POLITECNICO DI MILANO
Department of Aerospace Science and Technology
Doctoral Programme In Aerospace Engineering – XXVIII Cycle

Robust Shape Optimization of Fixed and Morphing Rotorcraft Airfoils

Doctoral Dissertation of:
Francesca Fusi

Supervisor:
Prof. Giuseppe Quaranta

Co-Supervisor:
Dr. Pietro Marco Congedo

Tutor:
Prof. Giuseppe Gibertini

The Chair of the Doctoral Program:
Prof. Luigi Vigevano

2016

Keywords:

Printed in Italy

Contents

Abstract

The thesis presents a methodology for robust optimization applied to the problem of finding optimal shapes of fixed and morphing airfoils for helicopter rotor blades. To this purpose, the employment of robust approaches is explored to demonstrate their capability of indicating interesting design in the application to helicopter rotor blade airfoils. Because robust optimization increases the computational cost of the optimization process, new uncertainty-based optimization methods are developed. The goal of this development is to improve the numerical efficiency, thereby making robust approaches more attractive to aerodynamic applications. The first method is based on a multi-fidelity approach to the estimation of the aerodynamic performance. The second strategy leverages an adaptive uncertainty quantification method to reduce the computational cost associated with poor design vectors inside the optimization loop. The methods are applied to algebraic test cases and to the robust optimization of fixed airfoils in the hovering condition. This first application demonstrates the effectiveness of the proposed strategies and it also provides the mean for a discussion on the impact of robustness criteria on airfoil design. The robust optimization framework is then used to tackle the forward flight case. In this flight condition, two representative position of the blade over the azimuth are considered, i.e. the advancing side and the retreating side. Deterministic and robust optimal shapes are compared and robust airfoils demonstrate to be able to trim the helicopter with performance close to the deterministic values. The post-processing analysis of this application includes the comparison of the estimate computed by the steady models used in the optimization loop throughout the thesis and the objective function estimates obtained from an unsteady CFD model. Finally, the application of morphing

airfoil is considered. The morphing strategy consists in a variable camber airfoil and it is conceived with the intent of changing its shape at the 1/rev frequency to enhance aerodynamic performance. The optimization of morphing airfoils presented in the thesis takes into account the aerodynamic performance, where technological aspects are accounted for by means of geometrical constraints. The gain of the morphing airfoil strategy is assessed not only from a deterministic point of view, but also when considering uncertainty in the operating conditions.

Sommario

La tesi presenta una metodologia per l'ottimizzazione aerodinamica robusta applicata alla ricerca dei profili ottimi fissi e morphing per le pale di rotore di elicottero. A tale scopo, l'utilizzo di approcci robusti è esplorato per dimostrare la loro capacità di indicare regioni dello spazio di design interessanti ai fini dell'applicazione elicotteristica. Poichè affrontare un problema di ottimizzazione robusta aumenta il costo computazionale del processo di ottimizzazione, la tesi presenta due nuovi metodi per l'ottimizzazione con incertezze. Lo scopo dello sviluppo di questi metodi è di migliorare l'efficienza numerica, permettendo così che l'ottimizzazione robusta venga applicata piú agevolmente a casi aerodinamici. Il primo metodo è basato su una strategia "multi-fidelity" che utilizza due modelli aerodinamici di differente complessità e costo. Il secondo metodo si basa su un raffinamento adattivo del piano stocastico per la stima delle statistiche della prestazione aerodinamica. I due metodi sono applicati ad alcuni casi algebrici e al problema di ottimizzazione robusta per profili fissi di pale di elicottero in volo a punto fisso. Questa prima applicazione dimostra l'efficacia delle strategie proposte e fornisce la base per una discussione sull'impatto dei criteri robusti sul progetto di profili. Nel caso di volo avanzato, vengono considerate due posizioni della pala lungo l'angolo azimutale: pala avanzante e pala retrocedente. I risultati ottenuti tramite l'ottimizzazione robusta e quella deterministica sono confrontati e i profili robusti dimostrano di poter garantire un adeguato livello di portanza con prestazioni analoghe a quelle dei profili deterministici. L'analisi dell'ottimizzazione in questa condizione include il confronto delle stime calcolate con i modelli stazionari usati nel loop di ottimizzazione con un modello di fluidodinamica computazionale instazionario. Infine, l'applicazione di profili morphing è

studiata per migliorare le prestazioni aerodinamiche. La strategia morphing consiste in un profilo in grado di modificare la sua linea media ed è pensato in modo da garantire una modifica della forma ad una frequenza pari alla 1/giro. L'ottimizzazione dei profili morphing presentata nella tesi prende in considerazione le prestazioni aerodinamiche, mentre gli aspetti tecnologici sono considerati attraverso vincoli geometrici. Il guadagno dei profili morphing è valutato non solo da un punto di vista deterministico, ma anche considerando incertezze nelle condizioni operative.

CHAPTER 1

Introduction

Helicopters play a unique role in modern aviation thanks to their ability to take-off and land vertically, fly forward or backward, climb or descend and hover. Thanks to this unique ability, helicopters are capable of missions that are unsuitable to fixed-wing aircraft, such as search and rescue in arduous conditions, police surveillance, offshore missions, and other specific military operations. This fostered the great employment of helicopter both in the civil and military aviation and the development of new technologies for this machine since the 1930s.

Aerodynamic forces to fly a helicopter are generated by the rotor, a system of rotating wings, called blades, that provide the lift to sustain flight, the propulsive force to move the aircraft through the air and the control to trim the aircraft or change its attitude. In classical configurations, a single main rotor is employed and a sideward, tail rotor is required to provide anti-torque and directional control. The presence of rotating elements no doubt increases the complexity of helicopters from both aerodynamic and mechanic points of view. With regard to aerodynamics, the flowfield around the helicopter is dominated by unsteady, three-dimensional effects and interactions among the helicopter bodies (rotor blades, main body and tail), as well as by vortices and wakes (see Fig. 1.1). This system is so complicated that, after many years of studies, it still defies an adequate and comprehensive description [1]. From a mechanic point of view, the presence of the rotor introduces the structural dynamics associated with the degrees of freedom of each blade. In conventional rotors with fully articulated hubs, helicopter blades are connected to the shaft by means of hinges

which permit flap and lead-lag motions. Pilot control is exerted through a change in the pitch angle of the blades, either cyclically or collectively. In addition, blades are slender bodies undergoing twisting and bending. The aeroelastic degrees of freedom, together with the remaining rigid motions, are coupled with the aerodynamic loads.

The aerodynamic flowfield, the blade dynamics and aeroelasticity, the pilot control are all relevant aspects of the environment surrounding the rotor, which must be taken into account in helicopter rotor design. This poses a great challenge to rotor designers, because the high-fidelity analysis of all these aspects and their mutual interaction requires the employment of tools with prohibitive cost. In comprehensive rotor analysis models, parsimonious models of each sub-system (e.g. aerodynamics, structural dynamics) are used to limit the cost necessary to estimate loads, displacements and overall performance of the rotor [1, p.805]. Examples are the CAMRAD code [2] developed by Johnson [3, 4] and MBDyn [5, 6] developed at the Department of Aerospace Science and Technology of the Politecnico di Milano [7]. A modest cost permits parametric calculations that are crucial in the design stage. To some extent, the employment of cheap models and the combined use of models of different fidelity are intrinsically rooted in the design of rotors.

Another major challenge of rotor design is the definition of the performance goals. In fact, helicopter rotors are required to fulfill many different mission requirements. This translates into the necessity for rotor blades to operate with satisfactory performance in extremely different operating conditions, depending on the flight

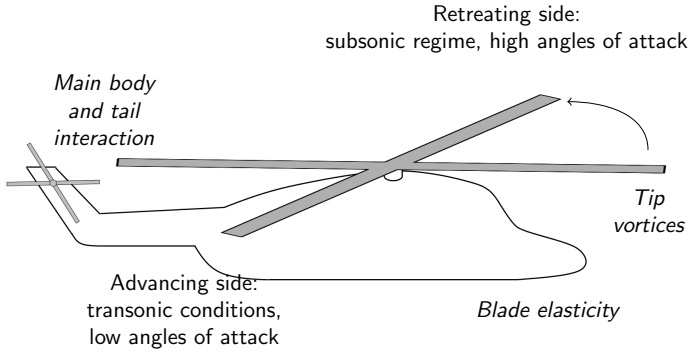

Figure 1.1: Characteristics of the complex flowfield around a helicopter rotor in forward flight.

condition, pitch control and advancing or climbing velocity. If the design process is tackled with an optimization approach, careful analysis is necessary in defining the interesting operating conditions and in setting the objectives. In most cases, a trade-off solution must be pursued, whose performance is sub-optimal in each condition over the flight envelope.

