

NM-MF: Non-Myopic Multifidelity Framework for Constrained Multi-Regime Aerodynamic Optimization

Original

NM-MF: Non-Myopic Multifidelity Framework for Constrained Multi-Regime Aerodynamic Optimization / DI FIORE, F., Mainini, L.. - In: AIAA JOURNAL. - ISSN 0001-1452. - 61:3(2023), pp. 1270-1280. [10.2514/1.J062219]

Availability:

This version is available at: 11583/2974676 since: 2023-02-20T20:18:10Z

Publisher:

AMER INST AERONAUTICS ASTRONAUTICS

Published

DOI:10.2514/1.J062219

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

AIAA preprint/submitted version e/o postprint/Author's Accepted Manuscript

© AIAA. NM-MF: Non-Myopic Multifidelity Framework for Constrained Multi-Regime Aerodynamic Optimization / DI FIORE, Francesco; Mainini, Laura published in : AIAA JOURNAL, 2023, 61, 3, 1270-1280, <http://dx.doi.org/10.2514/1.J062219>.

(Article begins on next page)

NM-MF: Non-Myopic Multifidelity Framework for Constrained Multi-Regime Aerodynamic Optimization

Francesco Di Fiore* and Laura Mainini †
Politecnico di Torino, Torino, Italy, 10129

The exploration and trade-off analysis of different aerodynamic design configurations requires solving optimization problems. The major bottleneck to assess the optimal design is the large number of time-consuming evaluations of high-fidelity computational fluid dynamics (CFD) models, necessary to capture the non-linear phenomena and discontinuities that occurs at higher Mach number regimes. To address this limitation, we introduce an original non-myopic multifidelity Bayesian framework aimed at including expensive high-fidelity CFD simulations for the optimization of the aerodynamic design. Our scheme proposes a novel two-step lookahead policy to maximize the improvement of the solution quality considering the rewards of future steps, and combines it with utility functions informed by the fluid dynamic regime and the information extracted from data, to wisely select the aerodynamic model to interrogate. We validate the proposed framework for the case of a constrained drag coefficient optimization problem of a NACA 0012 airfoil, and compare the results to other popular multifidelity and single-fidelity optimization frameworks. The results suggest that our strategy outperforms the other approaches, allowing to significantly reduce the drag coefficient through a principled selection of limited evaluations of the high-fidelity CFD model.

I. Introduction

AERODYNAMIC design optimization is an essential component for the design of resource-efficient air vehicles, and relates to airfoils, aircraft wings, turbine engine blades and combustion chambers among others [1]. Modern engineering approaches to the aerodynamic design involve the use of high-fidelity physics-based computer models, capable to capture non-linear physics and discontinuities in the flow-field with a high level of detail. In particular, accurate representations are necessary to simulate the compressibility effects of the transonic regime of flight, where the aerodynamic domain is characterized by a mixed subsonic-supersonic flow-field that requires the numerical solution of nonlinear partial differential equations to be modeled. However, the implementation of these accurate models in the aerodynamic optimization poses significant challenges due to their high computational cost (ranging from a few hours to weeks on high-performance computing clusters), which is not well suited with simulation-based optimization

*PhD. Student, Department of Mechanical and Aerospace Engineering, francesco.difiore@polito.it, AIAA Member

†Visiting/Adjunct Professor, Department of Mechanical and Aerospace Engineering, laura.mainini@polito.it, AIAA Associate Fellow

techniques that might require a large number of model evaluations during the search for the optimal design. Therefore, the aerodynamic optimization frameworks commonly rely on low-fidelity physics based models for the flow-field analysis, which introduce approximations of the aerodynamics phenomena to save computational cost. However, this models may not be adequate to predict nonlinear physics and compressibility effects as shock waves, flow separation or turbulence, that occur at higher flight speeds in the transonic regime. This motivates the interest for efficient optimization approaches to include costly high-fidelity physics-based models in the aerodynamic optimization process and closely capture the fluid domain while keeping the computational cost acceptable.

This work proposes an original computational framework for the multifidelity aerodynamic design and optimization process that aims to wisely use high-fidelity aerodynamic evaluations to improve the search and assessment of the design alternatives. Multifidelity methods are computational strategies that combine information from models at different levels of fidelity, leveraging fast low-fidelity data to accelerate the exploration of design configurations, and refine the accuracy of the solution through a limited number of high-fidelity evaluations [2–6]. Commonly, multifidelity strategies adopt Bayesian frameworks for the optimization of black-box functions. Bayesian optimization estimates the objective function through a probabilistic surrogate model (typically Gaussian processes) and repeatedly maximizes an acquisition function to select the new design to evaluate. The acquisition function defines a policy that allows to query the most promising input locations with a continuous trade-off between the exploration and exploitation of the design space. The multifidelity scenario allows us to query black box representations of the objective function at different levels of fidelity. Accordingly, multifidelity Bayesian optimization combines the information of the objective at different fidelities into a single surrogate model; the surrogate model is then used to compute the acquisition function that selects the location of the new design to evaluate and the associated level of fidelity to query.

Multifidelity Bayesian optimization has been applied to a wide range of aerodynamic design and optimization problems. Reischel and Allen [7] proposed a multifidelity Bayesian framework for the aeroelastic design and optimization, developing an adaptive method to achieve probabilistic convergence and enhance performance, and validating the approach for the case of a UAV wing design and optimization. Meliani et al. [8] addressed the airfoil shape design and optimization through a multifidelity framework for high-dimensional Bayesian optimization, to improve the efficiency of the methodology as the number of design variables increases. Mondal et al. [9] developed a multifidelity global-local scheme for the design and optimization of a highly-loaded transonic compressor rotor, and demonstrate it for the case of a transonic NASA rotor 37. Kontogiannis et al. [10] proposed a comparison between local trust region and global Bayesian multifidelity approaches applied to the test case of a RAE2822 airfoil optimization problem, demonstrating that the global approach allows to better explore the aerodynamic domain outperforming the local optimization strategy. All these works model the aerodynamic domain through two levels of fidelity, ranging from high-fidelity computational fluid dynamics (CFDs) to low-fidelity representations of the flow-field. Most works demonstrate the formulations for an aerodynamic domain that do not include variations in the fluid domain regime,

limiting the admissible range of the Mach number. Multi-regime optimization brings additional challenges associated with the selection of the accuracy of the aerodynamic model to closely predict non-linear phenomena that occurs during the cross-regime scenario. Moreover, these works use multifidelity acquisition functions that measure only the direct effect of the evaluation of the objective through an immediate measure of the solution quality, precluding greater improvement that may be captured considering future steps.

This motivates the interest for non-myopic formulations of the acquisition function that allow to unlock the informative gains from future optimization steps. Ginsbourger and Le Riche [11] and Lam et al. [12] formulated a non-myopic sampling strategy based on a partially observable Markov decision process (POMDP). However, the computational cost associated with POMDP makes this strategy suitable only for low dimensional problems. Several strategies have been proposed to approximate the solution of the POMDP problem. In particular, two-steps lookahead acquisition functions are particularly attractive because allow to improve the performance of Bayesian optimization while containing the computational cost that would raise by looking ahead into future steps [11, 13–16]. However, most of the non-myopic formulations of the acquisition function developed in literature are proposed for a single fidelity Bayesian framework, which would not be suitable for aerodynamic optimization problems that interrogate extensively high-fidelity fluid-dynamics models.

This paper introduces a novel multifidelity Bayesian optimization based on a non-myopic formulation to compute an efficient surrogate model for the aerodynamic design and optimization task in a multi-regime fluid domain. We adopt a multifidelity Gaussian process as the surrogate model of the objective function, which is progressively updated through an acquisition function based on the multifidelity expected improvement. Our computational approach proposes an original multifidelity acquisition function that comes with two novel properties: (i) a two-steps lookahead sampling policy that unlocks the improvement of the solution quality obtainable through a coordinated action across future evaluations; (ii) a strategy for selecting the level of fidelity of the aerodynamic model based not only on the combination of statistical properties of the surrogate model, but also on the flow-field data of the aerodynamic domain. Our acquisition function allows to sample the new design configuration to be evaluated through an optimal policy based on the solution of a dynamic programming (DP) problem [17, 18]: we minimize the expected loss achievable at the end of the next optimization step to quantify the reward in the design solution that could be collected at the same future step. Solving the DP problem requires to compute nested expectations and maximizations for which is not available a closed-form solution. Our framework bypasses the computation in closed form of the DP solution through a Monte Carlo strategy, to provide a robust estimate of the acquisition function. In addition, we formulate the acquisition function to grasp the evolution of the aerodynamic regimes through a domain-aware utility function, that encourages the evaluation of the flow-field using the highest level of fidelity as the Mach number approaches to the transonic regime. This element of our computational strategy allows to select the fidelity of the aerodynamic model accounting for discontinuities and separations of the flow-field that may not be captured by the low-fidelity model.

We demonstrate our strategy for the aerodynamic design and optimization problem of a NACA 0012 airfoil: the objective is the minimization of the drag coefficient subject to a constraint on the lift coefficient at a fixed flight altitude. The physics-based modeling consists of three levels of fidelity for the aerodynamics, including an high-fidelity CFD model to solve numerically the Reynolds Averaged Navier-Stokes (RANS) equations, a mid-fidelity CFD model implementing a coarser discretization to solve RANS, and a low-fidelity model based on the potential flow panel method. The results obtained with our two-steps lookahead multifidelity and multi-regime scheme are compared over a variety of single-fidelity and multifidelity acquisition functions widely used in literature.

The paper is organized as follows. Section II illustrates the physics based aerodynamic models and the formulation of the demonstrative optimization problem of a NACA 0012 airfoil. Section III presents our original two-steps multifidelity and multi-regime framework, and its implementation for the test-case optimization problem. Section IV discusses the results obtained from our numerical experiments, and Section V provides conclusions and remarks.

