scientific reports

OPEN

Non-myopic multipoint multifidelity Bayesian framework for multidisciplinary design

Francesco Di Fiore[1] & Laura Mainini[1,2]

The adoption of high-fidelity models in multidisciplinary design optimization (MDO) permits to enhance the identification of superior design configurations, but would prohibitively rise the demand for computational resources and time. Multifidelity Bayesian Optimization (MFBO) efficiently combines information from multiple models at different levels of fidelity to accelerate the MDO procedure. State-of-the-art MFBO methods currently meet two major limitations: (i) the sequential adaptive sampling precludes parallel computations of high-fidelity models, and (ii) the search scheme measures the utility of new design evaluations only at the immediate next iteration. This paper proposes a Non-Myopic Multipoint Multifidelity Bayesian Optimization (NM3-BO) algorithm to sensitively accelerate MDO overcoming the limitations of standard methods. NM3-BO selects a batch of promising design configurations to be evaluated in parallel, and quantifies the expected long-term improvement of these designs at future steps of the optimization. Our learning scheme leverages an original acquisition function based on the combination of a two-step lookahead policy and a local penalization strategy to measure the future utility achieved evaluating multiple design configurations simultaneously. We observe that the proposed framework permits to sensitively accelerate the MDO of a space vehicle and outperforms popular algorithms.

Optimization is becoming essential in science and engineering to empower the performance and sustainability of complex systems toward global challenges, such as green development and climate change. In many real-world applications, the optimization of advanced technologies involves multiple scientific disciplines characterized by complex relationships and couplings difficult to be tackled. Multidisciplinary design optimization (MDO) relates to the development of computational methodologies for the design and optimization of complex systems taking into account the interactions of multiple disciplines[1,2]. Those interactions can span different strength and width of the cross disciplinary couplings[3,4]. MDO approaches have been applied to the design and optimization of a broad range of engineering systems including aircraft[5], spacecrafts[6], launch vehicles[7], buildings[8], electric automobiles[9], ships[10], energy systems[11] and robots[12]. One of the major challenges addressed by MDO is represented by the possibility to use expensive high-fidelity disciplinary models, such as the ones that are given by large scale systems of equations for the numerical solutions of partially differential equations, directly in the simulation-based optimization process[13,14]. Indeed, the search of optimal design solutions would benefit from accurate representations of the system behaviour and physics including the couplings across the multiple disciplines. However, considering all those interactions and complex couplings would typically demand for large amount of evaluations of the high-fidelity disciplinary representations which would result in prohibitive computational costs for the overall MDO procedure.

To address these challenges, MDO literature proposed a variety of solutions that rely on the use of low-fidelity models to reduce the computational burden and complexity associated with disciplinary analysis and save computing resources. As discussed by Peherstorfer et al.[15], low-fidelity disciplinary solvers range from simplified models directly derived from the high-fidelity counterpart using expert knowledge[16], to projection-based models that identify a low-fidelity subspace retaining the essential features of the system[17], and to surrogate-based models where the input-output relationships of disciplines are derived from observations of the high-fidelity model[18]. Even if the evaluation of these low-fidelity representations could be sensitively reduced in computational cost, the former simplified models might not be adequate to depict complex non-linear phenomena that frequently characterize the disciplinary domain, while the latter projection-based and surrogate-based models might require a large amount of costly high-fidelity data for their construction.

[1]Departement of Mechanical and Aerospace Engineering, Politecnico di Torino, 10129 Turin, Italy. [2]Department of Aeronautics, Imperial College London, London SW7 2AZ, UK. ✉email: francesco.difiore@polito.it

1

Multifidelity methods acknowledge the opportunity offered by low-fidelity representations and offer approaches to address the research gap of including expensive high-fidelity disciplinary analysis into the MDO process[19]. Multifidelity methods combines data extracted from a library of disciplinary models that can be hierarchically ordered according to accuracy and computational cost[15,20]. The availability of multiple levels of fidelity can be exploited to support the search procedure through a principled elicitation of information: fast low-fidelity models are used to massively explore different design configurations, and expensive high-fidelity models sparingly refine the solution of the MDO problem. Multifidelity methods have been successfully applied to a variety of MDO applications ranging from aircraft[21,22] and space vehicles[23,24] to ships[25] and unmanned underwater vehicles[26], from electric[27] and hybrid[28] vehicles to green energy technologies[29,30]. In most cases, the complexity of the disciplinary analysis and couplings discourages the use of gradient-based optimization strategies: the computation of the derivatives might demand for massive high-fidelity data and increase the overall computational burden. Therefore, the MDO procedure commonly relies on a black-box approach where the disciplinary analyses are regarded as a pure input/output relationship, whose information about the mathematical properties and derivatives are not available.

Multifidelity Bayesian optimization (MFBO) provides a computational framework for black-box optimization and leverages disciplinary solvers at different levels of fidelity to accelerate the identification of promising design solutions[31–33]. MFBO realizes an adaptive sampling scheme based on a multifidelity acquisition function that targets the design improvement with a continuous trade-off between optimization performance and computational cost. Several MFBO approaches have been proposed to address a variety of design problems. Charayron et al.[34] applied an original multifidelity Bayesian framework for the MDO of a full electric drone accounting for a long-range surveillance mission. Serani et al.[35] proposed a MFBO based on stochastic radial-basis functions surrogate applied to the design optimization of a destroyer-type vessel. Reisenthel et al.[36] adopted an MFBO framework for the aeroelastic design optimization of a UAV wing. Tran et al.[37] implemented MFBO for the chemical design optimization of atomistic materials to identify the optimal bulk modulus.

Most state-of-the-art MFBO algorithms are greedy and sequential in nature since the multifidelity acquisition function (i) quantifies only how the design and level of fidelity selected at the current iteration affect the immediate next step, and (ii) selects only a single combination of design variables and level of fidelity to be evaluated at the next iteration. In particular, (i) the popular greedy approach precludes greater informative gains that can be acquired through the measure of the long-term reward obtained at future steps of the optimization, and might preclude superior accelerations of the MDO procedure; (ii) the sequential search might not be computationally efficient for MDO problems where the simulation of complex interdisciplinary relationships demands for a huge amount of high-fidelity data and associated computational expense.

To address these gaps, we propose a non-myopic multipoint multifidelity Bayesian optimization (NM3-BO) framework that aims to overcome both greedy and sequential limitations of standard MFBO methodologies. Our original multifidelity search defines an optimal sequence of decisions to (i) maximize the long-term reward as the improvement of the optimal design solution achieved at future steps of the optimization, and (ii) select a batch of design configurations and levels of fidelity to be evaluated simultaneously. To overcome the shortcomings of standard MFBO, our non-myopic multipoint and multifidelity learning scheme (i) is derived formalizing MFBO as a dynamic system under uncertainty addressed through a dynamic programming technique, (ii) defines an optimal policy as a sequence of decisions to maximize the two steps ahead utility obtained evaluating a design with a specific disciplinary model, and (iii) uses a local penalization strategy to enable multiple decisions as a batch of paired designs and levels of fidelity to query in parallel.

The performance of the proposed NM3-BO are illustrated and discussed in comparison with standard MFBO frameworks for the multidisciplinary design optimization of a re-entry vehicle. We adopt the Multidisciplinary Feasible (MDF) formulation to formalize the space-vehicle MDO problem, and capture the multidisciplinary nature of the system considering the contributions of the propulsion system, re-entry descend trajectory, aerothermodynamic effects, and thermo-structural interactions. In particular, the specific MDF architecture considers the coupling between the trajectory and the aerothermodynamic disciplines through the aerodynamic coefficients, and the coupling between the thermo-structural and aerothermodynamic disciplines through the thermal protection system wall temperature. This design problem is specifically selected to exemplify the marked cross-disciplinary scenario and strong couplings between disciplines that can be traced in the vast majority of MDO problems in science and engineering.

Methodology

The non-myopic multipoint multifidelity Bayesian optimization (NM3-BO) framework formalizes an adaptive sampling scheme that measures the long-term utility of a batch of design configurations evaluated simultaneously. Our formulation regards MFBO as a decision making problem affected by uncertainty: the decision task relates to the selection of promising design configurations – combination of design parameters – and the associated levels of fidelity to query; the uncertainty elements relate to the black-box nature of the objective function and the probabilistic prediction of the surrogate model.

In the following, the MDO problem setup and the BO single-fidelity and multifidelity frameworks are firstly introduced to provide an overview of the core background of our work. Then, we formalize the optimal policy for MFBO and illustrate how to robustly approximate it through a Monte Carlo technique. In addition, we propose a multiple decision making strategy to enable parallel computations of a batch of designs and associated levels of fidelity. Finally, the NM3-BO algorithmic framework is presented and discussed.

