in polar coordinates (\mathbf{n}, ρ):

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{T}^d)^2}{N} \int_0^\infty \rho^{d-1} \int_{S^{d-1}} \langle \mathcal{P}_S(\rho \mathbf{n}) \rangle \mathcal{P}_F(\rho \mathbf{n}) d\mathbf{n} d\rho, \quad (\text{E-11})$$

where S^{d-1} is a $(d-1)$ -dimensional unit sphere in \mathbb{R}^d . Since the sampling pattern is isotropic, we can replace $\langle \mathcal{P}_S(\omega) \rangle$ by its radial mean, denoted by $\check{\mathcal{P}}_S(\omega)$. Please note that the expected value is contained in this symbol. Then, the radial power spectrum no longer depends on \mathbf{n} , and thus:

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{T}^d)^2}{N} \int_0^\infty \rho^{d-1} \check{\mathcal{P}}_S(\rho) \int_{S^{d-1}} \mathcal{P}_F(\rho \mathbf{n}) d\mathbf{n} d\rho, \quad (\text{E-12})$$

where the inner integral corresponds to the definition of the radial mean of the integrand power spectrum (\mathcal{P}_F), denoted by $\check{\mathcal{P}}_F(\rho)$, giving the variance as:

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{T}^d)^2 \mu(S^{d-1})}{N} \int_0^\infty \rho^{d-1} \check{\mathcal{P}}_S(\rho) \check{\mathcal{P}}_F(\rho) d\rho, \quad (\text{E-13})$$

in terms of the *radial mean* spectra of both \mathbf{S} and F . Here, $\mu(S^{d-1}) = 2\sqrt{\pi^d}/\Gamma(d/2)$ is the surface area of the $(d-1)$ -dimensional unit sphere. The Lebesgue measure of the toroidal domain will be ignored in the rest of this paper when we refer to Eq. (E-13), since it is equal to one. In the next section, we show that a similar variance expression can also be obtained in for the sphere and hemisphere.

5 MC variance analysis for S^2 and \mathcal{H}^2

To perform variance analysis in the spherical domain (S^2), we follow the definition of MC integration from Eq. (2), with the domain of integration ($\mathbb{D} = S^2$) being the unit sphere (with the Lebesgue measure $\mu(S^2) = 4\pi$). This allows to rewrite the MC estimator as:

$$\mathbf{I}_N = \frac{\mu(S^2)}{N} \int_{S^2} \mathbf{S}(x) F(x) dx. \quad (\text{S-14})$$

Before proceeding, we review some background on spherical harmonics, used as an analog for the Fourier basis for performing spectral analysis on the sphere. We then represent the MC estimator (Eq. (S-14)) in spectral form using the spherical harmonics.

5.1 Background

The spherical harmonics (SH) are a basis for functions on the sphere [Groemer 1996; Wieczorek and Simons 2005]. The (l, m) -th basis function is given by:

$$Y_l^m(\theta, \phi) := \sqrt{(2 - \delta_{0m}) \frac{(2l+1)(l-m)!}{4\pi(l+m)!}} P_l^m(\cos \theta) e^{im\phi}. \quad (\text{S-15})$$

Here, δ_{ij} is the Kronecker delta function, $Y_l^m(\theta, \phi)$ is the complex spherical harmonic basis function of degree $l \geq 0$ and order $m \in [-l, l]$ with $P_l^m(x)$ denoting the *associated Legendre Polynomials*, for $x \in [-1, 1]$. The spherical harmonics are orthonormal and any integrable function F on S^2 can be decomposed into SH components. The inner product between any two arbitrary functions, $G(x)$ and $F(x)$ defined over a unit sphere, is related to its spectral coefficients by:

$$\int_{S^2} G(x) \overline{F(x)} dx = \sum_{l=0}^\infty \sum_{m=-l}^l \mathcal{S}_G(l, m) \cdot \overline{\mathcal{S}_F(l, m)}, \quad (\text{S-16})$$

where $\mathcal{S}_G(l, m) = \langle G, Y_l^m \rangle$ is the (l, m) -th spherical harmonic coefficients of $G(x)$.

5.2 Variance in spectral form in spherical domain

We are now ready to express the MC estimator from Eq. (S-14) in its spectral form using SH:

$$\mathbf{I}_N = \frac{\mu(S^2)}{N} \sum_{l=0}^\infty \sum_{m=-l}^l \mathcal{S}_S(l, m) \cdot \overline{\mathcal{S}_F(l, m)}. \quad (\text{S-17})$$

As in Sec. 4.2, by plugging the spherical MC estimator from Eq. (S-17) in the definition of variance, we obtain a closed-form expression of the variance in the spherical domain, similar to Eq. (E-4):

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{S}^2)^2}{N^2} \sum_{l,l'=0}^{\infty} \langle \mathcal{S}_{S,F}(l,l') \rangle - \|\langle \mathbf{I}_N \rangle\|^2, \quad (\text{S-18})$$

where:

$$\mathcal{S}_{S,F}(l,l') := \sum_{\substack{m=-l \\ m'=-l'}}^{l,l'} \mathcal{S}_S(l,m) \cdot \overline{\mathcal{S}_F(l,m)} \cdot \mathcal{S}_S(l',m') \cdot \overline{\mathcal{S}_F(l',m')}. \quad (\text{S-19})$$

Analogously, we simplify this variance expression (Eq. (S-18)) by restricting our analysis to homogeneous sampling patterns which are unbiased (*i.e.*, $\langle \mathcal{S}_S(l,m) \rangle = 0 \forall l, m \neq 0$ following Subr and Kautz [2013]). This makes the second term in Eq. (S-18)—*i.e.*, $\|\langle \mathbf{I}_N \rangle\|^2$ —non-zero only at the DC peak ($l = 0$) and which gets eliminated by the $\mathcal{S}_{S,F}(0,0)$ term. As a result, the contribution of components with $l = 0$ in the variance Eq. (S-18) becomes zero. This allows us to represent the variance in Eq. (S-18) using only the terms of $\mathcal{S}_{S,F}(l,l')$ with non-zero frequencies. Similar to the Euclidean case, we can leverage homogeneity by averaging over all rotations of the sampling pattern, $\tau(\mathbf{S})$:

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{S}^2)^2}{N^2} \frac{1}{\mu(\mathcal{S}^{SO(3)})} \sum_{l,l'=1}^{\infty} \left\langle \int_{SO(3)} \mathcal{S}_{\tau(\mathbf{S}),F}(l,l') d\tau \right\rangle, \quad (\text{S-20})$$

where $\tau \in SO(3)$ represents a rotation. To simplify the integral part in Eq. (S-20) we employ the *representation theory* which has been briefly explained in our *supplementary material, Section 2*. By using the fact that the span of the l -th frequency spherical harmonics, $V^l = \text{Span}\{Y_l^{-l}, \dots, Y_l^l\}$, is a $(2l+1)$ -dimensional irreducible representation for the $SO(3)$ group and applying *Propositions 2.1 and 2.2 from Section 2 of our supplementary material*, we have:

$$\frac{\int_{SO(3)} \langle \tau(\mathbf{G}), Y_l^{m'} \rangle \langle Y_l^m, \tau(\mathbf{G}) \rangle d\tau}{\mu(\mathcal{S}^{SO(3)})} = \delta_{ll'} \delta_{mm'} \mathcal{P}_G(l), \quad (\text{S-21})$$

for any function $G(x) \in \mathcal{S}^2$. Here, $\langle \cdot, \cdot \rangle$ is the (complex) inner-product operator on the space of function with $\mathcal{P}_G(l)$ as the angular power spectrum of the l -th spherical frequency of G , which is defined as $\mathcal{P}_G(l) := \frac{1}{2l+1} \sum_{m=-l}^l \|\mathcal{S}_G(l,m)\|^2$. Using the results from Eq. (S-21) and the definition of the *angular power spectrum* in (S-20), we get:

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{S}^2)^2}{N^2} \sum_{l,l'=1}^{\infty} \left\langle \mathcal{P}_S(l) \sum_{m=-l}^l \mathcal{S}_F(l,m) \cdot \overline{\mathcal{S}_F(l,m)} \right\rangle \quad (\text{S-22})$$

$$= \frac{\mu(\mathcal{S}^2)^2}{N^2} \sum_{l=1}^{\infty} \langle \mathcal{P}_S(l) \rangle \left(\sum_{m=-l}^l \|\mathcal{S}_F(l,m)\|^2 \right). \quad (\text{S-23})$$

Using the definition of angular power spectrum for F and normalizing $\langle \mathcal{P}_S(l) \rangle$ in Eq. (S-23) by $N/\mu(\mathcal{S}^2)$, we get:

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{S}^2)}{N} \sum_{l=1}^{\infty} (2l+1) \check{\mathcal{P}}_S(l) \mathcal{P}_F(l), \quad (\text{S-24})$$

where $\check{\mathcal{P}}_S(l)$ denotes the expected value of $\mathcal{P}_S(l)$, which is $\langle \mathcal{P}_S(l) \rangle$, from Eq. (S-22). This gives the final closed-form expression for variance in the spherical domain. As for Euclidean space, the variance depends on the *angular power spectrum* of \mathbf{S} and F .

