



HAL
open science

A Confidence-based Aerospace Design Approach Robust to Structural Turbulence Closure Uncertainty

Giulio Gori, Olivier Le Maitre, Pietro Marco Congedo

► **To cite this version:**

Giulio Gori, Olivier Le Maitre, Pietro Marco Congedo. A Confidence-based Aerospace Design Approach Robust to Structural Turbulence Closure Uncertainty. *Computers and Fluids*, 2022, 246, pp.105614. 10.1016/j.compfluid.2022.105614 . hal-03754147

HAL Id: hal-03754147

<https://hal.science/hal-03754147>

Submitted on 19 Aug 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Confidence-based Aerospace Design Approach Robust to Structural Turbulence Closure Uncertainty

G. Gori^{c,a,*}, O. Le Maître^{c,a,b}, P. M. Congedo^{c,a}

^a*INRIA, France*

^b*CNRS, France*

^c*Platon Team, Centre de Mathématiques Appliquées, École Polytechnique, IPP, Route de Saclay, 91128 Palaiseau, France*

Abstract

We propose a confidence-based design approach robust to turbulence closures model-form uncertainty in Reynolds-Averaged Navier-Stokes computational models. The Eigenspace Perturbation Method is employed to compute turbulence closure uncertainty estimates of the performance targeted by the optimizer. The magnitude of the uncertainty estimates is exploited to establish an indicator parameter associated to the credibility of numerical prediction. The proposed approach restricts the optimum search only to design space regions for which the credibility indicator suggests trustworthy RANS model predictions. In this way, we improve the efficiency of the design process, potentially avoiding designs for which the computational model is unreliable. The reference test case consists in a two-dimensional single element airfoil resembling a morphing wing section in a high-lift configuration. Results show that the prediction credibility constraint has a non negligible impact on the definition of the optimal design.

Keywords: Robust optimization, RANS model uncertainty, Eigenspace Perturbation Method, Turbulence closure uncertainty, Aerodynamic design

*Corresponding author, e-mail: giulio.gori@polimi.it, Permanent Address: Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156, Milano, Italy.

1. Introduction

The relevance of fluid dynamics sciences is undisputed in many industrial fields, from aeronautics to automotive, including naval applications, biology, medicine, civil engineering, energy production and many other. In this perspective, Computational Fluid Dynamics (CFD) aims at developing reliable numerical models that can predict reality with a reasonable level of accuracy. A popular approach in CFD consists in using the Reynolds-Averaged Navier-Stokes equations (RANS), an undetermined set of partial differential equations enforcing the conservation of fundamental quantities in time-averaged terms. In the RANS context, turbulence closures are needed to reconstruct the so-called Reynolds stress term. Unfortunately, the strong inherent model-form assumptions underlying turbulence closures question the credibility of RANS-based predictions for some particular flow configurations [1, 2, 3, 4, 5, 6, 7, 8]. Despite these evidences, the community still largely relies on RANS-based simulations for designing new products.

The objection driving the development of this paper is that the efforts pledged to the design process may be frustrated by the structural uncertainty inherent turbulence closures. Overall, the optimization process is inefficient, potentially exploring portions of the design space where the computational model is unreliable. Although an a posteriori validation is usually possible, it often requires the implementation of expensive experiments and, in any case, it provides nothing but a mere confirmation/rejection of the proposed design, depending on whether it delivers the expected performance or not. That is particularly true in aerospace applications, for which validation experiments involve operating costly wind tunnel facilities and for which high-fidelity simulations (Direct Numerical Simulations) of the full aircraft model are just prohibitive.

As stated in [9], design exploration and optimization under turbulence model-form uncertainty was practically never investigated prior to 2020. In the reference, the authors identify two hurdles. Namely, producing reliable estimates of the turbulence closure model-form uncertainty, and integrating this measure into traditional design optimization frameworks. Concerning the first issue, the high complexity of turbulence phenomena makes the direct quantification of the errors introduced by RANS closures difficult, if not intractable, requiring the implementation of formal and effective Uncertainty Quantification (UQ) techniques. For an updated and comprehensive review of the available approaches, we direct the reader to the Ref. [10].

In this paper, we take advantage of a non-parametric approach, the so-called Eigenspace Perturbation Method (EPM) [2, 3], a formal technique capable of providing an estimation of the uncertainty corresponding to turbulence modeling approximations in RANS closures. The EPM renders up a relatively cheap framework for the estimation of the model-form uncertainty relative to the structural deficiencies of the Reynolds stress tensor in RANS models. As stated in [3], note the deliberate use of the word *estimation*, rather than *quantification*, implying the computation of reasonable and informed uncertainty estimates rather than rigorous and provable bounds.

The aim of this paper is to provide a solution concerning the second hurdle exposed in [9]. Namely, integrating the measure of the turbulence model-form uncertainty into traditional design optimization frameworks by proposing an original approach devised to systematically account for the RANS model credibility limitations. Namely, the optimization problem is formulated under the constraint of that a credibility indicator, representative of the EPM turbulence closure uncertainty estimates, falls below a given threshold. The algorithm is expected to avoid exploring regions of the design space associated to unreliable numerical predictions, increasing its efficiency and ensuring the robustness of the optimal design w.r.t the reliability of RANS predictions. Naturally, the optimization procedure comes at an increased computational cost, inherent to the evaluation of the uncertainty estimates. For the sake of completeness, we remind here that the EPM approach can be formally applied also to obtain the uncertainty estimates associated to the closures provided for modeling the unresolved scales in LES simulation [11]. Therefore, the design approach described in our paper also formally applies to LES-based optimization frameworks.

The relevance of relying on credible computational predictions is undeniable in the aircraft preliminary design phase [12]. At the time of the writing, only three contributions are identified concerning the exploitation of the EPM for developing optimization strategies accounting for the epistemic uncertainty included in the RANS closure models. In Ref. [13], the authors optimize a turbine stator blade by weakly imposing a model error constraint as a penalty term included in the objective function. In Ref. [14], the authors maximize the aerodynamic performances of an airfoil targeting the lowest turbulence model-form uncertainty estimate. In Ref. [9], the optimization still targets the worst turbulence model-form uncertainty estimate, considering several geometries.

In this work, the proposed design strategy is applied to a single element

airfoil resembling a morphing wing section in a high-lift configuration. The paper is structured as follows. Section 2 describes the formulation of the optimization problem constrained to a threshold of some sort of estimator of the RANS model accuracy. Section 3 briefly summarizes the EPM approach, describing how it can be exploited to derive an estimator of the credibility of numerical simulations of turbulent flows. In particular, the computational model representing the baseline test case is presented and a preliminary EPM analysis is reported. Section 4 concerns a detailed description of the optimization procedure from the technical point of view, including the parametrization of the geometry. Sec. 5 reports the results achieved using the proposed optimization strategy, for different problem settings. Eventually, Sec. 6 summarizes the findings reported in this paper.

2. Formulation of the optimization problem

A mono-objective optimization problem is classically formulated as

$$\begin{aligned} \min_{\mathbf{x} \in \Omega \subset \mathbb{R}^d} \quad & f(\mathbf{x}), \\ \text{subject to} \quad & g_i(\mathbf{x}) = \tilde{g}_i, \quad \text{for } i = 1, \dots, N, \\ & h_j(\mathbf{x}) \leq \tilde{h}_j, \quad \text{for } j = 1, \dots, M, \end{aligned} \quad (1)$$

where $f(\mathbf{x})$ is the *objective function*, g_i a set of N equality constraints and h_j a set of M inequality constraints. The target function $f(\mathbf{x})$ depends upon a set of design parameters listed in vector $\mathbf{x} \in \Omega \subset \mathbb{R}^d$, where d is the dimension of the design space. Usually, the design space Ω is bounded by the physics of the problem or by limits set based on the designer's experience.

In this paper, the optimization problem is formulated to maximize the lift produced by a two-dimensional NACA airfoil. By defining the lift coefficient c_l , the optimization problem reads

$$\begin{aligned} \max_{\mathbf{x} \in [0,1]^d} \quad & c_l(\mathbf{x}), \\ \text{subject to} \quad & \Delta_{p_j}(\mathbf{x}) \leq \tilde{\Delta}_{p_j} \quad \text{for } j = 1, \dots, M, \end{aligned} \quad (2)$$

where $\mathbf{x} \in [0,1]^d$ is a vector including the d design parameters (min-max normalized) described later in Sec. 4.4. The lift coefficient $c_l(\mathbf{x})$ is the performance targeted by the optimizer, and it is computed using the baseline turbulence closure. The turbulence uncertainty estimate $\Delta_{p_j}(\mathbf{x})$ of a generic performance p_j is computed according to the EPM, described in Sec. 3. Note

that the designer may decide to apply as many credibility constraints as desired, depending on the specific application. For instance, one may chose to simultaneously apply credibility constraints on both the lift and the drag. The subscript j indicates indeed different performances. Hereinafter, we will only consider $M = 1$, therefore dropping the subscript and referring to $\Delta_p(\mathbf{x})$ directly. Briefly, the EPM ultimately results into the definition of several RANS models, each implementing a significantly different physics underlying turbulent phenomena. Predictions from the different RANS models can be exploited to define the trustworthiness indicator $\Delta_p(\mathbf{x})$, which basically measures their maximum discrepancy or, in other words, their dispersion. In the following, p will be either the lift (c_l) or the drag (c_d) coefficient, depending on the test case. We stress here again that the EPM only provides an estimation of the performance prediction uncertainty. As such, $\Delta_p(\mathbf{x})$ does not represent a bounded set of possible values but, rather than that, it can be interpreted as a trustworthiness indicator related to the turbulence closure. By constraining $\Delta_p(\mathbf{x})$ below a maximum threshold $\tilde{\Delta}_p$, we basically focus on designs for which the credibility of RANS prediction is endorsed, at least to an arbitrary extent.

At the time of the writing, three most notable contributions employing the EPM in the context of robust optimization are found [13, 14, 9]. In Ref. [13], the turbulence model error $\Delta_p(\mathbf{x})$ is treated as an arbitrarily weighted penalty term included in the objective function. Though effective for some problems, this approach does not allow for a stiff control over the magnitude of the performance uncertainty. Since no hard constraint is applied, the formulation does not prevent the optimizer from exploring and converging to designs associated to unreliable predictions. At the same time, the approach of maximizing (minimizing) a functional based on the worst (best) prediction among the diverse EPM RANS models, implemented in [14, 9], also leaves room to objections. Indeed, the EPM only provides reasonable and informed turbulence uncertainty estimates, and not rigorous and provable bounds [3]. In other words, there is absolutely no guarantee that the true performance will be higher (lower) than the worst (best) EPM prediction, making the proposed strategy misleading and, possibly, dangerous. Secondly, the approach does not prevent the optimizer from exploring regions of the design space associated to unreliable predictions, possibly wasting computational resources, and converging to not credible solutions, thus frustrating the optimization effort. Moreover, a theoretical analysis of the EPM [15] shows that there exist necessary and sufficient conditions that must be fulfilled to ensure the

physical plausibility (not realizability) of the CFD model. According to this analysis, relying on the worst/best EPM design may entail relying on a CFD solution associated to a non-plausible physics, with an not quantified risk of being either overly or poorly conservative.