MDO problem setup

This work considers the general formulation of the MDO problem as follows[4]:

$$\begin{aligned}
\text{minimize} \quad & f(\mathbf{x}) \\
\text{with respect to} \quad & \mathbf{x} \\
\text{subject to} \quad & c(\mathbf{x}) \geq 0 \\
& \mathscr{R}_i(\mathbf{x}) = 0 \quad \text{for } i = 1, ..., D
\end{aligned} \tag{1}$$

where the goal is to identify a set of design variables $\mathbf{x}$ that minimizes an objective function f subject to design constraints c, and the solution of governing equations in residual form $\mathscr{R}_i(\mathbf{x})$ for each i-th discipline.

The disciplinary analyses are usually performed through black-box simulations: computer codes operate independently and define relationships between inputs and outputs, and hide the procedure associated with their computation. High-fidelity disciplinary analyses involve the numerical solution of governing PDEs through expensive computational procedures, such as Computational Fluid Dynamics (CFD) techniques for the numerical solution of Navier-Stokes equations.

BO from single-fidelity to multifidelity

Bayesian optimization (BO) is an efficient computational strategy to address the MDO of expensive black-box objective functions[38,39]. To solve Eq. (1), BO uses two key components: a surrogate model of the objective function $f(\mathbf{x})$ and an acquisition function computed on the surrogate of f. The maximization of the acquisition function permits to select the most promising design configuration $\mathbf{x}'$ to query. BO uses the observed value $f(\mathbf{x}')$ to update the surrogate model and the process iterates until a certain termination criteria is reached. Popular formulations of Bayesian optimization and acquisition functions are overviewed by Frazier et al.[39] and Shahriari et al.[38].

The most widely used acquisition functions determine a greedy and sequential adaptive sampling scheme that considers only the immediate effect of evaluating the objective function for a single design, and do not consider the potential gains introduced in future evaluations. To address this type of greedy limitation, BO has been formalized as a partially observable Markov decision process[40], and several works[41–44] provide solutions to this process and formalize non-myopic multifidelity acquisition functions. In addition, multipoint formulations of the BO framework have been proposed to evaluate in parallel multiple designs with a single level of fidelity of the disciplinary model[45–47]. However, the combination of non-myopic and multipoint formulations in literature are conceived exclusively for a single-fidelity framework only: the optimization process relies on the responses of disciplinary models at one single fixed level of fidelity. In the MDO context, this single fidelity approach could hinder the expensive high-fidelity disciplinary models to be interrogated directly during the search, which otherwise would result in prohibitive computational costs. In addition, the computational cost becomes unmanageable as it scales exponentially when the disciplinary couplings are also considered during the process: the identification of an optimal design that satisfies all the interactions and couplings across the disciplines would require massive evaluations of high-fidelity disciplinary models with the associated growth of the computational demand.

In many applications, scientists and engineers might rely on different disciplinary models of the objective function and constraints with different degrees of accuracy and associated demand for computational resources. Multifidelity Bayesian Optimization (MFBO) combines disciplinary responses from this library of models $\{f^{(1)}, f^{(2)}, ..., f^{(L)}\}$ hierarchically ordered according to the level of fidelity $l = 1, ..., L$ to accelerate the solution of the MDO problem (Eq. 1). In this setting, the surrogate model synthesizes the outcomes computed with multiple models into a unique surrogate, and the acquisition function defines an adaptive sampling scheme that identifies the design configuration and the associated level of fidelity to query at each iteration.

Multifidelity Gaussian process
We adopt the Multifidelity Gaussian process (MFGP) as the surrogate model of the objective function given the successful application over a variety of methodologies and applications[47–49]. MFGP is formalized extending the Gaussian process[50] (GP) formulation to multiple levels of fidelity, and provides a prediction of the objective function over the design space $\mathscr{X}$ through the mean function $\mu^{(l)}(\mathbf{x}) = \mathbb{E}[f^{(l)}(\mathbf{x})]$, and the associated uncertainty of the prediction $\sigma^{2(l)}(\mathbf{x}) = \mathbb{E}[(f^{(l)}(\mathbf{x}) - \mu(\mathbf{x}))(f^{(l)}(\mathbf{x}') - \mu(\mathbf{x}'))]$ through a covariance function.

MFGP relies on an autoregressive scheme to synthesize responses from models at different levels of fidelity[51]:

$$f^{(l)} = \rho_{l-1} f^{(l-1)}(\mathbf{x}) + \delta^{(l)}(\mathbf{x}) \quad l = 2, ..., L \tag{2}$$

where ρ_{l-1} is a regression parameter that scales successive representations $f^{(l)}$ and $f^{(l-1)}$, and $\delta^{(l)}$ models the discrepancy between $f^{(l)}$ and the scaled model $\rho_{l-1} f^{(l-1)}$. This discrepancy term $\delta^{(l)}$ is assumed to be a Gaussian process characterized by mean function $\upsilon(\mathbf{x})^T \beta^{(l)}$ and covariance $\kappa^{(l)}(\mathbf{x}, \mathbf{x}')$, where υ is the vector of regression functions and $\beta^{(l)}$ are the associated weighting coefficients. We use the Gaussian correlation model as the covariance function of the MFGP surrogate:

$$\kappa(\mathbf{x}, \mathbf{x}') = \varsigma_l^2 \exp\left\{ -\sum_{m=1}^{M} \varpi_l^m (\mathbf{x}_m - \mathbf{x}'_m)^2 \right\} \tag{3}$$

where $\varpi = (\varpi_l^1, \varpi_l^2, ..., \varpi_l^M)$ is the roughness parameter, and ς_l^2 is the process variance of the l-th level of fidelity.

Thus, the posterior mean $\mu^{(l)}(\mathbf{x})$ and variance $\sigma^{2(l)}(\mathbf{x})$ of the objective are formalized through the covariance matrix $\mathbf{K}(i, j) = \left[\kappa\left((\mathbf{x}_i, l_i), (\mathbf{x}_j, l_j) \right) + \gamma_{ij} \sigma_\epsilon(\mathbf{x}_i) \right]$ and $\kappa_N^{(l)}(\mathbf{x}) = [\kappa((\mathbf{x}_i, l_i), (\mathbf{x}, l))]$:

3

$$\mu^{(l)}(\mathbf{x}) = \kappa_N^{(l)}(\mathbf{x})^T \mathbf{K}^{-1} \mathbf{y} \tag{4}$$

$$\sigma^{2(l)}(\mathbf{x}) = \kappa((\mathbf{x}, l), (\mathbf{x}, l)) - \kappa_N^{(l)}(\mathbf{x})^T \mathbf{K}^{-1} \kappa_N^{(l)}(\mathbf{x}) \tag{5}$$

where γ_{ij} is the Kronecker delta function. The hyperparameters $(\rho, \beta, \varpi, \varsigma)$ of the multifidelity Gaussian process surrogate model are estimated through maximum likelihood estimation methodology[52].

Multifidelity expected improvement
The multifidelity acquisition function U quantifies the utility of new design configurations with a trade-off between cost ad accuracy of the associated model to query. In this work, we will build our method onto the Multifidelity Expected Improvement (MFEI) acquisition function[31], given the popularity of this formulation across different science and engineering communities and the variety of adoptions documented in literature[36,53]. The MFEI is formulated as follows:

$$U_{MFEI}(\mathbf{x}, l) = U_{EI}(\mathbf{x})\alpha_1(\mathbf{x}, l)\alpha_2(\mathbf{x}, l)\alpha_3(l) \tag{6}$$

where U_{EI} is the expected improvement acquisition function computed with the highest level of fidelity. The utility functions α_1, α_2 and α_3 are defined as follows:

$$\alpha_1(\mathbf{x}, l) = corr[f^{(l)}(\mathbf{x}), f^{(L)}(\mathbf{x})] \tag{7}$$

$$\alpha_2(\mathbf{x}, l) = 1 - \frac{\sigma_\epsilon}{\sqrt{\sigma^{2(l)}(\mathbf{x}) - \sigma_\epsilon^2}} \tag{8}$$

$$\alpha_3(l) = \frac{\lambda^{(L)}}{\lambda^{(l)}}. \tag{9}$$

$\alpha_1(\mathbf{x}, l)$ reflects the decrease of accuracy associated with lower-fidelities quantified as the correlation between the l-th fidelity outcome $f^{(l)}(\mathbf{x})$ and the high-fidelity $f^{(L)}(\mathbf{x})$ disciplinary analysis. $\alpha_2(\mathbf{x}, l)$ brings awareness about the stochastic nature of the objective function through the measurement noise σ_ϵ, and considers the reduction of the uncertainty associated with the evaluation of design solutions with the l-th model. $\alpha_3(l)$ includes the computational cost $\lambda^{(l)}$ of the l-th disciplinary analysis in the sampling procedure, and balances the accuracy with the resources entailed for the objective evaluation.