II. Multi-regime Aerodynamic Optimization Problem

This work addresses the aerodynamic design and optimization task accounting for the variations of the physical properties that characterize the compressible fluid dynamics. Commonly, the evolution of the aerodynamic domain is denoted using the Mach number as a criterion. The flow-field is defined subsonic if the Mach number is less than 1 in all the fluid domain, and hence the flow velocity is everywhere less than the speed of sound. This flow is characterized by smooth streamlines and continuously varying properties. As the Mach number approaches the sonic condition, the fluid domain is characterized by local supersonic regions with possible shock waves, causing the separation of the viscous boundary layer from the surface where the shock impinges on the body. Such mixed subsonic-supersonic fluid domain is commonly referred to as the transonic flow, and usually the transonic regime emerges between 0.8 and 1.2 values of the Mach number of the free-stream. The interaction between shock wave and boundary layer is one of the most important aspects of the transonic flow as induces discontinuities in the fluid properties, a large rise of the drag force, a rapid shift of the center of pressure, and unsteady effects. Due to the different physical characteristics of the flow regimes, the modelling of the aerodynamic phenomena can be addressed thorough different approaches: the subsonic regime can be analyzed using simplified aerodynamic models that are capable to approximate the governing partial differential equation, while the modeling of the transonic regime requires the implementation of accurate and robust schemes to resolve shock waves, contact discontinuities and non-linearity that characterize the mixed type flow-field.

Our framework aims at including multiple representations of the fluid domain in the aerodynamic design and optimization process, which can model the different physical phenomena associated with the evolution of the fluid regime. In particular, we consider the case of the aerodynamic optimization of the drag coefficient of a 2D airfoil subject to a constraint on the lift coefficient. We use the popular NACA 0012 airfoil constrained optimization problem as the validation test case of our optimization framework due to the availability of the real-world solution of the optimization

problem and experimental data of the lift and drag coefficients [19, 20]. This aerodynamic problem is commonly adopted to validate multifidelity methods [21, 22] and surrogate modeling techniques [23, 24] enabling the assessment of the performance of the algorithms with respect to the experimental optimum and real-world aerodynamic data, which usually is not the case for real-world experiments where the optimum solution is not computable. The aerodynamic modeling consists of three representations at different levels of fidelity: the high-fidelity model is based on a CFD solver for the numerical approximation of the Reynolds Averaged Navier-Stokes equations (Section II.B); the mid-fidelity model is conceptually equivalent to the high-fidelity model, except for a coarser discretization of the aerodynamic domain (Section II.C); the low-fidelity model solves numerically the potential flow equation through the panel method (Section II.D). In principle, the experimental data available for the NACA 0012 airfoil can be used as an additional level of fidelity for the aerodynamic modeling through surrogate modeling techniques such as mixture of experts approach [25] or Bayesian hierarchical regression [26]. However, we prefer to rely on computer-based models since the limited number of experiments available may not guarantee the computation of a reliable representation of the flow-field.

A. Formulation of the Optimization Problem

The aerodynamic design optimization aims at computing the combination of the Mach number M and angle of attack α that jointly minimize the coefficient of drag C_d of a NACA 0012 airfoil subject to maintaining a minimum coefficient of lift C_l at a certain altitude h :

$$\begin{aligned} \min_{\mathbf{x} \in \mathcal{X}} \quad & C_d(\mathbf{x}) \\ \text{s.t.} \quad & C_{l0} - C_l(\mathbf{x}) \leq 0 \end{aligned} \tag{1a}$$

$$h - h_0 = 0 \tag{1b}$$

$$\mathcal{X} = I_M \times I_\alpha \tag{1c}$$

where $\mathbf{x} = [M, \alpha]$ is the design input, h is the flight altitude, $C_{l0} = 0.45$ is the constraint on the lift coefficient, $h_0 = 10000m$ is the constraint on the flight altitude, and $\mathcal{X} = I_M \times I_\alpha$ is the design space with $I_M = [0.6, 0.99]$ and $I_\alpha = [0^\circ, 4.5^\circ]$. The constraint on the lift coefficient (Equation 1a), on the flight altitude (Equation 1b), and on the domain boundaries (Equation 1c) are imposed to replicate the flight scenario of a civil aircraft during the cruise phase. In particular, the Mach number limits are defined to explore design configurations in both the subsonic and transonic regime with the appropriate level of fidelity to grasp design solutions that may be otherwise discarded.

B. High-fidelity Aerodynamic Model

The high-fidelity aerodynamic model computes the drag coefficient C_d and lift coefficient C_l , given the geometry of the airfoil, the flight altitude h , the Reynolds number Re , the free-stream Mach number M and the flight angle of attack α . We model the aerodynamic domain through the Reynolds Averaged Navier-Stokes (RANS) equations to capture the effects of turbulence that occur at higher regimes of Mach number and angle of attack. The formulation of the RANS model in differential form is:

$$\mathbf{R}(\mathbf{U}) = \frac{\partial(\mathbf{U})}{\partial t} + \nabla \cdot \mathbf{F}^c - \nabla \cdot \mathbf{F}^v - \mathbf{Q} = 0 \quad \text{in } \Omega, \quad t > 0 \quad (2)$$

where Ω is the computational domain, \mathbf{R} are the numerical residuals, \mathbf{Q} is the source term, $\mathbf{U} = (\rho, \rho v, \rho E)$ is the vector of the conservative variables, being $\rho = \rho(h)$ the air density, v the free-stream velocity and E the total energy, and \mathbf{F}^c and \mathbf{F}^v are the convective and viscous fluxes, respectively:

$$\mathbf{F}^c = \begin{pmatrix} \rho v \\ \rho v \otimes v + I p \\ \rho E v + p v \end{pmatrix} \quad (3)$$

$$\mathbf{F}^v = \begin{pmatrix} \cdot \\ \bar{\tau} \\ \bar{\tau} v + k \nabla T \end{pmatrix} \quad (4)$$

where $T = T(h)$ is the free-stream temperature, $p = p(h)$ is the free-stream static pressure, $k = k(h)$ is the thermal conductivity and $\bar{\tau}$ is the tensor of viscous stresses.

We use the finite volume method to discretize Equation 2 in space with a standard edge-based data structure where \mathbf{F}^c and \mathbf{F}^v are evaluated at the midpoint of the edges. The computational domain consists of a grid of 27125 quadrilateral elements with the farfield boundary extended to $500c$ away from the airfoil surface, where $c = 1m$ is the Reynolds length. Figure 1a provides details about the discretization of the computational domain considered in this model. The mesh is refined in proximity of the contour of the airfoil to capture the local shock waves that occur in the transonic regime, which are critical for the accurate estimate of the drag coefficient as they generate strong discontinuities in the flow that induce the separation of the fluid vein. The boundary conditions include the adiabatic wall on the contour of the airfoil and the free-stream conditions on the farfield of the computational domain.

We solve the Reynolds Averaged Navier-Stokes equations (Equation 2) using SU2 version 7.0.3 in RANS mode with the Spalart-Allmaras turbulence model [27], using the Jameson-Schmidt-Turkel model (JST) and the Scalar upwind

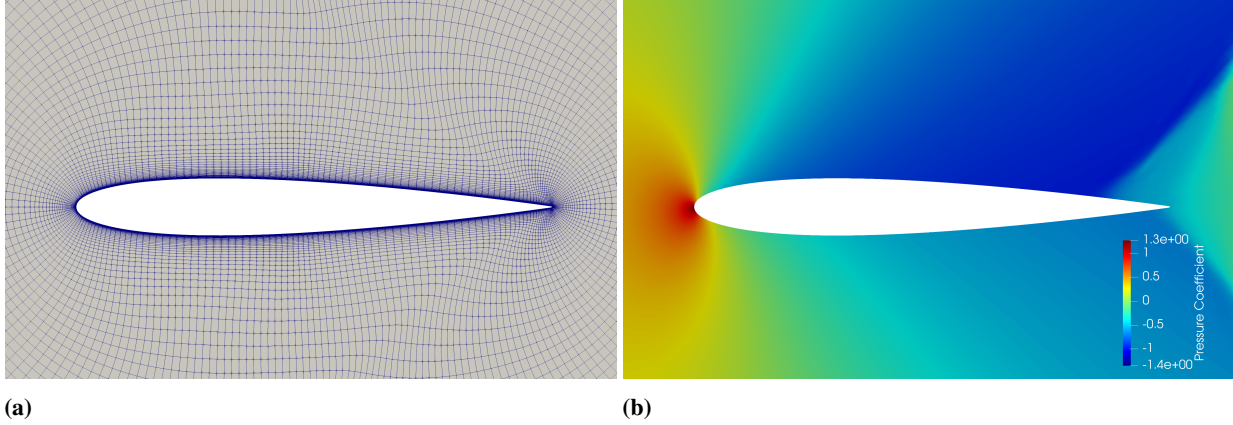


Fig. 1 (a) High-fidelity discretization of the computational domain with 27125 quadrilateral elements, and (b) pressure coefficient contours for a Mach number of 0.9, angle of attack of 3° and an altitude of 10000 m.

model as the flow and turbulence numerical method, respectively, and Euler implicit for the time discretization. The convergence criterion is set as satisfied for values of the computational residuals minor than 10^{-6} . We consider this model as the high-fidelity representation of the aerodynamic phenomena for the capability to provide an accurate estimate of aerodynamic coefficients C_d and C_l even in the presence of discontinuity and separation phenomena typical of the mixed flow condition that characterizes the transonic regime. Figure 1b illustrates the distribution of the pressure coefficient around the NACA 0012 airfoil computed with the high-fidelity aerodynamic model, for a Mach number of 0.9, angle of attack of 3° and an altitude of 10000m.

