



HAL
open science

Improvement of the cross-entropy method in high dimension for failure probability estimation through a one-dimensional projection without gradient estimation

Maxime El Masri, Jérôme Morio, Florian Simatos

► To cite this version:

Maxime El Masri, Jérôme Morio, Florian Simatos. Improvement of the cross-entropy method in high dimension for failure probability estimation through a one-dimensional projection without gradient estimation. *Reliability Engineering and System Safety*, 2021, 216, pp.107991. 10.1016/j.ress.2021.107991 . hal-03463928

HAL Id: hal-03463928

<https://hal.science/hal-03463928>

Submitted on 2 Dec 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Improvement of the cross-entropy method in high dimension for failure probability estimation through a one-dimensional projection without gradient estimation

Maxime El Masri^{a,b}, Jérôme Morio^a, Florian Simatos^b

^aONERA/DTIS, Université de Toulouse, F-31055 Toulouse, France

^bISAE-SUPAERO and Université de Toulouse, Toulouse, France

Abstract

Rare event probability estimation is an important topic in reliability analysis. Stochastic methods, such as importance sampling, have been developed to estimate such probabilities but they often fail in high dimension. In this paper, we propose a new cross-entropy-based importance sampling algorithm to improve rare event probability estimation in high dimension. We focus on the cross-entropy method with Gaussian auxiliary distributions and we suggest to update the Gaussian covariance matrix only in a one-dimensional subspace. For that purpose, the main idea is to consider the projection in the one-dimensional subspace spanned by the sample mean vector, which gives an influential direction for the variance estimation. This approach does not require any additional simulation budget compared to the basic cross-entropy algorithm and we show on different numerical test cases that it greatly improves its performance in high dimension.

Keywords: Rare event simulation, Failure probability, Importance sampling, Cross-entropy, High dimension, Projection.

1. Introduction

Classical algorithms for reliability analysis. In reliability analysis, an important goal is to determine the probability of failure of an engineering system, often represented as the probability that a function, called limit-state or performance function, exceeds a given threshold. Estimating directly this probability can be difficult because calls to the function can be computationally expensive and the failure of the system is often a rare event.

Various techniques have been developed to estimate such rare event probabilities [1, 2]. First/second order reliability methods (FORM/SORM [3]) seek the most probable failure point (or design point) which minimizes the distance, in the standard normal space, between the failure hypersurface and the origin, and then realize a linear (FORM) or second-order (SORM) approximation of the limit-state function at this design point in order to estimate the probability with this approximation. Sampling techniques based on Monte Carlo simulations take advantage of the law of large numbers for rare event probability estimations [4], where this probability is estimated as the average number of random samples which fall in the failure domain. Another sampling alternative is the

Email addresses: melmasri@onera.fr (Maxime El Masri), jerome.morio@onera.fr (Jérôme Morio), florian.simatos@isae.fr (Florian Simatos)

subset simulation method [5, 6, 7], or adaptive splitting, that expresses the probability as the product of several conditional probabilities.

In this article, we focus on the Importance Sampling (IS) method [8, 9] that is another well-known Monte Carlo method considered in rare event simulation. IS consists in sampling from an alternative density, called auxiliary density, for which the failure is more frequent. Different techniques exist to find an efficient auxiliary density, as this choice has a strong influence on the performance of the overall IS estimation scheme. The performance comparison of IS with other failure probability estimation techniques is beyond the scope of this article. However, some insights are proposed in the recent article [10] that analyses the results of classical Monte Carlo, subset simulation and importance sampling, combined with a surrogate model on different reliability problems. One of the most well-known method is the cross-entropy method [11, 12], which we will simply refer to as CE in the rest of the paper. CE aims to determine the best auxiliary density within a given, usually parametric, family of densities: here “best” refers to the density which minimizes the Kullback–Leibler (KL) divergence with the optimal auxiliary density, which is intractable as it involves the unknown probability. Many papers are still published on the CE method, aiming at improving or applying CE in various contexts ([13, 14, 15, 16]). When the goal is to estimate a probability, a variant of CE (called improved cross-entropy method and referred to as iCE in the sequel) has also been developed in [13] to improve its efficiency by using a smooth approximation of the indicator function of the failure event. CE and iCE, which are the basis of the present work, are recalled in the Appendix. CE belongs to a larger family of sampling algorithms called Adaptive Importance Sampling (AIS) [17, 18], where one seeks for a valuable auxiliary density in a parametric family by successive updates of the parameters. A non-parametric IS has also been proposed, considering kernel density estimation to approximate iteratively the optimal auxiliary density (see [19] or [20]).

New Challenges in importance sampling for reliability analysis. While the fundamentals of IS are relatively well established, IS suffers from several limitations that may hinder its applicability. First of all, IS often requires a computational budget of at least a thousand calls to the simulation code, which can be too prohibitive in practice. To this end, it is interesting to couple surrogate models and IS and different approaches have been described to optimally achieve this coupling [21, 22, 23]. The enrichment of the surrogate model and particularly the choice of the calculations [24] to be carried out or the characterisation of the uncertainty [25] due to the surrogate model on the probability estimation are among others emerging research questions. Unlike classical Monte Carlo approaches, IS is not immediately suitable for high-dimensional problems because the search space for an efficient auxiliary density becomes very large [20, 26, 27] and also because of the degeneracy of the IS weights [28]. To counter this limitation, new approaches for intrinsic dimension analysis and dimension reduction [29, 30] have recently been proposed to focus on the dimensions most influential on integration. We give more details about this issue on the specific case of CE in the next paragraph.

Rare event probability estimation with CE in high dimension. All these techniques face the curse of dimensionality and become inaccurate as the dimension increases. For CE this is due to the weight degeneracy phenomenon in high dimension, which involves the collapse of the estimated covariance matrix to 0 due to weight degeneracy [28, 31]. In [32], auxiliary densities focus on the most influential variables which are selected beforehand, to reduce the number of updated parameters and decrease the variance of the IS estimator. However the technique is not efficient when all variables are influential. [15] proposes

to directly generate a sample from the optimal unknown IS distribution with an MCMC method and then estimate the CE-optimal parameter. This one should be more accurate and thus lead to a better probability estimation, but in high dimension the MCMC step can require a large simulation budget to reach convergence. Another idea to improve CE in high dimension is to use different, lower-dimensional parametric families of auxiliary densities. For instance, [33] propose to use von Mises–Fisher mixtures and [13] to use von Mises–Fisher–Nakagami mixtures: both distributions are supported by the hypersphere which reduces the number of parameters compared to Gaussian mixtures [34] which are not efficient in high dimension. These methods are mainly well suited to multimodal problems, that will not be tackled in this paper, and require a large number of samples.

Closer to the spirit of our paper, the recent paper [30] proposes to construct a subspace, called the Failure-Informed Subspace (FIS), which identifies a low-dimensional structure of the problem and then applies CE on the FIS by projecting the samples: the resulting algorithm is a variation of iCE called iCEred. The FIS is chosen based on a theoretical bound which suggests that it induces a small KL divergence: however its computation requires the evaluation of the gradient of the limit-state function. This presents two drawbacks: 1) the limit-state function needs to be differentiable and 2) computing its gradient is often expensive (and sometimes even out of reach).

Novel methodological contribution. In the present paper, we propose two new algorithms, CE- m^* and iCE- m^* , which pursue the line of research recently initiated in [30], namely circumventing the curse of dimensionality faced by CE (and iCE) for the estimation of a failure probability. The main idea of this article is to project the generated samples on a suitable low-dimensional space. The proposed projection direction is the sample mean vector of the input conditioned on failure (see Section 3.2 for more details) whose computation does not require any additional model evaluations compared to CE or iCE as it is not based on the gradient. This direction is motivated by heuristic arguments presented in Section 3.3 and on all the numerical examples that we considered, in particular the representative ones that we report in Section 4, they improve the performance of CE. In many cases and even for simplistic linear performance functions, CE and iCE do not converge for dimensions larger than 30 whereas our algorithms provide accurate estimates up to dimensions 200 and above.

Organization of the paper. The following study is divided in four sections. Section 2 recalls the foundations of IS, CE and iCE. Section 3 explains why CE does not work in high dimension and then introduces the two proposed algorithms, CE- m^* and iCE- m^* . Their performance is numerically evaluated in Section 4 on four different examples. We conclude the paper in Section 5, giving a short summary and some perspectives.

2. The cross-entropy method

2.1. Rare event probability estimation

Let X be an n -dimensional random vector from a continuous probability distribution function f modeling the input variables. Define a failure domain $\mathcal{F} = \{x \in \mathbb{R}^n : \varphi(x) \geq 0\}$ and the associated probability of failure as:

$$P = \mathbb{P}_f(\varphi(X) \geq 0) = \mathbb{E}_f(\mathbb{I}_{\{\varphi(X) \geq 0\}}) = \int_{\mathbb{R}^n} \mathbb{I}_{\{\varphi(x) \geq 0\}} f(x) dx$$

where $\mathbb{I}_{\{\varphi(x) \geq 0\}} = 1$ when $\varphi(x) \geq 0$ and 0 otherwise. Throughout the paper \mathbb{P}_f and \mathbb{E}_f denote the probability and the expectation with respect to f . The function φ represents a computationally demanding engineering model and is considered as a black box, so that the probability P cannot be estimated analytically. Naive Monte Carlo methods with independent and identically distributed (i.i.d.) samples from f can be used to estimate P , but they require a large number of samples when the failure is a rare event.