In the strategy we propose, we explicitly limit the optimum search only to regions where we are confident we can obtain credible predictions, possibly improving the efficiency of the process. The optimum search is carried out using a state-of-the-art turbulence closure. At the same time, by enforcing a $\Delta_p(\mathbf{x}) \leq \tilde{\Delta}_p$ we require that the optimal design should perform consistently under RST model-form uncertainty. Practically, we aim at limiting the dispersion of prediction from the diverse EPM RANS model, which we recall implement a very different physics. In doing so, the relevance of the plausibility of the EPM realization diminishes, since the EPM predictions contribute only in establishing the credibility estimator. However, our strategy relies on the arbitrary definition of the maximum credibility estimator threshold $\tilde{\Delta}_p$, the choice of which is still left to the experienced designer, depending on the particular application and on the confidence required. In particular, a too restrictive $\tilde{\Delta}_p$ would lead to designs that are completely not sensitive to turbulent fluctuations. On the other hand, a loose constraint would allow selecting solutions associated to not credible RANS predictions. At the same time, we recall here that, within the realizability limits, the EPM methodology relies on arbitrary perturbation of the RST. Choices may be different and the designer may decide to explore extremal states or limit the magnitude of perturbation in order to explore the close proximity of the baseline solution. Naturally, the credibility threshold must be selected also in view of the applied perturbations. Moreover, it has to be stressed that the approach and the results presented in this paper are limited by the strong hypothesis of that the turbulence closure model-form error is the only source of uncertainty. Naturally, this assumption may be relaxed in future developments, extending the confidence-based strategy with the inclusion of other uncertainty sources.

3. Estimate of the turbulence model uncertainty

The Reynolds-Averaged Navier-Stokes model (RANS) is a popular approach for describing the motion of turbulent fluid flows. According to the Reynolds decomposition, the RANS model is obtained from the more general Navier-Stokes equations, by separating unsteady quantities into a time-

average term $\langle \phi \rangle$ plus a random fluctuation ϕ . The so-called Reynolds Stress Tensor (RST) is a byproduct of the Reynolds averaging operation, and it incorporates all the effects of turbulent motions. By definition, the RST is symmetric and turbulence closures are required for estimating the six independent components. A large body of literature deals with the problem of specifying the RST components. Nonetheless, it is widely acknowledged that the strong inherent model-form assumptions underlying turbulence closures limit the credibility of CFD predictions, at least for some particular flow configurations [1]. Unfortunately, the direct quantification of the errors introduced in turbulent modeling is practically intractable and the implementation of advanced UQ techniques is required.

In the following, we illustrate the EPM for estimating the turbulence model uncertainty $\Delta_p(\mathbf{x})$ and we apply it to an exemplary test case concerning a subsonic flow past a NACA0012 airfoil. This very same test case and its setting are used later in the optimization problems.

3.1. Eigenspace Perturbation Method

The Eigenspace Perturbation Method [2, 3] was devised to estimate the L2 [1] uncertainty arising from the process of relating the microscopic state of a flow to macroscopic quantities. By definition, the RST $\langle u_i u_j \rangle$, with $i, j = \{1, 2, 3\}$, must be positive semi-definite. Hence, it must fulfill a set of realizability conditions [16, 17]

$$\langle u_i u_i \rangle \geq 0, \quad \langle u_i u_i \rangle + \langle u_j u_j \rangle \geq |2 \langle u_i u_j \rangle|, \quad \det(\langle u_i u_j \rangle) \geq 0. \quad (3)$$

Notoriously, the RST is decomposed into an anisotropy and a deviatoric part

$$\langle u_i u_j \rangle = 2k \left(b_{ij} + \frac{\delta_{ij}}{3} \right). \quad (4)$$

being k the turbulent kinetic energy and δ_{ij} the Kronecker delta. Naturally, the realizability conditions (3) must apply also to the anisotropy tensor b_{ij} , which can be expressed in its spectral form $b_{ij} = v_{ik} \Lambda_{kl} v_{jl}$, being v_{ik} and v_{jl} the left and the right eigenvectors, whereas Λ_{kl} is a diagonal matrix containing the eigenvalues λ_i in a decreasing order.

The EPM consists in applying perturbations of finite amplitude to the RST eigenspace during the CFD solution iterations, yet fulfilling the realizability conditions (3). Hereinafter, we will take advantage of superscript * to point out a perturbed entity. In the most general approach, perturbations

consist in increasing/decreasing the amount of turbulent kinetic energy (k^*), in varying the anisotropy tensor spectral distribution (Λ_{kl}^*), or in changing the orientation of the anisotropy tensor basis (v_{ik}^*), obtaining

$$\langle u_i u_j \rangle^* = 2k^* \left(v_{ik}^* \Lambda_{kl}^* v_{jl}^* + \frac{\delta_{ij}}{3} \right). \quad (5)$$

By means of perturbations, the whole RST realizability space is explored, thus allowing for the estimation of turbulence closure uncertainty. The interested reader is referred to [2, 3] for a thorough overview of the EPM. In principle, perturbations are subject to a spatial variation, leading to a very high dimensionality hindering practical applications. To overcome the challenge, researchers follow diverse approaches e.g., see [3, 7, 18].

In this paper we adopt the approach described in [4]. Namely, we assume uniform spatial perturbations Λ_{kl}^* and v_{ik}^* , while neglecting the perturbation of the k value. According to the realizability conditions (3), there exist three limiting states of turbulence componentiality that can be attained by perturbing the spectral energy distribution i.e., the eigenvalues in Λ_{kl}^* . These extremal states, labeled $1C$, $2C$ and $3C$, correspond to the 1-, 2-, and 3-component (isotropic) turbulence. In our setting, the componentiality of turbulence is first perturbed towards one of these limiting states. After, we apply a perturbation to the orientation of the eigenvectors which, according to [4], corresponds to modulating the production of turbulent kinetic energy \mathcal{P} . Naturally, it is possible to perturb the eigenvectors in order to maximize or to minimize the production of turbulent kinetic energy.

Summarizing, a total of six combinations would arise considering the three extremal states of turbulent componentiality ($1C$, $2C$ and $3C$) and two extremal states associated to the production of turbulent kinetic energy (\mathcal{P}^{max} and \mathcal{P}^{min}). In practice, the extremal state \mathbf{x}_{3C} is associated with an isotropic turbulence componentiality. Therefore, \mathcal{P}_{3C} is invariant to the orientation eigenvector basis, thus reducing the number of possible extremal states to five. Hereinafter, we will adopt the following labels \mathcal{P}_{1C}^{max} (PA), \mathcal{P}_{2C}^{max} (PB), $3C$ (PC), \mathcal{P}_{1C}^{min} (PD), \mathcal{P}_{2C}^{min} (PE), and (BS) to indicate the baseline model. To estimate the turbulence closure uncertainty, the five EPM cases need to be simulated using a modified CFD model. The turbulent uncertainty estimates of a selected performance p are obtained by considering the max/min values resulting from the five EPM solutions. Therefore, the indicator

$$\Delta_p(\mathbf{x}) = \max(p^{PA}, p^{PB}, p^{PC}, p^{PD}, p^{PE}) - \min(p^{PA}, p^{PB}, p^{PC}, p^{PD}, p^{PE}), \quad (6)$$




Figure 1: NACA0012 baseline geometry. (a) Profile sketch. (b) Enlargement of the numerical grid in the close proximity of the airfoil.

represents the difference among the upper and lower turbulence uncertainty estimates and provides an estimation of the RANS model trustworthiness.

3.2. Application to the simulation of a NACA0012 airfoil

We describe the reference test case namely, the baseline geometry and the boundary conditions, specifying the set up of numerical CFD simulations. The very same numerical setting will be later used for solving the optimization problems. The EPM analysis of the reference test case is also reported, to highlight the uncertainty estimates related to the credibility of RANS predictions.

The NASA TMR (Turbulence Modeling Resource) [19, 20] provides validation test cases for turbulence models, including a database gathering experimental data from different sources and a large sequence of numerical grids. The test case investigated here is referred to as the 2DN00 in the NASA TMR database. Summarizing, the geometry consists in the classical NACA 0012 airfoil with a chord of 1 m, see Fig. 1(a). The airfoil is plunged into a subsonic free flow at Mach 0.15, with a static temperature of 300 K degrees. The Reynolds number (Re), computed considering the unit profile chord reference length and assuming air viscosity to correspond to that at 300 K, is equal to $6 \cdot 10^6$. In our computations, viscosity is computed locally according to the Sutherland law. Moreover, the thermal conductivity is also evaluated locally, ensuring a constant Prandtl number. At the far-field inlet,

a turbulent to laminar viscosity ratio of 10.0 and a 5% intensity of turbulent fluctuation are applied. Over the airfoil, an adiabatic no-slip boundary condition is imposed. The turbulence baseline closure (BS) is the Menter’s Shear Stress Transport (SST) closure [21].

CFD simulations are carried out using the open-source SU2 suite, see [22, 23]. SU2 was recently equipped with EPM capabilities [5] making SU2, at the time of the writing, the only viable CFD solver option for achieving the purposes pursued in this work. The targeted QoI namely, the airfoil *lift coefficient* (c_l), the *drag coefficient* (c_d), the *efficiency* ($EF = c_l/c_d$) and the *pitching coefficient* (c_m) are directly returned by the solver. Field solutions are obtained using a standard implicit time-marching approach taking advantage of a generalized Approximate Riemann solver of Roe type with a Monotone Upstream-centered Scheme for Conservation Laws (MUSCL) [24]. The Venkatakrishnan flux limiter is employed to modulate the second-order spatial accuracy. A two-level multi-grid approach is adopted to possibly hasten the convergence of the numerical solution.

According to the coding reported by [19, 20], we selected the 449x129 (257 points on the airfoil surface) grid from the FAMILY I. A mesh sensitivity analysis concerning the very same test case is reported in [25], showing that the selected mesh spacing is a reasonable trade-off for ensuring the capturing of the baseline flow gradients within the wall boundary layer, yet maintaining an affordable computational burden. Rigorously, mesh sensitivity studies should be carried out independently for the five EPM perturbations, to ensure that the resolution is adequate to accurately capture the turbulence uncertainty estimates. Nevertheless, since we are aiming at controlling a dispersion Δ_p rather than maximizing/minimizing the lower/upper bound, we accept a small accuracy loss in favor of a significant reduction of the computational burden. Anyways, enforcing a credibility constraint also promotes a similar sensitivity of the 5 EPM and the BS models to the grid resolution.

