

The square root rule for adaptive importance sampling

Art B. Owen Yi Zhou
Stanford University

January 2019

Abstract

In adaptive importance sampling, and other contexts, we have unbiased and uncorrelated estimates of a common quantity μ and the variance of the k 'th estimate is thought to decay like k^{-y} for an unknown rate parameter $y \in [0, 1]$. If we combine the estimates as though $y = 1/2$, then the resulting estimate attains the optimal variance rate with a constant that is too large by at most $9/8$ for any $0 \leq y \leq 1$ and any number K of estimates.

1 Introduction

A useful rule for weighting uncorrelated estimates appears in an unpublished technical report (Owen and Zhou, 1999). This paper states and proves that result and discusses some mild generalizations.

The motivating context for the problem is adaptive importance sampling, though there may be other uses. There is an unknown quantity $\mu \in \mathbb{R}$ and we have K unbiased and uncorrelated measurements of it, denoted $\hat{\mu}_k$ for $k = 1, \dots, K$. In adaptive importance sampling, $\hat{\mu}_k$ could be the result of an importance sampler chosen on the basis of the data that produced $\hat{\mu}_1, \dots, \hat{\mu}_{k-1}$. The estimates $\hat{\mu}_k$ are uncorrelated by construction but are not independent. The variance of $\hat{\mu}_k$ is affected by the prior sample values. We have in mind a setting where each $\hat{\mu}_k$ is based on the same number n of function evaluations but we do not need to use that n in our analysis. Our setup could also be reasonable in settings where n_k evaluations are used to construct $\hat{\mu}_k$.

In an adaptive method with one pilot estimate and one final estimate, $K = 2$ and n is usually large. In other settings, each $\hat{\mu}_k$ could be based on just one evaluation of the integrand (e.g., Ryu and Boyd (2014)) and then K would typically be very large. In intermediate settings we might have something like $K = 10$ estimates based on perhaps thousands of evaluations each. For instance, the cross-entropy method (De Boer et al., 2005) might be used this way.

We estimate μ by $\hat{\mu} = \sum_{k=1}^K \omega_k \hat{\mu}_k$ where $\sum_{k=1}^K \omega_k = 1$. Then $\mathbb{E}(\hat{\mu}) = \mu$ and to minimize $\text{Var}(\hat{\mu})$ we should take $\omega_k \propto \text{Var}(\hat{\mu}_k)^{-1}$. The problem is that we do

not know $\text{Var}(\hat{\mu}_k)$. We might have unbiased estimates $\widehat{\text{Var}}(\hat{\mu}_k)$ of $\text{Var}(\hat{\mu}_k)$, but taking $\omega_k \propto \widehat{\text{Var}}(\hat{\mu}_k)^{-1}$ will yield an estimate $\hat{\mu}$ that is no longer unbiased. The bias is potentially serious in rare event problems. When the k 'th sample fails to include the rare event much or at all, then $\hat{\mu}_k$ and $\widehat{\text{Var}}(\hat{\mu}_k)$ are both likely to be small, whereas obtaining more than the expected number of rare events makes it more likely that both are large. That is, we anticipate a negative correlation between $\hat{\mu}_k$ and $\widehat{\text{Var}}(\hat{\mu}_k)^{-1}$. Then the small values of $\hat{\mu}_k$ will be upweighted while the large ones will be downweighted resulting in a downward bias for $\hat{\mu}$. For background on importance sampling for rare events, see L'Ecuyer et al. (2009). The AMIS algorithm of Cornuet et al. (2012) uses a weighted combination of estimates with weights that are correlated with those estimates and that complicates even the task of proving consistency.

Suppose that $\text{Var}(\hat{\mu}_k) \propto k^{-y}$ for some unknown y with $0 \leq y \leq 1$. The value $y = 0$ is a model for importance sampling where adaptation brings no benefit. The value $y = 1$ is a model for a setting where importance sampling is working very well, roughly as well as quasi-Monte Carlo sampling (Dick and Pillichshammer, 2010). The optimal choice is $\omega_k \propto k^y$. Not knowing the true y we might take $\omega_k \propto k^x$ for some $x \in [0, 1]$. The square root rule from the appendix of Owen and Zhou (1999) takes $\omega_k \propto \sqrt{k}$. It never has variance more than 9/8 times that of the unknown best weighting rule when $\text{Var}(\hat{\mu}_k) \propto k^{-y}$ for some $0 \leq y \leq 1$.

