

RARE-EVENT SIMULATION WITHOUT STRUCTURAL INFORMATION: A LEARNING-BASED APPROACH

Zhiyuan Huang

Department of IOE
University of Michigan
1205 Beal Avenue
Ann Arbor, MI 48105, USA

Henry Lam

Department of IEOR
Columbia University
500 W. 120th Street
New York, NY 10027, USA

Ding Zhao

Department of Mechanical Engineering
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, USA

ABSTRACT

Importance sampling has been extensively studied as a variance reduction tool in rare-event simulation. The design and efficiency of this method often relies on structural knowledge about the target problem in hand. In this paper, we consider a simple rare-event setting, driven by Gaussian input variates, where there is no direct information about the system of interest, or where the system of interest can only be evaluated through a “black-box” model. This setting mimics, on a basic level, situations where the system of interest could be too complex to analyze as in the conventional importance sampling literature. We investigate an approach based on two-stage sampling, where the first stage learns the rare-event set of interest, and the second stage uses a change of measure that exploits the first-stage finding. We present some guarantees and discuss comparisons of this scheme with more naive approaches.

1 INTRODUCTION

Rare-event simulation concerns the estimation of a probability that is very small, which usually serves as a risk measurement of the occurrence of catastrophic events. Because of the small target probability, crude Monte Carlo typically bears a huge relative error, and consequently requires a huge number of samples to obtain a reasonably accurate estimate. To overcome this issue, importance sampling (IS) is commonly suggested as an effective variance reduction method in such contexts (early references include, e.g., Siegmund 1976; Glynn and Iglehart 1989; Sadowsky and Bucklew 1990 etc.).

To design an efficient IS estimator, existing approaches often involve analyzing the structures or dynamics of the system of interest. A large literature studies the use of large deviations to invent and mathematically justify IS procedures (see, e.g., the surveys Bucklew 2013; Asmussen and Glynn 2007; Rubinstein and Kroese 2016; Glasserman 2013; Juneja and Shahabuddin 2006; Blanchet and Lam 2012). This line of studies requires, in particular, thorough knowledge about the system or the underlying function at hand. Another line of studies use techniques such as Markov-chain Monte Carlo (MCMC) to sample from the rare-event set of interest, or approximately from the conditional distribution given the occurrence of the rare event (Botev et al. 2013; Grace et al. 2014; Chan and Kroese 2012).

In this paper, we study a setting where we do not have direct information about the underlying performance function. This happens when, e.g., the function is only accessible as a “black-box”, so that it can be revealed only at given input values chosen by the user. This situation, on a basic level, mimics the setting where the function may be too complex to mathematically analyze for designing IS schemes. It is motivated by several recent applications of IS in self-driving cars evaluation (Zhao et al. 2017; Huang et al. 2017a), neural network training (Bengio and Sen  cal 2008), and reinforcement learning (Frank, Mannor, and Precup 2008; Shelton 2001). This is in addition to more established applications in, e.g., finance (Glasserman et al. 1999), queueing (Dupuis et al. 2007; Dieker and Mandjes 2006; Blanchet et al. 2009), structural reliability (Au and Beck 1999; Dubourg et al. 2013; Grabaskas et al. 2016; Heidelberger 1995) and manufacturing (Kanj et al. 2006).

As a motivating example, suppose we have a black-box prediction model $\hat{y}(x)$ (built from, e.g., a machine learning model) for the response of a physical system $y(x)$ that is unknown and needs to be observed. Here x is the input that is random. Think, in the setting of self-driving car, that x encodes the surroundings and $y(x)$ is the movement of the surrounding objects in the next time unit. The action of the self-driving car relies on the prediction model, and there would be an accident or a serious loss if the prediction deviates from the physical outcome too much. Here, we can use $g(x) = d(\hat{y}(x), y(x))$ to capture this deviation, where $d(\cdot, \cdot)$ denotes some distance measure. The probability $P(g(x) \geq \gamma)$ captures the risk of a serious loss. Since a typical machine learning output $\hat{y}(x)$ can be highly unstructured and reminiscent of a black box, using IS to speed up the estimation of such type of probability falls into our study setup.

This paper focuses on a simple example of the above setting. In particular, we assume the input variates are Gaussian and that the rare-event set has one “dominating point” (Sadowsky and Bucklew 1990; Dieker and Mandjes 2006; Blanchet and Li 2011; Adler et al. 2012). These assumptions could be possibly relaxed, e.g., to Gaussian mixtures that can be used to model or approximate a wider class of distributions, and to rare-event sets that have multiple dominant points. We propose a two-stage sampling scheme in the above situations. In the first stage, we sample from the space using a space-filling design. This stage serves to learn information about the rare-event set, in particular the dominant point. In the second stage, we construct IS via an exponential tilting based on the estimated dominant point. From the viewpoint of sequential learning, one can think of the first stage as the “exploration” phase and the second stage as “exploitation”.

We present some theoretical error bounds on this approach. Our analysis, in particular, shows the impact of using exponential tilting based on a point that only approximates the dominating point in terms of their distance, which is in general exponential. However, there is a concentration-type behavior in our first-stage learning on the dominant point, which offsets this exponential dependence. One of the insights is that, under the particular setting we consider, the two-stage approach is advantageous against more naive approaches, such as pure uniform sampling, if the sample space is large compared to the location of the rare-event set.