Figure 1(b) reports an enlargement of the numerical grid around the airfoil. The convergence of each simulation is monitored by computing integrated quantities namely, the c_d coefficient, ensuring that its value remains within a 10^{-4} tolerance range for a significant number (100) of consecutive iterations. If convergence could not be achieved within the 8000 iterations, possibly because mean-flow instabilities are likely to be developed within the domain, simulations are stopped. In such cases, the amount of lift generated by the airfoil drops significantly, making the corresponding design a poor candidate for the optimizer. Moreover, the RANS model is known to be in-

accurate in predicting separated flows, therefore the five simulations from the EPM are expected to return significantly different predictions. Constraining the optimization procedure to a limited difference among EPM predictions helps ensuring that the not converged solution is discarded. Naturally, more refined criteria may be adopted to improve the optimization process.

Having clear the numerical setting, we present the EPM analysis of the reference test case. Figure 2(a) and Fig. 2(b) report, respectively, the lift and drag polar for the NACA0012 profile. The black continuous curves are obtained using the baseline closure [21] whereas the dashed curves correspond to the five EPM predictions. The gray shaded area indicates the max-min envelope returned by the EPM i.e., what we consider to be the turbulence uncertainty estimate of the aerodynamic performance.




Figure 2: baseline (BS) and EPM (PA-PE) polar curves. The gray envelope indicates the performance credibility $\Delta_p(\mathbf{x})$. (a) Lift coefficient c_l . (b) Drag Coefficient C_D .

Clearly, both Fig. 2(a) and Fig. 2(b) show that the turbulence closure uncertainty has negligible relevance for small-to-medium Angle of Attack (AoA, α), below approximately 12 degrees. This is coherent with the physics of a high Reynolds flow around a slender body, which typically presents a weak adverse pressure gradient and a limited streamlines' curvature. For large AoAs, the above conditions are no longer met and the RANS model loses credibility. This fact is highlighted by a rapid enlargement of the turbulence uncertainty estimates. In particular, the stall angle prediction varies significantly depending on the injected perturbation. The PC and PD model predict stall at around 14 degrees, whereas a significantly larger values results from PA and PB. Namely, the PA/PB perturbations are associated to a




Figure 3: baseline (BS) and EPM (PA-PE) polar curves. The gray envelope indicates the performance credibility $\Delta_p(\mathbf{x})$. (a) Efficiency c_l/C_D . (b) Pitching Coefficient C_M .

maximization of the turbulent kinetic energy production within the boundary layer. Coherently with experience, a highly energized boundary layer helps delaying stall. On the other hand, perturbations that suppress kinetic energy production (PC/PD) and the one that spreads k isotropically (PC) predict an earlier flow separation. Note that experimental data reveal that the NACA0012 undergoes stall at around 18 [deg], see Ref. [26]. Figure 3(a) and (b) report the efficiency and the pitching moment envelopes w.r.t. the AoA, as computed by SU2. In particular, it is interesting to point out that the baseline pitching moment polar, Fig. 3(b), exceeds the turbulence bound estimates at an AoA of about 16 [deg], thus confirming that the EPM does not return rigorous and provable bounds.

4. Optimization method

We take advantage of the Efficient Global Optimization (EGO) technique from the class of Bayesian methods, see [27]. Our implementation closely follows the algorithm described in [27], therefore in the following we just recall the fundamental aspects only. In Bayesian optimization, the design space is efficiently searched by taking advantage of the so-called acquisition function, an auxiliary mathematical object that aims at exploiting prior belief and available data. Different acquisition functions may be exploited, singularly or combined, to further improve the efficiency of the optimization, see [28]. In view of investigating possible solutions to the issue of integrating

the measure of turbulence model-form uncertainty into traditional design optimization frameworks, we explored several strategies for accounting for the credibility constraint when acquiring a new design point. Here, we rely on three alternatives namely, the *Expected Improvement* (EI) [27], the *Expected Constrained Improvement* (EIC) [29, 30] and the *Augmented Expected Improvement* (AEI) [31]. In all cases, the acquisition function is combined with surrogate modeling techniques i.e, a Gaussian Processes (GP), for the learning of the objective function.

4.1. Acquisition Functions

We detail here the three Acquisition Functions (\mathcal{AF}) exploited in our optimization process. The Expected Improvement (EI) infill criteria aims at minimizing the expected deviation from the true maximum when acquiring a new design point. The Improvement estimator is defined as $I(\mathbf{x}) = \max(0, f_{\max} - f(\mathbf{x}))$, being f_{\max} the maximum value of the cost function as observed up to the current step. Note that, in the computation of the Improvement, the full computational model $f(\mathbf{x})$ is substituted by a gaussian predictor $\hat{f}(\mathbf{x})$. Hence, $I(\mathbf{x})$ is a positive random scalar value which is larger than zero if, and only if, the prediction of \hat{f} at a selected design point is better than the best recorded value. According to the EI acquisition criteria, the new design point \mathbf{x}_n is found by solving the following constrained maximization problem $\mathbf{x}_n = \arg \max_{\mathbf{x} \in \Omega} EI(\mathbf{x})$, subject to $h_j(\mathbf{x}_n) \leq \tilde{h}_j$, being $EI(\mathbf{x}) = \mathbb{E}_{\hat{f}(\mathbf{x})}[I(\mathbf{x})]$ the so-called Expected Improvement and $\mathbb{E}_{\hat{f}(\mathbf{x})}$ the expectation operator related to the randomness of the gaussian predictor $\hat{f}(\mathbf{x})$. The $h_j(\mathbf{x})$ are the M inequality constraints to which the design process may be subject to. Note that in our problem we consider a single inequality constraint corresponding to $\Delta_p(\mathbf{x}) \leq \tilde{\Delta}_p$, but multiple constraints may be enforced at once. Similarly to the target function, the M constraints are modeled using M independent gaussian predictors $\hat{h}_j(\mathbf{x})$.

An analytic expression of the expected improvement is available, see [32, 33], reading

$$EI(\mathbf{x}) = \begin{cases} (\mu_{\hat{f}(\mathbf{x})} - f_{\max} - \epsilon)\Phi(Z) + \sigma_{\hat{f}(\mathbf{x})}\phi(Z), & \text{if } \sigma_{\hat{f}(\mathbf{x})} > 0, \\ 0, & \text{if } \sigma_{\hat{f}(\mathbf{x})} = 0, \end{cases} \quad (7)$$

with $Z = (\mu_{\hat{f}(\mathbf{x})} - f_{\max} - \epsilon)/\sigma_{\hat{f}(\mathbf{x})}$. The terms Φ and ϕ are, respectively, the cumulative and probability density functions of the standard normal distribution. The quantity $\sigma_{\hat{f}(\mathbf{x})}$ is the standard deviation associated to the

gaussian posterior prediction of the target function at \mathbf{x} . The ϵ parameter enables the control on the trade-off between exploration and exploitation. Following [34], ϵ is set to the value 0.01 scaled by the variance of the gaussian signal.

The second acquisition function we consider is the so-called Expected Constrained Improvement (hereinafter referred to as EIC), see [29] and [30]. Inequality constraints are no longer accounted in the EI maximization process, but they are directly included in a modified formulation of the improvement indicator reading $I_C(\mathbf{x}) = Q(\mathbf{x}) \max(0, f_{\max} - f(\mathbf{x})) = Q(\mathbf{x})I(\mathbf{x})$, which is known as the constrained improvement. The term $Q(\mathbf{x})$ is a feasibility indicator function, reading 1 if the constraint is satisfied, 0 if not. Each constraint $h_j(\mathbf{x})$ is again modeled using secondary conditionally independent gaussian predictor $\hat{h}_j(\mathbf{x})$. Due to the randomness of the gaussian predictors, $Q(\mathbf{x})$ becomes a Bernoulli random variable and the probability that M inequality constraints are satisfied can possibly be expressed according to

$$\mathcal{P}_C(\mathbf{x}) = \prod_{j=1}^M \mathcal{P}_j(h_j(\mathbf{x}) \leq \tilde{h}_j), \quad (8)$$

with

$$\mathcal{P}_j(h_j(\mathbf{x}) \leq \tilde{h}_j) = \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{\tilde{h}_j - \hat{h}_j(\mathbf{x})}{\hat{\sigma}_j^h \sqrt{2}} \right) \right], \quad (9)$$

being $\hat{\sigma}_j^h$ the standard deviation associated to the gaussian posterior prediction of the constraint value at \mathbf{x} . Hence, the expected constrained improvement can be expressed as $EIC(\mathbf{x}) = \mathbb{E}_{\hat{f}(\mathbf{x})}[I_C(\mathbf{x})] = \mathcal{P}_C(\mathbf{x})EI(\mathbf{x})$, being the new promising design point \mathbf{x}_n the argument maximizing $EIC(\mathbf{x})$. The EIC acquisition function has the advantage of being auto-adaptive since the variability of the surrogates modeling constraints impacts directly on the choice of \mathbf{x}_n .

The third acquisition function considered in this work is the so-called Augmented Expected Improvement (AEI)[31]. It consists in a modification of the classical EI in-fill criteria by means of a scaling factor $B(\mathbf{x})$:

$$AEI(\mathbf{x}) = B(\mathbf{x})EI(\mathbf{x}) = \left(1 - \frac{\sigma_\epsilon}{\sqrt{\sigma^2(\mathbf{x}) + \sigma_\epsilon^2}} \right) EI(\mathbf{x}). \quad (10)$$

where σ_ϵ^2 is the variance of the random noise added to the gaussian predictor. Note that, in the AEI in-fill criteria, the improvement I is computed w.r.t.

the effective best solution $f^* = f(\mathbf{x}^*)$ (not to be confused with the current best one f_{\max}) which is determined by means of an utility function $u(\mathbf{x})$, by solving the following maximization problem

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \Omega} u(\mathbf{x}). \quad (11)$$

In this work, $u(\mathbf{x}) = -\mu_{\hat{f}}(\mathbf{x}) - \sigma_{\hat{f}}(\mathbf{x})$

The scaling term B in Eq. 10 is designed to account the diminishing return of additional replicates as predictions become more accurate. In other words, it introduces a measure of the usefulness of an additional sample at a particular location, improving the balance between exploration and exploitation in the presence of a noisy response surface.

4.2. Surrogate Models

We rely on Gaussian Processes (GP) to mimic the behavior of the objective and the constraint functions w.r.t. the design parameters. GPs allow to retrieve the analytic expression of the EI and the EIC, significantly simplifying the implementation [35, 33]. We underline that the selection of a specific GP kernel ($k(\mathbf{x}, \mathbf{x}')$) is a very delicate task. Selecting a particular family of covariance functions introduces some strong assumptions inherent the form of the response surface e.g., smoothness properties. Naturally, these choices reflect on the acquisition function and impact the convergence of the optimizer. In this work, all the GP surrogates rely on a Stationary Squared Exponential Covariance kernel with the addition of a White Gaussian Noise term of uniform amplitude $k(\mathbf{x}, \mathbf{x}') = \sigma_n^2 \exp(-|\mathbf{x} - \mathbf{x}'|^2 / 2l^2) + \sigma_\epsilon^2 \delta(\mathbf{x} - \mathbf{x}')$.

In the kernel expression, $\sigma_n^2 \in [0.3, 10]$ (signal variance), $l \in [0.01, 10]$ (correlation length) and $\sigma_\epsilon \in [10^{-6}, 1]$ (gaussian noise level) are the hyperparameters. These are fit to the available data at each iteration of the optimization process using routines from the Scikit-learn Python library [36].