This paper is organized as follows. Section 2 presents our notation and some of the context from adaptive importance sampling. Section 3 states and proves our main result using two lemmas. Lemma 2 there has an inequality that we had originally described as ‘straightforward algebra’. Revisiting the problem now we find the inequality to be quite delicate. Section 4 includes some remarks on generalizations to settings where even better convergence rates, including exponential convergence, apply.

2 Notation

Step k of our adaptive importance sampler generates data X_k . We let $Z_k = (X_1, \dots, X_{k-1})$ denote the data from all prior steps, with Z_1 being empty. We assume that our estimates $\hat{\mu}_k$ satisfy

$$\mathbb{E}(\hat{\mu}_k | Z_k) = \mu, \quad \text{and} \quad \text{Var}(\hat{\mu}_k | Z_k) = \sigma_k^2 < \infty, \quad k = 1, \dots, K. \quad (1)$$

In Owen and Zhou (1999), $\hat{\mu}_k$ was an importance sampled estimate of an integral over the unit cube, sampling from a mixture of products of beta distributions whose parameters were tuned to the prior data in Z_k .

We combine the estimates via

$$\hat{\mu} = \sum_{k=1}^K \omega_k \hat{\mu}_k$$

where $\sum_{k=1}^K \omega_k = 1$. Then $\mathbb{E}(\hat{\mu}_k) = \mathbb{E}(\mathbb{E}(\hat{\mu}_k | Z_k)) = \mu$ so $\mathbb{E}(\hat{\mu}) = \mu$. For $\ell > k$,

$$\text{Cov}(\hat{\mu}_k, \hat{\mu}_\ell) = \mathbb{E}(\mathbb{E}((\hat{\mu}_k - \mu)(\hat{\mu}_\ell - \mu) | Z_\ell)) = \mathbb{E}((\hat{\mu}_k - \mu)\mathbb{E}(\hat{\mu}_\ell - \mu | Z_\ell)) = 0.$$

Therefore

$$\text{Var}(\hat{\mu}) = \sum_{k=1}^K \omega_k^2 \text{Var}(\hat{\mu}_k)$$

The conditional variances σ_k^2 are random. If stage k of importance sampling has 2 or more observations in it, then we can ordinarily construct conditionally unbiased variance estimates $\hat{\sigma}_k^2$. That is,

$$\mathbb{E}(\hat{\sigma}_k^2 | Z_k) = \sigma_k^2.$$

Then

$$\mathbb{E}\left(\sum_{k=1}^K \omega_k^2 \hat{\sigma}_k^2\right) = \sum_{k=1}^K \omega_k^2 \mathbb{E}(\hat{\sigma}_k^2) = \sum_{k=1}^K \omega_k^2 \text{Var}(\hat{\mu}_k) = \text{Var}(\hat{\mu}).$$

We can thus combine results from the K stages of the AIS algorithm and get an unbiased estimate of the variance. The underlying idea here is that $D_L = \sum_{k=1}^L \omega_k (\hat{\mu}_k - \mu)$ is a martingale in L (Williams, 1991). We will not make formal use of martingale arguments.