Variance in spectral form on \mathcal{H}^2 A similar mathematical framework can be developed for the hemisphere (\mathcal{H}^2) using spherical harmonics. For this purpose, we associate the hemisphere with *projective 2-space* (more details in Section 3 of the *supplementary material*). The only assumption here is that we restrict our analysis to functions on the hemisphere with $G(x) = G(-x)$ for all points on the equator. This allows us to represent all hemispherical signals as even functions on the sphere and study their spectral properties using spherical harmonics. In the following sections, we will use our mathematical framework to obtain convergence rates of various state-of-the-art sampling patterns.

6 Best and worst case variance

Using the framework derived in the Euclidean (Sec. 4.2) and spherical (Sec. 5.2) domains, we have shown that the variance in MC integration is related to the power spectra of the integrand and the sampling pattern. This implies that if we know the sampling and integrand power spectra then we can predict the variance and the corresponding variance convergence rate, as a function of the number of samples N , in MC integration. In practice however, the power spectrum of an integrand is unknown. Therefore, we restrict our analysis to a particular class of functions. We follow the work of Brandolini and colleagues [2001], and restrict our analysis to integrable functions of the form $F(x)\chi_{\Omega}(x)$ with $F(x)$ smooth and Ω a bounded domain with smooth boundary (where, $\chi_{\Omega}(x)$ is the characteristic function of Ω). We consider a *best-case* function and a *worst-case* function, both from this class of functions to derive the best- and worst-case variance convergence rate, as the number of samples N increases.

Conversely, analytical models for the sampling power spectra are often available [Gabrielli and Torquato 2004]. We show that these power spectra can be used in our framework to predict the variance of an integrand. We bound the variance convergence rate of these samplers—with known power spectra—by simply imposing bounds on the sampling power spectra in MC integration. For sampling power spectra with no analytical formulations, we provide tools for bounding the power spectra in Sec. 7.

To perform our variance convergence rate analysis, we consider power spectra that are homogeneous and isotropic, and that can bound any homogeneous sampling power spectrum. First, we perform our theoretical convergence analysis in the Euclidean (toroidal) domain, following Eq. (E-13), where the variance is characterized by the radial mean power spectra of both sampling pattern and integrand signal.

Euclidean domain First, we derive theoretical convergence rates of the variance for the best- and the worst-case functions. We define our *best-case* integrand directly in the spectral domain with the radial mean power spectrum profile $\check{\mathcal{P}}_F(\rho)$ which is a constant c_F for $(\rho < \rho_0)$, and zero elsewhere. The constant c_F comes from the fact that the power spectrum is bounded. We provide a proof of this statement in the *supplementary material, Section 4.1*. The variance corresponding to the best-case can be obtained using Eq. (E-13) as follows:

$$\text{Var}(\mathbf{I}_N) \leq \frac{\mu(\mathcal{S}^{d-1})}{N} c_F \int_0^{\rho_0} \rho^{d-1} \check{\mathcal{P}}_S(\rho) d\rho. \quad (\text{E-25})$$

For the *worst-case*, we consider our integrand to exhibit a radial mean power spectrum:

$$\check{\mathcal{P}}_F(\rho) \leq \begin{cases} c_F & \text{if } \rho < \rho_0 \\ c'_F \rho^{-d-1} & \text{otherwise.} \end{cases} \quad (\text{E-26})$$

where c_F and c'_F are non-zero positive constants. This spectral profile $\check{\mathcal{P}}_F(\rho)$ has a decay rate $O(\rho^{-d-1})$ for $\rho > \rho_0$. This matches the worst-case spectral decay rate of the given class of functions derived by Brandolini and colleagues [2001]. The corresponding worst-case variance can be obtained using Eq. (E-13) as follows:

$$\begin{aligned} \text{Var}(\mathbf{I}_N) &\leq \frac{\mu(\mathcal{S}^{d-1})}{N} c_F \int_0^{\rho_0} \rho^{d-1} \check{\mathcal{P}}_S(\rho) d\rho \\ &\quad + \frac{\mu(\mathcal{S}^{d-1})}{N} c'_F \int_{\rho_0}^{\infty} \rho^{-2} \check{\mathcal{P}}_S(\rho) d\rho. \end{aligned} \quad (\text{E-27})$$

We use Eq. (E-25) and (E-27) in the next section (Sec. 7), to analyze the variance convergence rate of various sampling patterns in the Euclidean domain and to develop a theoretical tool that can be used in the design of new sampling patterns. For our experiments (Sec. 8), we consider a $2D$ disk function, as it falls into the given class of functions. The characteristic function of a disk in $2D$ exhibits a radial power spectral profile with a decay rate of order $O(\rho^{-3})$. For a mathematical derivation of this decay rate, please refer to our *supplementary material, Section 4.2*.

Spherical domain Analogously, in \mathcal{S}^2 , we choose our *best-case* integrand with an angular power spectrum $\check{\mathcal{P}}_F(l)$, which is a constant a_F for $(l < l_0)$ and zero elsewhere. The corresponding variance can be obtained using Eq. (S-24):

$$\text{Var}(\mathbf{I}_N) \leq \frac{a_F \mu(\mathcal{S}^2)^2}{N} \sum_{l=1}^{l_0} \check{\mathcal{P}}_S(l), \quad (\text{S-28})$$

where a_F is a non-zero positive constant, that can be derived by plugging in the best-case angular power spectrum.

From the class of functions provided by Brandolini and colleagues [2001], we define our *worst-case* integrand to be a sum of zonal harmonics, $\mathcal{S}_F(l, m) = 0 \forall m \neq 0$, with spectral amplitude square zonal coefficients $\|\mathcal{S}_F(l, 0)\|^2$ exhibiting a decay rate of order $O(l^{-2})$ for $l > l_0$. The corresponding worst-case *angular power spectrum* profile exhibits a decay rate of order $O(l^{-3})$, which can be used in Eq. (S-24) to obtain the worst-case variance:

$$\text{Var}(\mathbf{I}_N) \leq \frac{\mu(\mathcal{S}^2)^2 a'_F}{N} \sum_{l=0}^{\infty} \frac{\check{\mathcal{P}}_S(l)}{l^2}, \quad (\text{S-29})$$

where a'_F is a non-zero positive constant that can be derived by plugging in the worst-case angular power spectrum. For our experiments, we consider a spherical cap as our worst-case integrand as it falls into the given class of functions. The corresponding power spectrum exhibits a decay rate of order $O(l^{-3})$, whose derivation is given in our *supplementary material, Section 4.3*.

7 Theoretical convergence analysis

In the previous section (Sec. 6), we consider a class of functions with the corresponding best- and worst-case (integrand) power spectra. We then obtain the best- and worst-case variance, in the Euclidean (Eq. (E-25) and (E-27)) and spherical (Eq. (S-28) and (S-29)) domains, as a function of the sampling power spectrum $\check{\mathcal{P}}_S(\cdot)$. In this section, we study the effects of different shapes

of $\check{\mathcal{P}}_S(\cdot)$ on the variance in MC integration for both the best- and worst-case integrands. We use simple shape profiles (quadratic, polynomial, step) for $\check{\mathcal{P}}_S(\cdot)$ and derive the variance convergence rate associated with each sampling power spectra. Then, we classify the existing state-of-the-art sampling power spectra—with respect to the shape of the low frequency zone of these power spectra—in terms of these profiles. We also obtain corresponding upper and lower bounds on the sampling power spectra in terms of these profiles.

To analyze the variance convergence rate of a sampling pattern we first identify, which part of our variance formulation (Eq. (E-25) and (E-27) and, Eq. (S-28) and (S-29)) depends on N . Since the distribution of distances in a sampling pattern gets affected by the sampling density, the *mean neighborhood distance* between samples can be approximated as the d -dimensional root of the representative area of each sample. In the Fourier domain, the corresponding frequency of this *mean distance* is $\sqrt[d]{N}$, which corresponds to the first peak in the power spectrum of a regular grid pattern, and also, approximately, to the first bump of the typical blue-noise spectrum. As N increases, the low frequency region of the sampling power spectrum becomes more and more significant with respect to the integrand power spectrum. This explains the noticeable variance reduction observed for sampling patterns exhibiting no low-frequency content. Mathematically, this means that $\check{\mathcal{P}}_S(\omega)$ is dependent on a factor $\sqrt[d]{N}$. In the following subsections, we use this factor ($\sqrt[d]{N}$) to derive variance convergence rates.

