Simultaneous estimation of complementary moment independent sensitivity measures for reliability analysis

Pierre Derennes[a,b], Jérôme Morio[b], Florian Simatos[c]

[a]*Université de Toulouse, UPS IMT, F-31062 Toulouse Cedex 9, France*
[b]*ONERA/DTIS, Université de Toulouse, F-31055 Toulouse, France*
[c]*ISAE-SUPAERO and Université de Toulouse, Toulouse, France*

Abstract

In reliability-based design, the estimation of the failure probability is a crucial objective. However, focusing only on the occurrence of the failure event may be insufficient to entirely characterize the reliability of the considered system. This paper provides a common estimation scheme of two complementary moment independent sensitivity measures, allowing to improve the understanding of the system's reliability. Numerical applications are performed in order to show the effectiveness of the proposed estimation procedure.

Keywords: Moment independent importance measure, reliability, target and conditional sensitivity analysis, Sequential Monte Carlo, Subset simulation, Maximum entropy principle.

1. Introduction

1.1. Context

Reliability analysis is broadly concerned with the failure analysis of physical systems. Given a part that plays a critical role for the system, a typical question in reliability concerns the study of its failure probability, for instance its estimation or how it is influenced by some of the system's components. Recently this scope has expanded, with practitioners being increasingly interested not only in the failure probability but also in the system's behavior upon failure. However, depending on the viewpoint considered, the same question may lead to widely different answers.

To illustrate this point, consider for instance the following simple toy model:

$$Y = X_1 + \mathbb{1}_{X_1 > 3} |X_2| \qquad (1)$$

where X_1 and X_2 are independent centered Gaussian random variables with respective variance 1 and 5. The random variables X_1 and X_2 are viewed as the system input and Y as the system output. Let us consider that for this system, $\{Y > 3\}$ is the failure event, and try to answer the following question: which out of X_1 and X_2 is more important from a reliability perspective? Actually, the answer depends on the viewpoint considered:

- if one is interested in the **impact of the input on the failure occurring or not**, then of course X_1 is highly influential and X_2, that only kicks in Y upon failure, plays no role;

- if one is now interested in the most influential input **upon failure** occurring, then X_2 should intuitively be more important than X_1 because of its higher variance.

Because of these different objectives, various reliability indices have been proposed, each with its own estimation strategy. In this paper we focus on moment-independent indices, that have recently attracted increasing attention in order to alleviate some of the limitations of classical variance-based indices. Our main message is that several of these indices can be efficiently estimated simultaneously with only one run of a Sequential Monte Carlo (SMC) algorithm, one of the most efficient methods for estimating probabilities of rare events. The crucial observation is that once we have run an SMC algorithm, no more calls to the (supposedly expensive) black box function $\mathcal{M}$ are needed.

1.2. Two complementary moment independent sensitivity measures

In this paper we focus on Borgonovo's indices originally proposed in [1], although our method can be generalized to more general indices as discussed in Section 4. Let in the sequel $Z \mid Z'$ denote a random variable with random distribution the distribution of Z conditioned on Z'. To measure the sensitivity of the output $Y = \mathcal{M}(\mathbf{X})$ with respect to one of its input X_i, where $\mathbf{X} = (X_1, \ldots, X_d)$, Borgonovo [1] proposed in the case where (X_i, Y) is absolutely continuous with respect to Lebesgue measure the index

$$\delta_i = \frac{1}{2} \mathbb{E} \left[\left\| f_Y - f_{Y|X_i} \right\|_{L^1(\mathbb{R})} \right], \tag{2}$$

i.e., half the average of the L^1 distance between the density of Y and the random density of Y conditioned on X_i. If X_i has a high influence on Y, the conditional density should be different from the non-conditioned one and δ_i should thus take large values. For further references and more details on δ-sensitivity measures the reader can consult [2].

In this paper we will adopt a more general definition of Borgonovo's index, which will make it possible to consider cases where (X_i, Y) is not absolutely continuous with respect to Lebesgue measure. The motivation stems from considering the influence of X_i not only on Y but on possibly discrete functions of Y such as $\mathbb{1}_{Y>S}$, which captures the influence of X_i on the failure occurring or not.

For this generalization, we see Borgonovo's index as a measure of dependency between X_i and Y. Namely, let $d_{\text{TV}}(Z_1, Z_2)$ denote the total variation distance between the distributions of the random variables Z_1 and Z_2. When Z_1 and Z_2 are absolutely continuous with respect to Lebesgue measure, we have $d_{\text{TV}}(Z_1, Z_2) = \frac{1}{2} \| f_{Z_1} - f_{Z_2} \|_{L^1(\mathbb{R})}$ and so we adopt the following generalization of Borgonovo's index:

$$\delta_i = \mathbb{E} \left[d_{\text{TV}} \left(Y, Y \mid X_i \right) \right] = d_{\text{TV}} \left((X_i, Y), (X_i, Y') \right) \tag{3}$$

where the second equality holds when (X_i, Y) is absolutely continuous with respect to some product measure $\lambda(\mathrm{d}x) \otimes \mu(\mathrm{d}y)$ (typically, (X_i, Y) is absolutely continuous with respect to Lebesgue measure, or X_i is and Y is a discrete random variable).

In a reliability context, we are interested in the impact of X_i not only on Y but also on the occurrence of some rare event which we write $\{Y > S\}$. This means that we are interested in the influence of X_i on the random variable $\mathbb{1}_{Y>S}$: the corresponding generalized Borgonovo's index is therefore given by

$$\eta_i = \mathbb{E}\left[d_{\mathrm{TV}}(\mathbb{1}_{Y>S}, \mathbb{1}_{Y>S} \mid X_i)\right] = \mathbb{E}\left[|\mathbb{P}(Y > S) - \mathbb{P}(Y > S \mid X_i)|\right] \tag{4}$$

which is actually twice the index proposed in Cui et al. [7]. One of the drawback of this index is that it is unnormalized as it is upper bounded by twice the rare event probability $2\mathbb{P}(Y > S)$. To obtain a $[0, 1]$-valued index, we use the relation

$$\eta_i = 2\mathbb{P}(Y > S) \times d_{\mathrm{TV}}(X_i, X_i \mid Y > S) \tag{5}$$

observed in [18] and that can be derived using Bayes' Theorem, to propose the $[0, 1]$-valued index

$$\bar{\eta}_i = d_{\mathrm{TV}}(X_i, X_i \mid Y > S) = \frac{1}{2}\left\|f_{X_i} - f_{X_i|Y>S}\right\|_{L^1(\mathbb{R})}. \tag{6}$$

