



Bayesian subset simulation

Julien Bect, Ling Li, Emmanuel Vazquez

► To cite this version:

| Julien Bect, Ling Li, Emmanuel Vazquez. Bayesian subset simulation. 2016. hal-01253706v2

HAL Id: hal-01253706

<https://hal.science/hal-01253706v2>

Preprint submitted on 9 Dec 2016 (v2), last revised 23 Aug 2017 (v4)

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.

Bayesian subset simulation*

Julien Bect[†], Ling Li[‡], and Emmanuel Vazquez[†]

Abstract. We consider the problem of estimating a probability of failure α , defined as the volume of the excursion set of a function $f : \mathbb{X} \subseteq \mathbb{R}^d \rightarrow \mathbb{R}$ above a given threshold, under a given probability measure on \mathbb{X} . In this article, we combine the popular subset simulation algorithm (Au and Beck, Probab. Eng. Mech. 2001) and our sequential Bayesian approach for the estimation of a probability of failure (Bect, Ginsbourger, Li, Picheny and Vazquez, Stat. Comput. 2012). This makes it possible to estimate α when the number of evaluations of f is very limited and α is very small. The resulting algorithm is called Bayesian subset simulation (BSS). A key idea, as in the subset simulation algorithm, is to estimate the probabilities of a sequence of excursion sets of f above intermediate thresholds, using a sequential Monte Carlo (SMC) approach. A Gaussian process prior on f is used to define the sequence of densities targeted by the SMC algorithm, and drive the selection of evaluation points of f to estimate the intermediate probabilities. Adaptive procedures are proposed to determine the intermediate thresholds and the number of evaluations to be carried out at each stage of the algorithm. Numerical experiments illustrate that BSS achieves significant savings in the number of function evaluations with respect to other Monte Carlo approaches.

Key words. Probability of failure, Computer experiments, Sequential design, Gaussian process, Stepwise uncertainty reduction Sequential Monte Carlo

AMS subject classifications. 62L05, 62K99, 62P30

1. Introduction. Probabilistic reliability analysis has become over the last thirty years an essential part of the engineer's toolbox (see, e.g., [19, 44, 47]). One of the central problems in probabilistic reliability analysis is the computation of the probability of failure

$$(1) \quad \alpha = \int_{\mathbb{X}} \mathbb{1}_{f \leq 0} d\mathbf{P}_{\mathbb{X}}$$

of a system (or a component in a multicomponent system; see, e.g., [48], where $\mathbf{P}_{\mathbb{X}}$ is a probability measure over some measurable space $(\mathbb{X}, \mathcal{B})$ representing all possible sources of uncertainty acting on the system—both epistemic and aleatory—and $f : \mathbb{X} \rightarrow \mathbb{R}$ is the so-called *limit-state function*, such that f takes positive values when the system behaves reliably, and negative values when the system behaves unreliably, or fails. It is assumed in this article that \mathbb{X} is a subset of \mathbb{R}^d —in other words, we consider reliability problems where all uncertain factors can be described as a d -dimensional random vector. Numerous examples of applications that fall into this category can be found in the literature (see, for instance, [5, 18, 34, 39, 54, 55]).

*This research was partially funded by the French *Fond Unique Interministériel* (FUI 7) in the context of the CSDL (Complex Systems Design Lab) project. Parts of this work were previously published in the proceedings of the PSAM 11 & ESREL 12 conference [41] and in the PhD thesis of the second author [40].

[†]Laboratoire des Signaux et Systèmes, CentraleSupélec, CNRS, Univ. Paris-Sud, Université Paris-Saclay. 3 rue Joliot-Curie, 92192 Gif-sur-Yvette, France. Email: firstname.lastname@centralesupelec.fr

[‡]Email: ling.li.supelec@gmail.com. *Present address:* Schlumberger Gould Research Center, Cambridge, UK.

Two major difficulties usually preclude a brute-force Monte Carlo (MC) approach, that is, using the estimator

$$\hat{\alpha}^{\text{MC}} = \frac{1}{m} \sum_{i=1}^m \mathbf{1}_{f(X_i) \leq 0}, \quad X_i \stackrel{\text{i.i.d.}}{\sim} \mathbf{P}_{\mathbb{X}},$$

which requires m evaluations of f . First, the evaluation of f for a given $x \in \mathbb{X}$ often relies on one or several complex computer programs (e.g., partial differential equation solvers) that take a long time to run. Second, in many applications, the failure region $\Gamma = \{x \in \mathbb{X} \mid f(x) \leq 0\}$ is a *rare event* under the probability $\mathbf{P}_{\mathbb{X}}$; that is, the probability of failure $\alpha = \mathbf{P}_{\mathbb{X}}(\Gamma)$ is small. When α is small, the standard deviation of $\hat{\alpha}^{\text{MC}}$ is approximately $\sqrt{\alpha/m}$. To estimate α by MC with a standard deviation of 0.1α thus requires approximately $100/\alpha$ evaluations of f . As an example, with $\alpha = 10^{-3}$ and 10 minutes per evaluation, this means almost two years of computation time.

The first issue—designing efficient algorithms to estimate α in the case of an expensive-to-evaluate limit-state function—can be seen as a problem of *design and analysis of computer experiments* (see, e.g., [50]), bearing some similarities to the problem of global optimization (see [53] and references therein). Several sequential design strategies based on Gaussian process models have been proposed in the literature, and spectacular evaluation savings have been demonstrated on various examples with moderately small α (typically, 10^{-2} or 10^{-3}); see [6] for a review of fully sequential strategies and [3, 27] for examples of two-stage strategies. The closely related problem of quantile estimation has also been investigated along similar lines [1, 13, 46].

A key idea to address the second issue—i.e., to estimate a small probability of failure—is to consider a decreasing sequence of events $\Gamma_1 \supset \Gamma_2 \supset \dots \supset \Gamma_T = \Gamma$ such that the conditional probabilities $\mathbf{P}_{\mathbb{X}}(\Gamma_t \mid \Gamma_{t-1})$ are reasonably large, and therefore easier to estimate than α itself. Then, sequential Monte Carlo simulations [21] can be used to produce estimates \hat{p}_t of the conditional probabilities $\mathbf{P}_{\mathbb{X}}(\Gamma_t \mid \Gamma_{t-1})$, leading to a product-form estimate $\prod_{t=1}^T \hat{p}_t$ for α . This idea, called *subset simulation*, was first proposed in [2] for the simulation of rare events in structural reliability analysis¹, but actually goes back to the much older *importance splitting* (or *multilevel splitting*) technique used for the simulation of rare events in Markovian models (see, e.g., [36] and references therein). Subset simulation has since then become one of the most popular techniques for the computation of small probabilities of failure, and the theoretical properties of several (most of the times idealized) variants of the algorithm have recently been investigated by several authors (see, e.g., [11, 17]). However, because of the direct use of a Monte Carlo estimator for \hat{p}_t at each stage t , the subset simulation algorithm is not applicable when f is expensive to evaluate.

In this article we propose a new algorithm, called *Bayesian subset simulation* (BSS), which tackles both issues at once using ideas from the sequential design of computer experiments and from the literature on sequential Monte Carlo methods. Section 2 reviews the subset simulation algorithm from the point of view of sequential Monte Carlo (SMC)

¹A very similar algorithm had in fact been proposed earlier by [23], but for a quite different purpose (estimating the probability of a rare event under the bootstrap distribution).

techniques to prepare the ground for the introduction of our new algorithm. Section 3 describes the algorithm itself and Section 4 presents numerical results. Finally, Section 5 concludes the article with a discussion.

2. Subset simulation: a sequential Monte Carlo algorithm. This section recalls the main ideas of the classical subset simulation algorithm [2], which, although not originally presented as such, can be seen as a sequential Monte Carlo sampler [17, 21].

2.1. Idealized subset simulation (with fixed levels and IID sampling). We consider the problem of estimating the probability α of a rare event Γ of the form $\Gamma = \{x \in \mathbb{X} : f(x) > u\}$, where $u \in \mathbb{R}$ and $f : \mathbb{X} \rightarrow \mathbb{R}$, using pointwise evaluations of f . Note that the limit-state function (see Section 1) can be defined as $x \mapsto u - f(x)$ with our notations. Assuming, for the sake of simplicity, that $\mathbb{P}_{\mathbb{X}}$ has a probability density function $\pi_{\mathbb{X}}$ with respect to Lebesgue's measure, we have

$$\alpha = \int_{\mathbb{X}} \mathbf{1}_{f(x) > u} \pi_{\mathbb{X}}(x) dx.$$

The key idea of the subset simulation algorithm is to introduce an increasing (finite) sequence of thresholds $-\infty = u_0 < u_1 < u_2 \cdots < u_T = u$, which determine a corresponding decreasing sequence of subsets:

$$\mathbb{X} = \Gamma_0 \supset \Gamma_1 \supset \cdots \supset \Gamma_T = \Gamma, \quad \Gamma_t := \{x \in \mathbb{X} : f(x) > u_t\},$$

of the input space \mathbb{X} . Let $\alpha_t = \mathbb{P}_{\mathbb{X}}(\Gamma_t)$. The decreasing sequence $(\alpha_t)_{0 \leq t \leq T}$ obeys the recurrence formula

$$(2) \quad \alpha_{t+1} = \alpha_t \mathbb{P}_{\mathbb{X}}(\Gamma_{t+1} | \Gamma_t) = \alpha_t \int \mathbf{1}_{\Gamma_{t+1}}(x) q_t(x) dx,$$

where q_t stands for the truncated density

$$(3) \quad q_t(x) = \frac{\mathbf{1}_{\Gamma_t}(x) \pi_{\mathbb{X}}(x)}{\int \mathbf{1}_{\Gamma_t}(y) \pi_{\mathbb{X}}(y) dy}.$$

The small probability $\alpha = \alpha_T$ can thus be rewritten as a product of conditional probabilities, which are larger (and therefore easier to estimate) than α :

$$\alpha = \prod_{t=1}^T p_t, \quad p_t := \mathbb{P}_{\mathbb{X}}(\Gamma_t | \Gamma_{t-1}).$$

Assume that, for each $t \in \{0, 1, \dots, T-1\}$, a sample $(Y_t^j)_{1 \leq j \leq m}$ of independent and identically distributed (IID) random variables from the truncated density q_t is available. Then, each conditional probability p_t can be estimated by the corresponding Monte-Carlo estimator $\hat{p}_t = \frac{1}{m} \sum_{j=1}^m \mathbf{1}_{\Gamma_t}(Y_{t-1}^j)$, and α can be estimated by the product-form estimator $\hat{\alpha}^{\text{ss}} = \prod_{t=1}^T \hat{p}_t$. By choosing the thresholds u_t in such a way that the conditional probabilities p_t are high, α can be estimated using fewer evaluations of f than what would have been necessary using a simple Monte Carlo approach (see Section 2.4 for a quantitative example).

2.2. Sequential Monte-Carlo simulation techniques. Generating exact IID draws from the densities q_t is usually not possible, at least not efficiently, even if a method to generate IID samples from $q_0 = \pi_{\mathbb{X}}$ is available. Indeed, although the accept-reject algorithm (see, e.g., [49], Section 2.3) could be used in principle, it would be extremely inefficient when t is close to T , that is, when $\mathbb{P}_{\mathbb{X}}\{\Gamma_t\}$ becomes small. This is where sequential Monte-Carlo (SMC) simulation techniques are useful.

Given a sequence $(q_t)_{0 \leq t < T}$ of probability density functions over \mathbb{X} , SMC samplers sequentially generate, for each target density q_t , a weighted sample $\mathbb{Y}_t = ((w_t^j, Y_t^j))_{1 \leq j \leq m}$, where $w_t^j \geq 0$, $\sum_j w_t^j = 1$ and $Y_t^j \in \mathbb{X}$. The random vectors Y_t^j are usually called *particles* in the SMC literature, and the weighted sample \mathbb{Y}_t is said to *target* the distribution q_t . They are, in general, neither independent nor distributed according to q_t , but when the sample size m goes to infinity, their empirical distribution $\mu_t^{(m)} = \sum_{j=1}^m w_t^j \delta_{Y_t^j}$ converges to the target distribution—that is, to the distribution with probability density function q_t —in the sense that

$$\int_{\mathbb{X}} h(x) d\mu_t^{(m)}(x) = \sum_{j=1}^m w_t^j h(Y_t^j) \rightarrow \int_{\mathbb{X}} h(x) q_t(x) dx,$$

for a certain class of integrable functions h .

In practice, each weighted sample \mathbb{Y}_t is generated from the previous one, \mathbb{Y}_{t-1} , using transformations; SMC algorithms are thus expected to be efficient when each density q_t is, in some sense, close to its predecessor density q_{t-1} . The specific transformations that are used in the subset simulation algorithm are described next. The reader is referred to [21, 42] and references therein for a broader view of SMC sampling techniques, and to [25] for some theoretical results on the convergence (law of large numbers, central limit theorems) of SMC algorithms.

2.3. Reweight/resample/move. We now describe the reweight/resample/move scheme that is used in the subset simulation algorithm to turn a weighted sample \mathbb{Y}_{t-1} targeting $q_{t-1} \propto \mathbb{1}_{\Gamma_{t-1}} \pi_{\mathbb{X}}$ into a weighted sample \mathbb{Y}_t targeting $q_t \propto \mathbb{1}_{\Gamma_t} \pi_{\mathbb{X}}$. This scheme, used for instance in [16], can be seen as a special case of the more general SMC sampler of [21]².

Assume a weighted sample $\mathbb{Y}_{t-1} = ((w_{t-1}^j, Y_{t-1}^j))_{1 \leq j \leq m}$ targeting q_{t-1} has been obtained at stage $t - 1$. The *reweight* step produces a new weighted sample $\mathbb{Y}_{t,0} = ((w_{t,0}^j, Y_{t-1}^j))_{1 \leq j \leq m}$ that targets q_t , by changing only the weights in \mathbb{Y}_{t-1} :

$$w_{t,0}^j \propto \frac{q_t(Y_{t-1}^j)}{q_{t-1}(Y_{t-1}^j)} w_{t-1}^j.$$

The *resample* and *move* steps follow the reweighting step. These steps aim at avoiding the degeneracy of the sequence of weighted samples—i.e., the accumulation of most of the probability mass on a small number of particles with large weights.

