Variance Analysis for Monte Carlo Integration: A Representation-Theoretic Perspective

Michael Kazhdan, Gurprit Singh, Adrien Pilleboue, David Coeurjolly, Victor Ostromoukhov

HAL Id: hal-01259838

https://hal.archives-ouvertes.fr/hal-01259838

Submitted on 21 Jan 2016

Variance Analysis for Monte Carlo Integration:
A Representation-Theoretic Perspective

Michael Kazhdan[1]   Gurprit Singh[2]   Adrien Pilleboue[2]   David Coeurjolly[3]   Victor Ostromoukhov[2,3]
[1]Johns Hopkins University   [2]Université Lyon 1   [3]CNRS/LIRIS UMR 5205

1 Overview

In this report, we revisit the work of Pilleboue et al. [2015], providing a representation-theoretic derivation of the closed-form expression for the expected value and variance in homogeneous Monte Carlo integration. We show that the results obtained for the variance estimation of Monte Carlo integration on the torus, the sphere, and Euclidean space can be formulated as specific instances of a more general theory. We review the related representation theory and show how it can be used to derive a closed-form solution.

2 Problem Statmement

We begin by reviewing some basic concepts from Monte Carlo integration. Next, we present a formal definition of homogeneity. And finally, we formulate the generalized problem statement.

Monte Carlo Integration

Definition Given a domain Ω and given two (complex-valued) functions $F, G : \Omega \to \mathbb{C}$, the *dot-product* of the functions is the integral of the product of F with the complex conjugate of G:

$$\langle F, G \rangle = \int_{\Omega} F(x) \cdot \overline{G(x)} \, \mathrm{d}x.$$

Definition Given a domain Ω and given $S = \{s_1, \cdots, s_N\} \in \Omega^N$, the *Monte Carlo estimate* of the integral of a function $F : \Omega \to \mathbb{C}$ is obtained by averaging the values of F at the N positions:

$$\mathrm{MC}(F, S) := \frac{1}{N} \sum_{i=1}^{N} F(s_i).$$

Treating S as the average of delta functions, centered at $\{s_i\}$:

$$S(x) \equiv \frac{1}{N} \sum_{i=1}^{N} \delta_{s_i}(x),$$

the Monte Carlo estimate becomes the dot-product of F and S:

$$\mathrm{MC}(F, S) = \langle F, S \rangle.$$

Definition Given a domain Ω and a positive integer N, a *sampling pattern* is a function $\mathcal{P} : \Omega^N \to \mathbb{R}$ over the set of all N-tuples of points in Ω, satisfying:

$$\int_{\Omega^N} \mathcal{P}(S) \, \mathrm{d}S = |\Omega| \qquad \text{and} \qquad \mathcal{P}(S) \geq 0, \quad \forall \, S \in \Omega^N,$$

where $|\Omega|$ is the measure of Ω.

Definition Given a sampling pattern $\mathcal{P}$ and a function $F : \Omega \to \mathbb{C}$, the *expected value* of the integral of F and the *variance* in the estimate of the integral are given by:

$$\mathsf{E}_{\mathcal{P}}[\langle F, \mathbf{S} \rangle] := \int_{\Omega^N} \langle F, S \rangle \cdot \mathcal{P}(S) \, \mathrm{d}S$$

$$\mathsf{Var}_{\mathcal{P}}(\langle F, \mathbf{S} \rangle) := \mathsf{E}_{\mathcal{P}}\left[\|\langle F, \mathbf{S} \rangle\|^2\right] - \|\mathsf{E}_{\mathcal{P}}[\langle F, \mathbf{S} \rangle]\|^2.$$

Homogeneity

In order to make the problem of estimating the variance in Monte Carlo integration tractable, we restrict ourselves to sampling patterns that are homogeneous. To make this formal, we first define a notion of a group action.

Definition We say that a group, Γ, *acts on* Ω if each element $\gamma \in \Gamma$ defines a maps $\gamma : \Omega \to \Omega$ the preserves the measure on Ω.

