

Importance Sampling for multi-constraints rare event probability.

Virgile Caron

Telecom ParisTech, 37-39 rue Dareau, 75014 Paris
virgile.caron@telecom-paristech.fr

January 14, 2014

Abstract

Improving Importance Sampling estimators for rare event probabilities requires sharp approximations of the optimal density leading to a nearly zero-variance estimator. This paper presents a new way to handle the estimation of the probability of a rare event defined as a finite intersection of subset. We provide a sharp approximation of the density of long runs of a random walk conditioned by multiples constraints, each of them defined by an average of a function of its summands as their number tends to infinity.

1 Introduction and context

In this paper, we consider efficient estimation of the probability of large deviations of a multivariate sum of independent, identically distributed, light-tailed and non-lattice random vectors. The probability to be estimated is defined as an intersection of s event where each of them is characterized by a sum of n i.i.d. random variable belonging to some countable union of intervals.

Consider $\mathbf{X}_1^n := (\mathbf{X}_1, \dots, \mathbf{X}_n)$ n i.i.d. random vectors with known common density $p_{\mathbf{X}}$ on \mathbb{R}^d , $d \geq 1$, copies of $\mathbf{X} := (\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(d)})$. The superscript (j) pertains to the coordinate of a vector and the subscript i pertains to replications. Consider also u a measurable function defined from \mathbb{R}^d to \mathbb{R}^s with $s < n$. The last condition on s forbid to have a zero-one probability event. Define $\mathbf{U} := u(\mathbf{X})$ with density $p_{\mathbf{U}}$ and

$$\mathbf{U}_{1,n} := \sum_{i=1}^n \mathbf{U}_i.$$

We intend to estimate for large but fixed n

$$P_n := P(\mathbf{U}_{1,n} \in nA) \tag{1}$$

where A is a non-empty measurable set of \mathbb{R}^s such as $E[u(\mathbf{X})] \notin A$.

For sake of clarity, the probability to be estimate can be written as follows

$$P_n = P \left[\bigcap_{j=1}^s \left\{ \sum_{i=1}^n u^{(j)}(\mathbf{X}_i^{(1)}, \dots, \mathbf{X}_i^{(d)}) \in nA^{(j)} \right\} \right] \tag{2}$$

where, for $j \in \{1, \dots, s\}$, $u^{(j)}$ is a measurable function from \mathbb{R}^d to \mathbb{R} and $A^{(j)}$ is a countable union of intervals of \mathbb{R} . According to the context, we will use either (1) or (2).

In [3], the authors consider in details the case where $d = s = 1$ and $A := A_n = (a_n, \infty)$ with a_n a convergent sequence.

The basic estimate of P_n is defined as follows: generate L i.i.d. samples $X_1^n(l)$ with underlying density $p_{\mathbf{X}}$ and define

$$\widetilde{P}_n := \frac{1}{L} \sum_{l=1}^L \mathbb{1}_{\mathcal{E}_n}(X_1^n(l))$$

where

$$\mathcal{E}_n := \left\{ (x_1, \dots, x_n) \in (\mathbb{R}^d)^n : (u(x_1) + \dots + u(x_n)) \in nA \right\}. \quad (3)$$

The Importance Sampling estimator of P_n with sampling density g on $(\mathbb{R}^d)^n$ is

$$\widehat{P}_n := \frac{1}{L} \sum_{l=1}^L \hat{P}_n(l) \mathbf{1}_{\mathcal{E}_n}(Y_1^n(l)) \quad (4)$$

where $\hat{P}_n(l)$ is called "importance factor" and can be written

$$\hat{P}_n(l) := \frac{\prod_{i=1}^n p_{\mathbf{X}}(Y_i(l))}{g(Y_1^n(l))} \quad (5)$$

and where the L samples $Y_1^n(l) := (Y_1(l), \dots, Y_n(l))$ are i.i.d. with common density g ; the coordinates of $Y_1^n(l)$ however need not be i.i.d.. It is known that the optimal choice for g is the density of $\mathbf{X}_1^n := (\mathbf{X}_1, \dots, \mathbf{X}_n)$ conditioned upon $(\mathbf{X}_1^n \in \mathcal{E}_n)$, leading to a zero variance estimator. We refer to [5] for the background of this section.

