Gaussian process modeling for stochastic multi-fidelity simulators, with application to fire safety

Rémi STROH [1,2] & Julien BECT [2] & Séverine DEMEYER [1] & Nicolas FISCHER [1] & Emmanuel VAZQUEZ [2]

[1] *LNE, Service Mathématiques et Statistiques;*
29, avenue Roger Hennequin, 78190, Trappes; prenom.nom@lne.fr
[2] *Laboratoire des Signaux et Systèmes*
CentraleSupélec, CNRS, Univ. Paris-Sud, Université Paris Saclay
3, rue Joliot-Curie, 91190, Gif-sur-Yvette; prenom.nom@lss.supelec.fr

Résumé. Pour évaluer les possibilités d'évacuation d'un bâtiment lors d'un incendie, une méthode standard consiste à simuler la propagation d'un incendie, au moyen de modèles de type différences finies, et en prenant en compte le comportement aléatoire du feu, de sorte que le résultat d'une simulation est non-déterministe. La finesse du maillage détermine la qualité du modèle numérique, ainsi que son coût de calcul. En fonction de la taille des mailles, une seule simulation peut durer entre quelques minutes et quelques semaines. Dans cet article, nous cherchons à prédire le comportement du simulateur à une maille fine, à partir de résultats moins coûteux, à des mailles plus grossières. Dans la littérature de la conception et de l'analyse d'expériences numériques, on parle d'approche multi-fidélité. Notre contribution est d'étendre au cas de simulateurs stochastiques du modèle bayésien multi-fidélité proposé par Picheny et Ginsbourger (2013) et Tuo et al. (2014), .

Mots-clés. Expériences numériques, Processus gaussien, Multi-fidélité, Sécurité incendie

Abstract. To assess the possibility of evacuating a building in case of a fire, a standard method consists in simulating the propagation of fire, using finite difference methods and takes into account the random behavior of the fire, so that the result of a simulation is non-deterministic. The mesh fineness tunes the quality of the numerical model, and its computational cost. Depending on the mesh fineness, one simulation can last anywhere from a few minutes to several weeks. In this article, we focus on predicting the behavior of the fire simulator at fine meshes, using cheaper results, at coarser meshes. In the literature of the design and analysis of computer experiments, such a problem is referred to as multi-fidelity prediction. Our contribution is to extend to the case of stochastic simulators the Bayesian multi-fidelity model proposed by Picheny and Ginsbourger (2013) and Tuo et al. (2014).

Keywords. Numerical experiments, Gaussian process, Multi-fidelity, Fire safety

1   Introduction

Fire Dynamics Simulator (FDS) is a numerical simulator developed by the National Institute of Standards and Technology, that is used to simulate the propagation of fire in a building, and assess its conformity to fire safety standards. Fire Dynamics Simulator is based on a finite difference method, that takes into account the random behavior of fire propagation. Consequently, the outputs of Fire Dynamics Simulator are stochastic. Using smaller mesh size increases the quality of a simulation with respect to the physical reality, but also increases the computational cost (see Table 1). Thus, the mesh size controls a trade-off between speed and fidelity. In other words, Fire Dynamics Simulator is a *multi-fidelity* simulator.

Our work aims at estimating the behavior of Fire Dynamics Simulator at a very fine mesh, with a limited computational budget, using the result of simulations carried out using coarser mesh sizes. To do this, we use a Bayesian approach, where we construct a model of the output of Fire Dynamics Simulator as a function of the mesh size. Following Kennedy and O'Hagan (2000) and others, our approach is based on Gaussian process modeling.

Section 2 shows how to extend Bayesian multi-fidelity models proposed by Picheny and Ginsbourger (2013) and Tuo et al. (2014) in the case of deterministic simulators, to deal with the case of stochastic simulator. Section 3 presents numerical results to assess the quality of our new model.

