



Sequential design of experiments to estimate a probability of exceeding a threshold in a multi-fidelity stochastic simulator

Rémi Stroh, Séverine Demeyer, Nicolas Fischer, Julien Bect, Emmanuel Vazquez

► To cite this version:

Rémi Stroh, Séverine Demeyer, Nicolas Fischer, Julien Bect, Emmanuel Vazquez. Sequential design of experiments to estimate a probability of exceeding a threshold in a multi-fidelity stochastic simulator. 61th World Statistics Congress of the International Statistical Institute (ISI 2017), Jul 2017, Marrakech, Morocco. <<http://www.isi2017.org>>. <hal-01569037>

HAL Id: hal-01569037

<https://hal-centralesupelec.archives-ouvertes.fr/hal-01569037>

Submitted on 26 Jul 2017

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.



Sequential design of experiments to estimate a probability of exceeding a threshold in a multi-fidelity stochastic simulator

Rémi STROH*

Laboratoire National de métrologie et d'Essais (LNE), Trappes, France - remi.stroh@lne.fr
Laboratoire des Signaux & Systèmes (L2S), CentraleSupélec, Univ. Paris-Sud, Centre National de la Recherche Scientifique (CNRS), Université Paris-Saclay, Gif-sur-Yvette, France

Séverine DEMEYER, Nicolas FISCHER
LNE, Trappes, France - name.surname@lne.fr

Julien BECT, Emmanuel VAZQUEZ
L2S, CentraleSupélec, Univ. Paris-Sud, CNRS, Université Paris-Saclay, Gif-sur-Yvette, France - name.surname@centralesupelec.fr

Abstract

In this article, we consider a stochastic numerical simulator to assess the impact of some factors on a phenomenon. The simulator is seen as a black box with inputs and outputs. The quality of a simulation, hereafter referred to as fidelity, is assumed to be tunable by means of an additional input of the simulator (e.g., a mesh size parameter): high-fidelity simulations provide more accurate results, but are time-consuming. Using a limited computation-time budget, we want to estimate, for any value of the physical inputs, the probability that a certain scalar output of the simulator will exceed a given critical threshold at the highest fidelity level. The problem is addressed in a Bayesian framework, using a Gaussian process model of the multi-fidelity simulator. We consider a Bayesian estimator of the probability, together with an associated measure of uncertainty, and propose a new multi-fidelity sequential design strategy, called Maximum Speed of Uncertainty Reduction (MSUR), to select the value of physical inputs and the fidelity level of new simulations. The MSUR strategy is tested on an example.

Keywords: Multi-fidelity; Sequential design; Bayesian analysis; Gaussian process.

1. Introduction

The objective of this article is to propose a new Bayesian algorithm for *sequential design of experiments* in the context of *multi-fidelity stochastic* numerical simulators. A numerical simulator is a computer program modeling a physical phenomenon or a system. When the simulator is deterministic, running the computer program twice using the same inputs yields the same outputs. In this case, the simulator can be viewed formally as a function. When the simulator is stochastic, running twice the simulator with the same inputs does not return the same output. Moreover, we assume that the simulator has a particular input called *fidelity parameter*—e.g., the mesh size of a finite-difference partial differential equation solver—that controls a trade-off between quality of simulation and computation time. A high-fidelity simulation provides an accurate result, but is time-consuming.

Let $t \in \mathbb{T}$ be the fidelity input of the simulator, and $x \in \mathbb{X}$ the vector of all other inputs. At fixed (x, t) , the output of the simulator Z follows a probability distribution $\mathbb{P}_{x,t}^Z$. We denote by t^{HF} the value of the fidelity input associated to the highest available level of fidelity: this is the level of interest for the user of the simulator. In this article, we focus on comparing the output Z to a critical threshold z^{crit} —typically, a level of output that should not be exceeded by the physical phenomenon or the system under study. This comparison can be studied from two points of view: either locally, by considering the function p that gives for each point x the probability of exceeding the threshold at this point,

$$p(x) = \mathbb{P}_{x,t^{\text{HF}}}^Z (Z > z^{\text{crit}}), \quad x \in \mathbb{X}, \quad (1)$$

or globally, by computing the global probability P of exceeding the threshold, which may be written as

$$P = \int_{\mathbb{X}} p(x) \mathbb{P}_{\mathbb{X}}(dx), \quad (2)$$

where $\mathbb{P}_{\mathbb{X}}$ is a probability distribution which models uncertainty about the value of the input factors.