The state-independent IS scheme for rare event estimation (see [6] or [12]), rests on two basic ingredients: the sampling distribution is fitted to the so-called dominating point (which is the point where the quantity to be estimated is mostly captured; see [11]) of the set to be measured; independent and identically distributed replications under this sampling distribution are performed. More recently, a state-dependent algorithm leading to a strongly efficient estimator is provided by [2] when $d = 1$, $u(x) = x$ and $A = (a; \infty)$ (or, more generally in \mathbb{R}^d , with a smooth boundary and a unique dominating point). Indeed, adaptive tilting defines a sampling density for the i -th r.v. in the run which depends both on the target event $(\mathbf{U}_{1,n} \in nA)$ and on the current state of the path up to step $i - 1$. Jointly with an ad hoc stopping rule controlling the excursion of the current state of the path, this algorithm provides an estimate of P_n with a coefficient of variation independent upon n . This result shows that nearly optimal estimators can be obtained without approximating the conditional density.

The main issue of the method described above is to find dominating point. However, when the dimension of the set A increases, finding a dominating point can be very tricky or even impossible. A solution will be to divide the set under consideration into smaller subset and, for each one of this subset, find a dominating point. Doing so makes the implementation of an IS scheme harder and harder as the dimension increases.

Our proposal is somehow different since it is based on a sharp approximation result of the conditional density of long runs. The approximation holds for any point conditioning of the form $(\mathbf{U}_{1,n} = nv)$. Then sampling v in A according to the distribution of $\mathbf{U}_{1,n}$ conditioned upon $(\mathbf{U}_{1,n} \in nA)$ produces the estimator. By its very definition this procedure does not make use of any dominating point, since it randomly explores the set A . Indeed, our proposal hints on two choices: first do not make use of the notion of dominating point and explore all the target set instead (no part of the set A is neglected); secondly, do not use i.i.d. replications, but merely sample long runs of variables under a proxy of the optimal sampling scheme.

We will propose an IS sampling density which approximates this conditional density very sharply on its first components y_1, \dots, y_k where $k = k_n$ is very large, namely $k/n \rightarrow 1$. However, but in the Gaussian case, k should satisfy $(n - k) \rightarrow \infty$ by the very construction of the approximation. The IS density on $(\mathbb{R}^d)^n$ is obtained multiplying this proxy by a product of a much simpler state-independent IS scheme following [13].

The paper is organized as follows. Section 2 is devoted to notations and hypothesis. In Section 3, we expose the approximation scheme for the conditional density of \mathbf{X}_1^k under $(\mathbf{U}_{1,n} = nv)$. In Section 4, we introduce our IS scheme. Simulated results are presented in Section 5 which enlighten the gain of the present approach over state-dependent Importance Sampling schemes.

We rely on [7] where the basic approximation (and proofs) used in the present paper can be found. The real case is studied in [4] and applications for IS estimators can be found in [3].

2 Notations and hypotheses

We consider approximations of the density of the vector \mathbf{X}_1^k on $(\mathbb{R}^d)^k$, when the conditioning event writes (2) and $k := k_n$ is such that

$$\lim_{n \rightarrow \infty} \frac{k}{n} = 1 \quad (\text{K1})$$

$$\lim_{n \rightarrow \infty} (n - k) = +\infty. \quad (\text{K2})$$

Therefore we may consider the asymptotic behavior of the density of the random walk on long runs.

Throughout the paper the value of a density $p_{\mathbf{Z}}$ of some continuous random vector \mathbf{Z} at point z may be written $p_{\mathbf{Z}}(z)$ or $p(\mathbf{Z} = z)$, which may prove more convenient according to the context.

Let p_{nv} denote the density of \mathbf{X}_1^k under the local condition $(\mathbf{U}_{1,n} = nv)$

$$p_{nv}(\mathbf{X}_1^k = Y_1^k) := p(\mathbf{X}_1^k = Y_1^k | \mathbf{U}_{1,n} = nv) \quad (6)$$

where Y_1^k belongs to $(\mathbb{R}^d)^k$ and v fixed in \mathbb{R}^s .

We will also consider the density p_{nA} of \mathbf{X}_1^k conditioned upon $(\mathbf{U}_{1,n} \in nA)$

$$p_{nA}(\mathbf{X}_1^k = Y_1^k) := p(\mathbf{X}_1^k = Y_1^k | \mathbf{U}_{1,n} \in nA). \quad (7)$$

The approximating density of p_{nv} is denoted g_{nv} ; the corresponding approximation of p_{nA} is denoted g_{nA} . Explicit formulas for those densities are presented in the next section.

3 Multivariate random walk under a local conditionning event.

Let ϵ_n be a postive sequence such as

$$\lim_{n \rightarrow \infty} \epsilon_n^2 (n - k) = \infty \quad (\text{E1})$$

$$\lim_{n \rightarrow \infty} \epsilon_n (\log n)^2 = 0 \quad (\text{E2})$$

It will be shown that $\epsilon_n (\log n)^2$ is the rate of accuracy of the approximating scheme.

