



A Non-Nested Infilling Strategy for Multi-Fidelity based Efficient Global Optimization

Matthieu Sacher, Olivier Le Maitre, Régis Duvigneau, Frédéric Hauville,
Mathieu Durand, C. Lothode

► To cite this version:

Matthieu Sacher, Olivier Le Maitre, Régis Duvigneau, Frédéric Hauville, Mathieu Durand, et al.. A Non-Nested Infilling Strategy for Multi-Fidelity based Efficient Global Optimization. International Journal for Uncertainty Quantification, 2021, 11 (1), pp.1-30. 10.1615/Int.J.UncertaintyQuantification.2020032982 . hal-02901774

HAL Id: hal-02901774

<https://inria.hal.science/hal-02901774>

Submitted on 17 Jul 2020

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.

A Non-Nested Infilling Strategy for Multi-Fidelity based Efficient Global Optimization

Matthieu Sacher^{a,*}, O.P. Le Maître^b, R. Duvinneau^c, F. Hauville^a, M. Durand^d, C. Lothodé^e

^aNaval Academy Research Institute - IRENav, BREST, France

^bCMAP, CNRS, Inria, École Polytechnique, Palaiseau, France

^cUniversité Côte d'Azur, Inria, CNRS, LJAD, Sophia-Antipolis, France

^dSirli Innovations, Pornichet, France

^eK-Epsilon, Valbonne, France

Abstract

Efficient Global Optimization (EGO) has become a standard approach for the global optimization of complex systems with high computational costs. EGO uses a training set of objective function values computed at selected input points to construct a statistical surrogate model, with low evaluation cost, on which the optimization procedure is applied. The training set is sequentially enriched, selecting new points, according to a prescribed infilling strategy, in order to converge to the optimum of the original costly model. Multi-fidelity approaches combining evaluations of the quantity of interest at different fidelity levels have been recently introduced to reduce the computational cost of building a global surrogate model. However, the use of multi-fidelity approaches in the context of EGO is still a research topic. In this work, we propose a new effective infilling strategy for multi-fidelity EGO. Our infilling strategy has the particularity of relying on non-nested training sets, a characteristic that comes with several computational benefits. For the enrichment of the multi-fidelity training set, the strategy selects the next input point together with the fidelity level of the objective function evaluation. This characteristic is in contrast with previous nested approaches, which require estimation all lower fidelity levels and are more demanding to update the surrogate. The resulting EGO procedure achieves a significantly reduced computational cost, avoiding computations at useless fidelity levels whenever possible, but it is also more robust to low correlations between levels and noisy estimations. Analytical problems are used to test and illustrate the efficiency of the method. It is finally applied to the optimization of a fully nonlinear fluid-structure interaction system to demonstrate its feasibility on real large-scale problems, with fidelity levels mixing physical approximations in the constitutive models and discretization refinements.

Keywords: Stochastic Preconditioner, Sampling Method, Domain Decomposition, Parallel Computation, Preconditioned Conjugate Gradient Method

1. Introduction

Efficient Global Optimization (EGO) strategies [1] have been used in various application domains, and many works demonstrated their applicability for the optimization of complex systems with costly objective function evaluation [2]. Examples of applications are the aerodynamic drag reduction [3], the vibration reduction for rotating aircraft [4], the optimization of Fluid-Structure Interactions (FSI) problems [5] and the sail trimming optimization [6]. In a nutshell, the EGO method builds a statistical surrogate model of the costly objective function f to be optimized, usually a Gaussian Processes [7] (GP), and a statistical criterion accounting for the Expected Improvement (EI), often referred to as *merit function*, is maximized to select a new design point. The objective function is then evaluated at the new-design point to enrich the training set, the surrogate is updated, and the algorithm proceeds with the selection of a new design point or eventually stops if it has converged.

*Corresponding author. Present affiliation: ENSTA Bretagne, CNRS, IRDL, Brest, France

Email addresses: matthieu.sacher@ensta-bretagne.fr (Matthieu Sacher), olivier.le-maitre@polytechnique.edu (O.P. Le Maître), Regis.Duvigneau@inria.fr (R. Duvinneau), frederic.hauville@ecole-navale.fr (F. Hauville)

When the evaluation of the cost function uses complex computational tools, it is usually possible to define a set of models of different fidelities and computational costs to estimate its value. As an example, the fidelity levels may correspond to finite element models using different mesh refinements, therefore having variable computational cost and accuracy. Thus, the global computational cost of the optimization can be reduced by using such different models in a multi-fidelity (MF) optimization framework. Several methods have been proposed in this perspective; these methods differ by the construction of the MF surrogate and the infilling criterion used to select a new evaluation point at a given fidelity level. The first set of approaches consists of the methods that do not make any assumption regarding a possible hierarchy between the models constituting the MF system. These approaches are especially attractive when an evident cost or accuracy hierarchy between the models is not known *a priori*, or when the fidelity of information sources varies over the design space [8, 9, 10]. This work is motivated by application to complex coupled mechanical systems, whose simulation relies on models of increasing complexity, involving progressively more and more physical phenomena. Further, these models have well-characterized and predictable complexities such that there exists a natural ordering to define a hierarchy of models. Consequently, to extend the EGO paradigm, we consider the second set of approaches consisting of hierarchical methods. The authors of [11] first proposed a MF surrogates combining estimations from a hierarchy of fidelity models. The authors of [12] introduced an EGO formulation in the MF framework. However, this approach accounts for two-fidelity levels only (also called *co-kriging*), and the merit function does not account for the fidelity level as the method uses the prediction of the highest fidelity surrogate to estimate the merit function. The work of [13] investigates the resolution of an optimization problem involving seven design variables and using two Computational Fluid Dynamics (CFD) models having different fidelity. The authors of the paper propose five infilling strategies, all focusing on the global reduction of the MF-surrogates uncertainty and the selection of the CFD model to use for the estimation of the objective function at the new point. Therefore, this approach is not a formal EGO method, as it does not account for the optimization objective. Moreover, infilling strategies based on the global reduction of the surrogate uncertainty are known to result in sub-optimal approaches as they waste resources to improve the surrogate in regions of poor objective value instead of focusing on areas with substantial expected improvement [14].

In [15], the authors have proposed a nested infilling strategy for the optimization based on MF-surrogates. The nested character means that the sequence of training sets at successive fidelity levels are nested; specifically, the set of training points at level l is a subset of the training points at level $l - 1$. In their work, the merit function uses the predictive variance of the highest fidelity surrogate. Then, an approximation of the integrated variance predictor [16] is employed to compute the potential variance reduction of the highest fidelity surrogate, after the enrichment of the training sets up to level l , with the same training point (thus requiring the evaluation of all models with fidelity level $\leq l$). This maximum fidelity level l is determined by balancing the resulting integrated variance reduction with the computational cost of the l model evaluation. To sum up, the method in [15] has the advantage of incorporating a fidelity level selection in the infilling strategy, but it has several drawbacks. First, this strategy is constrained to nested training sets, and so it may not be optimal. Besides, the method tends to favor low fidelity levels (with low computational cost) in addition to use a merit function only based on the variance that is known to be of poor efficiency for optimization problem applications.

In [17], a non-nested MF approach was proposed, based on a merit function defined to mitigate the correlation between a given level and the highest fidelity one with the relative amount of noise and the evaluation cost at that level. However, it does not account for the reduction of the predictive variance at the highest fidelity level.

Recently, the problem of applying EGO to airfoil shape optimization in a multi-fidelity framework was considered in [18]. Nested datasets are used to build MF-surrogates, and the Expected Improvement merit function is computed on the highest fidelity level. This work also proposed an infilling criterion for the selection of the fidelity level. The criterion corresponds to the ratio of the variance reduction to the associated squared cost of the highest fidelity level to evaluate. From our perspective, this strategy presents several drawbacks. Indeed, the use of nested datasets may not be optimal. Besides, the merit function and the fidelity level selection are decoupled. Finally, the ratio of variance to cost is not a dimensionless number, and that may raise difficulties if the costs are not well known.

Recent works have proposed several MF extensions. These extensions are not directly related to the present work but we briefly review them for the sake of completeness. In [19, 20], the MF model is extended to the prediction of vector-valued fields, using a multivariate recursive MF approach. This methodology has been successfully applied to uncertainty propagation and quantification problems, but it could be adapted to more general parametric studies and MF optimization. Other works have considered MF algorithms in the context of bandit optimization. The work in [21] presents a MF optimization framework for a continuous set of variable fidelity models. In [22], the authors propose

an alternative optimization strategy using a low-fidelity model to cheaply eliminate regions of poor performance and a high-fidelity model in the promising regions only.

In the present work, we propose to extend the original EI-based merit function of the EGO method to the case of objective functions estimated using multiple fidelity models at non-nested sets of training points. We propose a novel merit function that selects the next design point and the fidelity level of the model to be evaluated (enrichment level). This infilling method accounts for the computational cost and the potential reduction of the predictive variance of the highest fidelity surrogate, within the region of interest (susceptible to contain the optimum). Several advantages arise from this new approach. The main one concerns a potential reduction of the computational cost compared to the single-fidelity EGO (*i.e.*, all evaluations based on the highest fidelity model) and MF methods based on nested training sets. Using non-nested training sets presents the positive aspect of naturally selecting the optimal fidelity level, regarding cost/variance-reduction trade-off, while avoiding the computations of useless lower fidelity models. Besides, our numerical tests also demonstrate that our approach is more robust than the nested approach when the correlation between the successive levels is weak or in the presence of evaluation noise.

The paper structure is as follows. Section 2 concerns the optimization methods; we review the constituents of the classical single-fidelity EGO (SF-EGO) in Section 2.1. We detail the construction of MF surrogates in Section 2.2 and our infilling strategy for non-nested training sets in Section 3. Section 4 then presents several analytical test problems to compare the efficiency of the single-fidelity EGO with the nested and non-nested MF-EGO approaches. The optimization of a complex nonlinear FSI model is subsequently considered in Section 5, with the maximization of the aerodynamic performance of a Figaro racing yacht through the optimization of the two sail trimming variables. The multi-fidelity framework involves five numerical models for the computation of the objective function, which differ in their discretization levels and physical models simplifications of both the structural and fluid solvers involved in the FSI problem. Finally, the major conclusions of this work are proposed in Section 6.

2. GP-based optimization

Our objective is to determine the optimum design variable \mathbf{x}_{opt} , the solution of the following generic optimization problem

$$\mathbf{x}_{\text{opt}} = \arg \min_{\mathbf{x} \in \Omega} f(\mathbf{x}), \quad (1)$$

where \mathbf{x} is the vector of design variables, $\Omega \subset \mathbb{R}^d$ is the admissible optimization domain and $f : \Omega \mapsto \mathbb{R}$ is the objective function. The global optimum of f can be very costly to find when, in particular, individual evaluations of $f(\mathbf{x})$ are numerically expensive. A classical approach is then to use a surrogate model in place of f to reduce the computational burden related to the evaluations of f [2]. In this work, we consider the use of Gaussian processes (GP) [7], which, owing to their statistical nature, provide both a prediction of the objective function value and a measure of the uncertainty (variance) in this prediction. These features are appealing in the optimization context, as they can be exploited to derive rigorous optimization strategies based on the maximization of an Expected Improvement (EI) criterion, and leading to approaches globally referred to as Efficient Global Optimization (EGO) [1] methods.

We summarize in Section 2.1 the construction of single-fidelity GP surrogate of an objective function f , and the resulting SF-EGO strategy. Section 2.2 introduces the extension of the GP surrogate construction to the multi-fidelity case, while our new non-nested infilling strategy for MF-EGO is detailed in Section 3.

2.1. Single-fidelity EGO

2.1.1. SF-surrogate construction

Consider a set of n training points $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, each in Ω . The training points are associated to the vector $\mathbf{Y} = (y_1, \dots, y_n)$ of noisy observations of the objective function f . It is assumed that $y_i = f(\mathbf{x}_i) + \varepsilon_i$, where the ε_i are independent and identically distributed Gaussian random variables having zero-mean and variance σ_ϵ^2 , *i.e.*, $\epsilon_i \sim \mathcal{N}(0, \sigma_\epsilon^2)$.

The GP construction considers that $f(\mathbf{x})$ is a realization of a zero-mean multivariate Gaussian process with covariance function C . Without loss of generality, we consider here the multidimensional squared exponential covariance

functions defined by

$$C(\mathbf{x}, \mathbf{x}'; \Theta) \doteq \theta_1 \prod_{i=1}^d \exp\left(\frac{-(x_i - x'_i)^2}{2l_i^2}\right) + \theta_2, \quad (2)$$

where $\Theta = \{\theta_1, \theta_2, l_1, l_2, \dots, l_d\}$ is the vector of covariance hyper-parameters to be inferred from the observations. Denoting $\mathbf{C} \in \mathbb{R}^{n \times n}$ the symmetric covariance matrix of the training points, with entries $C_{i,j} \doteq C(\mathbf{x}_i, \mathbf{x}_j; \Theta)$, for $1 \leq i, j \leq n$, the joint Gaussian distribution of the noisy observations \mathbf{Y} and prediction of a new observation $Y(\mathbf{x})$ at $\mathbf{x} \notin \mathcal{X}$ is given by

$$\begin{pmatrix} \mathbf{Y} \\ Y(\mathbf{x}) \end{pmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{C} + \sigma_\epsilon^2 \mathbf{I}_n & \mathbf{k}(\mathbf{x}) \\ \mathbf{k}^T(\mathbf{x}) & \kappa(\mathbf{x}) + \sigma_\epsilon^2 \end{bmatrix}\right). \quad (3)$$

In (3) we have denoted $\kappa(\mathbf{x}) \doteq C(\mathbf{x}, \mathbf{x}; \Theta)$, $\mathbf{k}(\mathbf{x}) \doteq (C(\mathbf{x}, \mathbf{x}_1; \Theta) \dots C(\mathbf{x}, \mathbf{x}_n; \Theta))^T$ and \mathbf{I}_n the identity matrix of \mathbb{R}^n . From the conditional rules of joint Gaussian distributions [23], the *best* prediction $\hat{f}(\mathbf{x})$ of $f(\mathbf{x})$, i.e. the mean of $Y(\mathbf{x})$, and the prediction variance $\hat{\sigma}^2(\mathbf{x})$ are given by

$$\hat{f}(\mathbf{x}) = \mathbf{k}^T(\mathbf{x}) (\mathbf{C} + \sigma_\epsilon^2 \mathbf{I}_n)^{-1} \mathbf{Y}, \quad (4)$$

$$\hat{\sigma}^2(\mathbf{x}) = \kappa(\mathbf{x}) + \sigma_\epsilon^2 - \mathbf{k}^T(\mathbf{x}) (\mathbf{C} + \sigma_\epsilon^2 \mathbf{I}_n)^{-1} \mathbf{k}(\mathbf{x}). \quad (5)$$

In the following, we shall denote $\mathbf{K} \doteq (\mathbf{C} + \sigma_\epsilon^2 \mathbf{I}_n)$ the matrix of the GP model. This matrix depends only on the set \mathcal{X} of training points and the hyper-parameters of the covariance. The inverse of \mathbf{K} appears in the expressions of the prediction expectation and variance. As predictions may have to be evaluated at multiple points \mathbf{x} , the inverse of \mathbf{K} or its factorization is usually computed and stored for subsequent use. The hyper-parameters Θ of the GP covariance and the noise variance σ_ϵ^2 can be determined by maximizing the log-marginal likelihood of the GP model [23]; in the present work, we used an evolution strategy algorithm [24] to determine the optimal hyperparameters. More details on GP models can be found in [23].