Notation Given a group action of Γ on Ω, given $F : \Omega \to \mathbb{C}$, and given $\gamma \in \Gamma$, we denote by $\gamma(F) : \Omega \to \mathbb{C}$ the function obtained by applying the inverse of γ to the argument of F:

$$[\gamma(F)](x) := F\left(\gamma^{-1}(x)\right).$$

Here, inversion is required so that $(\gamma \circ \tilde{\gamma})(F) = \gamma(\tilde{\gamma}(F))$ for all $\gamma, \tilde{\gamma} \in \Gamma$.

Remark Since the map $\gamma : \Omega \to \Omega$ preserves the measure, the associated map on the space of functions is unitary. That is, for any functions $F, G : \Omega \to \mathbb{C}$ we have:

$$\langle F, G \rangle = \langle \gamma(F), \gamma(G) \rangle, \qquad \forall \, \gamma \in \Gamma.$$

Definition Given a group action of Γ on Ω, we say that a sampling pattern $\mathcal{P} : \Omega^N \to \mathbb{R}$ is *homogeneous with respect to* Γ if the probability of choosing a sampling pattern is the same as the probability of choosing any of its transformations by the group elements:

$$\mathcal{P}(S) = \mathcal{P}\left(\gamma(S)\right), \qquad \forall \, \gamma \in \Gamma.$$

(Note that we can either think of $\gamma(S)$ as the sampling pattern obtained by transforming the sample positions, $s_i \mapsto \gamma(s_i)$, or as the transform of the sum of delta functions – the two definitions are consistent.)

Remark If the group Γ is compact, one can always transform an initial sampling pattern $\mathcal{P}_0$ into a homogeneous sampling pattern $\mathcal{P}$ by averaging over the group elements:

$$\mathcal{P}(S) := \frac{1}{|\Gamma|} \int_{\Gamma} \mathcal{P}_0\left(\gamma(s)\right) \, \mathrm{d}\gamma.$$

Remark It is common to use the term *homogeneous* to refer to invariance to translation and the term *isotropic* to refer to invariance to rotations. As the general theory we present will not distinguish between the group actions, we will use the term *homogeneous* throughout.

Problem Statement

Thinking of the space of functions as a complex inner-product space, thinking of a sampling pattern as a real-valued function on this vector space, and using the fact that computing the Monte Carlo integral amounts to taking the dot-product of the integrand with the average of delta functions, we can view the problem of estimating variance in Monte Carlo integration as an instance of the following, more general, algebraic problem:

Assume we are given a complex inner-product space $(V, \langle \cdot, \cdot \rangle)$, a group Γ acting on V, and a homogeneous function $\mathcal{P} : V \to \mathbb{R}$. Then, given $w \in V$, compute the expected value and variance of the dot-product of w with the vectors $\mathbf{v} \in V$:

$$\mathsf{E}_{\mathcal{P}}[\langle w, \mathbf{v} \rangle] = \int_V \langle w, v \rangle \cdot \mathcal{P}(v) \, \mathsf{d}v$$

$$\mathsf{Var}_{\mathcal{P}}(\langle w, \mathbf{v} \rangle) = \mathsf{E}_{\mathcal{P}}\left[\|\langle w, \mathbf{v} \rangle\|^2 \right] - \|\mathsf{E}_{\mathcal{P}}[\langle w, \mathbf{v} \rangle]\|^2.$$

The advantage of formulating the problem in this manner is that it makes it easier to leverage representation theory to find a solution. To this end, we review some basic concepts from representation theory in the next section, as well as derive two lemmas describing how the average of the inner-products of vectors behave as we transform one of the arguments by the elements of the group. Using these, we present our closed-form expression for the expected value and variance, given in terms of the Fourier coefficients of the integrand F and the sampling patterns $\mathbf{S}$, in Section 4.