2.2. Importance sampling (IS)

IS [8, 9] consists in sampling the input variables from an alternative distribution g well suited to the failure domain, that is, such that the probability of failure with respect to g is not too rare. Let g be a probability density function such that $g(x) = 0 \Rightarrow \mathbb{I}_{\{\varphi(x) \geq 0\}} f(x) = 0$, then we have

$$P = \int_{\mathbb{R}^n} \mathbb{I}_{\{\varphi(x) \geq 0\}} L(x) g(x) dx = \mathbb{E}_g \left(\mathbb{I}_{\{\varphi(X) \geq 0\}} L(X) \right)$$

where $L = f/g$ is called the likelihood ratio (or importance weight), and g is the auxiliary density. The IS estimator of P is then

$$\hat{P} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}_{\{\varphi(X_i) \geq 0\}} L(X_i)$$

where X_1, \dots, X_N are i.i.d. random vectors drawn from g . The auxiliary density has a strong influence on the accuracy of the estimation. The optimal auxiliary density g^* is the distribution conditioned on the failure domain, i.e.,

$$g^*(x) = \frac{\mathbb{I}_{\{\varphi(x) \geq 0\}} f(x)}{P}, \quad x \in \mathbb{R}^n. \quad (1)$$

With this distribution, the probability estimate \hat{P} is exactly P and has zero variance. But g^* depends on P , the unknown probability, hence it cannot be used in practice. For this reason, one often looks for an approximation of g^* in a parametric family of densities $\mathcal{G} = \{g_\lambda : \lambda \in \Lambda\}$ called auxiliary density family. A natural choice for λ is the parameter λ^* that minimizes the variance, i.e.,

$$\lambda^* = \arg \min_{\lambda \in \Lambda} \text{Var}_{g_\lambda}(\hat{P}(\lambda)) = \arg \min_{\lambda \in \Lambda} \text{Var}_{g_\lambda} \left(\mathbb{I}_{\{\varphi(X) \geq 0\}} L_\lambda(X) \right)$$

with $\hat{P}(\lambda) = (1/N) \sum_{i=1}^N \mathbb{I}_{\{\varphi(X_i) \geq 0\}} L_\lambda(X_i)$ the IS estimator and $L_\lambda = f/g_\lambda$. The variance minimization problem is not convex with no analytical solution. An alternative and popular choice, which we discuss next, is the cross-entropy method which estimates the parameter λ that minimizes the Kullback-Leibler divergence between g^* and g_λ .

2.3. The cross-entropy method (CE)

The Kullback-Leibler (KL) divergence between two densities f and h with f absolutely continuous with respect to h is defined by:

$$D(f, h) = \mathbb{E}_f \left[\log \left(\frac{f(X)}{h(X)} \right) \right] = \int f(x) \log \left(\frac{f(x)}{h(x)} \right) dx.$$

The KL divergence measures a “distance” between the two densities. It is always non-negative, and it is zero only when $f = h$ almost everywhere. In order to find an efficient

auxiliary density, we can minimize the KL divergence between the optimal g^* and the given parametric auxiliary family and choose λ^* given by

$$\lambda^* = \arg \min_{\lambda \in \Lambda} D(g^*, g_\lambda).$$

Since $D(g^*, g_\lambda) = \mathbb{E}_{g^*} [\log(g^*(X))] - \mathbb{E}_{g^*} [\log(g_\lambda(X))]$ and the first term in the right-hand side is independent of λ , this problem reduces to maximizing the cross-entropy $\mathbb{E}_{g^*} [\log(g_\lambda(X))]$. Using finally that $g^* = f\mathbb{I}_{\{\varphi \geq 0\}}/P$ with P independent of λ , we get

$$\lambda^* = \arg \max_{\lambda \in \Lambda} \mathbb{E}_f \left[\log(g_\lambda(X)) \mathbb{I}_{\{\varphi(X) \geq 0\}} \right]. \quad (2)$$

This maximization problem is called the cross-entropy problem. Contrary to the variance minimization problem, the cross-entropy problem is generally concave and differentiable [11, 4], so the solution can be found by setting the gradient with respect to λ equal to zero. Another advantage is the existence of analytical solutions of the cross-entropy problem for exponential families.

From now on, we suppose that the initial density f is the standard normal density g_{0, I_n} and that the auxiliary density is a Gaussian distribution $\{g_{m, \Sigma} : m \in \mathbb{R}^n, \Sigma \in \mathcal{M}_n^+\}$, which belongs to the exponential family. Here and in the sequel, $g_{m, \Sigma}$ denotes the Gaussian density with mean $m \in \mathbb{R}^n$ and covariance matrix $\Sigma \in \mathcal{M}_n^+$ with $\mathcal{M}_n^+ \subset \mathbb{R}^{n \times n}$ the set of symmetric, positive-definite matrices, and I_n is the identity matrix in dimension n . It is convenient to work in the standard normal space, because it enables easier implementation and computation. By transforming the random inputs X with isoprobabilistic transformations [35, 36], we can in principle always come back to this case.

In this case, we have $\lambda = (m, \Sigma) \in \mathbb{R}^n \times \mathcal{M}_n^+$ and the cross-entropy maximization problem (2) can be solved explicitly and gives

$$m^* = \mathbb{E}_f(X \mid \varphi(X) \geq 0) \quad \text{and} \quad \Sigma^* = \mathbb{E}_f \left((X - m^*)(X - m^*)^\top \mid \varphi(X) \geq 0 \right), \quad (3)$$

see for instance [11, 34]. Since m^* and Σ^* depend on the unknown distribution conditional on failure g^* , we cannot compute them explicitly but, provided we have a relevant auxiliary density h , we can estimate them by IS through

$$\hat{m} = \sum_{i=1}^N w_i X_i \quad \text{and} \quad \hat{\Sigma} = \sum_{i=1}^N w_i (X_i - \hat{m})(X_i - \hat{m})^\top \quad (4)$$

where $w_i = w'_i / \sum_j w'_j$ and $w'_i = \mathbb{I}_{\{\varphi(X_i) \geq 0\}} f(X_i) / h(X_i)$ are the unnormalized importance weights, and the X_i are i.i.d. samples of density h .

To find a good auxiliary density h , we can consider sequential simulation techniques to define a sequence of parameters $(m_0, \Sigma_0), (m_1, \Sigma_1), \dots$. Provided (m_t, Σ_t) is given, the idea is to generate a sample according to g_{m_t, Σ_t} and estimate (m_{t+1}, Σ_{t+1}) from the empirical ρ -quantile ($\rho \in (0, 1)$) of $\varphi(X_1), \dots, \varphi(X_N)$. This is a classical algorithm often referred to as the multilevel cross-entropy algorithm (CE). For completeness, a detailed description of CE is provided in [Appendix A](#), Algorithm 3 (see also [11]). Without giving all the details recalled in the appendix, it is important to have in mind the following overview of one iteration of CE's main step:

1. Draw a sample X_1, \dots, X_N from g_{m_t, Σ_t} with (m_t, Σ_t) the current estimate of (m^*, Σ^*) ;

2. Compute γ_t the empirical ρ -quantile of the sample $\varphi(X_1), \dots, \varphi(X_N)$, and get a new estimation (m_{t+1}, Σ_{t+1}) of the conditional mean and variance by

$$m_{t+1} = \sum_{i=1}^N w_i X_i \quad \text{and} \quad \Sigma_{t+1} = \sum_{i=1}^N w_i (X_i - m_{t+1})(X_i - m_{t+1})^\top \quad (5)$$

where $w_i = w'_i / \sum_j w'_j$ and $w'_i = \mathbb{I}_{\{\varphi(X_i) \geq \gamma_t\}}(f/g_{m_t, \Sigma_t})(X_i)$.

Remark 1 (Smooth updating of the parameters). Instead of updating m_{t+1} and Σ_{t+1} from (5), it is sometimes more efficient to perform a smoothed updating rule, as follows :

$$\tilde{m}_{t+1} = \alpha m_{t+1} + (1 - \alpha) \tilde{m}_t \quad \text{and} \quad \tilde{\Sigma}_{t+1} = \alpha \Sigma_{t+1} + (1 - \alpha) \tilde{\Sigma}_t$$

where m_{t+1} and Σ_{t+1} are estimated in (5), \tilde{m}_t and $\tilde{\Sigma}_t$ are the smoothed parameters computed at the previous step "t", and $\alpha \in [0, 1]$ is the smoothing parameter. This procedure can avoid a too fast convergence to degenerate solutions, especially in discrete optimization problems (see [1, 11, 37] for instance). The smoothing parameter α can be fixed beforehand (static smoothing) or updated at each iteration (dynamic smoothing). In the numerical analysis of Section 4, we do not perform the smoothing procedure (or equivalently, we keep $\alpha = 1$) since it does not improve significantly the estimation results.

2.4. Improved cross-entropy (iCE) method

One of the limitations of CE lies in the estimations of equation (4) of the conditional distribution at each step of the algorithm. This estimation only uses a fraction ρ of the sample (recall that $\rho \in (0, 1)$ is a parameter of CE, see Section 2.3), and the idea of iCE is to replace the indicator function $\mathbb{I}_{\{\varphi(X) \geq 0\}}$ by a smooth approximation as suggested in [13, 38], for instance. The method proposed in [13] considers the approximation $\mathbb{I}_{\{\varphi(x) \geq 0\}} \approx F(\varphi(x)/\sigma)$ with F a cumulative distribution function, typically of the standard normal random variable, so that for $\varphi(x) \neq 0$ we have $F(\varphi(x)/\sigma) \rightarrow \mathbb{I}_{\{\varphi(x) \geq 0\}}$ as $\sigma \rightarrow 0$. We present this approach with Gaussian densities in Algorithm 4 of Appendix A.

2.5. Comment on the choice of auxiliary family

As mentioned above, we restrict our study to the Gaussian auxiliary family $\{g_{m, \Sigma} : m \in \mathbb{R}^n, \Sigma \in \mathcal{M}_n^+\}$. This choice is well suited for unimodal problems, which is the typical case that we have in mind, and in such cases CE is known to work well for small dimensions. Our goal here is to improve the efficiency of CE when the dimension increases. In particular, the heuristic arguments justifying CE- m^* presented in Section 3.3 work in this case and not necessarily in more general settings; moreover, we have only considered such test cases for the numerical applications. In case the failure domain is not unimodal, the auxiliary family needs to be updated, see for instance [13, 33, 34].