Our study here is perhaps closest to the setting of cross-entropy method (Rubinstein and Kroese 2004, De Boer et al. 2005), which can be used when there is no information about the underlying structure of the rare event. This approach relies on a sequential use of sample average approximation to minimize the Kullback-Leibler divergence between an IS distribution and the zero-variance distribution to update the IS parameters. Part of our motivation is that the cross-entropy method, though demonstrated to be very effective, typically requires an explicit identification of a threshold parameter controlling the rarity level that can be sequentially adjusted. In contrast, our method does not need this parameter, and instead we rely on a concentration-type exploration of the rare-event set in the first stage that informs our IS. Furthermore, our study is also related to a line of IS work based on learning the rare-event set, for instance, through first or second order approximation (Zhao and Ono 1999), quadratic response surface (Das and Zheng 2000), neural network (Papadrakakis and Lagaros 2002), support vector machine (Hurtado 2004; Bourinet et al. 2011; Huang et al. 2018), kriging or Gaussian process (Echard et al. 2011; OHagan 2006; Dubourg et al. 2013), and the utilization of monotonic features (Huang et al. 2017b). These methods are all heuristically

built on the use of data to approximate the rare-event set, and there has been no theoretical guarantees on their performances. Our study here could provide some basic mathematical insights on the errors and their impacts on Monte Carlo from such type of learning. Lastly, a rare-event setting where the system is partially known and learned is theoretically analyzed in Lam (2012), but only confined to the setting of level-crossing of random walks.

The remainder of this paper is as follows. Section 2 reviews the IS technique including in particular the Gaussian input case. Section 3 presents our precise setting and the proposed procedure. Section 4 shows numerical results and discusses some extensions.

2 BASIC BACKGROUND

Suppose that we are interested in estimating the probability of hitting a set \mathcal{R} , i.e., $p = P(X \in \mathcal{R})$. IS replaces the original probability measure P with \tilde{P} , such that the likelihood ratio $P(X)/\tilde{P}(X)$, given by the Radon-Nikodym derivative between $P(\cdot)$ and $\tilde{P}(\cdot)$, is well defined on the set \mathcal{R} . To adjust for the bias of this change, IS outputs

$$Z = I(X \in \mathcal{R}) \frac{P(X)}{\tilde{P}(X)}. \quad (1)$$

To measure the efficiency of an IS scheme, we introduce a rarity parameter, say n , that parametrizes the rare-event probability p_n such that $p_n \rightarrow 0$ as $n \rightarrow \infty$. Note that since the probability of interest is small, one should focus on the relative error of the Monte Carlo estimator with respect to the magnitude of this probability. To this end, we call an IS estimator Z_n for p_n asymptotically efficient if

$$\lim_{n \rightarrow \infty} \frac{\log \tilde{E}[Z_n^2]}{\log \tilde{E}[Z_n]} = 2. \quad (2)$$

The notion (2) is equivalent to saying that $\tilde{E}[Z_n^2]/\tilde{E}[Z_n]^2$ is at most polynomially growing in n . This ensures that the second moment, or the variance, does not explode exponentially relative to the probability of interest as n increases, thus preventing an exponentially large number of simulation replications to achieve a given relative accuracy. For further details, see, e.g., Asmussen and Glynn (2007).

As an example that will be used in this paper, consider a standard Gaussian random vector $X \sim N(0, I)$ and a rare-event set \mathcal{R} . We recall the notion of a dominating point a^* of \mathcal{R} , i.e., a point $a^* \in \mathcal{R}$ such that $\mathcal{R} \subseteq \{x : a^{*'}(x - a^*) \geq 0\}$. Suppose that \mathcal{R} is convex, then this condition is equivalent to saying that a^* optimizes $\min_{x \in \mathcal{R}} \|x\|^2$, as $2a^{*'}(x - a^*) = \nabla(\|x\|^2)'(x - a^*) \geq 0$ is precisely the first order condition of optimality for the convex program $\min_{x \in \mathcal{R}} \|x\|^2$. Moreover, since we consider standard Gaussian X , we have $a^* = \operatorname{argmax}\{\phi(x) : x \in \mathcal{R}\}$, where ϕ denotes the density of standard Gaussian distribution.

Suppose that a^* is big, so that $a^* \rightarrow \infty$ and \mathcal{R} becomes rare. A commonly considered IS scheme is to use exponential tilting, triggered by a mean shift of X from 0 to a^* , i.e., we use the estimator

$$Z = I(X \in \mathcal{R}) \frac{\phi(X)}{\phi(X - a^*)}. \quad (3)$$

The second moment is given by

$$\tilde{E}[Z^2] = \tilde{E} \left[\left(\frac{\phi(X)}{\phi(X - a^*)} \right)^2 ; X \in \mathcal{R} \right] = \tilde{E} \left[e^{-\|X\|^2 + \|X - a^*\|^2} ; X \in \mathcal{R} \right] = e^{-\|a^*\|^2} \tilde{E} \left[e^{-2a^{*'}(X - a^*)} ; X \in \mathcal{R} \right].$$

Since a^* is assumed a dominating point, we have $a^{*'}(x - a^*) \geq 0$ for any $x \in \mathcal{R}$, and thus the above is bounded from above by $e^{-\|a^*\|^2}$.