3 Representation Theory

The study of how the Fourier coefficients of a signal change as it is transformed by the elements of a group is best expressed in the language of representation theory. We review some basic concepts from this theory, before deriving the lemmas that lead to a closed-form expression for the expected value and variance of the Monte Carlo integral.

In what follows, we will assume a compact (closed and bounded) Lie group Γ.

Definition Given a complex inner-product space $(V, \langle \cdot, \cdot \rangle)$, we say that the (ρ, V) is a *representation of* Γ if ρ is a group homomorphism from Γ into the group of unitary transformations on V. That is:

$$\rho(\gamma \circ \tilde{\gamma}) = \rho(\gamma) \circ \rho(\tilde{\gamma}), \qquad \forall \, \gamma, \tilde{\gamma} \in \Gamma.$$

Notation Given a representation (ρ, V), a group element $\gamma \in \Gamma$, and a vector $v \in V$, we will write:

$$\gamma(v) := \rho(\gamma)(v).$$

Definition Given a vector space V, the *trivial representation* is the map ρ sending every group element to the identity:

$$\rho(\gamma) = \mathsf{Id.}, \qquad \forall \, \gamma \in \Gamma.$$

Definition Given a representation (ρ, V) and a subspace $W \subset V$, we say that W is a *sub-representation* if $\gamma(w) \in W$ for all $w \in W$ and all $\gamma \in \Gamma$.

Definition We say that (ρ, V) is an *irreducible representation* if the only sub-representations are $W = \{0\}$ and $W = V$.

Remark Since any subspace of a trivial representation is a sub-representation (as the identity maps all vectors to themselves), a trivial representation is irreducible if and only if it is one-dimensional.

Given a representation (ρ, V), Maschke's Theorem [Serre 1977; Fulton and Harris 1991] tells us that we can decompose V as the direct sum of finite-dimensional, irreducible representations:

$$V = \bigoplus_{\lambda \in \Lambda} V^\lambda$$

with V^λ and $V^{\tilde{\lambda}}$ perpendicular whenever $\lambda \neq \tilde{\lambda}$.

Choosing an orthonormal basis $\{b_\lambda^1, \cdots, b_\lambda^{n_\lambda}\}$ for each V^λ allows us to define a Fourier transform:

Definition Given a vector $v \in V$, and indices $\lambda \in \Lambda$ and $m \in [1, \cdots, n_\lambda]$, the (λ, m)-*th Fourier coefficient* of v, denoted $\widehat{v}_\lambda^m$, is the coefficient of v corresponding to the basis vector b_λ^m:

$$\widehat{v}_\lambda^m := \langle v, b_\lambda^m \rangle.$$

Remark Since the basis defining the Fourier coefficients is orthonormal, we can write the inner product of two functions $v, w \in V$ in terms of these coefficients as:

$$\langle v, w \rangle \equiv \sum_{\lambda \in \Lambda} \sum_{m=1}^{n_\lambda} \widehat{v}_\lambda^m \cdot \overline{\widehat{w}_\lambda^m}.$$

Lemma 3.1. *Given an irreducible representation (ρ, V) of a group Γ, for any $x, y, v, w \in V$, we have:*

$$\frac{1}{|\Gamma|} \int_\Gamma \langle \gamma(x), y \rangle \cdot \overline{\langle \gamma(v), w \rangle} \, \mathsf{d}\gamma = \frac{1}{dim(V)} \cdot \langle x, v \rangle \cdot \overline{\langle y, w \rangle}.$$

Corollary 3.2. *In particular, letting $\{b^1, \cdots, b^n\}$ be an orthonormal basis for V, taking $y = b^i$ and $w = b^j$, and fixing $x = v$, the above statement becomes:*

$$\int_\Gamma \widehat{\gamma(v)}^i \cdot \overline{\widehat{\gamma(v)}^j} = \frac{|\Gamma|}{dim(V)} \cdot \|v\|^2 \cdot \delta_{ij}.$$

That is, the Fourier coefficients of $\gamma(v)$, thought of as complex-valued functions on Γ, are orthogonal and the magnitude is independent of which Fourier coefficient we are considering.