4.3. The optimization algorithm

The work flow of the automated optimization framework here presented is summarized by Algo. 1. Namely, the optimizer is initialized by sampling the \mathbf{x}_0 point at random. The BS CFD model and the five EPM simulations are carried out at $\mathbf{x} = \mathbf{x}_0$, to evaluate the objective function c_l and the model trustworthiness indicator Δ_p . Note that these data, associated to a single point, are employed to initialize the GP_{c_l} and GP_{Δ_p} surrogates. Starting from this initialization, the real optimization process is unrolled. The

selected acquisition function (the EI is considered in Algo. 1) is maximized to select the most promising design point \mathbf{x}_n under the constraint of that $GP_{\Delta_p}(\mathbf{x}_n) \leq \tilde{\Delta}_p$. In this work, we rely on the Sequential Least Squares Programming (SLSQP) method implemented in the Python SciPy library [37], for solving the constrained maximization problem, or on the Limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS-B) method [37], for the unconstrained one. Once a new promising design point is identified, the full CFD baseline model simulation and the five EPM simulations are automatically performed. The produced data are then used to update the surrogates of the objective and the constrain surfaces. The optimizer is left searching the design space until the maximum budget of CFD simulations (100) is exhausted.

Algorithm 1: Pseudo-code describing the confidence-based Bayesian Optimization framework using the EI criteria.

Input: Random \mathbf{x}_0

- 1 Evaluate $c_l(\mathbf{x}_0)$ using the baseline CFD model;
- 2 Evaluate $\Delta_p(\mathbf{x}_0)$ using the EPM;
- 3 Build the surrogate models GP_{c_l} and GP_{Δ_p} ;
- 4 **do**
- 5 Find $\mathbf{x}_n = \operatorname{argmax}_{\mathbf{x}} EI^{GP_{c_l}}(\mathbf{x})$ s.t. $GP_{\Delta_p}(\mathbf{x}) \leq \tilde{\Delta}_p$;
- 6 Evaluate $c_l(\mathbf{x}_n)$ using the baseline CFD model;
- 7 Evaluate $\Delta_p(\mathbf{x}_n)$ using the EPM;
- 8 Update the surrogate models GP_{c_l} and GP_{Δ_p} ;
- 9 **if** $c_l(\mathbf{x}_n) > c_l^{BEST}$ **then**
- 10 $c_l^{BEST} = c_l(\mathbf{x}_n)$;
- 11 $it++$;
- 12 **while** $it < N_{it}^{MAX}$;
- 13 **return** Best design \mathbf{x} ;

4.4. Design Parametrization

We rely on the classical NACA 4 digit series codification to parametrize the shape of the airfoil. Though we acknowledge the existence of more flexible parametrization strategies, such as Class Shape Transformation [38] or B-splines curves/surfaces [39, 40], we here deem the NACA 4 digit series codification to be flexible enough to illustrate the relevant aspects of the

proposed optimization strategy. Besides, we also consider the AoA, defined as the angle between the airfoil chord and the direction of the freestream flow, to be a design parameter. Namely, a total of four design parameters, $\mathbf{x} \in \Omega \subset \mathbb{R}^4$, is considered: the AoA and the three entries defining the NACA x-x-xx profile, where each "x" represents one of the four digits. In this work, we refer to the first digit as n_1 , to the second as n_2 and to the third and fourth digits (considered as a single parameter) as n_3 . In the NACA codification, the first digit n_1 indicates the maximum camber as percentage of the chord. The second entry n_2 represents the distance, expressed as tenths of the chord, of the maximum camber from the leading edge. Eventually, n_3 expresses the maximum airfoil thickness as percent of the chord. Note that we will consider each digit as a floating point number, and not as an integer, thus extending the range of possible profiles which is possible to attain. Moreover, the four design parameters are bounded in order to avoid the selection of degenerate profiles or not admissible configurations $n_1 \in [0.0, 8.0]$, $n_2 \in [4.5, 8.0]$, $n_3 \in [1.0, 2.5]$ and $\alpha \in [0.0^\circ, 30.0^\circ]$. These bounds are selected through a dedicated heuristic a priori analysis not reported here for brevity reasons. Note that, in the actual optimization process, all the parameters are min-max normalized, so that $\mathbf{x} \in [0, 1]^4$. For each design, a deformation algorithm based on Radial Basis Functions (RBF) is used to adapt the baseline mesh to the different geometries. This procedure helps conserving a nice grid quality level while sparing the re-computing of the mesh.

5. Results

We consider three different optimization problems (A, B, C) targeting the maximization of the baseline lift coefficient c_l^{BS} . In test case A (Section 5.1), no shape deformation is considered and the AoA serves as the only design parameter. The optimum is searched both without and with enforcing the constraint $\Delta_{c_l} \leq \tilde{\Delta}_{c_l}$ limiting the optimizer to Confidence-Based (CB) predictions only. Section 5.2 extends the NACA airfoil optimization including the shape parametrization described in Sec. 4.4. The lift coefficient c_l is first targeted by an unconstrained optimization setting. Secondly, the credibility constraint $\tilde{\Delta}_{c_l}$ is restored. Eventually, Sec. 5.3 concerns the optimization of the airfoil geometry in order to achieve a prescribed value of the lift coefficient c_l with minimum drag, subject to a credibility constraint $\tilde{\Delta}_{c_d}$ acting on the drag prediction. For all tests, the set up of numerical simulation matches the one described in Sec. 3.2.

Case	CB	\mathcal{AF}	n_1	n_2	n_3	AoA	c_l^{BS}	c_d^{BS}	c_l^{BS}/c_d^{BS}	c_m^{BS}
A	N	EI	-	-	-	16.119	1.514	0.038	39.64	-0.0170
	A	EI	-	-	-	13.496	1.418	0.023	62.42	-0.0087
	A	EIC	-	-	-	14.286	1.465	0.026	57.19	-0.0121
	A	AEI	-	-	-	13.522	1.420	0.023	62.20	-0.0088
B	N	EI	8.00	8.00	1.66	18.014	2.341	0.074	31.73	0.2155
	A	EI	8.00	8.00	1.15	14.419	2.213	0.053	41.66	0.2463
	A	EIC	8.00	8.00	1.36	13.798	2.199	0.050	43.87	0.2519
	A	AEI	8.00	8.00	1.23	14.362	2.215	0.053	41.86	0.2476
C	A	EIC	5.01	4.92	1.33	3.13	0.994	0.009	107.20	0.148

Table 1: assessment of design performances. The CB column indicate the credibility constraint state (N, not active and A, Active) whereas \mathcal{AF} the acquisition function employed.

Case	CB	Δ_{c_l}	Δ_{c_d}	Δ_{c_l/c_d}	Δ_{c_m}	Δ_{c_l} %	Δ_{c_d} %	Δ_{c_l/c_d} %	Δ_{c_m} %
A	N	0.990	0.150	43.20	0.070	158 %	425 %	1236 %	-481 %
	A	0.199	0.007	19.38	0.014	16.2 %	33.0 %	46.1 %	-76.3 %
B	N	0.520	0.006	19.60	0.040	27.3 %	79.7 %	128.3 %	23.4 %
	A	0.196	0.005	6.57	0.032	9.2 %	9.5 %	17.9 %	13.9 %
C	A	0.008	0.003	24.58	0.002	0.9 %	33.9 %	32.8 %	1.5 %

Table 2: EPM assessment of design performances. The CB column indicate the credibility constraint state (N, not active and A, Active).

Table 1 summarizes the BS performance delivered by the optimal designs, as resulting from the different optimization strategies. CB indicate the (Active or Not active) state of the credibility constrain, whereas column \mathcal{AF} states the acquisition function employed. Table 2 reports instead the EPM analysis of the most representative optimal design as found considering the active/non active constraint and the different \mathcal{AF} , for the three test cases. The table reports the uncertainty turbulence estimates both in magnitude and in percentage, w.r.t. the EPM lowest prediction.

5.1. Test Case A

This test case aims at maximizing the lift generated by the NACA 0012 airfoil based on the sole variation of the AoA. Since stall results in an abrupt degradation of the targeted performance, the problem is formally equivalent

to maximizing the c_l w.r.t. α and subject to stall avoidance. Namely,

$$\arg \max_{\alpha} c_l(\alpha), \quad \text{s.t.} \quad \Delta_{c_l}(\alpha) \leq \tilde{\Delta}_{c_l} = 0.2. \quad (12)$$

where $\tilde{\Delta}_{c_l} = 0.2$ is the threshold selected for optimization problems enforcing a prediction credibility constraint.

5.1.1. Unconstrained Optimization

Figure 4(a) and (b) report, respectively, the surrogate model of the objective function c_l (black continuous curve) as learned from the evaluation of the first two design points, and as learned after the maximum computational budget is exhausted. The objective function is complemented by the one-standard deviation interval (σ , gray shaded area). The blue continuous curve represents the EI acquisition function evaluated on the whole parameter space. The next proposed design point is marked with a triangle and it corresponds the maximum of the EI function. At the beginning of the optimization process, the gaussian approximation of the c_l polar curve is poor, see Fig. 4(a). As the design space is explored/exploited, during the optimum search, the amount of data increases so to reasonably approximate the objective function. Indeed, the c_l polar in Fig. 4(b) closely resembles the baseline one reported in Fig. 2(a).




Figure 4: test case A, unconstrained optimization. Learning curve for the target function c_l , black line, and expected improvement function, blue line. Black points indicate sampled designs whereas the black triangle marks the proposed design. Shaded areas indicate the one standard deviation interval of gaussian predictors. (a) Initial steps. (b) final step.

Figure 5(a) reports the history of the targeted performance, predicted using the BS model, recorded as the optimizer explores the design space. The algorithm quickly converges to the optimum, located at $\text{AoA} \approx 16.12$ degrees, with an estimated c_l^{BS} of 1.51 [-], $c_d^{BS} \approx 0.038$, $c_l^{BS}/c_d^{BS} \approx 40$ and pitching moment c_m^{BS} coefficient of about -0.017, see Tab. 1.

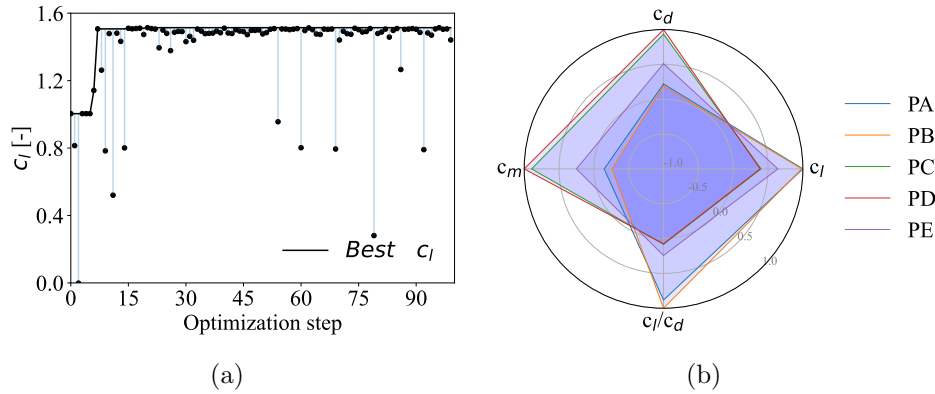


Figure 5: test case A, unconstrained optimization. (a) Optimization history. (b) EPM analysis of the optimal configuration.