3. New algorithms CE- m^* and iCE- m^* with projection

3.1. Inefficiency of CE in high dimension

The inefficiency of CE in high dimension is well documented: the general behavior, even outside a reliability context, is that the estimation of the covariance matrix collapses

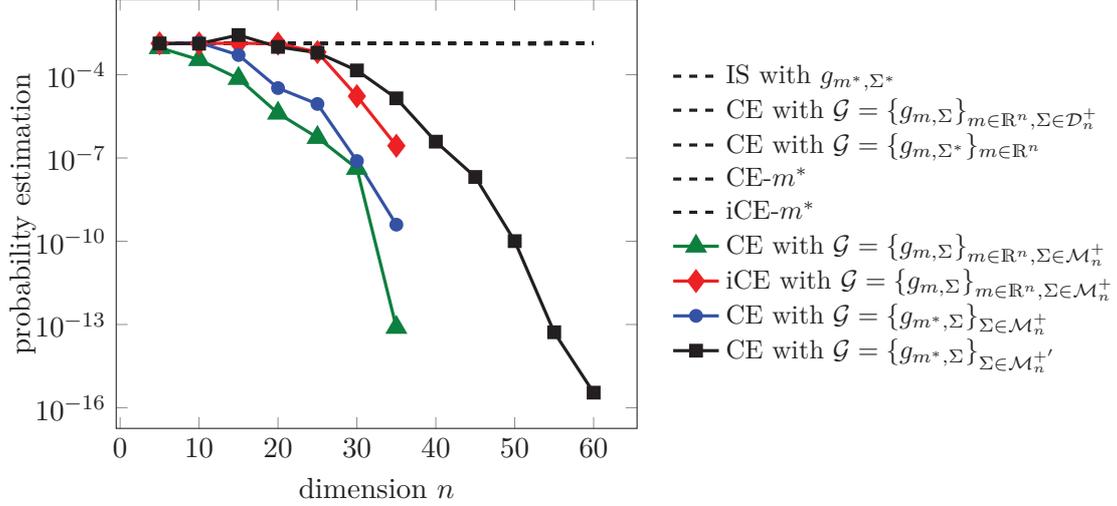


Figure 1: Evolution of the estimation of the probability $P = \mathbb{P}_f(\varphi(X) > 0)$ as the dimension increases for φ linear given by (6) and for different algorithms. Results are obtained with parameters set as in Section 4.1. In dashed are three algorithms which provide accurate and indistinguishable results: IS with the optimal density g_{m^*, Σ^*} , CE with auxiliary family $\mathcal{G} = \{g_{m, \Sigma}\}_{m \in \mathbb{R}^n, \Sigma \in \mathcal{D}_n^+}$ where one uses the optimal variance and estimates the mean, and CE with auxiliary family $\mathcal{G} = \{g_{m, \Sigma}\}_{m \in \mathbb{R}^n, \Sigma \in \mathcal{D}_n^+}$ where one estimates the mean and only the diagonal of the covariance matrix. In green with triangle is CE with $\{g_{m, \Sigma}\}_{m \in \mathbb{R}^n, \Sigma \in \mathcal{M}_n^+}$ (both the mean and the full covariance matrix are estimated). In red with diamonds is iCE with the same auxiliary family. In blue with circles is CE with $\{g_{m^*, \Sigma}\}_{\Sigma \in \mathcal{M}_n^+}$ (one uses the optimal mean m^* and estimates the full covariance matrix). In black with squares is CE with $\{g_{m^*, \Sigma}\}_{\Sigma \in \mathcal{M}_n^{+'}}$ (one uses the optimal mean m^* and estimates 3/4-th of the covariance matrix).

very suddenly to 0 [28, 32, 39]. To illustrate this, let us consider the simple yet illustrative example where φ is the following affine function:

$$\varphi : x = (x_1, \dots, x_n) \in \mathbb{R}^n \mapsto \sum_{j=1}^n x_j - 3\sqrt{n} \quad (6)$$

such that $P = \mathbb{P}_f(\varphi(X) \geq 0) = \mathbb{P}(Z \geq \beta) \simeq 1.35 \cdot 10^{-3}$ independently of n , with Z following the standard normal distribution on \mathbb{R} . The evolution of the mean probability estimation by CE and iCE is represented in Figure 1 (green triangles for CE, red diamonds for iCE). As the dimension increases, the estimation becomes less accurate for both methods. For CE, the error is significant even for dimensions as low as $n \approx 15$, while iCE delivers accurate estimations only up to dimensions $n \approx 25$, after what its accuracy dramatically drops. Thus even for moderately high dimensions ($n \approx 30$), both methods fail to provide accurate estimations.

The degradation is due to errors in the parameters estimation. Indeed, both methods aim at finding a good auxiliary density, close to g^* , in the family $\{g_{m, \Sigma}\}_{m \in \mathbb{R}^n, \Sigma \in \mathcal{M}_n^+}$: to do so, they use sequential simulation techniques to get estimates \hat{m} and $\hat{\Sigma}$ of the optimal parameters m^* and Σ^* . As $m^* \in \mathbb{R}^n$ and $\Sigma^* \in \mathbb{R}^{n \times n}$ is symmetric, this actually involves the estimation of $n + \frac{1}{2}n(n+1) = \frac{1}{2}n(n+3)$ parameters, a number which grows quickly (quadratically) as the dimension increases. Thus, although the error in each dimension may be low, their global effect may actually be quite significant in the final estimation. This is supported by the numerical results presented in Figure 1. First, we see indeed that IS with the optimal auxiliary density g_{m^*, Σ^*} works very well, so final estimation

errors indeed come from considering approximations \hat{m} and $\hat{\Sigma}$ of m^* and Σ^* . The four other curves related to CE are somehow in-between the case g_{m^*,Σ^*} where 0 parameter has to be estimated, and the case $\{g_{m,\Sigma}\}_{m \in \mathbb{R}^n, \Sigma \in \mathcal{M}_n^+}$ where $n + n(n+1)/2$ have to be estimated. In the case $\{g_{m,\Sigma^*}\}_{m \in \mathbb{R}^n}$ and $\{g_{m,\Sigma}\}_{m \in \mathbb{R}^n, \Sigma \in \mathcal{D}_n^+}$ one only has to estimate n and $2n$ parameters, respectively (\mathcal{D}_n^+ is the set of n -dimensional matrices with positive entries on the diagonal and 0's elsewhere), and CE with these auxiliary families behave very well. In the two other cases the number of parameters to be estimated grows quadratically in n where we have estimated P with four different methods: among them $\{g_{m^*,\Sigma}\}_{\Sigma \in \mathcal{M}_n^{+'}}$ with $3n(n+1)/8$ parameters to estimate behaves better than $\{g_{m^*,\Sigma}\}_{\Sigma \in \mathcal{M}_n^+}$ with $n(n+1)/2$ parameters to estimate ($\mathcal{M}_n^{+'}$ is the subset of matrices in \mathcal{M}_n^+ equal to the identity except for the north-west corner of size $3n/4$). In other words, Figure 1 suggests that the efficiency of a given auxiliary family depends on the number of parameters to be estimated: the efficiency improves as the number of parameters to be estimated decreases and a linear number of parameters to estimate behaves very well. This suggests the idea of the present paper, namely lowering the number of estimations needed for the covariance matrix by only estimating the variance in one direction.

3.2. Presentation of CE- m^* and iCE- m^*

From Figure 1 one could be tempted to use the auxiliary family $\{g_{m,\Sigma}\}_{m \in \mathbb{R}^n, \Sigma \in \mathcal{D}_n^+}$ where one only estimates diagonal entries: however, even if it behaves very well on the particular example of the sum (6), until dimension 60, this actually corresponds to a particular choice of auxiliary densities proscribing correlations between input which may in some cases be undesirable. Moreover, this choice involves more parameter estimation in non influential directions than the proposed method, which can result in inaccurate probability estimation or even the non-convergence of the algorithm in high dimension (see results in Section 4). Despite these limitations, we have studied this algorithm numerically as it fits within the framework of the present paper as it can be seen as a dimension-reduction algorithm, see Remark 4 below.

The proposition of this paper is to estimate the variance in the direction of the current estimation of m^* : similarly as the choice $\{g_{m,\Sigma}\}_{m \in \mathbb{R}^n, \Sigma \in \mathcal{D}_n^+}$ this presents the advantage of only necessitating a linear number of estimations (actually, $n+1$) but in contrast to $\{g_{m,\Sigma}\}_{m \in \mathbb{R}^n, \Sigma \in \mathcal{D}_n^+}$ this is a more generic choice that does not make any a priori assumption on the problem. In fact, an advantage of estimating the variance only in the direction of the updated mean vector is that we do not cancel all the correlation coefficients contrary to the diagonal form of the covariance matrix. The optimal auxiliary IS distribution has dependent components and thus a non diagonal covariance matrix is adapted. Thus, in low dimensions, the lack of flexibility of the proposed model compared to the diagonal covariance matrix is compensated by a better model assumption. The diagonal form assumption constrains the update of the variance in all the directions given by the canonical basis. Among the n directions, very few canonical directions have a variance different from 1 in high dimension and thus do really need to be updated. On the contrary, in the rare event context, the direction given by the sample mean vector is relevant since the variance has to decrease in this direction even if it is not always the direction that offers the best variance reduction. Having only 1 covariance matrix parameter to update whatever the dimension n becomes clearly an advantage in high dimensions and explains why CE with diagonal covariance matrix suffers more from the curse of dimensionality than CE- m^* , iCE- m^* . The precise algorithm is detailed in Algorithm 1 and its improved (in the sense of iCE) version in Algorithm 2. Its main step can be summarized as follows:

1. Draw a sample X_1, \dots, X_N from g_{m_t, Σ_t} with (m_t, Σ_t) the current estimate of (m^*, Σ^*) ;
2. Compute the empirical ρ -quantile γ_t of the sample $\varphi(X_1), \dots, \varphi(X_N)$ and get a new estimation m_{t+1} of the conditional mean m^* by

$$m_{t+1} = \sum_{i=1}^N w_i X_i \quad \text{with} \quad w_i = \frac{w'_i}{\sum_j w'_j} \quad \text{and} \quad w'_i = \mathbb{I}_{\{\varphi(X_i) \geq \gamma_t\}} \frac{f(X_i)}{g_{m_t, \Sigma_t}(X_i)};$$

3. Define $Y_i = R^\top X_i$ where $R^\top = m_{t+1}^\top / \|m_{t+1}\|$: the Y_i 's are the projections of the X_i 's on $\text{span}(m_{t+1})$, by linearity their weighted mean is equal to $R^\top m_{t+1} = \|m_{t+1}\|$;
4. Compute the variance \hat{v} in the direction of m_{t+1} through

$$\hat{v} = \sum_{i=1}^N w_i (Y_i - \|m_{t+1}\|)^2;$$

5. Get a new estimation Σ_{t+1} of Σ^* by considering a variance of \hat{v} in the direction of m_{t+1} and a unit variance I_{n-1} in its orthogonal: as will be discussed below, this amounts to taking $\Sigma_{t+1} = I_n + (\hat{v} - 1)RR^\top$.