Our objective is to estimate these quantities by observing the stochastic outcomes $z_1, z_2 \dots$ of simulations at points $(x_1, t_1), (x_2, t_2) \dots$. The estimators are built using a Bayesian model of the simulator. More precisely, we assume a prior model about $\mathbb{P}_{x,t}^Z$ and compute the posterior distribution of the quantity of interest (either p or P) given $\chi_n = (x_i, t_i; z_i)_{i \leq n}$. In this article, we suggest a new sequential design algorithm to select points $(x_1, t_1), (x_2, t_2) \dots$ in order to obtain a *fast* reduction of the uncertainty about p or P . Our methods deals both with the stochastic nature of the simulator and the tunable simulation cost.

The paper is organized as follows. Section 2 sets the Bayesian framework that we consider for stochastic multi-fidelity simulators. Section 3 describes our sequential algorithm and our new sampling criterion called Maximum Speed of Uncertainty Reduction (MSUR). Section 4 illustrates the method on a test problem, and Section 5 concludes the paper.

2. Bayesian framework

Following our previous work on multi-fidelity stochastic simulators (Stroh et al., 2016), assume that the output Z of a simulation at (x, t) follows a normal distribution:

$$Z \mid \xi, \lambda \sim \mathcal{N}(\xi(x, t), \lambda(x, t)). \quad (3)$$

Assume moreover that conditional on ξ and λ , all runs of the simulator are independent. Of course, the choice of a normal distribution here is a simplifying hypothesis, that must be verified in practice for each particular simulator¹.

Assume a Gaussian process prior for the mean function ξ :

$$\xi \mid m, k \sim \text{GP}(m, k((x, t), (x', t'))), \quad m \sim \mathcal{U}(\mathbb{R}), \quad (4)$$

where m is the (constant) mean of the Gaussian process, k its covariance function, and $\mathcal{U}(\mathbb{R})$ the (improper) uniform distribution on \mathbb{R} . For the sake of simplicity, the functions λ and k are assumed to be known. Under this prior, the posterior distribution of ξ conditionally to the observations χ_n is Gaussian:

$$\xi \mid \chi_n \sim \text{GP}(m_n(x, t), k_n((x, t), (x', t'))), \quad (5)$$

where m_n and k_n are given by the kriging formulas (see, e.g., Stein, 1999).

Since Z has a Gaussian distribution conditional on ξ and λ , for all $x \in \mathbb{X}$, the probability of exceeding a critical value z^{crit} can be written as

$$p(x) = \Phi \left(\frac{\xi(x, t^{\text{HF}}) - z^{\text{crit}}}{\sqrt{\lambda(x, t^{\text{HF}})}} \right), \quad (6)$$

with Φ the cumulative distribution function of the normal distribution. Then, denoting by $\mathbb{E}_n = \mathbb{E}(\cdot \mid \chi_n)$ and $\mathbb{V}\text{ar}_n = \mathbb{V}\text{ar}(\cdot \mid \chi_n)$ the posterior mean and variance operators, we have:

$$\mathbb{E}_n(p(x)) = \Phi(u_n(x)), \quad (7)$$

$$\mathbb{V}\text{ar}_n(p(x)) = \Phi_2(u_n(x), u_n(x); r_n(x, x)) - \Phi(u_n(x))^2, \quad (8)$$

where $\Phi_2(\cdot, \cdot; r)$ stands for the bivariate normal distribution function with correlation r , $V_n(x, t) = \lambda(x, t) + k_n((x, t), (x, t))$, $u_n(x) = \frac{m_n(x, t^{\text{HF}}) - z^{\text{crit}}}{\sqrt{V_n(x, t^{\text{HF}})}}$, and $r_n(x, x') = \frac{k_n((x, t^{\text{HF}}), (x', t^{\text{HF}}))}{\sqrt{V_n(x, t^{\text{HF}})V_n(x', t^{\text{HF}})}}$.