To provide an insight of the performance variability w.r.t. turbulence closure uncertainty, we plot a polar graph comparing the airfoil aerodynamic coefficients (c_l , c_d , c_m and efficiency EF) from the EPM analysis of the optimal design, see Fig. 5(b). Note that the values reported on the polar graph are normalized w.r.t. the EPM estimate of highest magnitude. Clearly, the EPM simulations return significantly different performances. Quantitatively, the EPM returns a $\Delta_{c_l} \approx 0.99$, $\Delta_{c_d} \approx 0.15$, $\Delta_{c_l/c_d} \approx 43.2$ and $\Delta_{c_m} \approx 0.07$. In percentage w.r.t. EPM minimum value, that makes a difference of, respectively, 158%, 425%, 1236% and -481%. Therefore, the EPM analysis reveals that RANS predictions of the NACA0012 airfoil at the optimum AoA of 16.12 degrees are not credible.

A more detailed analysis of flow predictions from the five EPM simulations helps providing some deeper physical insights. In the following, we show the turbulent kinetic energy (TKE) field as predicted by the PA, PC and PE EPM models, normalized w.r.t. the largest local value among the three solutions. Note that perturbations PA and PB give birth to very similar flows, as well as PC and PD. Figure 6(a), reporting the TKE field for PA, reveals an attached flow. In particular, the TKE field is homogeneous

and most of the energy remains confined within the wake and the boundary layer. Simulations PC and PD also predict similar performances, and the resulting TKE field for PC reveals a fully separated flow, see Fig. 6(b), with a recirculation bubble developing on the whole airfoil suction side. Eventually, Fig. 6(c) presents the PE flowfield, characterized by a semi-stalled condition. The TKE field shows that the boundary layer remains attached in the forward portion of the airfoil, while separation is triggered in the aft part of the profile. Naturally, these very diverse behaviors described by the five EPM simulations lead to an increase of the Δ_p indicators.




Figure 6: test case A, unconstrained optimization. EPM predictions of the TKE field developing around the airfoil in the optimal configuration. (a) PA. (b) PC. (c) PE.

5.1.2. Confidence-based optimization

The optimization is now constrained to credible CFD predictions only by requiring $\Delta_{c_l} < \tilde{\Delta}_{c_l} = 0.2$. Figure 7 reports, for the initial (left column) and for the final (right column) optimization steps, the surrogates of the objective function c_l (black curve), the acquisition function evaluated on the whole design space (blue curve) and the surrogate of the trustworthiness indicator Δ_{c_l} (red curve). Figures on different rows correspond to different \mathcal{AF} , respectively, the EI, the EIC and the AEI. Surrogates are complemented by a shaded area indicating the one-standard deviation interval. The horizontal red-shaded straight line indicates the $\tilde{\Delta}_{c_l} = 0.2$ threshold, whereas a vertical red-shaded line intersects the proposed design point. Note that, for the EIC, we also report the contributions of different factors composing the acquisition function (namely, the EI and the probability of respecting the credibility constraint \mathcal{P}_C , see Sec. 4.1).

During the optimization process, the algorithm explores designs, sometimes rejecting high performance solutions which are associated to not credible predictions. Depending on the \mathcal{AF} , a different balance among exploration

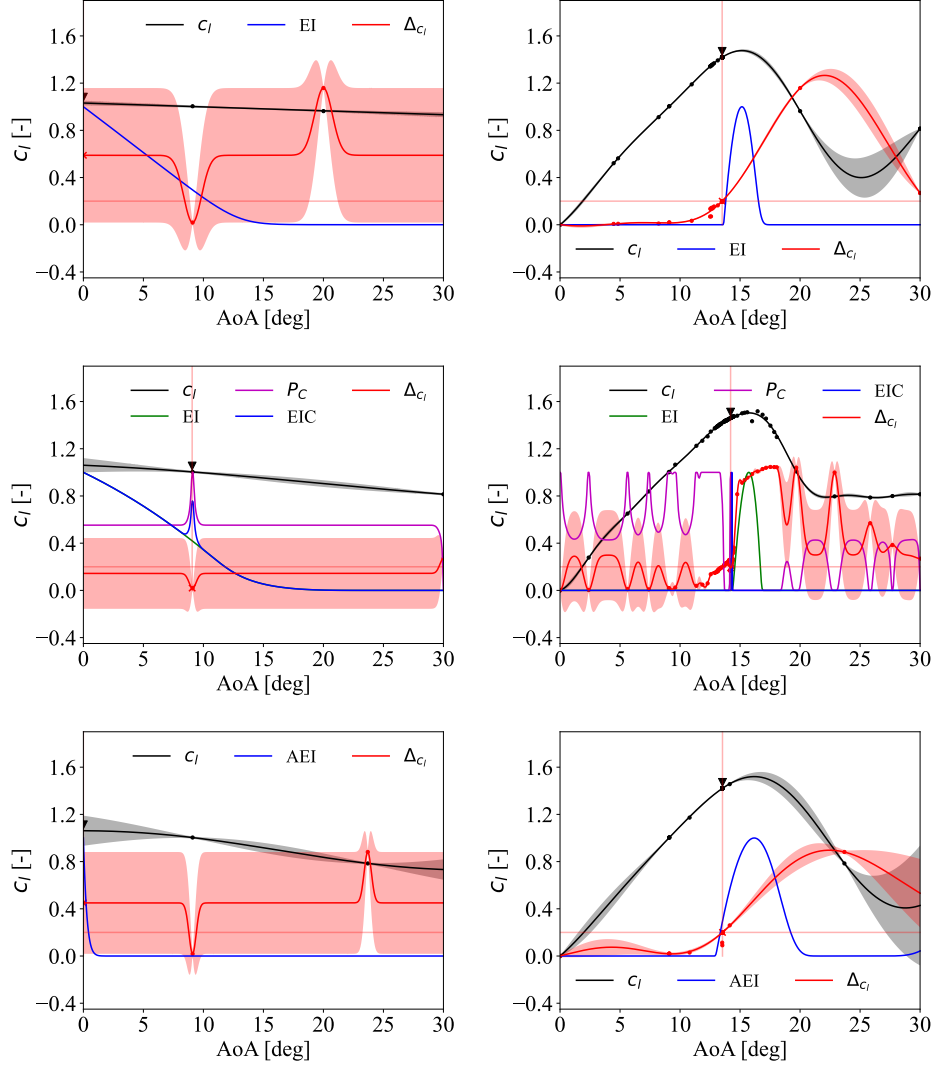


Figure 7: test case A, confidence-based optimization. Rows correspond to different acquisition functions, respectively EI, EIC and AEI. Black and red points indicate sampled designs whereas the black triangle marks the proposed design. Shaded areas indicate the one standard deviation interval of gaussian predictors. Left column: initial step. Right column: final step.

and exploitation is promoted, leading to a different path for reaching the optimum. Notably, the EIC allows for a more thorough exploration of the region of the design space associated with a large Δ_{c_l} . This is due to the fact that

the constraint is accounted in a probabilistic setting. That is, the algorithm is still allowed to explore regions of the domain likely subject to unreliable RANS predictions, if the credibility indicator function in such region is still not properly learned. The large amount of data available for the stalled region causes the gaussian approximator to exhibit an oscillatory behavior. As expected for low angles of attack, the EPM predicts a very limited Δ_{c_l} . Indeed, red dots are close to the zero line axis for an AoA smaller than about 12 degrees. While approaching the stall angle, the Δ_{c_l} indicator undergoes a quite rapid increase, with a significant steep gradient, almost a step, at about 15 degrees. Because of this rapid variation, which is not captured in the EI and the AEI settings, the resulting gaussian approximator must be endowed with a very small correlation length. Since the surrogate is built with respect to a variation about the mean of the available data, predictions tends to the mean as the distance of the query point from the available data is larger than the correlation length, giving birth to an oscillatory behavior and to a large prediction uncertainty. For the EI and AEI strategies, the high AoA region is poorly explored, the available data being concentrated within the pre-stall region where the target function has a smooth behavior. Therefore, surrogates are characterized by a larger correlation length. In particular, the same tendency to the mean is observed for AEI but, in this latter case, the larger correlation length limits makes it less noticeable. If the optimizer were to explore regions of large Δ_{c_l} prediction uncertainty, a new data point would inform the gaussian approximator, thus updating the belief (according to the Bayesian interpretation) about the unknown functional dependency of Δ_{c_l} w.r.t. the design parameter. Nevertheless, in all cases most points are found on the left hand side of the graph, indicating that eventually the applied constraint is actively forcing the optimizer to remain in the region of the design space where RANS predictions are credible i.e., low-to-medium AoA ahead of stall.

The role of the constraint is more evident by observing the optimizer target track records, for the three \mathcal{AF} , reported in the left column of Figure 8. The designs indicated by red crosses are rejected because the credibility constraint is violated. The rejection rates are 11/100 for the EI, 50/100 for EIC and 81/100 for AEI. The AEI has the highest rejection rate due to the fact that the optimizer straddles the maximum credibility limit while exploiting the region of AoA ≈ 13.5 degrees. In particular, the algorithm tends to evaluate designs which are slightly above the threshold. Similarly, the EI also promote exploitation over exploration, but with the largest number of




Figure 8: test case A, confidence-based optimization. Figure rows correspond to different acquisition functions, respectively EI, EIC and AEI. Left column: optimization history (red crosses are rejected because the credibility constraint is not fulfilled). Right column: EPM analysis of the optimal configuration.

sampling points located within the region of feasible designs. On the contrary, the EIC evaluates designs associated to a more variable performance, indicating a more exploratory behavior. Not surprisingly, the performance delivered by the constrained optimum design is lower than what achieved

in the unconstrained setting, see Tab. 1. Notably, the EIC \mathcal{AF} leads to a slightly different design (AoA is about 0.7 degrees larger than what obtained with the EI/AEI). This is related to the probabilistic approach employed in handling the constraint, which allows for a more thorough exploration of the design space. Figure 9 reports an enlargement of the Δ_{c_l} surrogate centered on the region in the close proximity of the optimum. Notably, the EI and the AEI functions (first and third rows) present a large number of sampling points accumulated at the intersection of the red shaded orthogonal lines indicating optimal angle and of the constraint threshold. The plot in the middle row (EIC) presents instead a more spread data set: the sampling spans a wider area and, because of this, a particular, non-monotone, trend of the Δ_{c_l} curve is consistently predicted in between 13.9 and 14.2 degrees. The reason why the curve is characterized by this non-monotone region is unclear. Possibly, the non-monotone behavior finds its roots in numerical aspect related to the indulgent convergence criteria employed for CFD simulations, but a thorough investigation is left for future works. However, all the considered simulations rely on the very same numerical configuration and grid, being the AoA the only free parameter. In the EI and in the AEI, the rigid enforcement of the credibility constraint quickly limits the optimum search to the portion of the design space located on the left hand side of the 13.53 degree threshold, preventing the learning of the Δ_{c_l} in the non-monotone region. In the EIC setting, the probabilistic treatment of the credibility constraint allows for a more thorough exploration of the design space, thus leading to the discovery of an optimal credible design characterized by slightly better performance. The capturing of the abrupt variation in the Δ_{c_l} curve, past 14 degrees, also explains why the surrogate is characterized by a correlation length much shorter than the EI and the AEI counterparts.