Alternative formulations of the multifidelity acquisition function have been proposed in literature, such as the Multifidelity Max-Value Entropy Search (MFMES)[32] and the Multifidelity Probability of Improvement[33] (MFPI). Both these acquisition functions define a sampling scheme sequential and greedy, where the design improvement is measured as the probability of lower the objective function according to the surrogate prediction (MFPI) or as the maximum decrease of differential entropy (MFMES). In this work, the proposed NM3-BO framework is compared against all those formulations of the acquisition functions for the MDO problem of a re-entry vehicle.

Optimal decision making process over the next two-step ahead

The multifidelity Bayesian optimization is regarded as a decision making problem and formulated as a Markov Decision Process (MDP). MDPs are discrete-time stochastic control processes that allow to model the sequential decision making process of a dynamic system under uncertainty[54,55]. The main methodological elements of MDPs are: (i) the Markov chains model the transitions of the dynamic system from the initial state to future states, (ii) a decision-making model makes a decision at each state transition of the system, and (iii) an utility function quantifies the reward achieved by a certain decision with reference to the given goal. The objective of MDPs is to identify the optimal set of decision to efficiently reach the given goal over time. Following the perspective of MFBO as a dynamic system under uncertainty, (i) the multifidelity Gaussian process predicts the transitions of the MFBO system to future states given the initial conditions, (ii) the MFBO system evolves making a decision at each transition on the next design and level of fidelity to be evaluated, (iii) the reward of each decision is measured through the multifidelity acquisition function that evaluates the benefits of each decision through a trade-off between obtaining utility from the current state and altering the opportunities to obtain utility in the future.

The solution of the MFBO Markov Decision Process requires a procedure to perform statistical inference on the system behaviour, and depict the transitions of the system from one state to the other after exploring every possible decision. We employ the dynamic programming[56,57] approach to solve this specific MDP. Dynamic programming (DP) partitions the optimization procedure in simpler sub-problems defined in a recursive way across several transitions. This permits to define the optimal policy as a sequence of rules that maximizes the cumulative reward achieved making decisions at future iterations. In the following, we formulate the optimization policy based on an original multifidelity acquisition function to measure the utility of decisions expected over future steps of the optimization procedure.

Let us consider the MFBO dynamic system fully characterized at each time step t by a state $s_t \in \mathscr{S}_t$, where $\mathscr{S}_t$ denotes a set of states that represent all the possible configurations of the system at each time step. Following the dynamic programming approach, we consider any future iteration of the optimization process as time steps $\{t, ..., T\}$ of the MFBO dynamics. For a generic time step t, the multifidelity Gaussian process conditioned on $\mathscr{D}_t = \{\mathbf{x}_n, y^{(l_n)}, l_n\}_{n=1}^N$ determines the posterior distribution of the objective function, where $\mathbf{x}$ is the design point and $y^{(l)}$ is the noisy observation of the objective function at the l-th level of fidelity. Thus, MFBO is fully

characterized by a state s_t defined through the dataset $\mathscr{D}_t \in \mathscr{S}_t$. Based on the simulated scenario defined by the surrogate model, MFBO makes an action $a_t = \{\mathbf{x}_{t+1}, l_{t+1}\}$ that activates the dynamic of the system, and defines the next design $\mathbf{x}_{t+1} \in \mathscr{X}$ and associated level of fidelity l_{t+1} to evaluate. This action is taken under an optimization policy $\pi_t : \mathscr{S}_t \rightarrow \mathscr{X}$ that maximizes the utility achieved in the future steps by mapping the state s_t to the action $a_t = \pi_t(s_t)$.

At the new time step $t + 1$, the value of $f^{(l_{t+1})}(\mathbf{x}_{t+1})$ is unknown and requires the evaluation of the objective function at the l_{t+1}-th level of fidelity. We can model this value as an uncertain quantity through the posterior distribution of the surrogate model conditioned on $\mathscr{D}_t$ at $\mathbf{x}_{t+1}$ and level of fidelity l_{t+1}. This simulated outcome is defined as a random disturbance normally distributed $w_{t+1}^{(l)} \sim \mathcal{N}(\mu_t^{(l)}(\mathbf{x}_{t+1}), \sigma_t^{2(l)}(\mathbf{x}_{t+1}))$ specified by the mean $\mu_t^{(l)}$ and variance $\sigma_t^{2(l)}$ of the multifidelity Gaussian process. Once the outcome is simulated, the system transitions to the new state s_{t+1} following its dynamics $\mathscr{F}$, which corresponds to the augmented dataset $\mathscr{D}_{t+1} = \mathscr{D}_t \cup \{\mathbf{x}_{t+1}, y^{(l_{t+1})}, l_{t+1}\}$:

$$\mathscr{D}_{t+1} = \mathscr{F}(\mathbf{x}_{t+1}, y^{(l_{t+1})}, l_{t+1}, \mathscr{D}_t) \tag{10}$$

At this point, we need to define a specific reward function for MFBO that measures the utility obtained from a simulated outcome $w_{t+1}^{(l)}$ when the action a_t is applied to the state s_t. This reward function can be formulated as the reduction of the objective function achieved at the time step $t + 1$ with respect to t:

$$r_t(\mathbf{x}_{t+1}, y^{(l_{t+1})}, l_{t+1}, \mathscr{D}_t) = (f_t^{*(L)} - f_{t+1}^{(L)})^+ \tag{11}$$

where $f_{t+1}^{(L)} = w_{t+1}^{(L)}$, and $f_t^{*(L)}$ is the minimum value of the objective function at the highest level of fidelity observed up to t. Thus, we follow the DP recursive strategy and define the expected reward at the generic time step t:

$$J_t^{\boldsymbol{\pi}}(\mathbf{x}_{t+1}, l_{t+1}, \mathscr{D}_t) = \mathbb{E}[r_t(\mathbf{x}_{t+1}, y^{(l_{t+1})}, l_{t+1}, \mathscr{D}_t) + J_{t+1}(\mathscr{F}(\mathbf{x}_{t+1}, y^{(l_{t+1})}, l_{t+1}, \mathscr{D}_t))] \tag{12}$$

where $\mathbb{E}[r_t(\cdot)] = U_{MFEI}(\mathbf{x}_{t+1}, l_{t+1})$ is the multifidelity expected improvement (Eq. 6), and $J_{t+1}(\mathscr{F}(\cdot))$ is the long-term expected reward. We formulate the two-step lookahead multifidelity acquisition function through an optimal policy $\boldsymbol{\pi}^*$ that maximizes the cumulative expected reward over two-step ahead of a pair of design configurations $\mathbf{x}_{t+2}$ and level of fidelity l_{t+2}:

$$U_t^{\boldsymbol{\pi}^*}(\mathbf{x}_{t+2}, l_{t+2}, \mathscr{D}_{t+1}) = U_{MFEI}(\mathbf{x}_{t+1}, l_{t+1}) + \mathbb{E}[\max(U_{MFEI}(\mathbf{x}_{t+2}, l_{t+2}))] \tag{13}$$

where we define the long term reward $J_{t+1} = \max(U_{MFEI}(\mathbf{x}_{t+2}, l_{t+2}))$ as the maximum of the multifidelity expected improvement conditioned on the dataset $\mathscr{D}_{t+1}$.

Robust approximation of the optimal decision making process

The evaluation of $U_t^{\boldsymbol{\pi}^*}$ (Eq. 13) requires the solution of nested expectations and maximizations computationally intractable. We adopt the Monte Carlo approach to avoid nested computations, and provide a robust estimate of the two-step lookahead multifidelity utility function. Let us formulate the observation of the objective function at the new design configuration $\mathbf{x}_{t+1}$ and level of fidelity l_{t+1} using the reparameterization strategy proposed by Wilson et al.[58]:

$$f^{(l)}(\mathbf{x}_{t+1}) = \mu_t^{(l)} + \mathbf{C}_t^{(l)}(\mathbf{x}_{t+1})Z \tag{14}$$

where $\mathbf{C}_t^{(l)}$ is the Cholesky decomposition of the covariance matrix $\mathbf{K}_t$, and Z is an independent standard normal random variable. We use Eq. (14) to compute the mean and variance of the multifidelity Gaussian process at $t + 1$ for the generic design configuration $\mathbf{x}$:

$$\mu_{t+1}^{(l)}(\mathbf{x}) = \mu_t^{(l)}(\mathbf{x}) + \mathbf{H}_t^{(l)}(\mathbf{x})Z \tag{15}$$

$$\sigma_{t+1}^{(l)}(\mathbf{x}) = \sigma_t^{(l)}(\mathbf{x}) - \mathbf{H}_t^{(l)}(\mathbf{x})\mathbf{H}_t^{(l)}(\mathbf{x})^T \tag{16}$$

where $\mathbf{H}_t^{(l)}(\mathbf{x}) = \kappa_t^{(l)}(\mathbf{x})\mathbf{C}_t^{(l)-1}(\mathbf{x})$.