In such a context, the goal of aerodynamic design is to improve the overall performance of the rotor by increasing thrust to accommodate for an increased payload, reducing the torque due to aerodynamic drag at a given thrust level, alleviating aerodynamic loads on the blade structure, enlarging the helicopter rotor flight envelope, reducing vibrations. In the following sections, a discussion is given on the basic features of rotorcraft aerodynamics. These basics found the principles of fixed and morphing airfoil design which are the subject of this thesis and are presented in the subsequent sections. Finally, the objectives and the structure of the thesis are detailed.

1.1 Basics of rotorcraft aerodynamics

Understanding rotor aerodynamics is key to predict helicopter performance and to improve helicopter aerodynamic design. The aerodynamic loads acting on the rotor blades depend on the local angle of attack and the local dynamic pressure of the blades. These conditions are influenced by the blade motion, the advancing velocity and the effects of wakes and flowstream across the rotor disk.

In fully articulated rotors, as the one presented Fig. 1.2, the blade is allowed to flap up and down freely about a “flapping hinge” at the blade root. In addition, a “lead-lag” hinge is also present to relieve the in-plane bending moments generated by the Coriolis forces caused by the flapping motion. In hingeless or bearingless rotor hub, the conventional flapping and lagging hinges are replaced by flexible structural elements. A great advantage of this hub is the elimination of hinges which leads to a great simplification of hub design and also reduced aerodynamic drag [8, p.279]. In both cases, the blades are allowed another degree of freedom at the blade root which is controlled by the pilot: the pitching motion. In fact, the magnitude and direction of the thrust produced by the rotor are controlled by changing the angle of attack of all the blades either collectively or cyclically. Collective pitch control changes the blade pitch angle of each blade in the same way in order to increase the level of thrust produced by the rotor. When cyclic pitch control is introduced, the blade pitch angle is varied periodically over the azimuth angle. Under azimuthally varying aerodynamic loads, the lift produced by the rotor tilts, so that a propulsive force is created to permit forward (sideward or rearward) flight. Such motion requires

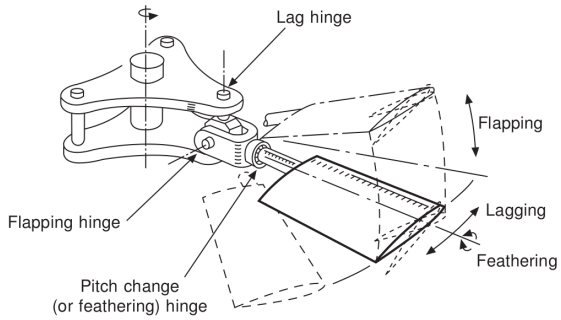

Figure 1.2: Motion in a fully articulated rotor hub (from Ref. [8, p.2]).

freedom of rotation of the blades about the quarter-chord line, which can be achieved by introducing a "feathering hinge" at the blade root. Thus, blade motion features three rigid degrees of freedom and it is subject to aerodynamic forces, centrifugal forces, structural forces and control forces.

Finally, at the tip of the blade, the local dynamic pressure is very high and strong vortices are produced [1]. Tip vortices are typically convected rearward of the blade and downward, producing helical trajectories around the rotor. In hovering flight, the vortices stay very close to the body and the rotor itself, whereas in forward flight they are convected to the tail of the helicopter. In any cases, these vortices interact with the blades themselves or with the other parts of the helicopter, modifying the aerodynamic loads acting on them. The same applies for the wakes trailed from the blades.

1.1.1 Hovering flight

Hovering flight is the condition in which the advancing and climbing velocities of the helicopter are null and the aerodynamic forces produced by the rotor are used to balance the helicopter weight and mantain its position. In this condition, the lift force produced by the rotor tends to draw the blade upward and places the tip path plane above the rotor hub as presented in Fig. 1.3a. The angle formed by the blade is called coning angle and it is essentially constant over the period of rotation of the rotor. The drag acting on the blades causes the tip of the blade to lag slightly behind the root of the blade. The centrifugal force acting radially on the rotating blade counteracts the coning and lag tendency.

The hover condition is essentially steady and azimuthally axisymmetric with regard to the blade. This means that each blade encounters the same velocity for each

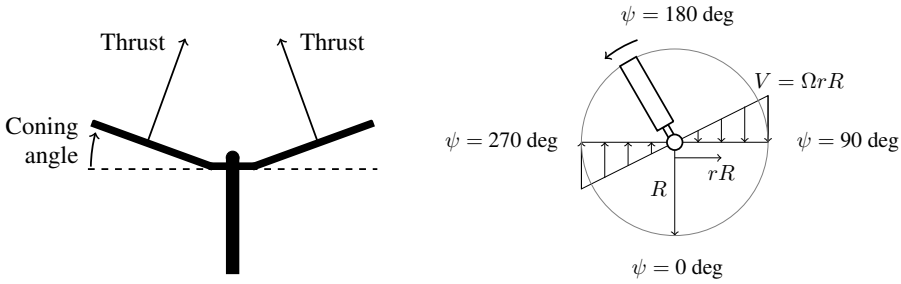

(a) Coning angle in forward flight.

(b) Scheme of incident speed on the rotor blade in hover.

Figure 1.3: Scheme of hovering condition.

azimuth angle over the period of rotation of the rotor as presented in Fig. 1.3b. Perturbations to this perfectly steady and axisymmetric condition may arise from the effects of tip and blade vortices, and small cyclic pitch control that may be necessary to trim the aircraft. The variation of the velocity V with the radius of the blade appears linear, because the only contribution is given by the rotor angular velocity Ω, so that

$$V = \Omega r R, \tag{1.1}$$

where r is the ratio between the local radius and the blade radius R. The local angle of attack is affected by the blade twist, the induced velocity V_i and the blade pitch angle, resulting from collective pitch control, possible small cyclic control and coning angle. The angle of attack remains constant with the azimuth angle, while it may vary radially due to twist and induced velocity distribution.

1.1.2 Forward flight

A helicopter rotor in forward flight encounters a periodically changing airflow. In fact, the advancing speed V_∞ of the helicopter adds to the rotational speed of the rotor, so that the local velocity at a specific radius rR is given by

$$V = V_\infty \sin\psi + \Omega r R, \tag{1.2}$$

where ψ is the azimuth angle which is equal to zero when the blade is pointing downstream as shown in Fig. 1.3b. The forward speed of the helicopter is expressed by

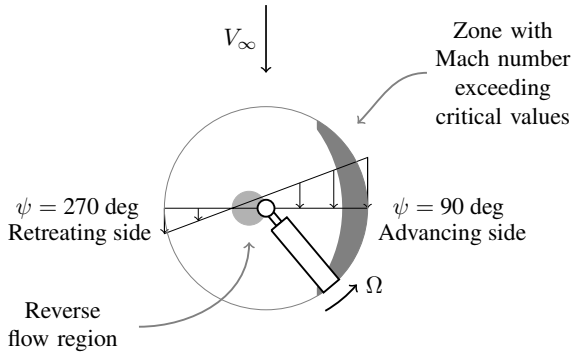

Figure 1.4: Scheme of incident speed on the rotor blade in forward flight.

means of the non-dimensional advance ratio μ, defined as follows

$$\mu = \frac{V_\infty}{\Omega R}, \tag{1.3}$$

or in other words as the ratio between the forward speed and the tip speed.

A typical distribution of the local velocity in forward flight is depicted in Fig. 1.4. In the advancing side, because of the summation of advancing speed and rotating speed, the local speed is so high that the outboard sections of the blade may encounter conditions exceeding the critical Mach number. In the retreating side (azimuth angle $\psi = 270$ deg) the inboard section may experience reverse flow. Here, the velocity is so low that the blades angle of attack is very high in order to develop a sufficient level of trust.

The change in the resultant angle of attack and incident velocity occurs at the rotational frequency of the rotor: 1/rev or once per revolution. Cyclic pitch is also applied with the same frequency. However, higher harmonics are introduced in this situation by the nonlinear blade dynamics and the nonlinear aerodynamic loads acting on the blades. A complete description of forward flight dynamics and aerodynamics can be found in Ref. [9, p.152] or Ref. [1, chap. 4].