Complementary to this approach, we may also be interested in the influence of X_i upon failure, which corresponds to considering δ_i but when all the random variables involved are conditioned upon the failure $Y > S$. Thus, this conditional index, denoted by δ_i^f, is given by

$$\delta_i^f = \mathbb{E}\left[d_{\mathrm{TV}}(Y \mid Y > S, Y \mid Y > S, X_i)\right]. \tag{7}$$

When (X_i, Y) is absolutely continuous, this is a particular case of (3) and so if we denote by $(\tilde{X}_i, \tilde{Y})$ a random variable distributed as (X_i, Y) conditioned on $Y > S$, then we have in this case

$$\delta_i^f = \frac{1}{2}\left\|f_{\tilde{X}_i, \tilde{Y}} - f_{\tilde{X}_i} f_{\tilde{Y}}\right\|_{L^1(\mathbb{R}^2)}. \tag{8}$$

Instead of focusing on Y, the indices η_i and $\bar{\eta}_i$ target a different output, namely $\mathbb{1}_{Y>S}$ and will thus be referred as target indices. Similarly, instead of working in the normal mode, the indices δ_i^f are concerned with the system conditioned upon failure and will thus be referred to as conditional indices. See Section 4 for more on this terminology.

For the toy model (1), we have $Y > S$ if and only if $X_1 > S$: this directly implies $\mathbb{P}(Y > S \mid X_1) = \mathbb{1}_{X_1>S}$ and $\mathbb{P}(Y > S \mid X_2) = \mathbb{P}(Y > S)$ and then

$$\bar{\eta}_1 = 1 - \mathbb{P}(X_1 > S) \approx 0.9987 \quad \text{and} \quad \bar{\eta}_2 = 0.$$

This confirms the intuition that, as far as we are concerned with the failure occurring or not, X_1 is highly influential and X_2, not at all. However, in this simple Gaussian case we can directly compute the δ_i^f's through numerical integration, which gives

$$\delta_1^f \approx 0.0781 \text{ and } \delta_2^f \approx 0.7686.$$

Thus upon failure, X_2 has become much more influential than X_1. This simple toy example illustrates the complementarity of the indices $\bar{\eta}_i$ and δ_i^f from a reliability perspective, and our goal in this paper is to show how they can be simultaneously and accurately estimated with only one run of an adaptive SMC algorithm, usually used in this context for estimating the rare event probability $\mathbb{P}(Y > S)$. In other words, we show that upon estimating this probability, one also gets "for free" an estimation of $\bar{\eta}_i$ and δ_i^f.

1.3. Organization of the paper

In Section 2 we present our simultaneous estimation scheme for $\bar{\eta}_i$ and δ_i^f measures and numerical applications are performed in Section 3 to assess its efficiency. Section 4 discusses how this scheme can be extended to a more general context. Our estimation scheme relies on the maximum entropy method which is recalled in Appendix A.

2. Simultaneous estimation of δ_i^f and $\bar{\eta}_i$

We consider throughout this article a general computer code $Y = \mathcal{M}(\mathbf{X})$ where the scalar output Y depends on a d-dimensional real valued random variable $\mathbf{X} = (X_1, \ldots, X_d) \in \mathbb{R}^d$ through a deterministic scalar function $\mathcal{M} : \mathbb{R}^d \to \mathbb{R}$ called "black box". Without loss of generality, it is assumed that the failure event corresponds to the exceeding of a critical threshold S by the output Y, i.e., is of the form $\{Y > S\}$.

We further assume that for every i, (X_i, Y) is absolutely continuous with respect to Lebesgue measure with density $f_{X_i, Y}$ and marginals f_{X_i} and f_Y. As above, we denote by $\tilde{\mathbf{X}} = (\tilde{X}_1, \ldots, \tilde{X}_d)$ a random variable distributed as $\mathbf{X}$ conditioned on $Y > S$ and define $\tilde{Y} = \mathcal{M}(\tilde{\mathbf{X}})$. Thus, $(\tilde{X}_i, \tilde{Y})$ is also absolutely continuous with respect to Lebesgue measure with density $f_{\tilde{X}_i, \tilde{Y}}$ with marginals $f_{\tilde{X}_i}$ and $f_{\tilde{Y}}$. Our simultaneous estimation scheme is obtained by combining state-of-the-art estimation techniques which we recall next.

2.1. Estimation of δ_i

Initial estimations of δ-sensitivity measures relied on their original definition in terms of total variation distance between conditional and unconditional distributions. Involving L^1 norms of differences of conditional and unconditional output probability density functions, this approach typically necessitates expensive double-loop estimation procedures with a prohibitive cost [1, 14, 15]. Alternative approaches were proposed in [21, 22], but these two methods rest on strong technical assumptions such as independence between input or approximation of the black box $\mathcal{M}$ within the cut-HDMR framework. An apparently efficient single-loop method was proposed in [20], but simulation results provided in [9] questioned its consistency. The interested reader is for instance referred to the introduction of [10] for a more detailed discussion on these estimation issues.

In the present paper, the estimation of δ_i is performed by using the method described in [10]: it does not rely on any assumption on the model and works in particular for dependent input. It rests on the copula-representation of δ_i noted in [19], namely

$$\delta_i = \frac{1}{2} \int_{0 \le u,v \le 1} |c_i(u,v) - 1| \, \mathrm{d}u\mathrm{d}v, \tag{9}$$

where c_i is the density copula of (X_i, Y), i.e., the density of $(F_{X_i}(X_i), F_Y(Y))$. Based on this representation, the approximation proposed in [10] uses a maximum entropy estimation $\hat{c}_i$ of c_i imposing estimated fractional moments as constraints, and then a Monte Carlo estimation $\frac{1}{2N'} \sum_{k=1}^{N'} |\hat{c}_i(U_1^k, U_2^k) - 1|$ of the integral with the (U_1^k, U_2^k) being i.i.d. random variables uniformly distributed on $[0,1]^2$.