The expectation term of our acquisition function in Eq. (13) is approximated through the prediction $\mu_{t+1}^{(l)}$ and associated uncertainty $\sigma_{t+1}^{(l)}$ of the multifidelity Gaussian process. This permits to estimate the multifidelity expected improvement at $t + 2$ as follows:

$$U_{MFEI}(\mathbf{x}_{t+2}, l_{t+2}) \sim \widehat{U}_{MFEI}(\mathbf{x}_{t+2}, l_{t+2}, Z) \tag{17}$$

At this point, the two-step lookahead multifidelity acquisition function is evaluated sampling the random variable Z and averaging many realizations of $\widehat{U}_t^{\boldsymbol{\pi}^*}$ through Eq. (17):

$$U_t^{\boldsymbol{\pi}^*}(\mathbf{x}_{t+2}, l_{t+2}, \mathscr{D}_{t+1}) \sim \mathbb{E}[\widehat{U}_t^{\boldsymbol{\pi}^*}(\mathbf{x}_{t+2}, l_{t+2}, Z)] \tag{18}$$

This Monte Carlo approach demands for a significant number of realizations of $\widehat{U}_t^{\boldsymbol{\pi}^*}$ – in the order of thousands – to provide a robust approximation of $U_t^{\boldsymbol{\pi}^*}$. However, we emphasize that the evaluation of Eqs. (15) and (16) is inexpensive and requires contained resources. This results in a total computational cost required by the Monte

Carlo procedure that is negligible if compared to the cost associated with the evaluation of the objective function through high-fidelity disciplinary analyses in complex MDO problems.

Enabling multiple decisions

Equation (13) provides a non-myopic multifidelity sampling scheme that sequentially selects the design configuration and the respective model to maximize the cumulative reward over two steps ahead. Di Fiore and Mainini[59] provided evidence that this strategy permits to achieve remarkable results over a variety of benchmark analytical problems with a reduction of the number of high fidelity interrogations. However, complex multidisciplinary design optimization problems in the form of Eq. (1) open major challenges for the intrinsic demand to scale: the required accurate evaluations of the objective function can dramatically upscale during the search for improved design solutions.

To address this complex multiphysics optimization scenario, we extend the optimal policy π^* to enable multiple decisions. This permits to define a decision making procedure that iteratively selects a batch of promising design points and associated levels of fidelity $\mathscr{B}_i^{n_b} = [(\mathbf{x}_{i,1}, l_{i,1}), ..., (\mathbf{x}_{i,n_b}, l_{i,n_b})]$ while improving the design solutions over future iterations. The potential of multipoint formulations has been illustrated by[60] for a greedy single-fidelity Bayesian framework, and motivates our proposal of a multipoint sampling strategy for the non-myopic MFBO in multidisciplinary settings. Accordingly, the formulation of our acquisition function $U_t^{\pi^*}$ is modified through a local penalization maximization as follows:

$$\mathbf{x}_{i,k}, l_{i,k} = \arg\max \left[U_t^{\pi^*}(\mathbf{x}_{t+2}, l_{t+2}, \mathscr{D}_{t+1}) \prod_{j=1}^{k-1} \psi(\mathbf{x}, \mathbf{x}_j) \right] \tag{19}$$

where $(\mathbf{x}_{i,k}, l_{i,k}) \in \mathscr{B}_i^{n_b}$, and ψ is the local penalty function which quantifies the probability that a design point $\mathbf{x}$ is a potential minimum not belonging to the hypersphere $\{\mathbf{x} \in \mathscr{X} : ||\mathbf{x}_j - \mathbf{x}|| \leq (\hat{f}^* - f^L(\mathbf{x}_j))/L\}$:

$$\psi(\mathbf{x}, \mathbf{x}_j) = \frac{1}{2} erfc \left[\frac{1}{\sqrt{2\sigma^{2(L)}(\mathbf{x}_j)}} \left(L||\mathbf{x}_j - \mathbf{x}|| - \hat{f}^* + \mu^{(L)}(\mathbf{x}_j) \right) \right] \tag{20}$$

where $erfc$ is the complementary error function, $\hat{f}^* = \min_{\mathbf{x} \in \mathscr{X}} \mu^{(L)}(\mathbf{x})$ is the minimum predicted by the surrogate model, and $L = \max_{\mathbf{x} \in \mathscr{X}} ||\mu_\nabla^{(L)}(\mathbf{x})||$ is the Gaussian process Lipschitz constant[60] defined as the maximum of the surrogate gradient. The rationale behind the formulation of ψ is to locally penalize the acquisition function and create exclusion zones whose amplitude is determined by the Lipschitz constant L. This results in small sized exclusion zones when the values of the mean function $\mu^{(L)}$ are closer to the predicted minimum $\hat{f}^*$. In contrast, regions of the design space where the mean function $\mu^{(L)}$ is far from the predicted minimum $\hat{f}^*$ produce larger exclusion zones. This penalization strategy emulates a sampling procedure over multiple iterations that would have been achieved by a sequential scheme if the previous evaluations of the designs in the batch were at disposal of the learner.

In addition, we provide an adaptive batch size formulation that determines the number of designs in a batch $n_b(i) = 1 + \beta/(\sqrt{2}i)$ as a function of the optimization iterations i and the initial batch size β. This strategy targets the efficient use of computational resources: the number of design evaluations increases at the beginning of the optimization to improve the knowledge of the objective function distribution over the domain, and is progressively reduced to catalyze the resources toward the analyses of optimal design solutions.

NM3-BO algorithm

Algorithm 1 illustrates the numerical implementation of our non-myopic multipoint multifidelity Bayesian optimization (NM3-BO) scheme. The initialization of the computations requires the definition of the design space $\mathscr{X} \in \mathbb{R}^d$ according to the number of variables d that characterize the design configuration $\mathbf{x}$, together with the library of models of the objective function $\{f^{(1)}, f^{(2)}, ..., f^{(L)}\}$. In addition, the multifidelity Gaussian process prior $GP(0, \kappa^{(l)}(\mathbf{x}, \mathbf{x}'))$ is defined for each level of fidelity, and represents the prior belief about the distribution of the objective function over the design space. Before the start of the iterative process, an initial subset of N_0 design configurations $\{\mathbf{x}_n\}_{n=1}^{N_0}$ and corresponding levels of fidelity $\{l_n\}_{n=1}^{N_0}$ are sampled through any design of experiments technique, and are used to compute the first observations of the objective function $\{y^{(l_n)}\}_{n=1}^{N_0}$. These data are collected into the dataset $\mathscr{D}_0 = \{\mathbf{x}_n, y^{(l_n)}, l_n\}_{n=1}^{N_0}$ which defines the initial state of the MFBO system, and induces the posterior distribution under the prior specified by the mean $\mu^{(l)}$ and variance $\sigma^{2(l)}$.

For a generic iteration i of the NM3-BO algorithmic flow, the surrogate model is updated through the observations of the objective function $f^{(l_{i,k})}(\mathbf{x}_{i,k})$ at each $n_b(i)$ pair of design configurations $\mathbf{x}_{i,k}$ and levels of fidelity $l_{i,k}$ that constitutes the batch $\mathscr{B}_i^{n_b(i)}$ selected at the previous iteration $i-1$. This represents the new state of the MFBO system $\mathscr{D}_i = \mathscr{D}_{i-1} \cup \{\mathbf{x}_{i,k}, f^{(l_{i,k})}(\mathbf{x}_{i,k}), l_{i,k}\}_{k=1}^{n_b(i)}$. At this stage, the algorithm selects the next design configurations and levels of fidelity $\mathscr{B}_{i+1}^{n_b(i+1)}$ to evaluate through the computation and maximization of our acquisition function $U_t^{\pi^*}$. Let now indicate with $t = i$ the current step of the optimization and with $t+1$ and $t+2$ the first and

---

Algorithm 1. NM3-BO: non-myopic multipoint multifidelity Bayesian optimization.