1.2 Airfoil design for helicopter rotor blades

The choice of airfoil sections for helicopter rotors requires special consideration because significant improvements in rotor performance can be realized with the optimal

selection of airfoil shapes [1]. In Ref. [10, p.88] a typical workflow of aerodynamic rotor design supplied by Westland Helicopters is presented. Given the basic helicopter specifications in terms of maximum flight speed, payload and range, the step-by-step design begins with the definition of the tip speed, followed by the blade area, rotor radius and finally the number of blades. The selection of the airfoil of the blade is regarded as an a priori choice which impacts the subsequent steps. For instance, tip speed for a specific rotor angular velocity is associated with the maximum flight speed and it is limited by the onset of compressibility effects in forward flight. In fact, a higher advancing speed in forward flight increases the local Mach number of the advancing blade, resulting in the onset of strong shock waves that produces excessive torque on the rotor. In this regard, the employment of advanced blade section can reduce drag rise and increase the maximum flight speed of the new helicopter, attracting more customers and building an advantage over competitors. The use of advanced section affects also the blade area: a higher blade-loading coefficient C_T/σ is obtained with advanced sections that provide higher lift coefficient in the retreating side, resulting eventually in lower rotor area. This is shown in Fig. 1.5 from Ref. [10, p.90] where the use of the RAE 9615 airfoil with respect to the classical NACA 0012 airfoil yields a gain close to 10 % in blade area which translates into rotor overall weight. Furthermore, it can be proved that the figure of merit of the rotor, i.e. the measure of rotor efficiency, is related to the lift-to-drag ratio of the blade sections in hovering condition and to the average drag coefficient of the airfoil sections in forward flight [1, pp.295–301].

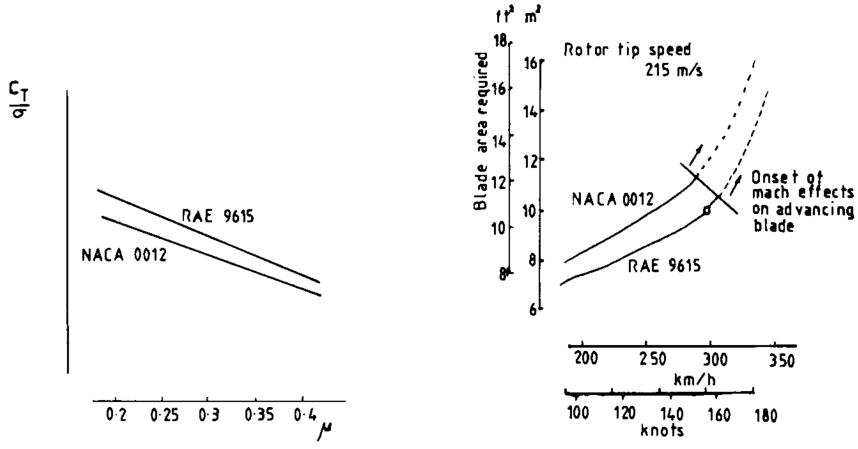

Figure 1.5: Impact of airfoil selection on blade required area (from Ref. [10, p.72]).

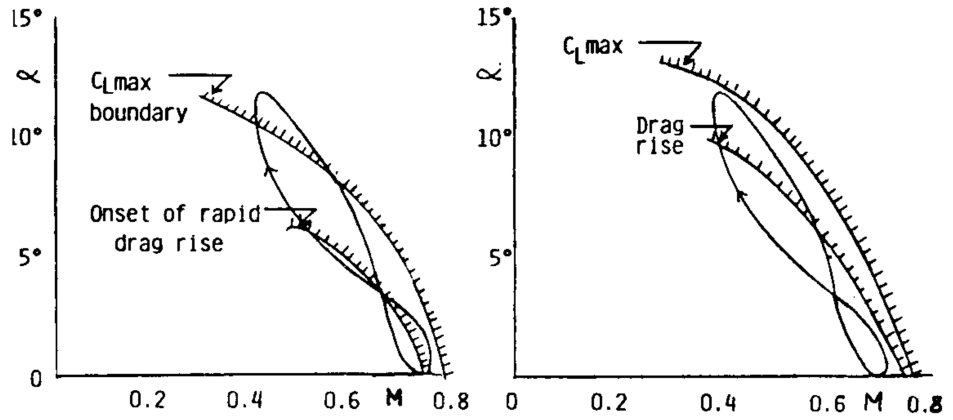

Figure 1.6: "Figure-of-eight" diagram for the NACA 0012 (left) and the NPL 9615 (right) airfoils with compressibility and stall boundaries (from Ref. [10, p.72]).

The selection of airfoil sections for helicopter rotor design is more complicated than in the case of fixed-wing applications, because rotors are not a "point" design [1, p.347]. In fact, a wide range of angle of attack and Mach number are encountered by a rotor blade section in the helicopter flight envelope [11]. A representation of the widely ranging conditions in forward flight is typically given by the "figure-of-eight" diagram which plots the operating conditions for a given blade radius rR at a particular value of the advancing ratio μ. Two representative diagrams for the NACA 0012 and the NPL 9615 airfoils are shown in Fig. 1.6, where the boundaries of static stall occurrence and drag rise are reported. In the advancing side, the blade section encounters high values of the Mach number and low values of the angle of attack, and may experience drag rise due to compressibility effects. Moving to the retreating side, the blade section encounters low values of the Mach number and high angle of attack which may cause separation and stall. Both retreating blade stall and the onset of strong shock waves on the advancing side limit the rotor performance and airfoil section design should address these issues with a careful airfoil selection. However, section with good performance in the advancing side are transonic airfoils, with modest thickness-to-chord ratio and low camber. On the other side, improving the behavior close to the stall limit typically means to increase the camber of the airfoil to cope with the high values of angle of attack. Figure 1.7 reports examples of sections used in helicopters that have been conceived to privilege either the advancing side or the retreating side. For instance, the SC1094-R8 airfoil was introduced because the original design of the UH-60 Black Hawk rotor blade based on the SC1095 airfoil for

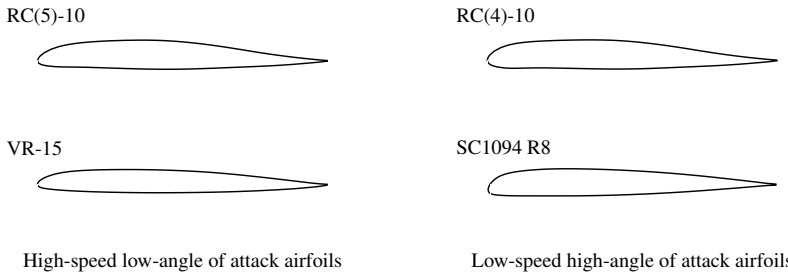

Figure 1.7: Airfoils designed for the advancing and retreating sides of helicopter rotors.

the entire blade radius could not meet the Army's manoeuvrability requirements [12]. Thus, a cambered nose section was added to the SC1095 airfoil from 0.47% to 0.85% of the blade radius to increase the maximum rotor thrust in manoeuvres. This modified SC1095 section is referred to as the SC1094 R8 and it provides higher maximum lift coefficient which clearly benefits the retreating side [1, p.372].

As a result, the design of blade airfoils appears as a trade-off problem between conflicting requirements. In addition to good stall and transonic characteristics, another important aspect in selecting airfoils for rotorcraft applications is the value of the aerodynamic moment coefficient. In fact, the pitching moment, increased either by cambered section or compressibility effects, has a negative effect on the load acting on the structure of the blade. Thus, a gain in maximum lift coefficient must always be considered in conjunction with a constraint on the moment coefficient.

In addition, in this context it is appropriate to measure the aerodynamic efficiency of the airfoil section. In Ref. [11] a general correspondence is observed between global performance parameters, such as the rotor lift-to-drag ratio and the airfoil performance parameters, such as the airfoil lift-to-drag ratio, the drag bucket boundaries and the divergence Mach number. This correspondence is pointed out also in Ref. [1, p.297], but in the case of Ref. [11] the author concludes that the relations observed between global performance and the airfoil environment are such that the emphasis in airfoil design should be for good lift-to-drag ratio, while the maximum lift coefficient performance is less important.

To sum up, the general, conflicting requirements for a good helicopter rotor airfoil are the following:

- high maximum lift coefficient
- high drag divergence Mach number

- good lift-to-drag ratio over a wide range of Mach number and angle of attack
- low pitching moment.