2   Model for multi-fidelity

To formalize, consider n input-output pairs of a stochastic simulator with a tuning parameter, $\left(\left(\boldsymbol{x_i}, t_i\right), Z_i\right)_{1 \leq i \leq n} \in \left(\mathbb{X} \times \mathbb{T}\right) \times \mathbb{R}$, where $\mathbb{X} \subset \mathbb{R}^d$ and $\mathbb{T} \subset \mathbb{R}^+$. The outputs Z_i are supposed to be realizations of random variables, following distributions $\mathrm{P}_{\boldsymbol{x_i}, t_i}$. The results are mutually independent. In order to simplify, the distributions $\mathrm{P}_{\boldsymbol{x_i}, t_i}$ are assumed to be Gaussian distributions, with means $\xi\left(\boldsymbol{x_i}, t_i\right)$, and variances $\lambda\left(\boldsymbol{x_i}, t_i\right)$:

$$Z_i \sim \mathcal{N}\left(\xi\left(\boldsymbol{x_i}, t_i\right), \lambda\left(\boldsymbol{x_i}, t_i\right)\right). \tag{1}$$

To simplify, λ is supposed to depend only on t : $\lambda\left(\boldsymbol{x}, t\right) = \lambda\left(t\right)$. Besides, we add a Gaussian prior distribution on the mean process, ξ. We assumed that $\xi \sim GP(m, k)$. The prior distributions of ξ and λ are independent.

Mesh size t (cm)	100	50	33.33	25	20
Duration of one simulation (h)	1/12	1	6	20	54

Table 1: Approximate duration of one run of Fire Dynamics Simulator, on the example presented in Section 3, as a function of the tuning parameter.

A popular model for the process ξ was developed by Kennedy and O'Hagan (2000) : it is a recursive model, built in the case of finite number of levels, which links two successive Gaussian processes by a autoregressive relationship, $AR(1)$. However, this model is not well-suited to a simulator with a continuous tuning parameter. Indeed, even if this model can be extended for any number of levels (Le Gratiet, 2013), the number of covariance parameters increases strongly with the number of levels. Also, this model does not actually use the value of the tuning parameter, t.

For these reasons, another modeling was recently developed by Picheny and Ginsbourger (2013) and Tuo et al. (2014). They supposed that an ideal simulator can be thought up by setting the tuning parameter to an extreme value $(t = 0)$. Then, the process ξ can be written as

$$\xi(\boldsymbol{x}, t) = \xi_0(\boldsymbol{x}) + \varepsilon(\boldsymbol{x}, t), \tag{2}$$

where ξ_0 models this ideal simulator, and ε represents a deterministic numerical error, independent of ξ_0, which decreases when t tends to 0 : $\lim_{t \to 0} \mathbb{E}\left[\varepsilon(\boldsymbol{x}, t)^2\right] = 0$.

Following this decomposition, the covariance function of ξ, k, is the sum of two covariance functions : $k_{\xi_0}(\boldsymbol{x}, \boldsymbol{x}') + k_\varepsilon((\boldsymbol{x}, t), (\boldsymbol{x}', t'))$. Two assumptions are made : first, the covariance function k_ε is separable, $k_\varepsilon((\boldsymbol{x}, t), (\boldsymbol{x}', t')) = r(t, t') k_{X\varepsilon}(\boldsymbol{x}, \boldsymbol{x}')$, then, the two spatial covariance functions, k_{ξ_0} and $k_{X\varepsilon}$ are stationary. We choose anisotropic Matérn covariance functions for both. The correlation r is chosen as a function of Brownian covariance function: $r(t, t') = \min\{t, t'\}^L$, L a positive real parameter.

The mean of ξ, m, is supposed constant, with an improper uniform prior distribution on $\mathbb{R}$.

Finally, in order to improve the estimation of the observation variances, particularly on costly levels, we add a prior distribution on $(\lambda(t))_{t \in \mathbb{T}}$. This prior distribution describes two ideas : values of $\lambda(t)$ are not precisely-known a priori, so $\mathbb{V}\text{ar}[\lambda(t)]$ are large; but the variances are alike, so $\mathbb{V}\text{ar}[\lambda(t)|\lambda(t')]$ are small. Finally, a log-normal prior is chosen :

$$(\ln[\lambda(t)])_{t \in \mathbb{T}} \sim \mathcal{N}\left(\ln(\lambda_{prior})\mathbf{1}, \varsigma^2 I_d + s^2 \mathbb{1}\right), \tag{3}$$

with λ_{prior} equals to 1% of the range of the output; $\mathbf{1}$ is the vector of ones; I_d the identity matrix; $\mathbb{1}$ the square matrix of ones; $s^2 = \ln(10)^2 \gg \varsigma^2 = (\ln(2)/3)^2$.