At this point we stress an important point: all these estimation techniques assume that one can sample from the input distribution $\mathbf{X}$. As explained in the introduction however, estimating δ_i^f amounts to applying these techniques when the input distribution is that of $\mathbf{X}$ conditioned on failure, which is in general unknown. Thus, before applying these methods one needs to be able to sample from $\tilde{\mathbf{X}}$.

2.2. Sampling from $\tilde{\mathbf{X}}$

The most naive method for generating failure samples is the rejection method. For a given sample $(\mathbf{X}^1, \ldots, \mathbf{X}^N)$ i.i.d. with common distribution $f_{\mathbf{X}}$, a subsample is obtained by recording samples which satisfy $\mathcal{M}(\mathbf{X}^k) > S$. However, this approach leads to a huge computational cost when the failure probability is low. When some information is known on the failure event, this cost can be reduced by leveraging "good" auxiliary distributions in importance sampling techniques. In reliability, a method widely used for designing auxiliary distributions is shifting the input distribution to a design point, which may be determined thanks to FORM/SORM methods. Nevertheless, finding an efficient auxiliary distribution remains challenging when the failure domain is disconnected or when the input dimension is significant.

When no good auxiliary distribution is available, one has to resort to Markov chain Monte Carlo methods. Concerning the estimation of rare event probabilities, one of the most efficient method is the adaptive Sequential Monte Carlo (SMC) procedure proposed and studied in [5] that we present next. In what follows, by duplicating a finite set $\{x_k\}$ into N, we mean drawing N times independently and uniformly from $\{x_k\}$. The algorithm parameters are N_x, ρ, A_x and T, corresponding respectively to the number of particles, the threshold for the quantile, the number of steps of the Metropolis–Hastings sampler, and the exploration (or proposal) kernel in this sampler.

Initialization: set $p = 0$, generate $(\mathbf{X}_p^1, \ldots, \mathbf{X}_p^{N_x})$ i.i.d. according to $f_{\mathbf{X}}$ and compute $Y_p^k = \mathcal{M}(\mathbf{X}_p^k)$ for $k = 1, \ldots, N_x$;

Selection: let γ_p be the ρ-quantile of the Y_p^k: if $\gamma_p > S$, then stop, otherwise duplicate the ρN_x particles with $Y^k > \gamma_p$ into N_x particles;

Mutation: apply A_x times the Metropolis–Hastings algorithm with exploration kernel T and target distribution $\mathbf{X} \mid \mathcal{M}(\mathbf{X}) > \gamma_p$ to each of the N_x particles, denote by $(\mathbf{X}_{p+1}^1, \ldots, \mathbf{X}_{p+1}^{N_x})$ the newly obtained particles with corresponding $Y_{p+1}^k = \mathcal{M}(\mathbf{X}_{p+1}^k)$, increment p and go back to the selection step.

The black box is called for every particle at every step of the Metropolis–Hastings sampler in order to compute the acceptance probability, so that if m denotes the (random) number of steps of this algorithm, then the number of calls to the black box $\mathcal{M}$ is equal to $N_x(1 + mA_x)$.

As noted in [5], at the end of this algorithm the $(\mathbf{X}_m^1, \ldots, \mathbf{X}_m^{N_x})$ are approximatively distributed according to $\mathbf{X} \mid Y > \gamma_m$ but are not independent. To improve independence and tune the final size of the sample, an additional step is considered. There are thus two additional parameters, the size N of the sample and the number of steps A of the Metropolis–Hastings sampler in this additional step.

Sampling: duplicate the N_x particles $(\mathbf{X}_m^1, \ldots, \mathbf{X}_m^{N_x})$ into N particles, and apply A times to each particle the Metropolis–Hastings algorithm with exploration kernel T and target distribution $\mathbf{X} \mid \mathcal{M}(\mathbf{X}) > S$.

This adds $N \times A$ calls to the black box, and the output of this algorithm is a sample $(\tilde{\mathbf{X}}^1, \ldots, \tilde{\mathbf{X}}^N)$ which is approximately i.i.d. according to $\tilde{\mathbf{X}} = \mathbf{X} \mid \mathcal{M}(\mathbf{X}) > S$ together with the corresponding values $\tilde{Y}^k = \mathcal{M}(\tilde{\mathbf{X}}^k)$.

2.3. Simultaneous estimation of δ_i^f and $\bar{\eta}_i$

We now explain how to combine the method for estimating δ_i with the adaptive SMC sampler described above to have a simultaneous estimation of δ_i^f and $\bar{\eta}_i$.

Step 1 - Input realizations generation. Using the adaptive SMC procedure of Section 2.2, obtain $(\tilde{\mathbf{X}}^1, \ldots, \tilde{\mathbf{X}}^N)$ approximately i.i.d. from $f_{\tilde{\mathbf{X}}}$ and their corresponding value $\tilde{Y}^k = \mathcal{M}(\tilde{\mathbf{X}}^k)$ by $\mathcal{M}$.

Step 2 - Density estimation. Use the sample $((\tilde{X}_i^k, \tilde{Y}^k), k = 1, \ldots, N)$ to obtain estimates $\hat{f}_{\tilde{X}_i}$ and $\hat{c}_i$ of the density $f_{\tilde{X}_i}$ of $\tilde{X}_i$ and of the copula c_i of $(\tilde{X}_i, \tilde{Y})$, respectively, using for both the maximum entropy method with estimated fractional moments (see Appendix A);

Step 3 - Indices estimation. Use the estimates $\hat{f}_{\tilde{X}_i}$ and $\hat{c}_i$ to obtain estimates of $\bar{\eta}_i$ and δ_i^f as follows:

- for $\bar{\eta}_i$, estimate the one-dimensional integral $\|f_{X_i} - \hat{f}_{\tilde{X}_i}\|_{L^1(\mathbb{R})}$ either by direct numerical approximation, or, if for instance f_{X_i} is unknown but can be sampled from, by Monte Carlo method via

$$\hat{\bar{\eta}}_i = \frac{1}{N'} \sum_{k=1}^{N'} \left| \frac{\hat{f}_{\tilde{X}_i}(X_i^k)}{f_{X_i}(X_i^k)} - 1 \right|$$

where the X_i^k are i.i.d. with common distribution X_i;

- for δ_i^f, generate $((U_1^k, U_2^k), k = 1, \ldots, N')$ i.i.d. uniformly distributed on $[0, 1]^2$ and estimate δ_i^f by

$$\hat{\delta}_i^f = \frac{1}{2N'} \sum_{k=1}^{N'} |\hat{c}_i(U_1^k, U_2^k) - 1| . \tag{10}$$

It has to be pointed out that the proposed procedure can be applied to output models with correlated inputs and, as promised, provides simultaneous estimation of both δ_i^f and $\bar{\eta}_i$ from one common SMC procedure: indeed, after the first step no more call to the black box $\mathcal{M}$ is needed. In particular, the (random) number of calls to the black box is $N_x + mA_x N_x + AN$ as explained in Section 2.2.