2.1.2. SF-infilling strategy

Let $\hat{\mathbf{x}}_{\text{opt}}$ be the optimum of $\hat{f}(\mathbf{x})$. It is expected that $\hat{\mathbf{x}}_{\text{opt}} \approx \mathbf{x}_{\text{opt}}$ if the approximation error $\hat{f} - f$ is small enough, in particular in the neighborhood of \mathbf{x}_{opt} . The advantage of minimizing \hat{f} instead of f is that GP models are usually inexpensive to evaluate compared to the original objective function. In practice, a sequential enrichment of the training set, with new training points $\mathbf{x}_i \in \Omega$ selected in areas of interest where the optimum is likely to be found, reduces the approximation error. A deterministic optimization procedure would choose the next training point \mathbf{x}_{n+1} as the optimal point of \hat{f} . However, the GP model provides probabilistic information that can be exploited to propose more robust strategies based on merit functions combining the prediction $\hat{f}(\mathbf{x})$ and its variance $\hat{\sigma}^2(\mathbf{x})$. In this work, we use the merit functions based on the Augmented Expected Improvement (AEI) [25]. The AEI estimates the expected progress in the objective value at a new point, given the prediction variance $\hat{\sigma}^2$ and observations noise in \mathbf{Y} . The AEI is written as

$$\text{AEI}(\mathbf{x}) = \text{EI}(\mathbf{x}) \left(1 - \frac{\sigma_\epsilon}{\sqrt{\hat{\sigma}^2(\mathbf{x}) + \sigma_\epsilon^2}}\right), \quad (6)$$

where the Expected Improvement $\text{EI}(\mathbf{x})$ is defined by

$$\text{EI}(\mathbf{x}) \doteq \hat{\sigma}(\mathbf{x}) [u(\mathbf{x}) \Phi(u(\mathbf{x})) + \phi(u(\mathbf{x}))], \quad u(\mathbf{x}) \doteq \frac{\hat{f}(\mathbf{x}_{\text{best}}) - \hat{f}(\mathbf{x})}{\hat{\sigma}(\mathbf{x})}, \quad (7)$$

with Φ and ϕ are the cumulative and density functions of the standard Gaussian distribution, and $\mathbf{x}_{\text{best}} \in \mathcal{X}$ is the current effective best solution (see [25]). The optimum point of the AEI is added to the training set \mathcal{X} , and f is evaluated at this new point to complement the observation vector \mathbf{Y} . A new iteration can then start updating the GP

model with the new training point. We remark that when the observation noise is zero, $\sigma_\epsilon^2 \rightarrow 0$, the AEI reduces to the EI while, for finite observation noise, the multiplicative factor tends to penalize areas of Ω where the prediction variance is comparatively small ($\hat{\sigma}^2(\mathbf{x}) \ll \sigma_\epsilon^2$). Overall, each iteration requires one computation of the objective function f and the resolution of two optimization problems: a first one for the hyper-parameters of the GP model, and a second one to find the AEI optimum. These iterations of the GP-based optimization continue until a stopping criterion is satisfied or the resources allocated to the optimization problem have been exhausted.

2.2. Multi-fidelity surrogates

We consider the availability of a sequence of L models $f^{(1)}, \dots, f^{(L)}$ of increasing fidelity for the evaluation of the objective function, such that

$$f^{(l)}(\mathbf{x}) \xrightarrow{l \rightarrow L} f(\mathbf{x}), \quad \forall \mathbf{x} \in \Omega.$$

Typically, the computational cost associated to the evaluation of $f^{(l)}(\mathbf{x})$ increases with the fidelity level l . The objective of the MF-surrogates is to incorporate observations corresponding to fidelity levels $l < L$ (thus associated with a lower computational cost) to improve the predictive capability of the surrogate model of $f^{(L)}(\mathbf{x})$.

In this work, we use the recursive formulation proposed in [26], which reduces the complexity of the original MF approach proposed in [11]. The recursive formulation breaks down the determination of the MF-surrogate in successive GP process constructions for each intermediate level. The construction at a level involves the observations of that level only, contrary to the direct approach of [11], where all observations at all levels are treated at once, leading to the inversions of much larger covariance matrices. The recursive approach also drastically eases the determination of the hyper-parameters associated with the GP models, by breaking down the global optimization in a sub-problem per level. Also, the extension to non-nested sets of training points of the recursive approach is immediate and has the advantage of providing explicit expressions for the predictive variances of the MF-surrogates at every intermediate level. We exploit the latter advantage in the following Section 3 to derive a novel non-nested infilling strategy.

We first consider the case of nested training sets. For $l = 1, \dots, L$, let us denote $\mathcal{X}^{(l)}$ the set of training points for model $f^{(l)}$, such that

$$\mathcal{X}^{(l+1)} \subseteq \mathcal{X}^{(l)}, \quad l = 1, \dots, L-1. \quad (8)$$

As a result, the number of observations of a model decreases with its fidelity level. We shall denote $n_{(l)} = |\mathcal{X}^{(l)}|$ the size of the training set. Let $\mathbf{Y}^{(l)}$ be the associated vector of noisy observations of $f^{(l)}$ at the corresponding training points in $\mathcal{X}^{(l)}$, and $\sigma_{\epsilon(l)}^2$ the variance of the observation noise on $f^{(l)}$.

2.2.1. MF-recursive construction

The key idea of the recursive construction [26] is to define the MF-surrogate at level l , denoted $Y^{(l)}$, as a correction of the MF-surrogate at level $l-1$. Specifically, we write for $l = 1, \dots, L$

$$Y^{(l)}(\mathbf{x}) \doteq \rho_{(l-1)} Y^{(l-1)}(\mathbf{x}) + \delta Y^{(l)}(\mathbf{x}), \quad Y^{(0)}(\mathbf{x}) \doteq 0, \quad (9)$$

where $\rho_{(l-1)} \in \mathbb{R}$ accounts for the correlation between the two successive levels, and $\delta Y^{(l)}(\mathbf{x})$ is the Gaussian correction process orthogonal to $Y^{(l-1)}(\mathbf{x})$. This particular form ensures a Markovian property between levels [11]. In particular, it can be shown that the covariance between $Y^{(l)}(\mathbf{x})$ and $Y^{(l-1)}(\mathbf{x}')$, conditioned on $Y^{(l-1)}(\mathbf{x})$, vanishes for all $\mathbf{x}, \mathbf{x}' \in \Omega^2$.

The MF-surrogate $Y^{(l)}$ is then a GP that depends on the previous level MF-surrogate and on the noisy observations of $f^{(l)}(\mathbf{x})$, that is

$$\left[Y^{(l)}(\mathbf{x}) | Y^{(l-1)}(\mathbf{x}), \mathbf{Y}^{(l)} \right] \sim \mathcal{N} \left(\hat{f}^{(l)}(\mathbf{x}), \hat{\sigma}_{(l)}^2(\mathbf{x}) \right), \quad (10)$$

with mean and variance respectively given by

$$\hat{f}^{(l)}(\mathbf{x}) = \rho_{(l-1)} \hat{f}^{(l-1)}(\mathbf{x}) + \mathbf{k}_{(l)}^T(\mathbf{x}) \left(\mathbf{C}_{(l)} + \sigma_{\epsilon(l)}^2 \mathbf{I}_{n_{(l)}} \right)^{-1} \Delta \mathbf{Y}^{(l)}(\rho_{(l-1)}), \quad (11)$$

$$\hat{\sigma}_{(l)}^2(\mathbf{x}) = \rho_{(l-1)}^2 \hat{\sigma}_{(l-1)}^2(\mathbf{x}) + \kappa_{(l)}(\mathbf{x}) + \sigma_{\epsilon(l)}^2 - \mathbf{k}_{(l)}^T(\mathbf{x}) \left(\mathbf{C}_{(l)} + \sigma_{\epsilon(l)}^2 \mathbf{I}_{n_{(l)}} \right)^{-1} \mathbf{k}_{(l)}(\mathbf{x}). \quad (12)$$

In the previous expressions, we have denoted $\Delta \mathbf{Y}^{(l)}$ the vector of noisy observation residuals at training points in $\mathcal{X}^{(l)}$ and defined as

$$\Delta \mathbf{Y}^{(l)}(\rho_{(l-1)}) = \begin{cases} \mathbf{Y}^{(l)}, & l = 1, \\ \mathbf{Y}^{(l)} - \rho_{(l-1)} \mathbf{Y}_{\mapsto(l)}^{(l-1)}, & l \geq 1, \end{cases} \quad (13)$$

where $\mathbf{Y}_{\mapsto(l)}^{(l-1)}$ is the subpart vector of $l - 1$ level associated with noisy observations in $\mathbf{Y}^{(l-1)}$, corresponding to the elements associated with training points also belonging to $\mathcal{X}^{(l)}$. Therefore, $\Delta \mathbf{Y}^{(l)}$ is the vector of noisy observation residuals (relative to the previous level observations) at the training points in $\mathcal{X}^{(l)}$. We also reused in (11) and (12) the notations of the SF case, adding subscripts (l) to stress the dependencies on the considered level of $\mathbf{k}_{(l)}(\mathbf{x})$, $\mathbf{C}_{(l)}$, $\kappa_{(l)}(\mathbf{x})$ and $\sigma_{\epsilon_{(l)}}$. These dependencies arise from the variable size of the training sets, which decreases with the level l , but also because of the hyper-parameters $\Theta_{(l)}$, which are allowed to change from a level to another, as discussed below.

Observing the expression of the mean in (11), one recognizes the contribution from the mean prediction at the previous fidelity level scaled by the correlation coefficient $\rho_{(l-1)}$, while the second contribution is due to the observation residual $\Delta \mathbf{Y}^{(l)}$. Similarly, the expression of the variance of the MF-surrogate prediction, in (12), makes evident the contribution of the variance at the previous level (scaled by $\rho_{(l-1)}^2$) which, consistently with the Markovian assumption, directly adds up with the predictive variance of the correction $\delta Y^{(l)}(\mathbf{x})$. In fact, an expression of the two-point covariance of the MF-surrogate at level l can be derived in terms of the lower level covariances. It comes [11]

$$\text{cov} \left(Y^{(l)}(\mathbf{x}), Y^{(l)}(\mathbf{x}') \right) = \sum_{j=1}^l \left(\prod_{i=j}^{l-1} \rho_{(i)}^2 \right) C(\mathbf{x}, \mathbf{x}'; \Theta_{(j)}). \quad (14)$$

As mentioned before, one advantage of the recursive approach is that the hyper-parameters $\Theta_{(l)}$, noise variances $\sigma_{\epsilon_{(l)}}^2$ and regression coefficients $\rho_{(l-1)}$, must not be determined all together for all fidelity levels. Instead, owing to the Markovian structure of the construction, they are sequentially determined, from $l = 1$ to $l = L$. As in the SF case, the hyper-parameters of level l are selected by maximizing the log-marginal likelihood of $Y^{(l)}$ [23],

$$\mathcal{L}(\rho_{(l-1)}, \Theta_{(l)}, \sigma_{\epsilon_{(l)}}^2) = -\frac{n_{(l)}}{2} \log(2\pi) - \frac{1}{2} \log \left| \mathbf{C}_{(l)}(\Theta_{(l)}) + \sigma_{\epsilon_{(l)}}^2 \mathbf{I}_{n_{(l)}} \right| \quad (15)$$

$$- \frac{1}{2} (\Delta \mathbf{Y}^{(l)}(\rho_{(l-1)}))^T \left(\mathbf{C}_{(l)}(\Theta_{(l)}) + \sigma_{\epsilon_{(l)}}^2 \mathbf{I}_{n_{(l)}} \right)^{-1} \Delta \mathbf{Y}^{(l)}(\rho_{(l-1)}). \quad (16)$$

2.2.2. Elementary example

As an elementary illustrative example, taken from [12], consider the case of two fidelity-levels with functions $f^{(1)}$ and $f^{(2)}$ defined as

$$f^{(1)}(x) = \frac{1}{2}(6x - 2)^2 \sin(12x - 4) + 10(x - 1), \quad f^{(2)}(x) = 2f^{(1)}(x) - 20(x - 1). \quad (17)$$

The two models are reported with the solid blue ($f^{(1)}$) and red ($f^{(2)}$) curves shown in Fig. 1. The low fidelity model $f^{(1)}$ is evaluated at 11 points (blue circles), while the high fidelity model $f^{(2)}$ is only evaluated at a subset of 4 points (red squares). In Fig. 1, the dashed green curve is the SF-surrogate based only on the 4 observations of $f^{(2)}$. The SF-surrogate is seen to badly approximate $f^{(2)}$: the location and value of the SF-surrogate minimum are far from those of $f^{(2)}$. In contrast, the MF-surrogate constructed as described below, and shown with a dashed black line, provides a much better estimation of $f^{(2)}$ with improved minimum value and location.