Now we introduce a model

$$\text{Var}(\hat{\mu}_k) = \tau^2 k^{-y} \tag{2}$$

for $0 \leq y \leq 1$ and a constant of proportionality $\tau \in (0, \infty)$. Our estimate is

$$\hat{\mu} = \hat{\mu}(x) = \frac{\sum_{i=1}^K i^x \hat{\mu}_i}{\sum_{i=1}^K i^x}.$$

The best choice is $\hat{\mu}(y)$ but y is unknown. Our variance using x is

$$\text{Var}(\hat{\mu}(x)) = \tau^2 \frac{\sum_{i=1}^K i^{2x-y}}{\left(\sum_{i=1}^K i^x\right)^2}$$

and we measure the inefficiency of our choice by

$$\rho_K(x | y) \equiv \frac{\text{Var}(\hat{\mu}(x))}{\text{Var}(\hat{\mu}(y))} = \frac{(\sum_{i=1}^K i^{2x-y})(\sum_{i=1}^K i^y)}{(\sum_{i=1}^K i^x)^2}. \tag{3}$$

If the variance of $\hat{\mu}_k$ decays as k^{-y} and, knowing that, we use $x = y$, then $\text{Var}(\hat{\mu}) = O(K^{-y-1})$. For the pessimistic value $y = 0$, the variance decays at the usual Monte Carlo rate in the number K of steps. For an optimistic value $y = 1$, the variance decays at the rate $O(K^{-2})$, slightly better than $O(K^{-2+\epsilon})$ (any $\epsilon > 0$) which holds for randomly shifted lattice rules applied to functions of bounded variation in the sense of Hardy and Krause. See L'Ecuyer and Lemieux (2000) for background on randomized lattice rules.

3 Square root rule

Owen and Zhou (1999) proposed to split the difference between the optimistic estimate $\hat{\mu}(1)$ and the pessimistic one $\hat{\mu}(0)$ by using $\hat{\mu}(1/2)$. The result is an unbiased estimate that attains the same convergence rate as the unknown best estimate with only a modest penalty in the constant factor.

Theorem 1. For $\rho_K(x | y)$ given by equation (3),

$$\sup_{1 \leq K < \infty} \sup_{0 \leq y \leq 1} \rho_K\left(\frac{1}{2} | y\right) \leq \frac{9}{8}. \quad (4)$$

If $x \neq 1/2$ and $K \geq 2$ then

$$\sup_{0 \leq y \leq 1} \rho_K(x | y) > \sup_{0 \leq y \leq 1} \rho_K\left(\frac{1}{2} | y\right). \quad (5)$$

Equation (4) shows that the square root rule has at most 12.5% more variance than the unknown best rule. Equation (5) shows that the choice $x = 1/2$ is the unique minimax optimal one, when $K > 1$. When $K = 1$ then $\hat{\mu}(x)$ does not depend on x and in that case, $\hat{\mu}(x) = \hat{\mu}(y)$. Before proving the theorem, we establish two lemmas.

Lemma 1. For $\rho_K(x | y)$ given by equation (3), with $K \geq 1$ and $0 \leq x \leq 1$,

$$\sup_{0 \leq y \leq 1} \rho_K(x | y) = \begin{cases} \rho_K(x | 1), & x \leq 1/2 \\ \rho_K(x | 0), & x \geq 1/2. \end{cases}$$

Proof. The result holds trivially for $K = 1$ because $\rho_1(x | y) = 1$. Now suppose that $K \geq 2$. Then

$$\frac{\partial^2}{\partial y^2} \rho_K(x | y) = C^{-2} \times \sum_{i=1}^K \sum_{j=1}^K i^{2x-y} j^y (\log(i) - \log(j))^2$$

for $C = \sum_{i=1}^K i^x$. Thus $\rho_K(x | y)$ is strictly convex in y for any x when $K \geq 2$, and so $\sup_{0 \leq y \leq 1} \rho_K(x | y) \in \{\rho_K(x | 1), \rho_K(x | 0)\}$.