7.1 Convergence analysis in Euclidean space

We start our analysis in the Euclidean space, and use our framework to derive the convergence rate of existing stochastic sampling methods. We cover classical stochastic samplers, from the simplest one (White noise) to the most sophisticated ones [Schlömer et al. 2011; de Goes et al. 2012]. We classify different sampling patterns with respect to the shape of their corresponding power spectra as follows:

Constant power spectra The simplest sampling profile is the constant function $\check{\mathcal{P}}_S(\rho) := \gamma$. Since the power spectrum is constant and $\check{\mathcal{P}}_F(\rho)$ does not depend on N , we can simplify Eq. (E-13) to obtain a convergence rate of order $O(N^{-1})$:

$$\begin{aligned} \text{Var}(\mathbf{I}_N) &= \frac{\mu(\mathcal{S}^{d-1})\gamma}{N} \int_0^{\infty} \rho^{d-1} \check{\mathcal{P}}_F(\rho) d\rho \\ &= O\left(\frac{1}{N}\right). \end{aligned} \quad (\text{E-30})$$

Step power spectrum The step profile corresponds to a constant power spectrum without the low-frequency part from $\rho = 0$ to $\alpha \sqrt[d]{N}$ for a given $\alpha \in \mathbb{R}^+ / \{0\}$, where α is used to quantify the range of energy-free frequency with respect to the mean frequency:

$$\check{\mathcal{P}}_S(\rho) := \begin{cases} 0, & \text{if } \rho < \alpha \sqrt[d]{N} \\ \gamma, & \text{otherwise.} \end{cases} \quad (\text{E-31})$$

For $N > \rho_0^d / \alpha^d$, the variance of MC integration of the best-case F is zero. For the worst-case, we obtain:

$$\begin{aligned} \text{Var}(\mathbf{I}_N) &\leq \frac{\mu(\mathcal{S}^{d-1})c'_F}{N} \int_{\alpha \sqrt[d]{N}}^{\infty} \frac{1}{\rho^2} d\rho \\ &= O\left(\frac{1}{N \sqrt[d]{N}}\right). \end{aligned} \quad (\text{E-32})$$

More generally, the variance generated by a step profile only depends on the rate of decay of the integrand in the frequency domain. If the

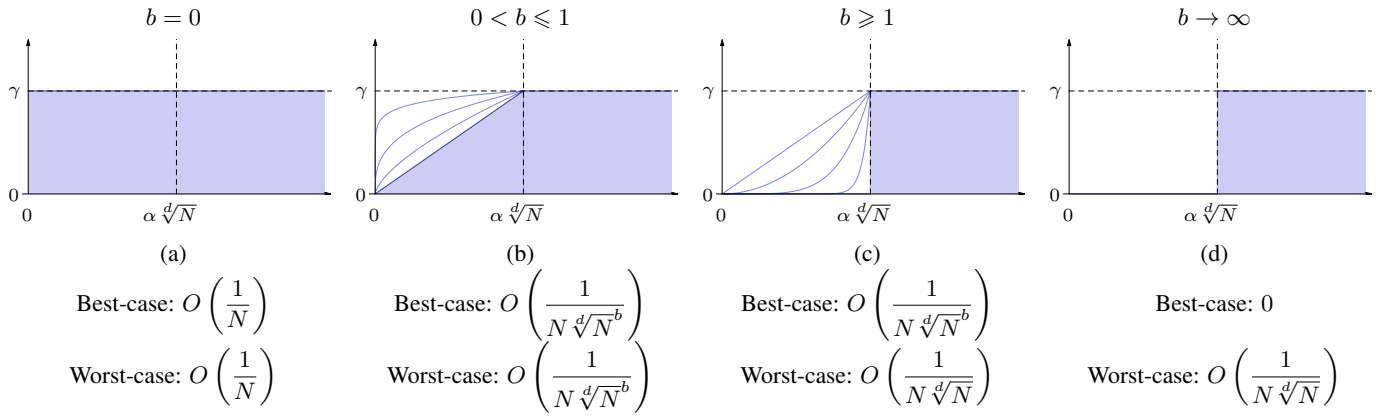


Figure 2: Summary of theoretical power spectra and their convergence rate for the best case and worst cases of integration. **From left to right:** Constant power spectrum, polynomial power spectrum with b less than 1, polynomial power spectrum with b greater than 1 and step power spectrum.

signal is bandwidth-limited, then it can be perfectly integrated. If not, the variance decreases as a function of the rate of decay of the power spectrum of the integrand.

Polynomial power spectra A more general radial power spectrum can be constructed from a polynomial ramp: from the frequency zero to $\alpha \sqrt[d]{N}$, the power spectrum is a polynomial ramp of degree b . Above frequency $\alpha \sqrt[d]{N}$, the radial power spectrum is a constant function, γ , similar to the constant and step power spectrum. More formally:

$$\check{\mathcal{P}}_S(\rho) := \begin{cases} \frac{\gamma \rho^b}{\alpha^b \sqrt[d]{N}^b} & \text{if } \rho < \alpha \sqrt[d]{N} \\ \gamma & \text{otherwise.} \end{cases} \quad (\text{E-33})$$

For $b = 0$, the polynomial power spectrum is equivalent to the constant power spectrum, and for $b \rightarrow \infty$, the power spectrum is equivalent to the step power spectrum. Following Eq. (E-25), the variance of the best-case is given by:

$$\begin{aligned} \text{Var}(\mathbf{I}_N) &\leq \frac{\mu(\mathcal{S}^{d-1}) \gamma c_F \rho_0^{d+b}}{\alpha^b N^{\frac{1}{d\sqrt{N}^b}}} \\ &= O\left(\frac{1}{N^{\frac{1}{d\sqrt{N}^b}}}\right). \end{aligned} \quad (\text{E-34})$$

For the worst-case integrand, we obtain

$$\begin{aligned} \text{Var}(\mathbf{I}_N) &\leq \frac{\mu(\mathcal{S}^{d-1}) \gamma c_F}{\alpha^b N^{\frac{1}{d\sqrt{N}^b}}} \int_0^{\rho_0} \rho^{d-1+b} d\rho \\ &\quad + \frac{\mu(\mathcal{S}^{d-1}) \gamma c'_F}{\alpha^b N^{\frac{1}{d\sqrt{N}^b}}} \int_{\alpha \sqrt[d]{N}}^{\rho_0} \rho^{b-2} d\rho \\ &\quad + \frac{\mu(\mathcal{S}^{d-1}) \gamma c'_F}{N} \int_{\alpha \sqrt[d]{N}}^{\infty} \frac{1}{\rho^2} d\rho. \end{aligned} \quad (\text{E-35})$$

Hence, when b is between 0 and 1, the convergence rate for the worst-case is of order $O(N^{-1-\frac{b}{d}})$. Additionally, when b is greater than 1, the convergence speed of the worst-case is of order $O(N^{-1-\frac{1}{d}})$.

7.2 Convergence analysis in \mathcal{S}^2 space

We perform a similar theoretical convergence analysis in order to obtain the variance convergence rates of various state-of-the-art

sampling patterns. We classify the sampling patterns in terms of the shape of their corresponding power spectra, as follows:

Constant power spectra Here, $\check{\mathcal{P}}_S(l) := \gamma$, where γ is a constant value. As the power spectrum is independent of N , we use Eq. (S-28) and Eq. (S-29) to deduce that the best- and worst-case variance convergence rates are of order $O(N^{-1})$.

Step power spectra We use Eq. (E-31) as our step power spectra ($\check{\mathcal{P}}_S(l)$), with l the frequency parameter and $d = 2$. We normalize the l term with $\alpha \sqrt{N}$, where α is a positive constant and N represents the number of samples. This allows us to bring all power spectra, shown in Fig. 5, to the same scale for all N . As in the Euclidean space, for $N > l^2/\alpha^2$ the variance of integration in the best-case is zero. The expression for the variance, $\text{Var}(\mathbf{I}_N)$, due to step power spectrum in the worst-case can be approximated using Eq. (S-29):

$$\text{Var}(\mathbf{I}_N) \leq \frac{\mu(\mathcal{S}^2)^2 a'_F}{N} \sum_{l=\lceil \alpha \sqrt{N} \rceil}^{\infty} \frac{1}{l^2}. \quad (\text{S-36})$$

where a'_F is a the hidden constant from $O(l^{-2})$. The summation term in the above equation can be symbolically solved to deduce the worst-case variance convergence rate of sampling patterns with a step power spectra, which is of order $O(N^{-1.5})$.

Quadratic power spectra In \mathcal{S}^2 , it is not trivial to obtain a general expression for a polynomial profile. Therefore, we directly study the polynomial profile defined as in Eq. (E-33) for $b = 2$ of a quadratic polynomial. Using this quadratic profile power spectrum, in the best-case variance Eq. (S-28), we obtain:

$$\text{Var}(\mathbf{I}_N) \leq \frac{\mu(\mathcal{S}^2)^2}{N} \sum_{l=1}^{l_0} \frac{\gamma a_F l^2}{(\alpha \sqrt{N})^2} \quad (\text{S-37})$$

$$= O(N^{-2}). \quad (\text{S-38})$$

Using the quadratic profile power spectrum for the worst-case variance Eq. (S-29), we obtain:

$$\text{Var}(\mathbf{I}_N) \leq \frac{\mu(\mathcal{S}^2)^2}{N} \left(\sum_{l=1}^{\lfloor \alpha \sqrt{N} \rfloor} \frac{\gamma a_F (l/(\alpha \sqrt{N}))^2}{l^2} + \sum_{l=\lceil \alpha \sqrt{N} \rceil}^{\infty} \frac{\gamma a'_F}{l^2} \right), \quad (\text{S-39})$$

which can be symbolically solved to obtain a convergence rate of order $O(N^{-1.5})$.