²See in particular Section 3.1.1, Remark 1, and Section 3.3.2.3.

The simplest variant of resampling is the *multinomial resampling* scheme. It produces a new weighted sample $\mathbb{Y}_{t,1} = ((w_t^j, Y_{t,1}^j))_{1 \leq j \leq m}$, where the new particles $Y_{t,1}^j$ have equal weights $w_t^j = \frac{1}{m}$, and are independent and identically distributed according to the empirical distribution $\sum_{j=1}^m w_{t,0}^j \delta_{Y_{t-1}^j}$. In this work, we use the slightly more elaborate *residual resampling* scheme (see, e.g., [42]), which is known to outperform multinomial resampling ([24], Section 3.2). As in multinomial resampling, the residual resampling scheme produces a weighted sample with equal weights $w_t^j = \frac{1}{m}$.

The resampling step alone does not prevent degeneracy, since the resulting sample contains copies of the same particles. The move step restores some diversity by moving the particles according to a Markov transition kernel K_t that leaves q_t invariant:

$$\int q_t(x) K_t(x, dx') = q_t(x') dx';$$

for instance, a random-walk Metropolis-Hastings (MH) kernel (see, e.g., [49]) or the MMA (modified Metropolis algorithm) kernel used in Au and Beck's original paper [2, Section 3.1].

Remark 1. In the special case of the subset simulation algorithm, all weights are actually equal *before* the reweighting step and, considering the inclusion $\Gamma_t \subset \Gamma_{t-1}$, the reweighting formula takes the form

$$w_{t,0}^j \propto \mathbf{1}_{\Gamma_t}(Y_{t-1}^j).$$

In other words, the particles that are outside the new subset Γ_t are given a zero weight, and the other weights are simply normalized to sum to one. Note also that the resampling step discards particles outside of Γ_t (those with zero weight at the reweighting step).

Remark 2. Note that Au and Beck's original algorithm [2] uses a slightly different (but essentially similar) sampling scheme to populate each level: assuming that $L_t = m/m_t$ is an integer, where m_t denote the number of particles from stage $t-1$ that belong to Γ_t , they start m_t independent Markov chains of length L_t from each of the particles (called “seeds”).

Remark 3. In the general version of the reweight/resample/move procedure, the resampling step is carried out only when some degeneracy criterion—such as the expected sample size (ESS)—falls below a threshold (see, e.g., [21, 22]).

2.4. Practical subset simulation: adaptive thresholds. It is easy to prove that the subset simulation estimator $\hat{\alpha}^{\text{ss}} = \prod_{t=1}^T \hat{p}_t$ is unbiased. Moreover, according to Proposition 3 in [17], it is asymptotically normal in the large-sample-size limit:

$$(4) \quad \sqrt{m} \frac{\hat{\alpha}^{\text{ss}} - \alpha}{\alpha} \xrightarrow[m \rightarrow \infty]{\mathcal{D}} \mathcal{N}(0; \sigma^2)$$

where $\xrightarrow{\mathcal{D}}$ denotes convergence in distribution and

$$(5) \quad \sigma^2 \approx \sum_{t=1}^T \frac{1 - p_t}{p_t},$$

when the MCMC kernel has good mixing properties (see [17] article for the exact expression of σ^2). For a given number T of stages, the right-hand side of (5) is minimal when all conditional probabilities are equal; that is, when $p_t = \alpha^{1/T}$.

In practice however, the value of α is of course unknown, and it is not possible to choose the sequence of threshold beforehand in order to make all the conditional probabilities equal. Instead, a value p_0 is chosen—say, $p_0 = 10\%$ —and the thresholds are tuned in such a way that, at each stage t , $\hat{p}_t = p_0$. A summary of the resulting algorithm is provided in Table 1.

Equations (4) and (5) can be used to quantify the number of evaluations of f required to reach a given coefficient of variation with the subset simulation estimator $\hat{\alpha}^{\text{ss}}$. Indeed, in the case where all conditional probabilities are equal, we have

$$(6) \quad \text{var}(\hat{\alpha}^{\text{ss}}/\alpha) \approx \frac{T}{m} \frac{1-p_0}{p_0}.$$

with $T = \log(\alpha)/\log(p_0)$. For example, take $\alpha = 10^{-6}$. With the simple Monte Carlo estimator, the number of evaluations of f is equal to the sample size m : approximately $n = \delta^{-2} \alpha^{-1} = 10^8$ evaluations are required to achieve a coefficient of variation $\delta = \text{std}(\hat{\alpha}^{\text{MC}})/\alpha = 10\%$. In contrast, with $p_0 = 10\%$, the subset simulation algorithm will complete in $T = \log(\alpha)/\log(p_0) = 6$ stages, thus achieving a coefficient of variation $\delta = \text{std}(\hat{\alpha}^{\text{ss}})/\alpha = 10\%$ with $m = \delta^{-2} T(1-p_0)/p_0 = 5400$ particles. Assuming that the move step uses only one evaluation of f per particle, the corresponding number of evaluations would be $n = m + (T-1)(1-p_0)m = 29700 \ll 10^8$.

Remark 4. The value $p_0 = 0.1$ was used in the original paper of Au and Beck, on the ground that it had been “found to yield good efficiency” [see 2, Section 5]. Based on the approximate variance formula (6), Zuev and co-authors [56] argue that the variance is roughly proportional for a given total number of evaluations to $(1-p_0)/(p_0(\log(p_0))^2)$, and conclude³ that any $p_0 \in [0.1; 0.3]$ should yield quasi-optimal results, for any α .

3. Bayesian subset simulation.

3.1. Bayesian estimation and sequential design of experiment. Our objective is to build an estimator of α from the evaluations results of f at some points $X_1, X_2, \dots, X_N \in \mathbb{X}$, where N is the total budget of evaluations available for the estimation. In order to design an efficient estimation procedure, by which we mean both the design of experiments and the estimator itself, we adopt a Bayesian approach: from now on, the unknown function f is seen as a sample path of a random process ξ . In other words, the distribution of ξ is a *prior* about f . As in [6, 15, 52], the rationale for adopting a Bayesian viewpoint is to design a good estimation procedure in an average sense. This point of view has been largely

³Their analysis is based on the observation that the total number of evaluations is equal to mT —in other words, that m new samples must be produced at each stage. Some authors [e.g., 11] consider a variant where the particles that come from the previous stage are simply copied to the new set of particles, untouched by the Move step. In this case, a similar analysis suggests that 1) the optimal value of p_0 actually depends on α , and is somewhere between 0.63 (for $\alpha = 0.01$) and 1.0 (when $\alpha \rightarrow 0$); and 2) the value of δ^2 is only weakly dependent on p_0 , as long as p_0 is not too close to 0 (say, $p_0 \geq 0.1$).

Table 1: Subset simulation algorithm with adaptive thresholds

Prescribe $m_0 < m$ a fixed number of “succeeding particles”. Set $p_0 = \frac{m_0}{m}$.	
1. Initialization (stage 0)	
(a)	Generate an m -sample $Y_0^j \stackrel{\text{i.i.d.}}{\sim} \mathbf{P}_{\mathbf{X}}$, $1 \leq j \leq m$, and evaluate $f(Y_0^j)$ for all j .
(b)	Set $u_0 = -\infty$ and $t = 1$.
2. Repeat (stage t)	
(a) Threshold adaptation	
•	Compute the $(m - m_0)$ -th order statistic of $(f(Y_{t-1}^j))_{1 \leq j \leq m}$ and call it u_t^0 .
•	If $u_t^0 > u$, set $u_t = u$, $T = t$ and go to the estimation step.
•	Otherwise, set $u_t = u_t^0$ and $\Gamma_t = \{x \in \mathbf{X}; f(x) > u_t\}$.
(b) Sampling	
•	<i>Reweight</i> : set $m_t = \text{card}\{j \leq m: Y_{t-1}^j \in \Gamma_t\}$ and $w_{t,0}^j = \frac{1}{m_t} \mathbf{1}_{Y_{t-1}^j \in \Gamma_t}$.
•	<i>Resample</i> : generate a sample $(\tilde{Y}_t^j)_{1 \leq j \leq m}$ from the distribution $\sum_{j=1}^m w_{t,0}^j \delta_{Y_{t-1}^j}$.
•	<i>Move</i> : for each $j \leq m$, draw $Y_t^j \sim K(\tilde{Y}_t^j, \cdot)$. (NB: here, f is evaluated.)
(c)	Increment t .
3. Estimation	Let m_u be the number of particles such that $f(Y_{T-1}^j) > u$. Set
	$\hat{\alpha}^{\text{SS}} = \frac{m_u}{m} p_0^{T-1}.$

explored in the literature of computer experiments (see, e.g., [50]), and that of Bayesian optimization (see [32] and references therein).

For the sake of tractability, we assume as usual that, under the prior probability that we denote by \mathbf{P}_0 , ξ is a Gaussian process (possibly with a linearly parameterized mean, whose parameters are then endowed with a uniform improper prior; see [6] Section 2.3, for details).

Denote by \mathbf{E}_n (resp. \mathbf{P}_n) the conditional expectation (resp. conditional probability) with respect to $X_1, \xi(X_1), \dots, X_n, \xi(X_n)$, for any $n \leq N$ and assume, as in Section 2, that $\mathbf{P}_{\mathbf{X}}$ has a probability density function $\pi_{\mathbf{X}}$ with respect to Lebesgue’s measure. Then, a natural (mean-square optimal) Bayesian estimator of $\alpha = \mathbf{P}_{\mathbf{X}}(\Gamma)$ using n evaluations is the posterior mean

$$(7) \quad \mathbf{E}_n(\alpha) = \mathbf{E}_n \left(\int_{\mathbf{X}} \mathbf{1}_{\xi(x) > u} \pi_{\mathbf{X}}(x) dx \right) = \int_{\mathbf{X}} \tilde{g}_{n,u}(x) \pi_{\mathbf{X}}(x) dx,$$

where $\tilde{g}_{n,u}(x) := \mathbf{E}_n(\mathbf{1}_{\xi(x) > u}) = \mathbf{P}_n(\xi(x) > u)$ is the coverage function of the random set Γ (see, e.g., [14]). Note that, since ξ is Gaussian, $\tilde{g}_{n,u}(x)$ can be readily computed for any x using the kriging equations (see, e.g., [6], Section 2.4).

Observe that $\tilde{g}_{n,u} \approx \mathbf{1}_{\Gamma}$ when the available evaluation results are informative enough to classify most input points correctly (with high probability) with respect to u . This suggests that the computation of the right-hand side of (7) should not be carried out using a brute force Monte Carlo approximation, and would benefit from an SMC approach similar to the subset simulation algorithm described in Section 2. Moreover, combining an SMC approach

with the Bayesian viewpoint is also beneficial for the problem of choosing (sequentially) the sampling points X_1, \dots, X_N . In our work, we focus on a *stepwise uncertainty reduction* (SUR) strategy [6, 52]. Consider the function $L : \widehat{\Gamma} \mapsto \mathbb{P}_{\mathbb{X}}(\Gamma \triangle \widehat{\Gamma})$, which quantifies the loss incurred by choosing an estimator $\widehat{\Gamma}$ instead of the excursion set Γ , where \triangle stands for the symmetric difference operator. Here, at each iteration n , we choose the estimator $\widehat{\Gamma}_{n,u} = \{x \in \mathbb{X} \mid \tilde{g}_{n,u}(x) > 1/2\}$. A SUR strategy, for the loss L and the estimators $\widehat{\Gamma}_{n,u}$, consists in choosing a point X_{n+1} at step n in such a way to minimize the expected loss at step $n+1$:

$$(8) \quad X_{n+1} = \underset{x_{n+1} \in \mathbb{X}}{\operatorname{argmin}} J_n(x_{n+1}),$$

where

$$(9) \quad J_n(x_{n+1}) := \mathbb{E}_n(\mathbb{P}_{\mathbb{X}}(\Gamma \triangle \widehat{\Gamma}_{n+1,u}) \mid X_{n+1} = x_{n+1}).$$

For computational purposes, J_n can be rewritten as an integral over \mathbb{X} of the expected probability of misclassification $\tau_{n+1,u}$ (see [6] for more details):

$$(10) \quad J_n(x_{n+1}) = \int_{\mathbb{X}} \mathbb{E}_n(\tau_{n+1,u}(x) \mid X_{n+1} = x_{n+1}) \pi_{\mathbb{X}}(x) dx.$$

where

$$(11) \quad \tau_{n,u}(x) := \mathbb{P}_n(x \in \Gamma \triangle \widehat{\Gamma}_{n,u}) = \min(\tilde{g}_{n,u}(x), 1 - \tilde{g}_{n,u}(x)).$$

For moderately small values of α , it is possible to use a sample from $\mathbb{P}_{\mathbb{X}}$ both for the approximation of the integral in the right-hand side of (10) and for an approximate minimization of J_n (by exhaustive search in the set of sample points). However, this simple Monte Carlo approach would require a very large sample size to be applicable for very small values of α ; a subset-simulation-like SMC approach will now be proposed as a replacement.

3.2. A sequential Monte Carlo approach. Assume that α is small and consider a decreasing sequence of subsets $\mathbb{X} = \Gamma_0 \supset \Gamma_1 \supset \dots \supset \Gamma_T = \Gamma$, where $\Gamma_t = \{x \in \mathbb{X} : f(x) > u_t\}$, as in Section 2. For each $t \leq T$, denote by $\widehat{\alpha}_t^{\mathbb{B}}$ the Bayesian estimator of $\alpha_t = \mathbb{P}_{\mathbb{X}}(\Gamma_t)$ obtained from n_t observations of ξ at points X_1, \dots, X_{n_t} :

$$(12) \quad \widehat{\alpha}_t^{\mathbb{B}} := \mathbb{E}_{n_t}(\alpha_t) = \int_{\mathbb{X}} g_t d\mathbb{P}_{\mathbb{X}},$$

where $g_t(x) := \tilde{g}_{n_t,u_t}(x) = \mathbb{P}_{n_t}(\xi(x) > u_t)$.