3. Non-Nested MF-infilling strategy

3.1. Non-Nested MF-surrogate construction

As discussed in the introduction, we prefer to consider non-nested sets of points for the MF-EGO, in order to reduce the computational cost and time associated with the resolution of lower fidelity models.




Figure 1: One-variable multi-fidelity example. The MF-surrogate uses both low ($\mathbf{Y}^{(1)}$) and high ($\mathbf{Y}^{(2)}$) fidelity observations.

The first step is then to extend the previous MF-surrogate construction to non-nested training sets. This turns to be quite natural within the recursive formulation of the MF-surrogate, simply substituting in the definition of the observation residual (13) the previous level observations $\mathbf{Y}_{\rightarrow(l)}^{(l-1)}$ with the corresponding vector of surrogate predictions $\hat{f}_{(l-1)}(\mathbf{x} \in \mathcal{X}^{(l)})$. This change affects the (16) and (11), for which we now have

$$\Delta \mathbf{Y}^{(l)}(\rho_{(l-1)}) = \begin{cases} \mathbf{Y}^{(l)}, & l = 1, \\ \mathbf{Y}^{(l)} - \rho_{(l-1)} \hat{f}_{(l-1)}(\mathbf{x} \in \mathcal{X}^{(l)}), & l \geq 1. \end{cases} \quad (18)$$

3.2. Non-Nested MF merit function

We now present our infilling method for MF-EGO that uses non-nested multi-fidelity surrogates. We start from the Augmented Expected Improvement (AEI) [25] merit function discussed in the SF-EGO section. We recall that this merit function, in (6), consists of the product of two terms: the expected improvement EI and a penalization term favoring areas with substantial prediction uncertainty.

A straight-forward MF-extension of (6) consists in defining the EI on the highest fidelity surrogate, *i.e.* using $\hat{f}^{(L)}(\mathbf{x})$, to obtain

$$\text{EI}(\mathbf{x}) = \hat{\sigma}_{(L)}(\mathbf{x}) [u(\mathbf{x})\Phi(u(\mathbf{x})) + \phi(u(\mathbf{x}))], \quad u(\mathbf{x}) = \frac{\hat{f}^{(L)}(\mathbf{x}_{\text{best}}) - \hat{f}^{(L)}(\mathbf{x})}{\hat{\sigma}_{(L)}(\mathbf{x})}, \quad (19)$$

with, as before, Φ and ϕ the cumulative and density functions of the standard Gaussian distribution, and where $\mathbf{x}_{\text{best}} \in \cup_l \mathcal{X}^{(l)}$ the current effective best solution of all the sets of training points. Concerning the augmentation factor, one can compare the variance of the MF-prediction variance with the inferred noise in the observation residual at the highest-fidelity level ($\Delta Y^{(L)}$). This would lead to a first expression of merit function:

$$\text{AEI}_{(L)}(\mathbf{x}) = \text{EI}(\mathbf{x}) \left(1 - \frac{\sigma_{\epsilon(L)}}{\sqrt{\hat{\sigma}_{(L)}^2 + \sigma_{\epsilon(L)}^2}} \right). \quad (20)$$

Contrary to the SF case, we have to select a new point \mathbf{x}_* and the fidelity level l_* of the model to be evaluated. The SF case would correspond to selecting systematically the highest level (L).

A possible point of view for selecting the fidelity level, whence the new point \mathbf{x}_* is decided, consists in answering the question: what is the effect of updating the MF-surrogate with a new observation of $f^{(l)}$ at \mathbf{x}_* , with $l < L$, compared to an evaluation of the highest fidelity model $f^{(L)}$? To address this question, we have to specify the measure of this effect. In the present work, we propose to measure the impact of the new observation \mathbf{x}_* , and its associated fidelity level l_* , through the induced reduction of the MF-surrogate predictive variance. To this end, we denote $\hat{\sigma}_{(L)}^2(\mathbf{x}|\mathbf{x}_*, l_*)$ the predictive variance at the highest resolution level, given a new observation $f^{(l_*)}(\mathbf{x}_*)$. The interest of using the predictive variance as a criterion for the fidelity level selection is that the evaluation of $\hat{\sigma}_{(L)}^2(\mathbf{x}|\mathbf{x}_*, l_*)$ does not need to evaluate $f^{(l_*)}(\mathbf{x}_*)$, provided that the hyper-parameters and regression parameters of the MF-GPs are not too affected by the new observation. In this case, we can exploit the recursive expression of the predictive variance in (12). In particular, the predictive variances $\hat{\sigma}_{(l)}^2$ of models $f^{(l)}$ with $l < l_*$ are not affected by a new observation of $f^{(l_*)}$.

We first express the change in the predictive variance of $f^{(l_*)}$ induced by the new observation. From (12) we see that $\hat{\sigma}_{(l_*)}^2(\mathbf{x})$ has 3 contributions: a first contribution, $\rho_{(l_*-1)}\hat{\sigma}_{(l_*-1)}^2(\mathbf{x})$, from the predictive variance at the previous level; a second contribution, $\kappa_{(l_*)}(\mathbf{x}) + \sigma_{\epsilon(l_*)}^2$, independent of the observation points in $\mathcal{X}^{(l_*)}$; and finally a third contribution $-\mathbf{k}_{(l)}^T(\mathbf{x}) \left(\mathbf{C}_{(l)} + \sigma_{\epsilon(l)}^2 \mathbf{I}_{n(l)} \right)^{-1} \mathbf{k}_{(l)}(\mathbf{x})$ which depends on the observation points in $\mathcal{X}^{(l_*)}$. Only the last contribution is affected by adding a new observation, and the updated prediction variance of $f^{(l_*)}(\mathbf{x})$ can be expressed as

$$\hat{\sigma}_{(l_*)}^2(\mathbf{x}|\mathbf{x}_*, l_*) = \hat{\sigma}_{(l_*)}^2(\mathbf{x}) + \Delta\hat{\sigma}_{(l_*)}^2(\mathbf{x}|\mathbf{x}_*, l_*), \quad (21)$$

where $\Delta\hat{\sigma}_{(l_*)}^2(\mathbf{x}|\mathbf{x}_*, l_*)$ is the predictive variance update. Reusing the previous notations the predictive variance update can be written as

$$\Delta\hat{\sigma}_{(l_*)}^2(\mathbf{x}|\mathbf{x}_*, l_*) = [\mathbf{k}_{(l)}^T(\mathbf{x}) \quad C_{(l)}(\mathbf{x}, \mathbf{x}_*)] \begin{bmatrix} \mathbf{K}_{(l)} & \mathbf{k}_{(l)}(\mathbf{x}_*) \\ \mathbf{k}_{(l)}^T(\mathbf{x}_*) & \kappa_{(l)}(\mathbf{x}_*) + \sigma_{\epsilon(l)}^2 \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{k}_{(l)}(\mathbf{x}) \\ C_{(l)}(\mathbf{x}, \mathbf{x}_*) \end{bmatrix} - \mathbf{k}_{(l)}^T(\mathbf{x}) \mathbf{K}_{(l)}^{-1} \mathbf{k}_{(l)}(\mathbf{x}), \quad (22)$$

where, consistently with the notations of Section 2.1.1, we have denoted the GP model matrix

$$\mathbf{K}_{(l)} = \left[\mathbf{C}_{(l)} + \sigma_{\epsilon(l)}^2 \mathbf{I}_{n(l)} \right].$$

Since for fixed hyper-parameters a new observation can only reduce the predictive variance of the GP, it holds for all \mathbf{x} and \mathbf{x}_* :

$$0 \leq -\Delta\hat{\sigma}_{(l_*)}^2(\mathbf{x}|\mathbf{x}_*, l_*) \leq \hat{\sigma}_{(l_*)}^2(\mathbf{x}).$$

Returning to the recursive expression for the predictive variance, we find that the updated MF predictive variance of $f^{(L)}(\mathbf{x})$, for the new observation $f^{(l_*)}(\mathbf{x}_*)$, is

$$\hat{\sigma}_{(L)}^2(\mathbf{x}|\mathbf{x}_*, l_*) = \hat{\sigma}_{(L)}^2(\mathbf{x}) + R_{l_*}^2 \Delta\hat{\sigma}_{(l_*)}^2(\mathbf{x}|\mathbf{x}_*, l_*), \quad R_{l_*}^2 = \begin{cases} \prod_{l=l_*}^{L-1} \rho_{(l)}^2, & l_* < L, \\ 1, & l_* = L. \end{cases}$$

To introduce a criteria on the fidelity level in the merit function, we can augment (20) to favor couples (\mathbf{x}_*, l_*) achieving significant reduction of the predictive variance. The reduction can be measured globally, integrating $R_{l_*}^2 \Delta\hat{\sigma}_{(l_*)}^2(\mathbf{x}|\mathbf{x}_*, l_*)$ for \mathbf{x} in the whole design space Ω , for instance using a Monte-Carlo quadrature or an adaptive sampling approach (see [9] for more details). To avoid this computational overhead, we restrict ourselves to a simpler reduction criteria and compare the prior and updated predictive variances locally at \mathbf{x}_* , where the reduction is the largest. Specifically, we consider the variance ratio $\hat{\sigma}_{(L)}^2(\mathbf{x}|\mathbf{x}_*, l)/\hat{\sigma}_{(L)}^2(\mathbf{x})$ to augment $\text{AEI}_{(L)}$ into the MF-merit function

$$\text{MF}^{\text{MF}}(\mathbf{x}, l) \propto \text{AEI}_{(L)}(\mathbf{x}) \max \left(0, \left(1 - \frac{\hat{\sigma}_{(L)}^2(\mathbf{x}|\mathbf{x}_*, l)}{\hat{\sigma}_{(L)}^2(\mathbf{x})} \right) \right) = \text{AEI}_{(L)}(\mathbf{x}) \max \left(0, \left(-R_l^2 \frac{\Delta\hat{\sigma}_{(l)}^2(\mathbf{x}|\mathbf{x}_*, l)}{\hat{\sigma}_{(L)}^2(\mathbf{x})} \right) \right). \quad (23)$$

Note that, in principle, $0 \leq \hat{\sigma}_{(L)}^2(\mathbf{x}|\mathbf{x}, l) \leq \hat{\sigma}_{(L)}^2(\mathbf{x})$ such that the max operator is theoretically not necessary. In practice, however, we found it necessary to ensure non negative merit value. In the limit $\sigma_\epsilon \rightarrow 0$ we have $\Delta\hat{\sigma}_{(l)}^2(\mathbf{x}|\mathbf{x}, l) \rightarrow -\hat{\sigma}_{(l)}^2(\mathbf{x})$, since the GP processes become interpolatory. Thus, in this limit, the augmented term favors models whose variances yield the largest predictive variance through the recursive rule. Also, models of fidelity level $l < L$ having large variance are only favored if they correlate well with $f^{(L)}$, i.e., if R_l^2 is significant.

The merit function above will generally favor an observation at the highest fidelity level L , for which $R_L^2 = 1$, as it is usually the most effective way to reduce the predictive variance. The interest in using lower-fidelity models emerges when the merit function accounts for the computational costs. To this end, let us denote $W_{(l)} \in \mathbb{R}_+$ a representative computational cost for a single evaluation of $f^{(l)}(\mathbf{x})$. We shall assume that the actual computational cost is not departing much from $W_{(l)}$ for all $\mathbf{x} \in \Omega$. These costs can be used to weight the expected reduction of the predictive variance, finally yielding the proposed (Non-Nested) MF-merit function

$$M^{\text{NN-MF}}(\mathbf{x}, l) \doteq \text{AEI}_{(L)}(\mathbf{x}) \frac{W_{(L)}}{W_{(l)}} \max\left(0, \left(1 - \frac{\hat{\sigma}_{(L)}^2(\mathbf{x}|\mathbf{x}, l)}{\hat{\sigma}_{(L)}^2(\mathbf{x})}\right)\right). \quad (24)$$

3.3. Nested case

We extend the merit function proposed above to the case of nested sets of training points. In this situation, for a new point \mathbf{x}_* with fidelity level l_* , the corresponding reduction of the predictive variance becomes:

$$\Delta_N \hat{\sigma}_{(L)}^2(\mathbf{x}|\mathbf{x}_*, l_*) = \sum_{l=1}^{l_*} R_l^2 \Delta \hat{\sigma}_{(l)}^2(\mathbf{x}|\mathbf{x}_*, l).$$

In words, the reduction for the nested approach is given by the sum of all reductions induced by levels $l \leq l_*$. Regarding the cost of the nested case, the new evaluations total cost is $\sum_{l \leq l_*} W_{(l)}$, which have to be compared to the cost of selecting systematically the highest fidelity level $l_* = L$, namely $\sum_{l \leq L} W_{(l)}$. The nested MF-merit function thus expresses as

$$M^{\text{N-MF}}(\mathbf{x}, l) \doteq \text{AEI}_{(L)}(\mathbf{x}) \frac{\sum_{l' \leq L} W_{(l')}}{\sum_{l' \leq l} W_{(l')}} \max\left(0, \left(-\frac{\sum_{l'=1}^l R_{l'}^2 \Delta \hat{\sigma}_{(l')}^2(\mathbf{x}|\mathbf{x}, l')}{\hat{\sigma}_{(L)}^2(\mathbf{x})}\right)\right). \quad (25)$$

An alternative expression for the Nested MF-EGO merit function is

$$M^{\text{N-MF}}(\mathbf{x}, l) = \text{AEI}_{(L)}(\mathbf{x}) \frac{\sum_{l' \leq L} W_{(l')}}{\sum_{l' \leq l} W_{(l')}} \max\left(0, \left(1 - \frac{\hat{\sigma}_{(L)}^2(\mathbf{x}|\mathbf{x}, l)}{\hat{\sigma}_{(L)}^2(\mathbf{x})}\right)\right), \quad (26)$$

where the updated variance $\hat{\sigma}_{(L)}^2(\mathbf{x}|\mathbf{x}, l)$ is given in the nested case by

$$\hat{\sigma}_{(L)}^2(\mathbf{x}|\mathbf{x}, l) = \hat{\sigma}_{(L)}^2(\mathbf{x}) + \sum_{l' \leq l} R_{l'}^2 \Delta \hat{\sigma}_{(l')}^2(\mathbf{x}|\mathbf{x}, l'). \quad (27)$$

3.4. Computational details

We see that the MF merit functions call for the resolution of a mixed-type non-linear programming problem for finding the optimum of the functional in (24) and (25). In the present work, where L remains small, it is possible to go through over all possible l and optimize the corresponding new point $\mathbf{x}_*(l)$ to select the best couple. The non-nested approach presents the advantage of making feasible the parallel optimizations of $\mathbf{x}_*(l)$. A parallel search of $\mathbf{x}_*(l)$ is also feasible in the nested approach, but it presents the inconvenience to go through all levels $l' \leq l$ to compute $M^{\text{N-MF}}(\mathbf{x}, l)$. As a result, the determination of the new point and maximal fidelity level is computationally more demanding in the nested case. The computational complexity in the nested case is further impacted by the number of models to evaluate, for each new point, and the update of all the MF-surrogates. Algorithm 1 sketches the two MF-EGO methods, highlighting their main similarities and differences in the updating of the MF-surrogates after the enrichment of the training set(s). Finally, these MF-EGO iterations continue until a stopping criterion is satisfied.