7.3 Discussion

The proposed classification of power spectra with respect to their shape can be used as a theoretical tool to bound arbitrary power spectra using simple profiles (e.g., linear, quadratic, etc.). Thus, the corresponding variance convergence rate of the sampling patterns can be deduced from the bounds of their respective profiles, summarized in Fig. 2. Below we discuss analytical formulations for some of the state-of-the-art sampling patterns whose variance convergence rate is derived using our framework. We discuss details regarding other sampling patterns in Sec. 8.

White noise The radial power spectrum of white noise has a constant profile with $\gamma = 1$ exhibiting a variance convergence rate of order $O(N^{-1})$, which is well known in literature.

Jittered Sampling Gabrielli and Torquato [2004] derived an analytical expression for the power spectrum of the jittered sampling in $d = 1$. By using the same formalism, one can obtain the analytical expression of jittered sampling in arbitrary d dimensions:

$$\mathcal{P}_S(\omega) = 1 - \prod_{i=1}^d \frac{\sin(\pi\omega_i)^2}{\pi^2\omega_i^2}, \quad (\text{E-40})$$

where ω_i is the i -th element of the vector ω . Despite the presence of the inverse function in this expression, the power spectrum is continuous everywhere and as $\omega_i \rightarrow 0$, the power goes to zero. Furthermore, we have $\lim_{\omega_i \rightarrow 0} \frac{\partial^2 \mathcal{P}_S(\omega_i)}{\partial \omega_i^2} = \frac{2\pi^2}{3}$, suggesting that the power spectrum has quadratic behavior near zero. In fact, we have numerically observed that \mathcal{P}_S can be upper bounded by a quadratic ramp ($b = 2$) with $\gamma = 1$ and $\alpha = \sqrt{3}/\pi$, and lower bounded by a quadratic ramp with $\gamma = 1/2$ and $\alpha = \sqrt{3}/\pi$, as shown in Fig. 4. Consequently, in the Euclidean domain, the variance convergence rate of MC integration with jittered sampling is of order $O(N^{-1-\frac{1}{d}})$ in the worst-case and $O(N^{-1-\frac{2}{d}})$ in the best-case. These results agree with the convergences rates found by Mitchell [1996].

Poisson Disk For Poisson Disk distributions, Torquato and colleagues [2006] observed that for a sufficiently large N , the minima of the mean power spectrum can be asymptotically seen near the DC peak, with a value of $1/20$. A similar observation can be made in the spherical domain. Following our theoretical analysis, the power spectrum of the Poisson disk distribution is lower bounded by a constant profile with $\gamma = 1/20$ for both Euclidean and spherical domains. Consequently, the variance convergence rate of the Poisson Disk sampling patterns is similar to the convergence rate of the white noise sampler, that is, of order $O(N^{-1})$. An illustration of the variance convergence rate is shown in Fig. 4 and Fig. 5.

8 Experiments and Results

In this section, we present several experiments comparing existing sampling methods including: white noise, jittered sampling, Poisson Disk and capacity constraint methods. First, we look at the implementation details and then proceed to our case study.

Implementation details In the toroidal domain (\mathcal{T}^d), we use the implementation by Gamito and Maddock [2009] for Poisson Disk sampling, for capacity constraint methods we choose the method of

de Goes and colleagues [2012]. We also include sampling patterns by Schlömer et al. [2011] and step noise as described by Heck and colleagues [2013]. In the spherical domain, we perform sampling directly on the sphere using the *Healpix* data structure [Górski et al. 2005] which is an equal area quadrangulation of the sphere. To implement jittered sampling, we use the *Healpix* quads as strata and randomly place samples in each stratum. For Poisson Disk sampling, there exist many algorithms in the literature that mimic the dart throwing approach on the surface of the sphere [Cline et al. 2009; Peyrot et al. 2013; Gamito and Maddock 2009]. In our implementation, we first generate a dense set of samples on the sphere, as in the work of Li and colleagues [2010], and then reject samples that are too close, as in the work of Bridson [2007]. We achieve a distribution close to the boundary sampling described in the work of Dunbar and Humphreys [2006]. We also implement a spherical version of the CCVT algorithm by Balzer and colleagues [2009] using the geodesic distance on the sphere. All sphere samplers are illustrated in Fig. 3 and are generated on a multiprocessor machine, with an Intel Core i7 CPU 980@3.3g GHz processor.

To plot the sampling power spectra, we compute the average power spectrum over 1000 realizations, as suggested by Schlömer and colleagues [2011]. We plot power spectra for $N = 16384$ in the toroidal domain and $N = 4096$ in the spherical domain. Later, we use these power spectra to estimate lower and upper bounds on the variance of our test functions in both the toroidal (Fig. 4) and spherical (Fig. 5) domains.

Case studies To support our theoretical claims for the variance convergence rate, we choose a 2D circular disk function in the Euclidean space and a spherical cap function in the spherical domain as our worst-case signals [Brandolini et al. 2001]. We also consider functions such as the Gaussian (in \mathcal{T}^2) and the spherical harmonic Y_l^m ($l = 4, m = 0$) basis function, with their respective parameters explained in our *supplementary material, Section 5*. Even though the Gaussian function does not belong to the given class of functions, we consider it a simple case of integration as it is a rapidly decreasing function with fast power spectral decay rate (due to its C^∞ nature) whereas the Y_l^m basis function is bandwidth-limited. We compute variance of all the above integrands with different sampling patterns. The variance was computed over 1000 trials for cases where the samplers (white noise and jittered) are not prohibitively slow. For spherical CCVT and Poisson Disk, we kept the number of trials between 200 to 1000, depending on the speed of the sampler. As the spectral profile of these test integrands is known, our variance prediction model can be used to estimate the bounds on the variance in integration of each function by simply using the bounds derived from the corresponding sampling power spectra.

We illustrate the variance computed via our framework in Fig. 4 and 5. Our experimental results are consistent with our variance predictive model (Sec. 7), and the variance is correctly bounded in both the toroidal and the spherical domains. In Fig. 4, we observe that tighter bounds on the sampling power spectra can result in tighter bounds on the variance. For example, the upper bound on the jittered sampling power spectrum is tighter than the lower bound which results in a tighter upper bound on the associated variance. For Poisson Disk sampling, the spectral profile is bounded by a constant profile and consequently, the convergence rate of the variance is similar to white noise. Also, as the lower bound on the Poisson Disk sampling is tighter than the upper bound, the corresponding variance achieves a tight lower bound. Surprisingly, for a given experimental setup, we observe an offset of 0.006935 for [Schlömer et al. 2011] in its spectral profile. This offset imposes bounds on the corresponding power spectrum with constant profiles, as in Poisson Disk. The variance generated by Schlömer and colleagues [2011] is one order of magnitude below Poisson disk sampling but the




Figure 3: Illustration of spherical sampling methods: Healpix underlying structure (a) used in regular sampling (b), stratified sampling (c) and our implementation of CCVT (d). Finally, (e) illustrates our implementation of Poisson disk sampling (see details in Sec. 8).

convergence rate is still $O(N^{-1})$. For the method of de Goes and colleagues [2012], the lower bound has a polynomial profile. Finally, instead of bounding the step power spectrum in Fig. 4(m) (in blue)—proposed by Heck and colleagues [2013]—we compare it with the theoretical step power spectral profile Eq. (E-31) (green, hidden under the blue curve). In Fig. 4(n) and 4(o), we demonstrate that the variance convergence rate of the step power spectrum generated by Heck and colleagues is clearly worse than its spectral target (step profile).

Similar behavior can also be observed in Fig. 5 for the sphere, where all the power spectral profiles are plotted as $(\check{\mathcal{P}}_F \text{ vs } \omega)$ with $\omega = l/(\alpha\sqrt{N})$. Variance bounds due to the white noise are also shown in Fig. 5. Note that, in Fig. 4 and 5, we only bound the variance convergence rate. This is why the variance values are normalized (with the variance value for $N = 1$) for all of our test integrands.

Rendering results We also implement ambient occlusion in a Cornell box scene to study the behavior of various state-of-the-art sampling patterns (shown in Fig. 3 of the supplementary material). As most state-of-the-art samplers are not directly extendable to the hemisphere, we sample a complete sphere and consider samples only from the visible part. We designed the Cornell box scene to enhance the noise level in the whole box. For rendering, we chose a regular sampling pattern on the image plane to keep the aliasing coherent throughout all samplers. This is why our reference image, shown as an inset in Fig. 6(f), shows some structural artifacts. The reference Cornell box image is computed by shooting one ray from the center of each pixel followed by shooting 32k secondary rays from each primary hit point. The secondary ray directions were sampled via a jittered sampling pattern. A similar process was adopted to generate the Cornell box images using other samplers. This is done to compare the variance coming through the hemispherical sampling patterns from each primary ray hit point.