C. Mid-fidelity Aerodynamic Model

The mid-fidelity aerodynamic model adopts the same governing equations and the same solution method of the high-fidelity model (Section II.B), with the grid density being far coarser than that of the high-fidelity one. In particular, we discretize the computational domain with 14336 quadrilateral elements, which leads to a lower computational time required to estimate the aerodynamic coefficients C_d and C_l in comparison to the high-fidelity model. Figure 2a illustrates the discretization of the computational domain for the mid-fidelity aerodynamic model. As a result of the coarser mesh, this model is not suitable for the computation of the aerodynamic field for Mach numbers close to the sonic condition, as it may not capture the discontinuities of the flow that occur at transonic speed regimes. Therefore, we consider this representation as the mid-fidelity aerodynamic model since the accuracy of the solution is not always satisfactory over the transonic regime with respect to the high-fidelity one. However, the mid-fidelity representation can still provides close approximations of the aerodynamic coefficients for Mach number regimes far from the sonic condition. Figure 2b presents the distribution of the pressure coefficient around a NACA 0012 airfoil in output from the mid-fidelity aerodynamic model, for a Mach number of 0.9, angle of attack of 3° and an altitude of 10000m.

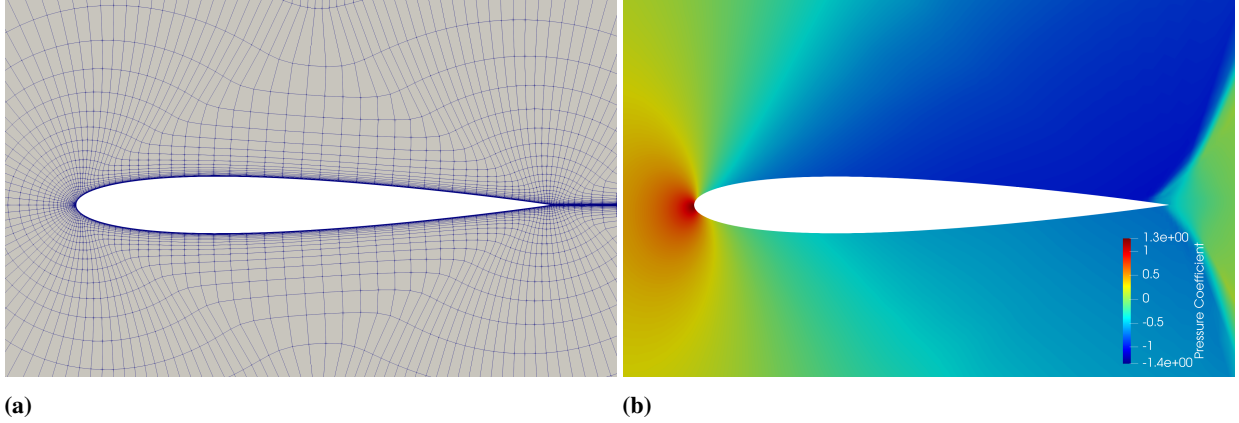


Fig. 2 (a) Mid-fidelity discretization of the computational domain with 14336 quadrilateral elements, and (b) pressure coefficient contours for a Mach number of 0.9, angle of attack of 3° and an altitude of 10000 m.

D. Low-fidelity Aerodynamic Model

The low-fidelity aerodynamic model permits to compute the aerodynamic coefficients of lift C_l and drag C_d given the free-stream Mach number M , the angle of attack α , the Reynolds number Re , and the geometry of the airfoil. We analyze the flow-field around the airfoil through the potential flow panel method with an integral boundary layer formulation and an approximate e^N envelope method to calculate the aerodynamic transition. We use XFOIL [28] to solve the potential flow equation through the panel method with the Karman-Tsien compressibility correction to rapidly predict the airfoil performances. The code is developed to evaluate the aerodynamic coefficients at low Reynolds numbers and is capable to calculate the viscous pressure distribution, capturing the influence of limited trailing edge separation and laminar separation bubbles. With the e^N envelope method, XFOIL tracks only the most amplified frequency on the airfoil downstream from the point of instability to obtain the amplitude of the disturbance, and transition is assumed if an empirical determined amplitude value is reached. The convergence is achieved through the iteration between the outer and inner flow solutions on the boundary layer displacement thickness. We use 200 panels to discretize the airfoil contour and set 200 iterations as the convergence criteria to contain the computation time.

We consider the aerodynamic coefficients computed with this model being the low-fidelity approximation of the real values, as the potential flow panel method and the compressibility correction are not capable to achieve an accurate representation of the flow-field in the more turbulent flow at higher values of Mach number ($M > 0.65$) and angle of attack ($\alpha > 4^\circ$). Moreover, the combination of the panel method with the Karman-Tsien correction causes XFOIL to overestimate lift and underestimate drag [29]. Figure 3a illustrates the comparison between the drag coefficient of a NACA 0012 airfoil computed with the low-fidelity model over the design space, and the wind tunnel prediction of the drag coefficient obtained interpolating with a spline function the experimental data [30]. Wind tunnel data are used in this work as the reference data against which the high-fidelity CFD model (Section II.B) has been validated. This allows to evidence the increase of the discrepancy between the high-fidelity and the low-fidelity outcomes when the flow-field

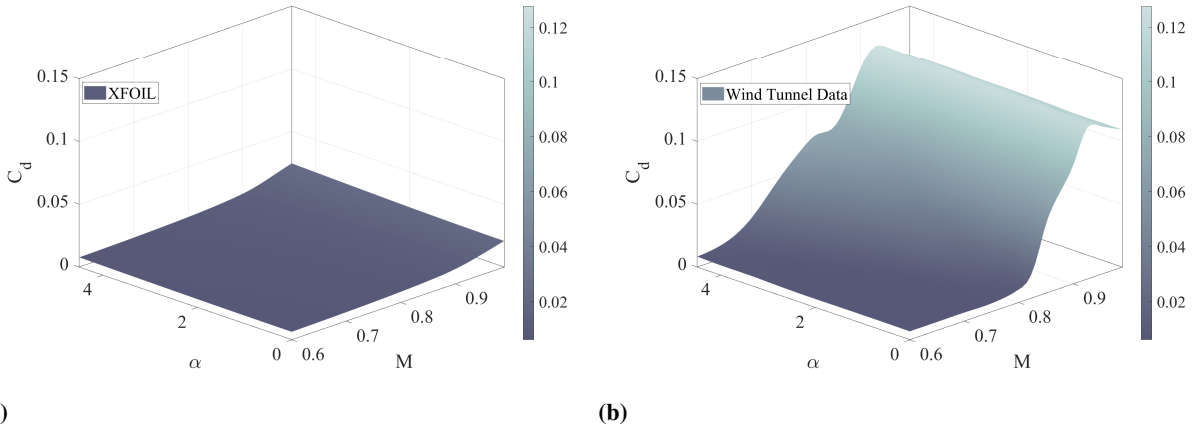


Fig. 3 Prediction of the drag coefficient C_d as a function of Mach number M and angle of attack α obtained by interpolating data computed by the low-fidelity model (3a) and experimental wind tunnel data (3b) for the NACA0012 airfoil [30]

approaches the transonic regime. Moreover, we use these experimental data to assess the goodness of our results in Section IV dedicated to results and discussion.

III. Non-Myopic Multifidelity Bayesian Optimization for Multi-regime Aerodynamics

Our optimization strategy relies on a Bayesian framework for the optimization of black-box objective functions. To optimize the objective function $f : \chi \rightarrow \mathbb{R}$, Bayesian optimization (BO) learns a probabilistic surrogate model to predict the values of the objective across the design space χ and the associated uncertainty in the prediction. Commonly, BO is based on Gaussian Process (GP) to approximate the objective function. The main feature of GP is that in addition to the prediction, it provides a reliable estimate of the uncertainty of the surrogate model response [31]. Based on these information, BO defines through an acquisition function (AF) a policy for the selection of the next design \mathbf{x} to evaluate, which usually involves a trade-off between the exploration of the surrogate and the exploitation toward the optimum design. By maximizing the acquisition function, BO selects the most promising design and updates the surrogate model with the observation of the objective at that configuration.

In the multifidelity context, the objective function is evaluated through multiple informative sources $f^{(1)}(\mathbf{x}), \dots, f^{(L)}(\mathbf{x})$, where the higher the level of fidelity l , the more accurate (yet costly) the query of $f^{(l)}(\mathbf{x})$. A common practice to extend BO to a multifidelity setting is to combine the different sources of information into a single surrogate model and define an acquisition function that determines not only the next design to sample, but also identifies at which level of fidelity evaluate the objective function [32, 33]. Popular multifidelity Bayesian frameworks rely on acquisition functions that measure only the direct reward of an immediate evaluation of the objective function. Examples include the Multifidelity Expected Improvement (MFEI) [34], Multifidelity Predictive Entropy Search (MFES) [35] or Multifidelity Max-value Entropy Search [36]. By myopically maximizing the immediate improvement in the solution quality, these acquisition

functions may preclude greater informative gains obtainable through the next design evaluations that can significantly reduce the computational cost of the optimization process. Moreover, in the context of multi-regime aerodynamic analysis, the optimization problem brings additional challenges associated with the need for the multifidelity scheme to account for the evolution of the physical domain. This motivates the interest in non-myopic multifidelity Bayesian frameworks, also known as lookahead Bayesian optimization, capable to acquire information from multiple fidelity sources looking further into future steps of the optimization, and deal with dynamic and complex physical phenomena.

This work proposes a non-myopic multifidelity Bayesian framework for the aerodynamic optimization in a multi-regime domain. Our strategy aims to effectively combine multiple aerodynamic informative sources at different levels of fidelity to obtain an efficient and accurate estimate of the drag coefficient of an airfoil subject to a lift coefficient constraint. We use a Gaussian process extended to a multifidelity scenario as the probabilistic surrogate model, that is updated iteratively during the optimization process through the evaluation of the acquisition function. Specifically, we conceived a two-step lookahead multifidelity acquisition function that is informed about the dynamic evolution of the aerodynamic domain. At each iteration of the optimization, our framework selects a new set of design variables considering the immediate reward in the solution quality and the improvement in the future steps, while the level of fidelity of the aerodynamic model is determined according to the surrogate properties and the fluid domain features.

In the following, we will present our original multifidelity framework and its peculiar elements as follows. Section III.A introduces the multifidelity Gaussian process surrogate model and the Gaussian regression to approximate the objective function at different levels of fidelity. Section III.B presents the domain-aware multifidelity acquisition providing details about its formulation informed about the fluid regime. In Section III.C, we illustrate the lookahead strategy that permits to incorporate a non-myopic sampling policy in the domain-aware scheme, and formulate our two-step lookahead multifidelity acquisition function for multi-regime optimization. Finally, Section III.D presents the integration of our original acquisition function and the surrogate model in the multifidelity Bayesian framework conceived to solve the multi-regime constrained aerodynamic optimization problem of Section II.A.