Algorithm 1: Multi-fidelity EGO algorithm.

Input: Initial training sets and observation vectors $\{\mathcal{X}^{(l)}, \mathbf{Y}^{(l)}, l = 1, \dots, L\}$, computational costs $W_{(1, \dots, L)}$
Output: Updated training sets and observation vectors $\{\mathcal{X}^{(l)}, \mathbf{Y}^{(l)}, l = 1, \dots, L\}$, MF-surrogate models

```

for  $l \leftarrow 1$  to  $L$  do
    Find  $(\rho_{(l-1)}, \Theta_{(l)}, \sigma_{\epsilon_{(l)}}^2)$  minimizing (16)                                ! Set initial MF-surrogate
    Factorize  $\mathbf{K}_{(l)} = (\mathbf{C}_{(l)} + \sigma_{\epsilon_{(l)}}^2 \mathbf{I}_{n_{(l)}})$                                 ! Factorize GP matrix
end
while Not converged do
    if Non-Nested Multifidelity then
        Find  $(\mathbf{x}_*, l_*)$  minimizing  $M^{\text{NN-MF}}(\mathbf{x}, l)$  in (24)                                ! Find new point and level
         $\mathcal{X}^{(l_*)} \leftarrow \mathcal{X}^{(l_*)} \cup \{\mathbf{x}_*\}, \mathbf{Y}^{(l_*)} \leftarrow \mathbf{Y}^{(l_*)} \cup \{f^{(l_*)}(\mathbf{x}_*)\}$                                 ! Update the training set
        for  $l \leftarrow l_*$  to  $L$  do
            Find  $(\rho_{(l-1)}, \Theta_{(l)}, \sigma_{\epsilon_{(l)}}^2)$  minimizing (16)                                ! Update the MF-surrogate
            Factorize  $\mathbf{K}_{(l)} = (\mathbf{C}_{(l)} + \sigma_{\epsilon_{(l)}}^2 \mathbf{I}_{n_{(l)}})$                                 ! Factorize the GP matrix
        end
    else
        Find  $(\mathbf{x}_*, l_*)$  minimizing  $M^{\text{N-MF}}(\mathbf{x}, l)$  in (25)
        for  $l \leftarrow 1$  to  $L$  do
            if  $l \leq l_*$  then
                 $\mathcal{X}^{(l)} \leftarrow \mathcal{X}^{(l)} \cup \{\mathbf{x}_*\}, \mathbf{Y}^{(l)} \leftarrow \mathbf{Y}^{(l)} \cup \{f^{(l)}(\mathbf{x}_*)\}$                                 ! Update the training sets
            end
            Find  $(\rho_{(l-1)}, \Theta_{(l)}, \sigma_{\epsilon_{(l)}}^2)$  minimizing (16)                                ! Update MF-surrogate of  $f^{(l)}$ 
            Factorize  $\mathbf{K}_{(l)} = (\mathbf{C}_{(l)} + \sigma_{\epsilon_{(l)}}^2 \mathbf{I}_{n_{(l)}})$                                 ! Factorize the GP matrix
        end
    end
end

```

This stopping criterion can be based, for instance, on the value of the expected improvement or the exhaustion of a computational budget.

As a final word on computational complexity, we observe that finding the optimum of the merit function calls for multiple estimation of the updated variance $\hat{\sigma}_{(l)}(\mathbf{x}|\mathbf{x}, l)$, which in turn calls for the computation of variance reduction terms, such as $\Delta\hat{\sigma}_{(l)}(\mathbf{x}|\mathbf{x}, l)$, defined by (22). The computationally intensive part of this evaluation is the inversion of the GP matrices in (22). The second one, $\mathbf{K}_{(l)}$, is not depending on the new point \mathbf{x}_* such that it can be factorized and inverted once for all during the search of the couple (\mathbf{x}_*, l_*) maximizing the merit function. The first matrix, on the contrary, depends on the new points, and many inversions can be necessary during the search for \mathbf{x}_* , with a high computational burden as a result. The matrix inversion lemma can alleviate this computational burden. Indeed, dropping temporarily the fidelity level index l , we have

$$\begin{bmatrix} \mathbf{K} & \mathbf{k}(\mathbf{x}_*) \\ \mathbf{k}^T(\mathbf{x}_*) & \kappa(\mathbf{x}_*) + \sigma_\epsilon^2 \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{K}^{-1} + r(\mathbf{x}_*)\mathbf{K}^{-1}\mathbf{k}(\mathbf{x}_*)\mathbf{k}^T(\mathbf{x}_*)\mathbf{K}^{-1} & -r(\mathbf{x}_*)\mathbf{K}^{-1}\mathbf{k}(\mathbf{x}_*) \\ -r(\mathbf{x}_*)\mathbf{k}^T(\mathbf{x}_*)\mathbf{K}^{-1} & r(\mathbf{x}_*) \end{bmatrix}, \quad (28)$$

where

$$r(\mathbf{x}_*) \doteq \left(\kappa(\mathbf{x}_*) + \sigma_\epsilon^2 - \mathbf{k}^T(\mathbf{x}_*)\mathbf{K}^{-1}\mathbf{k}(\mathbf{x}_*) \right)^{-1}. \quad (29)$$

With this expression, the evaluation of $\Delta\hat{\sigma}_{(l)}(\mathbf{x}|\mathbf{x}, l)$ involves only the inverse of the GP model matrix $\mathbf{K}_{(l)}$, which is already assembled and inverted (or factorized) during the MF-surrogate construction (see Section 2.1.1). Using (28),

therefore, leads to significant computational savings as it avoids the inversion of the extended system of size $n_{(l)} + 1$ for every new point proposed. We emphasize this aspect in Algorithm 1 by making clear the factorizations of the GP matrices.

3.5. Elementary example

Figure 2 illustrates the non-nested MF-EGO on the univariate case with functions $f^{(1)}$ and $f^{(2)}$ of (17). The initial training sets are the ones used in Section 2.2.2. Each plot shows a different iteration EGO iteration N_{iter} the MF-prediction $\hat{f}_2(x)$ (black line) and the merit function $M^{\text{NN-MF}}(x, l)$ for $l = 1$ (blue dashed line) and $l = 2$ (red dashed line), respectively. The assumed computing costs are $W_{(1)} = 1$ and $W_{(2)} = 10$. At the first NN-MF-EGO iteration ($N_{\text{iter}} = 1$, Fig. 2(a)), the algorithm selects $(x_* \approx 0.6, l_* = 1)$, which yield a clear improvement of the prediction $\hat{f}^{(2)}(x)$ close to the global optimum (see Fig. 2(b)). After the MF-surrogates updates, the merit function drastically reduces for $l = 1$ and increases for $l = 2$ in the neighborhood of the next point, also it is still presenting a global maximum for $l = 1$. However, the procedure does select $l_* = 2$ at the subsequent iterations $N_{\text{iter}} = 3$ and $N_{\text{iter}} = 4$. After $N_{\text{iter}} = 10$ iterations, the merit function $M^{\text{NN-MF}}(x, l)$ is small and $f^{(2)}$ is well approximated by the MF-surrogate $\hat{f}_2(x)$ at the optimum point (see black and red continuous curves in Fig. 2(f)).




Figure 2: Non-nested EGO iterations for the univariate case of Section 2.2.2. The blue and red symbols indicate the training sets of $f^{(1)}$ and $f^{(2)}$, respectively, while the dashed lines depicts $M^{\text{NN-MF}}(x, l)$ for $l = 1$ and 2. The solid black line is the MF-surrogate $\hat{f}_2(x)$, while the solid red line in the last plot is the function $f^{(2)}(x)$.

4. Analytical test problems

The objective of this section is to illustrate the effectiveness of the proposed MF-EGO methods and demonstrate the interest of using non-nested training sets. To this end, we apply our non-nested method (NN-MF) on different

variants of an analytical problem, and contrast its performances with optimization results obtained with the nested (N-MF) and single fidelity (SF) methods. A comparison of our non-nested method (NN-MF) with another non-nested method proposed in [17] is also given in the last Section 4.5.

4.1. Problem settings

The analytical test problem is the minimization of the 6-dimensional Hartmann function [14], defined as

$$f(\mathbf{x}) = \sum_{i=1}^4 \alpha_i \exp \left(- \sum_{j=1}^6 A_{ij} (x_j - P_{ij})^2 \right), \quad (30)$$

for $\mathbf{x} \in \Omega = [0, 1]^6$, and where $\boldsymbol{\alpha} = (1, 1.2, 3, 3.2)^T$ and the matrix \mathbf{A} and \mathbf{P} are

$$\mathbf{A} = \begin{pmatrix} 10 & 3 & 17 & 3.5 & 1.7 & 8 \\ 0.05 & 10 & 17 & 0.1 & 8 & 14 \\ 3 & 3.5 & 1.7 & 10 & 17 & 8 \\ 17 & 8 & 0.05 & 10 & 0.1 & 14 \end{pmatrix}, \quad \mathbf{P} = 10^{-4} \begin{pmatrix} 1312 & 1696 & 5569 & 124 & 8283 & 5886 \\ 2329 & 4135 & 8307 & 3736 & 1004 & 9991 \\ 2348 & 1451 & 3522 & 2883 & 3047 & 6650 \\ 4047 & 8828 & 8732 & 5743 & 1091 & 381 \end{pmatrix}.$$

The global optimum is $\mathbf{x}_{\text{opt}} = [0.20169, 0.150011, 0.476874, 0.275332, 0.311652, 0.6573]$.

We considered the following sequence to define lower fidelity versions of the Hartmann function,

$$U_{k+1}(\mathbf{x}) = \frac{1}{2} \left(\frac{f(\mathbf{x})^2}{U_k(\mathbf{x})} + U_k(\mathbf{x}) \right), \quad U_0(\mathbf{x}) = u_0 \in \mathbb{R}. \quad (31)$$

In the following we initiate the sequence with $u_0 = -5$; the sequence has then $f(\mathbf{x})$ as a contractive fixed-point: $U_k \xrightarrow{k \rightarrow \infty} f(\mathbf{x})$. Using the sequence of U_k , we define the multi-fidelity approximations of $f(\mathbf{x})$ as

$$f^{(l)}(\mathbf{x}) \doteq U_{k_l}(\mathbf{x} + \delta \mathbf{x}_{k_l}), \quad l = 1, \dots, L, \quad (32)$$

for some increasing sequence of indexes $k_l \in \mathbb{N}$, and where the translation vector $\delta \mathbf{x}_k$ is given by:

$$(\delta \mathbf{x}_k)_i = \delta/k, \quad i = 1, \dots, 6.$$

The translation by $\delta \mathbf{x}_k$ prevents the functions $f^{(l)}(\mathbf{x})$ in (32) to have all the same minimum \mathbf{x}_{opt} ; instead, the minimum of $f^{(l)}$ is translated to a distance $\sqrt{6}|\delta|/k_l$ and we shall use δ to control this distance. We remark, however, that the translation does not affect the convergence of U_k to f for finite $\delta \in \mathbb{R}$. Figure 3 illustrates the effects of increasing δ on the multi-fidelity approximations of $f(\mathbf{x})$. It is seen that increasing δ reduces the correlation between $U_k(\mathbf{x})$ and $U_{k+1}(\mathbf{x})$, while the correlation increases with k .

In all the section, we fixed $L = 3$ and use the multi-fidelity sequence $f^{(1,2,3)}$ defined by $k_1 = 1$, $k_2 = 3$ and $k_3 = k_L = \infty$; in other words $f^{(L)}(\mathbf{x}) = f(\mathbf{x})$. Regarding the initialization of the EGO procedures, all the tests use the same initial training sets. For the MF-methods, the initial lower fidelity set $\mathcal{X}^{(1)}$ consists of $n_{(1)} = 20$ training points drawn randomly in Ω using Latin Hypercube Sampling (LHS). The next training sets $\mathcal{X}^{(2)}$ and $\mathcal{X}^{(3)}$ are obtained by selecting recursively, and at random, $n_{(2)} = 15$ and $n_{(3)} = 10$ elements of their respective predecessor sets ($\mathcal{X}^{(1)}$ and $\mathcal{X}^{(2)}$ respectively). The initial training sets are nested, even for the NN-MF-EGO. For the SF-EGO, the initial training set is selected as $\mathcal{X} = \mathcal{X}^{(1)}$. Finally, unless stated otherwise, we assume $W_{(1)} = 1$, $W_{(2)} = 100$ and $W_{(L)} = 1,000$ for the computational costs of the different fidelity models, and set $N_{\text{iter}} = 400$ for the number of EGO-iterations in the all tests. Note that the total computational cost W_{tot}

$$W_{\text{tot}} = \sum_{l=1}^L W_{(l)} n_{(l)}, \quad (33)$$

increases with N_{iter} , but changes from a method to another depending on the level of the new training points.