The only difference between CE (Algorithm 3) and CE- m^* (Algorithm 1) is step 9 of CE- m^* : in CE, one estimates the whole covariance matrix through

$$\Sigma_{t+1} = \sum_{i=1}^N w_i (X_i - m_{t+1})(X_i - m_{t+1})^\top$$

whereas in CE- m^* , one only estimates a scalar

$$\hat{v} = \sum_{i=1}^N w_i (Y_i - \|m_{t+1}\|)^2$$

and then uses $\Sigma_{t+1} = I_n + (\hat{v} - 1)m_{t+1}m_{t+1}^\top / \|m_{t+1}\|^2$ (actually, with an additional εI_n term, see Remark 2 below). The number of parameters to be estimated in CE- m^* thus reduces from $\frac{1}{2}n(n+3)$ in CE to $n+1$.

Having the variance updated in only one direction $d \in \mathbb{R}^n$ amounts to seeking an auxiliary density of the form $X = \mu + \sigma Gd + R_\perp G_\perp$ with $G \sim \mathcal{N}(0, 1)$ and $G_\perp \sim \mathcal{N}(0, I_{n-1})$ independent standard Gaussian vectors in dimension one and $n-1$, respectively, σ^2 the variance of X in the direction of d and $R_\perp : \mathbb{R}^{n-1} \rightarrow \mathbb{R}^n$ that lifts a vector expressed in the basis of $\text{span}(d)^\perp$ (the orthogonal of the vector space spanned by d) to $\text{span}(d)^\perp$, which satisfies $dd^\top + R_\perp R_\perp^\top = I_n$. From the expression $X = \mu + \sigma Gd + R_\perp G_\perp$, one gets that X follows an n -dimensional Gaussian distribution with mean μ and covariance

$$\text{Var}(X) = \sigma^2 dd^\top + R_\perp R_\perp^\top = I_n + (\sigma^2 - 1)dd^\top$$

which justifies the update rule for Σ_{t+1} in step 9 of CE- m^* .

Remark 2. As can be seen in Algorithms 1 and 2, we actually update Σ_{t+1} through $\Sigma_{t+1} = (1+\varepsilon)I_n + (\hat{v}-1)RR^\top$: the additional εI_n term is here to avoid degeneracy problems due to numerical rounding errors. This numerical trick is also used in the implementation of CE and iCE available on the authors' website [40]. In all our simulations the value of ε is set to $\varepsilon = 10^{-6}$.

Algorithm 1: Proposed algorithm: CE- m^* = CE + projection on $\text{span}(m^*)$.

Data: parameter $\rho \in (0, 1)$, sample size N

Result: Estimation \hat{P} of the failure probability $\mathbb{P}_f(\varphi(X) \geq 0)$

- 1 Initialization: set $t = 0$, $m_t = 0$ and $\Sigma_t = I_n$;
 - 2 Generate X_1, \dots, X_N independently according to g_{m_t, Σ_t} and define $L_t = f/g_{m_t, \Sigma_t}$;
 - 3 Compute $\varphi_i = \varphi(X_i)$ for $i = 1, \dots, N$;
 - 4 Sort them in ascending order : $\varphi_{(1)} \leq \dots \leq \varphi_{(N)}$, and set $\gamma_t = \varphi_{(\lfloor (1-\rho)N \rfloor)}$;
 - 5 Compute the weights $w_i = w'_i / \sum_j w'_j$ with $w'_i = \mathbb{I}_{\{\varphi_i \geq \gamma_t\}} L_t(X_i)$;
 - 6 **while** $\gamma_t < 0$ **do**
 - 7 Compute $m_{t+1} = \sum_{i=1}^N w_i X_i$;
 - 8 Define $R^\top = m_{t+1}^\top / \|m_{t+1}\| \in \mathbb{R}^{1 \times n}$ the projection on $\text{span}(m_{t+1})$;
 - 9 Compute the projected samples $Y_i = R^\top X_i$ and their conditional variance $\hat{v} = \sum_{i=1}^N w_i (Y_i - \|m_{t+1}\|)^2$ and set $\Sigma_{t+1} = (\hat{v} - 1)RR^\top + (1 + \varepsilon)I_n$;
 - 10 Set $t \leftarrow t + 1$;
 - 11 Repeat Steps 2, 3, 4 and 5 ;
 - 12 **end**
 - 13 Return $\hat{P} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}_{\{\varphi(X_i) \geq 0\}} L_t(X_i)$.
-

Remark 3. As for the CE and iCE algorithms, it is possible to use a smoothed updating procedure (see remark 1) in CE- m^* and iCE- m^* . In this paper, the parameters are not updated with the smoothing rule in the simulations proposed in this article, since it does not improve the results significantly.

Remark 4. The auxiliary family $\{g_{m, \Sigma}\}_{m \in \mathbb{R}^n, \Sigma \in \mathcal{D}_n^+}$ amounts to only evaluating variance terms in the n canonical directions of \mathbb{R}^n . As such, it is of the same nature than the algorithms CE- m^* and iCE- m^* that we propose here, i.e., it seeks to improve results by projecting. However, the details differ significantly: although we propose here to project on one adaptive direction, using $\{g_{m, \Sigma}\}_{m \in \mathbb{R}^n, \Sigma \in \mathcal{D}_n^+}$ amounts to projecting on n static directions. As will be seen below, CE- m^* and iCE- m^* behave better than $\{g_{m, \Sigma}\}_{m \in \mathbb{R}^n, \Sigma \in \mathcal{D}_n^+}$, which suggests that it is better to have a smaller number of carefully chosen directions, rather than using many directions but which are not adapted to the samples.

3.3. Justification

The idea of CE- m^* , and other projection algorithms such as [30], is to only estimate variance parameters in relevant directions. In [30] the choice of the directions is suggested by a theoretical result which provides an upper bound on the KL divergence. However, the computation of these directions requires to compute the gradient of the limit-state function which presents two drawbacks: 1) the limit-state function needs to be differentiable and 2) computing its gradient is often expensive (and sometimes even out of reach). In comparison, the advantage of CE- m^* is twofold: 1) the gradient need not be computed: this can save a significant amount of simulation budget, and makes the method applicable even if the performance function is not differentiable (which is for instance the case

Algorithm 2: Proposed algorithm: iCE- m^* = iCE + projection on $\text{span}(m^*)$.

Data: parameter δ_{target} , sample size N
Result: Estimation \hat{P} of the failure probability $\mathbb{P}_f(\varphi(X) \geq 0)$

- 1 Initialization: set $t = 0$, $m_t = 0$, $\Sigma_t = I_n$ and $\sigma_t = \infty$;
- 2 Generate X_1, \dots, X_N independently according to g_{m_t, Σ_t} and define $L_t = f/g_{m_t, \Sigma_t}$;
- 3 Compute $\varphi_i = \varphi(X_i)$ for $i = 1, \dots, N$;
- 4 Compute $\widehat{c\hat{v}}$ the empirical coefficient of variation of the $\mathbb{I}_{\{\varphi_i \geq 0\}}/F(\varphi_i/\sigma_t)$;
- 5 **while** $\widehat{c\hat{v}} \geq \delta_{\text{target}}$ **do**
- 6 Compute $\sigma_{t+1} = \arg \min(\widehat{\delta}_t(\sigma) - \delta_{\text{target}})^2$ where the minimum is taken over $\sigma \in (0, \sigma_t)$ and $\widehat{\delta}_t(\sigma)$ is the coefficient of variation of the $F(\varphi_i/\sigma)L_t(X_i)$ with $L_t = f/g_{m_t, \Sigma_t}$;
- 7 Compute the weights $w_i = w'_i / \sum_j w'_j$ with $w'_i = F(\varphi_i/\sigma_{t+1})L_t(X_i)$;
- 8 Compute $m_{t+1} = \sum_{i=1}^N w_i X_i$ and set $R = m_{t+1} / \|m_{t+1}\|$;
- 9 Compute the projected samples $Y_i = R^\top X_i$ and their conditional variance $\hat{v} = \sum_{i=1}^N w_i (Y_i - \|m'_{t+1}\|)^2$ and set $\Sigma_{t+1} = (\hat{v} - 1)RR^\top + (1 + \varepsilon)I_n$;
- 10 Set $t \leftarrow t + 1$;
- 11 Repeat Steps 2, 3 and 4;
- 12 **end**
- 13 Return $\hat{P} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}_{\{\varphi(X_i) \geq 0\}} L_t(X_i)$.

for the financial application considered in Section 4.3); and 2) it is extremely simple, in particular there is no additional model evaluation compared to CE. Last but not least, in high dimension where CE behaves poorly, it can significantly improve the performance.

At a high level, a good direction in which to project the random variable $X|\varphi(X) \geq 0$ (and thus estimate its variance) is a direction in which its variance is significantly different from the variance of the initial density of X , i.e 1 : indeed, this means that the conditioning by the limit state function plays a role on the resulting variance, in contrast to what happens in directions where the variance remains close to a unit value. We notice in the simulations that it is indeed the case for the direction m^* , i.e., in all the test cases that we considered, the variance in the direction of m^* is systematically smaller than one. The reason for that lies in the very light tail of the Gaussian distribution which makes the variance of the overshoot very small. More precisely, if X is normally distributed, then conditionally on $X > S$ we have for large S the asymptotic expansion $X \approx S + E/S$ with E an exponential random variable. More formally, by considering Laplace transforms, it is easy to establish that $S(X - S)$ conditionally on $X > S$ converges in distribution to an exponential random variable from which we can show that the variance of X conditional on $X > S$ goes to 0 as $S \rightarrow \infty$. We illustrate this situation in Figure 2, where we simulate with acceptance-rejection 20 samples of a standard normal variable X knowing that $X > S$ for an increasing sequence of thresholds S . The variance of the resulting samples tends to 0 as S increases.

This behavior is relevant for our purposes because by construction, m^* is the direction to which the CE “pushes” the sample, and so heuristically this direction points toward the failure domain. Because of the light tail property of the normal distribution and because the updated variance is only computed on the “overshoots”, i.e., the points exceeding the border of the failure domain, it is natural that the variance in the direction of m^* decreases.

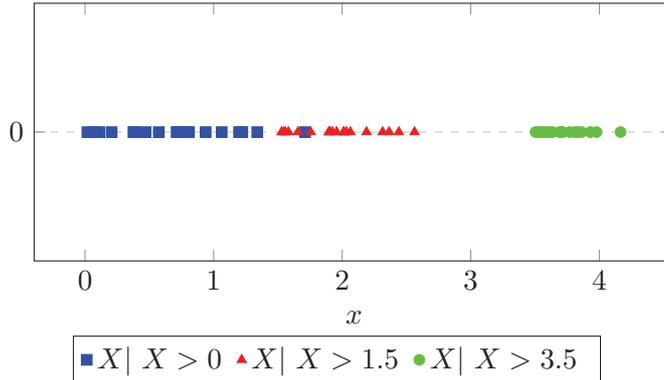


Figure 2: 20 standard normal samples conditionally to be greater than three different thresholds in dimension 1.