By symmetry, $\rho_K(x | 2x - y) = \rho_K(x | y)$. So if $x \geq 1/2$, then $\rho_K(x | 0) = \rho_K(x | 2x) \geq \rho_K(x, 1)$. The last step follows because $2x \geq 1$ and $\rho_K(x | y)$ is a convex function of y with its minimum at $y = x \leq 1$. This establishes the result for $x \geq 1/2$ and a similar argument holds for $x \leq 1/2$. For $x = 1/2$, $\rho_K(x | 0) = \rho_K(x | 1)$. \square

We will use the following integral bounds, for integers $K \geq 1$. If $0 \leq x \leq 1$, then

$$\sum_{i=1}^K i^x > \int_{1/2}^{K+1/2} v^x dv = \frac{(K+1/2)^{x+1} - (1/2)^{x+1}}{x+1}, \quad \text{and} \quad (6)$$

$$\sum_{i=1}^K i^x < \int_1^{K+1} v^x dv = \frac{(K+1)^{x+1} - 1}{x+1}. \quad (7)$$

Equation (6) uses concavity of v^x and is much sharper than the bound one gets by integrating over $0 \leq v \leq K$.

Lemma 2. For $\rho_K(x | y)$ given by equation (3), $\rho_{K+1}(1/2 | 1) > \rho_K(1/2 | 1)$ holds for any integer $K \geq 1$.

Proof. Let $S_K = \sum_{i=1}^K i^{1/2}$. Then

$$\begin{aligned} \frac{\rho_{K+1}}{\rho_K} &= \frac{(K+1)(K+2)S_K^2}{K^2(S_K + \sqrt{K+1})^2} = \frac{(K+1)(K+2)}{K^2\left(1 + \frac{\sqrt{K+1}}{S_K}\right)^2} \\ &> \frac{(K+1)(K+2)}{K^2\left(1 + \frac{\sqrt{K+1}}{((K+1/2)^{3/2} - 2^{-3/2})/(3/2)}\right)^2} \\ &= \frac{(K+1)(K+2)}{\left(K + \frac{3}{2} \frac{K}{f(K) - 2^{-3/2}/\sqrt{K+1}}\right)^2}, \end{aligned}$$

for

$$f(K) = (K + 1/2)^{3/2} / \sqrt{K+1}.$$

Now $g(K) = f(K)/(K + 1/4)$ has $g'(K) < 0$ and $\lim_{K \rightarrow \infty} g(K) = 1$. Therefore, if $K > 7$,

$$\frac{\rho_{K+1}}{\rho_K} > \frac{(K+1)(K+2)}{\left(K + \frac{3}{2} \frac{K}{K+1/4 - 2^{-3/2}/\sqrt{8}}\right)^2} = \frac{(K+1)(K+2)}{\left(K + \frac{3}{2} \frac{K}{K+1/8}\right)^2}. \quad (8)$$

The numerator in (8) is $K^2 + 3K + 2$ while the denominator is

$$K^2 + 3K \frac{K}{K+1/8} + \frac{9}{4} \frac{K^2}{(K+1/8)^2}.$$

The numerator is larger than the denominator, establishing the theorem for $K > 7$. For $K \leq 7$ we compute that $\rho_{K+1}(1/2 | 1)/\rho_K(1/2 | 1) > 1.0038$ for $1 \leq K \leq 7$. \square

Proof of Theorem 1. Applying Lemmas 1 and 2,

$$\sup_{1 \leq K < \infty} \sup_{0 \leq y \leq 1} \rho_K\left(\frac{1}{2} | y\right) = \lim_{K \rightarrow \infty} \rho_K\left(\frac{1}{2} | 1\right) = \lim_{K \rightarrow \infty} \frac{K^2(K+1)/2}{\left(\sum_{i=1}^K i^{1/2}\right)^2}.$$

Using the integral bounds (6) and (7) we can show that the above limit is $9/8$, establishing equation (4). Next, if $x < 1/2$, then

$$\sup_{0 \leq y \leq 1} \rho_K(x | y) = \rho_K(x | 1) = \frac{\left(\sum_{i=1}^K i^{2x-1}\right)\left(\sum_{i=1}^K i\right)}{\left(\sum_{i=1}^K i^x\right)^2} = \frac{K(K-1)}{2} \frac{\left(\sum_{i=1}^K i^{2x-1}\right)}{\left(\sum_{i=1}^K i^x\right)^2}.$$