Since $\tilde{E}[Z]^2 = P(X \in \mathcal{R}) = \operatorname{poly}(a^*)e^{-\|a^*\|^2}$, where $\operatorname{poly}(a)$ denotes a polynomial in a^* , the scheme (3) is asymptotically efficient as defined by (2).

3 IMPORTANCE SAMPLING FOR UNKNOWN RARE-EVENT SETS

Motivated by the discussion and example given in the introduction, we consider the situation where the rare-event set is unknown. To do so, we refine our setting in Section 2. We assume \mathcal{R} is unknown, but, that we know \mathcal{R} is convex, so that the rare-event set has an unknown dominating point a^* . Moreover, assume it is known that $\mathcal{R} \subset [0, B]^m$ for some $B > 0$ (these could be the physical boundary of the given problem). Given a point $x \in \mathbb{R}^m$, we have the capability to observe whether $x \in \mathcal{R}$, but this action would incur a cost that is equivalent to running a simulation copy of $I(X \in \mathcal{R})$.

The basic challenge is to attempt to learn \mathcal{R} so that we can then apply IS effectively. Given a budget n on the number of evaluations, we investigate a two-stage IS scheme as follows. In the first stage, we use a space filling design (e.g. uniform sampling) over the design space $[0, B]^m$ and collect n_1 observations $(X_i, I(X_i \in \mathcal{R}))_{i=1}^{n_1}$. In the second stage, we construct an IS distribution by shifting the mean of the standard Gaussian distribution to $\hat{a} = \operatorname{argmin}_{X_i, i=1, \dots, n_1} \{\|X_i\|^2 : X_i \in \mathcal{R}\}$, or equivalently $\operatorname{argmax}_{X_i, i=1, \dots, n_1} \{\phi(X_i) : X_i \in \mathcal{R}\}$. In other words, each IS output in the second stage is

$$Z = I(X \in \mathcal{R}) \frac{\phi(X)}{\phi(X - \hat{a})}. \quad (4)$$

We use n_2 simulation replication in the second stage to obtain a sample mean as our final estimate of $P(X \in \mathcal{R})$. In the case that \hat{a} does not exist, i.e., none of the first-stage samples fall into \mathcal{R} , we perform a “back-up” procedure in the second stage that has a second moment of, say, M (this can be, e.g., a uniform sampling on the whole space). For convenience of analysis, we use $n_1 = n_2$. We summarize our procedure in Algorithm 1.

Algorithm 1: The two-stage IS scheme for unknown rare-event sets.

Input: Sample budget n_1, n_2 .

Output: IS estimator Z .

Stage 1:

- 1 Use space filling design (e.g. uniform sampling) over the design space $[0, B]^m$ to collect n_1 observations $(X_i, I(X_i \in \mathcal{R}))_{i=1}^{n_1}$;
- 2 Find approximate dominating point by solving $\hat{a} = \operatorname{argmin}_{X_i, i=1, \dots, n_1} \{\|X_i\|^2 : X_i \in \mathcal{R}\}$;

Stage 2:

- 3 **If** \hat{a} exists **do**
 - 4 Obtain the IS estimator as the sample average of n_2 copies of $Z = I(X \in \mathcal{R}) \frac{\phi(X)}{\phi(X - \hat{a})}$;
 - 5 **Else do**
 - 6 A back-up procedure (e.g. use space filling design to collect n_2 samples and output the associated IS estimate);
 - 7 **End;**
-

3.1 Error When Using an Approximate Dominating Point

The second-stage IS is precisely the classical scheme introduced in Section 2, except that we replace the unknown a^* with an estimator \hat{a} . The following approximates the error from such an estimation:

Theorem 1 (Error from using approximate dominating point) Consider estimating $P(X \in \mathcal{R})$ where $\mathcal{R} \in \mathbb{R}^m$ is a convex set and $X \sim N(0, I)$. Denote a^* as a dominating point of \mathcal{R} , and $d = \|\hat{a} - a^*\|$ for a given point \hat{a} . Suppose we use the IS scheme (4). Then its second moment satisfies $\tilde{E}[Z^2] \leq e^{-\|a^*\|^2 + d^2}$.

Proof. Consider the second moment of Z given by

$$\begin{aligned}
 & \tilde{E}[Z^2] \\
 &= \tilde{E}[e^{-\|X\|^2 + \|X - \hat{a}\|^2}; X \in \mathcal{R}] \\
 &= e^{-\|a^*\|^2} \tilde{E}[e^{\|\hat{a}\|^2 + \|a^*\|^2 - 2X'\hat{a}}; X \in \mathcal{R}] \\
 &= e^{-\|a^*\|^2 + \|\hat{a} - a^*\|^2} \tilde{E}[e^{-2(X - a^*)'a^* - 2(X - a^*)'(\hat{a} - a^*)}; X \in \mathcal{R}] \\
 &\leq e^{-\|a^*\|^2 + \|\hat{a} - a^*\|^2} \tilde{E}[e^{-2(X - a^*)'(\hat{a} - a^*)}; X \in \mathcal{R}] \\
 &\quad \text{since } (x - a^*)'a^* \geq 0 \text{ for any } x \in \mathcal{R} \text{ by the definition of dominating point} \\
 &\leq e^{-\|a^*\|^2 + \|\hat{a} - a^*\|^2} \tilde{E}[e^{-2(X - a^*)'(\hat{a} - a^*)}] \\
 &= e^{-\|a^*\|^2 + \|\hat{a} - a^*\|^2}
 \end{aligned}$$

where the last step follows by noting that $X - a^* \sim N(\hat{a} - a^*, I)$ under \tilde{P} , and $\tilde{E}[e^{-2(X - a^*)'(\hat{a} - a^*)}]$ is the moment generating function of $X - a^*$ at $-2(\hat{a} - a^*)$, so that $\tilde{E}[e^{-2(\hat{a} - a^*)'(X - a^*)}] = e^{-2\|\hat{a} - a^*\|^2 + 4\|\hat{a} - a^*\|^2/2} = 1$. This concludes the theorem. \square