A. Multifidelity Gaussian Process

Our framework is based on the Gaussian process regression to compute the surrogate model of the objective function. Let us consider a set of design inputs $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]$ and outputs $\mathbf{y} = [y(\mathbf{x}_1), \dots, y(\mathbf{x}_n)]$, where we assume that the observations of the objective functions are noisy quantities $y(\mathbf{x}) = f(\mathbf{x}) + \varepsilon$ where $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon)$ is the normally distributed measurement noise with standard deviation σ_ε . For the case of the aerodynamic problem described in Section II.A, the design inputs consists of combinations of the Mach number M and angle of attack α while the outputs are the observations of the drag C_d and lift C_l coefficients. The Gaussian process assumes that the output is distributed according to a multivariate Gaussian distribution $p(\mathbf{y}|\mathbf{X}) = \mathcal{N}(\mathbf{y}|\boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\kappa}(\mathbf{x}, \mathbf{x}'))$, where $\boldsymbol{\mu}$ is the mean function and $\boldsymbol{\kappa}$ is the covariance function. According to the Bayesian approach, we define a prior distribution of the objective function

as a Gaussian process with mean function $\mu(\mathbf{x}) = 0$ and covariance function κ , and relocate probabilities based on the observed data using Bayes' rule. Due to the multi-variate Gaussian form, given a new design input \mathbf{x}_{n+1} , the posterior distribution of the output $p(f(\mathbf{x}_{n+1}) | \mathbf{x}_{n+1}, \mathbf{X}, \mathbf{y})$ is a Gaussian process completely specified by the posterior mean μ and variance σ^2 .

The Gaussian process regression can be extended to the multifidelity scenario combining multiple representations of the objective function into a single probabilistic model. This is the case of the NACA0012 optimization problem, where we consider three aerodynamic models at different levels of fidelity hierarchically ordered according to the accuracy and cost of the prediction of the aerodynamic coefficients. We assume that the prediction of the objective at a certain level of fidelity l can be derived recursively through an autoregressive scheme [37], modeling the lowest-fidelity prior $f^{(1)}$ as a Gaussian process with mean function $\mu^{(l)} = 0$ and covariance $\kappa_1(\mathbf{x}, \mathbf{x}')$:

$$f^{(l)} = \rho f^{(l-1)}(\mathbf{x}) + \delta^{(l)}(\mathbf{x}) \quad l = 2, \dots, L \quad (5)$$

where ρ is a constant scaling factor and $\delta^{(l)}$ is the discrepancy between the two levels of fidelity defined as a Gaussian process with mean function $\mu^{(l)} = 0$ and covariance function $\kappa^{(l)}(\mathbf{x}, \mathbf{x}')$. The posterior distribution of the multifidelity Gaussian process is defined by the mean and variance:

$$\mu^{(l)}(\mathbf{x}) = \kappa_n^{(l)}(\mathbf{x})^T (\mathbf{K} + \sigma_\epsilon \mathbf{I})^{-1} \mathbf{y} \quad (6)$$

$$\sigma^{2(l)}(\mathbf{x}) = \kappa((\mathbf{x}, l), (\mathbf{x}, l)) - \kappa_n^{(l)}(\mathbf{x})^T (\mathbf{K} + \sigma_\epsilon \mathbf{I})^{-1} \kappa_n^{(l)}(\mathbf{x}) \quad (7)$$

where $\kappa_n^{(l)}$ is defined as $\kappa_n(\mathbf{x}) \doteq (\kappa((\mathbf{x}, l), (\mathbf{x}_0, l_0)), \dots, \kappa((\mathbf{x}, l), (\mathbf{x}_n, l_n)))$, the term \mathbf{K} is the kernel matrix such that $\mathbf{K}(i, j) = \kappa((\mathbf{x}_i, l_i), (\mathbf{x}_j, l_j))$, the vector $\mathbf{y} \doteq (y^{(l_1)}(\mathbf{x}_1), \dots, y^{(l_n)}(\mathbf{x}_n))^T$ collects the noisy observations and \mathbf{I} is the n -dimensional identity matrix. The posterior mean and variance evaluated at any point of the feasible set of input χ represent the prediction and associated uncertainty of the model at the point \mathbf{x} , respectively.

B. Multifidelity Acquisition Function for Domain-Awareness

Thus far, we have described the probabilistic surrogate model used to represent our belief about the unknown objective function when multiple representations at different levels of fidelity are available. Based on the posterior mean and variance, the multifidelity acquisition function determines the policy that decides the next design to evaluate and which aerodynamic model to query at each iteration of the optimization process. Our Multifidelity Domain-Aware (MFDA) acquisition function builds up onto multifidelity expected improvement [34] and can be formulated for the multi-regime aerodynamic optimization problem:

$$MFDA(\mathbf{x}, l, M) = EI(\mathbf{x})\alpha_1(\mathbf{x}, l)\alpha_2(\mathbf{x}, l)\alpha_3(l)\alpha_4(M, l). \quad (8)$$

In Equation (8) $EI(\mathbf{x})$ is the expected improvement [38] evaluated at the highest level of fidelity:

$$EI(\mathbf{x}) = \begin{cases} \left(\mu^{(L)}(\mathbf{x}) - f(\mathbf{x}^+) \right) \Phi \left(\frac{\mu^{(L)}(\mathbf{x}) - f(\mathbf{x}^+) - \xi}{\sigma^{(L)}(\mathbf{x})} \right) + \sigma^{(L)}(\mathbf{x}) \phi \left(\frac{\mu^{(L)}(\mathbf{x}) - f(\mathbf{x}^+) - \xi}{\sigma^{(L)}(\mathbf{x})} \right) & \text{if } \sigma^{(L)}(\mathbf{x}) > 0 \\ 0 & \text{if } \sigma^{(L)}(\mathbf{x}) = 0 \end{cases} \quad (9)$$

where $f(\mathbf{x}^+)$ is the minimum value of the objective function evaluated at the best design \mathbf{x}^+ identified at the current iteration, $\Phi(\cdot)$ and $\phi(\cdot)$ are, respectively, the cumulative distribution function and the probability density function of $\mathcal{N}(0, 1)$, and the parameter ξ allows to balance the global exploration of the design space and the local exploitation of the regions where the optimum design is supposed to be; larger values of ξ result in a major influence of the posterior uncertainty $\sigma^{(L)}$ encouraging the exploration of uncertain regions of χ while decreasing ξ leads to an increased importance of the posterior prediction $\mu^{(L)}$ incentivizing the sampling toward the optimum. The formulation of the multifidelity acquisition function includes the utility functions α_1 , α_2 , α_3 , and α_4 defined as follows:

$$\alpha_1(\mathbf{x}, l) = \text{corr}(f^{(l)}, f^{(L)}) = \frac{\kappa((\mathbf{x}, l), (\mathbf{x}, L))}{\sqrt{\sigma^{2(l)}\sigma^{2(L)}}} \quad (10)$$

$$\alpha_2(\mathbf{x}, l) = 1 - \frac{\sigma_\varepsilon}{\sqrt{\sigma^{2(l)}(\mathbf{x}) + \sigma_\varepsilon^2}} \quad (11)$$

$$\alpha_3(l) = \frac{\lambda^{(L)}}{\lambda^{(l)}} \quad (12)$$

$$\alpha_4(M, l) = \begin{cases} 1 & \text{if } l = 1, \dots, L-1 \\ \frac{M_s}{M_s - M} & \text{if } l = L \quad M_s = 1 \end{cases} \quad (13)$$

where $\alpha_1(\mathbf{x}, l)$ accounts for the reduction of the uncertainty of the aerodynamic output as the level of fidelity increases and is defined as the correlation between the evaluation of the objective at l -th fidelity and the high-fidelity evaluation; $\alpha_2(\mathbf{x}, l)$ is conceived to consider the reduction of the standard deviation of the posterior GP model after a new observation of the objective function; $\alpha_3(l)$ takes into account the computational cost corresponding to the fidelity of the aerodynamic

model and is defined as the ratio between the cost per evaluation of the high-fidelity model and the cost associated with the l -th fidelity model. For the specific case of the NACA 0012 drag coefficient optimization (Section II.A), we set $\lambda^{(L=3)} = 1$ for the high-fidelity model (Section II.B), $\lambda^{(l=2)} = 1/2$ for the mid-fidelity model (Section II.C), and $\lambda^{(l=1)} = 1/1000$ for the low-fidelity model (Section II.D).

The utility function $\alpha_4(M, l)$ is specifically designed to consider a multi-regime aerodynamic domain: as the Mach number gets closer to the speed of sound, the flow around the airfoil moves to the transonic regime generating local shock waves that can cause a large-scale separation downstream with an increase of the aerodynamic drag and unsteadiness in the flow around the airfoil. Therefore, we conceived α_4 to capture these phenomena with an high degree of accuracy, encouraging the selection of the highest-fidelity aerodynamic model when the Mach number approaches the sonic condition. As the values of M becomes close to the transonic region ($M = 0.8$), α_4 achieves large values for the highest-fidelity model while is imposed at 1 for the $l = 1, \dots, L - 1$ fidelity models. This allows to increase the values of the acquisition function when the objective function is observed with the L fidelity model in the transonic regime, ensuring the maximum accuracy of the airflow modeling.

To summarize, each utility function is beneficial to the performance of our multifidelity framework. If α_1 is excluded from the formulation, the algorithm queries always the low-fidelity aerodynamic model to contain the total computational cost of the optimization process. If α_2 is not considered, the algorithm samples regions of the design space already explored consuming the computational budget to evaluate unnecessary replicates. If α_3 is discarded, the computational cost of the aerodynamic models does not affect the selection of the level of fidelity to query, making more advantageous the selection of the high-fidelity model. If α_4 is not included, the algorithm is not informed about the evolution of the aerodynamic domain, and potentially selects the level of fidelity not appropriate to predict the physical phenomena that characterize a specific aerodynamic regime within the multi-regime scenario of optimization.