The main idea of our new algorithm is to use an SMC approach to construct a sequence of approximations $\widehat{\alpha}_t^{\text{BSS}}$ of the Bayesian estimators $\widehat{\alpha}_t^{\mathbb{B}}$, $1 \leq t \leq T$ (as explained earlier, the particles of these SMC approximations will also provide suitable candidate points for the optimization of a sequential design criterion). To this end, consider the sequence of probability density functions q_t defined by

$$(13) \quad q_t(x) := \frac{\pi_{\mathbb{X}}(x) g_t(x)}{\int \pi_{\mathbb{X}}(y) g_t(y) dy} = \frac{1}{\widehat{\alpha}_t^{\mathbb{B}}} \pi_{\mathbb{X}}(x) g_t(x).$$

We can write a recurrence equation for the sequence of Bayesian estimators $\hat{\alpha}_t^B$, similar to that used for the probabilities α_t in (2):

$$(14) \quad \hat{\alpha}_{t+1}^B = \int g_{t+1}(x) \pi_{\mathbb{X}}(x) dx = \hat{\alpha}_t^B \int \frac{g_{t+1}(x)}{g_t(x)} q_t(x) dx.$$

This suggests to construct recursively a sequence of estimators $(\hat{\alpha}_t^{\text{BSS}})$ using the following relation:

$$(15) \quad \hat{\alpha}_{t+1}^{\text{BSS}} = \hat{\alpha}_t^{\text{BSS}} \sum_{j=1}^m w_t^j \frac{g_{t+1}(Y_t^j)}{g_t(Y_t^j)}, \quad 0 \leq t < T,$$

where $(w_t^j, Y_t^j)_{1 \leq j \leq m}$ is a weighted sample of size m targeting q_t (as in Section 2.2) and $\hat{\alpha}_0^{\text{BSS}} = 1$. The final estimator can be written as:

$$(16) \quad \hat{\alpha}_T^{\text{BSS}} = \prod_{t=0}^{T-1} \frac{\hat{\alpha}_{t+1}^{\text{BSS}}}{\hat{\alpha}_t^{\text{BSS}}} = \prod_{t=0}^{T-1} \sum_{j=1}^m w_t^j \frac{g_{t+1}(Y_t^j)}{g_t(Y_t^j)}.$$

Remark 5. The connection between the proposed algorithm and the original subset simulation algorithm is clear from the similarity between the recurrence relations (2) and (14), and from the use of SMC simulation in both algorithms to construct recursively a product-type estimator of the probability of failure (see also this type of estimator is mentioned in a very general SMC framework).

Our choice for the sequence of densities q_1, \dots, q_T also relates to the original subset simulation algorithm. Indeed, note that $q_t(x) \propto \mathbb{E}_{n_t}(\mathbb{1}_{\xi > u_t} \pi_{\mathbb{X}})$, and recall from Equation (3) that $q_t \propto \mathbb{1}_{\xi > u_t} \pi_{\mathbb{X}}$ is the target distribution used in the subset simulation algorithm at stage t . This choice of instrumental density is also used by [27, 28] in the context of a two-stage adaptive importance sampling algorithm. This is indeed a quite natural choice, since $\tilde{q}_t \propto \mathbb{1}_{\xi > u_t} \pi_{\mathbb{X}}$ is the optimal instrumental density for the estimation of α_t by importance sampling [see, e.g., 49, Theorem 3.12].

3.3. The Bayesian subset simulation (BSS) algorithm. The algorithm consists of a sequence of stages (or iterations). For the sake of clarity, assume first that the sequence of thresholds (u_t) is given. Then, each stage $t \in \mathbb{N}$ of the algorithm is associated to a threshold u_t and the corresponding excursion set $\Gamma_t = \{f > u_t\}$.

The initialization stage ($t = 0$) starts with the construction of a *space filling* set of points $\{X_1, \dots, X_{n_0}\}$ in \mathbb{X} ⁴, and an initial Monte Carlo sample $\mathbb{Y}_0 = \{Y_0^1, \dots, Y_0^m\}$, consisting of a set of independent random variables drawn from the density $q_0 = \pi_{\mathbb{X}}$.

After initialization, each subsequent stage $t \geq 1$ of BSS involves two phases: an *estimation phase*, where the estimation of Γ_t is carried out, and a *sampling phase*, where a

⁴See Section 4.2.1 for more information on the specific technique used in this article. Note that it is of course possible, albeit not required to use the BSS algorithm, to perform first a change of variables in order to work, e.g., in the standard Gaussian space. Whether this will improve the performance of the BSS algorithm is very difficult to say in general, and will depend on the example at hand.

sample \mathbb{Y}_t targeting the density q_t associated to u_t is produced from the previous sample \mathbb{Y}_{t-1} using the reweight/resample/move SMC scheme described in Section 2.3.

In more details, the estimation phase consists in making $N_t \geq 0$ new evaluations of f to refine the estimation of Γ_t . The number of evaluations is meant to be much smaller than the size m of the Monte Carlo sample—which would be the number of evaluations in the classical subset simulation algorithm. The total number of evaluations at the end of the estimation phase at stage t is denoted by $n_t = n_{t-1} + N_t$. The total number of evaluations used by BSS is thus $n_T = n_0 + \sum_{t=1}^T N_t$. New evaluation points $X_{n_{t-1}+1}, X_{n_{t-1}+2}, \dots, X_{n_t}$ are determined using a SUR sampling strategy⁵ targeting the threshold u_t , as in Section 3.1 (see Supplementary Material SM1 for details about the numerical procedure).

In practice, the sequence of thresholds is not fixed beforehand and adaptive techniques are used to choose the thresholds (see Section 3.4) and the number of points per stage (see Section 3.5).

The BSS algorithm is presented in pseudo-code form in Table 2.

Remark 6. Algorithms involving Gaussian-process-based adaptive sampling and subset simulation have been proposed by Dubourg and co-authors [26, 29] and by Huang et al. [37]. Dubourg’s work addresses a different problem (namely, reliability-based design optimization). Huang et al.’s paper, published very recently, addresses the estimation of small probabilities of failure. We emphasize that, unlike BSS, none of these algorithms involves a direct interaction between the selection of evaluation points (adaptive sampling) and subset simulation—which is simply applied, in its original form, to the posterior mean of the Gaussian process (also known as kriging predictor).

3.4. Adaptive choice of the thresholds u_t . As discussed in Section 2.4, it can be proved that, for an idealized version of the subset simulation algorithm with fixed thresholds $u_0 < u_1 < \dots < u_T = u$, it is optimal to choose the thresholds to make all conditional probabilities $P_{\mathbb{X}}(\Gamma_{t+1} | \Gamma_t)$ equal. This leads to the idea of choosing the thresholds adaptively in such a way that, in the product estimate

$$\hat{\alpha}_T^{\text{ss}} = \prod_{t=1}^T \frac{1}{m} \sum_{i=1}^m \mathbb{1}_{\Gamma_t}(Y_{t-1}^i),$$

each term but the last is equal to some prescribed constant p_0 . In other words, u_t is chosen as the $(1 - p_0)$ -quantile of \mathbb{Y}_{t-1} . This idea was first suggested by ([2], Section 5.2), on the heuristic ground that the algorithm should perform well when the conditional probabilities are neither too small (otherwise they are hard to estimate) nor too large (otherwise a large number of stages is required).

Consider now an idealized BSS algorithm, where a) the initial design of experiment is independent of \mathbb{Y}_0 , b) the SUR criterion is computed exactly, or using a discretization scheme that does not use the \mathbb{Y}_t ’s; c) the minimization of the SUR criterion is carried out independently of the \mathbb{Y}_t ’s and d) the particles Y_t^j are independent and identically dis-

⁵Other sampling strategies (also known as “sequential design”, or “active learning” methods) could be used as well. See [6] for a review and comparison of sampling criterions.

Table 2: Bayesian subset simulation algorithm

1. Initialization (stage 0)
(a) Evaluate f on a set of points $\{X_1, \dots, X_{n_0}\}$, called the <i>initial design</i> (see Section 4.2.1 for details)
(b) Generate an IID sample $\mathbb{Y}_t = \{Y_0^1, \dots, Y_0^m\}$ from $\mathbf{P}_{\mathbb{X}}$.
(c) Choose a prior \mathbf{P}_0 (see Sections 3.1 and 4.2.1 for details).
(d) Set $u_0 = -\infty$, $g_0 = \tilde{g}_{0,-\infty} = \mathbb{1}_{\mathbb{X}}$, $n = n_0$ and $t = 1$.
2. Repeat (stage t)
(a) Estimation
• Set $k = 0$ and repeat
– Select a threshold $\tilde{u}_{t,k}$ by solving Equation (20) for u_t (with $n_t = n$)
– Stop if the condition (21) is met, with $n_t = n$ and $u_t = \tilde{u}_{t,k}$.
– Select X_{n+1} using the SUR strategy (8)–(11) with respect to $\tilde{u}_{t,k}$.
– Evaluate f at X_{n+1} . Increment n and k .
• Set $N_t = k$, $n_t = n$, $u_t = \tilde{u}_{t,k}$ and $g_t = \tilde{g}_{n_t, u_t} = \mathbf{P}_{n_t}(\xi(\cdot) > u_t)$.
(b) Sampling
• <i>Reweight</i> : calculate weights $w_{t,0}^j \propto g_t(Y_{t-1}^j)/g_{t-1}(Y_{t-1}^j)$, $1 \leq j \leq m$.
• <i>Resample</i> : generate a sample $(\tilde{Y}_t^j)_{1 \leq j \leq m}$ from the distribution $\sum_{j=1}^m w_{t,0}^j \delta_{Y_{t-1}^j}$.
• <i>Move</i> : for each $j \leq m$, draw $Y_t^j \sim K(\tilde{Y}_t^j, \cdot)$.
(c) Increment t .
3. Estimation – The final probability of failure is estimated by
$\hat{\alpha}_T^{\text{BSS}} = \prod_{t=0}^{T-1} \left(\frac{1}{m} \sum_{j=1}^m \frac{g_{t+1}(Y_t^j)}{g_t(Y_t^j)} \right).$

tributed according to q_t . Assumptions a)–c) ensure that the sequence of densities $(q_t)_{1 \leq t \leq T}$ is deterministic given ξ . Then (see Appendix A),

$$(17) \quad \text{var} \left(\frac{\hat{\alpha}_T^{\text{BSS}}}{\hat{\alpha}_T^{\text{B}}} \middle| \xi \right) = \frac{1}{m} \sum_{t=1}^T \kappa_t + O \left(\frac{1}{m^2} \right),$$

where

$$(18) \quad \kappa_t := \frac{\int_{\mathbb{X}} g_t^2 / g_{t-1} \pi_{\mathbb{X}}}{(\hat{\alpha}_t^{\text{B}})^2 / \hat{\alpha}_{t-1}^{\text{B}}} - 1.$$

Minimizing the leading term $\frac{1}{m} \sum_{t=1}^T \kappa_t$ in (17) by an appropriate choice of thresholds is not as straightforward as in the case of the subset simulation algorithm. Assuming that $g_{t-1} \approx 1$ wherever g_t is not negligible (which is a reasonable assumption, since $g_t(x) = \mathbf{P}_{n_t}(\xi(x) > u_t)$ and $u_t > u_{t-1}$), we get

$$\int_{\mathbb{X}} g_t^2 / g_{t-1} \pi_{\mathbb{X}} \approx \int_{\mathbb{X}} g_t^2 \pi_{\mathbb{X}} \leq \int_{\mathbb{X}} g_t \pi_{\mathbb{X}} = \hat{\alpha}_t^{\text{B}},$$

and therefore the variance (17) is approximately upper-bounded by

$$(19) \quad \frac{1}{m} \sum_{t=1}^T (1 - \hat{p}_t^B) / \hat{p}_t^B,$$

where $\hat{p}_t^B := \hat{\alpha}_t^B / \hat{\alpha}_{t-1}^B$. Minimizing the approximate upper-bound (19) under the constraint

$$\prod_{t=1}^T \hat{p}_t^B = \hat{\alpha}_T^B$$

leads to choosing the thresholds u_t in such a way that \hat{p}_t^B is the same for all stages t —that is $\hat{p}_t^B = (\hat{\alpha}_T^B)^{1/T}$. As a consequence, we propose to choose the thresholds adaptively using the condition that, at each stage (but the last), the natural estimator $\hat{\alpha}_t^{\text{BSS}} / \hat{\alpha}_{t-1}^{\text{BSS}}$ of \hat{p}_t^B is equal to some prescribed probability p_0 . Owing to (15), this amounts to choosing u_t in such a way that

$$(20) \quad \frac{1}{m} \sum_{i=1}^m \frac{g_t(Y_{t-1}^i)}{g_{t-1}(Y_{t-1}^i)} = p_0.$$

should be satisfied.

Equation (20) is easy to solve, since the left-hand side is a strictly decreasing and continuous function of u_t (to be precise, continuity holds under the assumption that the posterior variance of ξ does not vanish on one of the particles). In practice, we solve (20) each time a new evaluation is made, which yields a sequence of intermediate thresholds (denoted by $\tilde{u}_{t,0}, \tilde{u}_{t,1}, \dots$ in Table 2) at each stage $t \geq 1$. The actual value of u_t at stage t is only known after the last evaluation of stage t .

Remark 7. Alternatively, the effective sample size (ESS) could be used to select the thresholds, as proposed by [22]. This idea will not be pursued in this paper. The threshold selected by the ESS-based approach will be close to the threshold selected by Equation (20) when the all the ratios $g_t(Y_{t-1}^i) / g_{t-1}(Y_{t-1}^i)$, or most of them, are either close to zero or close to one.