3. Numerical applications

In this section, the proposed estimation scheme is applied on three output models. Firstly, we consider two analytical cases for which the unconditional and conditional output distributions are known so that theoretical values of the importance measures δ_i^f and $\bar{\eta}_i$ are available by using numerical integration. We then consider a single degree of freedom oscillator with $d = 6$ independent and lognormally distributed inputs.

Table 1: Estimates of δ_i^f and $\bar{\eta}_i$ of example 1. Set of parameters for the adaptive SMC algorithm: $N_x = 500$, $A_x = 3$, $\rho = 0.3935$, $a = 0.5$, $A = 5$ and $N = 3{,}000$.

Input	Theoretical value δ_i^f (rank)	Estimation $\hat{\delta}_i^f$		
		Mean (rank)	Sd	RD
X_1	**0.0781 (2)**	0.0930 (2)	0.0101	-0.1908
X_2	**0.7686 (1)**	0.7200 (1)	0.0077	0.0632

Input	Theoretical value $\bar{\eta}_i$ (rank)	Estimation $\hat{\bar{\eta}}_i$		
		Mean (rank)	Sd	RD
X_1	**0.9987 (1)**	0.9997 (1)	0.0095	-0.001
X_2	**0 (2)**	0.0315 (2)	0.0103	-

For each example, computation time and number of model calls are given to assess the efficiency of the proposed method. Results are obtained with a computer equipped with a 3.5 GHz Intel Xeon 4 CPU.

When the input $\mathbf{X} \sim N(\boldsymbol{\nu}, \boldsymbol{\Sigma})$ is normally distributed, mutation steps in the adaptive SMC algorithm are performed by using the natural exploration kernel so-called Crank Nicholson shaker and defined by

$$T(\mathbf{x}, \cdot) \sim \mathbf{L}\left(\sqrt{1-a} \times \mathbf{L}^{-1}(\mathbf{x} - \boldsymbol{\nu}) + \sqrt{a}Z\right) + \boldsymbol{\nu} \ ,$$

where $a \in (0,1)$ is a parameter of the kernel, $Z \sim N(0, I_d)$ and $\mathbf{L}$ is the lower triangular matrix in the Cholesky decomposition of $\boldsymbol{\Sigma}$, i.e., $\boldsymbol{\Sigma} = \mathbf{L}\mathbf{L}^T$.

Standard deviations (Sd) of estimators are computed by performing 100 runs of the proposed scheme in order to study its variability. When a theoretical value θ is available, the accuracy of an estimator $\hat{\theta}$ is measured by the mean of the relative difference (RD) $\frac{\theta - \hat{\theta}}{\theta}$.

3.1. Example 1: back to the toy model of the introduction

We go back to the toy model (1) of the introduction, i.e., $Y = X_1 + \mathbb{1}_{X_1 > S}|X_2|$ where $S = 3$, X_1 and X_2 are independent, $X_1 \sim N(0,1)$ and $X_2 \sim N(0,5)$. We compare in Table 1 theoretical values with estimates obtained with the proposed method. In average, runs last 317 seconds and make 34,640 calls to the black box. From the different relative differences, one can see that δ_i^f and $\bar{\eta}_i$ estimates are close to their respective reference values and present reasonable variability with regard to the budget allocated to the estimation.

3.2. Example 2: an analytical test case

Let us consider the following output model:

$$Y = X_1 + X_2^2$$

where X_1 and X_2 are i.i.d. standard Gaussian random variables and the failure event is $\{Y > 15\}$. This model is in the same vein as the previous toy model but slightly more realistic. Unconditional and conditional output distributions are known: $Y \mid X_1$ follows a χ^2-distribution shifted by X_1 and $Y \mid X_2$ is normally

Table 2: Estimates of δ_i^f and $\bar{\eta}_i$ of example 2. Set of parameters for the adaptive SMC algorithm: $N_x = 300$, $A_x = 3$, $\rho = 0.5507$, $a = 0.5$, $A = 3$ and $N = 5{,}000$.

Input	Theoretical value δ_i (rank)	Theoretical value δ_i^f (rank)	Estimation $\hat{\delta}_i^f$ Mean (rank)	Sd	RD
X_1	0.4930 (1)	0.001 (2)	0.0721 (2)	0.0266	-71.1
X_2	0.3049 (2)	0.4136 (1)	0.3998 (1)	0.0343	0.0334

Input	\	Theoretical value $\bar{\eta}_i$ (rank)	Estimation $\hat{\bar{\eta}}_i$ Mean (rank)	Sd	RD
X_1	\	0.2093 (2)	0.2066 (1)	0.0605	0.0129
X_2	\	0.9969 (1)	0.9723 (2)	0.0567	0.0247

distributed with unit variance and mean X_2. We thus have the following expressions for the densities:

$$f_Y(y) = \int_0^\infty \frac{e^{-\frac{(y-t)^2}{2} - \frac{t}{2}}}{2\pi\sqrt{t}} \mathrm{d}t$$

and

$$f_{Y|X_1}(y) = \frac{e^{-\frac{(y-X_1)}{2}}}{\sqrt{2\pi(y-X_1)}} \mathbb{1}_{y \geq X_1} \quad \text{and} \quad f_{Y|X_2}(y) = \frac{e^{-\frac{(y-X_2)^2}{2}}}{\sqrt{2\pi}}$$

for the conditional densities. Thus, theoretical values of sensitivity measures (δ_1, δ_2), (δ_1^f, δ_2^f) and $(\bar{\eta}_1, \bar{\eta}_2)$ are available via numerical integration. The failure probability can also be evaluated to 1.2387×10^{-4}. We gathered in Table 2 the estimates of δ_i^f and $\bar{\eta}_i$ obtained from the proposed method. In average, runs need 350 seconds to compute all the δ and η-indices and make 25,200 calls to the black box.