Proof Fixing $y, w \in V$, let $B_{y,w} : V \times V \to \mathbb{C}$ be the map:

$$B_{y,w}(x, v) = \int_\Gamma \langle \gamma(x), y \rangle \cdot \overline{\langle \gamma(v), w \rangle} \, \mathsf{d}\gamma.$$

It is not hard to show that this map is linear in the first argument, conjugate-linear in the second, and Γ-equivariant. (That is, for any $\gamma \in \Gamma$ we have $B_{v,w}(\gamma(x), \gamma(y)) = B_{v,w}(x, y)$). Thus, by Schur's Lemma [Serre 1977; Fulton and Harris 1991], $B_{y,w}$ is a scalar multiple of the inner-product on V:

$$B_{y,w}(x, v) = \lambda_{y,w} \cdot \langle x, v \rangle.$$

Noting that $B_{y,w}(x, v) = \overline{B_{x,v}(y, w)}$, it follows that:

$$B_{y,w}(x, v) = \lambda \cdot \langle x, v \rangle \cdot \overline{\langle y, w \rangle},$$

for some constant $\lambda \in \mathbb{C}$ that is independent of v and w.

Thus, we are left with the problem of determining λ. As it is independent of x, y, v, and w, it suffices to determine the value $B_{v,v}(v, v)$ for some $v \neq 0$. More generally, letting $\{b^1, \ldots, b^n\}$ be an orthonormal basis we can get an expression for λ in terms of the integrated square norm of the trace of $\rho(\gamma)$:

$$\int_\Gamma \|\mathsf{Tr}(\rho(\gamma))\|^2 \, \mathsf{d}\gamma = \int_\Gamma \left\| \sum_{i,j=1}^n \langle \gamma(b^i), b^j \rangle \right\|^2 \, \mathsf{d}\gamma$$

$$= \sum_{i=1}^n B_{b^i, b^j}(b^i, b^j)$$

$$= \dim(V) \cdot \lambda.$$

Since the trace is the character of the representation, it follows by the orthogonality of characters [Serre 1977; Fulton and Harris 1991] that $\int_\Gamma \|\mathsf{Tr}(\rho(\gamma))\|^2 \, \mathsf{d}\gamma = |\Gamma|$, giving $\lambda = |\Gamma|/\dim(V)$.

Thus, as desired, we get:

$$\frac{1}{|\Gamma|} \int_\Gamma \langle \gamma(x), y \rangle \cdot \overline{\langle \gamma(v), w \rangle} \, \mathsf{d}\gamma = \frac{1}{\dim(V)} \cdot \langle x, v \rangle \cdot \overline{\langle y, w \rangle}.$$

$\square$

Lemma 3.3. *Leveraging Schur's Lemma in a similar manner, it follows that if (ρ_1, V_1) and (ρ_2, V_2) are two irreducible representations that are not isomorphic, then for any $v_1, w_1 \in V_1$ and $v_2, w_2 \in V_2$:*

$$\int_\Gamma \langle \gamma(v_1), w_1 \rangle \cdot \overline{\langle \gamma(v_2), w_2 \rangle} \, d\gamma = 0.$$

4 Variance Estimation

Using the above theory, we are now prepared to estimate the variance in Monte Carlo integration. We begin by presenting a general expression for the expected value and variance and then consider the specific cases of the torus, the sphere, and Euclidean space.

