Proceedings of the 2022 Winter Simulation Conference
B. Feng, G. Pedrielli, Y. Peng, S. Shashaani, E. Song, C.G. Corlu, L.H. Lee, E.P. Chew, T. Roeder, and P. Lendermann, eds.

RARE-EVENT SIMULATION WITHOUT VARIANCE REDUCTION: AN EXTREME VALUE THEORY APPROACH

Yuanlu Bai

Industrial Engineering and Operations Research
Columbia University
500 West 120th Street
New York, NY 10027, USA

ABSTRACT

In estimating probabilities of rare events, crude Monte Carlo (MC) simulation is inefficient which motivates the use of variance reduction techniques. However, these latter schemes rely heavily on delicate analyses of underlying simulation models, which are not always easy or even possible. We propose the use of extreme value analysis, in particular the peak-over-threshold (POT) method which is popularly employed for extremal estimation of real datasets, in the simulation setting. More specifically, we view crude MC samples as data to fit on a generalized Pareto distribution. We test this idea on several numerical examples. The results show that our POT estimator appears more accurate than crude MC and, while crude MC can easily give a trivial probability estimate 0, POT outputs a non-trivial estimate with a roughly correct magnitude. Therefore, in the absence of efficient variance reduction schemes, POT appears to offer potential benefits to enhance crude MC estimates.

1 INTRODUCTION

A major goal of rare-event simulation is to estimate tiny probabilities that are triggered by rare but catastrophic events. In using Monte Carlo (MC) to estimate rare-event probabilities, a main challenge is that, by its own nature, the target rare events seldom occur in the simulation experiments. Since sufficient hits on the target events are required to achieve meaningfully accurate estimation, this makes crude MC computationally costly as the required simulation size to attain enough accuracy becomes enormous.

To address the inefficiency of crude MC, a range of variance reduction techniques have been developed. Despite demonstratably powerful in many problems, in order to attain good performances, these techniques often rely on tractable problem structures that allow careful algorithmic design. Unfortunately, this requirement could be difficult or even impossible to meet in complex practical applications. Thus, the main goal of this paper is to study *an approach to improve upon the efficiency of crude MC in the absence of variance reduction schemes*.

More specifically, we resort to extreme value theory (EVT) (Embrechts et al. 1997), which has been a prominent approach in extreme event analysis for real data. One major approach, which we would borrow here, is the peak-over-threshold (POT) method. This method is based on the Pickands-Balkema-de Haan theorem which states that, under suitable assumptions, the distribution function of the so-called excess loss above a threshold converges to the generalized Pareto distribution (GPD) as this threshold increases. In this sense, GPD is a justified model for tail data fitting. Analogically, we propose applying the POT method on the crude MC simulation data. Although the POT method is well-established, to our best knowledge, POT has not been considered in the rare-event simulation context.

2 METHODOLOGY

We consider a simulation model that outputs a random vector $X \in \mathbb{R}^d$ under probability measure P, and we are interested in estimating $p = P(X \in E)$ where $E \subset \mathbb{R}^d$ is a rare-event set, i.e., p is a tiny number. Suppose we generate n i.i.d. simulation samples $X_1, \dots, X_n$. Then the crude MC estimator is simply

$$\hat{p}_n^{MC} = \frac{1}{n}\sum_{i=1}^{n} I(X_i \in E).$$

In this rare-event setting, we would like to control the discrepancy between $\hat{p}_n^{MC}$ and p, *relative* to the magnitude of p itself, i.e., the probability $P(|\hat{p}_n^{MC} - p| > \delta p)$ for some fixed $\delta < 1$. By Chebyshev's inequality, $n \geq 1/(\varepsilon\delta^2 p)$ guarantees that $P(|\hat{p}_n^{MC} - p| > \delta p) \leq \varepsilon$ for any $\varepsilon > 0$. This reveals that, when p is tiny, the required simulation size which is reciprocal in p can be enormous. Therefore, as widely known in the literature, crude MC is inefficient for rare-event simulation.

As discussed in the introduction, we are interested in scenarios where crude MC is the only available simulation format. In this case, we study POT to obtain a better estimate from crude MC without overwhelming the computational effort. We assume that the target rare event $\{X \in E\}$ can be formulated as $\{f(X) \geq a\}$ where f is a real-valued function and a is a constant. Under this setting, we estimate the target probability p with the following procedure:

1. Generate i.i.d. simulation samples $X_1, \dots, X_n$ and compute $Y_i = f(X_i), i = 1, \dots, n$;
2. Pick a threshold $u < a$ following certain criterion;
3. Fit a GPD G with the excess data $\{Y_i - u: 1 \leq i \leq n, Y_i > u\}$;
4. Output

$$\hat{p}_n^{POT} = \left(\frac{1}{n}\sum_{i=1}^{n} I(Y_i > u)\right)\big(1 - G(a - u)\big).$$

3 NUMERICAL EXPERIMENTS

We apply this method to several numerical examples and compare the results with crude MC, including examples where variance reduction techniques cannot be applied easily. We find that, in a suitably wide parameter range, our POT estimator achieves smaller variance than crude MC. Moreover, with limited simulation samples, while crude MC often outputs a trivial estimate of 0, POT can output an estimate of a roughly correct magnitude. In Table 1, we show some results for the sample mean large deviations example with Pareto distribution. The ground truth is $p \approx 1.00 \times 10^{-7}$.

Table 1: Statistics of each estimator in Example 1 with Pareto(1, 2) distribution and a = 1000.

Method	Mean	Std	Min	P25	P50	P75	Max
MC	9.70e-8	3.22e-7	0	0	0	0	2.00e-6
POT	1.12e-7	7.22e-8	8.17e-8	6.10e-8	9.41e-8	1.45e-7	6.42e-7

4 CONCLUSION

POT outperforms crude MC in a reasonably wide spectrum of problems. It performs especially well when the simulation size is relatively large, the rare event boundary is not far from the threshold, and the tail of the distribution is heavy. Compared to crude MC, the POT estimator usually has smaller standard deviation with equal simulation size, and it gives an estimate of a roughly correct magnitude instead of a trivial estimate 0. Therefore, if efficient variance reduction techniques are not available, then our POT procedure can be used to refine and improve the crude MC estimate.

REFERENCES

Embrechts, P., T. Mikosch, and C. Kl¨uppelberg. 1997. Modelling Extremal Events: for Insurance and Finance. New York: Springer-Verlag.