3.5. Adaptive choice of the number N_t of evaluation at each stage. In this section, we propose a technique to choose adaptively the number N_t of evaluations of f that must be done at each stage of the algorithm.

Assume that $t \geq 1$ is the current stage number; at the beginning of the stage, n_{t-1} evaluations are available from previous stages. After several additional evaluations, the number of available observations of f is $n \geq n_{t-1}$. We propose to stop adding new evaluations when the expected error of estimation of the set Γ_t , measured by $\mathbb{E}_n \left(\mathbb{P}_{\mathbb{X}} \left(\Gamma_t \triangle \hat{\Gamma}_{n,u_t} \right) \right)$, becomes smaller than some prescribed fraction η_t of its expected volume $\mathbb{E}_n (\mathbb{P}_{\mathbb{X}} (\Gamma_t))$ un-

der $P_{\mathbb{X}}$. Writing these two quantities as

$$\begin{aligned} E_n(P_{\mathbb{X}}(\Gamma_t)) &= \int_{\mathbb{X}} \tilde{g}_{n,u_t}(x) \pi_{\mathbb{X}}(x) dx = \hat{\alpha}_{t-1}^B \int_{\mathbb{X}} \frac{\tilde{g}_{n,u_t}(x)}{g_{t-1}(x)} q_{t-1}(x) dx, \\ E_n(P_{\mathbb{X}}(\Gamma \triangle \hat{\Gamma}_{n,u_t})) &= \int_{\mathbb{X}} \tau_{n,u_t}(x) \pi_{\mathbb{X}}(x) dx = \hat{\alpha}_{t-1}^B \int_{\mathbb{X}} \frac{\tau_{n,u_t}(x)}{g_{t-1}(x)} q_{t-1}(x) dx, \end{aligned}$$

where \tilde{g}_{n,u_t} and τ_{n,u_t} have been defined in Section 3.1, and estimating the integrals on the right-hand side using the SMC sample \mathbb{Y}_{t-1} , we end up with the stopping condition

$$\frac{1}{m} \sum_{i=1}^m \frac{\tau_{n,u_t}(Y_{t-1}^i)}{g_{t-1}(Y_{t-1}^i)} \leq \eta_t \cdot \frac{1}{m} \sum_{i=1}^m \frac{\tilde{g}_{n,u_t}(Y_{t-1}^i)}{g_{t-1}(Y_{t-1}^i)}.$$

which, if u_t is re-adjusted after each evaluation using Equation (20), boils down to

$$(21) \quad \sum_{i=1}^m \frac{\tau_{n,u_t}(Y_{t-1}^i)}{g_{t-1}(Y_{t-1}^i)} \leq \eta_t m p_0.$$

Remark 8. In the case where several evaluations of the function can be carried out in parallel, it is possible to select evaluation points in batches during the sequential design phase of the algorithm. A batch-sequential version of the SUR strategy (8)–(11) has been proposed by [15].

Remark 9. The stopping criterion (21) is slightly different from the one proposed earlier by [41]: $\sum_{i=1}^m \tau_{n,u_t}(Y_{t-1}^i) \leq \eta' m$. If we set $\eta' = \eta_t p_0$ and assume (quite reasonably) that $g_{t-1}(Y_{t-1}^i) \approx 1$ for the particles where $\tau_{n,u_t}(Y_{t-1}^i)$ is not negligible, then it becomes clear that the two criteria are essentially equivalent. As a consequence, the left-hand side of (21) can also be interpreted, approximately, as the expected number of misclassified particles (where the expectation is taken with respect to ξ , conditionally on the particles).

4. Numerical experiments. In this section, we illustrate the proposed algorithm on three classical examples from the structural reliability literature and compare our results with those from the classical subset simulation algorithm and the ²SMART algorithm [10, 20]. These examples are not actually expensive to evaluate, which makes it possible to analyse the performance of the algorithms through extensive Monte Carlo simulations, but the results are nonetheless relevant to case of expensive-to-evaluate simulators since performance is measured in terms of number of function evaluations (see Section 4.3.2 for a discussion).

The computer programs used to conduct these numerical experiments are freely available from <https://github.com/jbect/bss> under the LGPL licence [33]. They are written in the Matlab/Octave language and use the STK toolbox [7] for Gaussian process modeling. For convenience, a software package containing both the code for the BSS algorithm itself and the STK toolbox is provided as Supplementary Material.

4.1. Test cases. For each of the following test cases, the reference value for the probability α has been obtained from one hundred independent runs of the subset simulation algorithm with sample size $m = 10^7$ (see Table 3).

Table 3: Summary of test cases

Example	Name	d	α_{ref}
4.1.1	Four-branch series system	2	$5.596 \cdot 10^{-9}$
4.1.2	Deviation of a cantilever beam	2	$3.937 \cdot 10^{-6}$
4.1.3	Response of a nonlinear oscillator	6	$1.514 \cdot 10^{-8}$

4.1.1. Four-branch series system. Our first example is a variation on a classical structural reliability test case (see, e.g., [30], Example 1, with $k = 6$), where the threshold u is modified to make α smaller. The objective is to estimate the probability $\alpha = \mathbb{P}_{\mathbf{X}}(f(\mathbf{X}) < u)$, where

$$(22) \quad f(x_1, x_2) = \min \left\{ \begin{array}{l} 3 + 0.1(x_1 - x_2)^2 - (x_1 + x_2)/\sqrt{2}, \\ 3 + 0.1(x_1 - x_2)^2 + (x_1 + x_2)/\sqrt{2}, \\ (x_1 - x_2) + 6/\sqrt{2}, \\ (x_2 - x_1) + 6/\sqrt{2} \end{array} \right\}$$

and $X_1, X_2 \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. Taking $u = -4$, the probability of failure is approximately $5.596 \cdot 10^{-9}$, with a coefficient of variation of about 0.04%. Figure 1 (left panel) shows the failure domain and a sample from the input distribution $\mathbb{P}_{\mathbf{X}}$.

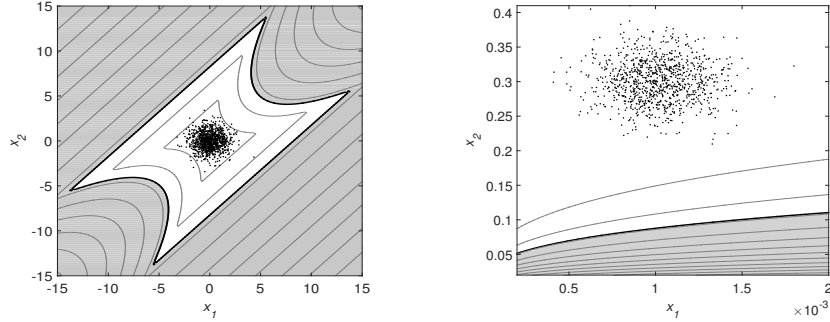


Figure 1: Contour plots of f in Example 4.1.1 (left) and Example 4.1.2 (right), along with a sample of size $m = 10^3$ from $\mathbb{P}_{\mathbf{X}}$ (dots). A failure happens when x is in the gray area.

4.1.2. Deviation of a cantilever beam. Consider a cantilever beam, with a rectangular cross-section, subjected to a uniform load. The deflection of the tip of the beam can be written as

$$(23) \quad f(x_1, x_2) = \frac{3L^4}{2E} \frac{x_1}{x_2^3},$$

where x_1 is the load per unit area, x_2 the thickness of the beam, $L = 6$ m and $E = 2.6 \cdot 10^4$ MPa. The input variable X_1 and X_2 are assumed independent, with $X_1 \sim \mathcal{N}(\mu_1, \sigma_1^2)$, $\mu_1 = 10^{-3}$ MPa, $\sigma_1 = 0.2\mu_1$, and $X_2 \sim \mathcal{N}(\mu_2, \sigma_2^2)$, $\mu_2 = 0.3$ m, $\sigma_2 = 0.1\mu_2$. A failure occurs

Table 4: Example 4.1.3: Means and standard deviations of the input variables.

Variable	x_1	x_2	x_3	x_4	x_5	x_6
μ_i	1	1	0.1	0.5	0.45	1
σ_i	0.05	0.1	0.01	0.05	0.075	0.2

when f is larger than $u = L/325$. The probability of failure is approximately $3.937 \cdot 10^{-6}$, with a coefficient of variation of about 0.03%. Note that the distribution of X_2 has been modified, with respect to the usual formulation (see, e.g., [35]), to make α smaller. Figure 1 (right panel) shows a contour plot of f , along with a sample of the input distribution.

4.1.3. Response of a nonlinear oscillator. In this example (see, e.g., [31]), the input variable is six-dimensional and the cost function is:

$$(24) \quad f(x_1, x_2, x_3, x_4, x_5, x_6) = 3x_4 - \left| \frac{2x_5}{x_1 w_0^2} \sin\left(\frac{w_0 x_6}{2}\right) \right|$$

where $w_0 = \sqrt{\frac{x_2 + x_3}{x_1}}$. The input variables are assumed independent, normal, with mean and variance parameters given in Table 4. A failure happens when the cost function is lower than the threshold $u = 0$. The probability of failure is approximately $1.514 \cdot 10^{-8}$, with a coefficient of variation of about 0.04%. This variant of the problem corresponds exactly to the harder case in [31].

4.2. Experimental settings.

4.2.1. BSS algorithm.

Initial design of experiments. We start with an initial design of size $n_0 = 5d$ (see [43] for a discussion on the size of the initial design in computer experiments), generated as follows. First, a compact subset $\mathbb{X}_0 \subset \mathbb{X}$ is constructed⁶:

$$\mathbb{X}_0 = \prod_{j=1}^d [q_\varepsilon^j; q_{1-\varepsilon}^j]$$

where q_ε^j and $q_{1-\varepsilon}^j$ are the quantiles of order ε and $1 - \varepsilon$ of the j^{th} input variable. Then, a “good” LHS design on $[0; 1]^d$ is obtained as the best design according the maximin criterion [38, 45] in a set of Q random LHS designs, and then scaled to \mathbb{X}_0 using an affine mapping. The values $\varepsilon = 10^{-5}$ and $Q = 10^4$ have been used in all our experiments.

Stochastic process prior. A Gaussian process prior with an unknown constant mean and a stationary anisotropic Matérn covariance function with regularity $5/2$ is used as our prior information about f (see Supplementary Material SM2 for more details). The unknown mean is integrated out as usual, using an improper uniform prior on \mathbb{R} ; as a consequence, the posterior mean coincides with the so-called “ordinary kriging” predictor.

⁶A similar technique is used by Dubourg and co-authors in a context of reliability-based design optimization [26, 29].

The remaining hyper-parameters (variance and range parameters of the covariance function) are estimated, following the empirical Bayes philosophy, by maximization of the marginal likelihood⁷. The hyper-parameters are estimated first on the data from the initial design, and then re-estimated after each new evaluation. In practice, we recommend to check the estimated parameters every once in a while using, e.g., leave-one-out cross-validation.

SMC parameters. Several values of the sample size m will be tested to study the relation between the variance of the estimator and the number of evaluations: $m \in \{500, 1000, 2000, \dots\}$. Several iterations of an adaptive Gaussian Random Walk Metropolis-Hastings (RWMH) algorithm, fully described in Supplementary Material SM3, are used for the move step of the algorithm.

Stopping criterion for the SUR strategy. The number of evaluations selected using the SUR strategy is determined adaptively, using the stopping criterion (21) from Section 3.5, with $\eta_t = 0.5$ for all $t < T$ (i.e., for all intermediate stages) and $\eta_T = 0.1 \hat{\delta}_{m,T}$ where $\hat{\delta}_{m,T}$ is the estimated coefficient of variation for the SMC estimator $\hat{\alpha}_T^{\text{BSS}}$ of α_T^{B} (see Appendix A). Furthermore, we require for robustness a minimal number N_{\min} of evaluations at each stage, with $N_{\min} = 2$ in all our simulations.

Adaptive choice of the thresholds. The successive thresholds u_t are chosen using the adaptive rule proposed in Section 3.4, Equation (20), with $p_0 = 0.1$. This value has been found experimentally to be neither too large (to avoid having a large number of stages) not too small (to avoid losing too many particles during the resampling step)⁸.

4.2.2. Subset simulation algorithm. The parameters used for the subset simulation algorithm are exactly the same, in all our simulations, as those used in the “SMC part” of the BSS algorithm (see Section 4.2.1). In particular, the number m_0 of surviving particles at each stage is determined according to the rule $p_0 = \frac{m_0}{m} = 0.1$ (see Table 1), and the adaptive MCMC algorithm described in Supplementary Material SM3 is used to move the particles. The number of evaluations made by the subset simulation algorithm is considered to be $m + (T - 1)(1 - p_0)m$, as explained in Section 2.4—in other words, in order to make the comparison as fair as possible, the additional evaluations required by the adaptive MCMC procedure are not taken into account.

4.2.3. ²SMART algorithm. ²SMART [10, 20] is another algorithm for the estimation of small probabilities, which is based on the combination of subset simulation with Support Vector Machines (SVM). We will present results obtained using the implementation of ²SMART proposed in the software package FERUM 4.1 [8], with all parameters set to their default values (which are equal to the values given in [10]).

Remark 10. Several other methods addressing the estimation of small probabilities of failure for expensive-to-evaluate functions have appeared recently in the structural reliability literature [4, 9, 12, 31, 37]. ²SMART was selected as a reference method due to the

⁷Used in combination with a uniform prior for the mean, for this specific model, the MML method is equivalent to the Restricted Maximum Likelihood (ReML) method advocated by [51], Section 6.4.

⁸Note that the considerations of Remark 4 on the choice of p_0 are not relevant here, since the computational cost of our method is mainly determined by the number of function evaluations, which is not directly related to the number of particles to be simulated. See also Supplementary Material SM4.

availability of a free software implementation in FERUM. A more comprehensive benchmark is left for future work.