One can see that estimates $\{\hat{\delta}_i^f\}$ respect the good importance ranking, namely $X_2 > X_1$. However, the estimation of δ_1^f exhibits an important difference between average values and reference ones. This difference is due to the fact that the samples $\{\mathbf{X}^k\}$ obtained with the SMC procedure are not completely independent and distributed from $f_{\bar{\mathbf{X}}}$ since only $A = 3$ steps of the Metropolis–Hastings sampler are performed in the final sampling step. Indeed, increasing A from 3 to 30 leads to average values of $\hat{\delta}_1^f$ of 0.0206 with a standard deviation of 0.0091.

In this example, the indices $\{\delta_i^f\}$ enable to detect a drastic change in the importance ranking. Indeed, the contribution of the first input X_1 becomes negligible at the failure of the system whereas it is the most influential under nominal operation. The indices $\{\bar{\eta}_i\}$ lead to the same conclusion, namely that the influence of the input X_2 at the failure predominates with $\bar{\eta}_2$ close to 1.

3.3. Example 3: a single Degree of Freedom (SDOF) oscillator

In this subsection, a non linear SDOF oscillator [4] made of a mass m and two springs with free length r and respective stiffness c_1 and c_2 is considered. It is subjected to a rectangular load pulse with random duration t and amplitude F. The model output is defined as

$$Y = -3r + \left| \frac{2F}{c_1 + c_2} \sin\left(\sqrt{\frac{c_1 + c_2}{m}} \frac{t}{2} \right) \right|,$$

Table 3: Distribution parameters (the mean and the standard deviation of the associated normal distribution) of input variables of the SDOF oscillator.

Input	Mean	Sd
c_1	2	0.2
c_2	0.2	0.02
r	0.6	0.05
m	1	0.05
t	1	0.2
F	1	0.2

i.e., the difference between the maximum displacement response of the system and $3r$. The six input variables c_1, c_2, r, m, t and F are assumed to be independent and lognormally distributed with respective parameters given in Table 3. The failure of the system is achieved when the output Y exceeds the threshold 0 and the associated failure probability is approximately equal to 9×10^{-5}.

We gathered in Table 4 the estimates of δ_i^f and $\bar{\eta}_i$ obtained from the proposed method. The δ_i's are obtained with the method described in [10]. In average, runs last 960 seconds and make 51,725 calls to the black box. It appears that the global importance ranking $X_4 < X_2 < X_5 < X_1 < X_6 < X_3$ drastically differs from the importance ranking provided by the conditional sensitivity indices δ_i^f. Especially, the most influential input $X_3 = r$ becomes negligible conditionally on the failure event. Changes are more nuanced as far as target indices are concerned. Indeed, target sensitivity indices $\bar{\eta}_i$ give approximately the same ranking, except that X_1 and X_2 predominate.

As in the previous example, variability of obtained estimates is non negligible. Here, inputs are lognormally distributed and there is no natural exploration kernel like in the Gaussian case. We can find in [6] a discussion about implementation issues for the choice of the exploration kernel. In the current example, a candidate is drawn by adding a Gaussian noise with the same standard deviation as inputs. With this choice, it appears that we respect standard practice which is to tune the proposal distribution to get around 20%–25% acceptance rate [17]. Then, the only way to improve previous results is to increase the budget allocated to Metropolis–Hastings steps by increasing parameter A and decreasing the parameter ρ which regulates values of thresholds involved in the SMC procedure. From Table 5 which displays associated results, one can see that previous observed variability has been reduced. The new computation budget is about 262,500 calls to the model, which is quite substantial. Nevertheless, it remains substantially less expensive than the budget required by a classical Monte Carlo procedure. Furthermore, associated computational cost may be reduced by using a surrogate model. For instance, AK-SS method [11] combining Kriging and SMC simulation enables to assess small probabilities while replacing the expensive black box $\mathcal{M}$ by a less time-consuming function.

4. Generalization: target and conditional sensitivity analysis

Following the approach of [16], we explain here how to generalize our estimation scheme in two directions: (1) considering a more general notion of distance between

Table 4: Estimates of δ_i^f and $\bar{\eta}_i$ for the SDOF oscillator. Set of parameters for the adaptive SMC algorithm: $N_x = 500$, $A_x = 3$, $\rho = 0.4866$, $A = 10$ and $N = 3,000$.

Input	Estimation $\hat{\delta}_i$		Estimation $\hat{\delta}_i^f$		Estimation $\hat{\bar{\eta}}_i^f$	
	Mean (rank)	Sd	Mean (rank)	Sd	Mean (rank)	Sd
$X_1 = c_1$	0.0769 (3)	0.0066	0.0995 (1)	0.0210	0.8332 (2)	0.0653
$X_2 = c_2$	0.0231 (5)	0.0050	0.0322 (6)	0.0090	0.1352 (5)	0.0340
$X_3 = r$	0.4441 (1)	0.0063	0.0329 (5)	0.0117	0.6494 (3)	0.0690
$X_4 = m$	0.0219 (6)	0.0051	0.0343 (4)	0.0101	0.1306 (6)	0.0874
$X_5 = t$	0.0751 (4)	0.0075	0.0474 (3)	0.0150	0.3312 (4)	0.0710
$X_6 = F$	0.1554 (2)	0.0074	0.0871 (2)	0.0191	0.9078 (1)	0.0317

Table 5: Estimates of δ_i^f and $\bar{\eta}_i$ for the SDOF oscillator with a higher budget allocated to the adaptive SMC algorithm. Set of parameters for the adaptive SMC algorithm: $N_x = 500$, $A_x = 10$ (instead of 10), $\rho = 0.1813$, $A = 10$ and $N = 3,000$.

Input	Estimation $\hat{\delta}_i^f$		Estimation $\hat{\bar{\eta}}_i^f$	
	Mean (rank)	Sd	Mean (rank)	Sd
c_1	0.0674 (2)	0.0150	0.7949 (2)	0.0325
c_2	0.0275 (5)	0.0064	0.1131 (5)	0.0173
r	0.0346 (4)	0.0089	0.5651 (3)	0.0375
m	0.0267 (6)	0.0056	0.0459 (6)	0.0196
t	0.0366 (3)	0.0074	0.2812 (4)	0.0200
F	0.1147 (1)	0.0164	0.9205 (1)	0.0149

distributions; (2) assessing the impact of X_i on functions of Y.