Input: Design space $\mathcal{X} \in \mathbb{R}^d$, multifidelity models of the objective function $f^{(l)}(\boldsymbol{x})$ and the multifidelity Gaussian process prior $GP(0, \boldsymbol{\kappa}^{(l)}(\boldsymbol{x}, \boldsymbol{x}'))$

Output: $f^* = \min f(\boldsymbol{x})$

 1:   $\mathcal{D}_0 \leftarrow \{\boldsymbol{x}_n, y^{(l_n)}, l_n\}_{n=1}^{N_0}$ collect initial observations
 2:   $\mu_0^{(l)}, \sigma_0^{2(l)} \leftarrow$ learn the initial multifidelity GP
 3:   $i \leftarrow 1$
 4:   **repeat**
 5:      Load the k pairs of $\boldsymbol{x}_{i,k}$ and associated $l_{i,k}$ from the batch $\mathcal{B}_{i-1}^{n_b(i-1)}$
 6:      Evaluate $f^{(l_{i,k})}(\boldsymbol{x}_{i,k}) \ \forall n_b(i-1)$
 7:      $\mathcal{D}_i = \mathcal{D}_{i-1} \cup \{\boldsymbol{x}_{i,k}, f^{(l_{i,k})}(\boldsymbol{x}_{i,k}), l_{i,k}\}_{k=1}^{n_b(i)}$
 8:      $\mu_i^{(l)}, \sigma_i^{2(l)} \leftarrow$ update the multifidelity Gaussian process
 9:      $t \leftarrow i$
10:      Compute $U_{MFEI}(\boldsymbol{x}_{t+1}, l_{t+1})$
11:      **for** $j \leftarrow 1, N_{MC}$ **do**
12:         $Z_j \leftarrow \mathcal{N}(0,1)$
13:         $\mu_{t+1}^{(l)}, \sigma_{t+1}^{2(l)} \leftarrow$ estimate the multifidelity GP
14:         Compute $\widehat{U}_{MFEI}(\boldsymbol{x}_{t+2}, l_{t+2}, Z_j)$
15:      **end for**
16:      **return** $\widehat{U}_t^{\pi^*} = \{\widehat{U}_{z_j}^{\pi^*}\}_{j=1}^{N_{MC}}$
17:      $U_i^{\pi^*} = \mathbb{E}\left[\widehat{U}_t^{\pi^*}\right]$
18:      **for** $k \leftarrow 1, n_b(i)$ **do**
19:         $\boldsymbol{x}_{i,k}, l_{i,k} = \max\left[U_t^{\pi^*}(\boldsymbol{x}_{t+2}, l_{t+2}, \mathcal{D}_{t+1}) \prod_{j=1}^{k-1} \psi(\boldsymbol{x}, \boldsymbol{x}_j)\right]$
20:      **end for**
21:      **return** $\mathcal{B}_i^{n_b(i)} = [(\boldsymbol{x}_{i,1}, l_{i,1}), ..., (\boldsymbol{x}_{i,n_b(i)}, l_{i,n_b(i)})]$
22:      $i+1 \leftarrow i$
23:   **until** $B_i \leq B_{max}$
24:   **return** x^* that minimize $f(\boldsymbol{x})$ over $\mathcal{D}_i$

---

the second step ahead, respectively. The first element of $U_t^{\pi^*}$ is determined using the information extracted from the surrogate model updated at the current state of the system $\mathcal{D}_t = \mathcal{D}_i$. The second element requires our Monte Carlo technique to compute the nested expectation and maximization, and quantify the informative gains at future iterations. Accordingly, the algorithm samples independently a random variable Z_j normally distributed for the j-th Monte Carlo realization, and simulates the future optimization scenario through the estimate of the mean $\mu_{t+1}^{(l)}$ and variance $\sigma_{t+1}^{2(l)}$ of the surrogate model by the computation of Eqs. (15) and (16). This provides an estimate of the multifidelity acquisition function $\widehat{U}_{t_j}^{\pi^*}$ as the expectation taken over the realizations $\{\widehat{U}_t^{\pi^*}\}_{j=1}^{N_{MC}}$ (Eq. 18). Then, the penalized maximization of the acquisition function determines the next batch $\mathcal{B}_{i+1}^{n_b(i+1)}$ of design configurations and the levels of fidelity to be evaluated in parallel at the next iteration. This optimization procedure iterates until a maximum computational budget $B_i = B_{max}$ is reached, where B_i is the cumulative computational cost adopted until iteration i.

Space vehicle multidisciplinary design optimization
MDO problem setup
The design of a space re-entry vehicle is a multidisciplinary optimization problem that well carries the computational challenges associated with the design of complex engineering systems[61,62]. This paper uses this demanding MDO applications to demonstrate our NM3-BO scheme and discuss it in comparison with popular standard MFBO algorithms. The space vehicle MDO problem captures the multi-physics nature of the atmospheric re-entry and involves several disciplinary analyses, namely the contributions of the propulsion system, the re-entry descend trajectory, the aerothermodynamic effects that occurs during the descend path, and the thermo-structural interaction between the re-entry flow-field and the thermal protection system. Figure 1 illustrates the concept of operations of the re-entry mission. This involves several phases, namely a maneuver sequence to introduce a thrust component that shapes the re-entry trajectory, the heat peak along the descent caused by the hypersonic aerothermodynamic phenomena, and the deployment of the parachutes during the landing phase.

Figure 2 illustrates the design structure matrix[63] of the re-entry vehicle optimization problem. We adopt the multidisciplinary feasible architecture[4] to address the MDO problem through a single optimization procedure where the design variables and constraints are under the direct control of the optimizer. The disciplinary analyses flow follows the diagonal of the DSM, while the feed forward flows are represented on the upper triangle and the couplings between disciplines are reported on the lower side. The propulsion system is modeled according to the chemical rocket theory, and comprises primary and secondary chemical thrusters fueled by an hypergolic propellant. The trajectory solver models the descend trajectory as a bi-dimensional orbit propagated through the numerical integration of the non-linear re-entry planetary ordinary differential equations. The aerothermodynamic analysis consists of two disciplinary solvers at different levels of fidelity. The high-fidelity model simulates
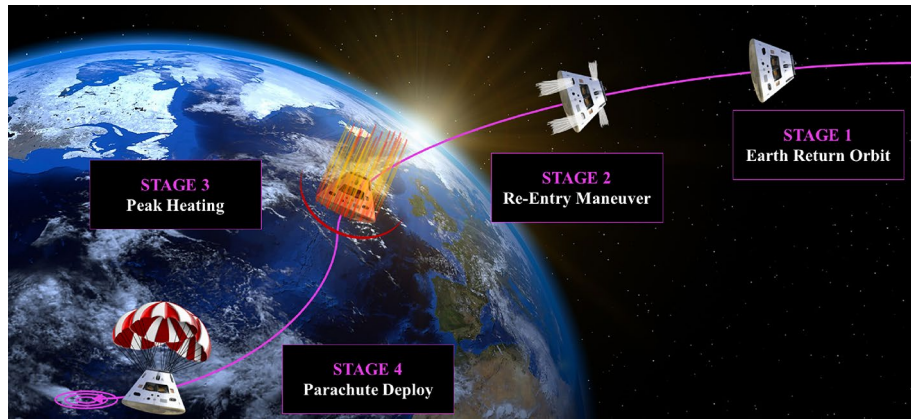
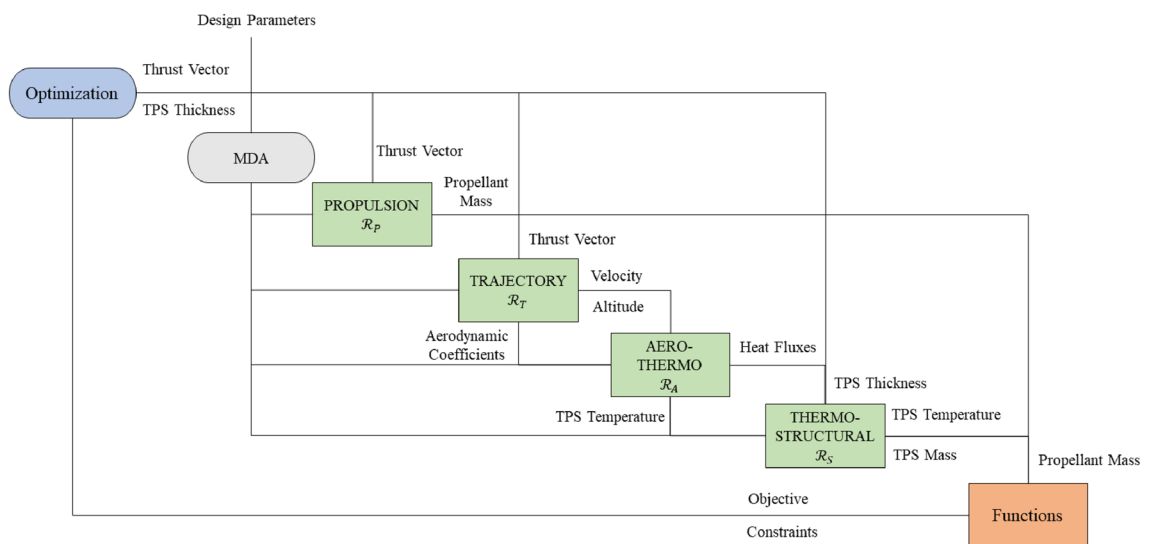**Figure 1.** Re-entry mission concept of operations.

Figure 2. Design structure matrix of the space vehicle MDO problem.

the full order aerothermodynamic physics through the numerical solution of the Reynolds-Averaged Navier-Stokes equations. The low-fidelity model uses the Oswatitsch Mach number independence principle jointly with the Tauber-Sutton and Sutton-Grave formulations to provide an approximated representation of the aerothermodynamic domain. The high-fidelity model requires hours of computation on an high performance computing cluster, while the low-fidelity analysis is three orders of magnitude faster on a standard computing platform. The thermo-structural analysis models the interaction between the flow-field and the structure of the Thermal Protection System (TPS) through the thermo-elastic equations. Further details about the disciplinary models are provided in the Supplementary Material of this manuscript consistent with what proposed in literature[64].