Note that this heuristic argument relies on the unimodal nature of the failure domain (see Section 2.5), since otherwise it can easily be put in default (consider for instance the one-dimensional case $\varphi(x) = |x| - 10$ where the conditional variance actually increases). Note also that this unimodal underlying assumption also intuitively guarantees that the conditional mean m^* is non-zero, thereby leading to a proper definition of the projection on its span.

Consistent with our choice of auxiliary family, it is therefore natural that the variance decreases in the direction of m^* , which makes it a good direction to estimate the variance. This does not mean that it is always the best direction, and we provide indeed an example in Section 4.4 where, for instance, $m^* = \|m^*\|e_1$ while the best direction (the one in which the variance evolves the most) is e_2 . However, even in this case, CE- m^* is more efficient than CE in high dimension (updating the full covariance matrix or only the diagonal), see Section 4.4 for more details.

4. Numerical Results

We compare the two proposed algorithms CE- m^* and iCE- m^* to CE and iCE, with full covariance matrix and diagonal covariance, on four examples. The cross-entropy algorithms updating only the diagonal of the covariance are named CE_d and iCE_d, while the ones updating the full covariance are still called CE and iCE. The first example is the linear model already considered above, which illustrates the performance of our algorithms on a simple case where CE has already been shown to fail. The second example considers a modified Ackley function [41], which is a classical test case in optimization. The third example considers an example from finance (see [15, 42]), with a limit-state function which is not differentiable. Finally, the last example considers a simple parabolic function that illustrates potential limits of our methods, and thus also ways of improvement.

In all simulations we use $\rho = 0.1$ for CE and $\delta = 1.5$ for iCE, which are classical choices recommended by the authors in [12] and [13], $\rho = 0.1$ for CE- m^* and CE_d, and $\delta = 3$ for iCE- m^* and iCE_d. For our proposed algorithms, the parameters were selected empirically as the parameters giving the most accurate results. As mentioned in Remarks 1 and 3, all the algorithms are performed without smoothing. The sample size may vary from one algorithm to another because the number of iterations may be different, and it is chosen so that the total average simulation budget of the different algorithms is the same (around

8,000). Each algorithm is repeated 100 times, which produces a sample $\hat{P}_1, \dots, \hat{P}_{100}$ of 100 independent copies of \hat{P} . In the following tables we report five quantities: the mean estimation, the coefficient of variation, the relative bias, the sample size N and the average simulation budget of these 100 samples. When P is known they are defined as

$$\text{Mean} = \frac{1}{100} \sum_{k=1}^{100} \hat{P}_k, \quad \text{C.o.V.} = \frac{\sqrt{\frac{1}{100} \sum_{k=1}^{100} (\hat{P}_k - P)^2}}{P} \quad \text{and} \quad \text{Relative bias} = \frac{\text{Mean} - P}{P},$$

and when P is unknown then it is first-hand estimated by a crude Monte Carlo method with a high simulation budget.

Moreover, we set the maximum number of iterations of the algorithm to 10: if after 10 iterations the algorithm did not stop, we report NC (No Convergence) in the tables. As will be seen below, this is systematically the case for CE and iCE for dimensions 100 and above. Heuristically, if the algorithm does not converge, this means that it did not manage to build an auxiliary density which samples sufficiently often in the failure domain. This may essentially be due to three reasons:

1. either because the estimated mean m_t does not bring the sample close enough to the failure domain: this can typically happen in all six algorithms because the variance decreases, so that increments $m_{t+1} - m_t$ become small and the algorithm gets stuck;
2. or because the weights are collapsing quickly to 0, hence the estimations in (4) are not even possible because of divisions by 0; this mainly happens in CEd and iCEd for very high dimension in our examples;
3. or m_t gets close to the failure domain, but the variance does not “fit” the failure domain: this can typically happen in CE- m^* and iCE- m^* because m^* is not good enough as a direction of projection (see Remark 5 in Section 4.4).

We have to note that the non-convergence of these algorithms for the two first reasons can sometimes be avoided with a smoothing procedure or by varying the parameters ρ (for the CE-based methods), and δ (for the iCE-based methods), but the estimation results stay very inaccurate. Indeed, with a smoothing parameter α between 0.5 and 0.7 the algorithms can converge in less than 10 iterations, and with a simulation budget around 8000, but the probability estimation is far from the true probability (often less than 10^{-10} instead of $\simeq 10^{-3}$). In the same way, decreasing ρ (or increasing δ for iCE) can make the algorithms converge quickly but the probabilities are almost equal to 0, since the estimations are computed with very small samples (e.g. ρN with $\rho = 0.01$ and $N = 4000$). Conversely, increasing ρ (or decreasing δ) does not allow to converge and adds some unnecessary iterations.

4.1. Toy example: sum of independent variables

First we come back to the linear case (6) of Section 3.1 where $\varphi(x) \approx \sum_{j=1}^n x_j - 3\sqrt{n}$ and $P = \mathbb{P}(\varphi(X) \geq 0) = 1.35 \cdot 10^{-3}$, for all dimension n . We already saw in Figure 1 that the performance of CE and iCE degrades significantly as the dimension n increases. Figure 1 also presents results for CE- m^* , iCE- m^* , CEd and iCEd (in dotted lines) which are indistinguishable from those obtained by, for instance, CE with the optimal auxiliary density g_{m^*, Σ^*} . This shows that projecting on m^* stabilizes the estimation up to the dimension $n = 60$ considered in this figure. However, CEd and iCEd have also similar performance on this test case.

		iCE	iCE- m^*	iCEd	CE	CE- m^*	CEd
$n = 30$ $P = 1.35 \cdot 10^{-3}$	Mean ($\times 10^{-3}$)	0.011	1.35	1.36	$5.8 \cdot 10^{-5}$	1.34	1.35
	C.o.V.	99%	2.6%	5.1%	100%	4.6%	5.8%
	Relative bias	-99%	-0.1%	0.9%	-100%	-0.5%	0.4%
	Sample size	1,000	2,700	3,700	1,000	2,700	3,400
	Simu. budget	8,500	8,000	8,000	8,300	8,100	8,000
$n = 100$ $P = 1.35 \cdot 10^{-3}$	Mean ($\times 10^{-3}$)	NC	1.35	1.36	NC	1.35	1.37
	C.o.V.	NC	4.6%	6.8%	NC	4.4%	7.4%
	Relative bias	NC	0.04%	0.9%	NC	-0.08%	1.1%
	Sample size	NC	2,900	3,700	NC	2,700	3,600
	Simu. budget	NC	8,000	8,100	NC	8,100	8,100
$n = 300$ $P = 1.35 \cdot 10^{-3}$	Mean ($\times 10^{-3}$)	NC	1.35	1.36	NC	1.34	1.31
	C.o.V.	NC	7.3%	11.2%	NC	12.6%	27.6%
	Relative bias	NC	-0.2%	0.6%	NC	-0.8%	-3.2%
	Sample size	NC	4,000	3,700	NC	2,700	3,700
	Simu. budget	NC	8,000	8,000	NC	8,100	8,000

Table 1: Simulation results when φ is given by the sum of coordinates (6). Recall that NC means no convergence, see beginning of Section 4.

We report in Table 1 quantitative results for dimensions $n = 30, 100$ and 300 . As suggested by Figure 1, even for dimension $n = 30$ both CE and iCE provide inaccurate estimates with almost -100% relative bias, meaning that almost no sample falls in the failure domain, whereas for the same dimension, the relative bias of the other algorithms is less than 1% with coefficients of variation between 2 and 6% . In dimensions $n = 100$ and 300 , CE and iCE fail to converge. In sharp contrast, the performance of CE- m^* and iCE- m^* remains very good, with relative bias less than 1% and coefficient of variation between 4 and 13% . The C.o.V. and the relative bias of CE_d and iCE_d are slightly larger than for CE- m^* and iCE- m^* respectively, in dimension $n = 100$. For $n = 300$, iCE_d is again slightly less efficient than iCE- m^* (C.o.V. equal to 11.2% against 7.3% , and similar relative bias), and CE_d is clearly less accurate than CE- m^* since the coefficient of variation is twice larger. We suggest that CE- m^* is better than CE_d in high dimension because CE_d updates more parameters which are not influential on the estimation and involve more estimation errors whereas CE- m^* updates the variance in the optimal direction. We note that on this example, iCE- m^* and iCE_d perform better than CE- m^* and CE_d respectively, with lower or comparable coefficient of variation and lower relative bias. Moreover, iCE- m^* converges faster (with less iterations) which makes its sample size larger in dimension $n = 300$.

4.2. Application in optimization: modified Ackley function

The Ackley function [41] is a non-convex function widely used to test optimization algorithms that we use here to test the performance of CE- m^* and iCE- m^* against the dimension. Thus we consider

$$\varphi(x) = 20 \exp \left(-0.2 \sqrt{\frac{1}{n} \sum_{j=1}^n (a_j x_j - 3)^2} \right) + \exp \left(\frac{1}{n} \sum_{j=1}^n \cos(2\pi(a_j x_j - 3)) \right) - c_n \quad (7)$$

		iCE	iCE- m^*	iCEd	CE	CE- m^*	CEd
$n = 30$ $P = 1.64 \cdot 10^{-3}$	Mean ($\times 10^{-3}$)	1.23	1.62	1.65	0.92	1.62	1.66
	C.o.V.	89%	7.6%	9.0%	75%	13.2%	9.9%
	Relative bias	-24.7%	-1.1%	0.6%	-43.9%	-1.2%	0.9%
	Sample size	1,000	2,700	2,700	2,700	2,200	2,700
	Simu. budget	8,100	8,000	8,100	8,100	8,000	8,100
$n = 100$ $P = 1.18 \cdot 10^{-3}$	Mean ($\times 10^{-3}$)	NC	1.17	1.19	NC	1.22	1.20
	C.o.V.	NC	13.0%	10.0%	NC	29.4%	21.4%
	Relative bias	NC	-1.2%	0.9%	NC	3.4%	1.6%
	Sample size	NC	2,700	2,700	NC	2,200	2,700
	Simu. budget	NC	8,100	8,100	NC	8,000	8,100
$n = 200$ $P = 1.72 \cdot 10^{-3}$	Mean ($\times 10^{-3}$)	NC	1.71	1.69	NC	1.72	1.82
	C.o.V.	NC	9.1%	11.6%	NC	26.9%	120%
	Relative bias	NC	-0.5%	-1.5%	NC	-0.05%	5.7%
	Sample size	NC	2,700	2,700	NC	2,700	3,700
	Simu. budget	NC	8,100	8,000	NC	8,100	8,000