4.1. More general distance

As explained in the introduction, Borgonovo's index is the total variation distance between (X_i, Y) and (X_i, Y') with Y' independent from X_i. In the absolutely continuous case, this corresponds to the L_1 distance between the joint density $f_{X_i, Y}$ and the product $f_{X_i} f_Y$ of its marginals, which reflects that this index is a measure of dependency between X_i and Y. Of course, many other dependency measures exist, for instance the Csiszár dependency measure.

Let $\phi : \mathbb{R}_+ \to \mathbb{R} \cup \{+\infty\}$ be a convex function with $\phi(1) = 0$: then the Csiszár divergence between two probability measures P and Q is given by

$$\mathrm{div}_\phi(P, Q) = \int \phi\left(\frac{\mathrm{d}P}{\mathrm{d}Q}\right) \mathrm{d}Q$$

where P is assumed to be absolutely continuous with Radon-Nikodym derivative $\frac{\mathrm{d}P}{\mathrm{d}Q}$ with respect to Q. For instance, for $\phi(x) = \frac{1}{2}|1 - x|$ this is the total variation distance, and for $\phi(x) = -\log(x)$ this is the Kullback–Leibler divergence. From this divergence, we can then define the Csiszár dependency measure (CDM_ϕ) between two random variables Z_1 and Z_2 as

$$\mathrm{CDM}_\phi(Z_1, Z_2) = \mathrm{div}_\phi\left((Z_1, Z_2), (Z_1, Z_2')\right)$$

with Z_2' equal in distribution to Z_2 and independent from Z_1 (identifying in the above a random variable and its distribution). Because the total variation distance

corresponds to the case $\phi(\cdot) = \frac{1}{2}|1 - \cdot|$, we recover Borgonovo's index with this choice, i.e., we have $\mathrm{CDM}_{\frac{1}{2}|1-\cdot|}(X_i, Y) = \delta_i$. Moreover, we note that this dependency measure can still be expressed in a straightforward manner from the copula c of (Z_1, Z_2) provided it exists, namely

$$\mathrm{CDM}_\phi(Z_1, Z_2) = \int \phi(c(u, v)) \mathrm{d}u \mathrm{d}v, \tag{11}$$

thereby generalizing the relation (9) at the heart of our estimation scheme for δ_i^f.

4.2. Impact on a function of Y

Consider any function $w : \mathcal{M}(\mathbb{R}^d) \to \mathbb{R}_+$ such that $w(Y)$ is integrable and let $\tilde{\mathbb{P}}^w$ be the probability measure which is absolutely continuous with respect to $\mathbb{P}$ with Radon-Nikodym derivative $w(Y)$. Thus, $\tilde{\mathbb{P}}^w$ is the unique probability measure defined on $(\Omega, \mathcal{F})$ such that

$$\tilde{\mathbb{P}}^w(A) = \frac{\mathbb{E}(w(Y)\mathbb{1}_A)}{\mathbb{E}(w(Y))}$$

for any measurable set $A \in \mathcal{F}$. Adopting the terminology of [16], we can generalize the two problems laid out in the introduction as follows:

Target sensitivity analysis: what is the influence of X_i on $w(Y)$ (rather than on Y)?

Conditional sensitivity analysis: what is the influence of X_i on Y under $\tilde{\mathbb{P}}^w$ (rather than under $\mathbb{P}$)?

What we have done before corresponds to the case $w(y) = \mathbb{1}_{y>S}$. Indeed, for this choice of w the measure $\tilde{\mathbb{P}}^w \circ \mathbf{X}^{-1}$ is the law of $\tilde{\mathbf{X}}$ as defined earlier:

$$\tilde{\mathbb{P}}^w(\mathbf{X} \in A) = \mathbb{P}(\mathbf{X} \in A \mid Y > S) = \mathbb{P}(\tilde{\mathbf{X}} \in A).$$

Thus, we generalize $\tilde{\mathbf{X}}$ to $\tilde{\mathbf{X}}^w = (X_1^w, \ldots, X_d^w)$ by defining it as a random variable with law $\tilde{\mathbb{P}}^w \circ \mathbf{X}^{-1}$, and we define $\tilde{Y}^w = \mathcal{M}(\tilde{\mathbf{X}}^w)$.

4.3. Generalization

In view of the Equations (4), (6) and (7) defining η_i, $\bar{\eta}_i$ and δ_i^f, respectively, the above extensions suggest the following more general version of these indices:

$$\eta_i^{\phi,w} = \mathrm{CDM}_\phi\left(X_i, w(Y)\right), \quad \bar{\eta}_i^{\phi,w} = \mathrm{div}_\phi\left(\tilde{X}_i^w, X_i\right) \quad \text{and} \quad \delta_i^{\phi,w} = \mathrm{CDM}_\phi\left(\tilde{X}_i^w, \tilde{Y}^w\right).$$

We will assume that (X_i, Y) is absolutely continuous with respect to Lebesgue measure with density $f_{X_i,Y}$, and that $(X_i, w(Y))$ is absolutely continuous with respect to the product measure $\mathrm{d}x\mu(\mathrm{d}a)$ with μ a measure on $\mathcal{M}(\mathbb{R}^d)$ with density $f_{X_i,w(Y)}$. If $w(Y)$ takes values in $\mathbb{R}$, one should typically think of μ as Lebesgue measure, but this more general formalism also makes it possible to encompass the important case where $w(Y)$ follows a discrete distribution: in this case, μ should simply be the counting measure and $(X_i, w(Y))$ is automatically absolutely continuous (with respect to $\mathrm{d}x\mu(\mathrm{d}a)$).

Under these assumptions, we have that:

- $\eta_i^{\phi,w} = \mathbb{E}\left[\operatorname{div}_\phi(w(Y), w(Y) \mid X_i)\right]$;

- $(\tilde{X}^w, \tilde{Y}^w)$ is absolutely continuous with respect to Lebesgue measure with density

$$f_{\tilde{X}_i^w, \tilde{Y}^w}(x, y) = \frac{w(y) f_{X_i, Y}(x, y)}{\mathbb{E}(w(Y))}.$$

For $w(y) = \mathbb{1}_{y>s}$ and $\phi(x) = |1 - x|$, we have the relation (5) between $\eta_i^{\phi,w}$ and $\bar{\eta}_i^{\phi,w}$ which reads

$$\eta_i^{\phi,w} = \mathbb{E}(w(Y)) \times \bar{\eta}_i^{\phi,w}.$$

However, this relation does not seem to hold outside this case, and so in general it is not clear whether $\eta_i^{\phi,w}$ and $\bar{\eta}_i^{\phi,w}$ can be easily related. Guided by the choice made in the case $w(y) = \mathbb{1}_{y>S}$, we consider in the sequel the index $\bar{\eta}_i^{\phi,w}$ even though it may seem at first glance less natural than $\eta_i^{\phi,w}$.