The design optimization problem targets the best design configuration $\mathbf{x} = [F_V, F_N, s_{TPS}]$ of the re-entry vehicle in terms of thrust capabilities $F = [F_V, F_N]$ and TPS structural thickness s_{TPS} that jointly minimizes the temperature T_{TPS} reached by the TPS frame, the overall structural mass m_{TPS} of the TPS, and the mass of propellant m_P burned during the re-entry maneuver. The multidisciplinary design optimization problem is formulated as follows:
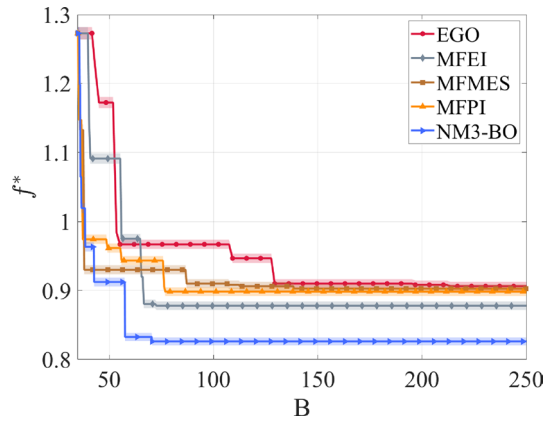
Figure 3. Statistics over 25 experiments of the minimum of the objective function f^* obtained with the competing algorithms.

$$\begin{aligned}
\text{minimize} \quad & f(\mathbf{x}) = 0.4\frac{m_{TPS}(\mathbf{x})}{m_{TPS0}} + 0.4\frac{T_{TPS}(\mathbf{x})}{T_{TPS0}} + 0.2\frac{m_P(\mathbf{x})}{m_{P0}} \\
\text{with respect to} \quad & \mathbf{x} = [F_V, F_N, s_{TPS}] \\
\text{subject to} \quad & 100km \leq h^*(\mathbf{x}) \leq 125km \\
& \mathscr{R}_P(\mathbf{x}) = 0 \\
& \mathscr{R}_T(\mathbf{x}) = 0 \\
& \mathscr{R}_A^{(l=1)}(\mathbf{x}) = 0 \\
& \mathscr{R}_A^{(l=2)}(\mathbf{x}) \leq 10^{-6} \\
& \mathscr{R}_S(\mathbf{x}) = 0
\end{aligned} \tag{21}$$

where the objectives are evaluated with reference to the baseline values for the TPS mass $m_{TPS0} = 700\ kg$ and temperature $T_{TPS0} = 1000\ K$, and for the mass of propellant $m_{P0} = 150\ kg$ derived from similar re-entry capsules[65]. The search is bounded by the move limits of the design space $\mathscr{X} = \mathscr{X}_{F_V} \times \mathscr{X}_{F_N} \times \mathscr{X}_{s_{TPS}}$, where the thrust capabilities tangential $\mathscr{X}_{F_V} = [29.2\ kN, 146\ kN]$ and normal $\mathscr{X}_{F_N} = [0.48\ kN, 2.4\ kN]$ to the trajectory are defined according to the propulsion system specifications, and the limits on the TPS thickness $\mathscr{X}_{s_{TPS}} = [0.03\ m, 0.1\ m]$ are imposed from expert knowledge. The MDO problem requires a specific range of altitudes h^* for the re-entry maneuver to simulate a real-world mission. Additional constraints include the feasibility of the physics-based models at each iteration of the optimization procedure, namely the complete resolution of the propulsion system model $\mathscr{R}_P(\mathbf{x}) = 0$, the trajectory model $\mathscr{R}_T = 0$, the low-fidelity aerothermodynamic model $\mathscr{R}_A^{(l=1)}(\mathbf{x}) = 0$, the high-fidelity aerothermodynamic model ensured reducing the computational residuals below $\mathscr{R}_A^{(l=2)}(\mathbf{x}) \leq 10^{-6}$, and the thermo-structural model $\mathscr{R}_S(\mathbf{x}) = 0$.

Results and discussion

The capabilities of the proposed non-myopic multipoint multifidelity algorithm NM3-BO are compared with standard Multifidelity Bayesian Optimization frameworks. All those MFBO algorithms rely on the multifidelity Gaussian process surrogate model and implement different formulations of the acquisition function, including the Multifidelity Expected Improvement[31] (MFEI), Multifidelity Max-value Entropy Search[32] (MFMES), and Multifidelity Probability of Improvement[33] (MFPI). In addition, we report the outcomes achieved with the Efficient Global Optimization[66] (EGO) algorithm using only high-fidelity queries to provide a comparison with a popular single fidelity Bayesian optimization methodology.

The performance of the competing algorithms are evaluated in terms of the minimum of the objective function $f^*(\mathbf{x}^*) = \min_{\mathbf{x} \in \mathscr{X}} f(\mathbf{x})$ as a function of the computational budget $B = \sum \lambda_i^{(l)}$ at each iteration i of the optimization procedure. The computational costs for the aerothermodynamic analyses are imposed at $\lambda^{(2)} = 1$ for the high-fidelity model and $\lambda^{(1)} = 0.001$ for the low-fidelity model; these specific values reflect the time required to complete the aerothermodynamic simulation adopting either the CFD solver or the low-fidelity formulations. We consider a statistics over 25 experiments for each algorithm, and initialize the searches with random initial samples collected through a Latin hypercube sampling scheme. This experimental methodology permits to quantify the influence of different initialization on the algorithms performance. In particular, the multifidelity algorithms are initialized with 1000 design configurations evaluated with the low-fidelity model, and 34 design points computed with the high-fidelity analysis. The single-fidelity algorithm starts the search with an initial set of 35 designs evaluated with the high-fidelity model.

Figure 3 illustrates the median values (solid line) of the minimum of the objective function f^* along with the observations falling between the 25-th and 75-th percentiles (shaded area). All the initial samples correspond to design configurations that score worse than the baseline solution ($f > 1$), and all the competing algorithms

B	f^*_{EGO}	f^*_{MFEI}	f^*_{MFMES}	f^*_{MFPI}	f^*_{NM3-BO}
50	1.1723 (– 17.2 %)	1.0910 (– 9.10 %)	0.9301 (6.99 %)	0.9616 (3.84 %)	0.9123 (8.77 %)
75	0.9670 (3.30 %)	0.8779 (12.21 %)	0.9301 (6.99 %)	0.9434 (5.66 %)	0.8260 (17.4 %)
100	0.9670 (3.30 %)	0.8779 (12.21 %)	0.9101 (8.99 %)	0.8984 (10.16 %)	0.8260 (17.4 %)
150	0.9101 (8.99 %)	0.8779 (12.21 %)	0.9027 (9.73 %)	0.8984 (10.16 %)	0.8260 (17.4 %)
200	0.9082 (9.18 %)	0.8779 (12.21 %)	0.9026 (9.73 %)	0.8984 (10.16 %)	0.8260 (17.4 %)
250	0.9062 (9.38 %)	0.8779 (12.21 %)	0.9026 (9.73 %)	0.8984 (10.16 %)	0.8260 (17.4 %)

Table 1. Median values of the minimum of the objective function f^* and corresponding design improvement (·) obtained with the competing algorithms.

Method	$f^*(\mathbf{x}^*)$	$\mathbf{x}^* = [F^*_V, F^*_N, s^*_{TPS}]$	m^*_{TPS}	T^*_{TPS}	m^*_P
EGO	0.8999 (10.01 %)	$\mathbf{x}^* = [33.63\ kN, 0.969\ kN, 0.0396\ m]$	476.6 kg	1320 K	74.61 kg
MFEI	0.8717 (12.83 %)	$\mathbf{x}^* = [35.67\ kN, 1.561\ kN, 0.0341\ m]$	410.35 kg	1326 K	80.06 kg
MFMES	0.8963 (10.37 %)	$\mathbf{x}^* = [35.97\ kN, 2.046\ kN, 0.0373\ m]$	447.96 kg	1329 K	81.52 kg
MFPI	0.8921 (10.79 %)	$\mathbf{x}^* = [35.40\ kN, 0.691\ kN, 0.0377\ m]$	453.4 kg	1322 K	77.97 kg
NM3-BO	0.8202 (17.98 %)	$\mathbf{x}^* = [29.53\ kN, 0.807\ kN, 0.0304\ m]$	365.17 kg	1310 K	65.45 kg

Table 2. Comparison between the best design solutions identified with the competing algorithms.
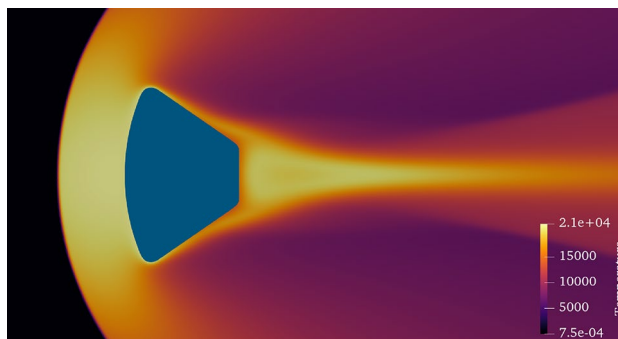

Figure 4. Temperature contours at the heat peak condition evaluated with the high-fidelity aerothermodynamic model considering the best design solution achieved with our NM3-BO algorithm.

improve the baseline design solution within the maximum computational budget available $B = 250$. A significant achievement is that the multifidelity frameworks outperform the single-fidelity EGO algorithm. This indicates that the combination of data from multiple information sources allows to efficiently explore the design space and contain the computational cost. However, our NM3-BO algorithm obtains superior design solutions – larger reductions of the objective function – with a fraction of the computational cost required by the competing algorithms to identify suboptimal designs. This outcome suggests that the combination of the non-myopic scheme and the multiple decision making process capitalizes from the design evaluations adopting different sources of information, and effectively accelerates the search toward optimal design solutions.