So far, only static characteristics of airfoils have been discussed. However, rotor blades operate in an unsteady environment which strongly affects the airfoil behavior. Thus, an assessment of the unsteady characteristics of the airfoils selected in design is of utmost importance to get reliable estimate of the rotor performance. In particular, airfoils for helicopters should be tested to assess the occurrence of dynamic stall, namely the stall encountered by an airfoil undergoing a time-varying change in the angle of attack. Dynamic stall exhibits with a hysteresis loop of the aerodynamic loads, as presented in Fig. 1.8. In short, as the angle of attack increases beyond the static stall point, the lift coefficient keeps increasing owing to a delay caused by the kinematics of the pitch rate. Eventually, flow separation develops at the leading edge, creating a vortex on the upper surface that travels downstream. As the vortex moves along the chord, lift is still generated, but a strong pitching moment is produced as well. When the vortex reaches the trailing edge region, the flow is fully separated and lift breaks down. As the angle of attack decreases and becomes sufficiently low (lower than the angle of static stall) the flow reattaches front to back.

Dynamic stall is a very complex phenomenon that has been extensively studied due to its importance [13, 14], but its accurate prediction is still a formidable challenge [15]. Simplified models of dynamic stall have been developed to implement straightforward corrections to the linear steady or unsteady models which are capable of representing the nonlinear behavior of oscillating airfoils. This simplified models used in aeromechanical comprehensive solvers for rotors are either semi-analytic methods, combining qualitative feature of the phenomenon with quantitative test data, or empirical methods based on experimental data. With regard to semi-analytic methods, great development was carried out in the 70s with the work of Gross and Harris [16], Gormont [17] and Bielawa [18]: the former two references present the Boeing-Vertol method, a correction of the unsteady thin airfoil theory of Theodorsen [19], whereas in the latter a modification of the aerodynamic loads is developed based on the steady airfoil data.

1.2.1 State-of-the art on optimization of helicopter airfoils

Helicopter airfoil design has been tackled by means of optimization methods in several works. Some of them only deal with hover optimization, others treat the hovering case as a starting point for the forward flight case. For instance, in Ref. [20]

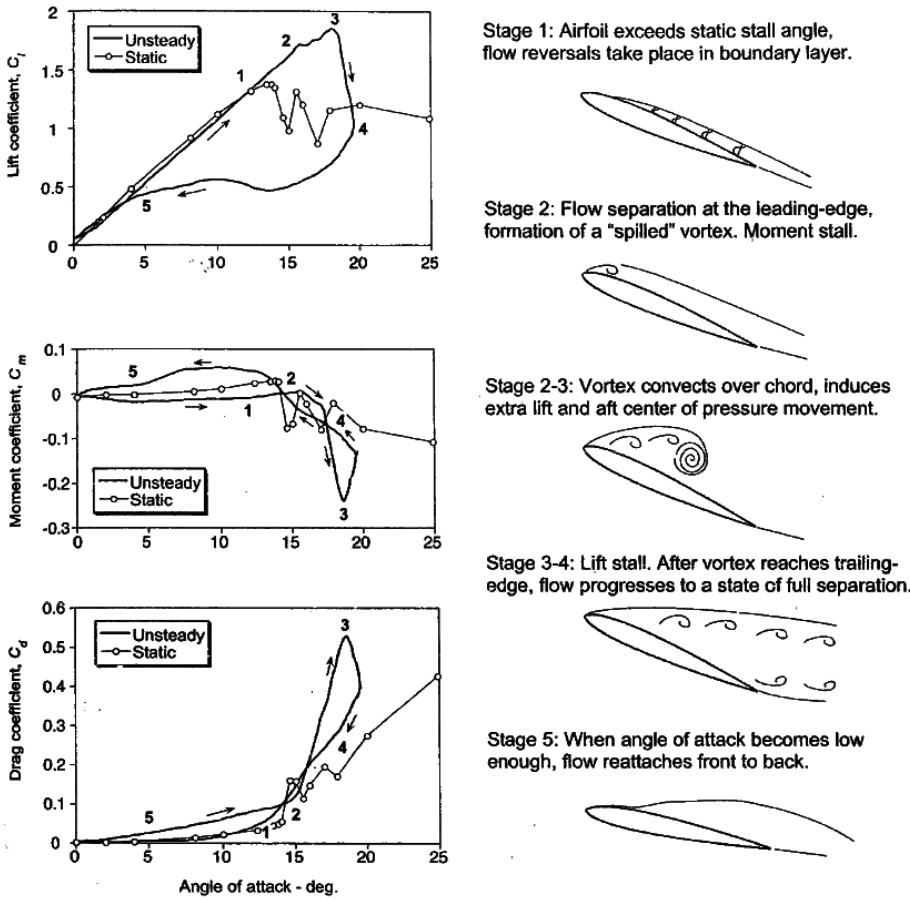

Figure 1.8: Scheme of the essential steps of dynamic stall (from Ref. [1]).

a two-dimensional rotor airfoil optimization is performed for a steady condition that is representative for hover. In Refs. [21, 22, 23] the optimization of airfoils for helicopter rotor blades is considered either in the hovering condition or with a multi-point approach for forward flight with two or three representative conditions. In these cases each condition is treated as a steady simulation.

A notable effort in introducing unsteady effects in the optimization loop is the work presented in Ref. [24] which presents an adjoint-based optimization for a pitching airfoil. This work shows that for low values of the pitching frequency a multi-point optimization with steady simulations yields similar results with respect to an unsteady computation. However, as pitching frequency increases the discrepancy between the optimal airfoil shapes grows.

Optimization of airfoil sections for the Caradonna-Tung rotor blade is performed in Ref. [25] by means of an adjoint method on unstructured mesh. References [26, 27] present the optimization of blade surface by changing the airfoil shapes. The optimization of the section shape is performed in conjunction with other blade parameters (such as twist and chord distribution) in Ref. [28]. Finally, optimizations of blade parameters (twist, chord, sweep, anhedral/dihedral) are presented in Refs. [29, 30, 31, 32, 33].

Reference [34] points out that no algorithmic "silver bullet" has emerged to make the solution of helicopter optimization problems as computationally efficient and robust as for traditional structural optimization problems. In the same reference, the need for a robustness analysis of the optimum is raised to determine the sensitivity of the optimal solution to modeling inaccuracies.

As a matter of fact, uncertainty in the modelling used within the optimization procedure no doubt affects the optimum. If the operating condition for which a blade is optimized during design do not correspond to the conditions in which the blade actually operates in the real world, the performance of the optimized blade could be unsatisfactory. Off-design performance analysis and optimization considering uncertainty may cope with these situations. Indeed, robustness concepts directly impact the design process. As pointed out in Ref. [1, p.278], in the design of a new helicopter, performance guarantees are made to the customer based on various agreed specifications. Any failure of the manufacturer to achieve the negotiated performance may result in substantial costs or loss of customers. It is then important to have a high level of confidence that the performance guarantees can be met. In this regard, robustness criteria in the design stage can reduce the likelihood of undesired behavior during flight or experimental tests at the late stage of design and the need to redesign a new machine. Few works consider uncertainty in the field of helicopter design. In Ref. [35] a study is presented of the impact on the helicopter performance of structural and aerodynamic uncertainties, such as blade chord, rotor radius, two-dimensional lift-curve slope, blade profile drag coefficient, rotor angular velocity, blade pitch angle, and blade twist rate per radius of the rotor. The uncertain variables are propagated through an aerodynamic model based on blade element and momentum theory, and a 20-25% excess power is found to be required to the helicopter due to uncertainties. Reference [36] investigates the effects of spatially uncertain material properties on the aeroelastic response predictions (e.g., rotating frequencies, vibratory loads, etc.) of composite helicopter rotor and it shows that spatial uncertainty has considerable influence on the vibratory hub load predictions.

1.3 Morphing airfoils for helicopter rotor blades

As presented in the preceding sections, airfoils of helicopter rotor blades are required to operate in different conditions and good performance in one condition does not transfer to another one. A very attractive solution to this problem is the employment of a blade capable of changing its shape during flight, that is a morphing blade.

With regard to helicopters, morphing rotor design is a challenging and active area of research [37]. In helicopter applications, the urgent need to improve the boundaries of the flight envelope leads research and industry to experiment extreme configurations and technological solutions. Examples of this are tilt-rotors (e.g. the Bell XV-15) or hybrid configurations with a rotor and one or two propellers (such as Sikorsky X2 or Airbus Helicopter X3). In addition, the internal structure of the blades is made by a single D-shaped spar close to the leading edge and honeycomb filler for the remaining part. Thus, it seems suited to accommodate for a morphing actuation mechanism.