4.3. Results.

4.3.1. Illustration. We first illustrate how BSS works using one run of the algorithm on Example 4.1.1 with sample size $m = 1000$. Snapshots of the algorithm at stages $t = 1$, $t = 5$ and $t = T = 9$ are presented on Figure 2. Observe that the additional evaluation points selected at each stage using the SUR criterion (represented by black triangles) are located in the vicinity of the current level set. The actual number of points selected at each stage, determined by the adaptive stopping rule, is reported in Table 5. Observe also that the set of particles (black dots in the right column) is able to effectively capture the bimodal target distribution. Finally, observe that a significant portion of the evaluation budget is spent on the final stage—this is again a consequence of our adaptive stopping rule, which refines the estimation of the final level set until the bias of the estimate is (on average under the posterior distribution) small compared to its standard deviation.

Table 5: Number of evaluations per stage on Example 4.1.1 (four-branch series system). For the BSS algorithm, recall that the number of evaluations at each stage is chosen adaptively (see Section 3.5) and is therefore random: the numbers shown here correspond to the run with $m = 1000$ that is shown on Figure 2. For the subset simulation algorithm, the number of evaluations is directly related to the m , $q_0 = 1 - p_0$ and T (see Section 2.4).

stage number t	0	1	2	3	4	5
BSS ($m = 1000$)	$2d = 10$	2	6	3	2	3
subset simulation	m	q_0m	q_0m	q_0m	q_0m	q_0m

stage number t	6	7	8	9	total
BSS ($m = 1000$)	2	3	2	28	61
subset simulation	q_0m	q_0m	q_0m	0	$m + (T - 1)q_0m$

4.3.2. Average results. This section presents average results over one hundred independent runs for subset simulation, BSS and ²SMART.

Figure 3 represents the average number of evaluations used by the BSS algorithm as a function of the sample size m . The number of evaluations spent on the initial design is constant, since it only depends on the dimension d of the input space. The average number of evaluations spent on the intermediate stages ($t < T$) is also very stable⁹ and independent of the sample size m . Only the average number of evaluations spent on the final stage—i.e., to learn the level set of interest—is growing with m . This growth is necessary if one wants the estimation error to decrease when m increases: indeed, the variance of $\hat{\alpha}_T^{\text{BSS}}$ automatically goes to zero at the rate $\frac{1}{m}$, but the bias $\hat{\alpha}_T^{\text{B}} - \alpha$ does not unless additional evaluations are added at the final level to refine the estimation of Γ .

Figure 4 represents the relative Root-Mean-Square Error (RMSE) of all three algorithms, as a function of the average number of evaluations. For the subset simulation

⁹Actually, for Examples 4.1.2 and 4.1.3, it is equal to $T N_{\min}$ for all runs; in other words, the adaptive stopping rule only came into play at intermediate stages for Example 4.1.1.

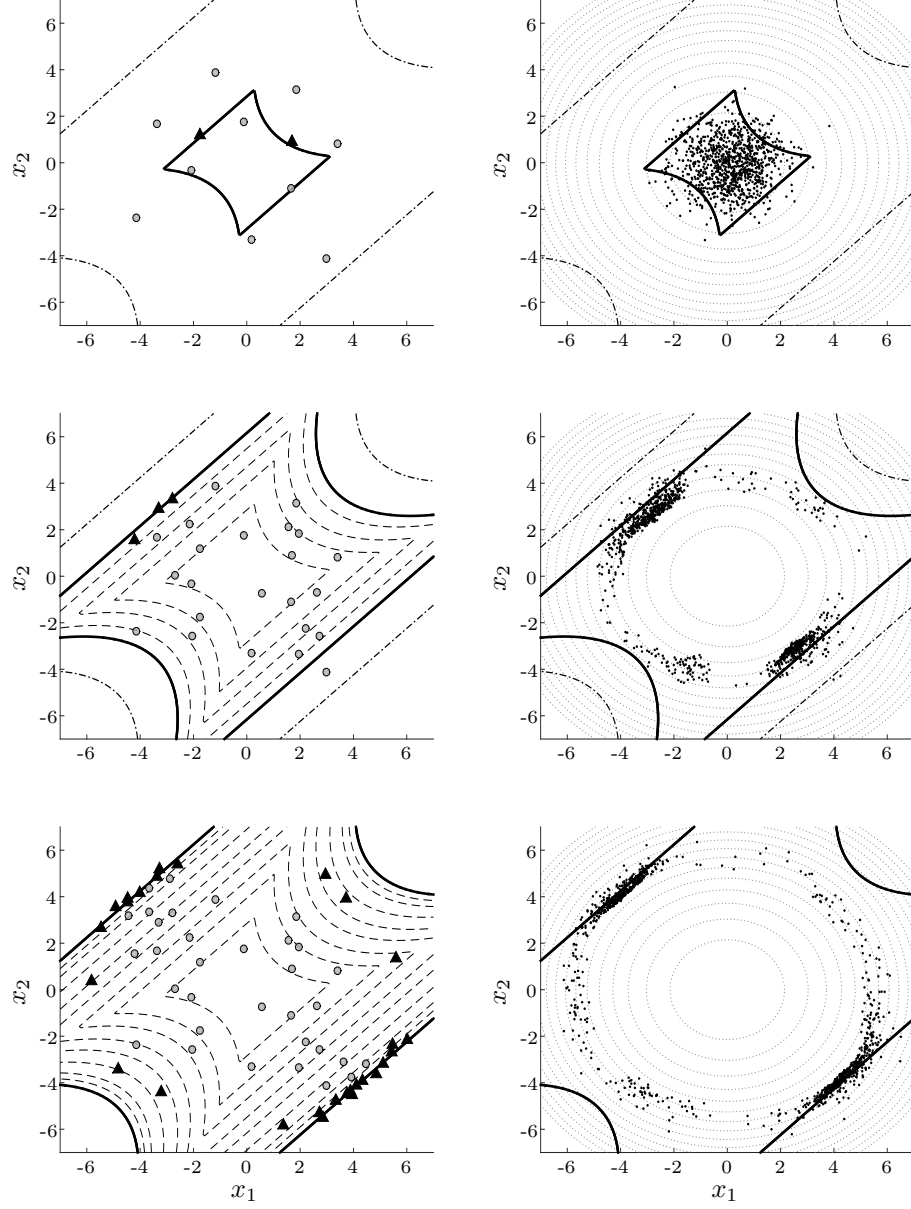


Figure 2: Snapshots of the BSS algorithm running on Example 4.1.1 (four branch series system) with sample size $m = 1000$. The first, second and third row correspond respectively to the end of the first stage ($t = 1$), the fifth stage ($t = 5$) and the last stage ($t = T = 9$). The true level set corresponding current target level u_t is represented by a thick line and, in the left column, true level sets corresponding to previous levels (u_s , $s < t$) are recalled using dashed contours. Evaluation points from previous stages are represented by gray disks (in particular, the initial design of experiment of size $n_0 = 10$ is visible on the top-left panel) and new evaluations performed at the current levels are marked by black triangles. In the right column, the sample points Y_{t-1}^j , $1 \leq j \leq m$ and the level sets of the input density $\pi_{\mathbb{X}}$ (corresponding to probabilities $1 - 10^{-k}$, $k = 1, 2, \dots$) are represented respectively by black dots and dotted lines.

algorithm, the number of evaluations is directly proportional to m and the RMSE decreases as expected like $\frac{1}{m}$ (with a constant much smaller than that of plain Monte Carlo simulation). ²SMART clearly outperforms subset simulation, but offers no simple way of tuning the accuracy of the final estimate (which is why only one result is presented, using the default settings of the algorithm). Finally, BSS clearly and consistently outperforms both ²SMART and subset simulation on these three examples: the relative RMSE goes to zero at a rate much faster than subset simulation's (a feature that is made possible by the smoothness of the limit-state function, which is leveraged by the Gaussian process model), and the sample size m is the only tuning parameters that needs to be acted upon in order change the accuracy of the final estimate. Figure 5 provides more insight into the error of the BSS estimate by confirming that, as intended by design of the adaptive stopping rule, variance is the main component of the RMSE (in other words, the bias is negligible in these examples).

Finally, note that the BSS estimation involves a computational overhead with respect to subset simulation. A careful analysis of the run times of BSS on the three examples (provided as Supplementary Material SM4.1) reveals that the computational overhead of BSS is approximately equal to $C_0 + C_1 m N_{\text{SUR}}$, where N_{SUR} the total number of evaluations selected using the SUR criterion (i.e., all evaluations except the initial design of experiments). This shows that, in our implementation, the most time-consuming part of the algorithm is the selection of additional evaluation points using the SUR criterion. However, in spite of its computational overhead, BSS is preferable to the subset simulation algorithm in terms of computation time, on the three test cases, as soon as the evaluation time τ_{sim} of f is large enough—larger than, say, 10 ms for the considered range of relative RMSE (see Supplementary Material SM4.2 for details). Consider for instance Example 4.1.1 with $\tau_{\text{sim}} = 1$ s: BSS with sample size $m = 8000$ achieves a relative RMSE of approximately 10% in about 3 minutes¹⁰ while subset simulation requires about 19 hours to achieve a comparable accuracy.

5. Discussion. We propose an algorithm called Bayesian subset simulation for the estimation of small probabilities of failure—or more generally the estimation of the volume of excursion of a function above a threshold—when the limit-state function is expensive to evaluate. This new algorithm is built upon two key techniques: the SMC method known as subset simulation or adaptive multilevel splitting on the one hand, and the Bayesian (Gaussian process based) SUR sampling strategy on the other hand. SMC simulation provides the means for evaluating the Bayesian estimate of the probability of failure, and to evaluate and optimize the SUR sampling criterion. In turn, the SUR sampling strategy makes it possible to estimate the level sets of the (smooth) limit-state function using a restricted number of evaluations, and thus to build a good sequence of target density for SMC simulation. Our numerical experiments show that this combination achieves significant savings in evaluations on three classical examples from the structural reliability literature.

An adaptive stopping rule is used in the BSS algorithm to choose the number of evalua-

¹⁰This computation time can be further decomposed as follows: 1 minute of evaluation time, corresponding to $\bar{N} = 63.2$ evaluations on average (see Figure 3) and 2 minutes of algorithm overhead.

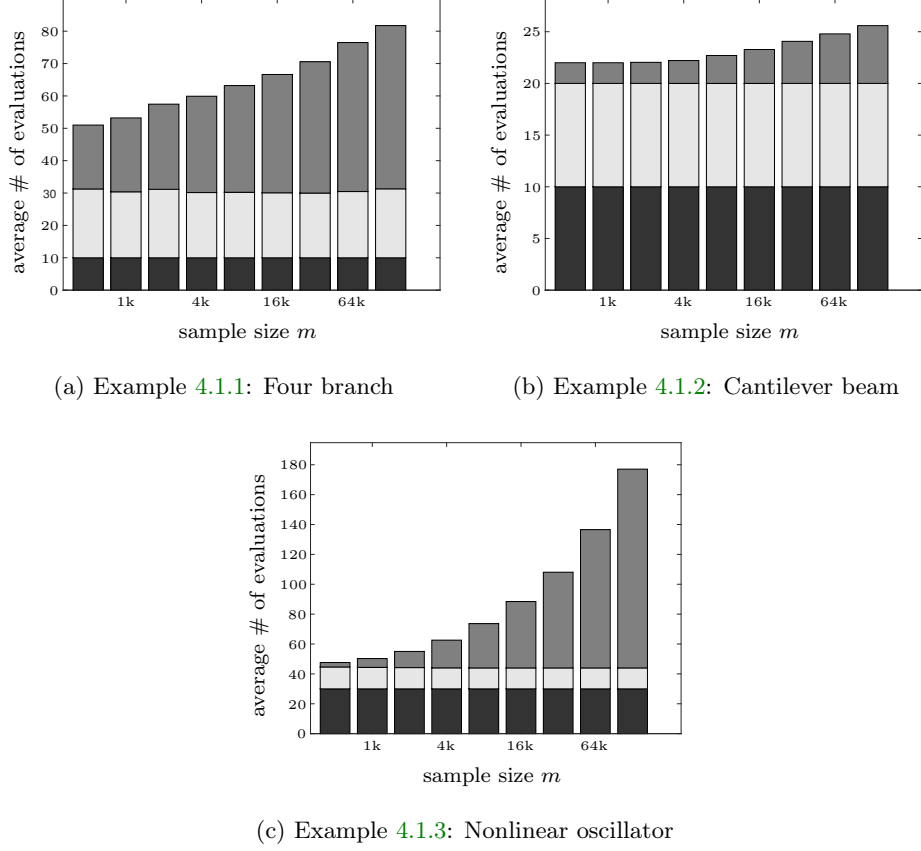


Figure 3: Average number of evaluations used by the BSS algorithm, over 100 independent runs, as a function of the sample size m on Examples 4.1.1–4.1.3. The total number of evaluations is split in three parts: the size n_0 of the initial design (dark gray), the number $\sum_{t=1}^{T-1} N_t$ of evaluations in intermediate stages (light gray) and the number of evaluations N_T in the final stage (middle gray).

tion added by the SUR sampling strategy at each stage. Evaluations at intermediate stages are not directly useful to refine the final probability estimate, but their importance must not be overlooked: they make it possible to learn in a robust way the level sets of the limit-state function, and therefore to build a sequence of densities that converges to the boundary of the failure region. Achieving a better understanding of the connection between the number of evaluations spent on intermediate level sets and the robustness of the algorithm is an important perspective for future work. In practice, if the budget of evaluations permits, we recommend running several passes of the BSS algorithm, with decreasing tolerances for the adaptive stopping rule, to make sure that no failure mode has been missed.

The adaptive stopping rule also makes it possible to refine the estimation of the final level set to make sure that the posterior model is good enough with respect to the SMC sample size. Other settings of the stopping rule could of course be considered. For instance, BSS could stop when the bias is expected to be of the same order as the standard deviation.