Comparisons To provide a side by side comparison of the variance convergence rate of different samplers, we plot the convergence rate curves in Fig. 6 for both the Euclidean and spherical domains. As predicted by our model, in the Euclidean domain, we can see that the asymptotic behavior of de Goes and colleagues [2012] and jittered sampling is better than both the Poisson Disk sampling and the sampling method of Schlömer and colleagues [2011]. The jittered sampling convergence rate overtakes the Poisson Disk sampling convergence rate just after 200 samples and for the sampling method proposed by Schlömer and colleagues [2011], jittered sampling overtakes after 3072 samples (as shown in Fig. 6). Similar observations are made in the spherical domain for the jittered sampling patterns. CCVT [Balzer et al. 2009], jittered, and regular sampling patterns have asymptotically similar behaviors in the spherical domain, though jittered sampling patterns show higher variance.

9 Conclusion and future work

In this work, we develop a framework to analyze the effects of various sampling patterns on the variance of MC integration. We use the notion of *homogeneous* sampling patterns that allows manifestation of error only in terms of variance during MC integration. We show that the variance is directly related to the sampling and integrand power spectra in both the Euclidean and the spherical domains. We also show that for a given class of functions, we obtain the same worst-case variance convergence rates in both the domains.

One of the limitations of our framework is that it is designed for only stochastic samplers and assumes some knowledge of the power spectrum of the sampling pattern under study. Therefore, sampling patterns with varying power spectra (e.g. Sobol), that is, which depend on the number of samples, cannot be handled by our framework. We would like to explore this issue in future. Another promising direction for future research is to develop a similar framework for deterministic sampling methods like, Sobol or other low-discrepancy sequences.

Additionally, our framework can be used to study the effects of various sampling patterns with known and/or unknown closed-form power spectral formulation. We show that by using simple shapes (quadratic, constant, etc.) for the power spectra, existing state-of-the-art sampling power spectra can be easily bounded. This also results in bounding the associated variance for MC integration. We use our framework to theoretically derive best- and worst-case variance convergence rates—for a given class of functions—of various state-of-the-art sampling methods.

Design principles In our theoretical framework (Sec. 7), we propose some principles to design new sampling patterns. We also advocate the importance of the low frequency zone of the corresponding sampling power spectra. We emphasize the fact that an ideal sampling power spectrum must attain zero value as the frequency parameter tends to zero. This property ensures that the distribution is sufficiently uniform. In addition, the shape of the power spectrum near the zero frequency zone is also very important and must have at least a linear profile behavior. Moreover, the flatter the shape of the power spectral profile in this low frequency zone, the faster the variance converges in MC integration for favorable integrands. Ideally, the power spectrum must contain a range of low frequencies that are perfectly zero, such as the step power spectrum profile, shown in Fig. 2(d).

Our experimental results in the Euclidean and spherical domains show that no existing methods fulfill these design principles. Sampling methods based on packing and hard local conditions like Poisson Disk and [Schlömer et al. 2011], do not have a power spectrum converging to zero near zero frequency. This can be explained by the lack of uniformity in the distribution which can be corrected using proper subdivision of the domain through tile-based methods.