The DARPA Mission Adaptive Rotor (MAR) is a research project on adaptive rotors. In 2009 DARPA planned to fly adaptive rotors by 2018, with the goal of developing a rotor that can change its configuration before a mission and in flight, between mission segments and with every revolution [37]. Potential approaches to reach this goal are variable blade diameter, sweep, chord, variable-camber airfoils, varying blade twist, anhedral/dihedral, tip speed, stiffness, and damping. These different strategies are expected to increase rotorcraft payload by 30% and range by 40%, while reducing the acoustic detection range by 50% and vibration by 90% over fixed geometry rotors. Some studies in the field have demonstrated the aerodynamic and acoustic benefits of active twist rotor [38, 39], rotor power reduction using trailing-edge flaps or envelope expansion using extendable chord sections [37]. Geometry variations concerned span variation by means of a telescopic approach (leveraging centrifugal forces to alleviate actuation). Probably the most developed type of morphing in the field is related to twist, obtained with aeroelastic effects produced by trailing-edge control surfaces and “smart spar”, i.e. variable stiffness spar to create aeroelastic twisting.

Airfoil morphing is mainly concerned with camber variations: to a large extent, the camber of an airfoil describes its asymmetry, and is typically used to control its zero-lift angle of attack [40]. Adding camber, for example, will tend to increase the amount of lift produced at a given angle of attack of the airfoil, although this is limited by stall and separation. Almost all modern aircraft use discrete control surfaces, such as flaps, ailerons, or sometimes slats, to adjust the camber of the wing. Trailing

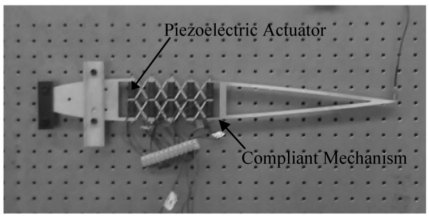

(a) Prototype of conformable camber airfoil with piezoelectric actuators developed in Ref. [41].

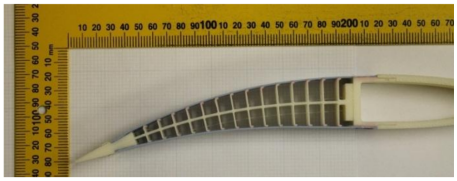

(b) FishBAC structure [45].

Figure 1.9: Technologies for variable camber morphing.

edge devices are control surfaces typically hinged at 75% of the chord which rotate to change their angle and also translate in some cases to increase chord as well as camber. In helicopter rotor blades, trailing edge flaps have been extensively considered for vibration reduction over the last decades [41]. For instance, flaps are considered in the Boeing Smart Materials Actuated Rotor Technology (SMART) rotor project [42] focused on the development of controllable rotors with smart-material actuators [43, 44]. However, the camber change is almost always discrete in that after actuation of the control surface, there is a sudden transition of camber in the chordwise direction. This causes a sudden change in the pressure distribution over the corner created at the hinge line, causing a drag penalty and the possibility of separation. As a result, a smooth variation of camber obtained by airfoil morphing is an attractive solution to increase the lift coefficient in particular conditions of flight.

Variable camber airfoils have been adopted in fixed-wing applications [46, 47] and this solution has recently been explored for helicopter applications, as well. To this extent, a continuously changing rotor blade airfoil is seen as a successor to rotor blade trailing-edge flaps [41]. In Ref. [48] an adaptive airfoil with an aerodynamically smooth variable leading-edge droop is used to postpone dynamic stall. A variable-camber airfoil based on a compliant mechanism and piezo-electric actuators is developed in Ref. [41] (see Fig.1.9a). The morphing device is used for active helicopter vibration reduction. In this case the shape of the airfoil is fixed, but an optimization is performed to find the optimal material distribution to maximize the trailing-edge deflection under actuation forces and minimize the deflection under external aerodynamic loads. Another important study in the field is the development of the FishBone Active Camber (FishBAC) concept introduced and developed in Refs. [49, 45, 36]. The FishBAC is made of a compliant spine and stringer skele-

ton, pre-tensioned elastomeric matrix composite skin, an antagonistic tendon drive linked to a non-backdrivable spooling pulley, and a non-morphing,"rigid", main spar (see Fig. 1.9b).

Another solution in the literature is proposed in Ref. [50]. In this case a variable droop leading edge is used to improve the sectional lift-to-drag ratio by alleviating dynamic stall on the retreating blade while simultaneously reducing the transonic drag rise of the advancing blade. A leading edge variation combined with a trailing edge camber variation is the solution presented in Ref. [51] to define shock-free airfoils for helicopter rotors.

Because morphing structures typically imply increased weight, complexity and cost, an in-depth assessment of the potential aerodynamic benefit of such an approach is strongly recommended [45]. As a final remark, optimization of the morphing airfoil shape to improve the aerodynamic performance is typically run separately to the morphing mechanism design [40]. To this extent, the optimization approach is conceived as a two-level approach [46]: (i) the best deformed airfoil configuration is determined according to a desired aerodynamic performance, and (ii) in the second level, the best internal structural configuration is obtained using a topology optimization suited to the chosen compliant structure able to adapt itself in order to match the optimal shape coming out from the first level.

1.4 Objective of the PhD thesis

The thesis addresses the problem of designing an airfoil of a helicopter blade section with novel approaches. In the first instance, a problem of design under uncertainty is formulated to cope with the uncertainty affecting the operating conditions. Secondly, a morphing airfoil is considered to cope with the different conditions encountered in forward flight. The possibility of a morphing strategy is also explored from a robust perspective.

To clarify the application, a template of the aerodynamic design optimization problems tackled throughout the thesis is depicted in Fig. 1.10. A "classical" optimization problem requires the definition of an aerodynamic system, composed of the following elements: (i) a model of the system, (ii) the design input, i.e. the airfoil shape, and (iii) the output of interest, which is typically a function of the flowfield (e.g. lift coefficient, drag coefficient, pressure in the separation region). The aerodynamic model is in general a partial differential equation of the type $\mathcal{L} = 0$, where $\mathcal{L}$ is the proper differential operator. The model describes the flowfield around the airfoil, starting from suited initial and boundary conditions, such as unsteady pitching

angle or freestream operating conditions. The model computes the performance, i.e. the aerodynamic loads for a given airfoil shape. However, many of the variables that need to be specified for this design problem are affected by sources of uncertainty.

With regard to operating conditions, sources of uncertainty may arise from environmental perturbations or lack of knowledge in the modelling assumptions. Specifically, the flow conditions encountered by rotor blades may vary due to the effects of the blade dynamics and flexibility, as well as the interaction with vortices and wakes trailed from preceding blades. It follows that the flow conditions considered at the design stage for the definition of the optimal two-dimensional section are usually affected by uncertainty, not only because they are estimates of the actual condition, but also because modelling involves some level of approximation of the real flowfield. These sources of uncertainty impact the local angle of attack and incident velocity at which the airfoil of the blade operates. In the application to morphing structures, a high level of uncertainty affects the modelling of the environmental conditions encountered by the airfoil, and therefore robustness becomes critical to define reliable design.

In such a context, taking into account the variability of the modelling parameters in the optimization procedure is a possible way to result in a more reliable optimal design. This is the goal of uncertainty-based optimization. In particular, robust optimization seeks an optimal design that is minimally sensitive with respect to changes of the uncertain parameters that define the problem. The result is a robust design that maintains satisfactory performance also in off-design conditions.

In such a context, several are the objectives of the PhD thesis. First, a methodology for robust optimization is developed to tackle the problem of designing fixed and

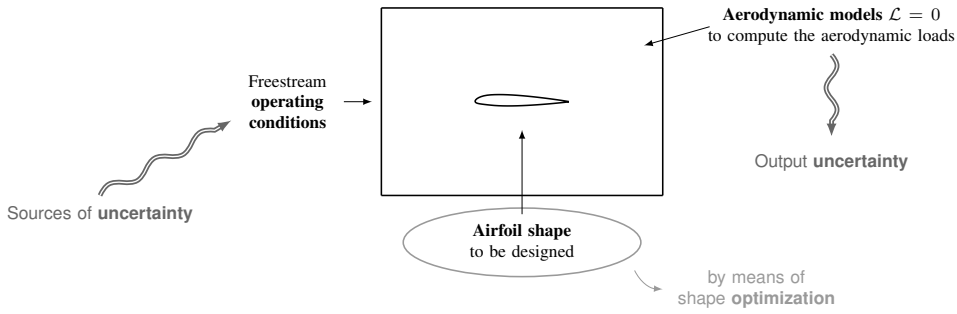

Figure 1.10: Representative scheme of the optimization problem of the airfoil shape under uncertainty.

morphing airfoils of helicopter rotor blades. To this extent, the employment of robust approaches is explored to demonstrate their capability of indicating interesting design in the application to helicopter rotor blade airfoils. In addition, the impact of robustness criteria on airfoil design is discussed. Because robust optimization increases the computational cost of the optimization process, new uncertainty-based optimization methods are developed. The goal of this development is to improve the numerical efficiency, thereby making robust approaches more attractive to aerodynamic applications. Finally, the gain of the morphing airfoil strategy is assessed not only from a deterministic point of view, but also when considering uncertainty in the operating conditions. In such a way, a more realistic estimate of the gain of morphing strategies can be defined.