In order to generalize our estimation scheme, we first need a generalization of the adaptive SMC algorithm of Section 2.2. To sample from the tilted distribution $\tilde{\mathbb{P}}^w$, usual particle algorithms can be used such as the Metropolis–Hastings sampler with input target density $w(\mathcal{M}(\cdot)) f_{\mathbf{X}}(\cdot)/\mathbb{E}(w(Y))$. In the case $w(y) = \mathbb{1}_{y>S}$ it is hard to sample directly from $\tilde{\mathbb{P}}^w$ and intermediate distributions, say $\tilde{\mathbb{P}}^{w_p}$ with $w_p = \mathbb{1}_{y>\gamma_p}$, are needed. In this case and with a general w, one can for instance use the sequential Monte Carlo samplers proposed in [8].

Assume now that one is given a sample $(\tilde{\mathbf{X}}^{w,1}, \ldots, \tilde{\mathbf{X}}^{w,N})$ approximately i.i.d. with common distribution $\tilde{\mathbf{X}}^w$ and their values $\tilde{Y}^{w,k} = \mathcal{M}(\tilde{\mathbf{X}}^{w,k})$ by $\mathcal{M}$. As discussed above, in the case $w(y) = \mathbb{1}_{y>S}$ this is precisely the purpose of the adaptive SMC algorithm of Section 2.2. Then Step 2 of our estimation scheme remains unchanged and leads to:

- an estimate $\hat{f}_{\tilde{X}_i^w}$ of the density $f_{\tilde{X}_i^w}$ of $\tilde{X}_i^w$;

- an estimate $\hat{c}^w$ of the copula c^w of $(\tilde{X}_i^w, \tilde{Y}^w)$.

Using (11) we then have the following two estimations of $\bar{\eta}_i^{\phi,w}$ and $\delta_i^{\phi,w}$: for $\bar{\eta}_i^{\phi,w}$, an estimation $\hat{\bar{\eta}}_i^{\phi,w}$ can be obtained by numerically integrating the one-dimensional integral

$$\operatorname{div}_\phi(\tilde{X}_i^w, X_i) = \int \phi\left(\frac{f_{\tilde{X}_i^w}(x)}{f_{X_i}(x)}\right) f_{X_i}(x) \mathrm{d}x$$

or by a Monte Carlo approximation:

$$\hat{\bar{\eta}}_i^{\phi,w} = \frac{1}{N'} \sum_{k=1}^{N'} \phi\left(\frac{f_{\tilde{X}_i^w}(X_i^k)}{f_{X_i}(X_i^k)}\right)$$

with the X_i^k i.i.d. with common distribution f_{X_i}. For $\delta_i^{\phi,w}$, draw i.i.d. random variables (U_1^k, U_2^k) uniformly distributed on $[0,1]^2$ and consider

$$\hat{\delta}_i^{\phi,w} = \frac{1}{N'} \sum_{k=1}^{N'} \phi\left(\hat{c}^w(U_1^k, U_2^k)\right).$$

Appendix A. Maximum entropy principle

Appendix A.1. General principle

The maximum entropy principle was introduced by Jaynes [12], and the reader is for instance referred to [13] for more details. Let $\mathcal{P}_d(S)$ be the set of probability density functions on $S \subset \mathbb{R}^d$, and for $f \in \mathcal{P}_d(S)$ let $H(f)$ be its differential entropy, defined as

$$H(f) = -\int_S \log f(x) f(x) \mathrm{d}x \in [-\infty, +\infty].$$

In order to choose a density satisfying some constraints $\mathcal{C} \subset \mathcal{P}_d(S)$ (for instance, prescribed first and second moments), the maximum entropy principle asserts to choose among these densities the one with highest entropy, i.e.,

$$\begin{aligned} \underset{f \in \mathcal{P}_d(S)}{\arg\min} \quad & H(f) \\ \text{subject to} \quad & f \in \mathcal{C} \end{aligned} \tag{A.1}$$

When the constraints are linear equality constraints, i.e., are of the form $\mathcal{C} = \{f \in \mathcal{P}_d(S) : \int \varphi(x) f(x) \mathrm{d}x = \mu\}$ for some $\varphi : \mathbb{R}^d \to \mathbb{R}^d$ and $\mu \in \mathbb{R}^d$, then the above optimization problem is convex and a solution is of the form $f(x) = c e^{-\langle \Lambda^*, \varphi(x) \rangle} \mathbb{1}_S(x)$ where $\langle \cdot, \cdot \rangle$ denotes the inner product in $\mathbb{R}^d$, c is the normalization constant and Λ^* is a feasible solution of the dual optimization problem

$$\Lambda^* = \underset{\Lambda \in \mathbb{R}^n}{\arg\min} \ \left\{ \langle \Lambda, \mu \rangle + \log \left(\int_S e^{-\langle \Lambda, \varphi(x) \rangle} \mathrm{d}x \right) \right\}, \tag{A.2}$$

see for instance [3] for more details. The above objective function is strictly convex on the set of feasible points and so admits respectively a unique minimum which can been found using standard convex optimization techniques, for instance interior-point algorithms.

The above method can be used to estimate a given density f_0: if one knows some moments of the sought density f_0, then the idea is simply to put this information as constraints in (A.1).

Appendix A.2. Application to Step 2 of our estimation scheme

In our case, we want to apply the above maximum entropy principle in Step 2 of our estimation scheme (see Section 2.3) to estimate the density $f_{\tilde{X}_i}$ of $\tilde{X}_i$, and the density c_i of $(F_{\tilde{X}_i}(\tilde{X}_i), F_{\tilde{Y}}(\tilde{Y}))$. Ideally, we would like to consider solutions to (A.1) with linear equality constraints but the problem is that moments of the sought distributions are unknown. To circumvent this difficulty, we use the sample $((\tilde{X}_i^k, \tilde{Y}^k), k = 1, \dots, N)$ provided by the first step to estimate these moments. Also, for reasons discussed in [10] we consider fractional moments for the constraints.