We assume that $\mathbf{U} := u(\mathbf{X})$ has a density $p_{\mathbf{U}}$ (with p.m. $P_{\mathbf{U}}$) absolutely continuous with respect to Lebesgue measure on \mathbb{R}^s . Futhermore, we assume that u is such that the characteristic function of \mathbf{U} belongs to L^r for some $r \geq 1$.

Denote $\underline{0}$ is the vector of \mathbb{R}^s with all coordinates equal to 0 and $V(\underline{0})$ a neighborhood of $\underline{0}$.

We assume that \mathbf{U} satisfy the Cramer condition, meaning

$$\Phi_{\mathbf{U}}(t) := E[\exp \langle t, \mathbf{U} \rangle] < \infty, \quad t \in V(\underline{0}) \subset \mathbb{R}^s.$$

and denote

$$m(t) := {}^t\nabla \log(\Phi_{\mathbf{U}}(t)), \quad t \in V(\underline{0}) \subset \mathbb{R}^s$$

and

$$\kappa(t) := {}^t\nabla \nabla \log(\Phi_{\mathbf{U}}(t)), \quad t \in V(\underline{0}) \subset \mathbb{R}^s.$$

as the mean and the covariance matrix of the tilted density defined by

$$\pi_u^\alpha(x) := \frac{\exp \langle t, u(x) \rangle}{\Phi_{\mathbf{U}}(t)} p_{\mathbf{X}}(x). \quad (8)$$

where t is the only solution of $m(t) = \alpha$ for α in the conxev hull of $P_{\mathbf{U}}$. Conditions on $\Phi_{\mathbf{U}}(t)$ which ensure existence and uniqueness of t are referred to stepness properties (see [1], p153 ff, for all properties of moment generating funtion used in this paper).

Let $v \in \mathbb{R}^s$. We now state the general form of the approximating density. Denote

$$g_0(y_1 | y_0) := \pi_u^v(y_1) \quad (9)$$

with an arbitrary y_0 and π_u^v defined in (8).

For $1 \leq i \leq k-1$, we recursively define $g(y_{i+1}|y_1^i)$. Set $t_i \in \mathbb{R}^s$ to be the unique solution to the equation

$$m(t_i) = m_{i,n} := \frac{n}{n-i} \left(v - \frac{u_{1,i}}{n} \right) \quad (10)$$

where $u_{1,i} = u(y_1) + \dots + u(y_i)$.

Denote

$$\kappa_{(i,n)}^{j,l} := \frac{d^2}{dt^{(j)} dt^{(l)}} \left(\log E_{\pi_{\mathbf{U}}^{m_{i,n}}} \exp \langle t, \mathbf{U} \rangle \right) (0)$$

and

$$\kappa_{(i,n)}^{j,l,m} := \frac{d^3}{dt^{(j)} dt^{(l)} dt^{(m)}} \left(\log E_{\pi_{\mathbf{U}}^{m_{i,n}}} \exp \langle t, \mathbf{U} \rangle \right) (0).$$

for j, l and m in $\{1, \dots, s\}$

Denote

$$g(y_{i+1}|y_1^i) := C_i \mathbf{n}_s(u(y_{i+1}); \beta\alpha + v, \beta) p_{\mathbf{X}}(y_{i+1}) \quad (11)$$

where C_i is a normalizing factor, $\mathbf{n}_s(u(y_{i+1}); \beta\alpha + v, \beta)$ is the normal density at $u(y_{i+1})$ with mean $\beta\alpha + v$ and covariance matrix β with α and β defined by

$$\alpha := \left(t_i + \frac{\kappa_{(i,n)}^{-2} \gamma}{2(n-i-1)} \right) \beta := \kappa_{(i,n)}(n-i-1)$$

where γ is defined by

$$\gamma := \left(\sum_{j=1}^s \kappa_{(i,n)}^{j,j,p} \right)_{1 \leq p \leq s}.$$

Then

$$g_{nv}(y_1^k) := g_0(y_1|y_0) \prod_{i=1}^{k-1} g(y_{i+1}|y_1^i) \quad (12)$$

Theorem 1 Assume (E1), (E2), (K1) and (K2).