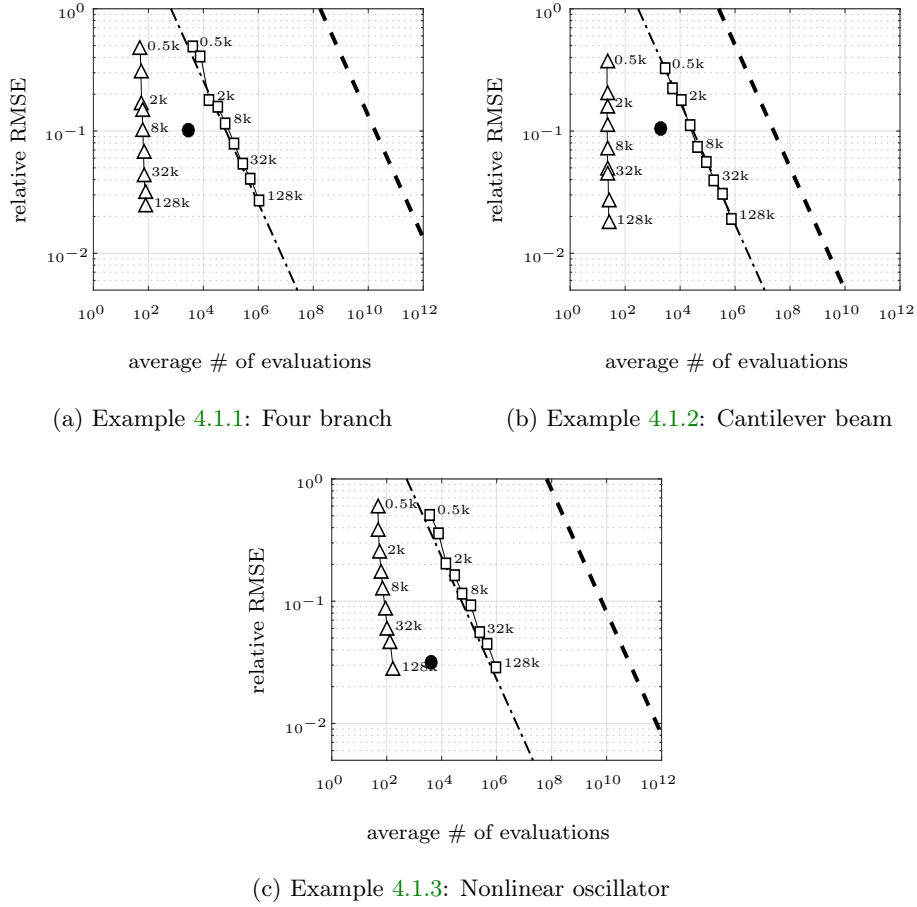


Figure 4: Relative root-mean-square error (RMSE) as a function of the average number of evaluations, over 100 independent runs, on Examples 4.1.1–4.1.3. For the subset simulation algorithm (squares) and for the BSS algorithm (triangles), the results are provided for several values of the sample size ($m \in \{500, 1000, 2000, \dots\}$). For the 2 SMART algorithm (filled circles), only one result is presented, corresponding to the default settings of the algorithm. The expected performance of plain Monte Carlo sampling is represented by a dashed line. The mixed line represents a simple approximation of the relative RMSE for the subset simulation algorithm: $\frac{T_\alpha}{m} \frac{1-p_0}{p_0}$, where $T_\alpha = \lceil \frac{\log \alpha}{\log p_0} \rceil$.

Future work will focus on fully automated variants on the BSS algorithm, where the number of evaluations and the SMC sample size would be controlled in order to achieve a prescribed error level.

Appendix A. Computation of the variance in the idealized setting.

This section provides a derivation of Equations (17) and (18), together with an explicit expression of the estimated coefficient of variation $\hat{\delta}_{m,T}$ used in Section 4.2.1. Both are obtained in the setting of the idealized BSS algorithm described in Section 3.4, where the samples $\mathbb{Y}_t = \{Y_t^1, \dots, Y_t^m\}$ are assumed IID (with density q_t) and mutually independent.

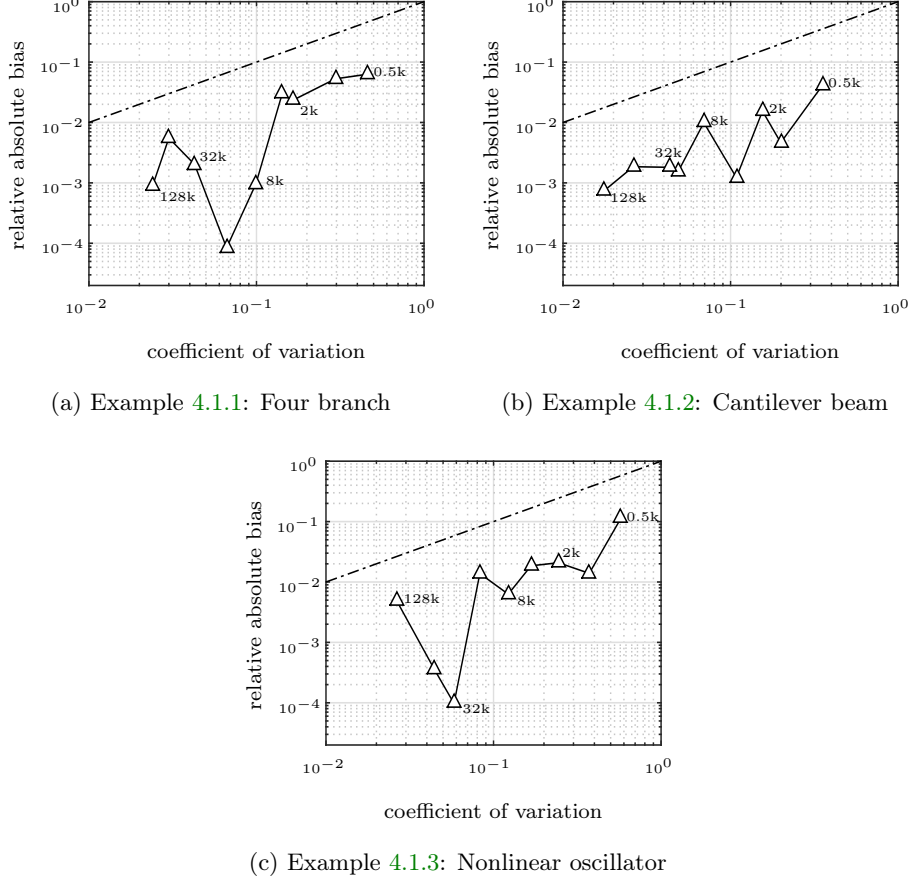


Figure 5: Relative absolute bias of the BSS estimator as a function of its coefficient of variation, estimated using 100 independent runs. The relative absolute bias is estimated using, for each test case, the reference value α_{ref} provided in Table 3.

Recall from Equation (16) that the BSS estimator can be written as

$$(25) \quad \hat{\hat{\alpha}}_T^{\text{BSS}} = \prod_{t=1}^T \frac{\hat{\hat{\alpha}}_t^{\text{BSS}}}{\hat{\hat{\alpha}}_{t-1}^{\text{BSS}}} = \prod_{t=1}^T \left(\frac{1}{m} \sum_{j=1}^m \frac{g_t(Y_{t-1}^j)}{g_{t-1}(Y_{t-1}^j)} \right) = \prod_{t=1}^T \hat{\hat{p}}_t^{\text{BSS}},$$

where we have set, for all $t \in \{1, \dots, T\}$,

$$\hat{\hat{p}}_t^{\text{BSS}} = \frac{1}{m} \sum_{j=1}^m \frac{g_t(Y_{t-1}^j)}{g_{t-1}(Y_{t-1}^j)}.$$

Observe that the random variables $\widehat{p}_t^{\text{BSS}}$ are independent, with mean

$$\mathbb{E} \left(\widehat{p}_t^{\text{BSS}} \right) = \int_{\mathbb{X}} \frac{g_t}{g_{t-1}} q_{t-1} = \int_{\mathbb{X}} \frac{g_t}{g_{t-1}} \frac{g_{t-1} \pi_{\mathbb{X}}}{\widehat{\alpha}_{t-1}^{\text{B}}} = \frac{\widehat{\alpha}_t^{\text{B}}}{\widehat{\alpha}_{t-1}^{\text{B}}} = \widehat{p}_t^{\text{B}}$$

and variance

$$\begin{aligned} \text{var} \left(\widehat{p}_t^{\text{BSS}} \right) &= \frac{1}{m} \text{var} \left(\frac{g_t (Y_{t-1}^1)}{g_{t-1} (Y_{t-1}^1)} \right) \\ &= \frac{1}{m} \left[\int_{\mathbb{X}} \frac{g_t^2}{g_{t-1}^2} \frac{g_{t-1} \pi_{\mathbb{X}}}{\widehat{\alpha}_{t-1}^{\text{B}}} - \left(\frac{\widehat{\alpha}_t^{\text{B}}}{\widehat{\alpha}_{t-1}^{\text{B}}} \right)^2 \right] \\ &= \frac{1}{m} \left[\frac{1}{\widehat{\alpha}_{t-1}^{\text{B}}} \int_{\mathbb{X}} \frac{g_t^2}{g_{t-1}} \pi_{\mathbb{X}} - \left(\frac{\widehat{\alpha}_t^{\text{B}}}{\widehat{\alpha}_{t-1}^{\text{B}}} \right)^2 \right] \\ &= \frac{1}{m} (\widehat{p}_t^{\text{B}})^2 \kappa_t, \end{aligned}$$

where κ_t is defined by (18). Therefore, the coefficients of variation $\delta_{m,t}$ of the sequence of estimators $\widehat{\alpha}_t^{\text{BSS}}$ obey the recurrence relation $\delta_{m,t}^2 = \frac{1}{m} \kappa_t + \left(1 + \frac{1}{m} \kappa_t\right) \delta_{m,t-1}^2$, and we conclude that

$$\begin{aligned} \delta_{m,T}^2 &= \frac{1}{m} \kappa_T + \left(1 + \frac{1}{m} \kappa_T\right) \frac{1}{m} \kappa_{T-1} \\ &\quad + \left(1 + \frac{1}{m} \kappa_T\right) \left(1 + \frac{1}{m} \kappa_{T-1}\right) \frac{1}{m} \kappa_{T-2} + \cdots \\ &= \frac{1}{m} \sum_{t=1}^T \kappa_t + O\left(\frac{1}{m^2}\right), \end{aligned}$$

which proves Equations (17)–(18). We construct an estimator of the coefficient of variation recursively, using the relation

$$\widehat{\delta}_{m,t}^2 = \frac{1}{m} \widehat{\kappa}_t + \left(1 + \frac{1}{m} \widehat{\kappa}_t\right) \widehat{\delta}_{m,t-1}^2,$$

with

$$\widehat{\kappa}_t = \left(\widehat{p}_t^{\text{BSS}}\right)^{-2} \cdot \frac{1}{m} \sum_{j=1}^m \left(\frac{g_t (Y_{t-1}^j)}{g_{t-1} (Y_{t-1}^j)} - \widehat{p}_t^{\text{BSS}} \right)^2.$$

References.

- [1] A. ARNAUD, J. BECT, M. COUPLET, A. PASANISI, AND E. VAZQUEZ, *Évaluation d'un risque d'inondation fluviale par planification séquentielle d'expériences*, in 42èmes Journées de Statistique (JdS 2010), Marseille, France, May 24–28, 2010.

- [2] S. K. AU AND J. BECK, *Estimation of small failure probabilities in high dimensions by subset simulation*, Probab. Eng. Mech., 16 (2001), pp. 263–277.
- [3] Y. AUFRAY, P. BARBILLON, AND J.-M. MARIN, *Bounding rare event probabilities in computer experiments*, Comput. Statist. Data Anal., 80 (2014), pp. 153–166.
- [4] M. BALESDENT, J. MORIO, AND J. MARZAT, *Kriging-based adaptive Importance Sampling algorithms for rare event estimation*, Struct. Saf., 44 (2013), pp. 1–10.
- [5] M. J. BAYARRI, J. O. BERGER, E. S. CALDER, K. DALBEY, S. LUNAGOMEZ, A. K. PATRA, E. B. PITMAN, E. T. SPILLER, AND R. L. WOLPERT, *Using statistical and computer models to quantify volcanic hazards*, Technometrics, 51 (2009), pp. 402–413.
- [6] J. BECT, D. GINSBOURGER, L. LI, V. PICHENY, AND E. VAZQUEZ, *Sequential design of computer experiments for the estimation of a probability of failure*, Stat. Comput., 22 (2012), pp. 773–793.
- [7] J. BECT, E. VAZQUEZ, ET AL., *STK: a Small (Matlab/Octave) Toolbox for Kriging. Release 2.4 (to appear)*, 2016, <http://kriging.sourceforge.net>.
- [8] J.-M. BOURINET, *Ferum 4.1 user's guide*. <http://www.ifma.fr/FERUM>, 2010.
- [9] J.-M. BOURINET, *Rare-event probability estimation with adaptive support vector regression surrogates*, Reliab. Eng. Syst. Saf., 150 (2016), pp. 210–221.
- [10] J.-M. BOURINET, F. DEHEGER, AND M. LEMAIRE, *Assessing small failure probabilities by combined subset simulation and support vector machines*, Struct. Saf., 33 (2011), pp. 343–353.
- [11] C.-E. BRÉHIER, L. GOUDENÈGE, AND L. TUDELA, *Central limit theorem for adaptive multilevel splitting estimators in an idealized setting*, in Monte Carlo and Quasi-Monte Carlo Methods (MCQMC 2014), R. Cools and D. Nuyens, eds., Springer, 2016.
- [12] F. CADINI, F. SANTOS, AND E. ZIO, *An improved adaptive kriging-based importance technique for sampling multiple failure regions of low probability*, Reliab. Eng. Syst. Saf., 131 (2014), pp. 109–117.
- [13] C. CANNAMELA, J. GARNIER, AND B. IOOSS, *Controlled stratification for quantile estimation*, Annals of Applied Statistics, 2 (2008), pp. 1554–1580.
- [14] C. CHEVALIER, J. BECT, D. GINSBOURGER, AND I. MOLCHANOV, *Estimating and quantifying uncertainties on level sets using the Vorob'ev expectation and deviation with gaussian process models*, in mODa 10 — Advances in Model-Oriented Design and Analysis, Contributions to Statistics, Springer, 2013, pp. 35–43.
- [15] C. CHEVALIER, J. BECT, D. GINSBOURGER, E. VAZQUEZ, V. PICHENY, AND Y. RICHET, *Fast parallel kriging-based stepwise uncertainty reduction with application to the identification of an excursion set*, Technometrics, 56 (2013), pp. 455–465.
- [16] N. CHOPIN, *A sequential particle filter method for static models*, Biometrika, 89 (2002), pp. 539–552.
- [17] F. CÉROU, P. DEL MORAL, T. FURON, AND A. GUYADER, *Sequential Monte Carlo for rare event estimation*, Stat. Comput., 22 (2012), pp. 795–808.
- [18] E. DE ROCQUIGNY, N. DEVICTOR, S. TARANTOLA, ET AL., *Determination of the risk due to personal electronic devices (PEDs) carried out on radio-navigation systems aboard aircraft*, in [19], 2008, ch. 5, pp. 65–80.
- [19] E. DE ROCQUIGNY, N. DEVICTOR, S. TARANTOLA, ET AL., *Uncertainty in industrial*