¹Other (possibly non-parametric) families of distributions could be considered as well. Note, however, that the convenient conjugacy property of the Gaussian process prior with respect to the Gaussian likelihood would be lost by doing so.

3. Proposed algorithm

3.1 A SUR criterion to estimate a probability of exceeding a threshold in the case of stochastic outputs

The goal of sequential design of experiments is to select a simulation point (x_{n+1}, t_{n+1}) , using the result of the previous observations χ_n , to obtain a good improvement of our knowledge about a given quantity of interest. Our algorithm is based on the idea of Stepwise Uncertainty Reduction (SUR) strategies (Vazquez and Bect, 2009; Bect et al., 2012). The main idea of SUR strategies is to define a measure of uncertainty H_n about the estimator of a quantity of interest. Then, assuming that there is no fidelity parameter t , the next observation point is selected in order to minimize the expected uncertainty using this new observation, whose outcome is random:

$$x_{n+1} = \operatorname{argmin}_{x \in \mathbb{X}} \mathbb{E}_n (H_{n+1} \mid X_{n+1} = x). \quad (9)$$

Several measures of uncertainty H_n have been proposed in the literature (Bect et al., 2012; Chevalier et al., 2014; Azzimonti et al., 2016) for the problem of estimating the volume or contour of an excursion set of a deterministic function. In our framework, however, we deal with a stochastic simulator, and we propose a new measure of uncertainty for this case. Let L be the L^2 loss function

$$L(f, g) = \|f - g\|_{L^2(\mu)}^2 = \int_{\mathbb{X}} (f(x) - g(x))^2 \mu(dx), \quad (10)$$

where μ denotes a positive measure on \mathbb{X} . We suggest to measure uncertainty by the L^2 loss incurred when estimating p by $\hat{p}_n = \mathbb{E}_n(p)$:

$$H_n = \mathbb{E}_n(L(\hat{p}_n, p)) = \int_{\mathbb{X}} \operatorname{Var}_n(p(x)) \mu(dx). \quad (11)$$

Note that in the case where the goal is to estimate P , H_n provides an upper bound on the posterior variance for any $\mathbb{P}_{\mathbb{X}}$ that admits a density g with respect to μ : $\operatorname{Var}_n(P) \leq G H_n$, where $G = \int g^2 d\mu$ and $g = d\mathbb{P}_{\mathbb{X}}/d\mu$. Consequently, a sequential design strategy that aims to reduce H_n will be useful not only to estimate p itself but also to estimate P for a large class of probability distributions $\mathbb{P}_{\mathbb{X}}$.

Let $J_n(x, t) = \mathbb{E}_n(H_{n+1} \mid X_{n+1} = x, T_{n+1} = t)$. Since the output of the simulator has a Gaussian distribution, $J_n(x, t)$ can be written as

$$J_n(x, t) = \int_{\mathbb{X}} \left(\Phi_2(u_n(y), u_n(y); r_n(y, y)) - \Phi_2\left(u_n(y), u_n(y); \frac{k_n((y, t^{\text{HF}}), (x, t))^2}{V_n(y, t^{\text{HF}}) V_n(x, t)}\right) \right) \mu(dy). \quad (12)$$

3.2 Dealing with tunable fidelity

The SUR strategy (9) is relevant when there is no fidelity parameter t . In a multi-fidelity context, however, the simulator becomes more and more time-consuming as the fidelity parameters comes closer to t^{HF} . We propose a new sequential design algorithm which takes into account the variable simulation cost.

Let $C(x, t) > 0$ be the cost of observing the simulator at (x, t) . The idea is to balance the benefit of an observation against its cost. For us, the benefit of an observation is the reduction of uncertainty. Hence, we propose the following sequential strategy called Maximum Speed of Uncertainty Reduction (MSUR):

$$(x_{n+1}; t_{n+1}) = \operatorname{argmax}_{(x, t) \in \mathbb{X} \times \mathbb{T}} \frac{H_n - J_n(x, t)}{C(x, t)}. \quad (13)$$

This idea is similar to the one proposed by Huang et al. (2006) in the context of optimization, where the expected improvement is divided by the cost of an observation. Likewise, for the purpose of providing a surrogate model of the simulator, Le Gratiet and Cannamela (2015) suggested to choose the new level of fidelity by comparing the cost of the level and the reward in terms of reduction of uncertainty.