1.4.1 Structure of the PhD thesis

The thesis is outlined as follows. Chapter 2 presents the problem of design under uncertainty, with particular attention to robust design optimization. The methodology of robust optimization is presented and the numerical ingredients necessary to perform an optimization under uncertainty are detailed. In particular, multi-objective optimization algorithms and non-intrusive uncertainty quantification methods are presented.

A major problem in design under uncertainty is the numerical efficiency, because considering uncertainty results in an increase in computational cost. Thus, in Chapter 3 two novel strategies are developed to improve the efficiency of robust design methods. The methods are applied to a bi-objective algebraic optimization in the first instance to assess the impact on the global robust optimization cost.

Chapter 4 is devoted to the optimization of airfoils for the hovering flight condition. The methods presented in Chapter 3 are applied to the hovering optimization problem to assess the performance of the methods in an engineering case. Results obtained with the methods are also discussed from a physical point of view to analyze the aerodynamic performance of the optimal shapes found with the robust approach.

In Chapter 5 the forward flight condition is considered. In this case, two different classes of optimization are considered. First, a “classical”, deterministic airfoil optimization in forward flight is performed, with a minimization of the drag coefficient subject to a target lift coefficient for trimming purposes. The second type is the robust optimization in forward flight condition, namely an optimization that seeks a shape with a lift-to-drag ratio robust against uncertain operating conditions. The results of the two optimization problems are compared and their performance over the azimuth

angle is assessed. Because steady models are used throughout the thesis, a comparison between quasi-steady loads and unsteady loads is presented for the reference airfoils.

Finally, in Chapter 6 the aerodynamic optimization of morphing airfoils for forward flight application is considered. The morphing strategy consists of a conformable camber airfoil which changes over the period of rotation of the blade to cope with the variable flow conditions encountered in forward flight. A robust optimization is considered also in this case to cope with uncertainty in the operating conditions.

CHAPTER 2

Robust optimization

In most engineering applications, progress in science and technology increased demand for design with better performance, higher reliability and robustness, and lower cost and risk [52]. To effectively address these competing objectives, designers generally employ design and optimization methods considering all relevant aspects of the system operating cycle, in which there inherently exist numerous sources of uncertainties. For instance, system operating conditions are affected by the status of the surrounding environment which may change during the system life cycle. Moreover, some model parameters used to represent the physical system during design stage are affected by a lack of knowledge of the system itself. Figure 2.1 presents possible sources of uncertainty that may arise in aircraft design, production and service. It is clear that these uncertainties may alter the system's expected performance, jeopardize the safety, and even result in mission failure. Thus, it is beneficial to take into account this aspect in the optimization tools used during design.

Two main concepts are related to design in the presence of uncertainty: robustness and reliability. In the former case, design solutions are sought with improved performance and reduced sensitivity to variations. In the latter case, the optimal design guarantees that the occurrence of undesired events is avoided or kept under a given threshold. A comparison between the optimal solution of a traditional optimization and the solutions obtained considering robustness and reliability criteria is presented in Fig. 2.2. In the first case, the robust minimum guarantees that in the event of a perturbation Δx of the design variable x_1, the perturbation of the objective

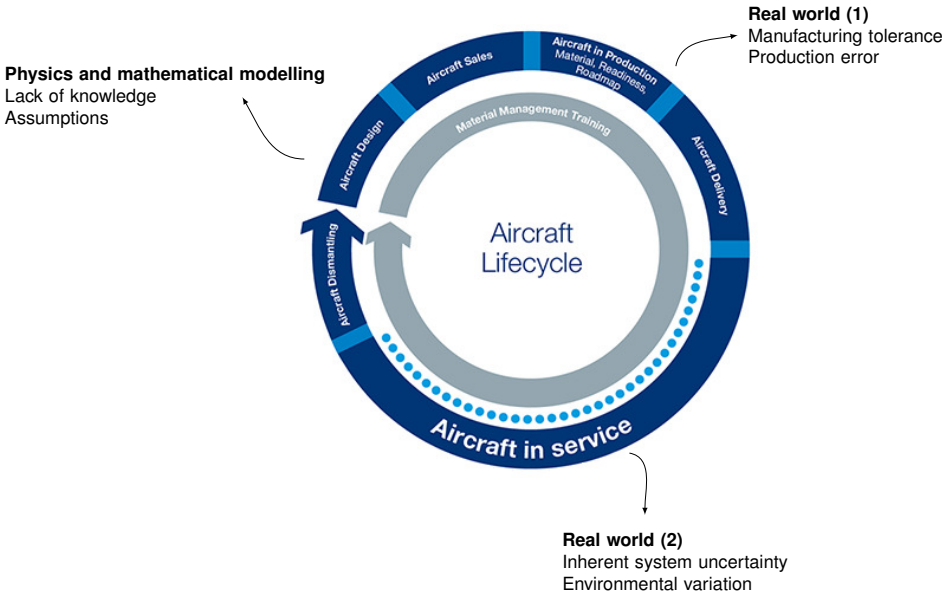

Figure 2.1: Sources of uncertainty in aircraft design, production and service (picture of aircraft life cycle from *www.airbus.com*).

function Δu is minimal. With respect to the traditional minimum (called deterministic in Fig. 2.2a), the robust minimum presents a small performance penalty, but ensures smaller perturbation of the performance. In Fig. 2.2b a reliable minimum is presented, which is the point with the lower value of the objective function that is also distant a given safety factor from the unfeasible region imposed by the constraint.

Reliability and robustness are of great interest in aircraft design. Robustness criteria drive aerodynamic designers to favor airfoils that exhibit trailing edge stall, rather than leading edge stall airfoils. In fact, in the former case airfoils present a smoother transition to full flow separation and stall condition, owing to a larger leading edge radius and strong upper surface curvature. In the latter case, airfoils have moderate leading edge radii and a flat upper surface, which lead in most cases to abrupt stall [53]. With regard to reliability-based design, reliability concepts are of primary interest in structural and aeroelastic aircraft design to prevent wing structure damage [54] or wing flutter onset [55].

Traditional strategies used in engineering design to ensure robustness and reliability are local sensitivity analysis and safety margin factors. Both strategies are valid and widely employed, but they have some limitations. In the first case, classical sen-

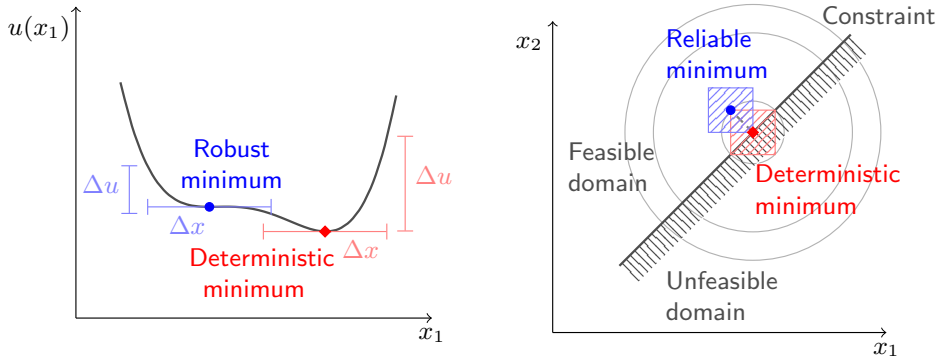

(a) Traditional (deterministic) and robust optimal solutions for the minimization of the performance u with respect to the design variable x_1.

(b) Traditional (deterministic) and reliable optimal solutions in the design space x_1-x_2. Circles indicate isolines of the performance.