Note that Theorem 1 does not assume how \hat{a} is obtained. It is applicable for any \hat{a} . An implication of the theorem is that the IS scheme (4) using \hat{a} is asymptotically efficient, except for an extra multiplicative error e^{d^2} in the ratio $\tilde{E}[Z^2]/\tilde{E}[Z]^2$ that, depending on how close \hat{a} is to a^* , could potentially be significant.

3.2 Efficiency of the Two-stage Procedure

Intuitively, if we use large enough n_1 in the first stage, we could get a good estimate of a^* , i.e., $d = \|\hat{a} - a^*\|$ is small probabilistically. To further analyze this error, let us denote

$$\alpha = \frac{\text{Area}(\mathcal{R})}{\text{Area}([0, B]^m)}$$

We recall that the first-stage observations of X are sampled uniformly on $[0, B]^m$ and used to obtain \hat{a} . Consider $P(\|\hat{a} - a^*\|^2 > y)$. Note that $\|x\|^2 = \|x - a^*\|^2 + \|a^*\|^2 + 2(x - a^*)'a^* \geq \|x - a^*\|^2 + \|a^*\|^2 = \|x - a^*\|^2 + \eta^2$ for any $x \in \mathcal{R}$, where we denote $\eta = \|a^*\|$. We have

$$\begin{aligned}
 P(\|\hat{a} - a^*\|^2 > y) &\leq P(\|\hat{a}\|^2 > y + \eta^2) \text{ since } \hat{a} \in \mathcal{R}, \text{ if exists} \\
 &= P(\|X_i\|^2 > y + \eta^2 \text{ or } X_i \notin \mathcal{R})^{n_1} \\
 &\quad \text{by our construction of } \hat{a} \text{ as the minimizer of } \|X_i\|^2 \text{ among } X_i \in \mathcal{R} \\
 &= (1 - \alpha P(\|X_i\|^2 - \eta^2 \leq y | X_i \in \mathcal{R}))^{n_1} \tag{5}
 \end{aligned}$$

To proceed, we make several simplifying assumptions. We let $m = 2$. Moreover, we consider an additional assumption that

$$\frac{\text{Area}(\{x \in \mathcal{R} : \|x\|^2 - \eta^2 \leq y_1\})}{\text{Area}(\{x \in \mathcal{R} : \|x\|^2 - \eta^2 \leq y_2\})} \geq k \frac{y_1}{y_2}$$

for $y_1 < y_2$ running from 0 to a constant $c > 0$, where k is a constant such that $0 < k \leq 1$. This assumption holds for $m = 2$ if the rare-event set roughly looks like a cone with rays generated from a^* . This includes, in particular, shapes in the form of a semi-disk.

Under the above assumptions, from (5), we have

$$P(\|\hat{a} - a^*\|^2 > y) \leq \left(1 - \alpha k \frac{y}{c}\right)^{n_1}$$

Consider the second moment of the overall estimator of the two-stage procedure, which is $e^{-\|\alpha^*\|^2} E[e^{d^2}]$ where $E[\cdot]$ is under the first-stage uniform sampling. Here we define $e^{d^2} = e^{\|\alpha^*\|^2} M$ if \hat{a} and hence d does not exist (Recall that if \hat{a} does not exist, then we use a back-up procedure giving a second moment M in each second-stage replication). Now, denote $\bar{F}(y) = P(d^2 > y)$, and consider

$$\begin{aligned} E[e^{d^2}] &= - \int_0^c e^y d\bar{F}(y) + (1 - \alpha)^{n_1} e^{\eta^2} M \\ &\leq 1 + (1 - \alpha k)^{n_1} (e^{\eta^2} M - e^c) + \int_0^c \left(1 - \alpha k \frac{y}{c}\right)^{n_1} e^y dy \end{aligned} \tag{6}$$

Consider the term

$$\begin{aligned} \int_0^c \left(1 - \alpha k \frac{y}{c}\right)^{n_1} e^y dy &= c \int_0^1 (1 - \alpha kz)^{n_1} e^{cz} dz \\ &= c \int_0^1 e^{n_1 \log(1 - \alpha kz) + cz} dz \\ &\leq c \int_0^1 e^{-bn_1 \alpha kz + cz} dz \end{aligned} \tag{7}$$

for some constant $b > 0$, where the last step follows since we can find such a b such that $\log(1 - \alpha kz) \leq -b\alpha kz$ for all $0 \leq z \leq 1$.