Table 2: Simulation results when φ is given by the modified Ackley function (7).

with $a = (a_1, \dots, a_n) = \frac{2}{n}(0, 1, 2, \dots, n-1) \in \mathbb{R}^n$ a fixed vector and c_n a real-valued constant depending on the dimension that we adjust in order to choose the value of the probability. The standard Ackley function corresponds to $a_i = 1$ for all i : for this choice, symmetry considerations imply that $m^* = \|m^*\|(1, \dots, 1)$ is the vector where all coordinates are equal. This is the same as in the previous example, and we introduce the additional parameter a in order to break this symmetry. Indeed, we can numerically estimate $m^* = (m_1^*, \dots, m_n^*)$, for instance by brute force Monte Carlo with a large simulation budget, and we get $m_1^* \approx 0.01 < m_2^* < \dots < m_n^* \approx 0.46$. In addition, the function does not depend on the first variable x_1 , since $a_1 = 0$. We plot the 2-dimensional modified Ackley function in Figure 3, and its projection in the plane $\{x_1 = 0\}$. The failure domain corresponds to the points (x_1, x_2) such that $\varphi(x_1, x_2) \geq 0$, i.e. values of x_2 such that the blue curve is above the red line on the right graph of Figure 3, since the function does not depend on x_1 .

The results are given in Table 2 and are qualitatively similar to those of the previous example with the sum: for dimension $n = 30$, CE and iCE give poor results (although better than for the sum, the absolute relative bias being between 25 and 45%) while CE- m^* , iCE- m^* , CEd, and iCEd give good estimation with a relative bias of 1% or less and coefficient of variation between 7 and 13%. As the dimension increases CE and iCE do not converge and the results for CE- m^* , iCE- m^* , and iCEd stay qualitatively similar: the relative bias remains small, between 1 and 3%, and the coefficient of variation between 9 and 30%. CEd is also very close to CE- m^* in dimension $n = 30$ and 100, and even a little more accurate (since the C.o.V. and the relative bias are smaller), however, in dimension $n = 200$, CEd has large coefficient of variation (120%) and relative bias (5.7%) contrary to CE- m^* . Then, CEd seems to be less efficient than CE- m^* in high dimension because it estimates too many non-influential parameters. Similarly as for the sum, we note that the iCE-based methods perform better than the ones based on CE.

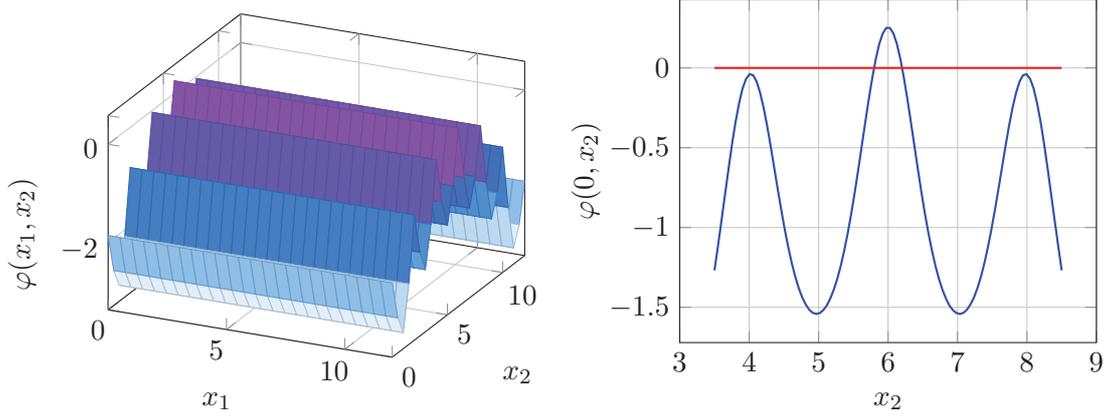


Figure 3: 2-dimensional modified Ackley function 7 (left) and the projection in the plane $\{x_1 = 0\}$ (right). The failure threshold is represented by the red line.

4.3. Application in finance: large portfolio losses

		iCE	iCE- m^*	iCEd	CE	CE- m^*	CEd
$n = 30$ $P = 4.29 \cdot 10^{-3}$	Mean ($\times 10^{-3}$)	0.0028	4.29	4.38	2.7	4.24	4.32
	C.o.V.	100%	10.1%	8.1%	54%	7.2%	9.0%
	Relative bias	-100%	-0.06%	2.1%	-37%	-1.3%	0.8%
	Sample size	1,000	2,700	3,200	2,600	2,700	2,700
	Simu. budget	8,400	7,900	8,000	8,000	8,100	8,100
$n = 100$ $P = 1.8 \cdot 10^{-3}$	Mean ($\times 10^{-3}$)	NC	1.9	1.8	NC	1.8	1.8
	C.o.V.	NC	8.5%	7.1%	NC	8.4%	14.4%
	Relative bias	NC	3.0%	2.0%	NC	1.3%	-1.8%
	Sample size	NC	3,000	2,700	NC	2,700	2,700
	Simu. budget	NC	8,000	8,000	NC	8,100	8,100
$n = 250$ $P = 1.0 \cdot 10^{-5}$	Mean ($\times 10^{-5}$)	NC	1.01	0.049	NC	0.94	NC
	C.o.V.	NC	110%	98%	NC	60%	NC
	Relative bias	NC	1.0%	-95%	NC	-5.9%	NC
	Sample size	NC	1,600	1,500	NC	2,100	NC
	Simu. budget	NC	8,000	8,000	NC	8,100	NC

Table 3: Simulation results when φ is given by the portfolio loss function (8).

Next we consider a financial application taken from [15, 42]. In this example, the failure probability is $P = \mathbb{P}(L(Z) > bn)$, where L is a loss function defined by

$$L(z) = \sum_{j=1}^n \mathbb{I}_{\{z_j \geq 0.5\sqrt{n}\}}.$$

We choose $b = 0.45$ when $n = 30$, and $b = 0.25$ when $n = 100$ and 250. The entries Z_j are dependent real variables defined by

$$Z_j = \left(qU + (1 - q^2)^{1/2} \eta_j \right) \mu^{-1/2}$$

where $U \sim \mathcal{N}(0, 1)$, $\eta_j \sim \mathcal{N}(0, 9)$, $j = 1, \dots, n$, $\mu \sim \text{Gamma}(6, 6)$ are all independent variables, and $q = 0.25$. The same realisation of U and μ is considered for the computation of all the variables Z_j . In our methods, as we work only with normal distributions and our initial density f has to be standard normal, we set $\eta_j = 3\tilde{\eta}_j$, for all $j = 1, \dots, n$ and $\mu = F_\Gamma^{-1}(F_\mathcal{N}(\tilde{\mu}))$ with $\tilde{\eta}_j, \tilde{\mu}$ independent standard normal variables and $F_\Gamma, F_\mathcal{N}$ the cumulative distribution functions of the Gamma(6, 6) and $\mathcal{N}(0, 1)$ respectively. Hence, if we set $X_1 = U, X_2 = \tilde{\mu}$ and $(X_3, \dots, X_{n+2}) = \tilde{\eta} \in \mathbb{R}^n$ and we define

$$\varphi(X) = \sum_{j=1}^n \mathbb{I}_{\{\psi(U, \tilde{\mu}, \tilde{\eta}_j) \geq 0.5\sqrt{n}\}} - bn \quad (8)$$

with

$$\psi(U, \tilde{\mu}, \tilde{\eta}_j) = \left(qU + 3(1 - q^2)^{1/2} \tilde{\eta}_j \right) \left[F_\Gamma^{-1}(F_\mathcal{N}(\tilde{\mu})) \right]^{-1/2}$$

then the sought probability $P = \mathbb{P}(L(Z) > bn)$ gets rewritten as $P = \mathbb{P}(\varphi(X) > 0)$ with X a standard Gaussian vector in dimension $n+2$. Then we test the performance of CE- m^* and iCE- m^* against CE, iCE, CE_d and iCE_d for dimension $n = 30, 100$ and 250 . Results are gathered in Table 3.

Similarly as in the previous examples, iCE fails in dimension $n = 30$ while CE leads to an inaccurate estimate with a relative bias of -37% and a coefficient of variation of 54% . For the same dimension and with a similar simulation budget, CE- m^* , iCE- m^* , CE_d and iCE_d give an estimation with relative bias less than 2% and coefficient of variation less than 10% . In dimension $n = 100$, CE and iCE fail to converge while the four other algorithms remain accurate with a relative bias less than 3% and a coefficient of variation of 8% , except CE_d which has 14.4% of C.o.V. and is slightly less accurate. Although CE- m^* and iCE- m^* behave similarly in dimension $n = 30$ and 100 , in dimension $n = 250$ CE- m^* behaves better than iCE- m^* with a comparable relative bias (6% vs. 1%) but a coefficient of variation which is twice smaller (60% vs. 110%). On the other hand, in dimension $n = 250$, we observe that CE_d doesn't converge because of the fast weight degeneracy to 0, and iCE_d fail to estimate accurately the probability with a relative bias of -95% .