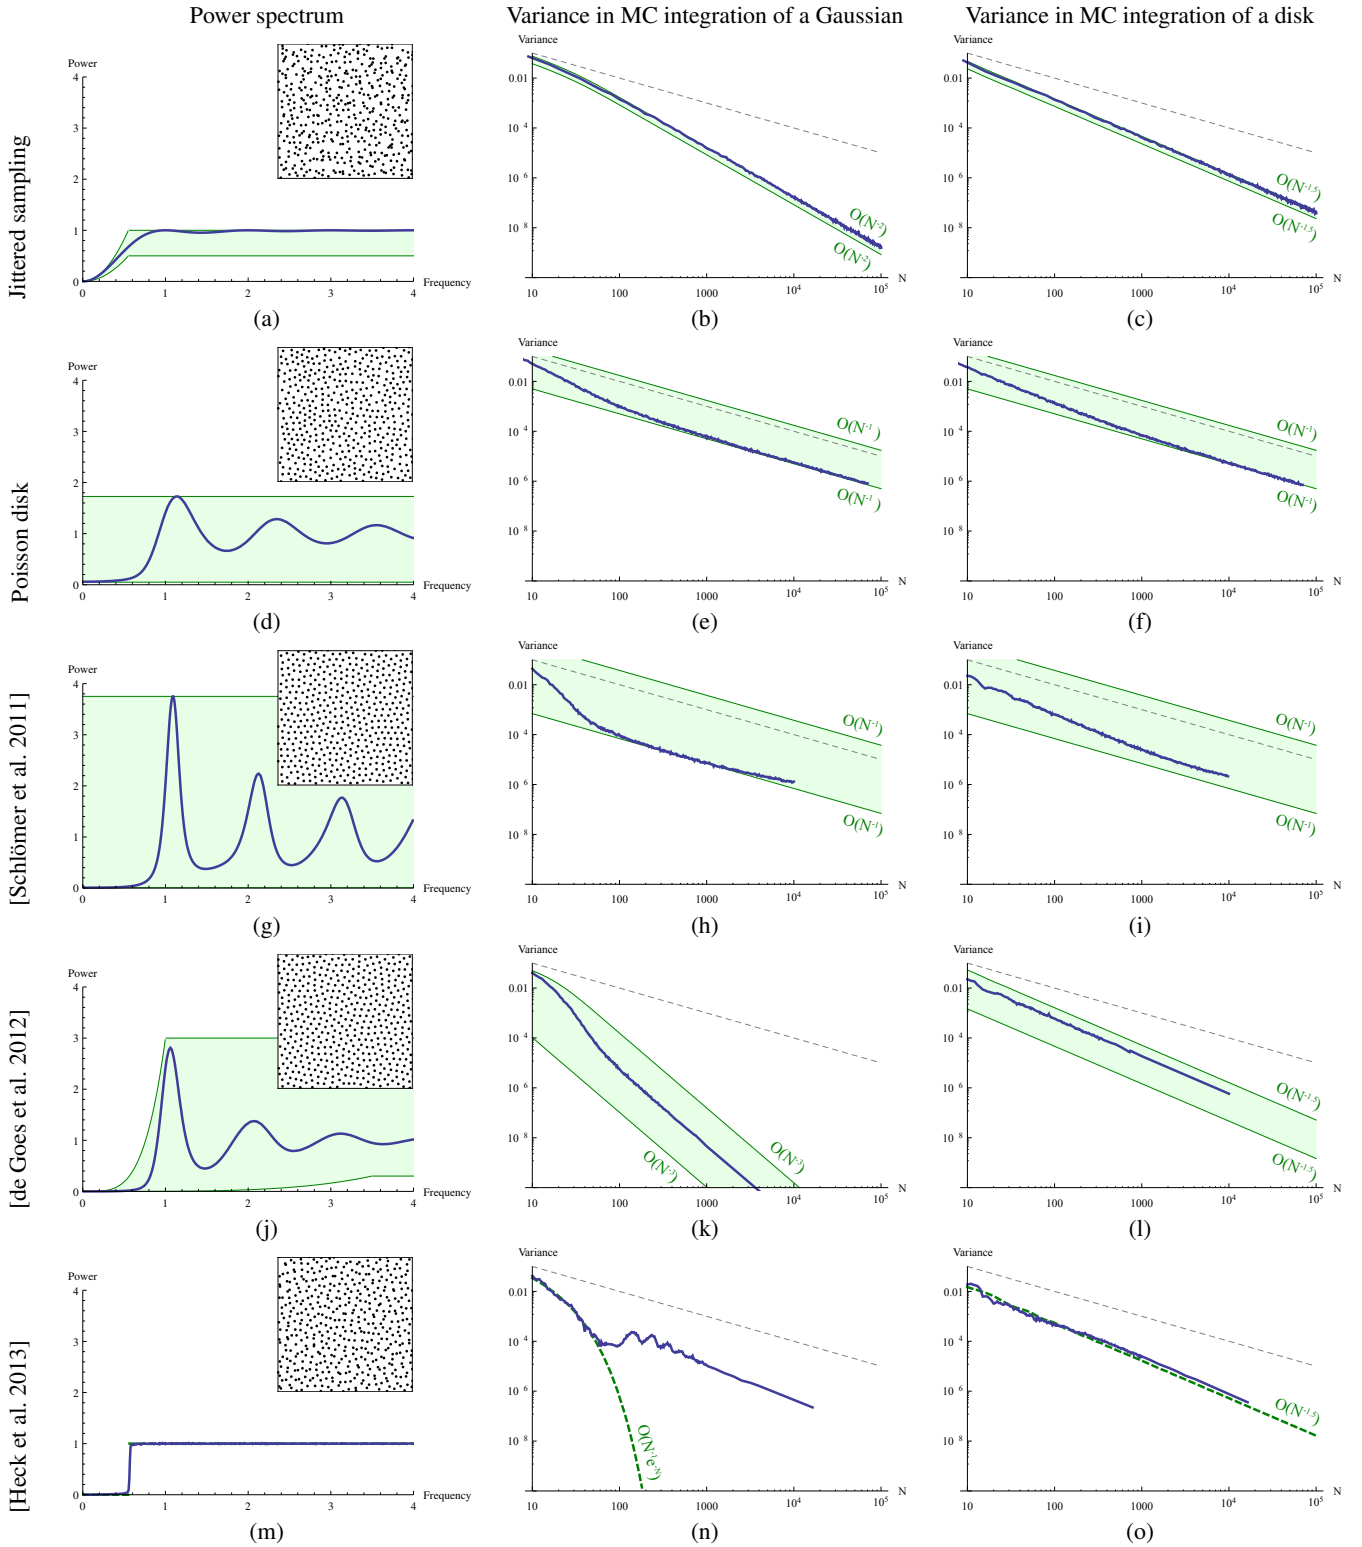


Figure 4: Bounds on the power spectra and on the variance convergence rate of our test integrands in the toroidal domain. Log-Log plots are shown in the Center and Right columns. **Left:** Power spectra of different sampling patterns (blue), bounded by a theoretical spectral profile (green). On the frequency axis, units corresponds to the frequency $\sqrt[4]{N}$. (d) and (g) have bounds with constant profiles while (a) and (j) have polynomial profiles. The corresponding parameters (γ, α, b) from Sec. 7, for both the upper $(\gamma_u, \alpha_u, b_u)$ and lower $(\gamma_l, \alpha_l, b_l)$ bounds are as follows: (a) $\gamma_u = 1, \alpha_u = \sqrt{3}/\pi, b_u = 2, \gamma_l = 1/2, \alpha_l = \sqrt{3}/\pi$ and $b_l = 2$, (d) $\gamma_u = 1.75$ and $\gamma_l = 1/20$, (g) $\gamma_u = 3.78$ and $\gamma_l = 0.006935$, (j) $\gamma_u = 3, \alpha_u = 1, b_u = 4, \gamma_l = 0.3, \alpha_l = 3.5$ and $b_l = 4$, and (m) is approximated by a step profile with $\gamma = 1$ and $\alpha = 1/\sqrt{\pi}$. **Center:** Gaussian function (in blue) with variance bounds (in green) computed using the bounds of the corresponding sampling power spectrum (Left). Variance generated by white noise is shown in dashed gray curve. **Right:** Similar visualization for a disk function.

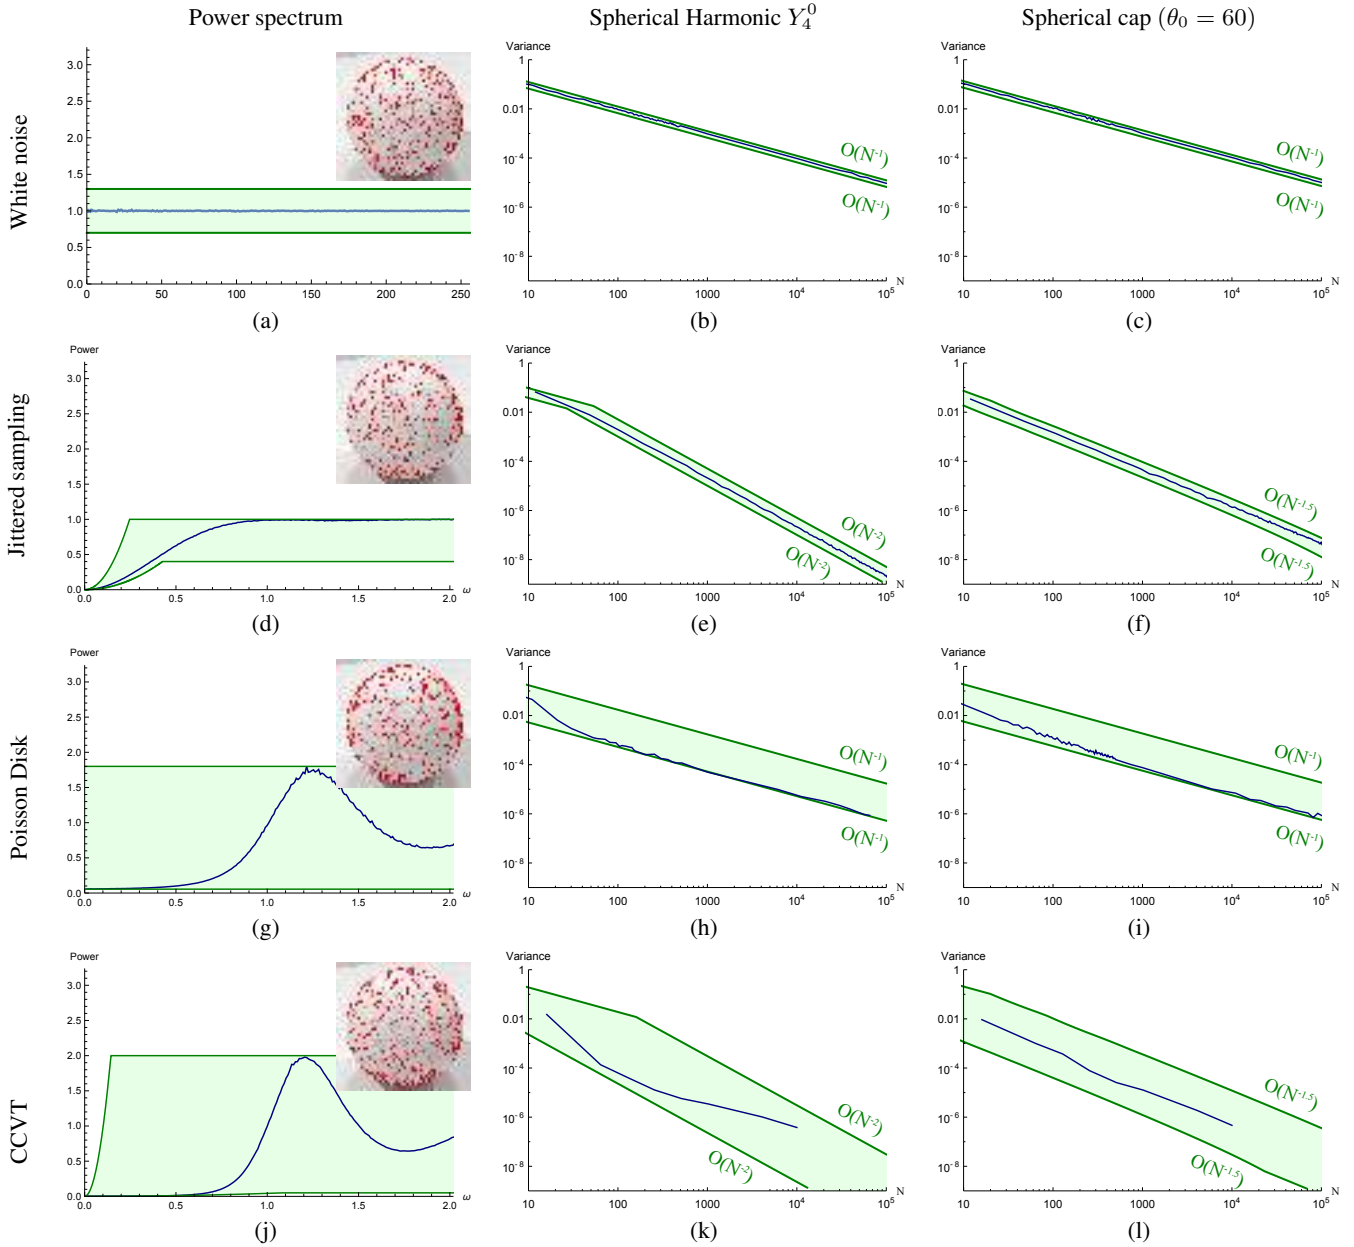


Figure 5: Bounds on the power spectra and on the variance convergence rate of our test integrands in \mathcal{S}^2 . Log-Log plots are shown in the **Center** and **Right** columns. **Left:** Power spectra of different sampling patterns (blue), bounded by a theoretical spectral profile (green). On the frequency axis, units corresponds to the frequency $\alpha\sqrt{N}$. (a) and (g) have bounds with constant profiles while (d) and (j) have quadratic ($b = 2$) profiles. The corresponding parameters (α, γ) from Sec. 7.2, for both the upper (α_u, γ_u) and lower (α_l, γ_l) bounds are also provided. (a) $\alpha_u = 1.0, \gamma_u = 1.3, \alpha_l = 1.0, \gamma_l = 0.7$ and (g) $\alpha_u = \sqrt{2.75}, \gamma_u = 1.8, \alpha_l = \sqrt{2.75}, \gamma_l = 0.055$ bounded with constant profiles while (d) $\alpha_u = 0.2, \gamma_u = 1.0, \alpha_l = \sqrt{0.6}, \gamma_l = 0.4$ and (j) $\alpha_u = \sqrt{0.05}, \gamma_u = 2.0, \alpha_l = \sqrt{2.85}, \gamma_l = 0.05$ with quadratic ($b = 2$) **Center:** The variance curve of a spherical harmonic basis function with $l = 4, m = 0$, (blue), with bounds (green) computed using the bounds of the corresponding power spectrum. **Right:** The variance in integration of a spherical cap, using the same visualization.