C. Two-steps Lookahead Multifidelity Acquisition Function

The multifidelity acquisition function depicted in Equation (8) allows to include information about the fluid dynamic regime in the selection of the aerodynamic model to query at each iteration. However, the approach to the sampling process limits the evaluation of the rewards of the solution quality to the immediate step of optimization. In the following, we define a two-step lookahead sampling policy that is able to consider the current informative gains in the design solution given a potential candidate, and leverage this information to find an even better design on the next query. The goal is to explore the high uncertainty region of the design space looking ahead into future evaluations to find improvements in the solution quality, while containing the computational cost through the principled use of inexpensive aerodynamic analyses conducted with lower fidelity models.

Considering a generic iteration i of the optimization, let us index with 0 the quantities already collected at the current iteration of the optimization, and with 1 and 2 the quantities referred to the first step and second step, respectively. The

two-step lookahead acquisition function is defined following the dynamic programming principle: an optimal policy has the property that whatever the initial training set and utility function, the decision variables selected over the remaining steps must be optimal for the remaining problem, considering the training set of the early step as the initial condition. For our case, the objective is to select the optimal design variable at the second step lookahead \mathbf{x}_2 to minimize the expected loss evaluated at the first step \mathbb{L}_1 :

$$\mathbb{L}_1 = \min_{\mathbf{x}_2 \in \mathcal{X}} \mathbb{E}_1 [f_2^*] = f_1^* - \max_{\mathbf{x}_2 \in \mathcal{X}} MFDA_1(\mathbf{x}_2, l, M) \quad (14)$$

where \mathbb{E}_1 is the expectation computed with respect to the multifidelity Gaussian process updated at the first step ahead, f_1^* and f_2^* are respectively the values of the objective corresponding to the optimal samples at the first and second step ahead, and $MFDA_1$ is the domain-aware multifidelity acquisition function (Equation 8) computed on the posterior mean $\mu_1^{(L)}$ (Equation 6) and variance $\sigma_1^{2(L)}$ (Equation 7) of the multifidelity Gaussian process at the first step. Therefore, the expected value \mathbb{E}_0 achieved at the start of the first step is:

$$\begin{aligned} \mathbb{E}_0 [\mathbb{L}_1] &= \mathbb{E}_0 \left[f_1^* - \max_{\mathbf{x}_2 \in \mathcal{X}} MFDA_1(\mathbf{x}_2, l, M) \right] = \\ &= \mathbb{E}_0 \left[f_0^* - (f_0^* - \min_{\mathbf{x}_1} f(\mathbf{x}_1))^+ - \max_{\mathbf{x}_2 \in \mathcal{X}} MFDA_1(\mathbf{x}_2, l, M) \right] = \\ &= f_0^* - MFDA_0(\mathbf{x}_1, l, M) - \mathbb{E}_0 \left[\max_{\mathbf{x}_2 \in \mathcal{X}} MFDA_1(\mathbf{x}_2, l, M) \right] \end{aligned} \quad (15)$$

where f_0^* is the value of the objective at the optimal sample identified before the lookahead steps, and $MFDA_0$ is the domain-aware multifidelity acquisition function (Equation 8) under the multifidelity Gaussian process with mean $\mu_0^{(L)}$ (Equation 6) and variance $\sigma_0^{2(L)}$ (Equation 7) evaluated at the current iteration of the optimization.

Combining Equation 14 with Equation 15, we define the two-steps lookahead acquisition function as:

$$U(\mathbf{x}_1, l, M) = MFDA_0(\mathbf{x}_1, l, M) + \mathbb{E}_0 \left[\max_{\mathbf{x}_2 \in \mathcal{X}} MFDA_1(\mathbf{x}_2, l, M) \right] \quad (16)$$

Equation 16 cannot be computed in closed form as it requires to evaluate the nested expectation. To estimate U we use the Monte Carlo approach, assuming that $f^{(l)}(\mathbf{x}_1) = \mu_0^{(l)} + \mathbf{C}_0^{(l)}(\mathbf{x}_1)Z$, where l is the level of fidelity of the aerodynamic model, $\mathbf{C}_0^{(l)}$ is the Cholesky decomposition of \mathbf{K}_0 , and Z is the independent standard normal random variable [39]. Thus, we define the predictive distribution of the multifidelity Gaussian process at a generic input \mathbf{x} after the first step:

$$\mu_1^{(l)} = \mu_0^{(l)} + \mathbf{A}_0^{(l)}(\mathbf{x})Z \quad (17)$$

$$\sigma_1^{(l)} = \sigma_0^{(l)} - \mathbf{A}_0^{(l)}(\mathbf{x}) \mathbf{A}_0^{(l)}(\mathbf{x})^T \quad (18)$$

where $\mathbf{A}_0^{(l)}(\mathbf{x}) = \kappa_0^{(l)}(\mathbf{x}) \mathbf{C}_0^{(l)-1}$. To estimate $U(\mathbf{x}_1, l, M)$, we sample Z and compute $U(\mathbf{x}_1, l, M, Z) = MFDA_0(\mathbf{x}_2, l, M) + \mathbb{E}_0[\max_{\mathbf{x}_2 \in \mathcal{X}} MFDA_1(\mathbf{x}_2, l, M, Z)]$. Averaging many such replications provides a robust approximation of $U(\mathbf{x}_1, l, M)$.

D. Aerodynamic Constrained Optimization Algorithm

This section illustrates the main steps and features of our overall non-myopic domain-aware multifidelity scheme for constrained aerodynamic optimization. We consider the generic problem of optimizing a quantity of interest $f(\mathbf{x})$ that characterize the aerodynamic domain subject to constraints $f_C(\mathbf{x})$, where the design variables \mathbf{x} include the Mach number M defined within a range that denotes a multi-regime scenario (Section II) and multiple representations at different levels of fidelity $l = 1, \dots, L$ are available for both $f^{(l)}$ and $f_C^{(l)}$. During the optimization process, the objective function and the constraints are approximated at each iteration to compute the multifidelity acquisition function and select the most promising design configuration and level of fidelity to query. For this estimation task, our algorithm uses the multifidelity Gaussian process illustrated in Section III.A to provide a probabilistic representation of the objective function $f^{(l)} = GP(\mu^{(l)}, \sigma^{(l)})$ and of the constraints $f_C^{(l)} = GP(\mu_C^{(l)}, \sigma_C^{(l)})$. To handle the constraints during the optimization process, we adopt a conditional selection strategy that defines an additional surrogate model $GP(\tilde{\mu}^{(l)}, \tilde{\sigma}^{(l)})$ to represent the feasible regions of the design space, where $\tilde{\mu}^{(l)} = \mu^{(l)}(\mu_C^{(l)} > f_C)$ and $\tilde{\sigma}^{(l)} = \sigma^{(l)}(\mu_C^{(l)} > f_C)$. Estimated the feasible subset of design configurations, the next step of our multifidelity scheme is the identification of the next design configuration and the selection of the level of fidelity to query maximizing the two-steps lookahead multifidelity acquisition function presented in Section III.C computed on $\tilde{\mu}^{(l)}$ and $\tilde{\sigma}^{(l)}$.

With reference to the test-case optimization problem defined by Equation 1 and related to the design optimization of a 2D airfoil, the algorithmic flow (Algorithm 1) starts from a first sample set of feasible design points \mathcal{P} that satisfy the constraint on the lift coefficient C_l (Equation 1a) and on the flight altitude h (Equation 1b), together with the multifidelity objective function $C_d(\mathbf{x})$ and the prior distribution of the multifidelity Gaussian process $GP(0, \sigma^{(l)})$. The feasibility of the set \mathcal{P} is not a strong assumption but allows to improve the performance of optimization by enabling the initial sampling process in the regions of the design space that are already feasible, increasing the likelihood that the algorithm endures the sampling of feasible design configurations. Moreover, if the initial set of points \mathcal{P} is not feasible, the algorithm may require several iterations to accurately characterize the constraints consuming computational budget, for instance when the feasible regions of the design space are disconnected.

At the beginning of the optimization process, we select the initial sampling points $P_0 = [\mathbf{x}_1, \dots, \mathbf{x}_n]$ and the associated levels of fidelity $\mathcal{L}_0 = [l_1, \dots, l_n]$, which are used to compute the first observations of the drag coefficient $\mathbf{C}_d = [C_{d_1}^{(l_1)}, \dots, C_{d_n}^{(l_n)}]$ and of the lift coefficient $\mathbf{C}_l = [C_{l_1}^{(l_1)}, \dots, C_{l_n}^{(l_n)}]$ and hence learn the initial surrogate models of the

Algorithm 1: Constrained multifidelity aerodynamic optimization of an airfoil

Input: Feasible set $\mathcal{P} \in \mathbb{R}^2$, multifidelity objective function $C_d(\mathbf{x})$ and the multifidelity Gaussian process prior $GP(0, \kappa^{(l)}(\mathbf{x}, \mathbf{x}'))$

Output: Set of design variables \mathbf{x} that minimize $C_d(\mathbf{x})$

Draw initial sampling points $P_0 \subset \mathcal{P}$ and associated levels of fidelity \mathcal{L}_0

$\mathbf{C}_d \leftarrow$ evaluate $C_d^{(l)}$ at points P_0 and level of fidelity $l \in \mathcal{L}_0$

$\mathbf{C}_l \leftarrow$ evaluate $C_l^{(l)}$ at points P_0 and level of fidelity $l \in \mathcal{L}_0$