Now consider $n_1 = \omega(1/\alpha)$, or more precisely $n_1 = \rho/\alpha$ where $\rho \rightarrow \infty$. Then (7) is equal to

$$c \int_0^1 e^{-bk\rho z + cz} dz = \frac{c}{bk\rho - c} (1 - e^{-(bk\rho - c)})$$

Continuing with the same choice of $n_1 = \rho/\alpha$, the second term in (6) becomes

$$(1 - \alpha k)^{\rho/\alpha} (e^{\eta^2} M - e^c)$$

Thus, if we use $n_2 = n_1 = \rho/\alpha$ replications in the second stage, we get an overall second moment of the final estimator bounded from above by

$$\frac{e^{-\eta^2} \frac{c}{bk\rho - c} (1 - e^{-(bk\rho - c)}) + (1 - \alpha k)^{\rho/\alpha} (e^{\eta^2} M - e^c) + 1}{\rho/\alpha} \tag{8}$$

3.3 Comparison with Pure Uniform Sampling

To give some insights on (8), we compare our approach with a naive uniform sampling scheme. Since we have assumed that $\mathcal{R} \subset [0, B]^m$, using uniform sampling as IS is feasible. To compare fairly with our two-stage procedure, suppose we use a sample size $n = n_1 + n_2$ for this uniform IS. The second moment for each replication is given by

$$\tilde{E} \left[\left(\frac{\phi(X)}{1/V} \right)^2; X \in \mathcal{R} \right]$$

where $V = \text{Area}([0, B]^m)$. This is equal to

$$V^2 \tilde{E} [\phi(X)^2; X \in \mathcal{R}] = V \int_{x \in \mathcal{R}} \phi(x)^2 dx \approx \frac{V}{\eta^2} e^{-\eta^2} \tag{9}$$

if we let $m = 2$.

For further comparison, we set $n_1 = n_2 = \rho/\alpha$, so that the total replication size is $n = 2\rho/\alpha$. The second moment becomes approximately

$$\frac{V\alpha}{2\eta^2\rho}e^{-\eta^2} \tag{10}$$

On the other hand, suppose that, in our two-stage procedure, the back-up algorithm is a pure uniform sampling, then M is approximately replaced by (9), so that the overall second moment is approximately bounded by

$$e^{-\eta^2} \frac{\frac{c}{bk\rho-c}(1 - e^{-(bk\rho-c)}) + (1 - \alpha k)^{\rho/\alpha}(V/\eta^2 - e^c) + 1}{\rho/\alpha} \tag{11}$$

Now consider the case where α is bounded away from 1, i.e., the rare-event set does not occupy the whole space, and that $c > 0$ is fixed, i.e., the volume of the rare-event set does not grow (This can be relaxed to letting c grow slowly). Since $\frac{c}{bk\rho-c}(1 - e^{-(bk\rho-c)})$ is bounded, and $(1 - \alpha k)^{1/\alpha} \leq e^{-l}$ for some constant $l > 0$, (11) is bounded by

$$e^{-\eta^2} \frac{\alpha(f + e^{-l\rho}V/\eta^2)}{\rho} \tag{12}$$

for some constant $f > 0$.

Thus, comparing (10) and (12), we see that when V/η^2 is large or goes to ∞ , the second moment provided by the two-stage procedure is smaller, and thus it is preferable. Note that V/η^2 occurs when the space in consideration is large compared to the position of the rare-event set. Intuitively, in this situation most of the uniformly sampled points will miss the rare-event set. Thus a pure uniform sampling works poorly. On the other hand, the two-stage approach makes use of a concentration-type behavior coming from solving the optimization based on the few hitting samples in the first stage, in order to locate the dominating point. This operation makes the overall performance of the two-stage approach better.

4 NUMERICAL EXPERIMENTS

This section uses some simple problems to illustrate the performance of our proposed approach and to compare with uniform IS scheme. We test two types of problems. In the first type of problems, we assume that the rare-event set is convex. In this case, the assumption for Theorem 1 holds and the proposed approach has the analyzed efficiency. The second type of problems considers non-convex rare-event set that violates our assumption, which serves to illustrate potential extensions of our proposed approach.

4.1 Convex Rare-event Set Example

We consider the following problem setting. The design vector x has dimension $m = 2$ and the design space is $[0, 5]^m$, i.e. $B = 5$ according to the notation we used in previous sections. Note that X follows standard Gaussian distribution. We set the rare-event set to be $\mathcal{R} = \{x : (x_1 - 5)^2 + (x_2 - 5)^2 \leq 4\}$. We compare the performance of the proposed two-stage approach and the uniform IS in two sets of experiments.

In the first set of experiments, we vary the number of total sample budget from $n = 1000$ to $n = 50,000$. For each selection of n , we use $n_1 = n_2 = n/2$ samples in the two-stage approach. Results of these experiments are shown in Figures 1 and 2. Figure 1 presents the final probability estimates from the two approaches with different sample budgets. We observe that the two-stage approach (red solid line) is slightly more stable than the uniform IS (black dash line), as the two-stage approach has less fluctuation when the budget is greater than 30,000. This observation is confirmed in Figure 2, where the confidence interval half-width of the two-stage approach (red solid line) is shown to be always better than the uniform IS (black dash line). Roughly speaking, the two-stage approach appears to require only half of the budget of the uniform IS for achieving the same level of accuracy.

We specifically look at the case where $n = 50,000$. Figure 3 shows a trajectory of the confidence interval half-width of the two approaches in a single sample path. The blue horizontal line shows the final




Figure 1: Probability estimates under different sample budgets.




Figure 2: Confidence interval half-widths under different sample budgets.