4.1 General Framework

We assume that we are given a representation (ρ, V) of a compact group Γ, a homogeneous function $\mathcal{P} : V \to \mathbb{R}$ (i.e. $\mathcal{P}(v) = \mathcal{P}(\gamma(v))$ for all $v \in V$ and all $\gamma \in \Gamma$), and a vector $w \in V$. Our goal is to express the expected value and variance of the dot-product of w with the vectors $\mathbf{v} \in V$:

$$\mathsf{E}_\mathcal{P}[\langle w, \mathbf{v} \rangle] = \int_V \langle w, v \rangle \cdot \mathcal{P}(v) \, dv$$

$$\mathsf{Var}_\mathcal{P}(\langle w, \mathbf{v} \rangle) = \mathsf{E}_\mathcal{P}\left[\|\langle w, \mathbf{v} \rangle\|^2\right] - \|\mathsf{E}_\mathcal{P}[\langle w, \mathbf{v} \rangle]\|^2.$$

In deriving the closed form expression for the expected value and variance, we will assume that the decomposition of V into irreducible representations $\{V^\lambda\}_{\lambda \in \Lambda}$ contains the trivial representation. (If it does not, we can take the direct sum of V with a one-dimensional space on which Γ acts trivially.) We will denote this one-dimensional representation as V^0 and let b_0^1 be a unit-vector spanning this space (with $\gamma(b_0^1) = b_0^1$ for all $\gamma \in \Gamma$).

Finally, for simplicity, we will assume that the irreducible representations occur without multiplicity. That is, if $\lambda \neq \widetilde{\lambda}$ then V^λ and $V^{\widetilde{\lambda}}$ are not isomorphic, for all $\lambda, \widetilde{\lambda} \in \Lambda$.

The Expected Value

Using the Fourier coefficients, we can express expected value of the dot-product of w with the vectors $\mathbf{v} \in V$ as:

$$\mathsf{E}_\mathcal{P}[\langle w, \mathbf{v} \rangle] = \int_V \sum_{\lambda \in \Lambda} \sum_{m=1}^{n_\lambda} \widehat{w}_\lambda^m \cdot \overline{\widehat{v}_\lambda^m} \cdot \mathcal{P}(v) \, dv.$$

Using the homogeneity of $\mathcal{P}$, the expected value computed by integrating over V is the same as the expected value computed by integrating over $\gamma(V)$. In particular, we can express the expected value as the average:

$$\mathsf{E}_\mathcal{P}[\langle w, \mathbf{v} \rangle] = \frac{1}{|\Gamma|} \int_\Gamma \int_{\gamma(V)} \sum_{\lambda \in \Lambda} \sum_{m=1}^{n_\lambda} \widehat{w}_\lambda^m \cdot \overline{\widehat{v}_\lambda^m} \cdot \mathcal{P}(v) \, dv \, d\gamma$$

$$= \frac{1}{|\Gamma|} \int_\Gamma \int_V \sum_{\lambda \in \Lambda} \sum_{m=1}^{n_\lambda} \widehat{w}_\lambda^m \cdot \overline{\widehat{\gamma(v)}_\lambda^m} \cdot \mathcal{P}(\gamma(v)) \, dv \, d\gamma$$

$$= \int_V \sum_{\lambda \in \Lambda} \sum_{m=1}^{n_\lambda} \widehat{w}_\lambda^m \frac{1}{|\Gamma|} \left(\int_\Gamma \overline{\langle \gamma(v), b_\lambda^m \rangle} \, d\gamma \right) \cdot \mathcal{P}(v) \, dv.$$

Since b_0^1 is a unit vector on which γ acts as the identity, we have:

$$\mathsf{E}_\mathcal{P}[\langle w, \mathbf{v} \rangle] = \int_V \sum_{\lambda \in \Lambda} \sum_{m=1}^{n_\lambda} \widehat{w}_\lambda^m \cdot$$

$$\cdot \frac{1}{|\Gamma|} \left(\int_\Gamma \overline{\langle \gamma(v), b_\lambda^m \rangle \cdot \overline{\langle \gamma(b_0^1), b_0^1 \rangle}} \, d\gamma \right) \cdot \mathcal{P}(v) \, dv.$$