To clarify the results from the statistics, Table 1 summarizes the median values of f^* for discrete values of the computational budget B. We can observe that the NM3-BO achieves the higher design improvement (8.77%) after the consumption of a budget $B = 50$, whereas the other strategies score worst in terms of design upgrades. A remarkable outcome is the overall acceleration of the MDO procedure provided by NM3-BO: our framework converges for a computational budget below $B = 75$ and leads to a design upgrade of the 17.4%. This result is outstanding if compared with the design improvement of about the 10% obtained by the EGO, MFMES, and MFPI algorithms, and the design upgrade around the 12% achieved by the MFEI at convergence.

Table 2 compares the best design solutions obtained with the competing algorithms over the collected experiments. NM3-BO identifies an optimal design configuration of the re-entry vehicle that delivers an upgrade of the 17.98%, and privileges lower thrust capabilities and contained thickness of the thermal protection system. This determines a lower storage of propellant $m^*_P = 65.45\ kg$ on-board and permits to navigate a safe re-entry trajectory that contains the heat loads affecting the frame. As a result, the temperature of the TPS structure is kept below $T^*_{TPS} = 1310\ K$ with a total TPS structural mass of $m^*_{TPS} = 365.17\ kg$. Figure 4 provides details about the temperature distribution achieved adopting the best re-entry capsule computed with the NM3-BO at the

heat peak condition. It should be noticed that all the design solutions identified by the algorithms prioritize the reduction of both the TPS and propellant mass, and penalize the temperature reached by the heat shield. On one hand, this permits to contain the overall mass of the vehicle with consequent savings in terms of launch costs; on the other hand, the temperature peaks experienced by the structural frame are far below the thermal properties of the TPS material: this guarantees the survival of the vehicle during the atmospheric descent.

Conclusions

This paper proposes a Non-Myopic Multipoint Multifidelity Bayesian Optimization (NM3-BO) framework to significantly accelerate expensive multidisciplinary design optimization problems. NM3-BO combines two distinguishing features: i) a non-myopic decision making process maximizes the cumulative reward of design solutions over future iterations and ii) a penalization strategy enables multiple decisions as a batch of design configurations and associated level of fidelity to evaluate simultaneously. This search scheme identifies promising batches through the measure of their future utility, and leverages parallel computations to reduce the overall computational cost of the MDO procedure. The NM3-BO algorithm is demonstrated for the MDO problem of a space re-entry vehicle. The method permits substantial accelerations and identifies superior design solutions compared to state-of-the-art multifidelity and single-fidelity algorithms. In particular, NM3-BO delivers on average a space vehicle design improvement of the 17.4% with a fraction of the computational resources adopted by competing algorithms to identify suboptimal solutions. The results suggest that the non-myopic multiple decision making scheme can pave the way to major computational and energy efficiency gain for the multidisciplinary design of complex engineering systems.

Data availability

The datasets generated and analysed during the current study available from the corresponding author on reasonable request.

References

1. Sobieszczanski-Sobieski, J. Multidisciplinary design optimization: An emerging new engineering discipline. In *Advances in Structural Optimization* (ed. Sobieszczanski-Sobieski, J.) 483–496 (Springer, 1995).
2. Alexandrov, N. M. *et al. Multidisciplinary Design Optimization: State of the Art* (SIAM, 1997).
3. Agte, J. *et al.* Mdo: Assessment and direction for advancement-an opinion of one international group. *Struct. Multidiscip. Optim.* **40**, 17–33 (2010).
4. Martins, J. R. & Lambe, A. B. Multidisciplinary design optimization: A survey of architectures. *AIAA J.* **51**, 2049–2075 (2013).
5. Benaouali, A. & Kachel, S. Multidisciplinary design optimization of aircraft wing using commercial software integration. *Aerosp. Sci. Technol.* **92**, 766–776 (2019).
6. Fazeley, H., Taei, H., Naseh, H. & Mirshams, M. A multi-objective, multidisciplinary design optimization methodology for the conceptual design of a spacecraft bi-propellant propulsion system. *Struct. Multidiscip. Optim.* **53**, 145–160 (2016).
7. Brevault, L., Balesdent, M. & Hebbal, A. Multi-objective multidisciplinary design optimization approach for partially reusable launch vehicle design. *J. Spacecr. Rockets* **57**, 373–390 (2020).
8. Gerber, D. J. & Lin, S.-H.E. Designing in complexity: Simulation, integration, and multidisciplinary design optimization for architecture. *Simulation* **90**, 936–959 (2014).
9. Wang, W., Gao, F., Cheng, Y. & Lin, C. Multidisciplinary design optimization for front structure of an electric car body-in-white based on improved collaborative optimization method. *Int. J. Automot. Technol.* **18**, 1007–1015 (2017).
10. Wang, S. *et al.* Multidisciplinary design optimization of underwater glider for improving endurance. *Struct. Multidiscip. Optim.* **63**, 2835–2851 (2021).
11. Yan, J. *et al.* Multidisciplinary design optimization of distributed energy generation systems: The trade-offs between life cycle environmental and economic impacts. *Appl. Energy* **284**, 116197 (2021).
12. Hwang, J. T., Jain, A. V. & Ha, T. H. Large-scale multidisciplinary design optimization-review and recommendations. *AIAA Aviat. 2019 Forum* https://doi.org/10.2514/6.2019-3106 *(2019).*
13. Simpson, T. W. & Martins, J. R. Multidisciplinary design optimization for complex engineered systems design: State of the research and state of the practice-report from a national science foundation workshop. *Int. Des. Eng. Tech. Conf. Comput. Inf. Eng. Conf.* **54822**, 835–845 (2011).
14. Dubreuil, S., Bartoli, N., Gogu, C. & Lefebvre, T. Towards an efficient global multidisciplinary design optimization algorithm. *Struct. Multidiscip. Optim.* **62**, 1739–1765 (2020).
15. Peherstorfer, B., Willcox, K. & Gunzburger, M. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *SIAM Rev.* **60**, 550–591 (2018).
16. Piperni, P., DeBlois, A. & Henderson, R. Development of a multilevel multidisciplinary-optimization capability for an industrial environment. *AIAA J.* **51**, 2335–2352 (2013).
17. Mifsud, M., MacManus, D. G. & Shaw, S. A variable-fidelity aerodynamic model using proper orthogonal decomposition. *Int. J. Numer. Methods Fluids* **82**, 646–663 (2016).
18. Liem, R. P., Mader, C. A. & Martins, J. R. Surrogate models and mixtures of experts in aerodynamic performance prediction for aircraft mission analysis. *Aerosp. Sci. Technol.* **43**, 126–151 (2015).
19. Goertz, S. *et al.* Overview of collaborative multi-fidelity multidisciplinary design optimization activities in the DLR project Victoria. *AIAA Aviat. 2020 Forum* https://doi.org/10.2514/6.2020-3167 *(2020).*
20. Fernández-Godino, M. G., Park, C., Kim, N.-H. & Haftka, R. T. Review of multi-fidelity models. Preprint at http://arxiv.org/abs/1609.07196 (2016).
21. Mosca, V., Karpuk, S., Sudhi, A., Badrya, C. & Elham, A. Multidisciplinary design optimisation of a fully electric regional aircraft wing with active flow control technology. *Aeronaut. J.* **126**, 730–754 (2022).
22. Wu, N., Mader, C. A. & Martins, J. R. A gradient-based sequential multifidelity approach to multidisciplinary design optimization. *Struct. Multidiscip. Optim.* **65**, 1–20 (2022).
23. Van Nguyen, N., Tyan, M., Jin, S. & Lee, J.-W. Adaptive multifidelity constraints method for efficient multidisciplinary missile design framework. *J. Spacecr. Rockets* **53**, 184–194 (2016).

24. Shi, R., Liu, L., Long, T., Wu, Y. & Gary Wang, G. Multi-fidelity modeling and adaptive co-kriging-based optimization for all-electric geostationary orbit satellite systems. *J. Mech. Des.* **142**, 021404 (2020).