Compute posterior $\mu^{(l)}$ and $\sigma^{2(l)}$ based on \mathbf{C}_d

Compute posterior $\mu_C^{(l)}$ and $\sigma_C^{2(l)}$ based on \mathbf{C}_l

$i \leftarrow 0$

repeat

- `// Function and constraint evaluation`
- Load \mathbf{x}_i
- $C_{d_i}^{(l_i)} \leftarrow$ evaluate $C_d^{(l)}$ at the point \mathbf{x}_i and level of fidelity l_i
- $\mathbf{C}_d = [\mathbf{C}_d, C_{d_i}^{(l_i)}]$
- $C_{l_i}^{(l_i)} \leftarrow$ evaluate $C_l^{(l)}$ at the point \mathbf{x}_i and level of fidelity l_i
- $\mathbf{C}_l = [\mathbf{C}_l, C_{l_i}^{(l_i)}]$
- `// Update the surrogate models`
- Update posterior $\mu^{(l)}$ and $\sigma^{2(l)}$ based on \mathbf{y}
- Update posterior $\mu_C^{(l)}$ and $\sigma_C^{2(l)}$ based on \mathbf{y}_C
- `// Next iteration`
- Compute $\tilde{\mu}^{(l)} \subset \mu^{(l)} (\mu_C^{(l)} > C_l^0)$ and $\tilde{\sigma}^{(l)} \subset \sigma^{(l)} (\mu_C^{(l)} > C_l^0)$
- Estimate U based on $\tilde{\mu}^{(l)}$ and $\tilde{\sigma}^{(l)}$
- $[\mathbf{x}_{i+1}, l_{i+1}] \leftarrow \max_{\mathbf{x}} (U(\mathbf{x}, l, M)) \forall l$
- $i + 1 \leftarrow i$

until $B \leq B_{max}$;

return \mathbf{x} that minimize $C_d(\mathbf{x})$ over the whole dataset

objective function $GP(\mu^{(l)}, \sigma^{(l)})$ and of the constraint $GP(\mu_C^{(l)}, \sigma_C^{(l)})$.

Considering a generic iteration i of our optimization framework, we select the design point \mathbf{x}_i at which we compute the value of the objective function y_i and of the constraint y_{Ci} from the previous iteration $i - 1$. These observations are used to update the surrogate multifidelity models for both the drag and lift coefficient, and compute the constrained posterior mean $\tilde{\mu}^{(l)}$ and standard deviation $\tilde{\sigma}^{(l)}$ on which we define the acquisition function. At this point, the algorithm estimates and optimizes the two-steps lookahead multifidelity acquisition function through the procedure described in Section III.C, determining the next design \mathbf{x}_{i+1} to evaluate and associated fidelity of the aerodynamic model l_{i+1} . The process is then iterated until the cumulative cost of the evaluation of the aerodynamic models $B = \sum \lambda_i^{(l)}$ reaches a prescribed maximum budget B_{max} .

IV. Results and Discussion

This section discusses the results we obtained with the two-step lookahead domain-aware multifidelity framework (lookMFBO-DA) for the constrained aerodynamic optimization problem of a NACA 0012 airfoil subject to a multi-regime

flow-field (Section II.A). The proposed strategy is compared with Bayesian frameworks widely used in literature including efficient global optimization (EGO) [38], a two-step lookahead single fidelity framework (lookEGO) [11, 13–15], MFBO based on the multifidelity expected improvement acquisition function (MFBO) [34], MFBO implementing the domain-aware acquisition function only (MFBO-DA), and a two-step lookahead multifidelity framework (lookMFBO) realized excluding Equation (13) from our multifidelity acquisition function illustrated in Section III.C. All these competing algorithms have been implemented in-house by the authors to ensure a fair comparison and discussion of the results.

The multifidelity algorithms consider three levels of fidelity for the aerodynamic modeling including the high-fidelity CFD model (Section II.B), the mid-fidelity CFD model (Section II.C), and the low-fidelity model (Section II.D), while the single fidelity frameworks elicitate evaluations of the high-fidelity model only. For the numerical experiments discussed in this paper, the initial set of samples is determined through a Latin hypercube strategy and are used to compute the surrogate model at the first step (Section III.D). Specifically, the multifidelity searches are initialized with 13 initial design configurations among which 10 observations of the objective function are evaluated with the low-fidelity model, 2 are computed with the mid-fidelity model, and the last one is evaluated with the high-fidelity model, while for the single fidelity frameworks we consider 2 initial design configurations at which we compute high-fidelity observations. The computational cost B is evaluated as the sum of the computational costs required to estimate the aerodynamic domain with a certain level of fidelity. To total computational budget assigned is fixed at $B_{max} = 10$, that is equivalent to 10 evaluations of the high fidelity model ($\lambda^{(L)} = 1$, Section III.B). The overall budget includes also the cost for the initial samples. All the competing methods are based on the Gaussian process surrogate model [31] and its implementation in the multifidelity scenario (Section III.A). We adopt the square exponential kernel for the covariance function of the GP models and provide an estimate of the hyperparameters using the maximum likelihood estimation technique [40]. Considering the lookMFBO-DA and lookMFBO algorithms, we sample 2000 random variables to approximate the two-step lookahead acquisition function through the Monte Carlo strategy depicted in Section III.C, ensuring the robustness of the estimate of Equation (16) while keeping acceptable the required computational cost.

Figure 4a illustrates the values of the minimum drag coefficient $C_d^* = \min(C_d(\mathbf{x}))$ while Figure 4b provides details about the error of the surrogate model obtained from a randomized statistic of 20 experiments starting from different initial samples for all the algorithms. The evaluation of the reduction of C_d allows to assess the accuracy of the optimal solution identified by the competing methods, providing a quantification of the improvement in the solution of the optimization problem. The measure of the surrogate error permits to investigate how wisely the algorithms combine data at different levels of fidelity to improve the accuracy of the surrogate model; this reflects the capability of the optimization methods to better inform the acquisition function with a reliable estimate of the objective function, leading to superior optimization performance. The outcomes are reported in terms of the median values (solid line) together with the observations falling in the interval between the 25-th and 75-th percentiles (shaded area). The error of the

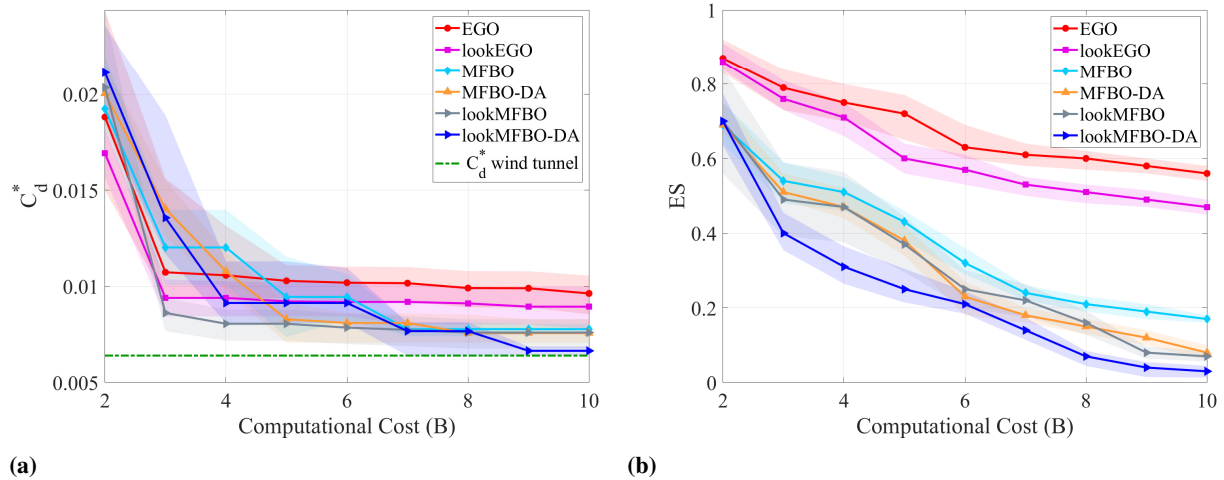


Fig. 4 Statistics over 20 runs of the minimum drag coefficient C_d^* (4a) and of the error of the surrogate model ES computed as the Gaussian process of the objective function (4b) obtained with efficient global optimization (EGO) [38], two-step lookahead single-fidelity Bayesian optimization (lookEGO) [11, 13–15], multifidelity Bayesian optimization (MFBO) [34], domain-aware multifidelity Bayesian optimization (MFBO-DA), two-step lookahead multifidelity framework (lookMFBO), and the proposed two-step lookahead domain-aware multifidelity framework (lookMFBO-DA).

surrogate model of the objective function ES is defined as the relative error between the values of the drag coefficient \tilde{C}_d predicted by the surrogate model of the objective function (Section III.A) and the values of $C_{d,wt}$ determined through real world NACA 0012 wind tunnel data [30], since the high fidelity model (Section II.B) has been validated with this experiments:

$$ES = \frac{1}{n_{wt}} \sum_{i=1}^{n_{wt}} \frac{\tilde{C}_d(\mathbf{x}_{i,wt}) - C_{d,wt}(\mathbf{x}_{i,wt})}{C_{d,wt}(\mathbf{x}_{i,wt})} \quad (19)$$

where $n_{wt} = 120$ is the number of experimental data available and $\mathbf{x}_{i,wt}$ is the i -th design configuration evaluated in the wind tunnel experiments. The experimental optimum is $C_{d,wt}^* = 0.0064$ (dashed line in Figure 4a) corresponding to the optimum design $\mathbf{x}_{wt}^* = [M_{wt}^*, \alpha_{wt}^*] = [0.6574, 3.013]$, and is also estimated through wind tunnel experiments. The results suggest that our optimization strategy is substantially more effective than competing methods, providing larger improvements in the solution quality (reduction of the drag coefficient and of the surrogate error) earlier on in the first stages of the optimization process. Moreover, when the optimizer runs out of budget, it can be observed that our framework permits to identify a design solution superior to other methods: the two-step lookahead formulation, the multifidelity strategy and the active awareness of the fluid regime jointly allow to efficiently learn the aerodynamic surrogate and minimize the drag coefficient. In addition, it can be observed that including forms of domain awareness allows to obtain sensitive improvements of the accuracy of the multifidelity surrogate thanks to the wise use high-fidelity models at higher mach number regimes. Table 1 depicts the minimum median values of the drag coefficient achieved by

| Method | mean C_d^* | best C_d^* | best M^* | best α^* | HF calls | MF calls | LF calls |
|---------------------|--------------|--------------|------------|-----------------|----------|----------|----------|
| EGO | 0.009631 | 0.008586 | 0.6210 | 2.622 | 10 | - | - |
| lookEGO | 0.008943 | 0.007590 | 0.6580 | 3.192 | 10 | - | - |
| MFBO | 0.007774 | 0.007423 | 0.6490 | 3.230 | 4 | 11 | 497 |
| MFBO-DA | 0.007568 | 0.007066 | 0.6360 | 3.078 | 5 | 8 | 988 |
| lookMFBO | 0.007597 | 0.006768 | 0.6334 | 2.959 | 4 | 10 | 997 |
| lookMFBO-DA | 0.006648 | 0.006427 | 0.6563 | 3.000 | 4 | 10 | 993 |
| Wind Tunnel Optimum | - | 0.0064 | 0.6574 | 3.013 | - | - | - |