Figure 2.2: Scheme of optimal solutions under uncertainty.

sitivity analysis is local and it does not account for the actual range of variations of the uncertain variables which may dramatically change the system response. Safety factors may yield too conservative design solutions, which in aerospace structures sizing translates into increased weight and cost. In both cases, a complete representation of the behavior under the estimated uncertainties is not provided. This has motivated the development of more advanced and accurate analytical approaches to tackle uncertainties systematically and rationally during design. These methods are ascribed to *uncertainty-based design optimization* or *design under uncertainty* and they are based on the coupling of an optimization method and an Uncertainty Quantification (UQ) method used to quantify the uncertainty of the system performance due to the uncertainty of the system.

Although the application of uncertainty-based analysis to aerospace vehicle design has been studied over the past two decades, a concerted effort is still needed to fully convey existing uncertainty-based design optimization methods from the research community to industry. One major challenge is the improvement of the efficiency of UQ methods and of the coupling between UQ and optimization methods to solve uncertainty-based optimization.

In this chapter, the concept of robust design is described in Sec. 2.1 using a general framework for design under uncertainty. Section 2.2 defines key elements of multi-objective optimization and presents the optimization algorithms used through-

out the thesis. Finally, the uncertainty modelling and uncertainty propagation methods are described in Sec. 2.3.

2.1 What is robust optimization?

The process of obtaining an optimal design according to a set of objectives and constraints is called design optimization. In most engineering applications, the objective of the optimization is a cost function, or quantity of interest, that defines the performance of the system to be improved. Constraints are necessary in design optimization to enforce compliance of geometrical specifications or avoidance of catastrophic events, for example.

Both objectives and constraints are typically computed by a mathematical model of the system under analysis. For this general description, let us consider a model for a time- and space-dependent system, which could be represented by the following equation

$$\mathcal{L}(\mathbf{x}, t, \mathbf{y}; \phi(\mathbf{x}, t, \mathbf{y})) = 0, \tag{2.1}$$

with appropriate initial and boundary conditions. To give a rigorous mathematical description, the operator $\mathcal{L}$ is defined on the domain $\mathcal{D} \times \mathcal{T} \times \mathcal{X}$, where $\mathbf{y} \in \mathcal{D}$ and $t \in \mathcal{T}$ are the spatial and temporal coordinates with $\mathcal{D} \subset \mathbb{R}^d$, $d \in \{1, 2, 3\}$, and $\mathcal{T} \subset \mathbb{R}$. The vector of the design variables $\boldsymbol{x}$ is defined on the domain $\mathcal{X} \subset \mathbb{R}^n$, where n is the number of design variables, and the solution of Eq. (2.1) is $\phi(\mathbf{x}, t, \mathbf{y})$.

In traditional design optimization, engineers assume that the variables and models used in the optimization problem are known with certainty. In this case, the optimization is referred to as *deterministic design optimization* and it may be formulated in mathematical terms as follows

$$\begin{aligned} \text{minimize/maximize:} \quad & u\left(\boldsymbol{x}\right) \\ \text{subject to:} \quad & \mathbf{g}\left(\boldsymbol{x}\right) \leq \mathbf{0} \\ \text{by changing:} \quad & \boldsymbol{x} \end{aligned} \tag{2.2}$$

where $\boldsymbol{x} \subset \mathbb{R}^n$ is the vector of n design variables, u is the performance, i.e. the objective function and vector $\mathbf{g} \subset \mathbb{R}^q$ describes the q constraints of the optimization problem. The dependence of the objective u and constraints $\mathbf{g}$ on vector $\boldsymbol{x}$ may be a direct dependence (e.g. a geometrical constraint), or it may stem from the dependence on the solution ϕ of the equation describing the system under consideration (Eq. (2.1)). Please note that the arguments $\mathbf{y}$ and t are dropped in the preceding

equation to highlight only the dependency on the design and uncertain variables. In addition, vector $\mathbf{g}$ typically contains also Eq. (2.1), which represents a constraint to the optimization problem [56].

Now, let us assume that the parameters or the model used in the optimization problem are affected by uncertainty. In this case, equation (2.1) is written as the following

$$\mathcal{L}(\mathbf{x}, t, \mathbf{y}, \boldsymbol{\xi}; \phi(\mathbf{x}, t, \mathbf{y}, \boldsymbol{\xi})) = 0. \tag{2.3}$$

The new element in this equation is vector $\boldsymbol{\xi}$ collecting the uncertain variables. Uncertain variables are not deterministic, or in other words thay cannot be assigned a single, deterministic value. Instead, they are typically described by means of a distribution describing the likelihood of the uncertain variable being equal to a specific value. Different mathematical theories are used to model uncertainties, from probability theory to interval analysis [57]. A discussion on these theories is presented in Sec. 2.3. Here, it is important to note that the solution ϕ of Eq. (2.3) and consequently the quantity of interest u depend on the uncertain variables $\boldsymbol{\xi}$; in particular, $u = u(\phi(\boldsymbol{x}, t, \mathbf{y}, \boldsymbol{\xi}))$. In other words, the performance of the system is now a function of a non-deterministic variable and it is a non-deterministic variable itself.

When uncertainty is considered, the goal of the optimization problem is typically to control the variability of the performance u with respect to the uncertainty in the system variables $\boldsymbol{\xi}$. This optimization problem is called *uncertainty-based optimization* or *optimization under uncertainty*. In general, a problem of uncertainty-based optimization is stated in mathematical terms as follows

$$\begin{aligned} \text{minimize/maximize:} \quad & \mathbf{f}\left(u\left(\boldsymbol{x}, \boldsymbol{\xi}\right)\right) \\ \text{subject to:} \quad & \mathbf{g}\left(u\left(\boldsymbol{x}, \boldsymbol{\xi}\right)\right) \geq \mathbf{0} \\ \text{by changing:} \quad & \boldsymbol{x} \in \mathcal{X}, \end{aligned} \tag{2.4}$$

where the m objective functions are collected in vector $\mathbf{f} \subset \mathbb{R}^m$, and $\mathbf{g} \subset \mathbb{R}^q$ are the q constraints of the optimization problem. The objective functions are deterministic measures of the variation of the performance due to the uncertain variables. For example, if probability theory is used to describe the uncertainty involved in the robust aerodynamic optimization of a transonic wing airfoil, the mean value and the variance of the drag coefficient are taken as objective functions [58]. Please note that the optimization problem stated in Eq. (2.4) applies to a generic problem of optimization under uncertainty. Depending on the objective functions chosen in Eq. (2.4) different optimization problems are formulated. For instance, Appendix B presents a formu-

lation for design under uncertainty based on the minimum and the mean value, that was explored in the thesis.

Specifically, **robust design optimization** [59] is the process of finding a design which provides an improved and satisfactory performance and at the same time minimizes its sensitivity due to uncertainty. In this way, it is guaranteed that the system will operate adequately even in off-design conditions. When considering robust design optimization, the term "robustness" has been used with various meanings depending on the desired objective [60]. It could indicate: (i) a design that is minimally sensitive to a variability either in the system or in the design itself, (ii) a design that uniformly improves the performance of the system over a set of operating conditions, (iii) a design with an improved worst-case performance, and (iv) a design that improves the overall performance over the entire range of operating conditions and/or lifetime. In this work, robustness means that the design parameters $\boldsymbol{x}$ allow simultaneous improvement in the overall performance and minimization of the sensitivity with respect to changes in the system operative conditions. Thus, in the case of robust design within a probabilistic framework, the objectives of the optimization are two low-order statistical moments of the performance of interest u, i.e. the mean value μ and the variance σ^2, and the optimization problem in Eq. (2.4) can be reformulated as a multi-objective, constrained optimization:

$$\begin{aligned} \text{minimize/maximize:} \quad & \mu\left(u\left(\boldsymbol{x},\boldsymbol{\xi}\right)\right) \\ \text{and minimize:} \quad & \sigma^2\left(u\left(\boldsymbol{x},\boldsymbol{\xi}\right)\right) \\ \text{subject to:} \quad & \mathbf{g}\left(u\left(\boldsymbol{x},\boldsymbol{\xi}\right)\right) \geq \mathbf{0} \\ \text{by changing:} \quad & \boldsymbol{x} \in \mathcal{X}, \end{aligned} \tag{2.5}$$

Unlike single-objective optimization where there is only one global optimum, in the case of multi-objective optimization there is a set of optimal solutions (the so-called Pareto front) corresponding to various trade-offs between the mean value and the variance.