The EPM analysis of the similar optimum designs resulting from the three \mathcal{AF} is reported on the right column of Fig. 8. In all cases, the five EPM simulations return similar polar graph (we recall the values reported on the polar graph are normalized w.r.t. the EPM upper estimate) confirming a quite high confidence in RANS predictions. Quantitatively, a significantly reduced performance variability is found namely, Δ_{c_l} is approximately 0.1996. The other performance trustworthiness estimators read $\Delta_{c_d} \approx 0.007$, $\Delta_{c_l/c_d} \approx 19.38$ and $\Delta_{c_m} \approx 0.014$. In percentage, the variability of the four performances is 16.2%, 33.0%, 46.12% and -76.33%, thus significantly reduced w.r.t. the unconstrained optimal solution.

Eventually, Fig. 10 reports the PA, PC and PE flowfields associated to

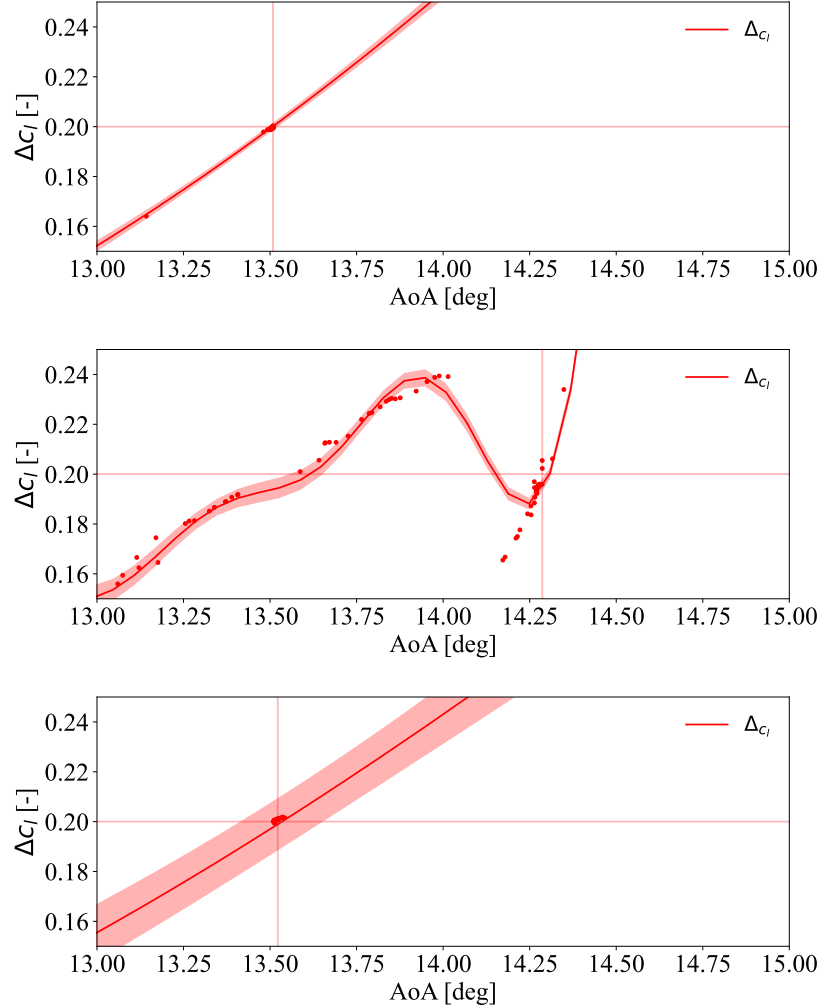


Figure 9: test case A, confidence-based optimization. Figure rows correspond to different acquisition functions, respectively EI, EIC and AEI. Plots reports an enlargement of the Δ_{c_l} surrogate centered on the region in the close proximity of the optimum (red vertical line) and of the credibility constraint threshold (red horizontal line).

the optimal configuration. Compared to results presented in Fig. 6, it is clear how the constraint drives the optimizer towards a credible solution. Indeed, all the five EPM simulations now predict an attached flow, thus excluding configurations characterized by separation or recirculation bubbles, for which we know the RANS model to be unreliable.




Figure 10: test case A, confidence-based optimization. EPM predictions of the TKE field developing around the airfoil in the optimal configuration. (a) PA. (b) PC. (c) PE.

5.2. Test Case B

This test case is concerned with maximizing the airfoil lift by varying its shape, according to the geometry parametrization presented in Sec. 4.4, and the AoA. Therefore, assuming $\mathbf{x} = (n_1, n_2, n_3, \alpha) \in \Omega \subset \mathbb{R}^4$, we seek

$$\arg \max_{\mathbf{x}} c_l(\mathbf{x}), \quad \text{s.t.} \quad \Delta_{c_l}(\mathbf{x}) \leq \tilde{\Delta}_{c_l} = 0.2. \quad (13)$$

5.2.1. Unconstrained Optimization

The optimum design parameters and performances are reported in Tab. 1. Figure 11(a) reports the record track of the four design parameters (min-max normalized). The red line indicates the parameters values of the optimal configuration. The picture shows that, after an initial exploration phase, the algorithm eventually sets to exploiting the optimum close proximity for n_1 , n_2 and α . On the contrary, the history associated to the n_3 parameter shows a more exploratory behavior.

Figure 11(b) shows the c_l^{BS} as resulting at each design point evaluated by the optimizer, revealing that the algorithm quickly finds high lift designs. The optimum is achieved applying the maximum camber at the most rearward point, see Fig. 12(a). That is, high-lift is achieved by applying a large deformation in the aft-part of the airfoil, resembling a plain flap device.

An a posteriori CFD assessment predicts $c_l^{BS} \approx 2.34$, $c_d^{BS} \approx 0.074$, $c_l^{BS}/c_d^{BS} \approx 31.7$ and $c_m^{BS} \approx 0.22$ for the optimal airfoil configuration, see Tab. 1. Figure 12(b) reports the polar graph comparing the performances predicted by the five EPM simulations. Not surprisingly, polar plots reveal a high prediction variability due to the turbulence closure uncertainty. In particular, the analysis returns a $\Delta_{c_l} \approx 0.52$, $\Delta_{c_d} \approx 0.006$, $\Delta_{c_l/c_d} \approx 19.6$ and $\Delta_{c_m} \approx 0.04$, with a relative variability, w.r.t. the lowest predicted value,




Figure 11: test case B, unconstrained optimization history. (a) Record track of the four design parameters (min-max normalized). (b) Record track of the targeted performance as predicted by the baseline computational model.




Figure 12: test case B, unconstrained optimization. (a) Comparison of the optimal (red) and the baseline (black) airfoil shape. (b) Comparison of the aerodynamic performances resulting from the EPM analysis of the optimal airfoil configuration.

of respectively 27.3%, 79.7%, 128.32% and 23.4%. Therefore, the optimal solution is associated to not credible RANS predictions.

Figure 13(a-c) shows, respectively the PA, PC and PE turbulent kinetic energy fields in the portion of the domain surrounding the airfoil. Clearly, all EPM simulations predict a highly turbulent separated flow in the close proximity trailing edge.




Figure 13: test case B, unconstrained optimization. EPM predictions of the TKE field developing around the airfoil in the optimal configuration. (a) PA. (b) PC. (c) PE.

5.2.2. Confidence-based Optimization

The c_l is now maximized under the constraint of that $\Delta_{c_l} < \tilde{\Delta}_{c_l} = 0.2$. The optimization history is reported Fig. 14, each row is associated to a different \mathcal{AF} , respectively EI, EIC and AEI. Similar optimal performances are achieved using the diverse learning functions, see Tab. 1. Despite the very similar performance delivered (within the 1% deviation), the EIC optimal designs slightly differs in n_3 and AoA. Possibly, these differences are related to the probabilistic treatment of the credibility constraint which leads, again, to a more exploratory behavior. Increasing the computational budget could lead to improving the EIC convergence to the optimum.

The record track of the four design parameters, reported on the left-hand column of Fig. 14, shows the different exploration/exploitation trade-off offered by the different \mathcal{AF} . Clearly, the AEI is endowed with an high initial exploratory behavior while favoring exploitation at later stages. The EI and the EIC reveal instead an alternating mixing of exploration/exploitation. On the right column of Fig. 14, the optimizer history reports red crosses to indicate designs that are rejected because the RANS model is deemed not reliable. Again, the design rejection rate strongly depends on the \mathcal{AF} employed (37/100 for EI, 75/100 for EIC, and 21/100 for AEI).

The a posteriori analysis of the EI optimum design predicts $c_l^{BS} \approx 2.21$, $c_d^{BS} \approx 0.053$, $c_l^{BS}/c_d^{BS} \approx 41.7$ and $c_m^{BS} \approx 0.246$. The EPM analysis predicts $\Delta_{c_l} \approx 0.196$ (0.1986), $\Delta_{c_d} \approx 0.005$, $\Delta_{c_l/c_d} \approx 6.57$ and $\Delta_{c_m} \approx 0.032$, with a relative variability, w.r.t. the lowest predicted value, of respectively 9.22%, 9.53%, 17.91% and 13.87%. Thanks to the credibility constraint, the variability of all the performances is now significantly reduced, in particular for the c_d and c_l/c_d coefficients. The polar graphs comparing the aerodynamic coefficients c_l , c_d , c_m and c_l/c_d from the EPM analysis of the optimal designs are




Figure 14: test case B, confidence-based optimization. Figures rows correspond to, respectively, the EI, EIC and AEI functions. Left column: record track of the design parameters. Right column: record track of the c_l^{BS} (red crosses indicate rejected designs).

reported in Fig. 15. With respect to the unconstrained optimization setting, the algorithm is trading a certain amount of performance for more reliable RANS predictions. Interestingly, the comparison against the unconstrained optimal design reveals a slightly thinner airfoil and a much lower AoA. Nev-

ertheless, the optimum airfoil shape still consists in a highly cambered profile resembling an airfoil with a plain flap device, see Fig 15.