- Let Y_1^k a sample of P_{nv} . Then

$$p\left(\mathbf{X}_1^k = Y_1^k | \mathbf{U}_{1,n} = nv\right) = g_{nv}(Y_1^k)(1 + o_{P_{nv}}(1 + \epsilon_n(\log n)^2)) \quad (13)$$

- Let Y_1^k a sample of G_{nv} . Then

$$p\left(\mathbf{X}_1^k = Y_1^k | \mathbf{U}_{1,n} = nv\right) = g_{nv}(Y_1^k)(1 + o_{G_{nv}}(1 + \epsilon_n(\log n)^2)) \quad (14)$$

Remark 2 The approximation of the density of \mathbf{X}_1^k is not performed on the sequence of entire spaces $(\mathbb{R}^s)^k$ but merely on a sequence of subsets of $(\mathbb{R}^s)^k$ which contains the trajectories of the conditioned random walk with probability going to 1 as n tends to infinity. The approximation is performed on typical paths. For sake of applications in Importance Sampling, Theorem 14 is exactly what we need. Nevertheless, as proved in [7], the extension of our results from typical paths to the whole space \mathbb{R}^k holds: convergence of the relative error on large sets imply that the total variation distance between the conditioned measure and its approximation goes to 0 on the entire space.

Remark 3 The rule which defines the value of k for a given accuracy of the approximation is stated in Section 5 of [7].

Remark 4 When the \mathbf{X}_i 's are i.i.d. Gaussian standard and $u(x) = x$, the result of the approximation theorem are true for $k = n - 1$ without the error term. Indeed, it holds $p(\mathbf{X}_1^{n-1} = x_1^{n-1} | \mathbf{U}_{1,n} = nv) = g_{nv}(x_1^{n-1})$ for all x_1^{n-1} in $(\mathbb{R}^d)^{n-1}$.

As stated above the optimal choice for the sampling density is p_{nA} . It holds

$$p_{nA}(x_1^k) = \int_A p_{nv}(\mathbf{X}_1^k = x_1^k) p(\mathbf{U}_{1,n}/n = v | \mathbf{U}_{1,n} \in nA) dv \quad (15)$$

so that, in contrast with [2] or [6], we do not consider the dominating point approach but merely realize a sharp approximation of the integrand at any point of A and consider the dominating contribution of all those distributions in the evaluation of the conditional density p_{nA} .

4 Adaptive IS Estimator for rare event probability

The IS scheme produces samples $Y := (Y_1, \dots, Y_k)$ distributed under g_{nA} , which is a continuous mixture of densities g_{nv} as in (12) with $p(\mathbf{U}_{1,n}/n = v | \mathbf{U}_{1,n} \in nA)$.

Simulation of samples $\mathbf{U}_{1,n}/n$ under this density can be performed through Metropolis-Hastings algorithm, since

$$r(v, v') := \frac{p(\mathbf{U}_{1,n}/n = v | \mathbf{U}_{1,n} \in nA)}{p(\mathbf{U}_{1,n}/n = v' | \mathbf{U}_{1,n} \in nA)}$$

turns out to be independent upon $P(\mathbf{U}_{1,n} \in nA)$. The proposal distribution of the algorithm should be supported by A .

The density g_{nA} is extended from $(\mathbb{R}^s)^k$ onto $(\mathbb{R}^s)^n$ completing the $n - k$ remaining coordinates with i.i.d. copies of r.v.'s Y_{k+1}, \dots, Y_n with common tilted density

$$g_{nA}(y_{k+1}^n | y_1^k) := \prod_{i=k+1}^n \pi_u^{m_k}(y_i) \quad (16)$$

with $m_k := m(t_k) = \frac{n}{n-k} \left(v - \frac{u_{1,k}}{n} \right)$ and

$$u_{1,k} = \sum_{i=1}^k u(y_i).$$

The last $n - k$ r.v.'s \mathbf{Y}_i 's are therefore drawn according to the state independent i.i.d. scheme in phase with Sadowsky and Bucklew [13].

We now define our IS estimator of P_n . Let $Y_1^n(l) := Y_1(l), \dots, Y_n(l)$ be generated under g_{nA} . Let

$$\widehat{P}_n(l) := \frac{\prod_{i=0}^n p_{\mathbf{X}}(Y_i(l))}{g_{nA}(Y_1^n(l))} \mathbb{1}_{\mathcal{E}_n}(Y_1^n(l)) \quad (17)$$

and define

$$\widehat{P}_n := \frac{1}{L} \sum_{l=1}^L \widehat{P}_n(l). \quad (18)$$

in accordance with (4).

Remark 5 In the real case and for $A = (a, \infty)$, the authors of [3] shows that under regularity conditions the resulting relative error of the estimator is proportional to $\sqrt{n - k_n}$ and drops by a factor $\sqrt{n - k_n}/\sqrt{n}$ with respect to the state independent IS scheme. Slight modification in the extension of g_{nA} allows to prove the strong efficiency of the estimator (18) using arguments from both [2] and [3]; see [8].