Note that in many applications, the cost C depends only on the level t : $C(x, t) = C(t)$. In this case, (13) can be divided into two steps: first, select the input x at each level that minimizes the SUR criterion; second, select the level t that maximizes the MSUR criterion. Our sequential design algorithm can thus be rewritten in this case as

$$\begin{cases} \tilde{x}(t) = \underset{x \in \mathbb{X}}{\operatorname{argmin}} J_n(x, t), \\ t_{n+1} = \underset{t \in \mathbb{T}}{\operatorname{argmax}} \frac{H_n - J_n(\tilde{x}(t), t)}{C(t)}, \\ x_{n+1} = \tilde{x}(t_{n+1}). \end{cases} \quad (14)$$

4. Illustration

We illustrate our algorithm on a two-dimensional example inspired by Au and Beck (2001). We consider a damped harmonic oscillator subject to random forcing, described by the stochastic differential equation

$$\dot{X}(u) + 2\zeta\omega_0\dot{X}(u) + \omega_0^2 X(u) = W(u), \quad X(0) = 0, \quad \dot{X}(0) = 0, \quad u \in [0; u_{\max} = 30 \text{ s}], \quad (15)$$

where W is a Gaussian white noise (with spectral density equal to one). We compute an approximation $(\hat{X}_n)_{n=0, \dots, \lfloor u_{\max}/\Delta \rfloor}$ of X by finite differences with a time step Δ . We use an explicit exponential Euler scheme (see, e.g., Jentzen and Kloeden, 2009). The output of our simulator is defined as

$$\begin{aligned} Z: [0; 30] \text{ rad s}^{-1} \times [0; 1] \times [0; 1] \text{ s} &\rightarrow \mathbb{R} \\ (\omega_0, \zeta, \Delta) &\mapsto \log \left(\max \{ \hat{X}_n, n \leq \lfloor \frac{u_{\max}}{\Delta} \rfloor \} \right). \end{aligned} \quad (16)$$

The pair (ω_0, ζ) corresponds to the input vector x and the step time Δ plays the role of the fidelity parameter. The critical threshold z^{crit} is set to -3 . The highest level of fidelity is set to $t^{\text{HF}} = 0.01 \text{ s}$. The computational cost is empirically linear with respect to the fidelity level: $C(\Delta) = a/\Delta + b$, with $a = 0.0098$ and $b = 0.02$ (coefficients chosen to have $C(t^{\text{HF}}) = 1$). Figure 1 represents the contour plots of the mean function, the standard deviation and the probability function at three different levels of fidelity.

We compute a reference value p^* of p on a 100×100 regular grid at the highest level of fidelity using 10^4 simulations of Z (see Figure 1i). We suppose that the input distribution $\mathbb{P}_{\mathbb{X}}$ is the uniform distribution on $[0; 30] \times [0; 1]$ and let $\mu = \mathbb{P}_{\mathbb{X}}$. The mean value $P^* = 83.3\%$ of p^* on the grid serves as a reference value for P .

We use the Bayesian model of Section 2 with fixed hyper-parameters, and ten fixed levels of fidelity: $\Delta = 1, 0.5, 0.33, 0.25, 0.2, 0.17, 0.1, 0.05, 0.02$ and 0.01 s . The initial design (initial simulations before applying our MSUR strategy) consists of $180 \times 60 \times 20 \times 10 \times 5$ nested observations on the five lowest fidelity levels (180 on the level $\Delta = 1 \text{ s}$, 60 on the level $\Delta = 0.5 \text{ s}$, ...). The initial design is set using the algorithm of Qian (2009).

In our experiments, we compare two strategies: five Single-Level SUR (SL-SUR) strategies and the MSUR strategy (14). For SL-SUR strategies, all new points are sequentially selected on a fixed level, using then SUR strategy (9). Each SL-SUR strategy corresponds to one level of fidelity. Each time a new observation point must be selected, we choose the point that achieves the best value of J_n among 500 candidate points per level drawn from $\mathbb{P}_{\mathbb{X}}$. We allocate a simulation-time budget of 20 for each strategy. All experiments are repeated 12 times. Integrals are approximated by a sum on the 100×100 regular grid.