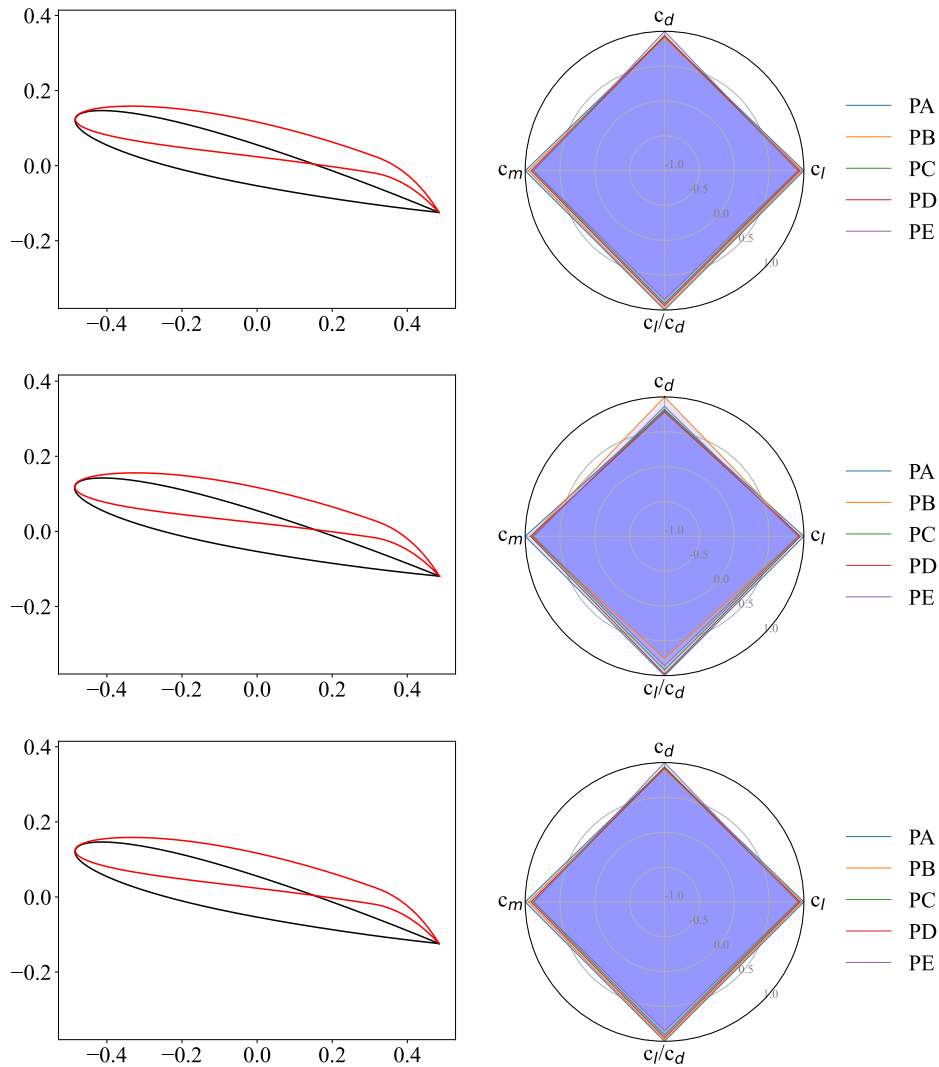


Figure 15: test case B, confidence-based optimization. Figure rows correspond to, respectively, the EI, EIC and AEI functions. Left column: optimal design (red) superimposed to the baseline one (black). Right column: EPM analysis of the optimal design.

Eventually, in Fig. 16 we report the PA, PC and PE flowfields. Not surprisingly, the flow now remains attached to the airfoil suction side for almost

the entire chord length and only a small region is separated. The effect of constraining the optimization to the RANS credibility constraint is particularly evident comparing Fig. 16 against its (unconstrained) counterpart Fig. 13.




Figure 16: test case B, confidence-based optimization (EI). EPM predictions of the TKE field developing around the airfoil in the optimal configuration. (a) PA. (b) PC. (c) PE.

5.3. Test Case C

In Sec. 5.2 we managed to achieve a design of maximum lift under a credibility constraint acting on c_l . We now slightly change the approach towards a more common engineering practice. Namely, we prescribe a target value $c_l^{\text{target}} = 1.0$ and we seek for a design capable of delivering the desired performance while also aiming at minimizing the drag. The credibility constraint is now applied to the drag coefficient, in order to select configurations for which the resistance force can be confidently predicted. Assuming again $\mathbf{x} = (n_1, n_2, n_3, \alpha) \in \Omega \subset \mathbb{R}^4$, we seek

$$\begin{aligned} \max_{\mathbf{x} \in [0,1]^d} & \quad - (c_l^{BS}(\mathbf{x}) - c_l^{\text{target}})^2 - \beta c_d^{BS}(\mathbf{x}), \\ \text{subject to} & \quad \Delta_{c_d}(\mathbf{x}) \leq \tilde{\Delta}_{c_d} = 0.005. \end{aligned} \quad (14)$$

The 0.005 threshold was selected according to the value that the $\Delta_{c_d}(\mathbf{x})$ indicator assumes for the baseline profile in the linear region of the polar, reported in Fig. 2(b), for an angle of attack lower than 14 degrees. The coefficient $\beta = 2$ weights the relevance of the drag performance contribution to the cost function, practically implementing a so-called scalarized multi-objective approach. Since we are practically solving a bi-objective scalarized optimization problem, the relevance of the relative weighting among the terms included in the objective function has to be stressed. Varying the relative weighting basically varies the trade-off among achieving the desired performance or obtaining a design of minimum drag, the choice of which is left to the experienced designer, based on the specific application.

In solving the optimization problem, we rely on the EIC acquisition function only. The goal is to achieve an airfoil shape of prescribed performance and minimum drag while relying on a credible prediction of c_d . In particular, the credibility of the drag force prediction should be comparable to the one associated to the baseline geometry operating in the linear part of the polar. In a broader context, one could consider the overall aircraft performances i.e., target not just the airfoil but the whole wing, fuselage, tail and so on. In this perspective, multiple credibility constraints may be applied simultaneously w.r.t. the diverse aerodynamic qualities of interest for the designer, so to allow for a reliable assessment of the aircraft qualities in a determined flight condition.

The $\Delta_{c_d}(\mathbf{x})$ is quite sensitive to variation of the design parameters, this aspect being particularly important for highly cambered profiles at a high angle of attack. Indeed, the strong curvature at the trailing edge makes the flow prone to a localized separation, which of course results into large uncertainty associated to the predicted c_d . For this test case only, we reduced the range of the AoA to 0-15 degrees, so to directly exclude a large portion of the design space associated to solutions which surely do not fulfill the credibility constraint. Note that we are considering the deformation of the airfoil and that the bounds of the parameters controlling the deformation of the airfoil are instead maintained as reported in Sec. 4.4.

Figure 17(a) reports the optimization history of the four design parameters, revealing a highly explorative behavior from the optimizer. Figure 17(b) reports instead the record track of the optimum of the cost function as specified in Eq. (14), which we generically indicate as J . As the exploration proceeds, configurations are deemed adequate or rejected due to failing in fulfilling the credibility constraint, with an overall rejection rate of about 51 %. The parameters of the constrained optimal design are $n_1 = 5.01$, $n_2 = 4.92$, $n_3 \approx 1.33$ and $\text{AoA} \approx 3.13$. The a posteriori analysis predicts $c_l^{BS} \approx 0.994$, $c_d^{BS} \approx 0.009$, $c_l^{BS}/c_d^{BS} \approx 107.20$ and $c_m^{BS} \approx 0.148$, see Tab. 1. The optimal airfoil shape is quite different from what achieved with the other optimization settings. Indeed, it is a thicker and much less cambered, see Tab. 1 and Fig. 18(a).

The EPM analysis, Fig. 18(b), predicts $\Delta_{c_l} \approx 0.008$, $\Delta_{c_d} \approx 0.003$, $\Delta_{c_l/c_d} \approx 24.58$ and $\Delta_{c_m} \approx 0.002$, with a relative variability, w.r.t. the baseline value, of respectively 0.9%, 33.9%, 32.8% and 1.5%. In percentage, the uncertainty estimates of c_d , and therefore the efficiency Δ_{c_l/c_d} still suffers from a large variability which is a consequence of the very low baseline




Figure 17: test case C, confidence-based optimization (EIC) optimization history. (a) Record track of the four design parameters (min-max normalized). (b) Record track of the optimum of the cost function (J) as specified in Eq. (14).




Figure 18: test case C, confidence-based optimization (EIC). (a) Comparison of the optimal (red) and the baseline (black) airfoil shape. (b) Comparison of the aerodynamic performances resulting from the EPM analysis of the optimal airfoil configuration.

drag coefficient value employed to compute the percentages. Nonetheless, the credibility constraint is fulfilled and confidence in RANS predictions can be assumed up to the imposed threshold. In other words, the optimization is feasible for $\Delta_{c_d}(\mathbf{x}) < 0.005$, which corresponds to requiring the credibility of predictions to be comparable to that relative to simulations concerning a symmetric airfoil at a low (even null) angle of attack.

As mentioned, in the setting we employ we seek for designs capable of delivering a prescribed lift by minimizing the distance of the predicted performance from the desired one, while aiming also at minimizing the drag in a scalarized multi-objective approach. Figure 19(a) and (b) report the record track of the aerodynamic performances, respectively the lift and the drag, associated to the optimal design, as predicted by the baseline computational model. Clearly, the optimizer quickly achieves designs capable of delivering a c_l in the proximity of the prescribed value, indicated as a blue line in Fig. 19(a), while c_d initially increases. As the optimizer explores/exploits the design space, new designs with $c_l \approx 1.0$ and a better value of the drag coefficient are found.




Figure 19: test case C, confidence-based optimization (EIC). Record track of the aerodynamic performances associated to the optimal design, as predicted by the baseline computational model. (a) Lift coefficient. (b) Drag coefficient.

Eventually, in Fig. 20(a-c) we report the PA, PC and PE flowfields as resulting from this constrained optimization exercise. In all cases, the flow remains attached to the surface and, because of the strict credibility constraint enforced, no relevant sign of separation can be found.

6. Conclusions

In this work, we propose a confidence-based optimization approach robust to the epistemic uncertainty affecting the mathematical structure of turbulence closures typically required to model the Reynolds Stress Tensor

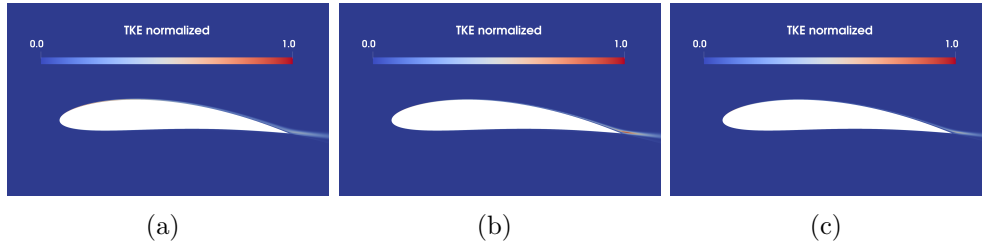


Figure 20: test case C, test case C, confidence-based optimization (EIC). EPM predictions of the TKE field developing around the airfoil in the optimal configuration. (a) PA. (b) PC. (c) PE.

appearing in the Reynolds-Averaged Navier-Stokes equations. Given that turbulence closures are formally equivalent to the models required to recover the contribution of the subfilter stresses in Large Eddy Simulations, the boundaries of application for the proposed approach could be straightforwardly extended. The optimization strategy attempts to maximize (or minimize) a targeted Quantity of Interest subject to the constraint of that the prediction credibility, due to turbulence closure uncertainty, satisfies a given threshold. That should not be intended as an approach to obtain aerodynamic shapes robust to turbulence uncertainty, but rather as a robust procedure that constrains the search for the optimum to the portion of the design space where the RANS model is expected to be credible. Since designs associated with unreliable predictions are avoided, the efficiency of the optimization process is improved.