25. Pellegrini, R., Serani, A., Broglia, R., Diez, M. & Harries, S. Resistance and payload optimization of a sea vehicle by adaptive multi-fidelity metamodeling. In *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 1904 (2018).

26. Alam, K., Ray, T. & Anavatti, S. G. Design optimization of an unmanned underwater vehicle using low-and high-fidelity models. *IEEE Trans. Syst. Man Cybern. Syst.* **47**, 2794–2808 (2015).

27. Wang, X., Liu, Y., Sun, W., Song, X. & Zhang, J. Multidisciplinary and multifidelity design optimization of electric vehicle battery thermal management system. *J. Mech. Des.* **140**, 094501 (2018).

28. Anselma, P., Niutta, C. B., Mainini, L. & Belingardi, G. Multidisciplinary design optimization for hybrid electric vehicles: component sizing and multi-fidelity frontal crashworthiness. *Struct. Multidiscip. Optim.* **62**, 2149–2166 (2020).

29. Sykes, V., Collu, M. & Coraddu, A. A flexible, multi-fidelity levelised cost of energy model for floating offshore wind turbines multidisciplinary design, analysis and optimisation approaches. *J. Phys. Conf. Ser.* **2265**, 042029 (2022).

30. Jasa, J., Bortolotti, P., Zalkind, D. & Barter, G. Effectively using multifidelity optimization for wind turbine design. *Wind Energy Sci.* **7**, 991–1006 (2022).

31. Huang, D., Allen, T. T., Notz, W. I. & Miller, R. A. Sequential Kriging optimization using multiple-fidelity evaluations. *Struct. Multidiscip. Optim.* **32**, 369–382 (2006).

32. Takeno, S. *et al.* Multi-fidelity Bayesian optimization with max-value entropy search and its parallelization. In *International Conference on Machine Learning* (eds Takeno, S. *et al.*) 9334–9345 (PMLR, 2020).

33. Ruan, X., Jiang, P., Zhou, Q., Hu, J. & Shu, L. Variable-fidelity probability of improvement method for efficient global optimization of expensive black-box problems. *Struct. Multidiscip. Optim.* **62**, 3021–3052 (2020).

34. Charayron, R., Lefebvre, T., Bartoli, N. & Morlier, J. Multi-fidelity Bayesian optimization strategy applied to overall drone design. In *AIAA SCITECH 2023 Forum*, 2366 (2023).

35. Serani, A. *et al.* Adaptive multi-fidelity sampling for CFD-based optimisation via radial basis function metamodels. *Int. J. Comput. Fluid Dyn.* **33**, 237–255 (2019).

36. Reisenthel, P. H. & Allen, T. T. Application of multifidelity expected improvement algorithms to aeroelastic design optimization. In *10th AIAA multidisciplinary design optimization conference*, 1490 (2014).

37. Tran, A., Tranchida, J., Wildey, T. & Thompson, A. P. Multi-fidelity machine-learning with uncertainty quantification and Bayesian optimization for materials design: Application to ternary random alloys. *J. Chem. Phys.* **153**, 074705 (2020).

38. Shahriari, B., Swersky, K., Wang, Z., Adams, R. P. & De Freitas, N. Taking the human out of the loop: A review of Bayesian optimization. *Proc. IEEE* **104**, 148–175 (2015).

39. Frazier, P. I. A tutorial on bayesian optimization. Preprint at http://arxiv.org/abs/1807.02811 (2018).

40. Ginsbourger, D. & Le Riche, R. Towards Gaussian process-based optimization with finite time horizon. In *mODa 9-Advances in Model-Oriented Design and Analysis* (eds Ginsbourger, D. & Le Riche, R.) 89–96 (Springer, 2010).

41. Osborne, M. A., Garnett, R. & Roberts, S. J. Gaussian processes for global optimization. In *3rd International Conference on Learning and Intelligent Optimization (LION3)*, 1–15 (Citeseer, 2009).

42. Lam, R., Willcox, K. & Wolpert, D. H. Bayesian optimization with a finite budget: An approximate dynamic programming approach. *Adv. Neural Inf. Process. Syst.* **29**, 883–891 (2016).

43. González, J., Osborne, M. & Lawrence, N. Glasses: Relieving the myopia of Bayesian optimisation. In *Artificial Intelligence and Statistics* (eds González, J. *et al.*) 790–799 (PMLR, 2016).

44. Wu, J & Frazier, P. Practical two-step lookahead Bayesian optimization. *Adv. Neural Inf. Process. Syst.* **32**, 9813–9823 (2019).

45. Wang, H., van Stein, B., Emmerich, M. & Back, T. A new acquisition function for Bayesian optimization based on the moment-generating function. In *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 507–512 (IEEE, 2017).

46. de Winter, R., van Stein, B. & Bäck, T. Multi-point acquisition function for constraint parallel efficient multi-objective optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 511–519 (2022).

47. Grassi, F., Manganini, G., Garraffa, M. & Mainini, L. Raal: Resource aware active learning for multifidelity efficient optimization. *AIAA J.* **61**, 2744–2753 (2023).

48. Meliani, M. *et al.* Multi-fidelity efficient global optimization: Methodology and application to airfoil shape design. *AIAA Aviat. 2019 Forum* https://doi.org/10.2514/6.2019-3236 (2019).

49. Brevault, L., Balesdent, M. & Hebbal, A. Overview of gaussian process based multi-fidelity techniques with variable relationship between fidelities, application to aerospace systems. *Aerosp. Sci. Technol.* **107**, 106339 (2020).

50. Williams, C. K. & Rasmussen, C. E. *Gaussian Processes for Machine Learning* Vol. 2 (MIT press Cambridge, 2006).

51. Kennedy, M. C. & O'Hagan, A. Predicting the output from a complex computer code when fast approximations are available. *Biometrika* **87**, 1–13 (2000).

52. Forrester, A., Sobester, A. & Keane, A. *Engineering Design via Surrogate Modelling: A Practical Guide* (Wiley, 2008).

53. Kontogiannis, S. G., Demange, J., Savill, A. M. & Kipouros, T. A comparison study of two multifidelity methods for aerodynamic optimization. *Aerosp. Sci. Technol.* **97**, 105592 (2020).

54. Puterman, M. L. Markov decision processes. *Handb. Oper. Res. Manag. Sci.* **2**, 331–434 (1990).

55. Papadimitriou, C. H. & Tsitsiklis, J. N. The complexity of Markov decision processes. *Math. Oper. Res.* **12**, 441–450 (1987).

56. Bellman, R. Dynamic programming. *Science* **153**, 34–37 (1966).

57. Eddy, S. R. What is dynamic programming?. *Nat. Biotechnol.* **22**, 909–910 (2004).

58. Wilson, J., Hutter, F. & Deisenroth, M. Maximizing acquisition functions for Bayesian optimization. *Adv. Neural Inf. Process. Syst.* **31**, 1–12 (2018).

59. Di Fiore, F. & Mainini, L. Non-myopic multifidelity bayesian optimization. Preprint at http://arxiv.org/abs/2207.06325 (2022).

60. González, J., Dai, Z., Hennig, P. & Lawrence, N. Batch Bayesian optimization via local penalization. In *Artificial Intelligence and Statistics* (eds González, J. *et al.*) 648–657 (PMLR, 2016).

61. Beran, P. S., Bryson, D., Thelen, A. S., Diez, M. & Serani, A. Comparison of multi-fidelity approaches for military vehicle design. *AIAA Aviat. 2020 Forum* https://doi.org/10.2514/6.2020-3158 (2020).

62. Mainini, L. *et al.* Analytical benchmark problems for multifidelity optimization methods. Preprint at http://arxiv.org/abs/2204.07867 (2022).

63. Steward, D. The design structure matrix: A method for managing the design of complex systems. *IEEE Trans. Eng. Manag.* **28**, 71–74 (1981).

64. Di Fiore, F., Maggiore, P. & Mainini, L. Multifidelity domain-aware learning for the design of re-entry vehicles. *Struct. Multidiscip. Optim.* **64**, 3017–3035 (2021).

65. Harris, R., Stewart, M. & Koenig, W. Thermal protection systems technology transfer from apollo and space shuttle to the orion program. In *2018 AIAA SPACE and Astronautics Forum and Exposition*, 5134 (2018).

66. Jones, D. R., Schonlau, M. & Welch, W. J. Efficient global optimization of expensive black-box functions. *J. Glob. Optim.* **13**, 455–492 (1998).

Author contributions

F.D.F. and L.M. contributed to the study conception and design, and performed material preparation, data collection and analysis. The first draft of the manuscript was written by F.D.F. and L.M. provided supervision, revised, edited and commented on all the versions of the manuscript. All authors approved the final manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1038/s41598-023-48757-3.

Correspondence and requests for materials should be addressed to F.D.F.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.