Using Lemmas 3.1 and 3.3, the facts that V^0 is orthogonal and not isomorphic to V^λ for all $\lambda \neq 0$, and that $\dim(V^0) = 1$, we have:

$$\mathsf{E}_\mathcal{P}[\langle w, \mathbf{v} \rangle] = \sum_{\lambda \in \Lambda} \sum_{m=1}^{n_\lambda} \int_V \widehat{w}_\lambda^m \cdot \overline{\langle v, b_0^1 \rangle} \cdot \langle b_\lambda^m, b_0^1 \rangle \cdot \mathcal{P}(v) \, dv$$

$$= \int_V \widehat{w}_0^1 \cdot \overline{\widehat{v}_0^1} \cdot \mathcal{P}(v) \, dv$$

$$= \widehat{w}_0^1 \cdot \overline{\int_V \widehat{v}_0^1 \cdot \mathcal{P}(v) \, dv}.$$

That is, the expected value of the dot-product is the trivial Fourier coefficient of w times the complex conjugate of the expected value of the trivial Fourier coefficient of the vectors $\mathbf{v} \in V$:

$$\boxed{\mathsf{E}_\mathcal{P}[\langle w, \mathbf{v} \rangle] = \widehat{w}_0^1 \cdot \overline{\mathsf{E}_\mathcal{P}[\widehat{\mathbf{v}}_0^1]}.} \qquad (1)$$

The Variance

Using Equation (1), we can express the variance of the dot-product of w with the vectors $\mathbf{v} \in V$ as:

$$\mathsf{Var}_\mathcal{P}(\langle w, \mathbf{v} \rangle) = \mathsf{E}_\mathcal{P}\left[\|\langle w, \mathbf{v} \rangle\|^2\right] - \left\|\widehat{w}_0^1\right\|^2 \cdot \left\|\mathsf{E}_\mathcal{P}\left[\widehat{\mathbf{v}}_0^1\right]\right\|^2$$

and we are left with the problem of computing $\mathsf{E}_\mathcal{P}\left[\|\langle w, \mathbf{v} \rangle\|^2\right]$.

Expressing the dot-product in terms of the Fourier coefficients gives:

$$\mathsf{E}_\mathcal{P}\left[\|\langle w, \mathbf{v} \rangle\|^2\right] = \int_V \left\| \sum_{\lambda \in \Lambda} \sum_{m=1}^{n_\lambda} \widehat{w}_\lambda^m \cdot \overline{\widehat{v}_\lambda^m} \right\|^2 \cdot \mathcal{P}(v) \, dv$$

$$= \int_V \sum_{\lambda, \widetilde{\lambda} \in \Lambda} \sum_{m=1}^{n_\lambda} \sum_{\widetilde{m}=1}^{n_{\widetilde{\lambda}}} \widehat{w}_\lambda^m \cdot \overline{\widehat{w}_{\widetilde{\lambda}}^{\widetilde{m}}} \cdot \overline{\widehat{v}_\lambda^m} \cdot \widehat{v}_{\widetilde{\lambda}}^{\widetilde{m}} \cdot \mathcal{P}(v) \, dv.$$

As with the expected value, homogeneity implies that we can average the integrals over all $\gamma(V)$, giving:

$$\mathsf{E}_\mathcal{P}\left[\|\langle w, \mathbf{v} \rangle\|^2\right] = \int_V \sum_{\lambda, \widetilde{\lambda} \in \Lambda} \sum_{m=1}^{n_\lambda} \sum_{\widetilde{m}=1}^{n_{\widetilde{\lambda}}} \widehat{w}_\lambda^m \cdot \overline{\widehat{w}_{\widetilde{\lambda}}^{\widetilde{m}}} \cdot$$

$$\cdot \frac{1}{|\Gamma|} \left(\int_\Gamma \overline{\widehat{\gamma(v)}_\lambda^m} \cdot \widehat{\gamma(v)}_{\widetilde{\lambda}}^{\widetilde{m}} \, d\gamma \right) \cdot \mathcal{P}(v) \, dv.$$