The strategies are compared based on the mean square error between estimations and references. Results are presented on Figure 2. Figure 2a represents the mean square error of the estimation \hat{P}_n of P , and Figure 2b represents the mean square error of the estimation \hat{p}_n of p , as a function of the computational cost. Each curve corresponds to an average over 12 experiments with the same method of sequential design. The plain red curve corresponds to the MSUR strategy. The dotted crossed blue-green curves correspond to SL-SUR strategies (one curve per level). We can see that the best level is achieved




Figure 1: Our simulator Z defined by (16). Each column corresponds to a level of fidelity (from left to right, $\Delta = 1, 0.1$ and 0.01 s). First row corresponds to the mean function. Second row corresponds to the standard deviation function. Third row corresponds to the probability of exceeding the threshold. The red bold line indicates the contour threshold.

at $\Delta = 0.05$ s. The lower fidelity levels are too biased to estimate the probabilities correctly, and the upper fidelity levels are too expensive to make it possible to carry out enough observations in order to estimate the probabilities accurately. Moreover, we can see that the MSUR strategy is as good as the SL-SUR strategy at $\Delta = 0.05$ s. Consequently, with the MSUR strategy, one does not need to know which level yields the best trade-off between accuracy and computational cost.

5. Conclusion

This article makes two contributions. First, we suggest a new SUR criterion to estimate a probability of exceeding a threshold in the case of a stochastic simulator. Second, we construct a sequential strategy called MSUR as an adaptation of the SUR strategies to deal with multi-fidelity simulators. Our first results are promising, because the MSUR strategy succeeds to get the better of all single-level SUR strategies without knowing which level provides the best compromise between speed and accuracy.




Figure 2: Mean square error of the quantities of interest get with the MSUR strategy and the Single-Level SUR strategies. The symbol $\langle \cdot \rangle$ means “average on the 12 repetitions”.

References

- Au, S.-K. and Beck, J. L. (2001). Estimation of small failure probabilities in high dimensions by subset simulation. *Probabilistic Engineering Mechanics*, 16(4):263–277.
- Azzimonti, D., Ginsbourger, D., Chevalier, C., Bect, J., and Richet, Y. (2016). Adaptive design of experiments for conservative estimation of excursion sets. *arXiv preprint arXiv:1611.07256*.
- Bect, J., Ginsbourger, D., Li, L., Picheny, V., and Vazquez, E. (2012). Sequential design of computer experiments for the estimation of a probability of failure. *Statistics and Computing*, 22(3):773–793.
- Chevalier, C., Bect, J., Ginsbourger, D., Vazquez, E., Picheny, V., and Richet, Y. (2014). Fast parallel kriging-based stepwise uncertainty reduction with application to the identification of an excursion set. *Technometrics*, 56(4):455–465.
- Huang, D., Allen, T. T., Notz, W. I., and Miller, R. A. (2006). Sequential kriging optimization using multiple-fidelity evaluations. *Structural and Multidisciplinary Optimization*, 32(5):369–382.
- Jentzen, A. and Kloeden, P. E. (2009). Overcoming the order barrier in the numerical approximation of stochastic partial differential equations with additive space–time noise. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 465(2102):649–667.
- Le Gratiet, L. and Cannamela, C. (2015). Cokriging-based sequential design strategies using fast cross-validation techniques for multi-fidelity computer codes. *Technometrics*, 57(3):418–427.
- Qian, P. Z. G. (2009). Nested latin hypercube designs. *Biometrika*, 96(4):957–970.
- Stein, M. L. (1999). *Interpolation of Spatial Data: Some Theory for Kriging*. Springer Series in Statistics. Springer, New York.
- Stroh, R., Bect, J., Demeyer, S., Fischer, N., and Vazquez, E. (2016). Gaussian process modeling for stochastic multi-fidelity simulators, with application to fire safety. In *48e Journées de Statistique de la SFdS (JdS 2016)*, Montpellier, France.
- Vazquez, E. and Bect, J. (2009). A sequential Bayesian algorithm to estimate a probability of failure. *IFAC Proceedings Volumes*, 42(10):546–550.