We have shown the potential of the proposed approach for a set of test cases of interest for aerospace applications namely, a single element airfoil resembling a morphing wing section in a high-lift configuration. Nonetheless, the potential of the proposed methodology should be envisioned beyond the test cases included here e.g., one may aim at optimizing air-intakes, engine cowls, propellers or the aerodynamic of a car. In general, the confidence-based approach leads to designs capable of delivering a performance lower than what achieved in the unconstrained setting. Nevertheless, the optimal designs are bounded to credible computational predictions.

The evaluation of turbulence uncertainty estimates, obtained by means of the Eigenspace Perturbation Method, comes at an significantly increased computational cost. At each iteration of the optimization process, six independent evaluation of the full computational model are required. Namely,

the baseline and the five EPM computations. The procedure also presents additional difficulties because each EPM simulation behaves differently, presenting faster/slower convergence rates or suffering from different numerical instabilities, thus requiring a careful fine tuning of the CFD solver configuration. Nevertheless, the proposed approach has a high potential that justifies the increased complexity, especially when the interest is pledged to the optimization of three-dimensional geometries where turbulence effects play an even more relevant role. Constraining the optimization to the domain of credibility of RANS (or LES) predictions avoids achieving (and building) geometries that will not perform as expected in reality. This could translate into a reduction of wind tunnel testing time, into significant cost savings, and into a shortened project lead time. Future work should be devoted to integrating the confidence-based approach with uncertainty quantification techniques capable of accounting for aleatory uncertainties such as, for instance, unknown boundary conditions or fluid properties. The potential of the proposed approach is foreseen to be even higher for unsteady problems i.e., helicopter rotor blade optimization, although the definition of a suitable unsteady credibility constraint must still be achieved.

Acknowledgements

This research was funded by the UTOPIAE Marie Curie Innovative Training Network, H2020-MSCA-ITN-2016, Grant Agreement number 722734. Experiments presented in this paper were carried out using the PlaFRIM experimental testbed, supported by Inria, CNRS (LABRI and IMB), Université de Bordeaux, Bordeaux INP and Conseil Régional d'Aquitaine (see <https://www.plafrim.fr/>)

References

- [1] K. Duraisamy, G. Iaccarino, and H. Xiao. Turbulence Modeling in the Age of Data. *Annual Review of Fluid Mechanics*, 51(1):357–377, 2019.
- [2] M. Emory, R. Pecnik, and G. Iaccarino. Modeling Structural Uncertainties in Reynolds-Averaged Computations of Shock/Boundary Layer Interactions. *AIAA 2011-479*, 2011.
- [3] M. Emory, J. Larsson, and G. Iaccarino. Modeling of structural uncertainties in Reynolds-averaged Navier-Stokes closures. *Physics of Fluids*, 25(11):110822, 2013.

- [4] G. Iaccarino, A.A. Mishra, and S. Ghili. Eigenspace perturbations for uncertainty estimation of single-point turbulence closures. *Phys. Rev. Fluids*, 2:024605, Feb 2017.
- [5] A.A. Mishra, J. Mukhopadhaya, G. Iaccarino, and J.J. Alonso. Uncertainty Estimation Module for Turbulence Model Predictions in SU2. *AIAA Journal*, 57(3):1066–1077, 2019.
- [6] C. Górlé, M. Emory, J. Larsson, and G. Iaccarino. Epistemic uncertainty quantification for rans modeling of the flow over a wavy wall. *Center or Turbulence Research Annual Briefs*, 2012.
- [7] C. Górlé and G. Iaccarino. A framework for epistemic uncertainty quantification of turbulent scalar flux models for Reynolds-averaged Navier-Stokes simulations. *Physics of Fluids*, 25(5):055105, 2013.
- [8] C. Górlé, S. Zeoli, M. Emory, J. Larsson, and G. Iaccarino. Epistemic uncertainty quantification for Reynolds-averaged Navier-Stokes modeling of separated flows over streamlined surfaces. *Physics of Fluids*, 31(3):035101, 2019.
- [9] A.A. Mishra, J. Mukhopadhaya, J.J. Alonso, and G. Iaccarino. Design exploration and optimization under uncertainty. *Physics of Fluids*, 32(8):085106, 2020.
- [10] H. Xiao and P. Cinnella. Quantification of model uncertainty in RANS simulations: A review. *Progress in Aerospace Sciences*, 108:1–31, 2019.
- [11] L. Jofre, S. Domino, and G. Iaccarino. A Framework for Characterizing Structural Uncertainty in Large-Eddy Simulation Closures. *Flow Turbulence Combust*, 100:341–363, 08 2018.
- [12] E. Torenbeek. *Advanced Aircraft Design: Conceptual Design, Analysis and Optimization of Subsonic Civil Airplanes*. John Wiley and Sons, Ltd, 2013.
- [13] N. Razaaly, G. Gori, G. Iaccarino, and P.M. Congedo. Optimization of an ORC supersonic nozzle under epistemic uncertainties due to turbulence models. In *GPPS 2019 - Global Power and Propulsion Society*, Zurich, Switzerland, January 2019.

- [14] L.W. Cook, A.A. Mishra, J.P. Jarrett, K.E. Willcox, and G. Iaccarino. Optimization under turbulence model uncertainty for aerospace design. *Physics of Fluids*, 31(10):105111, 2019.
- [15] A.A. Mishra and G. Iaccarino. Theoretical analysis of tensor perturbations for uncertainty quantification of reynolds averaged and subgrid scale closures. *Physics of Fluids*, 31(7):075101, 2019.
- [16] A.J. Simonsen and P.A. Krogstad. Turbulent stress invariant analysis: Clarification of existing terminology. *Physics of Fluids*, 17(8):088103, 2005.
- [17] U. Schumann. Realizability of Reynolds-stress turbulence models. *Physics of Fluids*, 20(5):721–725, 1977.
- [18] H. Xiao, J.L. Wu, J.X. Wang, R. Sun, and C.J. Roy. Quantifying and reducing model-form uncertainties in Reynolds-averaged Navier–Stokes simulations: A data-driven, physics-informed Bayesian approach. *Journal of Computational Physics*, 324:115–136, 2016.
- [19] NASA Langley Turbulence Modeling Resource website. <http://turbmodels.larc.nasa.gov>. Accessed: 2019-11-11.
- [20] D.C. Jespersen, T.H. Pulliam, and M.L. Childs. OVERFLOW: Turbulence Modeling Resource Validation Results. Technical Report NASA-2016-01, NASA Ames Research Center, Moffett Field, CA, 2010.
- [21] F. Menter. Zonal Two Equation k-w Turbulence Models For Aerodynamic Flows. *AIAA 1993-2906. 23rd Fluid Dynamics, Plasmadynamics, and Lasers Conference*, 1993.
- [22] F. Palacios, J.J. Alonso, K. Duraisamy, M. Colonno, J. Hicken, A. Aranake, A. Campos, S. Copeland, T.D. Economon, A. Lonkar, T. Lukaczyk, and T. Taylor. Stanford University Unstructured (SU²): An open-source integrated computational environment for multi-physics simulation and design. *AIAA 2013-287*, 2013.
- [23] T.D. Economon, D. Mudigere, G. Bansal, A. Heinecke, F. Palacios, J. Park, M. Smelyanskiy, J.J. Alonso, and P. Dubey. Performance optimizations for scalable implicit RANS calculations with SU2. *Computers & Fluids*, 129:146–158, 2016.

- [24] B. van Leer. Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov’s method. *Journal of Computational Physics*, 32(1):101–136, 1979.
- [25] G. Gori, O.P. Le Maître, and P.M. Congedo. On the sensitivity of structural turbulence uncertainty estimates to time and space resolution. *Computers & Fluids*, 229:105081, 2021.
- [26] C. Ladson. Effects of independent variation of mach and reynolds numbers on the low-speed aerodynamic characteristics of the NACA 0012 airfoil section. *NASA Technical Report*, (NASA-TM-4074):L-16472, NASA 1.15: 4074, 1988.
- [27] E. Brochu, V.M. Cora, and N. de Freitas. A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning. eprint arXiv:1012.2599, arXiv.org, December 2010.
- [28] M. Hoffman, E. Brochu, and N. de Freitas. Portfolio Allocation for Bayesian Optimization. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*, UAI’11, page 327–336, Arlington, Virginia, USA, 2011. AUAI Press.
- [29] J.R. Gardner, M.J. Kusner, Z. Xu, K.Q. Weinberger, and J.P. Cunningham. Bayesian Optimization with Inequality Constraints. In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32*, ICML’14, page II–937–II–945. JMLR.org, 2014.
- [30] M. Sacher. *Méthodes avancées d’optimisation par méta-modèles – Application à la performance des voiliers de compétition*. Theses, Ecole nationale supérieure d’arts et métiers - ENSAM, September 2018.
- [31] D. Huang, T.T. Allen, W.I. Notz, and N. Zeng. Global Optimization of Stochastic Black-Box Systems via Sequential Kriging Meta-Models. *J. of Global Optimization*, 34(3):441–466, March 2006.
- [32] J. Möckus, V. Tiesis, and A. Zilinskas. The Application of Bayesian Methods for Seeking the Extremum. *Towards Global Optimization*, 2(117-129):2, 1978.

- [33] D.R. Jones, M. Schonlau, and W.J. Welch. Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, 13(4):455–492, Dec 1998.
- [34] D.J. Lizotte. *Practical Bayesian Optimization*. PhD thesis, CAN, 2008. AAINR46365.
- [35] J. Mockus. Application of Bayesian approach to numerical methods of global and stochastic optimization. *Journal of Global Optimization*, 4(4):347–365, Jun 1994.
- [36] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [37] P. Virtanen, R. Gommers, and SciPy 1.0 Oliphant, T.E. et al. SciPy 1.0—Fundamental Algorithms for Scientific Computing in Python. *arXiv e-prints*, page arXiv:1907.10121, Jul 2019.
- [38] B. Kulfan and J. Bussoletti. *”Fundamental” Parametric Geometry Representations for Aircraft Component Shapes*.
- [39] J. Hoschek, D. Lasser, and L.L. Schumaker. *Fundamentals of Computer Aided Geometric Design*. A. K. Peters, Ltd., USA, 1993.
- [40] G. Farin. *8 - B-Spline Curves*. The Morgan Kaufmann Series in Computer Graphics. Morgan Kaufmann, San Francisco, fifth edition edition, 2002.