5 When the dimension becomes very high

This section compares the performance of the present approach with respect to the standard tilted one using i.i.d. replications under (8). Consider a random sample X_1, \dots, X_{100} where X_1 has a normal distribution $N(0.05, 1)$ and let

$$\mathcal{E}_{100} := \left\{ x_1^{100} : \frac{|x_1 + \dots + x_{100}|}{100} > 0.28 \right\}.$$

This example is in the same vein as the one developed in [9] or in [10]. Under the present proposal the distribution of the Importance Factor concentrates around P_{100} ; hence so-called "rogue path phenomenon" (see [9]) does not occur.

We explore the gain in relative accuracy when the dimension of the measured set increases. Let therefore $B := (\mathcal{E}_{100})^d$ which is the d -cartesian product of \mathcal{E}_{100} . The 100 r.v.'s X_i 's are i.i.d. random vectors in \mathbb{R}^d with common i.i.d. $N(0.05, 1)$ distribution. The dominating point has all coordinates equal 0.28. Rogue path curse produces an overwhelming loss in accuracy, imposing a very large increase in runtime to get reasonable results. Our interest is to show that in this simple dissymmetric case our proposal provides a good estimate, while the standard IS scheme ignores a part of the event B . The standard i.i.d. IS scheme introduces the dominating point $a = 0.28$ and the family of i.i.d. tilted r.v.'s with common $N(a, 1)$ distribution for each coordinates. It can be seen that a large part of B is never visited through the procedure, inducing a bias in the estimation.

This example is not as artificial as it may seem; indeed it leads to a 2^d dominating points situation which is quite often met in real life. Exploring at random the set of interest avoids any search for dominating points. Drawing L i.i.d. points v_1, \dots, v_L according to the distribution of $\mathbf{U}_{1,100}/100$ conditionally upon B we evaluate P_{100} with $k = 99$; note that in the Gaussian case Theorem 1 provides an exact description of the conditional density of X_1^k for all k between 1 and n . The following figure shows the gain in relative accuracy w.r.t. the state independent IS scheme according to the growth of d . The value of P_{100} is 10^{-2d} .




Figure 1: Relative Accuracy of the adaptive estimate (dotted line) w.r.t. i.i.d. twisted one (solid line) as a function of the dimension d for $L = 1000$.

6 Conclusion

In this paper, we explore a new way to estimate multi-constraints large deviation probability. In future work, the author will investigate the theoretical behaviour of the relative error of our proposed estimator.

References

- [1] Barndorff-Nielsen, OE (1978) Information and Exponential Families in Statistical Theory. Wiley, New-York.

- [2] Blanchet JH, Glynn PW, Leder K. (2009) Efficient simulation of light-tailed sums: an old-folk song sung to a faster new tune.... In P. L'Ecuyer and A.B. Owen, editors, Monte Carlo and quasi-Monte Carlo methods. Springer-Verlag, Berlin, pp 227-248.
- [3] Broniatowski M, Caron V, (2013) Towards zero variance estimators for rare event probabilities. ACM TOMACS:Special Issue on Monte Carlo Methods in Statistics 23(1), (article 7).
- [4] Broniatowski M, Caron V (2014) Long runs under a conditional limit distribution. Ann. Appl. Probab. To appear.
- [5] Bucklew JA (2004) Introduction to rare event simulation. Springer Series in Statistics, Springer-Verlag, New York.
- [6] Bucklew JA, Ney P, Sadowsky JS (1990) Monte Carlo simulation and large deviations theory for uniformly recurrent markov chains. J. Appl. Probab. 27(11):49-61.
- [7] Caron V (2013). Approximation of a multivariate conditional density. arxiv:0880506.
- [8] Caron V, Guyader A, Munoz Zuniga M, Tuffin B (2013) Some recent results in rare event estimation. To appear in ESAIM proceedings.
- [9] Dupuis P, Wang H (2004) Importance sampling, large deviations, and differential games. Stoch. Stoch. Rep., 76:481-508.
- [10] Glasserman P, Wang Y (1997) Counterexamples in importance sampling for large deviations probabilities. Ann. of Appl. Probab. 7(3):731-746.
- [11] Ney P (1983) Dominating points and the asymptotics of large deviations for random walk on \mathbb{R}^d . Ann. Probab., 11(1):158-167.
- [12] Sadowsky JS (1996) On Monte-Carlo estimation of large deviations probabilities. Ann. Appl. Probab., 9(2):493-503.
- [13] Sadowsky JS, Bucklew JA (1990) On large deviations theory and asymptotically efficient Monte Carlo estimation. IEEE Trans. Inform. Theo. 36(3):579-588.