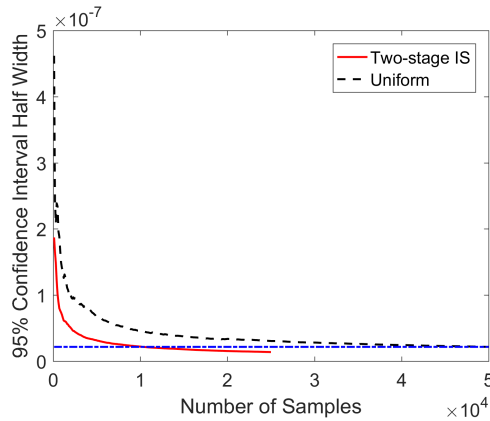


Figure 3: Confidence interval half-width in the IS stage when $n = 50,000$. The blue horizontal line shows the final confidence interval half-width of the uniform IS.

half-width of the uniform IS for easy comparison. It indicates that at roughly 10,000 samples in the second stage (35,000 samples in total), the two-stage approach obtains the same half-width as the uniform IS using all 50,000 samples. Overall, the two-stage approach performs visually better compared to the uniform IS.

Our second set of experiments serves to illustrate the analysis in Section 3.3. We consider a similar problem setting, but with a growing design space size (so that V/η^2 increases). For the design space $[0, B]^2$, we have B vary from $B = 4$ to $B = 7$ with 0.1 increments. We consider the rare-event set $\mathcal{R} = \{x : (x_1 - B)^2 + (x_2 - B)^2 \leq 4\}$ that has a stationary size as B changes. For each value of B , we use $n = 50,000$ samples for both approaches, where the two-stage approach uses $n_1 = n_2 = n/2$. Figure 4 shows that as the value of B increases, the probability estimate decays exponentially, and the two approaches give very similar values. However, in Figure 5, we observe that the width of confidence interval from the two-stage approach (red solid line) is smaller than the uniform approach (black dash line) for all B values. Figure 6 further presents the ratio between the width of confidence interval for the two-stage approach and the uniform approach. The ratio decays roughly linearly, which indicates that as the size of the design space increases, the two-stage approach has more advantages.

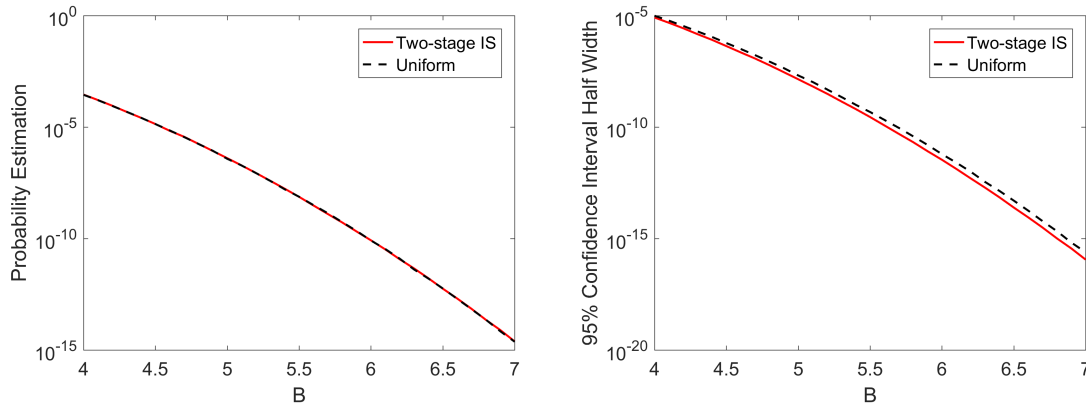


Figure 4: Probability estimates with different design space sizes. Figure 5: Confidence interval half-widths with different design space sizes.




Figure 6: Ratios of the confidence interval half-widths from the two-stage approach over the uniform IS.

These experiments indicate that the two-stage approach is efficient for problems with convex rare-event set compared to the uniform IS. Moreover, it shows more benefits when the design space is big compared to the location of the rare-event set, which is consistent with the findings in Section 3.3. Next, we investigate an extension of our approach to non-convex rare-event sets.

4.2 Extension to Non-convex Rare-event Sets

We consider a heuristic approach to generalize our scheme to rare-event set \mathcal{R} that is more complex and has multiple dominating points, say a_1^*, \dots, a_p^* , such that $\mathcal{R} \subseteq \cup_i \{x : a_i^{*'}(x - a_i^*) \geq 0\}$ and $a_i \in \mathcal{R}, i = 1, \dots, p$. In the first stage, we attempt to find a collection of approximate dominating points, denoted $\hat{a}_1, \dots, \hat{a}_{\hat{p}}$, where we note that \hat{p} may not necessarily equal p . In the second stage, we build an IS distribution that consists of a \hat{p} -mixture of Gaussian distributions, where each mixture component has mean \hat{a}_i and identity covariance.

In the first stage, we can find the approximate dominating points sequentially. We first find $\hat{a}_1 = \operatorname{argmin}_{x_i, i=1, \dots, n_0} \{\|x_i\|^2 : x_i \in \mathcal{R}\}$. Then we use a “cutting plane” to find $\hat{a}_2 = \operatorname{argmin}_{x_i, i=1, \dots, n_0} \{\|x_i\|^2 : x_i \in \mathcal{R}, \hat{a}_1'(x_i - \hat{a}_1) < 0\}$. We keep adding constraints to these optimization problems in a similar fashion, until we cover all the rare-event points.