The results reported in the following consist in the distance between the surrogate-based optimum and the true optimum $\|\mathbf{x}_{\text{opt}} - \hat{\mathbf{x}}_{\text{opt}}\|$, or the error on the optimal objective value $|f(\mathbf{x}_{\text{opt}}) - \hat{f}^{(L)}(\hat{\mathbf{x}}_{\text{opt}})|$, as functions of the cost W_{tot} . Therefore, we directly compare the efficiencies of the different methods. We consider the effects of the translation parameter δ in Section 4.2, the impact of observation noise in the middle-fidelity model $f^{(2)}$ in Section 4.3, and, finally, the costs per level in Section 4.4.




Figure 3: Random evaluations of 5000 samples of the multi-fidelity analytical test problem (32).

4.2. Impact of the shift

We start by solving the problem for $\delta = 0$ in the definition of the multi-fidelity functions. We report in Fig. 4 the convergence of the three EGO-methods: SF, N-MF, and NN-MF. We plot in the top row of Fig. 4 the distances to the optimum point and global minimum of f as functions of the computational cost W_{tot} . We see that the MF and SF methods have a fast decay of the errors initially (note the log-scale in the distances), although there is a slight shift in the SF method whose cost is higher. For the MF-methods, the fast decay continues until reaching the number of allowed EGO iterations. In contrast, the initial fast convergence stage of the SF-method is followed by a long plateau. After the plateau, the errors reduce again but at a much lower rate than initially. Eventually, all methods exhibit solutions with comparable accuracy, but the two MF-methods perform much better than the SF-method: the computational effort of the MF-methods is roughly 10 times less to reach a distance of 10^{-2} on the optimum. Also, the performances of the two MF-methods are very close; the N-MF-method appears to yield a slightly better estimate of the optimum for a slightly lower cost.

These differences in the convergence trends can be understood by inspecting the evolutions with W_{tot} of the training set sizes $n_{(l)}$ of the MF-methods, reported in the bottom row Fig. 4. The plot for the N-MF-method indicates that the first and second fidelity levels are more frequently selected. In contrast, NN-MF-method never selects and evaluates the first level until a cost of about $4 \cdot 10^4$, after which it becomes the most frequent level evaluated. As the paths of $n_{(2)}$ and $n_{(3)}$ are similar for the two MF-methods, one could expect a higher efficiency for the NN-MF-method, owing to saving evaluations of $f^{(1)}$, but in fact the distance and error for N-MF-method are slightly better at a given cost. This difference is explained by a slightly improved prediction of the nested multi-fidelity model. However, we underline that the two MF-methods have close efficiencies, and both do significantly better than pure SF-method.

We now turn to the case where the definition of the multi-fidelity functions has a non-zero δ . Figure 5 reports the results for the previous situation but with parameter values $\delta = 0.05$ (left) and $\delta = 0.1$ (right). The plots depict the distances to \mathbf{x}_{opt} and errors on f_{opt} for the different methods with $\delta = 0.05$ (left) and $\delta = 0.1$ (right). We see that for the two δ tested, the NN-MF-method (triangles) converges with a rate comparable to the case of $\delta = 0$ (see Fig. 4). This result proves the robustness of the NN-MF-method when the successive multi-fidelity functions are not much correlated, with important changes in the locations of their respective minimums. In contrast, the N-MF-method (squares) exhibits behavior that depends on δ . When $\delta = 0.05$, the convergence with the cost is similar to the NN-MF-method; but for $\delta = 0.1$, the errors on \mathbf{x}_{opt} and f_{opt} are stagnating to quite high values, about two orders of magnitude higher than for the NN-MF-method at the same cost. A closer analysis of the N-MF results indicates that the optimization procedure converges to a local minimum of the Hartmann function, and the number of EGO




Figure 4: Distances to optimum and errors on optimum value as functions of the computational cost for the NN-MF, N-MF and SF-EGO methods (top line), and training sample sizes of the MF-methods (bottom line). Case of $\delta = 0$.

iterations allowed in this example is not enough to learn that this minimum is not the global one.

The evolutions of the training-sets size with the computational cost, depicted in Fig. 6, help to better understand the differences between the behavior of the two MF methods. It is seen that δ has not a significant impact on the evolutions of $n_{(l)}$ for the two MF-method: the N-MF-method uses more low-level evaluations in the construction of the MF-surrogate, while the NN-MF-method considers lower level evaluations only at the end of the EGO iterations when the global optimum is already identified. These differences explain the higher robustness to the shift of the NN-MF-method since one can expect an overall improvement of the MF-surrogate when the estimation involves more correlated functions. Still, the NN-MF-EGO performs much better than the SF-EGO, by substituting several evaluations of $f^{(L)}$ with $f^{(2)}$ without impacting much the predictive quality of the model.

4.3. Noisy estimations

In the next tests, we investigate the impact of evaluation noise in the intermediate multi-fidelity functions on the MF-EGO and show that the NN-MF-method presents an advantage over the N-MF-method in that case. This situation

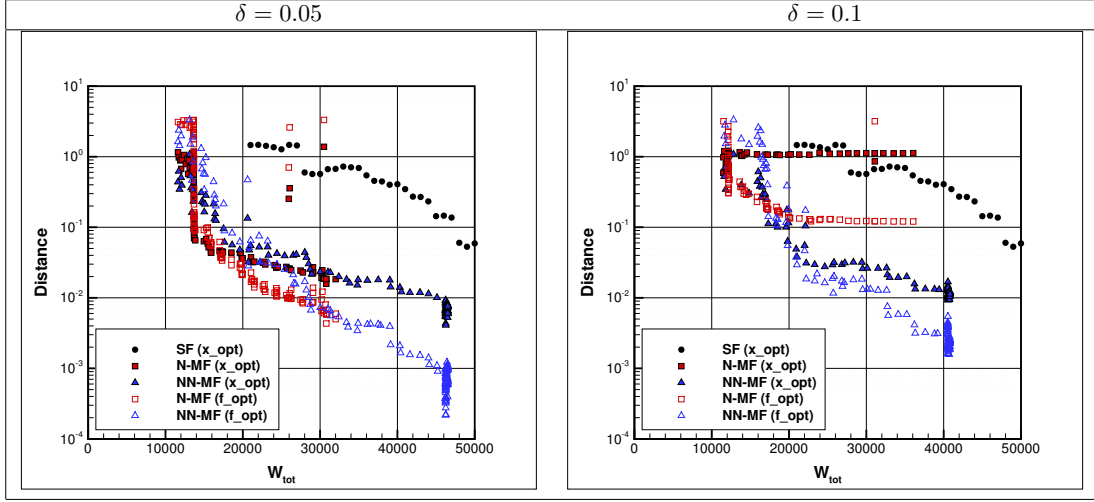


Figure 5: Distances to optimal solution as functions of the cost W_{tot} for the SF, N-MF and NN-MF methods. Cases of $\delta = 0.05$ (left) and $\delta = 0.1$ (right).

corresponds, for instance, to simulations using low-resolution meshes or very permissive convergence criteria in non-linear solvers, in order to obtain rough predictions at a low cost. One could also face this situation when using probabilistic predictors $f^{(l)}$. We then define the noisy version $f^{(2)}$ as

$$f_{\epsilon}^{(2)}(\mathbf{x}) = f^{(2)}(\mathbf{x}) [1 + \eta(\mathbf{x})], \quad (34)$$

where $\eta(\mathbf{x})$ is a white noise random field with uniform distribution in $[0, 0.1]$.

Figure 7 shows the distances and errors to the optimum as functions of the computational cost when solving problem (4.3) with $\delta = 0$ (left) and $\delta = 0.1$ (right). The NN-MF-method clearly outperforms the other approaches, determining the optimum at a cost that increases with δ . It is essential to notice that the SF-method is not affected by the noise since it affects the fidelity level $l = 2 < L$. In contrast, for the considered number of EGO-iterations, the N-MF-method is not improving over the SF-method, which stagnates at a local minimum. Two effects explain the poor performance of the N-MF-method in these cases. A first explanation is the nested character of the N-MF-method: in the present situation where evaluations of the terminal level are frequent, the nested nature forces to compute levels 1 and 2 when level 3 is selected, with an additional overhead compared to the NN-MF-method. Second, the construction in the nested cases uses actual observations at the previous level $l - 1$ to define the observation residual involved in the construction of the MF model $Y^{(l)}(\mathbf{x})$ at level l , see (13). Therefore, the noisy observations directly propagate from level 2 to level 3, and the terminal estimation sees the noise of level 2 in the N-MF-method. The NN-MF-method uses instead observation residuals based on the mean prediction at the previous level, see (18). Substituting $\mathbf{Y}_{\rightarrow(l)}^{(l-1)}$ by $\hat{f}_{(l-1)}(\mathbf{x} \in \mathcal{X}^{(l)})$ smooths-out most of the noise in the observations of level 2 and limits its propagation to subsequent levels. Due to higher noise propagation, the N-MF-EGO has difficulties in finding the true global optimum; it would need many more evaluations to reach the same prediction accuracy. In comparison the NN-MF-method evaluates $f^{(2)}$ infrequently compare to $f^{(L)}$. The number of evaluations can be appreciated from the evolutions of the training sets size reported in Fig. 8, which are very similar to the case without noise.

4.4. Impact of costs

The behaviors of the MF-methods are dependent on the costs associated with the multi-fidelity levels. To illustrate this dependence, we increase the cost associated with the second fidelity model and set in this section $W_{(2)} = 800$. This choice increases sharply the contrast between the costs associated with the evaluation of models $l = 1$ and $l = 2$ while reducing the computational gap significantly with the highest fidelity model.

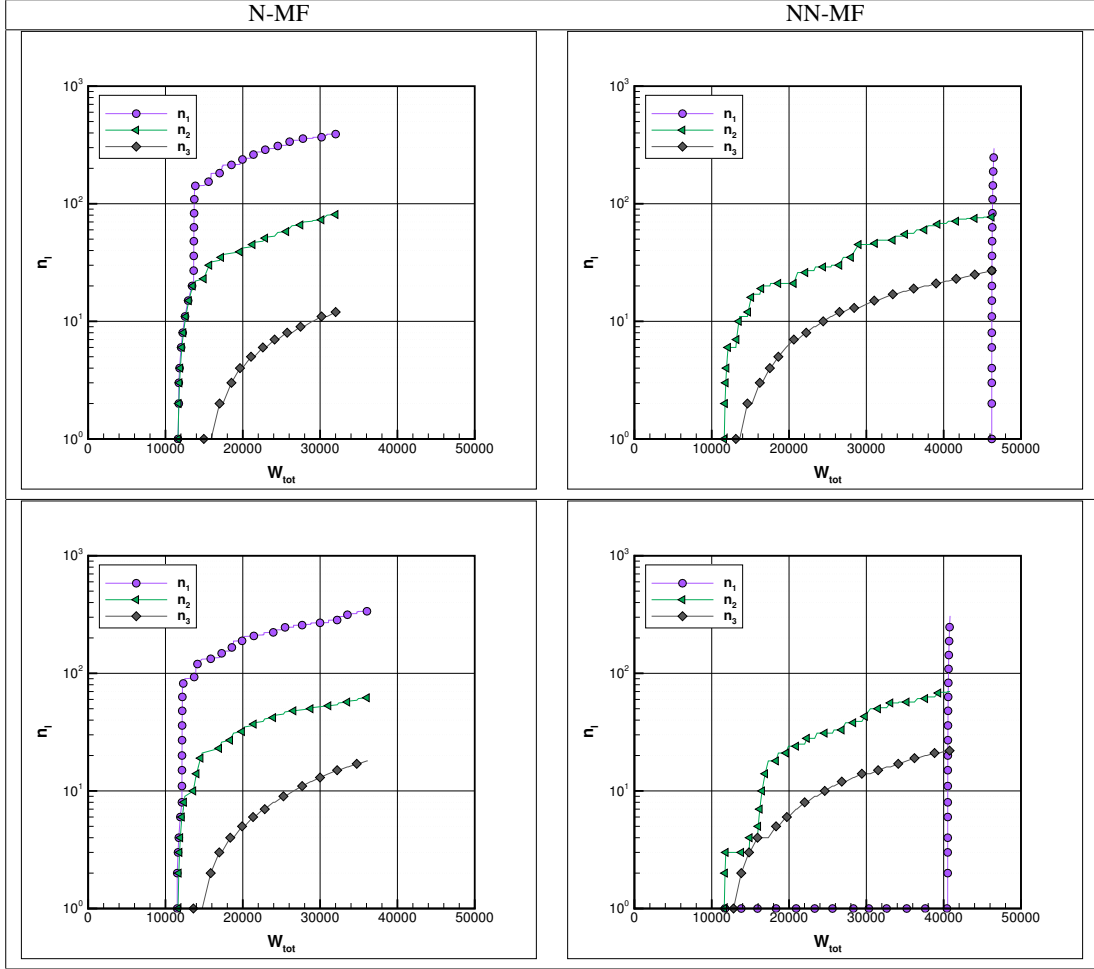


Figure 6: Training sets size as functions of the cost W_{tot} for the N-MF and NN-MF methods. Cases of $\delta = 0.05$ (top-row) and $\delta = 0.1$ (bottom-row).

Figure 9 shows, as functions of W_{tot} and for $\delta = 0$ (left) and $\delta = 0.1$ (right), the distances and errors to the optimum. Once more, it is seen that the NN-MF-method outperforms the N-MF and SF-methods and remains effective for the two values of δ tested. In fact, for the NN-MF-method, the comparison with the respective plots in Figures 4 and 5 shows no qualitative changes when $W_{(2)}$ rises, but an increase of the overall cost of the optimization. For the N-MF-method, more erratic behavior is reported for $\delta = 0$, while the method is again seen to plateau when $\delta = 0.1$, with no convergence to the solution in the prescribed number of EGO iterations.