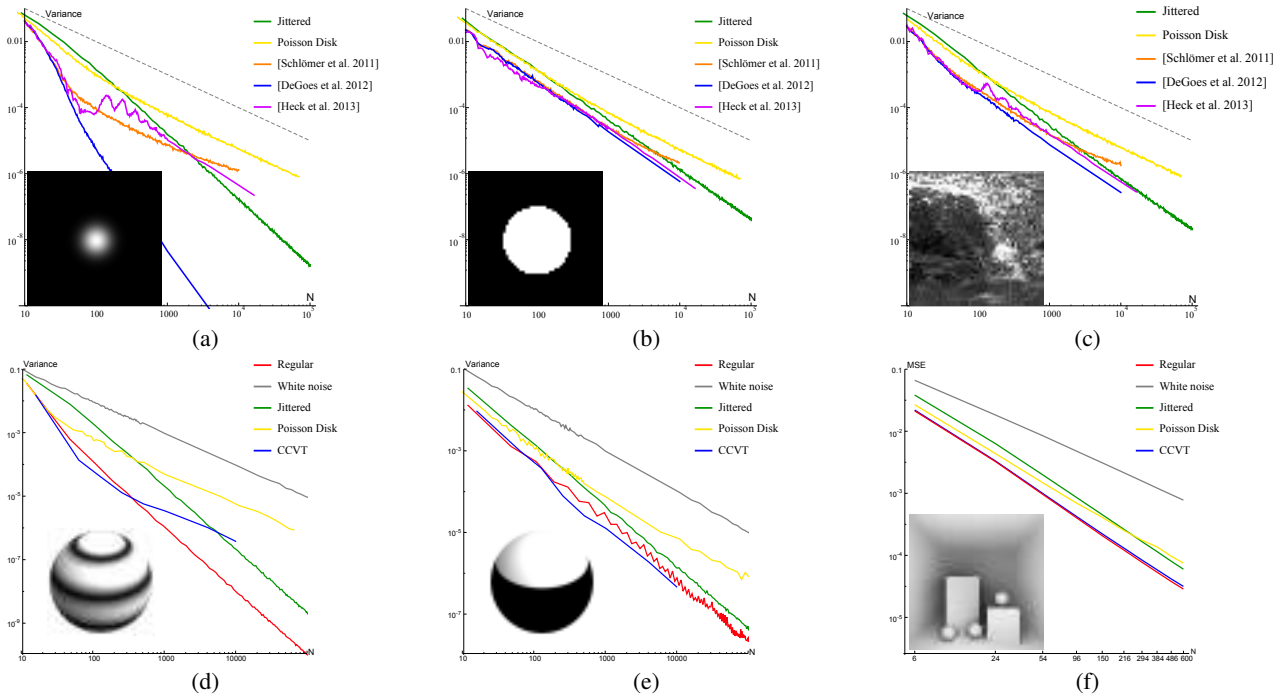


Figure 6: Comparison of the variance in MC integration for different integrand signals. Here we use experimental data from Fig. 4 and 5 in (a), (b), (d) and (e). **Top row** represents comparisons in the Euclidean space for (a) a Gaussian function, (b) a disk function and (c) an HDR image (16000×16000 pixels, taken from sIBL Archive, full image shown in the supplementary material). For reference, variance due to white noise is shown in the dashed gray line. **Bottom row** represents comparisons in the (hemi-)spherical domain : (d) a spherical harmonic basis function (Y_4^0). Inset illustrates the gray scale of the absolute values of the function, (e) a spherical cap function ($\theta_0 = 60$), where the white shade in the inset represents non-zero constant value region (f) a Cornell box scene.

For example, the tiled version of Poisson Disk sampling in Wachtel et al. [2014] has power spectrum converging to zero towards zero frequency. Similarly, all methods based on the capacity constraint [Balzer et al. 2009] seem to guarantee a power spectrum that converges to zero. However, these methods still have residual energy spread along low frequencies which can be improved.

Quadrature rules on the sphere In our framework, we perform Monte Carlo integration—instead of using optimal adaptive quadrature rules—on the spherical and hemispherical domains despite their low dimensional nature. We do this for a couple of reasons: First, our present framework is designed only for non-adaptive sampling patterns. Adaptive or importance-based sampling do not support the *homogeneity* of the sampling patterns, and designing a variance analysis framework in such a case needs further research; Second, as most production renderers are Monte Carlo based, our framework provides an analysis tool that can be employed in the production pipeline for choosing a sampling pattern that reduces the overall variance in integration of (hemi-)spherical signals. New sampling patterns can also be designed using our framework.

Future work Generating a sampling pattern with absolutely no energy in the low frequency range seems to be a very challenging problem. We expect to use the theoretical results and tools developed in this paper to help design such sampling patterns. Finally, taking into account more sophisticated sampling strategies, such as adaptive (or importance) sampling and filtering, could be an interesting avenue for future work.

10 Acknowledgements

This project was supported by the ANR excellence chair (ANR-10-CEXC-002-01) and digitalSnow/digitalfoam programs (ANR-11-BS02-009 and PALSE/2013/21). We thank the anonymous reviewers for their detailed feedback to improve the final version of the paper, Mathieu Desbrun, Katherine Breeden, Jonathan Dupuy, Nicolas Bonneel and Brian Wyvill for proof reading the paper and giving their insightful comments, Jean-Claude Iehl and Vincent Nivoliere for their active participation, and Kartic Subr for the discussions that ultimately led to this paper. Our special thanks to David Kazhdan for helping us in developing rigorous proofs.

A Proof: unbiased homogeneous sampling

In this appendix, we mathematically illustrate that the homogeneous sampling patterns are unbiased in nature. For this, we show that the expected value of the Monte Carlo estimator is equal to the integration of F . The expected value of the MC estimator can be written as:

$$\langle \mathbf{I}_N \rangle = \frac{\mu(\mathbb{D})}{N} \left\langle \int_{\mathbb{D}} \mathbf{S}(x) F(x) dx \right\rangle = \frac{\mu(\mathbb{D})}{N} \int_{\mathbb{D}} \langle \mathbf{S}(x) \rangle F(x) dx. \quad (41)$$

To make \mathbf{S} homogeneous, we perform averaging over the group of motions, H (translations for the torus and rotations for the sphere) over the whole domain, this gives us:

$$\langle \mathbf{I}_N \rangle = \frac{\mu(\mathbb{D})}{N} \int_{\mathbb{D}} \frac{1}{\mu(H)} \left\langle \int_H [\tau(\mathbf{S})](x) d\tau \right\rangle F(x) dx. \quad (42)$$

To simplify the inner integral, we use the fact that H acts transitively on \mathbb{D} , as a result, the average of $[\tau(S)](x)$ over all $\tau \in H$ is the average of $S(x)$ over all $x \in \mathbb{D}$:

$$\frac{1}{\mu(H)} \int_H [\tau(S)](x) d\tau = \frac{1}{\mu(\mathbb{D})} \int_{\mathbb{D}} S(x) dx. \quad (43)$$

Plugging Eq. (43) back into Eq. (42) (with a change of variable), we get:

$$\langle \mathbf{I}_N \rangle = \frac{\mu(\mathbb{D})}{N} \int_{\mathbb{D}} \left\langle \frac{1}{\mu(\mathbb{D})} \int_{\mathbb{D}} \mathbf{S}(y) dy \right\rangle F(x) dx \quad (44)$$

$$= \frac{1}{N} \int_{\mathbb{D}} \langle \mathbf{S}(y) \rangle dy \int_{\mathbb{D}} F(x) dx. \quad (45)$$