Table 1 Comparison across median and best values of the minimum drag coefficient C_d^* over 20 experiments for the competing algorithms, including the related optimum design variables M^* and α^* , and the queries of the high-fidelity (HF), mid-fidelity (MF) and low-fidelity (LF) model.

| Error C_d^* | EGO | lookEGO | MFBO | MFBO-DA | lookMFBO | lookMFBO-DA |
|---------------|-------|---------|-------|---------|----------|-------------|
| Mean | 50.5% | 39.7% | 21.5% | 18.2% | 18.7 % | 3.88% |
| Best | 34.2% | 18.6% | 15.9% | 10.4% | 5.75 % | 0.422% |

Table 2 Comparison of the percentage error of the estimate of the minimum drag coefficient C_d^* with respect to the wind tunnel optimum $C_{d,wt}^*$: mean and best values over the 20 trials.

the algorithms, and reports the drag coefficient along with the design configurations and the queries of the aerodynamic models corresponding to the experiments that give the best results. Our framework allows to achieve a design solution almost identical to the best design from wind tunnel data using 993 observations of the low-fidelity model, 10 evaluations of the mid-fidelity model and only 4 queries of the high-fidelity model. Table 2 indicates the percentage errors computed on the mean and on the best values of the drag coefficient determined by each competing methods with reference to the wind tunnel optimum. On average, the lookMFBO-DA strategy permits to achieve an error of about 3.88%, whereas the second-best lookMFBO algorithm achieves an error of the 5.75%. Moreover, our framework scores the overall maximum reduction of the error with a value of about 0.422%. This demonstrates the capability of the lookMFBO-DA to leverage multiple aerodynamic models that jointly with the lookahead property and the domain-awareness enhance the optimization performances.

Figure 5 depicts the design space obtained using a spline interpolation of wind tunnel data to illustrate the high-fidelity sampling process corresponding to the experiments that provide the best results in terms of reduction of the drag coefficient for all the algorithms considered. It can be observed that the multifidelity strategies allow to reduce the total number of high-fidelity observations respect to the single-fidelity algorithms, leveraging the information from the mid-fidelity and low-fidelity models. The proposed lookMFBO-DA framework uses the high-fidelity model to explore the regions of the design space characterized by higher Mach numbers and angles of attack thanks to the domain-awareness property, and effectively exploit the design space sampling closely to the optimum configuration through the lookahead feature (Figure 5f). These aspects permit to accurately predict the aerodynamic domain at

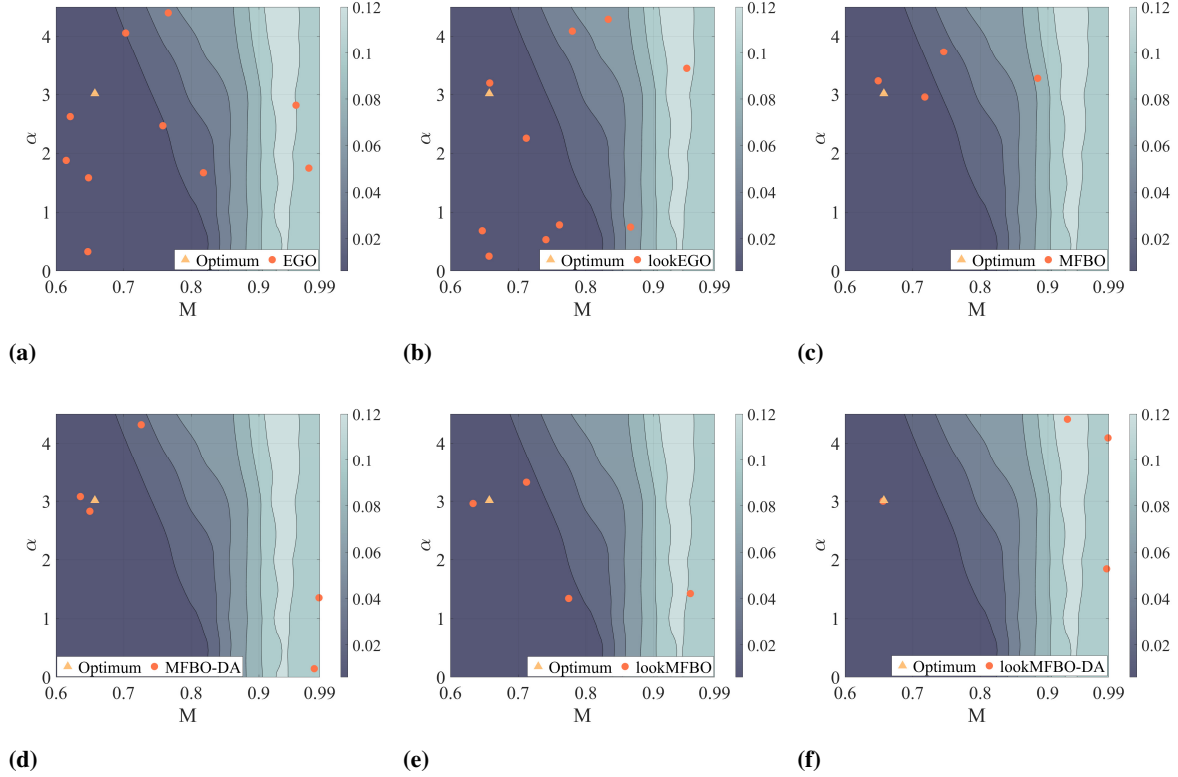


Fig. 5 High-fidelity evaluations of C_d called by the different algorithms: EGO (5a), lookEGO (5b), MFBO (5c), MFBO-DA (5d), lookMFBO (5e), and lookMFBO-DA (5f). The contour map indicates the C_d value obtained interpolating the wind tunnel data as a function of Mach number M and angle of attack α .

transonic regimes taking advantage of lower-fidelity models to explore design configurations at lower speeds and angle of attack. Also the MFBO-DA algorithm samples the design space with a similar approach thanks to the domain-aware formulation (Figure 5d). However, the exploitation is not effective as lookMFBO-DA demonstrating the significant role of the lookahead strategy to improve the optimization performance.

For all the numerical experiments presented in this paper, we run each test on a single core of a desktop PC with Intel Core i7-8700 (3.2 GHz) and 32 GB of RAM. The computational time of a single iteration of the lookMFBO-DA framework ranges from seconds to minutes, while the computation of the aerodynamic model takes 60 min, 30 min and 0.06 min for the high-fidelity, mid-fidelity and low-fidelity model, respectively.

The outcomes highlight the potential impact of our original framework for the aerodynamic design and optimization problems where the use of high-fidelity models is most of times computationally unfeasible, but essential to represent the full order complexity of the aerodynamic domain in certain fluid regimes. Therefore, a multifidelity framework that is conscious about the evolution of the physics domain can improve the wise inclusion of high-fidelity data during the optimization, and a non-myopic formulation can speed-up the search for the optimum containing the total computational cost.

V. Concluding remarks

This paper introduces a novel non-myopic multifidelity framework for aerodynamic optimization problems characterized by multi-regime scenarios where the range of variability of the Mach number determines a dynamic evolution of the fluid domain. Within the originality of the proposal, our framework allows to select the most promising design configuration through a two-step lookahead sampling policy that permits to measure the long term improvement of the solution of the optimization problem, and determines the level of fidelity to query considering the evolution of the aerodynamic regime. The objective is to propose a computational technique that allows to include expensive high-fidelity models in the multi-regime aerodynamic optimization, combining observations from multiple aerodynamic models at different levels of accuracy and associated expense to contain the total computational cost.

Our strategy is based on multifidelity Bayesian optimization that permits to synthesize observations of the objective function at different levels of fidelity into a unique surrogate model, and at each iteration maximizes a multifidelity acquisition function selecting the next design configuration and the associated level of fidelity of the aerodynamic model to compute. We adopt multifidelity Gaussian processes as the surrogate model, and illustrate an original formulation of the multifidelity acquisition function specifically conceived for the multi-regime constrained aerodynamic optimization. In particular, the distinguishing features of our formulation of the multifidelity acquisition function are: (i) an original optimal policy that maximizes the overall informative gain obtainable over two-steps of the optimization process evaluating a certain design configuration and associated level of fidelity, and (ii) a domain-aware utility function sensitive to the aerodynamic regime through the dependency of the Mach number. We observe these elements permit to enhance the performances of the optimization process computing an efficient aerodynamic surrogate model thanks to the wise evaluation of design configurations considering future scenarios, and the principled selection of the accuracy of the aerodynamic model encouraging the query of the high-fidelity representation when is required to reliably estimate compressibility and non-linear effects at high speed regimes.

We demonstrate our framework and validate our results against wind tunnel data for the multi-regime constrained aerodynamic optimization of an airfoil subject to maintain a minimum lift coefficient at a fixed flight altitude. Our non-myopic domain-aware multifidelity framework is compared against standard multifidelity and single-fidelity Bayesian frameworks. The results show that our strategy outperforms the competing algorithms allowing for a significant reduction of the drag coefficient with reduced computational expense through the construction and use of an efficient aerodynamic surrogate model. Specifically, the proposed framework allows to lower the error on the identification of the minimum drag coefficient down to the 3.88% on average respect to the experimental optimum, which goes down to 0.422% for the best test case over a statistic of 20 experiments conducted. This is achieved thanks to the two key features of the proposed multifidelity acquisition function which allow to consider the evolution of the fluid regime through the Mach number by including domain awareness in the formulation, and to account for the anticipated gain through the two-step lookahead scheme that permits the efficient exploration of the design space.