As expected, more significant changes occur in the size of the training sample sets when $W_{(2)}$ increases, see Fig. 10. A priori, reducing the gap between the costs of levels 2 and 3 will favor the selection of level 3 in the NN-MF-method because the computational saving at level 2 is less important. The effect in the N-MF-method is not as clear because selecting the level 3 implies to compute also levels 1 and 2. Overall, the plots of Fig. 10 indicate that for the N-MF-method, increasing $W_{(2)}$ favors a more frequent selection of $l = 3$, since the difference between n_2 and n_3 reduces compared to the previous cases $W_{(2)} = 100$. Regarding the first level, it is still infrequently selected during the initial phase of the N-MF-method, but it becomes the most selected level subsequently. The N-MF-method starts to plateau when $l = 1$ becomes the dominant level selected. In contrast, the differences between n_2

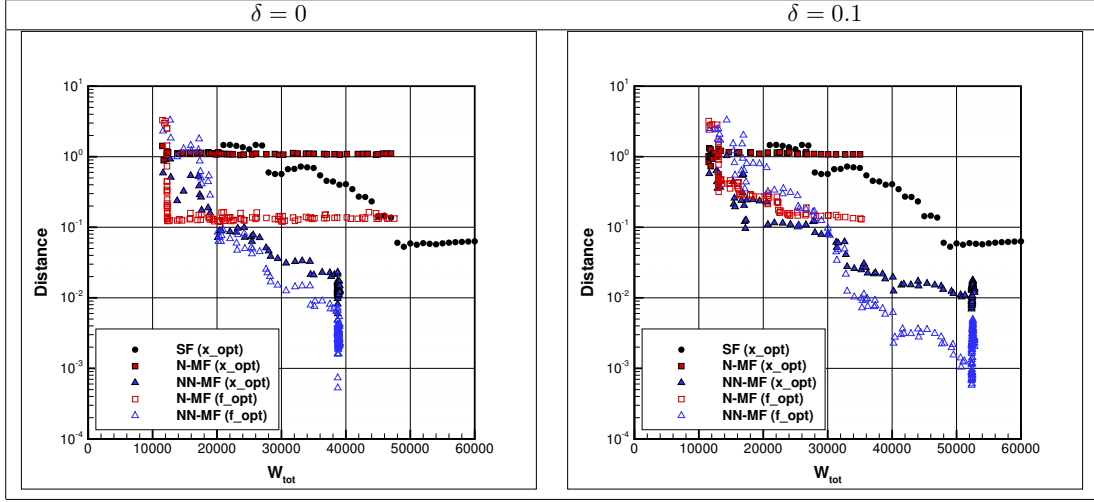


Figure 7: Distances to optimal solution as functions of the cost W_{tot} for the SF, N-MF and NN-MF methods. Cases of noisy estimations of $f^{(2)}$ (see (34)) with $\delta = 0$ (left) and $\delta = 0.1$ (right).

and n_3 increases when $W_{(2)}$ is raised in the NN-MF-method. In the case with $\delta = 0.1$, the algorithm systematically selects the last level $l = 3$ until a fairly converged solution is obtained. In other words, the NN-MF-method returns to the SF method (with lower error because it can exploit the initial training sets at lower levels).

4.5. Comparison with another non-nested method

The results above have contrasted the efficiency and robustness of the proposed method compared to the classical nested MF method. In this section we propose a comparison of our approach with the non-nested NN-MFSKO method proposed in [17]. For a fair comparison, we rely on the same multi-fidelity surrogate construction discussed in Section 2.2. Concerning the MF-infilling strategy, we contrast our approach detailed in Section 3 with the NN-MFSKO approach proposed in [17]. Specifically, the NN-MFSKO method uses the following merit function, recast in our notations:

$$M^{\text{NN-MFSKO}}(\mathbf{x}, l) \doteq \text{EI}_{(L)}(\mathbf{x}) \left(1 - \frac{\sigma_{\epsilon(l)}}{\sqrt{\hat{\sigma}_{(l)}^2(\mathbf{x}) + \sigma_{\epsilon(l)}^2}} \right) \frac{W_{(L)}}{W_{(l)}} \text{corr} \left(Y^{(l)}(\mathbf{x}), Y^{(L)}(\mathbf{x}) \right). \quad (35)$$

Compared to our merit function in (24), the merit function of [17] differs on two points. First, while we consider the augmented expected improvement at the terminal level, $\text{AEI}_{(L)}$, the authors of [17] define an augmentation specific to each level, based on the predictive variance $\hat{\sigma}_{(l)}^2$ and noise level $\sigma_{\epsilon(l)}$ of level l . Second, while our approach ranks the importance of the levels by the reduction of predictive variance for a new evaluation at level l , the merit function in (35) uses the correlation between the predictions at levels l and L . Although the authors of [17] do not discuss the choices leading to (35), it is clear that the point of view is different from ours. The objective of (35) is to exploit the predictive capability of lower fidelity levels in order to save computational resources. The merit function then mitigates the correlation between level l and the terminal one with the relative amount of noise and the evaluation cost *at that level*. Our merit function, on the contrary, focuses on the predictions at the terminal level L and aims at improving it (*i.e.* reducing the predictive variance of $\hat{\sigma}_{(L)}^2$) at the lower cost through a possible evaluation of a lower-level model. The two MF-merit functions were tested on the whole set of analytic test functions presented above; to limit the paper length, we only present a few representative cases of the behavior of the two approaches.

The results shown Figure 11 correspond to the problem defined in Section 4.2 with no shift ($\delta = 0$). In the left plot, we report the evolution with W_{tot} of the distance to the true optimum, for the SF method and NN-MF methods based




Figure 8: Training sets size as functions of the cost W_{tot} for the N-MF (left) and NN-MF (right). Cases of noisy estimations of $f^{(2)}$ (see (34)) with $\delta = 0$ (top) and $\delta = 0.1$ (bottom).

on merit functions in (24) and (35) (labeled NN-MFSKO). For this problem, the two MF-methods perform initially similarly. However, while our approach maintains a fast convergence rate to the optimum, the correlation-based merit function presents a long plateau before recovering a convergence similar to the SF-method. Our approach requires about 5 to 6 times less work to find an optimum at a distance of less than 10^{-2} . The right plot of Figure 11 reports the evolutions of the training sets size as functions of the work spent in the optimization. Our approach starts much earlier to evaluate the models at levels $l < L = 3$ compared to the other MF method. Evaluations at the second level ($l = 2$) appear from the very beginning of the optimization procedure and, after a total cumulated work $W_{\text{tot}} \approx 4 \cdot 10^5$, the procedure focuses entirely on the first level. In contrast, the introduction of lower levels is delayed and much more progressive for the NN-MFSKO merit function, resulting in a slower convergence with the computational effort. The initial mild correlation between $Y_{(L)}$ and $Y_{(l)}$ for $l < L$ explains the delayed selection of lower levels. Then, the reduction of the terminal level variance must be sufficient to engage other lower levels.

Following the test cases of the previous sections, Figure 12 compares the effects of the merit functions in presence of a shift $\delta = 0.1$ (left column), noise in the second level model (center column) and less contrasted work with $W_{(2)} = 800$ (right column). Shown are the converges to the optimum point (top row) and evolutions of the training

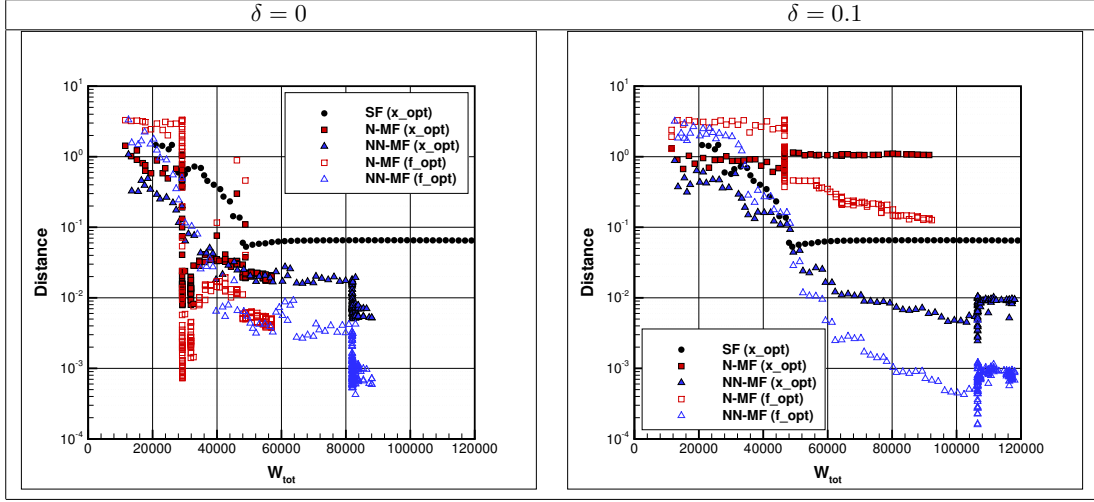


Figure 9: Distances to optimal solution as functions of the cost W_{tot} for the SF, N-MF and NN-MF-methods. Cases of $\delta = 0$ (left) and $\delta = 0.1$ (right), with cost $W_{(2)} = 800$.

sets size (bottom row) as functions of the total computational effort W_{tot} . Focusing first on the impact of a non-zero shift, we expect that $\delta \neq 0$ deteriorates the correlation between levels such that the NN-MFSKO merit function needs more iterations before evaluating models with $l < 3$, compared to the case with $\delta = 0$ (shown in Figure 11). The results indicate that this effect is marginal and $\delta > 0$ moderately slows down the convergence of optimization and the selection of levels $l < L$. However, the effect is significant enough to result in a MF optimization that is overall less effective than the SF-method in the sense that, for the same terminal cost, the NN-MFSKO merit function of [17] ends farther to the true optimum point. On the contrary, our approach maintains a high convergence rate and produce a much better estimate of the optimum compared to the SF method's for the same terminal cost.

Compared to the baseline case with $\delta = 0$, the introduction of noisy evaluations in the second model has a significant impact on the behavior of the optimization process with the merit function 35. For the number of optimization iterations considered in the tests, $N_{\text{iter}} = 400$, the NN-MFSKO method plateaus, quite similarly to the nested method (see Figure 7). In other words, the merit function (35) is not much robust to evaluation noise, in contrast to our function in (24). The center plot of the bottom row of Figure 12 explains these contrasted behaviors. In our approach, the optimization procedure quickly disregards the second level to focus on the first and last ones. The case of the NN-MFSKO merit function (35) is quite different. After selecting systematically the terminal level, the algorithm selects the second level with increasing frequency, and never select the first one. In other words, the algorithm focuses on reducing the predictive variance at level $l = 2$, without inducing a clear improvement on the prediction at the terminal level and the optimum estimation. This test case highlights the advantage of our approach which focuses on improving the prediction at level L . Interestingly enough, augmenting the EI by the estimated noise $\sigma_{\epsilon(l)}$ of the level is not able to penalize the selection of the noisy level. Perhaps, a different weighting of the noise between levels could help in these situations.

The final example shown concerns the case without shift ($\delta = 0$) and evaluation noise, but an increased work for the second level, such that the gap between the last 2 levels is reduced (see Section 4.4). While performing similarly to our method for the first iterations, the optimization with the merit function (35) subsequently stagnates and becomes even less effective than the SF-method. Looking to the evolutions with W_{tot} of the training sets size, it appears that the merit function just slightly delays the selection of the second level (compared to the case with $W_{(2)} = 100$ in Figure 11) which remains otherwise selected with high frequency. The frequent selection of the second level contrasts with the behavior of our merit function which delays *and* significantly reduces the selection frequency of the second level.




Figure 10: Training sets size as functions of the cost W_{tot} for the N-MF (left) and NN-MF (right) methods. Cases of $\delta = 0$ (top-row) and $\delta = 0.1$ (bottom-row), with $W_{(2)} = 800$.

4.6. Summary

From the tests presented in this section, we can conclude that the efficiency of the N-MF-EGO method deteriorates as the correlation between the successive levels is decreasing. The nested property is also detrimental in the presence of evaluation noise or when the contrast in the cost of the highest models is not progressive enough, to the point where the N-MF-method can be less effective than its direct SF counterpart. The non-nested MF-method with the proposed merit function (24), on the contrary, is robust and much less sensitive to low correlations between levels. Besides, it can entirely focus the effort on the last level if the lower levels are not useful to improve the prediction or when the contrast in computational costs is not significant enough. This behavior indicates that the proposed NN-MF merit function provides a robust and flexible strategy able to disregard useless fidelity models. This feature is particularly attractive in situations where the fidelity levels may not possess a clear hierarchy.

From the conducted tests on the NN-MFSKO approach proposed in [17], we can conclude that our NN-MF method is also more robust and less sensitive to low correlations between levels. Indeed, in comparative tests on shift effects, noisy observations in intermediate level, or cost intermediate level increases, the correlation-based merit function of [17] performs similarly to our method for the first iterations only. Then, the NN-MFSKO method presents a long

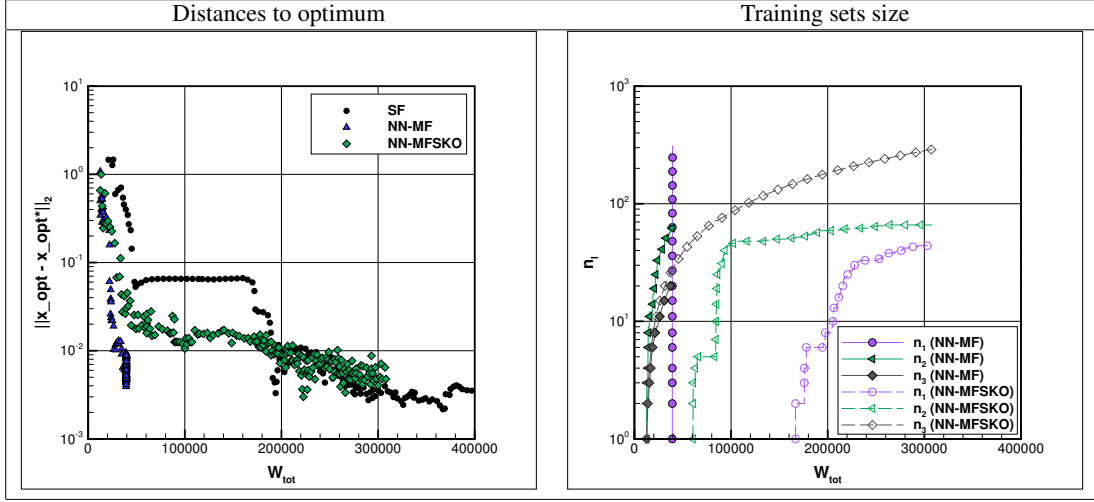


Figure 11: Comparison of the proposed NN-MF method (see (24)) with the NN-MFSKO method of [17] (see (35)). Problem of section 4.2 with $\delta = 0$.