Finally, our multi-fidelity non-stationary model is built as follows :

$$
\begin{aligned}
&(Z_i)_{1 \leq i \leq n}\Big|\xi, (\lambda(t))_{t \in \mathbb{T}} \sim \mathcal{N}\left((\xi(\boldsymbol{x}_i, t_i))_{1 \leq i \leq n}, \text{diag}\left\{(\lambda(t_i))_{1 \leq i \leq n}\right\}\right); \\
&\xi \sim GP(m, k); \\
&m(\boldsymbol{x}, t) = m \sim \mathcal{U}_{\mathbb{R}}; \\
&k((\boldsymbol{x}, t), (\boldsymbol{x}', t')) = k_{\xi_0}(\boldsymbol{x} - \boldsymbol{x}') + \min\{t, t'\}^L k_{X\varepsilon}(\boldsymbol{x} - \boldsymbol{x}'); \\
&(\ln[\lambda(t)])_{t \in \mathbb{T}} \sim \mathcal{N}\left(\ln(\lambda_{prior})\mathbf{1}, \varsigma^2 I_d + s^2 \mathbb{1}\right).
\end{aligned} \tag{4}
$$

3

Kind of design	Multi-fidelity design						High-fidelity design					
Property	Nested						Latin Hypercube					
Observations	t (cm)	100	50	33	25	20	t (cm)	100	50	33	25	20
	n_{points}	270	90	30	10	0	n_{points}	0	0	0	0	100
Speed	≈ 11 faster						1 (reference)					

Table 2: Summary of designs used for comparison. n_{points} : number of points.

The parameter m is integrated out analytically. All other parameters, $(\lambda(t))_{t \in \mathbb{T}}$, L and the hyper-parameters of k_{ξ_0} and $k_{X\varepsilon}$, are estimated by maximization of the joint posterior density (MAP estimation).

3 Numerical results

We consider a parallelepiped building, with two doors and two windows, simulated with Fire Dynamics Simulator. We study the maximal temperature in the building, $T_t^c(\boldsymbol{x})$, as function of $d = 8$ inputs (external temperature, fire area...).

The model presented in Section 2 is called multi-fidelity non-stationary model and denoted by M-F$_1$. Our objective here is to compare it to two other models. The first model is similar to (4), but uses a stationary anisotropic Matérn covariance function on $\mathbb{X} \times \mathbb{T}$. This model, called multi-fidelity stationary model and denoted by M-F$_2$, is a simplification of the multi-fidelity non-stationary model. The second model is built only from the most accurate level of the simulator. This model, called high-fidelity model and noted H-F., serves as a reference value.

Two datasets have been built for this numerical experiment (see Table 2). The first one is built with 400 simulations, at different mesh sizes. It is used to build both multi-fidelity models (M-F$_1$ and M-F$_2$). The second one consists of 100 simulations at $t = 20$ cm. It is used to build the high-fidelity model H-F., and also for validation.

The results of prediction are presented on Figure 1. On the left, figures show a comparison between predictions (posterior means) and observations. For the high-fidelity model, predictions are made by leave-one-out cross-validation. On the right, the densities of normalized residuals are compared with the probability density function of the normal distribution. Overall, the two multi-fidelity models present a goodness-of-fit similar to that of the high-fidelity model, which is our reference. On closer inspection, it appears that both multi-fidelity models—and most particularly the multi-fidelity stationary model—actually underestimate T^c for high values. However, the standard deviations of the densities of residuals are close to one, suggesting that posterior variances are neither too large, nor too thin.