Ignoring the factor $K(K-1)$ and taking the derivative yields

$$\frac{\sum_{i=1}^K \sum_{j=1}^K i^x j^{2x-1} (\log(j) - \log(i))}{(\sum_{i=1}^K i^{2x-1}) (\sum_{i=1}^K i^x)}. \quad (9)$$

The numerator in (9) is of the form $\sum_{\ell=1}^K \eta_\ell \log(\ell)$ where $\sum_{\ell=1}^K \eta_\ell = 0$. Furthermore $\eta_\ell = \sum_{i=1}^K (i^x \ell^{2x-1} - \ell^x i^{2x-1})$ is a decreasing function of ℓ while $\log(\ell)$ is increasing and so the expression in (9) is negative. Therefore $\rho_K(x | 1)$ is a decreasing function of x making $\rho_K(x | 1) > \rho_K(1/2 | 1)$ for $x < 1/2$. This establishes (5) for $x < 1/2$ and the case of $x > 1/2$ is similar. \square

4 Generalization

We can generalize some aspects of the square root rule to other power laws. Suppose that the true rate parameter y is known to satisfy $L \leq y \leq U$ for $0 \leq L \leq U < \infty$. We can then work with $\omega_k \propto k^x$ for $x = M \equiv (L+U)/2$. First we recall the definition of ρ_K from (3):

$$\rho_K(x | y) = \frac{(\sum_{i=1}^K i^{2x-y}) (\sum_{i=1}^K i^y)}{(\sum_{i=1}^K i^x)^2}.$$

Then $2M - U = L$ and

$$\rho_K(M | U) = \frac{(\sum_{i=1}^K i^L) (\sum_{i=1}^K i^U)}{(\sum_{i=1}^K i^M)^2}. \quad (10)$$

Lemma 3. For $\rho_K(x | y)$ given by equation (3), with $K \geq 1$ and $0 \leq L \leq x \leq U < \infty$,

$$\sup_{L \leq y \leq U} \rho_K(x | y) = \begin{cases} \rho_K(x | U), & x \leq M \\ \rho_K(x | L), & x \geq M, \end{cases}$$

for $M = (L+U)/2$.

Proof. The proof is the same as for Lemma 1 because $\rho_K(x | y)$ is convex in y for any $K \geq 1$ and any x and $\rho_K(x | 2x - y) = \rho_K(x | y)$. \square

The inequalities in Lemma 2 are rather delicate and we have not extended them to the more general setting. Equation (10) is easy to evaluate for integer values of L , M and U . For more general values, some sharper tools than the integral bounds in this paper are given by Burrows and Talbot (1984) who make a detailed study of sums of powers of the first K natural numbers. We do see numerically that $\rho_K(M | U)$ is nondecreasing with K in every instance we have inspected. We can easily find the asymptotic inefficiency

$$\lim_{K \rightarrow \infty} \rho_K(M | U) = \lim_{K \rightarrow \infty} \frac{\frac{K^L}{L+1} \frac{K^U}{U+1}}{\left(\frac{K^M}{M+1}\right)^2} = \frac{(M+1)^2}{(L+1)(U+1)}.$$

In cases with $L = 0$ and hence $M = U/2$ we get $(U/2 + 1)^2/(U + 1)$. For instance, an upper bound at the rate $y = U = 2$ corresponding roughly to asymptotic accuracy of scrambled net integration (Owen, 1997) leads to $x = M = 1$ and an asymptotic inefficiency of at most

$$\frac{(1 + 1)^2}{(0 + 1)(2 + 1)} = \frac{4}{3}.$$

There are problems in which the adaptive importance sampling error converges exponentially to zero. See for instance Kollman et al. (1999) as well as Kong and Spanier (2011). These examples involve particle transport problems through complicated media. It is reasonable to expect that each estimate $\hat{\mu}_k$ will require a large number of observations and that K will then be not too large.