Figure 12: Distances to optimum (top row) and training sets size (bottom row) as functions of W_{tot} .

plateau before recovering a convergence similar to the SF-method, or even, ends farther to the true optimum point in case of noisy intermediate level observations.

5. Application to nonlinear FSI problem

We now apply the MF-EGO methods on a complex problem corresponding to the optimization of the trimmings of a realistic system of sails. The physics of sail systems involve very complex phenomena, such as nonlinear Fluid-Structure Interaction (FSI) and aero-elastic instabilities. Moreover, the modeling of real sailing conditions is still an open challenge because of the significant uncertainties in the prediction of wind and sea states. To our knowledge, the sails optimization has thus been limited so far to somehow idealized situations. For instance, the authors in [27] performed sail-shape optimizations without accounting for the full FSI problem. Sail trimmings optimization based on the FSI problem was considered in [28], but for two-dimensional "sails" only. Regarding three-dimensional FSI problems, the authors of [29] mention optimization of the trimming of sails, but within an inviscid flow approximation, and the authors provide very few details on the optimization procedure they used. In our opinion, our work in [6] was the first one to consider the trimming optimization of a realistic sail system, and validate the numerical optimization using a comparison with the optimization of an experimental sail system in a wind tunnel. The optimization procedure used in [6] is the SF-EGO method.

The present section aims at demonstrating the interest of the NN-MF-EGO approaches for the optimization of complex sail systems with realistic physical models (e.g., nonlinear FSI and turbulent flows). In particular, we want to evidence the reduction of the optimization cost of the NN-MF method. Multi-fidelity models are natural for the optimization of flexible sails systems, where the models for the flow (fluid problem) and the sails deformation (structure problem) can both involve different levels of simplification with significantly differing computational costs. As a result, one can combine these different structures and fluid models to obtain different fidelity models of the FSI problem. To focus on the multi-fidelity aspects, we restrict ourselves to a problem with only two trimming parameters (the mainsail and jib sheet length) but for a costly FSI modeling including detailed nonlinear structural model and three-dimensional turbulent flow simulations (the highest fidelity model).

Indeed, the numerical models of sail systems are composed of several fidelity levels, that may concern both the structural and fluid models. The following optimization problem includes five fidelity levels (in structure and fluid) on a two-parameter objective function in order to allow easy surrogate visualizations. The application is also solved with the highest fidelity level alone to perform SF and MF efficiency comparisons. The optimization problem and the numerical solvers are presented in Section 5.1. Results are then given in Section 5.2.

5.1. Optimization of sails trimming

5.1.1. Problem definition

We consider the trimmings optimization of a *Figaro* mainsail and jib system. The *Incidence Sails* company provided the geometric and elastic characteristics of the two sails. The boat is assumed to cruise at a constant speed of 5.6 knots, in a constant wind with a true wind-speed of 8 knots (at 10 meters height) and a true wind-angle of 50 degrees. The trim, heel, and heading angles of the boat are fixed to 0, 10, and 0 (degrees), respectively. The optimization concerns two trimming parameters: the mainsail sheet (L_{main}) and jib sheet (L_{jib}) lengths (see Fig. 13).

The objective is to minimizing a composite objective function $f(\mathbf{x})$, consisting is a linear combination of the propulsion F_{prop} and the side F_{side} forces:

$$\mathbf{x}_{\text{opt}} = \arg \min_{\mathbf{x} \in \Omega} f(\mathbf{x}) = - (F_{\text{prop}}(\mathbf{x}) + 0.05 F_{\text{side}}(\mathbf{x})) . \quad (36)$$

Note that the propulsion force is counted negatively, for it is minimized, while the coefficient on F_{side} is set to penalize effects incurring from the side force, such as added hydrodynamic drag to counter the leeway-drift. The design domain is $\Omega = [1.5, 1.7] \times [2.45, 2.65]$, which correspond to a trimming range of 20 cm for the two sheets.

The MF framework uses $L = 5$ levels as reported in Table 1. The different levels combine different models of structure and fluid flow, each with increasing predictive capabilities. Concerning the structural differences, the sails and the rig will be considered rigid or flexible, depending on the level. Similarly, the flow models encompass inviscid and viscous assumptions, with different mesh refinements and boundary treatments in the latter case (levels $l = 4$ and 5). Specifically, the flow model at level $l = 5$ involves roughly 50% more cells, with resolved wall boundary conditions, compared to the preceding level, $l = 4$, which relies on a wall function. Table 1 also reports estimates of typical computational costs associated with the resolution of the different models. We see that the models based on the




Figure 13: Design variables: mainsail sheet (1) and jib sheet (2) lengths.

inviscid flow assumption are much less expensive to solve than for viscous flow models. Also, the mesh-refinement in the viscous flow models between levels 4 and 5 translates into a roughly ten times increase of the computational time.

Level	Sails model	Rig model	Fluid model	Fluid mesh size	Cost $W_{(L)}/W_{(1)}$
$l = 1$	Rigid	Rigid	Inviscid	–	1.0
$l = 2$	Flexible	Rigid	Inviscid	–	3.3
$l = 3$	Flexible	Flexible	Inviscid	–	4.2
$l = 4$	Flexible	Flexible	Viscous & Wall-law	2.3M	148.0
$l = 5$	Flexible	Flexible	Viscous & Resolved wall	3.6M	1430.4

Table 1: Structural and fluid model characteristics for the 5 fidelity levels considered.

5.1.2. Numerical solvers

We now briefly describe the different solvers involved in the optimization.

Structural solvers. We rely on the ARA software developed by K-Epsilon for the structural model of the elastic sails and rig. The definition of the geometric and elastic characteristics of the sails, as designed by Incidence Sails, are imported using the SAILPACK software developed by BSG Developments. The elastic model involves Timoshenko beams for the rig (mast, boom) and sail (battens) elements, cables (sail sheets), and Constant Strain Triangles (CST) membrane elements (cloth) of various types in large displacement formulation [30]. A wrinkle model [31] is added to capture local sail deformations with a reasonable number of membrane elements. See [32] or [33] for further details on CST elements and the wrinkle model. The nonlinear solutions (equilibrium under the fluid loads) are computed employing a Newton method with Aitken relaxation.

Flow solvers. Two different flow solvers implement the inviscid and viscous flow models. In the inviscid case, we use the software AVANTI [34]. This solver combines a Vortex Lattice Method (VLM) [31] at the sails surface, with a particle method for the dynamics of their wakes. In this approach, the sails surfaces are discretized in rectangular panels supporting doublets distributions, while particles carrying vorticity are released at the sails trailing edge to enforce the Kutta condition. The long-time limit of the unsteady simulation is considered as the steady inviscid solution. The inviscid model (see Fig. 14(a)) assumes slip conditions over all the sail surface and therefore cannot account for detached flow situations which can occur if the trimming parameters induce large angle of attacks or highly cambered sail profiles. However, the inviscid solver can accommodate substantial changes in the sails geometry since it requires only the discretization of the sail surfaces.

In the viscous case, we used ISIS-CFD software (of FINETM/Marine), which solves the incompressible Unsteady Reynolds-Average Navier-Stokes Equations (URANS). The turbulence model we used in this work is the two equations shear stress transport (SST) $k - \omega$ model [35]. This solver is based on a parallel finite-volume method, accommodating meshes with both structured and unstructured cells. Geometrical changes are less trivial to treat than for the inviscid case, and computationally more demanding, as the viscous solver requires a volume mesh fitted to the deformed sail geometries. In this work, we used the semi-automated mesh-generator HEXPRESSTM to construct meshes on reference geometries, and the Mesh Deformation Propagation (MDP) method [36] to adapt the mesh as the sails deform.

Coupled problem. We assume that the FSI problem has a steady solution for all values of the trimming parameters within the optimization range, and use a quasi-steady approach. Specifically, a quasi-monolithic algorithm [32] is employed to couple the structural and flow solvers, and produce steady FSI solutions. This algorithm relies on an implicit coupling strategy suited to partitioned solvers. Briefly, the resolution of the structural model is nested inside the iterations of the flow solver. This approach preserves the convergence and stability properties of the monolithic approach. More details on the coupling algorithms can be found in [32].

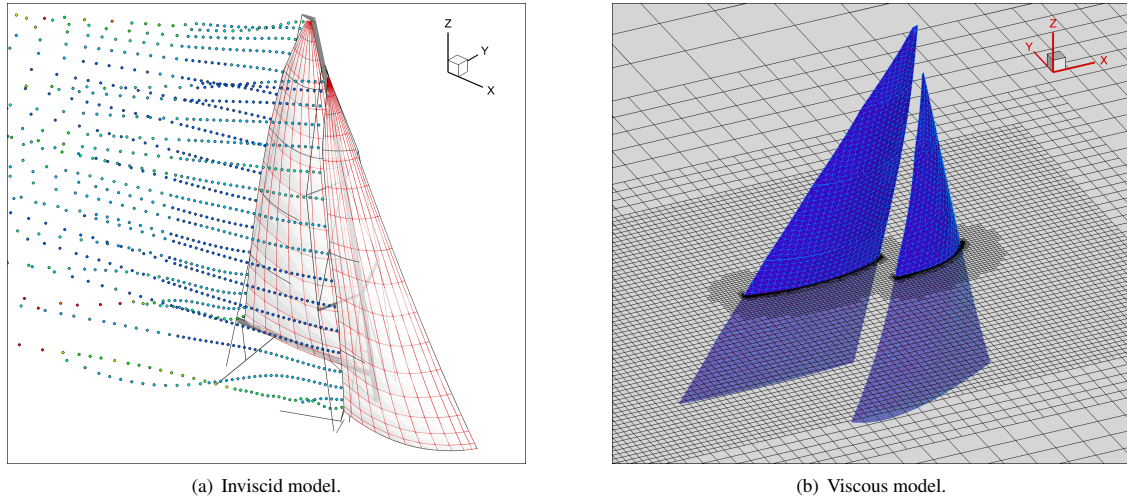


Figure 14: Illustration of the fluid meshes corresponding to the inviscid and viscous solvers.

5.2. Optimization results

The NN-MF-method is initialized from nested training sets $\mathcal{X}^{(l)}$ with $n_l = 1 + 2(L - l)$ elements. Therefore, the initialization uses a single observation of $f^{(L)}$, the model for which we wish to optimize the trimming parameters.

We start by presenting in Fig. 15 the MF-surrogates $\hat{f}^{(l)}$ obtained at the end of the optimization procedure. We first remark a clear separation between the predictions of the surrogates based on the inviscid flow models ($l = 1, 2, 3$), and the predictions of the surrogates incorporating viscous flow estimations ($l = 4, 5$). These differences are particularly important as the trimming variables depart from the optimum point, with a clear under-estimation of the objective value in the inviscid cases: $\hat{f}^{(l=1,2,3)}(\mathbf{x}) \ll \hat{f}^{(l=4,5)}(\mathbf{x})$. This behavior is due to the inability of the inviscid models to account for the separation of the flow from the sails, as discussed below. Perhaps more surprising, $\hat{f}^{(1)}$ is the inviscid MF-surrogate with predictions the closest to the viscous ones, albeit with the least dependences on the trimming parameters compared to the other inviscid surrogates $\hat{f}^{(2,3)}$, which exhibit more noticeable dependences on the optimization variables. In contrast, the viscous MF-surrogates $\hat{f}^{(4)}$ and $\hat{f}^{(5)}$ have predictions in close agreement to each other over the whole domain Ω . This proximity is not surprising since, compared to $\hat{f}^{(4)}$, the construction of

$\hat{f}^{(5)}$ incorporates just a few evaluations of the objective function $f^{(5)}$, based on a more refined mesh and a different boundary flow conditions, but the same physical models otherwise.




Figure 15: MF surrogates after $N_{\text{iter}} = 60$ iterations of the NN-MF optimization.

Eventhough, the MF-models $\hat{f}^{(l)}(x)$ reported above are only approximations of the MF-models $f^{(l)}(x)$, they provide a fair idea of the situation. In particular, they highlight potential difficulty of defining a clear hierarchy of fidelity level, when considering various physical approximations. Such situations clearly call for robust and agile MF-construction and infilling strategies, in order to minimize as much as possible the impact of *a priori* regarding the predictive capabilities of the models.

Figure 16 depicts the convergence of the optimization process up to $N_{\text{iter}} = 60$ iterations of the NN-MF approach. The value of $M^{\text{NN-MF}}$ is used to color the curve, and it indicates the expected improvement in the computed levels as a function of N_{iter} . Thus, we observe a decrease of the merit function and the procedure converges in $N_{\text{iter}} \approx 60$ iterations. We also qualitatively note that levels 2 and 3 are the ones that have been evaluated most of the time during the optimization.

Figure 17(a) reports the number of calls per fidelity level as a function of N_{iter} . The level $l = 3$ was selected most of the time, while level $l = 1$ has approximately the same number of calls than level 4 and 5. We also note that during the first 20 iterations, the level $l = 4$ is the second fidelity level of choice. These trends can be explained using Fig. 17(b), which depicts the mean prediction of the objective function $f_l(x)$ as a function of the two trimming parameters for all fidelity levels. The surface colors correspond to the colors used to differentiate the levels in Fig. 17(a). We note that the prediction f_1 is closer to the highest fidelity level f_5 , compared to the predictions f_2 and f_3 . However, the prediction of the objective function by the level 1 model has very low variabilities compared to the other levels. Consequently, although the computational cost of this level is low, the non-nested EGO approach does not compute a large number of observations at level $l = 1$ because the model is easy to learn, and additional points do not provide a significant reduction of the prediction variance at the highest fidelity level $l = 5$. Models at levels 2 and 3 overestimate the objective function, but they contain variabilities that correlate well with f_5 . The level 4 is




Figure 16: MF-EGO convergence when solving (36).

very close to the highest fidelity level in terms of mean response and variabilities. This proximity is not surprising, since both models use the same structural and flow solvers, and differences are only on the sails boundary conditions and mesh refinements. Overall, the non-nested approach explores Ω using low-cost levels and avoids useless fidelity levels by eventually refining promising areas with the two highest-fidelity levels.