- practice*, Wiley, 2008.
- [20] F. DEHEEGER, *Couplage mécano-fiabiliste: ²SMART – Méthodologie d'apprentissage stochastique en fiabilité*, PhD thesis, Université B. Pascal (Clermont-Ferrand II), 2008.
 - [21] P. DEL MORAL, A. DOUCET, AND A. JASRA, *Sequential Monte Carlo samplers*, J. the Royal Statistical Society: Series B (Statistical Methodology), 68 (2006), pp. 411–436.
 - [22] P. DEL MORAL, A. DOUCET, AND A. JASRA, *An adaptive sequential Monte Carlo method for approximate Bayesian computation*, Stat. Comput., 22 (2012), pp. 1009–1020.
 - [23] P. DIACONIS AND S. HOLMES, *Three examples of Monte-Carlo Markov chains: at the interface between statistical computing, computer science, and statistical mechanics*, vol. 72 of IMA volumes in Mathematics and its Applications, Springer, 1995, pp. 43–56.
 - [24] R. DOUC AND O. CAPPÉ, *Comparison of resampling schemes for particle filtering*, in Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis (ISPA), 2005, pp. 64–69.
 - [25] R. DOUC AND E. MOULINES, *Limit theorems for weighted samples with applications to sequential Monte Carlo methods*, The Annals of Statistics, 36 (2008), pp. 2344–2376.
 - [26] V. DUBOURG, *Adaptive surrogate models for reliability analysis and reliability-based design optimization*, PhD thesis, Université Blaise Pascal – Clermont II, 2011.
 - [27] V. DUBOURG, F. DEHEEGER, AND B. SUDRET, *Metamodel-based importance sampling for the simulation of rare events*, in 11th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP 11), 2011.
 - [28] V. DUBOURG, F. DEHEEGER, AND B. SUDRET, *Metamodel-based importance sampling for structural reliability analysis*, Probab. Eng. Mech., 33 (2013), pp. 47–57.
 - [29] V. DUBOURG, B. SUDRET, AND J.-M. BOURINET, *Reliability-based design optimization using kriging surrogates and subset simulation*, Struct. Multiscip. Optim., 44 (2011), pp. 673–690.
 - [30] B. ECHARD, N. GAYTON, AND M. LEMAIRE, *AK-MCS: An active learning reliability method combining kriging and Monte Carlo simulation*, Struct. Saf., 33 (2011), pp. 145–154.
 - [31] B. ECHARD, N. GAYTON, M. LEMAIRE, AND N. RELUN, *A combined importance sampling and kriging reliability method for small failure probabilities with time-demanding numerical models*, Reliab. Eng. Syst. Saf., 111 (2013), pp. 232–240.
 - [32] P. FELIOT, J. BECT, AND V. E., *A Bayesian approach to constrained single- and multi-objective optimization*, J. Global Optim., (in press), pp. 1–37.
 - [33] FREE SOFTWARE FOUNDATION, *GNU Lesser General Public License, version 2.1*, <http://www.gnu.org/licenses/old-licenses/lgpl-2.1.html>.
 - [34] E. GARCIA, *Electromagnetic compatibility uncertainty, risk, and margin management*, IEEE Trans. Electromag. Compat., 52 (2010), pp. 3–10.
 - [35] N. GAYTON, J. M. BOURINET, AND M. LEMAIRE, *CQ2RS: a new statistical approach to the response surface method for reliability analysis*, Struct. Saf., 25 (2003), pp. 99–121.
 - [36] P. GLASSERMAN, P. HEIDELBERGER, P. SHAHABUDDIN, AND T. ZAJIC, *Multilevel splitting for estimating rare event probabilities*, Oper. Res., 47 (1999), pp. 585–600.
 - [37] X. HUANG, J. CHEN, AND H. ZHU, *Assessing small failure probabilities by AK-SS: An*

- active learning method combining kriging and subset simulation*, Struct. Saf., 59 (2016), pp. 86–95.
- [38] M. E. JOHNSON, L. M. MOORE, AND D. YLVISAKER, *Minimax and maximin distance designs*, J. Statist. Plan. Inference, 26 (1990), pp. 131–148.
 - [39] S. N. JONKMAN, M. KOK, AND J. K. VRIJLING, *Flood risk assessment in the netherlands: A case study for dike ring south holland*, Risk Analysis, 28 (2008), pp. 1357–1374.
 - [40] L. LI, *Sequential Design of Experiments to Estimate a Probability of Failure*, PhD thesis, Supélec, May 2012, <https://tel.archives-ouvertes.fr/tel-00765457>.
 - [41] L. LI, J. BECT, AND E. VAZQUEZ, *Bayesian subset simulation: a kriging-based subset simulation algorithm for the estimation of small probabilities of failure*, in Proceedings of PSAM 11 & ESREL 2012, 25–29 June 2012, Helsinki, Finland., 2012.
 - [42] J. S. LIU, *Monte Carlo strategies in scientific computing*, Springer, 2008.
 - [43] J. L. LOEPPKY, J. SACKS, AND W. J. WELCH, *Choosing the sample size of a computer experiment: A practical guide*, Technometrics, 51 (2009), pp. 366–376.
 - [44] R. E. MELCHERS, *Structural Reliability: Analysis and Prediction. Second Edition.*, Wiley, 1999.
 - [45] M. D. MORRIS AND T. J. MITCHELL, *Exploratory designs for computational experiments*, J. Statist. Plan. Inference, 43 (1995), pp. 381–402.
 - [46] J. OAKLEY, *Estimating percentiles of uncertain computer code outputs*, J. Roy. Statist. Soc. Ser. C, 53 (2004), pp. 83–93.
 - [47] P. P. O’CONNOR AND A. KLEYNER, *Practical Reliability Engineering*, Wiley, 2012.
 - [48] M. RAUSAND AND A. HOYLAND, *System reliability theory: models and statistical methods (second edition)*, Wiley, 2004.
 - [49] C. P. ROBERT AND G. CASELLA, *Monte Carlo statistical methods, 2nd edition*, Springer Verlag, 2004.
 - [50] T. J. SANTNER, B. J. WILLIAMS, AND W. NOTZ, *The Design and Analysis of Computer Experiments*, Springer Verlag, 2003.
 - [51] M. L. STEIN, *Interpolation of Spatial Data: Some Theory for Kriging*, Springer, New York, 1999.
 - [52] E. VAZQUEZ AND J. BECT, *A sequential Bayesian algorithm to estimate a probability of failure*, in Proceedings of the 15th IFAC Symposium on System Identification, SYSID 2009, Saint-Malo France, 2009.
 - [53] J. VILLEMONTAIX, E. VAZQUEZ, AND E. WALTER, *An informational approach to the global optimization of expensive-to-evaluate functions*, J. Global Optim., 44 (2009), pp. 509–534.
 - [54] P. H. WAARTS, *Structural reliability using finite element methods*, PhD thesis, Delft University of Technology, 2000.
 - [55] E. ZIO AND N. PEDRONI, *Estimation of the functional failure probability of a thermal-hydraulic passive system by subset simulation*, Nuclear Eng. Design, 239 (2009), pp. 580–599.
 - [56] K. M. ZUEV, J. L. BECK, S.-K. AU, AND L. S. KATAFYGIOTIS, *Bayesian post-processor and other enhancements of subset simulation for estimating failure probabilities in high dimensions*, Computers & Structures, 92–93 (2012), pp. 283–296.

SUPPLEMENTARY MATERIALS

SM1. Approximation and optimization of the SUR criterion. This section discusses the numerical procedure that we use for the approximation and optimization of the SUR criterion used at each stage of the BSS algorithm (see Sections 3.1 and 3.3):

$$J_n(x_{n+1}) = \int_{\mathbb{X}} \mathbf{E}_{n,x_{n+1}}(\tau_{n+1,u_t}(x)) \pi_{\mathbb{X}}(x) dx, \quad n_{t-1} \leq n \leq n_t - 1,$$

where we have introduced the simplified notation $\mathbf{E}_{n,x_{n+1}} := \mathbf{E}_n(\cdot \mid X_{n+1} = x_{n+1})$. The numerical approach proposed here is essentially the same as that used by [SM6], with a more accurate way of computing the integrand, following ideas of [SM2].

Observing that

$$J_n(x_{n+1}) = \hat{\alpha}_{t-1}^B \int_{\mathbb{X}} \frac{\mathbf{E}_{n,x_{n+1}}(\tau_{n+1,u_t}(x))}{g_{t-1}(x)} q_{t-1}(x) dx,$$

the integral over \mathbb{X} can be approximated, up to a constant, using the weighted sample \mathbb{Y}_{t-1} :

$$\begin{aligned} J_n(x_{n+1}) &\propto \int_{\mathbb{X}} \frac{\mathbf{E}_{n,x_{n+1}}(\tau_{n+1,u_t}(x))}{g_{t-1}(x)} q_{t-1}(x) dx \\ (SM1) \quad &\approx \sum_{j=1}^m w_{t-1}^j \frac{\mathbf{E}_{n,x_{n+1}}(\tau_{n+1,u_t}(x))|_{x=Y_{t-1}^j}}{g_{t-1}(Y_{t-1}^j)}. \end{aligned}$$

Then, simple computations using well-known properties of Gaussian processes under conditioning allow to obtain an explicit representation of the integrand, in the spirit of [SM2], as a function of the Gaussian process posterior mean $\hat{\xi}_n$ and posterior covariance k_n :

$$\begin{aligned} \mathbf{E}_{n,x_{n+1}}(\tau_{n+1,u_t}(x)) &= \Phi\left(\frac{u - \hat{\xi}_n(x)}{\sigma_n(x)}\right) + \Phi\left(\frac{u - \hat{\xi}_n(x)}{s_n(x, x_{n+1})}\right) \\ (SM2) \quad &- 2\Phi_2\left(\begin{pmatrix} u \\ u \end{pmatrix}; \begin{pmatrix} \hat{\xi}_n(x) \\ \hat{\xi}_n(x) \end{pmatrix}, \begin{pmatrix} s_n^2(x, x_{n+1}) & s_n^2(x, x_{n+1}) \\ s_n^2(x, x_{n+1}) & \sigma_n(x)^2 \end{pmatrix}\right), \end{aligned}$$

where Φ is the cumulative distribution function of the normal distribution, Φ_2 the cumulative distribution function of the bivariate normal distribution, $\sigma_n^2(x) = k_n(x, x)$ and $s_n^2(x, x_{n+1}) = k_n(x, x_{n+1})^2 / \sigma_n(x_{n+1})^2$.

The main computational bottleneck, in a direct implementation of the approximation (SM1) combined with the representation (SM2), is in our experience the computation of the bivariate cumulative distribution function Φ_2 . Indeed, assume that the optimization of the approximate criterion is carried out by means of an exhaustive discrete search on $\{Y_{t-1}^j, 1 \leq j \leq m\}$. Then m^2 evaluations of Φ_2 are required in order select X_{n+1} . To mitigate this problem, we implemented the pruning idea proposed in Section 3.4 of [SM6]: only a subset of size $m_0 \leq m$ of the set of particles is actually used, both for the approximation the integral and for the optimization of the criterion. In this article, the size m_0

is determined automatically as follows: first, for each particle Y_{t-1}^j , the current weighted probability of misclassification

$$\tilde{\tau}_n^j := w_{t-1}^j \tau_{n,u_t} \left(Y_{t-1}^j \right) / g_{t-1} \left(Y_{t-1}^j \right)$$

is computed. Then, the particles are sorted according to the value of $\tilde{\tau}_n^j$, in decreasing order: $\tilde{\tau}_n^{\varphi(1)} \geq \tilde{\tau}_n^{\varphi(2)} \geq \dots \geq \tilde{\tau}_n^{\varphi(m)}$, and m_0 is set to $\min(m_0^{\max}, m_0(\tilde{\tau}_n))$, where $m_0(\tilde{\tau}_n)$ is the smallest integer such that

$$\sum_{j=1}^{m_0} \tilde{\tau}_n^{\varphi(j)} \geq \rho \sum_{j=1}^m \tilde{\tau}_n^j.$$

The values $m_0^{\max} = 1000$ and $\rho = 0.99$ have been used in all our simulations.