4.4. Example where m^* is not the best direction

We conclude this numerical section with an example showing the robustness of CE- m^* and iCE- m^* . This is an example where, although m^* is clearly not the best direction in which to estimate the variance, they do still improve on CE and iCE in high dimension. Consider the function

$$\varphi(x) = x_1 - 3x_2^2 - 3, \quad x \in \mathbb{R}^n \quad (9)$$

whose failure domain is represented in Figure 4. By symmetry it is clear that under the optimal auxiliary density g^* , which is the distribution conditional on failure, the average m^* is of the form $m^* = \alpha e_1$ with $\alpha \approx 3.4$ numerically evaluated and the variance Σ^* is equal to the identity, except for the variance terms associated to the first two coordinates, which can be numerically estimated to $\Sigma_{11}^* \approx 0.095$ and $\Sigma_{22}^* \approx 0.044$. Once the mean is within the failure domain, it is clear that one should decrease the variance in the direction of e_2 because the failure domain is quite narrow in this direction, as is reflected by the fact that $\Sigma_{22}^* < \Sigma_{11}^*$. And indeed, one can check that IS with such an auxiliary density behaves better. However, CE- m^* only updates the variance in the (approximate) direction of e_1

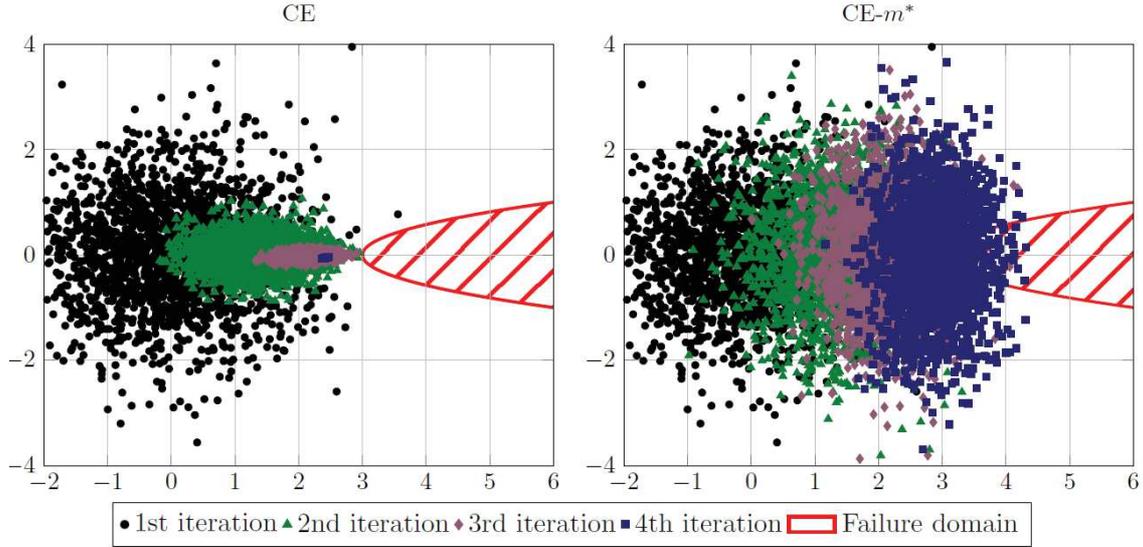


Figure 4: Projection on the (x_1, x_2) plane of the failure domain of the function $\varphi(x) = x_1 - 3x_2^2 - 3$ considered in (9) and the samples of every iteration of CE (left) and $CE-m^*$ (right), with dimension $n = 100$. Since φ does not depend on other coordinates, the full failure domain is a tube around this two-dimensional shape.

and not e_2 , so the density is not optimal to estimate the probability of failure, as we can see in Figure 4. Nevertheless, in high dimension $CE-m^*$ converges and is able to estimate the probability whereas CE doesn't because of the fast covariance matrix shrinkage. From a theoretical point of view, we do not fully understand why the covariance matrix collapses for CE and $CE-m^*$. Preliminary investigation suggests that this has to do with errors made in covariance terms, which are larger for CE than $CE-m^*$. The reason why $CE-m^*$ avoids the covariance matrix shrinkage problem constitutes in our view an interesting research question.

We can see it in the simulation results given in Table 4. Similarly to the previous examples, in dimension $n = 30$, CE and iCE already give inaccurate estimations with 100% of coefficient of variation and -100% of relative bias, whereas $CE-m^*$ and $iCE-m^*$ are very efficient (around 11% of coefficient of variation and between 0.5 and 1% of relative bias). iCEd is even more accurate with 5.2% of C.o.V. and 0.2% of relative bias but CED isn't so efficient since its coefficient of variation is 37% and its relative bias is -11% . For higher dimension, CE and iCE fail to converge. In dimension $n = 100$, $iCE-m^*$ and iCEd are still very accurate with a coefficient of variation between 7 and 12%, whereas $CE-m^*$ and CED gives 28.3% and 87% of C.o.V. respectively. CED is particularly inaccurate since its relative bias is -52% , while the three other convergent algorithms have a relative bias lower than 1.3%. For dimension $n = 300$, iCEd and CED doesn't converge anymore, because of the weight degeneracy, whereas $CE-m^*$ and $iCE-m^*$ are still able to estimate the probability with a small relative bias (-1.5% and 3.5% respectively), although the coefficient of variation can be large, especially for $CE-m^*$ (88% against 29% for $iCE-m^*$).

Hence, although m^* is not the optimal direction to estimate the covariance matrix, this example shows that $CE-m^*$ and $iCE-m^*$ are able to efficiently estimate the probability in high dimension while CE, iCE, CED and iCED fail.

Remark 5. If we consider a narrower parabola, (for example with $\varphi(x) = x_1 - 25x_2^2 - 3$), then $CE-m^*$ is not converging anymore. Indeed, in such cases the final density does not

		iCE	iCE- m^*	iCEd	CE	CE- m^*	CEd
$n = 30$ $P = 2.9 \cdot 10^{-4}$	Mean ($\times 10^{-4}$)	$2 \cdot 10^{-6}$	2.9	2.9	$9 \cdot 10^{-8}$	2.9	2.6
	C.o.V.	100%	11.4%	5.2%	100%	11.2%	37%
	Relative bias	-100%	1.0%	0.2%	-100%	0.5%	-11%
	Sample size	1,000	2,700	2,700	1,000	1,600	2,000
	Simu. budget	8,800	8,000	8,100	8,800	7,900	8,000
$n = 100$ $P = 2.9 \cdot 10^{-4}$	Mean ($\times 10^{-4}$)	NC	2.9	2.9	NC	3.0	1.4
	C.o.V.	NC	11.3%	7.2%	NC	28.3%	87%
	Relative bias	NC	1.1%	-0.3%	NC	1.3%	-52%
	Sample size	NC	2,700	2,600	NC	1,900	2,000
	Simu. budget	NC	8,000	8,000	NC	8,100	8,000
$n = 300$ $P = 2.9 \cdot 10^{-4}$	Mean ($\times 10^{-4}$)	NC	3.0	NC	NC	2.9	NC
	C.o.V.	NC	29.2%	NC	NC	88%	NC
	Relative bias	NC	3.5%	NC	NC	-1.5%	NC
	Sample size	NC	2,300	NC	NC	2,400	NC
	Simu. budget	NC	8,100	NC	NC	8,000	NC

Table 4: Simulation results when φ is given by the parabola function (9).

“fit” the failure domain and most of the samples fall outside the domain. This happens because m^* is not the best direction in which updating the covariance matrix. So this is a case where CE- m^* does not improve CE performance. However, CE, CEd, iCE and iCEd do not converge either in high dimension on this example. Moreover, in high dimension, we did not find any test case (satisfying the unimodal assumption) where CE- m^* degrades CE performance. And as shown in the above numerical examples, most of the times it greatly improves it.

5. Conclusion

In this paper, we present a new CE-based importance sampling method using a one-dimensional projection that provides accurate estimation of rare event probabilities in high dimension, provided the target optimal auxiliary density is unimodal. At each step of the algorithm, the idea is to only estimate variance and covariance terms in the direction of the current estimate of the best auxiliary parameter m^* : the light tail of the Gaussian distribution guarantees that the variance decreases in this direction which ensures that there will be more information than noise in estimating the variance in this direction. Empirically, our results show that, in high dimension, it is much more efficient to only estimate variance terms in one “good” direction than blindly in the whole space as is done in CE, or to estimate the variance diagonal terms (as in CEd). Although extremely simple, our algorithm turns out to be very efficient, all the more considered that it does not require gradient estimation. Indeed, we find that it cannot degrade CE performance and that in most cases, it actually greatly improves it, typically providing accurate estimates where CE simply failed to converge. Moreover, in moderately high dimension (say until 100), CE- m^* is as efficient as CEd, and it is often more accurate in higher dimension (more than 200).

This method opens many interesting research direction which we aim to pursue. First,

it would be interesting to get stronger theoretical justification of the accuracy and robustness of our algorithm. It seems in particular that it avoids the well-known shrinkage problem of CE because it estimates less covariance terms, and it would be interesting to confirm this hypothesis. Moreover, the intuition behind the choice of m^* as a direction to project could be deepened to select alternative, possibly better directions. The intuition being that the variance evolves in the direction of m^* , this naturally suggests to look at specific eigenvalues of the covariance matrix, which we are currently investigating. Finally, it would be interesting to extend our approach to the multi-modal case by trying to combine it with existing algorithms developed in this case in low dimensions.

Appendix A.

In this appendix we simply recall the basic CE and iCE algorithms.

Algorithm 3: CE method [11]

- Data:** parameter $\rho \in (0, 1)$, and sample size N
Result: Estimation \hat{P} of the failure probability $\mathbb{P}_f(\varphi(X) \geq 0)$
- 1 Initialization: set $t = 0$, $m_t = 0$ and $\Sigma_t = I_n$;
 - 2 Generate X_1, \dots, X_N independently according to g_{m_t, Σ_t} ;
 - 3 Compute $\varphi_i = \varphi(X_i)$;
 - 4 Sort the samples in ascending order: $\varphi_{(1)} \leq \dots \leq \varphi_{(N)}$, and set $\gamma_t = \varphi_{(\lfloor (1-\rho)N \rfloor)}$;
 - 5 Compute the weights $w_i = w'_i / \sum_{j=1}^n w'_j$ with $w'_i = \mathbb{I}_{\{\varphi_i \geq \gamma_t\}} L_t(X_i)$;
 - 6 **while** $\gamma_t < 0$ **do**
 - 7 Compute $m_{t+1} = \sum_{i=1}^N w_i X_i$ and $\Sigma_{t+1} = \sum_{i=1}^N w_i (X_i - m_{t+1})(X_i - m_{t+1})^\top$;
 - 8 Set $t \leftarrow t + 1$;
 - 9 Repeat Steps 2, 3, 4 and 5;
 - 10 **end**
 - 11 Return $\hat{P} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}_{\{\varphi(X_i) \geq 0\}} L_t(X_i)$.
-

Algorithm 4: Improved CE method of [13].