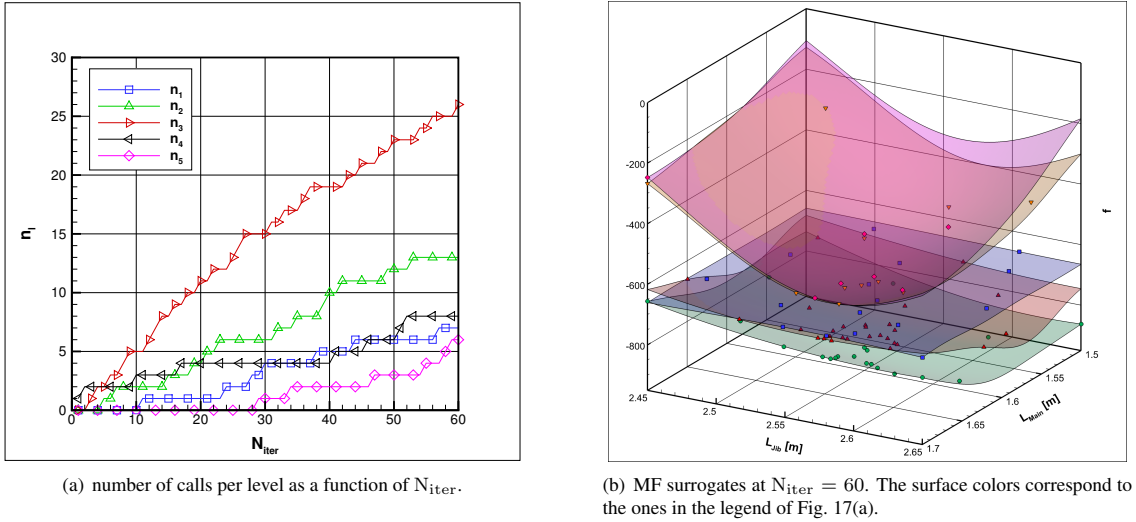


Figure 17: Number of call per level and MF surrogates at convergence when solving (36).

Figure 18 provides a (sailor) view of the optimally trimmed sails that are given by each fidelity level. The sails and rigs colors correspond to fidelity levels colors already used in Fig. 17(a). Those views allow us to understand the differences in objective responses of each level shown in Fig. 17(b). The two viscous models 4 and 5 are very close in terms of the jib and mainsail shapes. Differences are much more significant with the fidelity levels based on

the inviscid model. The mainsail drafts, in particular, are more pronounced for the viscous models. This difference comes from the inviscid flow model, which assumes an attached flow over the sails, leading to overestimations of aerodynamic forces and sail shapes with large drafts as a result.




Figure 18: Optimal trimming at each fidelity level when solving (36).

Computation cost saving: SF-EGO versus MF-EGO. We also solve the trimming optimization problem (36) with the SF-EGO approach on the highest fidelity level 5, using an initial LHS plan of 5 observations. Figure 19 shows the convergences of the SF and MF optima as functions of the cost. The NN-MF converges much faster than the SF method while producing optimums that are in very close agreement. Specifically, it takes approximately 3 times less computational resources (computational time) to converge for the NN-MF method than for the SF-EGO method. Regarding the estimated optimum point, the differences between SF-EGO and MF-EGO are ≈ 0.25 cm for the mainsail and ≈ 0.55 cm for the jib sheets lengths.

6. Conclusion

We have presented in this paper a new infilling formulation for multi-fidelity surrogate-based optimizations. The multi-fidelity approach allows using Efficient Global Optimization (EGO) with surrogate models based on observations at different fidelity levels. We proposed and compared two distinct approaches differing by the use of nested or non-nested training-sets. In the two cases, we proposed suitable merit functions, based on the concept of Expected Improvement, to select the new point and level in the EGO procedure. Specifically, the sequential approach finds the most promising design point and the fidelity level for the evaluation of the objective function at the new point. The selection criteria balance the computational cost at each level and the resulting uncertainty reduction in the multi-fidelity estimation of the objective function.

The proposed MF optimization approaches were validated and tested on several analytical problems. The results show that, compared to the nested formulation, the non-nested MF approach is more efficient in terms of numerical cost, but also much more robust against noisy evaluations of intermediate fidelity levels and low correlations between levels. It is also less sensitive to the computational costs associated with the levels.

Finally, the method was applied to the optimization of a fully nonlinear numerical fluid-structure interaction problem. Here we aimed at optimizing the sail trimmings of Figaro racing yacht. The problem has two parameters, and we considered five fidelity levels related both to structure and fluid models. In this example, the non-nested

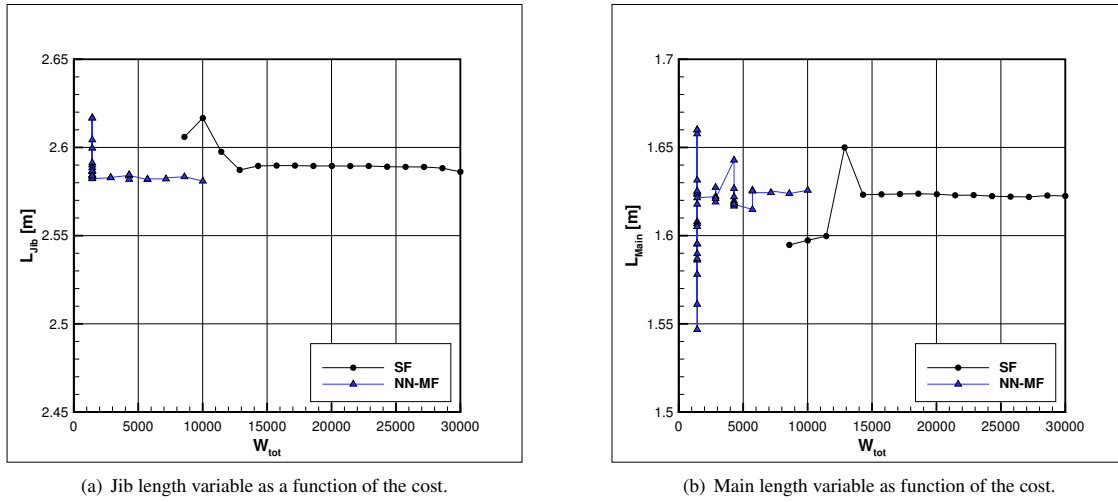


Figure 19: Optimal trimmings of SF and MF surrogates when solving (36).

multi-fidelity EGO was compared to the single-fidelity EGO only. This complex application highlighted the interest of using a non-nested multi-fidelity approach to reduce the computational cost significantly.

Future works will concern the formulation of multi-fidelity surrogate-based optimization in the parallel computing framework. Indeed, the computational costs and the multi-fidelity responses can be used to define a formulation of merit function, which allows several evaluations at each iteration of the EGO. Such a parallel extension would help of reducing, even more, the global cost of the optimization.

References

- [1] Jones, D.R., Schonlau, M., and Welch, W.J., Efficient Global Optimization of Expensive Black-Box Functions, *Journal of Global optimization*, 13(4):455–492, 1998.
- [2] Simpson, T., Poplinski, J., Koch, N.P., and Allen, J., Metamodels for computer-based engineering design: Survey and recommendations, *Engineering with Computers*, 17(2):129–150, 2001.
- [3] Jeong, S., Murayama, M., and Yamamoto, K., Efficient optimization design method using kriging model, *Journal of aircraft*, 42(2):413–420, 2005.
- [4] Glaz, B., Friedmann, P.P., and Liu, L., Helicopter vibration reduction throughout the entire flight envelope using surrogate-based optimization, *Journal of the American Helicopter Society*, 54(1):12007, 2009.
- [5] Aghajari, N. and Schäfer, M., Efficient shape optimization for fluid-structure interaction problems, *Journal of Fluids and Structures*, 57:298 – 313, 2015.
- [6] Sacher, M., Hauville, F., Duvigneau, R., Le Maître, O., Aubin, N., and Durand, M., Efficient optimization procedure in non-linear fluid-structure interaction problem: Application to mainsail trimming in upwind conditions, *Journal of Fluids and Structures*, 69:209 – 231, 2017.
- [7] Kleijnen, J.P., Kriging metamodeling in simulation: A review, *European Journal of Operational Research*, 192(3):707 – 716, 2009.
- [8] Lam, R., Allaire, D., and Willcox, K., Multifidelity optimization using statistical surrogate modeling for non-hierarchical information sources, In *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2015.
- [9] Poloczek, M., Wang, J., and Frazier, P., Multi-information source optimization, In *31st Conference on Neural Information Processing Systems (NIPS 2017)*, Long Beach, CA, USA., pp. 4288–4298, 2017.
- [10] Ghoreishi, S. and Allaire, D., Multi-information source constrained Bayesian optimization, *Structural and Multidisciplinary Optimization*, 59:977–991, 2019.
- [11] Kennedy, M.C. and O’Hagan, A., Predicting the output from a complex computer code when fast approximations are available, *Biometrika*, 87(1):1–13, 2000.
- [12] Forrester, A.I., Söbester, A., and Keane, A.J., Multi-fidelity optimization via surrogate modelling, *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 463(2088):3251–3269, 2007.
- [13] Elsayed, K., Optimization of the cyclone separator geometry for minimum pressure drop using co-kriging, *Powder Technology*, 269(Supplement C):409 – 424, 2015.
- [14] Picheny, V., Wagner, T., and Ginsbourger, D., A benchmark of kriging-based infill criteria for noisy optimization, *Structural and Multidisciplinary Optimization*, 48(3):607–626, Sep 2013.

- [15] Le Gratiet, L. and Cannamela, C. Kriging-based sequential design strategies using fast cross-validation techniques with extensions to multi-fidelity computer codes. working paper or preprint, October 2012.
- [16] Plutowski, M., Sakata, S., and White, H., Cross-validation estimates imse, In *Advances in neural information processing systems*, pp. 391–398, 1994.
- [17] Huang, D., Allen, T.T., Notz, W.I., and Miller, R.A., Sequential kriging optimization using multiple-fidelity evaluations, *Structural and Multidisciplinary Optimization*, 32:369–382, 2006.
- [18] Bartoli, N., Meliani, M., Morlier, J., Lefebvre, T., Bouhlel, M.A., and Martins, J., Multi-fidelity efficient global optimization: Methodology and application to airfoil shape design, In *AIAA Aviation 2019 Forum*, p. 3236, 2019.
- [19] Perdikaris, P., Venturi, D., Royset, J., and Karniadakis, G.E., Multi-fidelity modelling via recursive co-kriging and Gaussian–Markov random fields, In *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, number 471.2179, 2015.
- [20] Parussini, L., Venturi, D., Perdikaris, P., and Karniadakis, G., Multi-fidelity Gaussian process regression for prediction of random fields, *Journal of Computational Physics*, 336:36–50, 2017.
- [21] Kandasamy, K., Dasarathy, G., Schneider, J., and Barnabas, P., Multi-fidelity Bayesian optimisation with continuous approximations, In *Proceedings of the 34th International Conference on Machine Learning*, Vol. 70, 2017.
- [22] Kandasamy, K., Dasarathy, G., Oliva, J., Schneider, J., and Barnabas, P., Multi-fidelity Gaussian process bandit optimisation, *Journal of Artificial Intelligence Research*, 66:151–196, 2019.
- [23] Rasmussen, C.E. and Williams, C.K.I., *Gaussian Processes for Machine Learning*, MIT Press, 2006.
- [24] Hansen, N. The CMA Evolution Strategy: A Comparing Review. In *Towards a new evolutionary computation*, pp. 75–102. Springer, 2006.
- [25] Huang, D., Allen, T.T., Notz, W.I., and Zheng, N., Global Optimization of Stochastic Black-Box Systems via Sequential Kriging Meta-Models, *Journal of global optimization*, 34(3):441–466, 2006.
- [26] Le Gratiet, L. and Garnier, J., Recursive co-kriging model for Design of Computer experiments with multiple levels of fidelity, *International Journal for Uncertainty Quantification*, 4(5):365–386, 2014.
- [27] Rousselon, N., Optimization for Sail Design, In *modeFRONTIER Conference*, 2008.
- [28] Chapin, V.G., Neyhousser, R., Dulliand, G., and Chassaing, P., Design optimization of interacting sails through viscous CFD, In *INNOVSail, Innovation in high performance sailing Yacht*, Lorient, 2008.
- [29] Ranzenbach, R., Armitage, D., and Carrau, A., Mainsail Planform Optimization for IRC 52 Using Fluid Structure Interaction, In *The 21st Chesapeake Sailing Yacht Symposium, SNAME*, 2013.
- [30] Augier, B., Etudes expérimentales de l’interaction fluide-structure sur surface souple: application aux voiles de bateaux., PhD Thesis, Université de Bretagne Occidentale, 2012.
- [31] Katz, J. and Plotkin, A., *Low-Speed Aerodynamics*, Cambridge University Press, second edition, 2001.
- [32] Durand, M., Interaction fluide-structure souple et légère, application aux voiliers, PhD Thesis, Ecole Centrale de Nantes, 2012.
- [33] Nakashino, K. and Natori, M.C., Efficient modification scheme of stress-strain tensor for wrinkled membranes, *AIAA journal*, 43(1):206–215, 2005.
- [34] Roux, Y., Durand, M., Leroyer, A., Queutey, P., Visonneau, M., Raymond, J., Finot, J.M., Hauville, F., and Purwanto, A., Strongly coupled VPP and CFD RANSE code for sailing yacht performance prediction, In *3rd High Performance Yacht Design Conference, 2-4 December*, pp. 215–225, Auckland, 2008.
- [35] Menter, F.R., Kuntz, M., and Langtry, R., Ten Years of Industrial Experience with the SST Turbulence Model, *Turbulence, heat and mass transfer*, 4:625–632, 2003.
- [36] Durand, M., Leroyer, A., Lothodé, C., Hauville, F., Visonneau, M., Floch, R., and Guillaume, L., FSI investigation on stability of downwind sails with an automatic dynamic trimming, *Ocean Engineering*, 90:129–139, 2014.