SM2. Stochastic process prior. The stochastic process prior used for the numerical experiments in this article is a rather standard Gaussian process model. We describe it here in full detail for the sake of completeness. First, ξ is written as

$$\xi(x) = \mu + \xi^0(x),$$

where $\mu \in \mathbb{R}$ is an unknown constant mean and ξ^0 a zero-mean stationary Gaussian process with anisotropic covariance function

$$(SM3) \quad k(x, y) = \sigma^2 \kappa_\nu \left(\sqrt{\sum_{i=1}^d \frac{(x_{[i]} - y_{[i]})^2}{\rho_i^2}} \right), \quad x, y \in \mathbb{R}^d,$$

where $x_{[i]}, y_{[i]}$ denote the i^{th} coordinate of x and y , and κ_ν the Matérn correlation function of regularity ν (see [SM8], Section 2.10). The scale parameters ρ_1, \dots, ρ_d (characteristic correlation lengths) are usually called the *range* parameters of the covariance function.

In this article, the regularity parameter ν is set to the fixed value $\nu = 5/2$, leading to the following analytical expression for the Matérn correlation function:

$$(SM4) \quad \kappa_\nu(h) = \left(1 + \tilde{h} + \frac{1}{3} \tilde{h}^2 \right) \exp(-\tilde{h}), \quad \text{with } \tilde{h} = \sqrt{10} |h|.$$

As a consequence, ξ is twice differentiable in the mean-square sense, with sample paths almost surely in the Sobolev space $W^{s,2}$ for all $s < 5/2$ ([SM7], Theorem 3).

Remark SM1. The parameterization used in Equation (SM4) is the one advocated in [SM8]. Other parameterizations are sometimes used in the literature; e.g., [SM5] use $\tilde{h} = \sqrt{5} |h|$.

SM3. Adaptive Metropolis-Hastings algorithm for the move step. A fixed number S of iterations of a Gaussian Random Walk Metropolis-Hastings (RWMH) kernel is used for the move step, with adaptation of the standard deviations of the increments. More

precisely, for $s = 1, 2, \dots, S$, starting with the set of particles $(Y_{t,1}^{(j)})_{1 \leq j \leq m}$ produced by the resampling step, we first produce perturbed particles:

$$\tilde{Y}_{t,s}^{(j)} = Y_{t,s}^{(j)} + \Sigma_{\text{RW},t,s} U_{t,s}, \quad 1 \leq j \leq m,$$

where $\Sigma_{\text{RW},t,s}$ is a diagonal matrix and $U_{t,s}$ a d -dimensional standard normal vector. The perturbed particle is accepted as $Y_{t,s+1}^{(j)}$ with probability

$$a_{t,s}^{(j)} = 1 \wedge \frac{q_t(\tilde{Y}_{t,s}^{(j)})}{q_t(Y_{t,s}^{(j)})},$$

and $Y_{t,s}^{(j)}$ is kept otherwise. The k^{th} diagonal element $\sigma_{\text{RW},t,s}^{(k)}$ of $\Sigma_{\text{RW},t,s}$ is initialized with

$$(SM5) \quad \sigma_{\text{RW},1,1}^{(k)} = C_{\sigma,\text{init}} \sigma_{\mathbb{X}}^{(k)},$$

where $\sigma_{\mathbb{X}}^{(k)}$ is the standard deviation of the k^{th} marginal of $P_{\mathbb{X}}$, and then updated using

$$(SM6) \quad \log \sigma_{\text{RW},t,s}^{(k)} = \begin{cases} \log \sigma_{\text{RW},t,s-1}^{(k)} + \Delta_{\sigma}/s & \text{if } \bar{a}_{t,s} > a_{\text{target}} \\ \log \sigma_{\text{RW},t,s-1}^{(k)} - \Delta_{\sigma}/s & \text{otherwise,} \end{cases}$$

where $\sigma_{\text{RW},t,-1}^{(k)} = \sigma_{\text{RW},t-1,S}^{(k)}$, $\bar{a}_{t,s} = \frac{1}{m} \sum_{j=1}^m a_{t,s}^{(j)}$ is the average acceptance probability and a_{target} some prescribed target value.

The following parameter values have been used for the numerical simulations presented in this article: $S = 10$, $C_{\sigma,\text{init}} = 2/\sqrt{d}$, $\Delta_{\sigma} = \log(2)$ and $a_{\text{target}} = 30\%$.

Remark SM2. The $1/\sqrt{d}$ scaling for the constant $C_{\sigma,\text{init}}$ in (SM5) is motivated by the well-known theorem of [SM6], which provides the optimal covariance matrix $2.38^2 \Sigma/d$ for a Gaussian target with covariance matrix Σ in high dimension.

Remark SM3. The reader is referred to [SM9] and references therein for more information on adaptive MCMC algorithms, including adaptive scaling algorithms such as (SM6). Note, however, that the adaptation scheme proposed here is not strictly-speaking an adaptive MCMC scheme, since we use the entire population of particles to estimate the acceptance probability.

Remark SM4. Our adaptive scheme (SM5)–(SM6) is admittedly very simple, but works well in the examples of the paper. More refined adaptation schemes might be needed to deal with harder problems. For instance, it might become necessary to implement a truly anisotropic adaptation of the covariance matrix, or to use different kernels in different regions of the input space. Ideas from the adaptive SMC literature (see, e.g., [SM1, SM3, SM4]) could be leveraged to achieve these goals, which fall out of the scope of the present paper.

SM4. Run time of the BSS algorithm.

SM4.1. Overhead of the BSS algorithm. The median run times of the Bayesian subset simulation (BSS) and subset simulation (SS) algorithms on the three test cases of Section 4 are reported in Table SM1. Since the function f in each of these examples is actually very fast to evaluate, these run times provide a good indication of the numerical complexity of the algorithms. Consequently, they provide a measure of the *overhead* of the algorithm if it were actually run on an expensive-to-evaluate function (i.e., the fraction of the computation time not dedicated to evaluating the function).

As expected, because of the computations related to Gaussian process modeling, the overhead of BSS is larger than the one of subset simulation, for a given sample size m . In our implementation, this overhead, which we denote by $\tau_{\text{BSS},0}$, can be explained by a simple linear model:

$$(SM7) \quad \tau_{\text{BSS},0} \approx C_0 + C_1 m N_{\text{SUR}}$$

where m is the sample size and N_{SUR} the total number of evaluations selected using the SUR criterion (i.e., all evaluations except the initial design of experiments). On our standard Intel-Nehalem-based workstation, an ordinary least-square regression using Equation (SM7), on the $100 \times 3 \times 8 = 2400$ points of our data set (100 runs on 3 cases with 8 sample sizes), yields $C_0 \approx 27.0$ and $C_1 \approx 5.2 \cdot 10^{-5}$, with a coefficient of determination of 97.2%. We conclude that the most time-consuming part of the algorithm is the selection of additional evaluation points using the SUR criterion.

Remark SM5. Note that the computation time of BSS would grow as a function of $m^2 N_{\text{SUR}}$, instead of $m N_{\text{SUR}}$, without the pruning idea explained in Section A.

SM4.2. Extrapolation to expensive-to-evaluate functions. Let us now extrapolate from the available data, obtained for cheap-to-evaluate test functions, to the case of a non-negligible evaluation time. To this end, assume that each evaluation actually costs τ_{sim} in computation time. Then, the total run time of BSS becomes

$$(SM8) \quad \tau_{\text{BSS}} = \tau_{\text{BSS},0} + \tau_{\text{sim}} N,$$

where N is the total number of evaluations. For a given test case and a given sample size m , denote by \bar{N} the average total number of evaluation, $\bar{\tau}_{\text{BSS},0}$ the average overhead and rRMSE the relative root-mean-square error (computed using 100 runs of BSS). The efficiency of BSS with respect to the subset simulation algorithm can be measured by the ratio

$$(SM9) \quad \varepsilon_{\text{BSS}|\text{SS}} = \frac{\tau_{\text{SS}}}{\bar{\tau}_{\text{BSS}}}$$

where $\bar{\tau}_{\text{BSS}} = \bar{\tau}_{\text{BSS},0} + \tau_{\text{sim}} \bar{N}$ is the average total run time for BSS, and τ_{SS} is an approximation of the time it would take for subset simulation to reach the same relative RMSE, computed as:

$$(SM10) \quad \tau_{\text{SS}} = \tau_{\text{sim}} m (1 + (T_\alpha - 1)(1 - p_0)),$$

m		0.5k	1k	2k	4k	8k	16k	32k	64k	128k
Ex 4.1.1	BSS:	15.9	28.3	64.7	99.2	117.8	142.6	182.1	266.5	453.8
	SS:	0.1	0.1	0.1	0.1	0.2	0.3	0.7	1.3	3.1
Ex 4.1.2	BSS:	3.3	3.4	4.1	5.9	10.7	20.7	41.5	70.7	120.3
	SS:	0.0	0.1	0.1	0.1	0.2	0.3	0.6	1.2	2.5
Ex 4.1.3	BSS:	5.9	7.9	10.4	16.4	30.9	64.8	156.5	378.9	994.4
	SS:	0.1	0.1	0.1	0.2	0.3	0.6	1.3	2.8	6.4

Table SM1: Run time in seconds of the Bayesian subset simulation (BSS) and subset simulation (SS) algorithms on an Intel-Nehalem-based workstation (without parallelization). For each test case and each sample size m , the reported computation time is the median over 100 hundred runs of the algorithm.

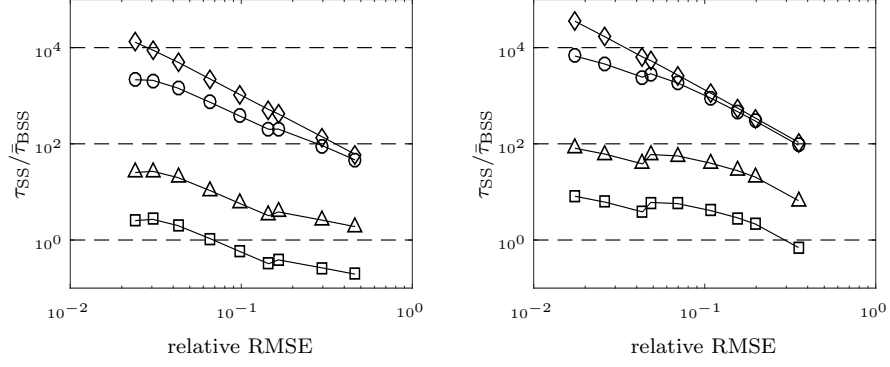
where $T_\alpha = \lceil \frac{\log \alpha}{\log p_0} \rceil$ and m is the sample size that gives the same relative RMSE according to the relation

$$(SM11) \quad \text{rRMSE} \approx \frac{T_\alpha}{m} \frac{1 - p_0}{p_0}$$

(cf. Figure 4). The efficiency $\varepsilon_{\text{BSS}|\text{SS}}$ is represented on Figure SM1, for the three test cases, as a function of the relative RMSE. It appears clearly for the three examples that, in spite of its computational overhead due to Gaussian process modeling, BSS is able to provide a significant time saving as soon as the evaluation time τ_{sim} of f is large enough—larger than, say, 10 ms for the considered range of relative RMSE.

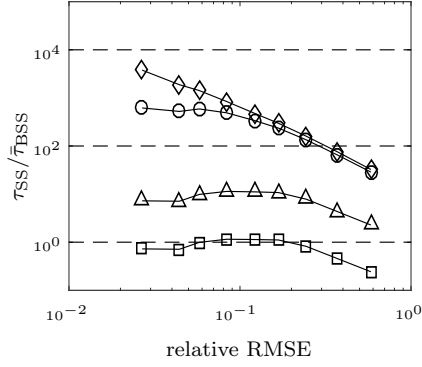
References for the Supplementary Materials.

- [SM1] B. L. CHANG, A. DOUCET, AND V. B. TADIC, *Optimisation of particle filters using simultaneous perturbation stochastic approximation*, in Proceedings of the 2003 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2003, pp. 681–684.
- [SM2] C. CHEVALIER, J. BECT, D. GINSBOURGER, E. VAZQUEZ, V. PICHENY, AND Y. RICHET, *Fast parallel kriging-based stepwise uncertainty reduction with application to the identification of an excursion set*, Technometrics, 56 (2013), pp. 455–465.
- [SM3] J. CORNEBISE, É. MOULINES, AND J. OLSSON, *Adaptive methods for sequential importance sampling with application to state space models*, Stat. Comput., 18 (2008), pp. 461–480.
- [SM4] P. FEARNHEAD AND B. M. TAYLOR, *An adaptive sequential monte carlo sampler*, Bayesian analysis, 8 (2013), pp. 411–438.
- [SM5] C. E. RASMUSSEN AND C. K. I. WILLIAMS, *Gaussian Processes for Machine Learning*, MIT Press, 2006.
- [SM6] G. O. ROBERTS, A. GELMAN, AND W. R. GILKS, *Weak convergence and optimal scaling of random walk Metropolis algorithms*, The Annals of Applied Probability, 7 (1997), pp. 110–120.
- [SM7] M. SCHEUERER, *Regularity of the sample paths of a general second order random field*, Stochastic Processes and their Applications, 120 (2010), pp. 1879–1897.



(a) Example 4.1.1: Four branch

(b) Example 4.1.2: Cantilever beam



(c) Example 4.1.3: Nonlinear oscillator

Figure SM1: Efficiency of BSS with respect to subset simulation, as a function of the relative RMSE, for several values of the virtual duration τ_{sim} of a single evaluation of the function: $\tau_{\text{sim}} = 1\text{ms}$ (squares), $\tau_{\text{sim}} = 10\text{ms}$ (triangles), $\tau_{\text{sim}} = 1\text{s}$ (circles) and $\tau_{\text{sim}} = 1\text{min}$ (diamonds).

[SM8] M. L. STEIN, *Interpolation of Spatial Data: Some Theory for Kriging*, Springer, New York, 1999.

[SM9] M. VIHOLA, *On the convergence of unconstrained adaptive Markov chain Monte Carlo algorithms*, PhD thesis, University of Jyväskylä, 2010.