Future developments will consider the extension of the proposed framework for the multidisciplinary design optimization of aerospace systems that involves the coupling between the fluid domain and other disciplines, as aeroelastic or thermofluid interactions. An additional aspect to explore relates to the implementation of multifidelity domain-aware and non-myopic strategies to reduce the cost of multidisciplinary design optimization problems where the expensive aerodynamic analysis frequently constitute the bottleneck of the optimization process. Further research would explore refined formulations of the domain-aware utility function to tailor the framework to specific aerodynamic design and optimization problems where the knowledge of the fluid domain could sensitively bias the dynamic selection of the level of fidelity of the aerodynamic model to interrogate.

Acknowledgments

The work has been partly conducted under the Visiting Professor Program of Politecnico di Torino. Additional acknowledgments to the Univeristy's Doctoral Scholarship supporting Mr. Francesco Di Fiore. The authors thank Prof. Paolo Maggiore at Politecnico di Torino for the support.

References

- [1] Slotnick, J. P., Khodadoust, A., Alonso, J., Darmofal, D., Gropp, W., Lurie, E., and Mavriplis, D. J., "CFD vision 2030 study: a path to revolutionary computational aerosciences," 2014.
- [2] Kennedy, M. C., and O'Hagan, A., "Predicting the output from a complex computer code when fast approximations are available," *Biometrika*, Vol. 87, No. 1, 2000, pp. 1–13.
- [3] Forrester, A. I., Sóbester, A., and Keane, A. J., "Multi-fidelity optimization via surrogate modelling," *Proceedings of the royal society a: mathematical, physical and engineering sciences*, Vol. 463, No. 2088, 2007, pp. 3251–3269.
- [4] Fernández-Godino, M. G., Park, C., Kim, N.-H., and Haftka, R. T., "Review of multi-fidelity models," *arXiv preprint arXiv:1609.07196*, 2016.
- [5] Park, C., Haftka, R. T., and Kim, N. H., "Remarks on multi-fidelity surrogates," *Structural and Multidisciplinary Optimization*, Vol. 55, No. 3, 2017, pp. 1029–1050.
- [6] Peherstorfer, B., Willcox, K., and Gunzburger, M., "Survey of multifidelity methods in uncertainty propagation, inference, and optimization," *Siam Review*, Vol. 60, No. 3, 2018, pp. 550–591.
- [7] Reienthel, P. H., and Allen, T. T., "Application of multifidelity expected improvement algorithms to aeroelastic design optimization," *10th AIAA Multidisciplinary Design Optimization Conference*, 2014, p. 1490.
- [8] Meliani, M., Bartoli, N., Lefebvre, T., Bouhlel, M.-A., Martins, J. R., and Morlier, J., "Multi-fidelity efficient global optimization: Methodology and application to airfoil shape design," *AIAA aviation 2019 forum*, 2019, p. 3236.

- [9] Mondal, S., Joly, M. M., and Sarkar, S., “Multi-fidelity global-local optimization of a transonic compressor rotor,” *Turbo Expo: Power for Land, Sea, and Air*, Vol. 58585, American Society of Mechanical Engineers, 2019, p. V02DT46A020.
- [10] Kontogiannis, S. G., Demange, J., Savill, A. M., and Kipouros, T., “A comparison study of two multifidelity methods for aerodynamic optimization,” *Aerospace Science and Technology*, Vol. 97, 2020, p. 105592.
- [11] Ginsbourger, D., and Le Riche, R., “Towards Gaussian process-based optimization with finite time horizon,” *mODa 9—Advances in Model-Oriented Design and Analysis*, Springer, 2010, pp. 89–96.
- [12] Lam, R., Willcox, K., and Wolpert, D. H., “Bayesian optimization with a finite budget: An approximate dynamic programming approach,” *Advances in Neural Information Processing Systems*, Vol. 29, 2016, pp. 883–891.
- [13] Osborne, M. A., Garnett, R., and Roberts, S. J., “Gaussian processes for global optimization,” *3rd international conference on learning and intelligent optimization (LION3)*, Citeseer, 2009, pp. 1–15.
- [14] González, J., Osborne, M., and Lawrence, N., “GLASSES: Relieving the myopia of Bayesian optimisation,” *Artificial Intelligence and Statistics*, PMLR, 2016, pp. 790–799.
- [15] Wu, J., and Frazier, P., “Practical two-step lookahead Bayesian optimization,” *Advances in neural information processing systems*, Vol. 32, 2019, pp. 9813–9823.
- [16] Ghoreishi, S. F., and Allaire, D., “Multi-information source constrained Bayesian optimization,” *Structural and Multidisciplinary Optimization*, Vol. 59, No. 3, 2019, pp. 977–991.
- [17] Powell, W. B., *Approximate Dynamic Programming: Solving the curses of dimensionality*, Vol. 703, John Wiley & Sons, 2007.
- [18] Bertsekas, D., *Dynamic programming and optimal control: Volume I*, Vol. 1, Athena scientific, 2012.
- [19] Thibert, J., Grandjacques, M., Ohman, L., et al., “NACA 0012 airfoil,” *AGARD Advisory Report*, Vol. 138, 1979.
- [20] Rumsey, C., Smith, B., and Huang, G., “2D NACA 0012 airfoil validation case,” *Turbulence Modeling Resource, NASA Langley Research Center*, 2014, p. 33.
- [21] Lam, R., Allaire, D. L., and Willcox, K. E., “Multifidelity optimization using statistical surrogate modeling for non-hierarchical information sources,” *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2015, p. 0143.
- [22] Ghoreishi, S. F., and Allaire, D. L., “A fusion-based multi-information source optimization approach using knowledge gradient policies,” *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2018, p. 1159.
- [23] Bertram, A., and Zimmermann, R., “Theoretical investigations of the new Cokriging method for variable-fidelity surrogate modeling,” *Advances in Computational Mathematics*, Vol. 44, No. 6, 2018, pp. 1693–1716.
- [24] Zhou, Q., Wu, Y., Guo, Z., Hu, J., and Jin, P., “A generalized hierarchical co-Kriging model for multi-fidelity data fusion,” *Structural and Multidisciplinary Optimization*, Vol. 62, No. 4, 2020, pp. 1885–1904.

- [25] Liem, R. P., Mader, C. A., and Martins, J. R., “Surrogate models and mixtures of experts in aerodynamic performance prediction for aircraft mission analysis,” *Aerospace Science and Technology*, Vol. 43, 2015, pp. 126–151.
- [26] Ghosh, S., Jacobs, R., and Mavris, D. N., “Multi-source surrogate modeling with bayesian hierarchical regression,” *17th AIAA Non-Deterministic Approaches Conference*, 2015, p. 1817.
- [27] Economon, T. D., Palacios, F., Copeland, S. R., Lukaczyk, T. W., and Alonso, J. J., “SU2: An open-source suite for multiphysics simulation and design,” *Aiaa Journal*, Vol. 54, No. 3, 2016, pp. 828–846.
- [28] Drela, M., “XFOIL: An analysis and design system for low Reynolds number airfoils,” *Low Reynolds number aerodynamics*, Springer, 1989, pp. 1–12.
- [29] Barrett, R., and Ning, A., “Comparison of airfoil precomputational analysis methods for optimization of wind turbine blades,” *IEEE Transactions on Sustainable Energy*, Vol. 7, No. 3, 2016, pp. 1081–1088.
- [30] Harris, C. D., “Two-dimensional aerodynamic characteristics of the NACA 0012 airfoil in the Langley 8 foot transonic pressure tunnel,” 1981.
- [31] Rasmussen, C. E., “Gaussian processes in machine learning,” *Summer school on machine learning*, Springer, 2003, pp. 63–71.
- [32] Viana, F. A., Haftka, R. T., and Watson, L. T., “Efficient global optimization algorithm assisted by multiple surrogate techniques,” *Journal of Global Optimization*, Vol. 56, No. 2, 2013, pp. 669–689.
- [33] Lam, R., Allaire, D. L., and Willcox, K. E., “Multifidelity optimization using statistical surrogate modeling for non-hierarchical information sources,” *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2015, p. 0143.
- [34] Huang, D., Allen, T. T., Notz, W. I., and Miller, R. A., “Sequential kriging optimization using multiple-fidelity evaluations,” *Structural and Multidisciplinary Optimization*, Vol. 32, No. 5, 2006, pp. 369–382.
- [35] Zhang, Y., Hoang, T. N., Low, B. K. H., and Kankanhalli, M., “Information-based multi-fidelity Bayesian optimization,” *NIPS Workshop on Bayesian Optimization*, 2017.
- [36] Takeno, S., Fukuoka, H., Tsukada, Y., Koyama, T., Shiga, M., Takeuchi, I., and Karasuyama, M., “Multi-fidelity Bayesian optimization with max-value entropy search and its parallelization,” *International Conference on Machine Learning*, PMLR, 2020, pp. 9334–9345.
- [37] Kennedy, M. C., and O’Hagan, A., “Predicting the output from a complex computer code when fast approximations are available,” *Biometrika*, Vol. 87, No. 1, 2000, pp. 1–13.
- [38] Jones, D. R., Schonlau, M., and Welch, W. J., “Efficient global optimization of expensive black-box functions,” *Journal of Global optimization*, Vol. 13, No. 4, 1998, pp. 455–492.
- [39] Wilson, J. T., Hutter, F., and Deisenroth, M. P., “Maximizing acquisition functions for Bayesian optimization,” *arXiv preprint arXiv:1805.10196*, 2018.

- [40] Sobester, A., Forrester, A., and Keane, A., *Engineering design via surrogate modelling: a practical guide*, John Wiley & Sons, 2008.