Figure 7: Probability estimates under different sample budgets.




Figure 8: Confidence interval half-widths under different sample budgets.

We use similar problem settings as in Section 4.1 to compare the performances of the two-stage approach with the uniform IS. We modify the rare-event set in Section 4.1 to $\mathcal{R} = \{x : (x_1 - 5)^2 + (x_2 - 5)^2 \leq 4\} \cup \{x : (x_1 - 5)^2 + (x_2 - 1)^2 \leq 1\}$. Note that this set can no longer be contained by any half space in the form of $\{x : a'(x - a)\}$.

We again consider different sample budgets from $n = 1,000$ to $n = 50,000$. Figure 7 presents the probability estimates of the two approaches with different sample budgets. In this figure, it is not that obvious which approach is more stable. Nonetheless, Figure 8 shows that the two-stage approach (red solid line) is better than the uniform approach (black dash line) in terms of shorter confidence interval half-widths. These results are quite similar to the convex set case, and hints that the two-stage approach shows better performances for such types of rare-event problems as well.

5 CONCLUSION

We have considered the use of IS to estimate rare-event probabilities that lack structural information that is typically needed to design efficient schemes. We have considered a two-stage approach where the first, exploration, stage generates a space-filling set of points to learn the rare event of interest, and the second, exploitation, stage estimates the dominating point to guide an exponential tilting. Our derivation indicates that our two-stage IS estimator is efficient when, roughly speaking, the volume of the space is large compared to the position of the rare-event set, when compared to an IS based on a naive uniform sampling. This phenomenon is confirmed in two sets of numerical experiments. In future studies, we will extend our efficiency analysis to more general settings and compare with other potential methods.

ACKNOWLEDGMENTS

We gratefully acknowledge support from the National Science Foundation under grants CMMI-1542020 and CAREER CMMI-1653339/1834710.

REFERENCES

Adler, R. J., J. H. Blanchet, J. Liu et al. 2012. “Efficient Monte Carlo for High Excursions of Gaussian Random Fields”. *The Annals of Applied Probability* 22(3):1167–1214.

Asmussen, S., and P. W. Glynn. 2007. *Stochastic Simulation: Algorithms and Analysis*, Volume 57. New York: Springer Science & Business Media.

Au, S., and J. L. Beck. 1999. “A New Adaptive Importance Sampling Scheme for Reliability Calculations”. *Structural Safety* 21(2):135–158.

- Bengio, Y., and J.-S. Senécal. 2008. "Adaptive Importance Sampling to Accelerate Training of A Neural Probabilistic Language Model". *IEEE Transactions on Neural Networks* 19(4):713–722.
- Blanchet, J., and C. Li. 2011. "Efficient Simulation for the Maximum of Infinite Horizon Discrete-time Gaussian Processes". *Journal of Applied Probability* 48(2):467–489.
- Blanchet, J., and H. Lam. 2012. "State-dependent Importance Sampling for Rare-Event Simulation: An Overview and Recent Advances". *Surveys in Operations Research and Management Science* 17(1):38–59.
- Blanchet, J., P. Glynn, and H. Lam. 2009. "Rare Event Simulation for A Slotted Time M/G/s Model". *Queueing Systems* 63(1-4):33.
- Botev, Z. I., P. LEcuyer, and B. Tuffin. 2013. "Markov Chain Importance Sampling with Applications to Rare Event Probability Estimation". *Statistics and Computing* 23(2):271–285.
- Bourinet, J.-M., F. Deheeger, and M. Lemaire. 2011. "Assessing Small Failure Probabilities by Combined Subset Simulation and Support Vector Machines". *Structural Safety* 33(6):343–353.
- Bucklew, J. 2013. *Introduction to Rare Event Simulation*. New York: Springer Science & Business Media.
- Chan, J. C., and D. P. Kroese. 2012. "Improved Cross-entropy Method for Estimation". *Statistics and Computing* 22(5):1031–1040.
- Das, P., and Y. Zheng. 2000. "Cumulative Formation of Response Surface and Its Use in Reliability Analysis". *Probabilistic Engineering Mechanics* 15(4):309–315.
- De Boer, P.-T., D. P. Kroese, S. Mannor, and R. Y. Rubinstein. 2005. "A Tutorial on the Cross-entropy Method". *Annals of Operations Research* 134(1):19–67.
- Dieker, A. B., and M. Mandjes. 2006. "Fast Simulation of Overflow Probabilities in A Queue with Gaussian Input". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 16(2):119–151.
- Dubourg, V., B. Sudret, and F. Deheeger. 2013. "Metamodel-based Importance Sampling for Structural Reliability Analysis". *Probabilistic Engineering Mechanics* 33:47–57.
- Dupuis, P., A. D. Sezer, H. Wang et al. 2007. "Dynamic Importance Sampling for Queueing Networks". *The Annals of Applied Probability* 17(4):1306–1346.
- Echard, B., N. Gayton, and M. Lemaire. 2011. "AK-MCS: An Active Learning Reliability Method Combining Kriging and Monte Carlo Simulation". *Structural Safety* 33(2):145–154.
- Frank, J., S. Mannor, and D. Precup. 2008. "Reinforcement Learning in the Presence of Rare Events". In *Proceedings of the 25th International Conference on Machine Learning*, edited by McCallum, A. et al., 336–343. New York: ACM.
- Glasserman, P. 2013. *Monte Carlo Methods in Financial Engineering*, Volume 53. New York: Springer Science & Business Media.
- Glasserman, P., P. Heidelberger, and P. Shahabuddin. 1999. "Asymptotically Optimal Importance Sampling and Stratification for Pricing Path-Dependent Options". *Mathematical Finance* 9(2):117–152.
- Glynn, P. W., and D. L. Iglehart. 1989. "Importance Sampling for Stochastic Simulations". *Management Science* 35(11):1367–1392.
- Grabaskas, D., M. K. Nakayama, R. Denning, and T. Aldemir. 2016. "Advantages of Variance Reduction Techniques in Establishing Confidence Intervals for Quantiles". *Reliability Engineering & System Safety* 149:187–203.
- Grace, A. W., D. P. Kroese, and W. Sandmann. 2014. "Automated State-dependent Importance Sampling for Markov Jump Processes via Sampling from the Zero-Variance Distribution". *Journal of Applied Probability* 51(3):741–755.
- Heidelberger, P. 1995. "Fast Simulation of Rare Events in Queueing and Reliability Models". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 5(1):43–85.
- Huang, Z., H. Lam, D. J. LeBlanc, and D. Zhao. 2017a. "Accelerated Evaluation of Automated Vehicles Using Piecewise Mixture Models". *IEEE Transactions on Intelligent Transportation Systems*: 1–11.