In particular, when $\mathbf{S}(y) = \sum \delta(y - s_k)$ for $k = 1 : N$, we get:

$$\langle \mathbf{I}_N \rangle = \int_{\mathbb{D}} F(x) dx. \quad (46)$$

□

References

- ARVO, J. 1995. Stratified sampling of spherical triangles. In *Proc. SIGGRAPH '95*, ACM, 437–438.
- ARVO, J. 2001. Stratified sampling of 2-manifolds. *SIGGRAPH 2001 Course Notes* 29, 2.
- BALZER, M., SCHLÖMER, T., AND DEUSSEN, O. 2009. Capacity-constrained point distributions: A variant of Lloyd's method. *ACM Trans. on Graphics* 28, 3, 86:1–8.
- BOWERS, J., WANG, R., WEI, L.-Y., AND MALETZ, D. 2010. Parallel Poisson disk sampling with spectrum analysis on surfaces. In *Proc. SIGGRAPH Asia '10*, ACM, 166:1–166:10.
- BRANDOLINI, L., COLZANI, L., AND TORLASCHI, A. 2001. Mean square decay of Fourier transforms in euclidean and non euclidean spaces. *Tohoku Math. J. (2)* 53, 3, 467–478.
- BRAUCHART, J., SAFF, E., SLOAN, I., AND WOMERSLEY, R. 2014. QMC designs: optimal order Quasi Monte Carlo integration schemes on the sphere. *Mathematics of Computation*.
- BRIDSON, R. 2007. Fast Poisson disk sampling in arbitrary dimensions. In *Proc. SIGGRAPH '07 Sketches*, ACM, Proc. SIGGRAPH '07.
- CHOIRAT, C., AND SERI, R. 2013. Computational aspects of cuifreen statistics for equidistribution on the sphere. *Mathematics of Computation* 82, 284, 2137–2156.
- CLINE, D., JESCHKE, S., WHITE, K., RAZDAN, A., AND WONKA, P. 2009. Dart throwing on surfaces. In *Proc. EGSR '09*, Eurographics Association, 1217–1226.
- COHEN, M., SHADE, J., HILLER, S., AND DEUSSEN, O. 2003. Wang tiles for image and texture generation. *ACM Trans. on Graphics* 22, 3, 287–294.
- COOK, R. L. 1986. Stochastic sampling in computer graphics. *ACM Trans. Graph.* 5, 1, 51–72.
- CROW, F. C. 1977. The aliasing problem in computer-generated shaded images. *Commun. ACM* 20, 11, 799–805.
- CUI, J., AND FREEDEN, W. 1997. Equidistribution on the sphere. *SIAM Scientific Computing* 18, 2, 595–609.
- DE GOES, F., BREEDEN, K., OSTROMOUKHOV, V., AND DESBRUN, M. 2012. Blue noise through optimal transport. *Proc. SIGGRAPH Asia '12* 31, 171:1–171:10.
- DIPPÉ, M. A. Z., AND WOLD, E. H. 1985. Antialiasing through stochastic sampling. In *Proc. SIGGRAPH '85*, ACM, 69–78.
- DUNBAR, D., AND HUMPHREYS, G. 2006. A spatial data structure for fast Poisson-disk sample generation. In *Proc. SIGGRAPH '06*, ACM, 503–508.
- DURAND, F. 2011. A frequency analysis of Monte-Carlo and other numerical integration schemes. *MIT CSAIL Technical report TR-2011-052*.
- GABRIELLI, A., AND TORQUATO, S. 2004. Voronoi and void statistics for superhomogeneous point processes. *Physical Review E* 70, 4, 041105.
- GAMITO, M. N., AND MADDOCK, S. C. 2009. Accurate multidimensional Poisson-disk sampling. *ACM Trans. on Graphics* 29, 1, 8.
- GÓRSKI, K. M., HIVON, E., BANDAY, A. J., WANDEL, B. D., HANSEN, F. K., REINECKE, M., AND BARTELMANN, M. 2005. Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal* 622, 2, 759.
- GROEMER, H. 1996. *Geometric Applications of Fourier Series and Spherical Harmonics*. Cambridge University Press. Cambridge Books Online.
- HANSEN, J.-P., AND McDONALD, I. R. 1990. *Theory of simple liquids*. Elsevier.
- HECK, D., SCHLÖMER, T., AND DEUSSEN, O. 2013. Blue noise sampling with controlled aliasing. *ACM Trans. on Graphics* 32, 3, 25:1–25:12.
- HESSE, K., SLOAN, I., AND WOMERSLEY, R. 2010. Numerical integration on the sphere. In *Handbook of Geomathematics*, W. Freeden, M. Nashed, and T. Sonar, Eds. Springer Berlin Heidelberg, 1185–1219.
- JAROSZ, W., CARR, N. A., AND JENSEN, H. W. 2009. Importance sampling spherical harmonics. *Computer Graphics Forum (Proceedings of Eurographics)* 28, 2 (Apr.), 577–586.
- KAUTZ, J., SLOAN, P.-P., AND SNYDER, J. 2002. Fast, arbitrary brdf shading for low-frequency lighting using spherical harmonics. In *Proc. of the 13th Eurographics Workshop on Rendering*, Eurographics Association, 291–296.
- KAZHDAN, M. 2007. An approximate and efficient method for optimal rotation alignment of 3d models. *IEEE Trans. Pattern Anal. Mach. Intell.* 29, 7, 1221–1229.
- KELLER, A., PREMOZE, S., AND RAAB, M. 2012. Advanced (quasi) Monte Carlo methods for image synthesis. In *SIGGRAPH '12 Courses*, ACM, New York, USA, 21:1–21:46.
- KOPF, J., COHEN-OR, D., DEUSSEN, O., AND LISCHINSKI, D. 2006. Recursive wang tiles for real-time blue noise. *ACM Trans. on Graphics* 25, 3, 509–518.
- LEMIEUX, C. 2009. *Monte Carlo and Quasi Monte Carlo Sampling*. Springer.
- LI, H., WEI, L.-Y., SANDER, P. V., AND FU, C.-W. 2010. Anisotropic blue noise sampling. In *Proc. SIGGRAPH Asia '10*, ACM, 167:1–167:12.

- MARQUES, R., BOUVILLE, C., RIBARDIRE, M., SANTOS, L. P., AND BOUATOUCH, K. 2013. Spherical Fibonacci point sets for illumination integrals. *Computer Graphics Forum* 32, 8, 134–143.
- MCEWEN, J., AND WIAUX, Y. 2011. A novel sampling theorem on the sphere. *Signal Processing, IEEE Trans. on* 59, 12, 5876–5887.
- MITCHELL, D. P. 1987. Generating antialiased images at low sampling densities. In *Proc. SIGGRAPH '87*, 65–72.
- MITCHELL, D. 1991. Spectrally optimal sampling for distributed ray tracing. In *Proc. SIGGRAPH '91*, vol. 25, 157–164.
- MITCHELL, D. P. 1996. Consequences of stratified sampling in graphics. In *Proc. SIGGRAPH '96*, ACM, 277–280.
- NIEDERREITER, H. 1992. *Random Number Generation and Quasi-Monte-Carlo Methods*. SIAM.
- OSTROMOUKHOV, V. 2007. Sampling with polyominoes. In *ACM Trans. on Graphics*, vol. 26, 78.
- ÖZTIRELI, A. C., AND GROSS, M. 2012. Analysis and synthesis of point distributions based on pair correlation. *ACM Trans. Graph.* 31, 6, 174:1–174:6.
- PEYROT, J.-L., PAYAN, F., AND ANTONINI, M. 2013. Feature-preserving direct blue noise sampling for surface meshes. In *Eurographics 2013*, 4 pages.
- RAMAMOORTHI, R., AND HANRAHAN, P. 2001. A signal-processing framework for inverse rendering. In *Proc. SIGGRAPH '01*, ACM, 117–128.
- RAMAMOORTHI, R., ANDERSON, J., MEYER, M., AND NOWROUZSAHRAI, D. 2012. A theory of monte carlo visibility sampling. *ACM Trans. on Graphics* 31, 5, 121:1–121:16.
- SCHLÖMER, T., HECK, D., AND DEUSSEN, O. 2011. Farthest-point optimized point sets with maximized minimum distance. In *Proc. Symp. High Performance Graphics '11*, ACM, 135–142.
- SHIRLEY, P. 1991. Discrepancy as a quality measure for sample distributions. In *Proc. Eurographics '91*, 183–194.
- SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *Proc. SIGGRAPH '02*, ACM, 527–536.
- SUBR, K., AND KAUTZ, J. 2013. Fourier analysis of stochastic sampling strategies for assessing bias and variance in integration. *ACM Trans. on Graphics* 32, 4, 128:1–128:12.
- SUBR, K., NOWROUZSAHRAI, D., JAROSZ, W., KAUTZ, J., AND MITCHELL, K. 2014. Error analysis of estimators that use combinations of stochastic sampling strategies for direct illumination. *Computer Graphics Forum (Proceedings of EGSR)* 33, 4 (June), 93102.
- TORQUATO, S., UCHE, O., AND STILLINGER, F. 2006. Random sequential addition of hard spheres in high euclidean dimensions. *Physical Review E* 74, 6, 061308.
- ULICHNEY, R. 1987. *Digital Halftoning*. MIT Press.
- UREÑA, C., FAJARDO, M., AND KING, A. 2013. An area-preserving parametrization for spherical rectangles. *Computer Graphics Forum* 32, 4, 59–66.
- WACHTEL, F., PILLEBOUE, A., COEURJOLLY, D., BREEDEN, K., SINGH, G., CATHELIN, G., DE GOES, F., DESBRUN, M., AND OSTROMOUKHOV, V. 2014. Fast tile-based adaptive sampling with user-specified Fourier spectra. *ACM Trans. on Graphics* 33, 4, 56:1–56:11.
- WIECZOREK, M. A., AND SIMONS, F. J. 2005. Localized spectral analysis on the sphere. *Geophysical Journal International* 162, 3, 655–675.
- XU, Y., HU, R., GOTSMAN, C., AND LIU, L. 2012. Blue noise sampling of surfaces. *Computers & Graphics* 36, 4, 232–240.
- ZHOU, Y., HUANG, H., WEI, L.-Y., AND WANG, R. 2012. Point sampling with general noise spectrum. *ACM Transactions on Graphics (TOG)* 31, 4, 76.