Suppose that $\text{Var}(\hat{\mu}_k) = \tau^2 \exp(-yk)k$ for some $y > 0$. Then the desired combination is

$$\hat{\mu}(y) = \frac{\sum_{i=1}^K \exp(iy) \hat{\mu}_i}{\sum_{i=1}^K \exp(iy)}.$$

Not knowing y we use

$$\hat{\mu}(x) = \frac{\sum_{i=1}^K \exp(ix) \hat{\mu}_i}{\sum_{i=1}^K \exp(ix)},$$

for some $x > 0$. If $y \neq x/2$, then our inefficiency is

$$\gamma_K(x | y) \equiv \frac{\sum_{i=1}^K \exp((2x - y)i) \sum_{i=1}^K \exp(yi)}{\left(\sum_{i=1}^K \exp(xi)\right)^2} = \frac{\left(\frac{e^{2x-y}(e^{K(2x-y)} - 1)}{e^{2x-y} - 1}\right) \left(\frac{e^y(e^{Ky} - 1)}{e^y - 1}\right)}{\left(\frac{e^x(e^{Kx} - 1)}{e^x - 1}\right)^2}.$$

If $y = x/2$ then the first factor in the numerator is K . It can be disastrously inefficient to use $x \ll y$. Some safety is obtained in the $x \rightarrow \infty$ limit where $\hat{\mu} = \mu_K$. That is, in that limit, one simply uses the final and presumably best estimate. Then

$$\lim_{x \rightarrow \infty} \gamma_K(x | y) = \frac{e^y(1 - e^{-Ky})}{e^y - 1} \leq \frac{e^y}{e^y - 1}.$$

For instance, if the variance is halving at each iteration, then $y = \log(2)$ and then taking $x \rightarrow \infty$ is inefficient by at most a factor of 2. Repeated ten-fold variance reductions correspond to $y = \log(10)$ and a limiting inefficiency of at most $10/9 \doteq 1.11$. The greatest inefficiency from using only the final iteration arises in the limit $y \rightarrow 0$ where the factor is K . In this setting, the user is not getting a meaningful exponential convergence and even there the loss factor is at most K and, as remarked above, that K is not likely to be large.

Acknowledgments

This work was supported by the US NSF under grants DMS-1521145 and IIS-1837931.

References

- Burrows, B. L. and Talbot, R. F. (1984). Sums of powers of integers. *The American Mathematical Monthly*, 91(7):394–403.
- Cornuet, J., Marin, J.-M., Mira, A., and Robert, C. P. (2012). Adaptive multiple importance sampling. *Scandinavian Journal of Statistics*, 39(4):798–812.
- De Boer, P.-T., Kroese, D. P., Mannor, S., and Rubinstein, R. Y. (2005). A tutorial on the cross-entropy method. *Annals of operations research*, 134(1):19–67.
- Dick, J. and Pillichshammer, F. (2010). *Digital sequences, discrepancy and quasi-Monte Carlo integration*. Cambridge University Press, Cambridge.
- Kollman, C., Baggerly, K., Cox, D., and Picard, R. (1999). Adaptive importance sampling on discrete Markov chains. *Annals of Applied Probability*, pages 391–412.
- Kong, R. and Spanier, J. (2011). Geometric convergence of adaptive Monte Carlo algorithms for radiative transport problems based on importance sampling methods. *Nuclear Science and Engineering*, 168(3):197–225.
- L’Ecuyer, P. and Lemieux, C. (2000). Variance reduction via lattice rules. *Management Science*, 46(9):1214–1235.
- L’Ecuyer, P., Mandjes, M., and Tuffin, B. (2009). Importance sampling and rare event simulation. In Rubino, G. and Tuffin, B., editors, *Rare event simulation using Monte Carlo methods*, pages 17–38. John Wiley & Sons, Chichester, UK.
- Owen, A. B. (1997). Scrambled net variance for integrals of smooth functions. *Annals of Statistics*, 25(4):1541–1562.
- Owen, A. B. and Zhou, Y. (1999). Adaptive importance sampling by mixtures of products of beta distributions. Technical report, Stanford University.
- Ryu, E. K. and Boyd, S. P. (2014). Adaptive importance sampling via stochastic convex programming. Technical report, arXiv:1412.4845.
- Williams, D. (1991). *Probability with martingales*. Cambridge University Press, Cambridge.