Using the fact that $\widehat{\gamma(v)}_\lambda^m = \langle \gamma(v), b_\lambda^m \rangle$ in conjunction with Lemmas 3.1 and 3.3 and letting $\pi_\lambda : V \to V^\lambda$ be the projection from V onto the irreducible representation V^λ, the summation simplifies to:

$$\mathsf{E}_\mathcal{P}\left[\|\langle w, \mathbf{v} \rangle\|^2\right] = \int_V \sum_{\lambda \in \Lambda} \sum_{m=1}^{n_\lambda} \|\widehat{w}_\lambda^m\|^2 \cdot \frac{\|\pi_\lambda(v)\|^2}{\dim(V^\lambda)} \cdot \mathcal{P}(v) \, dv$$

$$= \sum_{\lambda \in \Lambda} \frac{\|\pi_\lambda(w)\|^2 \cdot \mathsf{E}_\mathcal{P}\left[\|\pi_\lambda(\mathbf{v})\|^2\right]}{\dim(V^\lambda)}.$$

This gives a closed-form expression for the variance as:

$$\mathsf{Var}_{\mathcal{P}}(\langle w, \mathbf{v}\rangle) = \sum_{\lambda \in \Lambda \setminus \{0\}} \frac{\|\pi_\lambda(w)\|^2 \cdot \mathsf{E}_{\mathcal{P}}\left[\|\pi_\lambda(\mathbf{v})\|^2\right]}{\dim(V^\lambda)}. \qquad (2)$$

Note that by taking the summation over all irreducible representations *except for the trivial one* we subtract off the square-norm of the expected value.

4.2 The Torus

In this case, the domain of integration and the group of motions are both the d-dimensional torus, $\Omega = \Gamma = [0, 2\pi)^d$, and the representation is defined on the space of complex-valued functions on the torus, $V = L^2(\Omega, \mathbb{C})$, with an element $\gamma \in \Gamma$ acting on a function by translation:

$$[\gamma(F)](p) := F(p - \gamma).$$

The irreducible representations are all one-dimensional (since the group is commutative) and are indexed by points on the d-dimensional integer lattice, $\Lambda = \mathbb{Z}^d$. Specifically, the space V^λ is spanned by a complex exponential with frequency $\lambda \in \mathbb{Z}^d$:

$$V^\lambda = \mathrm{Span}\left\{b_\lambda^0(p) = \frac{e^{i\langle p, \lambda\rangle}}{(2\pi)^{d/2}}\right\}.$$

Thus, Equation (1) gives the expected value of the integral of F as the product of the DC component of F times the complex conjugate of the expected value of the DC component of $\mathbf{S}$. As we are considering Monte Carlo integration, the elements of $\mathbf{S}$ are all the average of N delta-functions, so that:

$$\mathsf{E}_{\mathcal{P}}[\langle F, \mathbf{S}\rangle] = \int_\Omega F(p)\, dp,$$

and the estimate is unbiased.

From Equation (2) the variance in the estimate of the integral can be obtained by taking the power spectrum of F, multiplying (frequency-wise) by the expected power spectrum of $\mathbf{S}$, and summing over all non-zero frequencies:

$$\mathsf{Var}_{\mathcal{P}}(\mathrm{MC}(F, \mathbf{S})) = \sum_{\lambda \in \mathbb{Z}^d \setminus \{0\}} \|\widehat{F}_\lambda\|^2 \cdot \mathsf{E}_{\mathcal{P}}\left[\|\widehat{\mathbf{S}}_\lambda\|^2\right],$$

where $\widehat{F}_l$ is the l-th Fourier coefficient of F.