More precisely, consider $\tilde{n}, n \in \mathbb{N}$ and real numbers $\alpha_1 < \cdots < \alpha_{\tilde{n}}$ and $\beta_1 < \cdots < \beta_n$, and let

$$\hat{M}_{r,s} := \frac{1}{N} \sum_{k=1}^N \left(\hat{F}_{\tilde{X}_i}(\tilde{X}_i^k) \right)^{\alpha_r} \left(\hat{F}_{\tilde{Y}}(\tilde{Y}^k) \right)^{\alpha_s}, \ r, s = 1, \dots, \tilde{n},$$

where $\tilde{F}_{\tilde{X}_i}$ and $\tilde{F}_{\tilde{Y}}$ are the empirical cumulative distribution functions of $\tilde{X}_i$ and $\tilde{Y}$, respectively, obtained from the sample $((\tilde{X}_i^k, \tilde{Y}^k), k = 1, \ldots, N)$, and

$$\hat{M}_t^i := \frac{1}{N} \sum_{k=1}^{N} (\tilde{X}_i^k)^{\beta_t}, t = 1, \ldots, n.$$

Then the estimates $\hat{f}_{\tilde{X}_i}$ and $\hat{c}_i$ of $f_{\tilde{X}_i}$ and c_i, respectively, are given by

$$\hat{f}_{\tilde{X}_i} = \underset{f \in \mathcal{P}_1(\mathrm{Supp}(\tilde{X}_i))}{\arg\min} H(f)$$

$$\text{subject to} \qquad \int_{\mathrm{Supp}(\tilde{X}_i)} x^{\beta_t} f(x) \mathrm{d}x = \hat{M}_t^i, \ t = 1, \ldots, n,$$

and

$$\hat{c}_i = \underset{f \in \mathcal{P}_2([0,1]^2)}{\arg\min} H(f)$$

$$\text{subject to} \quad \int_{[0,1]^2} x^{\alpha_r} y^{\alpha_s} f(x,y) \mathrm{d}x \mathrm{d}y = \hat{M}_{r,s}, \ r, s = 1, \ldots, \tilde{n}.$$

These solutions are obtained by the method described above. Note that the number of constraints is then n for estimating $f_{\tilde{X}_i}$ and $\tilde{n}^2$ for estimating c_i.

[1] Borgonovo, E. (2007). A new uncertainty importance measure. Reliability Engineering & System Safety, 92(6):771–784.

[2] Borgonovo, E. and Plischke, E. (2016). Sensitivity analysis: a review of recent advances. European Journal of Operational Research, 248(3):869–887.

[3] Boyd, S. and Vandenberghe, L. (2004). Convex optimization. Cambridge University Press, Cambridge.

[4] Bucher, C., Chen, Y., and Schuëller, G. (1989). Time variant reliability analysis utilizing response surface approach. In Reliability and Optimization of Structural Systems 88, pages 1–14. Springer.

[5] Cérou, F., Del Moral, P., Furon, T., and Guyader, A. (2012). Sequential Monte Carlo for rare event estimation. Statistics and computing, 22(3):795–808.

[6] Chib, S. and Greenberg, E. (1995). Understanding the Metropolis-Hastings algorithm. The American statistician, 49(4):327–335.

[7] Cui, L., Lü, Z., and Zhao, X. (2010). Moment-independent importance measure of basic random variable and its probability density evolution solution. Science China Technological Sciences, 53(4):1138–1145.

[8] Del Moral, P., Doucet, A., and Jasra, A. (2006). Sequential Monte Carlo samplers. J. R. Stat. Soc. Ser. B Stat. Methodol., 68(3):411–436.

[9] Derennes, P., Morio, J., and Simatos, F. (2017). A nonparametric importance sampling estimator for moment independent importance measures. Reliability Engineering & System Safety.

[10] Derennes, P., Morio, J., and Simatos, F. (2018). Estimation of the moment independent importance sampling measures using copula and maximum entropy framework. (Accepted to) Winter Simulation Conference.

[11] Huang, X., Chen, J., and Zhu, H. (2016). Assessing small failure probabilities by AK–SS: an active learning method combining kriging and subset simulation. Structural Safety, 59:86–95.

[12] Jaynes, E. T. (1957). Information theory and statistical mechanics. Physical review, 106(4):620.

[13] Kapur, J. N. and Kesavan, H. K. (1992). Entropy optimization principles with applications. Academic Press, Inc., Boston, MA.

[14] Liu, Q. and Homma, T. (2009). A new computational method of a moment-independent uncertainty importance measure. Reliability Engineering & System Safety, 94(7):1205–1211.

[15] Plischke, E., Borgonovo, E., and Smith, C. L. (2013). Global sensitivity measures from given data. European Journal of Operational Research, 226(3):536–550.

[16] Raguet, H. and Marrel, A. (2018). Target and conditional sensitivity analysis with emphasis on dependence measures. arXiv preprint arXiv:1801.10047.

[17] Sherlock, C. and Roberts, G. (2009). Optimal scaling of the random walk metropolis on elliptically symmetric unimodal targets. Bernoulli, pages 774–798.

[18] Wang, Y., Xiao, S., and Lu, Z. (2018). A new efficient simulation method based on Bayes' theorem and importance sampling Markov chain simulation to estimate the failure-probability-based global sensitivity measure. Aerospace Science and Technology.

[19] Wei, P., Lu, Z., and Song, J. (2014). Moment-Independent Sensitivity Analysis Using Copula. Risk Analysis, 34(2):210–222.

[20] Wei, P., Lu, Z., and Yuan, X. (2013). Monte Carlo simulation for moment-independent sensitivity analysis. Reliability Engineering & System Safety, 110:60–67.

[21] Zhang, L., Lu, Z., Cheng, L., and Fan, C. (2014). A new method for evaluating Borgonovo moment-independent importance measure with its application in an aircraft structure. Reliability Engineering & System Safety, 132:163–175.

[22] Zhang, X. and Pandey, M. D. (2013). Structural reliability analysis based on the concepts of entropy, fractional moment and dimensional reduction method. Structural Safety, 43:28–40.