Data: parameter δ_{target} , sample size N
Result: Estimation \hat{P} of the failure probability $\mathbb{P}_f(\varphi(X) \geq 0)$

- 1 Initialization: set $t = 0$, $m_t = 0$, $\Sigma_t = I_n$ and $\sigma_t = \infty$;
- 2 Generate X_1, \dots, X_N independently according to g_{m_t, Σ_t} ;
- 3 Compute $\varphi_i = \varphi(X_i)$ for $i = 1, \dots, N$;
- 4 Compute \widehat{cv} the empirical coefficient of variation of the $\mathbb{I}_{\{\varphi_i \geq 0\}}/F(\varphi_i/\sigma_t)$;
- 5 **while** $\widehat{cv} \geq \delta_{\text{target}}$ **do**
- 6 Compute $\sigma_{t+1} = \arg \min(\widehat{\delta}_t(\sigma) - \delta_{\text{target}})^2$ where the minimum is taken over $\sigma \in (0, \sigma_t)$ and $\widehat{\delta}_t(\sigma)$ is the coefficient of variation of the $F(\varphi_i/\sigma)L_t(X_i)$ with $L_t = f/g_{m_t, \Sigma_t}$;
- 7 Compute the weights $w_i = w'_i / \sum_j w'_j$ with $w'_i = F(\varphi_i/\sigma_{t+1})L_t(X_i)$;
- 8 Compute $m_{t+1} = \sum_{i=1}^N w_i X_i$ and $\Sigma_{t+1} = \sum_{i=1}^N w_i (X_i - m_{t+1})(X_i - m_{t+1})^\top$;
- 9 Set $t \leftarrow t + 1$;
- 10 Repeat Steps 2, 3 and 4;
- 11 **end**
- 12 Return $\hat{P} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}_{\{\varphi(X_i) \geq 0\}} L_t(X_i)$.

References

- [1] J.-M. Bourinet, “Reliability analysis and optimal design under uncertainty-focus on adaptive surrogate-based approaches,” 2018. Université Clermont Auvergne. Habilitation à diriger des recherches.
- [2] J. Morio, M. Balesdent, D. Jacquemart, and C. Vergé, “A survey of rare event simulation methods for static input–output models,” *Simulation Modelling Practice and Theory*, vol. 49, pp. 287–304, Dec. 2014.
- [3] A. Der Kiureghian, “First-and second-order reliability methods,” *Engineering design reliability handbook*, vol. 14, 2005. CRC Press Boca Raton, FL.
- [4] R. Y. Rubinstein and D. P. Kroese, *Simulation and the Monte Carlo method*. Wiley series in probability and statistics, Hoboken, New Jersey: Wiley, third edition ed., 2017. OCLC: 994882892.
- [5] S.-K. Au and J. L. Beck, “Estimation of small failure probabilities in high dimensions by subset simulation,” *Probabilistic Engineering Mechanics*, vol. 16, pp. 263–277, Oct. 2001.
- [6] S. Song, Z. Lu, and H. Qiao, “Subset simulation for structural reliability sensitivity analysis,” *Reliability Engineering & System Safety*, vol. 94, no. 2, pp. 658–665, 2009.
- [7] Z.-M. Jiang, D.-C. Feng, H. Zhou, and W.-F. Tao, “A recursive dimension-reduction method for high-dimensional reliability analysis with rare failure event,” *Reliability Engineering & System Safety*, p. 107710, 2021.
- [8] A. Owen and Y. Zhou, “Safe and Effective Importance Sampling,” *Journal of the American Statistical Association*, vol. 95, pp. 135–143, Mar. 2000.

- [9] S. Agapiou, O. Papaspiliopoulos, D. Sanz-Alonso, and A. M. Stuart, “Importance Sampling : Intrinsic Dimension and Computational Cost,” *Statistical Science*, vol. 32, no. 3, pp. 405–431, 2017.
- [10] M. Moustapha, S. Marelli, and B. Sudret, “A generalized framework for active learning reliability: survey and benchmark,” *arXiv preprint arXiv:2106.01713*, 2021.
- [11] R. Y. Rubinstein and D. P. Kroese, *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation and machine learning*. New York; London: Springer, 2011. OCLC: 1064531163.
- [12] P.-T. de Boer, D. P. Kroese, S. Mannor, and R. Y. Rubinstein, “A Tutorial on the Cross-Entropy Method,” *Annals of Operations Research*, vol. 134, pp. 19–67, Feb. 2005.
- [13] I. Papaioannou, S. Geyer, and D. Straub, “Improved cross entropy-based importance sampling with a flexible mixture model,” *Reliability Engineering & System Safety*, vol. 191, p. 106564, Nov. 2019.
- [14] Q. D. Cao and Y. Choe, “Cross-entropy based importance sampling for stochastic simulation models,” *Reliability Engineering & System Safety*, vol. 191, p. 106526, 2019.
- [15] J. C. C. Chan and D. P. Kroese, “Improved cross-entropy method for estimation,” *Statistics and Computing*, vol. 22, pp. 1031–1040, Sept. 2012.
- [16] C. Mattrand and J.-M. Bourinet, “The cross-entropy method for reliability assessment of cracked structures subjected to random markovian loads,” *Reliability Engineering & System Safety*, vol. 123, pp. 171–182, 2014.
- [17] M. F. Bugallo, V. Elvira, L. Martino, D. Luengo, J. Miguez, and P. M. Djuric, “Adaptive Importance Sampling: The past, the present, and the future,” *IEEE Signal Processing Magazine*, vol. 34, pp. 60–79, July 2017.
- [18] F. Portier and B. Delyon, “Asymptotic optimality of adaptive importance sampling,” *Advances in Neural Information Processing Systems*, vol. 31, pp. 3134–3144, 2018.
- [19] J. Morio, “Non-parametric adaptive importance sampling for the probability estimation of a launcher impact position,” *Reliability engineering & system safety*, vol. 96, no. 1, pp. 178–183, 2011.
- [20] J. C. Neddermeyer, “Computationally efficient nonparametric importance sampling,” *Journal of the American Statistical Association*, vol. 104, no. 486, pp. 788–802, 2009.
- [21] S. Xiao, S. Oladyshkin, and W. Nowak, “Reliability analysis with stratified importance sampling based on adaptive kriging,” *Reliability Engineering & System Safety*, vol. 197, p. 106852, 2020.
- [22] X. Zhang, L. Wang, and J. D. Sørensen, “Akois: an adaptive kriging oriented importance sampling method for structural system reliability analysis,” *Structural Safety*, vol. 82, p. 101876, 2020.

- [23] M. Munoz Zuniga, A. Murangira, and T. Perdrizet, “Structural reliability assessment through surrogate based importance sampling with dimension reduction,” *Reliability Engineering & System Safety*, vol. 207, p. 107289, 2021.
- [24] X. Yang, X. Cheng, T. Wang, and C. Mi, “System reliability analysis with small failure probability based on active learning kriging model and multimodal adaptive importance sampling,” *Structural and Multidisciplinary Optimization*, pp. 1–16, 2020.
- [25] M. Menz, S. Dubreuil, J. Morio, C. Gogu, N. Bartoli, and M. Chiron, “Variance based sensitivity analysis for monte carlo and importance sampling reliability assessment with gaussian processes,” *Structural Safety*, vol. 93, p. 102116, 2021.
- [26] L. Katafygiotis and K. Zuev, “Geometric insight into the challenges of solving high-dimensional reliability problems,” *Probabilistic Engineering Mechanics*, vol. 23, pp. 208–218, Apr. 2008.
- [27] S. Au and J. Beck, “Important sampling in high dimensions,” *Structural Safety*, vol. 25, pp. 139–163, Apr. 2003.
- [28] T. Bengtsson, P. Bickel, and B. Li, *Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems*, vol. Volume 2 of *Collections*, pp. 316–334. Beachwood, Ohio, USA: Institute of Mathematical Statistics, 2008.
- [29] O. Zahm, T. Cui, K. Law, A. Spantini, and Y. Marzouk, “Certified dimension reduction in nonlinear bayesian inverse problems,” *arXiv preprint arXiv:1807.03712*, 2018.
- [30] F. Uribe, I. Papaioannou, Y. M. Marzouk, and D. Straub, “Cross-entropy-based importance sampling with failure-informed dimension reduction for rare event simulation,” *arXiv:2006.05496 [stat]*, June 2020. arXiv: 2006.05496.
- [31] Y. El-Laham, V. Elvira, and M. F. Bugallo, “Robust covariance adaptation in adaptive importance sampling,” *IEEE Signal Processing Letters*, vol. 25, no. 7, pp. 1049–1053, 2018.
- [32] R. Y. Rubinstein and P. W. Glynn, “How to Deal with the Curse of Dimensionality of Likelihood Ratios in Monte Carlo Simulation,” *Stochastic Models*, vol. 25, pp. 547–568, Nov. 2009.
- [33] Z. Wang and J. Song, “Cross-entropy-based adaptive importance sampling using von Mises-Fisher mixture for high dimensional reliability analysis,” *Structural Safety*, vol. 59, pp. 42–52, Mar. 2016.
- [34] S. Geyer, I. Papaioannou, and D. Straub, “Cross entropy-based importance sampling using Gaussian densities revisited,” *Structural Safety*, vol. 76, pp. 15–27, Jan. 2019.
- [35] P.-L. Liu and A. Der Kiureghian, “Multivariate distribution models with prescribed marginals and covariances,” *Probabilistic Engineering Mechanics*, vol. 1, no. 2, pp. 105–112, 1986.
- [36] M. Hohenbichler and R. Rackwitz, “Non-normal dependent vectors in structural safety,” *Journal of the Engineering Mechanics Division*, vol. 107, no. 6, pp. 1227–1238, 1981.

- [37] A. Costa, O. D. Jones, and D. Kroese, “Convergence properties of the cross-entropy method for discrete optimization,” *Operations Research Letters*, vol. 35, pp. 573–580, Sept. 2007.
- [38] L. Katafygiotis and K. Zuev, “Estimation of small failure probabilities in high dimensions by adaptive linked importance sampling,” *COMPADYN 2007*, 2007.
- [39] Y. El-Laham, V. Elvira, and M. F. Bugallo, “Robust Covariance Adaptation in Adaptive Importance Sampling,” *IEEE Signal Processing Letters*, vol. 25, pp. 1049–1053, July 2018.
- [40] “Cross entropy method and improved cross entropy method.” <https://www.bgu.tum.de/era/software/software00/cross-entropy-method-and-improved-cross-entropy-method/>.
- [41] “Virtual library of simulation experiments, ackley function.” <https://www.sfu.ca/~ssurjano/ackley.html>.
- [42] A. Bassamboo, S. Juneja, and A. Zeevi, “Portfolio Credit Risk with Extremal Dependence: Asymptotic Analysis and Efficient Simulation,” *Operations Research*, vol. 56, pp. 593–606, June 2008.