- Huang, Z., H. Lam, and D. Zhao. 2017b. “An Accelerated Testing Approach for Automated Vehicles with Background Traffic Described by Joint Distributions”. In *Proceedings of 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, 933–938. Piscataway, New Jersey: IEEE.
- Huang, Z., Y. Guo, H. Lam, and D. Zhao. 2018. “A Versatile Approach to Evaluating and Testing Automated Vehicles based on Kernel Methods”.
- Hurtado, J. E. 2004. “An Examination of Methods for Approximating Implicit Limit State Functions from the Viewpoint of Statistical Learning Theory”. *Structural Safety* 26(3):271–293.
- Juneja, S., and P. Shahabuddin. 2006. “Rare-event Simulation Techniques: An Introduction and Recent Advances”. *Handbooks in Operations Research and Management Science* 13:291–350.
- Kanj, R., R. Joshi, and S. Nassif. 2006. “Mixture Importance Sampling and Its Application to the Analysis of SRAM Designs in the Presence of Rare Failure Events”. In *Proceedings of Design Automation Conference, 2006 43rd ACM/IEEE*, 69–72. Piscataway, New Jersey: IEEE.
- Lam, H. 2012. “Efficient Importance Sampling under Partial Information”. In *Proceedings of the 2012 Winter Simulation Conference (WSC)*, edited by Laroque, C. et al., 1–12. New Jersey: IEEE.
- OHagan, A. 2006. “Bayesian Analysis of Computer Code Outputs: A Tutorial”. *Reliability Engineering & System Safety* 91(10-11):1290–1300.
- Papadrakakis, M., and N. D. Lagaros. 2002. “Reliability-based Structural Optimization Using Neural Networks and Monte Carlo Simulation”. *Computer Methods in Applied Mechanics and Engineering* 191(32):3491–3507.
- Rubinstein, R. Y., and D. P. Kroese. 2004. *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation and Machine Learning*. New York: Springer.
- Rubinstein, R. Y., and D. P. Kroese. 2016. *Simulation and the Monte Carlo Method*, Volume 10. New Jersey: John Wiley & Sons.
- Sadowsky, J. S., and J. A. Bucklew. 1990. “On Large Deviations Theory and Asymptotically Efficient Monte Carlo Estimation”. *IEEE transactions on Information Theory* 36(3):579–588.
- Shelton, C. R. 2001. *Importance Sampling for Reinforcement Learning with Multiple Objectives*. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Siegmund, D. 1976. “Importance Sampling in the Monte Carlo Study of Sequential Tests”. *The Annals of Statistics*: 673–684.
- Zhao, D., H. Lam, H. Peng, S. Bao, D. J. LeBlanc, K. Nobukawa, and C. S. Pan. 2017. “Accelerated Evaluation of Automated Vehicles Safety in Lane-change Scenarios Based on Importance Sampling Techniques”. *IEEE Transactions on Intelligent Transportation Systems* 18(3):595–607.
- Zhao, Y.-G., and T. Ono. 1999. “A General Procedure for First/second-order Reliability Method (FORM/SORM)”. *Structural Safety* 21(2):95–112.

AUTHOR BIOGRAPHIES

ZHIYUAN HUANG is a third-year Ph.D. student in the Department of Industrial and Operations Engineering at the University of Michigan, Ann Arbor. His research interests include simulation and stochastic optimization. His email address is zhyhuang@umich.edu.

HENRY LAM is an Associate Professor in the Department of Industrial Engineering and Operations Research at Columbia University. His research focuses on Monte Carlo simulation, uncertainty quantification, risk analysis, and stochastic and robust optimization. His email address is khl2114@columbia.edu.

DING ZHAO is an Assistant Professor in the Department of Mechanical Engineering at Carnegie Mellon University. His research focuses on Connected and Automated Vehicles (CAVs) using synthesized approaches rooted in advanced statistics, modeling, optimization, dynamic control, and big data analysis. His email address is dingzhao@cmu.edu.