4.3 The Sphere

In this case, the domain of integration is the 2-sphere, $\Omega = S^2$, the group of motions is the group of rotations in 3D, $\Gamma = SO(3)$, and the representation is defined on the space of complex-valued functions on the sphere, $V = L^2(\Omega, \mathbb{C})$, with an element $\gamma \in \Gamma$ acting on a function by rotation:

$$[\gamma(F)](p) := F\left(\gamma^{-1}(p)\right).$$

In this case, the irreducible representations are indexed by the non-negative integers, $\Lambda = [0, \cdots, \infty)$, and the irreducible representation V^λ is a $(2\lambda + 1)$-dimensional space:

$$V^\lambda = \mathrm{Span}\left\{Y_\lambda^{-\lambda}(\theta, \phi), \cdots, Y_\lambda^\lambda(\theta, \phi)\right\},$$

with $Y_l^m(\theta, \phi)$ the spherical harmonic of frequency l and index m.

As with the torus the integrator is unbiased, and the variance can be computed by summing, over each non-zero spherical frequency, the product of the power of F and the expected power of $\mathbf{S}$ in that frequency, divided by the dimension of the frequency space:

$$\mathsf{Var}_{\mathcal{P}}(\mathrm{MC}(F, \mathbf{S})) = \sum_{l=1}^\infty \frac{\displaystyle\sum_{m=-l}^l \|\widehat{F}_l^m\|^2 \cdot \sum_{m=-l}^l \mathsf{E}_{\mathcal{P}}\left[\|\widehat{\mathbf{S}}_l^m\|^2\right]}{2l+1},$$

where $\widehat{F}_l^m$ is the (l, m)-th spherical harmonic coefficient of F.

4.4 Euclidean Space

In this case, the domain of integration is d-dimensional Euclidean space, $\Omega = \mathbb{R}^d$, the group is the group of Euclidean motions, $\Gamma = SE(d) = \mathbb{R}^d \times SO(d)$, and the representation is defined on the space of complex-valued functions on Euclidean space, $V = L^2(\Omega, \mathbb{C})$, with an element $\gamma = (\tau, \sigma) \in \Gamma$ acting on a function by a combination of translation and rotation:

$$[\gamma(F)](p) := F\left(\sigma^{-1}(p - \tau)\right).$$

Unfortunately, the analysis in Section 3 does not apply to this context because we assumed that the group is compact. None-the-less, we can formally carry over the results, replacing the notion of "dimension" with the the "size" of the irreducible representations.

In this case, the irreducible representations are indexed by the non-negative real numbers, $\Lambda = \mathbb{R}^{\geq 0}$ [Vilenkin 1978] and the space V^λ is the "span" of complex exponentials whose frequency has norm λ:

$$V^\lambda = \mathrm{Span}_{|q|=\lambda}\left\{b_\lambda^q(p) = e^{i\langle p, q\rangle}\right\}.$$

As above, the integrator is unbiased and, using the fact that the size of the λ-th irreducible representation is the size of of the $(d-1)$-dimensional sphere with radius λ, we get:

$$\mathsf{Var}_{\mathcal{P}}(\mathrm{MC}(F, \mathbf{S})) =$$
$$= \int_0^\infty \frac{\displaystyle\int_{|q|=\lambda} \|\widehat{F}_q\|^2\, dq \cdot \int_{|q|=\lambda} \mathsf{E}_{\mathcal{P}}\left[\|\widehat{\mathbf{S}}_q\|^2\right]\, dq}{\lambda^{d-1} \cdot |S^{d-1}|}\, d\lambda -$$
$$- \|\widehat{F}_0\|^2 \cdot \mathsf{E}_{\mathcal{P}}\left[\|\widehat{\mathbf{S}}_0\|^2\right],$$

where $\widehat{F}_p$ is the p-th Fourier coefficient of F.

References

FULTON, W., AND HARRIS, J. 1991. *Representation Theory: A First Course*. Springer-Verlag, New York.

PILLEBOUE, A., SINGH, G., COEURJOLLY, D., KAZHDAN, M., AND OSTROMOUKHOV, V. 2015. Variance analysis for Monte Carlo integration. *Transactions on Graphics (SIGGRAPH)*.

SERRE, J. 1977. *Linear representations of finite groups*. Springer-Verlag, New York.

VILENKIN, N. 1978. *Special Functions and the Theory of Group Representations*. American Mathematical Society.