Finally, we consider the problem of estimating the probability $\mathbb{P}_{\mathbb{X}}(T_{20cm}^c(\boldsymbol{x}) > 60°C)$ that the output temperature exceeds a critical threshold (here, $60°C$), where $\mathbb{P}_{\mathbb{X}}$ is a
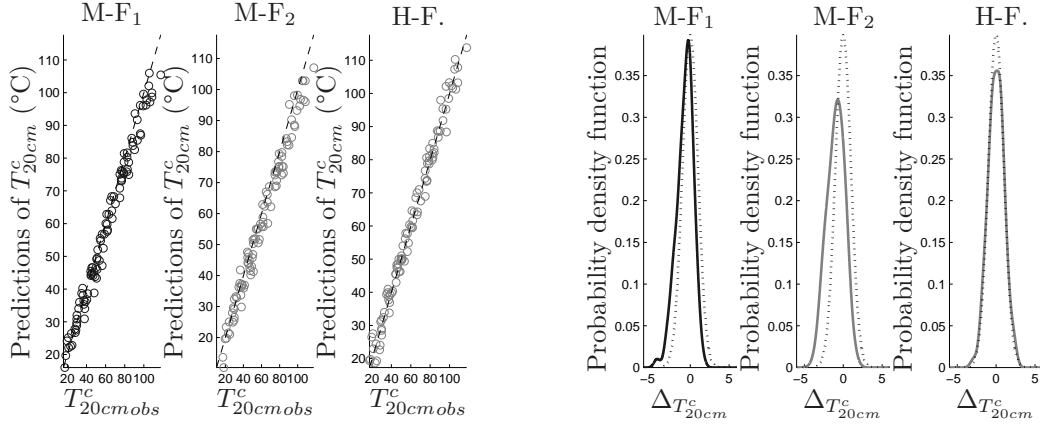
Figure 1: Results on prediction of level 20 cm, by the three models. From left to right, multi-fidelity non-stationary model (M-F$_1$); multi-fidelity stationary model (M-F$_2$); high-fidelity model (H-F.). Left : predictions versus observations; dashed line : first bisector ($y = x$). Right : Estimation of the probability density function of the normalized residuals $\Delta_{T^c_{20cm}}$; dashed line : probability density function of normal distribution.
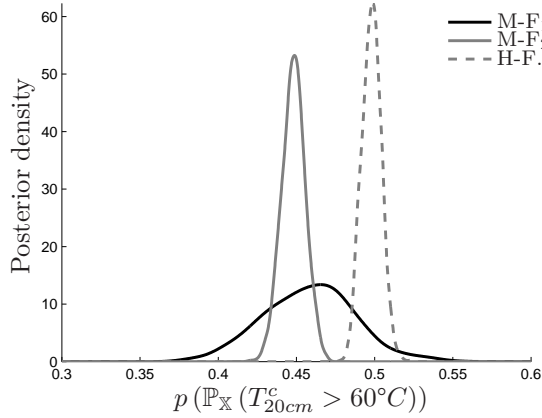

Figure 2: Estimations of the posterior function of $\mathbb{P}_{\mathbb{X}}(T^c_{20cm} > 60°C)$. These densities are estimated with $n_{sim} = 1000$ conditional random simulations on $n_{pts} = 5000$ inputs.

probability distribution of the input space $\mathbb{X}$. Such a probability is useful for fire engineers to assess the safety of the building. Figure 2 presents different estimations of a posterior probability density function of probability of exceeding the threshold. These densities are estimated with $n_{sim} = 1000$ conditional random simulations on $n_{pts} = 5000$ inputs. The high-fidelity model yields the narrowest posterior distribution. The multi-fidelity stationary model also yields a small posterior uncertainty, but the support of the density of the probability of exceeding the threshold does not agree with that of the high-fidelity model. Our model has a larger posterior variance, but is compatible with the reference model, and has been obtained using less computational resources.

4   Conclusion

To conclude, we have proposed an extension of the model of Picheny and Ginsbourger (2013) and Tuo et al. (2014) to the case of stochastic simulators. Our numerical results show that the proposed model makes it possible to predict the behavior of a stochastic multi-fidelity simulator at high fidelity, from simulations at low fidelities. We believe that this is a promising approach, particularly in the domain of fire safety. Future research will concentrate on fully Bayesian estimation of the parameters of the model, and sequential design of experiment in order to achieve a more accurate estimation of the probability of exceeding the threshold, using a limited computational budget for additional simulations.

References

Marc C Kennedy and Anthony O'Hagan. Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1):1–13, 2000.

Loic Le Gratiet. *Multi-fidelity Gaussian process regression for computer experiments*. PhD thesis, Université Paris-Diderot-Paris VII, 2013.

Victor Picheny and David Ginsbourger. A nonstationary space-time gaussian process model for partially converged simulations. *SIAM/ASA Journal on Uncertainty Quantification*, 1(1):57–78, 2013.

Rui Tuo, C.F. Jeff Wu, and Dan Yu. Surrogate modeling of computer experiments with different mesh densities. *Technometrics*, 56(3):372–380, 2014.