Robust design optimization is inextricably linked with the name of Taguchi [61] who initiated the robust design philosophy in manufacturing quality control. Taguchi's idea was to minimize a signal-to-noise ratio defined as the variance of the discrepancy between the measured quality and the desired quality. In his pioneering work, Taguchi did not use a direct optimization procedure in order to minimize the signal-to-noise ratio. Instead, Taguchi's method was based on a methodology similar to a design of experiments technique in order to evaluate different design vectors for different realizations of the noisy variables [62]. Thanks to the advent of high-speed

computers and its exponentially progress, robust design criteria have been included inside optimization loops and gained interest in several industrial fields.

Robust design can be intrinsically more fit than classical deterministic approaches for very preliminary design phases, since its objective can be interpreted as the search for optimal solutions that have a low sensitivity to variations of the parameters in the range of uncertainty. In turn this may reduce the probability of redesign at later stage of the design process due to lack of adequate performance. Indeed, Taguchi's ground-braking idea was to control the quality of the product (i.e. the compliance to given specifications) before design, and not only by inspection after the production stage. For this reason, robust approaches have been studied in the field of aerospace vehicles with application to conceptual, structural and aerodynamic design [52].

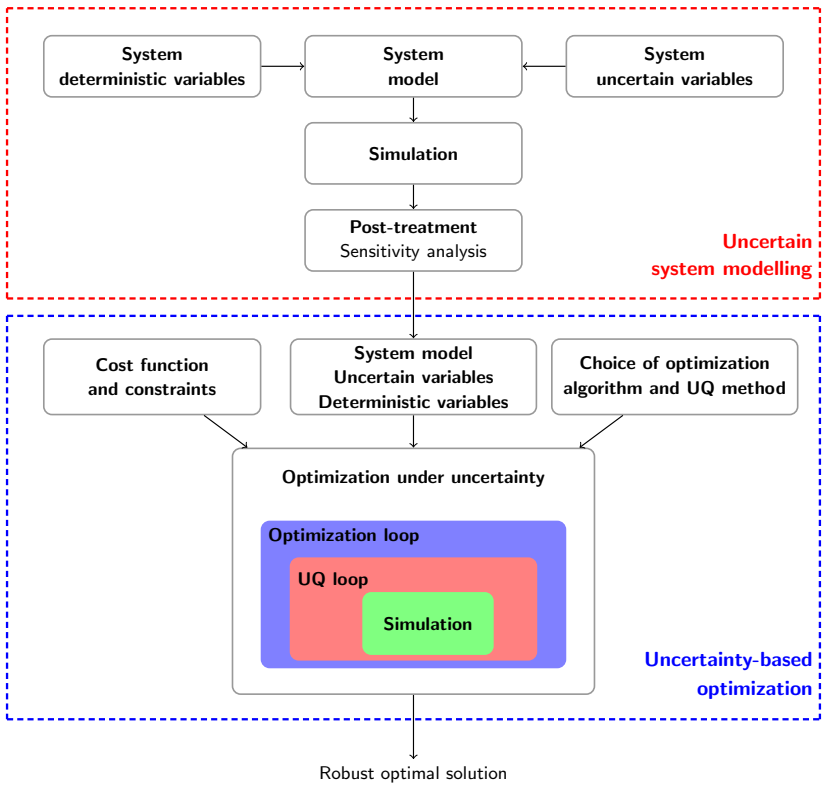

Figure 2.3: Scheme of uncertainty-based optimization.

2.1.1 Methodology for uncertainty-based optimization

A generic uncertainty-based optimization problem requires several ingredients that should be carefully analysed and determined. With reference to Fig. 2.3, the first step requires the modelling of the uncertain system, by defining the following elements.

- *System model*: a mathematical model such as the one presented in Eq. (2.1) is required to compute the objective functions and/or the contraints, and to propagate the uncertainty onto the performance u.

- *Deterministic variables*: among the variables used in the optimization problem it is necessary to distinguish between certain, deterministic variables and uncertain variables. The first set encompasses all the variables either in the system modelling or in the optimization problem that can be assigned with certainty.

- *Uncertain variables*: these are the variables that are affected by a source of uncertainty, such as operating or environmental conditions, modelling coefficients of a state law, or design variables themselves. The uncertain variables must be identified and characterized according to the nature of the problem. In practical design problems, a very large number of variables may be affected by uncertainty, because these variables stem from assumptions and semplifications of the real-world system. However, increasing the number of uncertain variables affect the computational cost of the uncertainty-based optimization procedure. Using a large number of uncertain variables may require an excessive computational effort. Therefore, it is useful to perform a sensitivity analysis, such as the one presented in Ref. [63], on a reference or initial design, in order to identify the uncertain variables whose effect on the performance is most significant. The uncertain variables whose effect can be disregarded may then be treated as deterministic variables for the optimization problem. Nevertheless, it is noted that a sensitivity analysis applied on a reference design may not be sufficient, because non-influential variables for this reference design could turn out to produce significant effects for other designs encountered when exploring the design space. With regard to the characterization of the uncertain variables, the issue of modelling uncertainty is tackled in detail in Sec. 2.3.1.

- *Cost function and constraints*: an optimization problem always requires a set of objective and constraints in order to discriminate between a good and feasible solution and a poor and unfeasible design. The cost and the constraints

functions need to be formulated to take into account the variability introduced by the uncertain variables.

In addition to these elements, the following numerical methods must be defined to proceed with the uncertainty-based optimization.

- *Simulation*: a numerical implementation of the system model is necessary to solve the equations describing the system and compute the performance. For instance, in aerodynamic design a Computational Fluid Dynamics (CFD) solver is used to simulate the flowfield past a wing or an airfoil and to compute the aerodynamic loads.

- *Optimization algorithm*: an optimization method is chosen according to the optimization problem formulated in the first steps. In robust design, optimization problems are typically bi-objective and constrained. An optimization algorithm is typically an iterative method performing a loop: at each iteration a new design is computed according to specific operators for the chosen optimizer and the objective function for the new design is evaluated. In traditional deterministic optimization, the evaluation consists of a single simulation of the system model, providing the performance. Instead, in uncertainty-based optimization, the objective functions are computed by means of an uncertainty quantification method

- *Uncertainty quantification method*: these methods propagate the uncertainty affecting the non-deterministic variables into the performance of the system. To compute the objectives of the current design in the optimization loop, uncertainty quantification methods require either a group of several simulations or the solution of an augmented version of the system model. More details on this are given in Sec 2.3.2.

Once the elements previously described are defined, it is necessary to build a framework to solve the problem presented in Eq. (2.4). A method for optimization under uncertainty consists of two nested loops: the optimization loop and the UQ loop. A representation of the coupling of an optimization algorithm and a UQ method is given in Fig. 2.4. The outer loop is performed by the optimization algorithm which generates a new design $\mathbf{x}_i$ based on the history of the optimization procedure. The new design is passed on to the inner loop in which a UQ method computes an approximation of the statistics $\mathbf{f}(\mathbf{x}_i)$ of the performance of interest. This computation is performed by resorting on the simulation of the system. The resulting statistics are

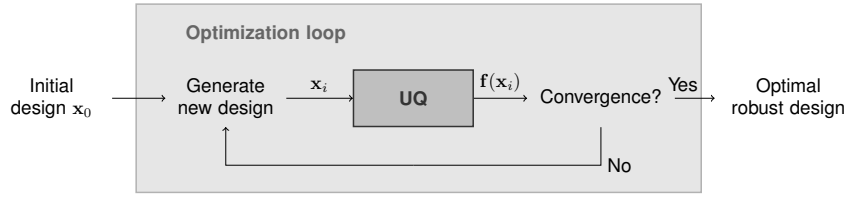

Figure 2.4: Block diagram of the uncertainty-based optimization loop.

the objective functions of the optimization loop, which are passed on to the optimization algorithm, which in turn checks if convergence is reached. If the convergence criterion is not met, the loop goes on with a new design.

Classical approaches for robust design optimization, called *decoupled* approaches, are based on two different and independent methods for solving the optimization problem and the reconstruction of the statistics, respectively. Decoupled approaches are widely employed in the uncertainty-based optimization literature [64, 65] and they are very expensive. In fact, performing a deterministic optimization generally requires several computations of the performance, since many design vectors have to be tested in the loop. This means that many simulations of the equations describing the system under consideration (Eq. (2.1)) are required. In the case of uncertainty-based optimization, the cost is even higher. In fact, to reconstruct the variability of the performance with respect to the uncertain variables in the system two strategies can be used: (i) several simulations of the system model are computed which correspond to different realizations of the input uncertain parameters, (ii) an augmented version of the system equation is considered to reproduce the dependance of the solution on the uncertainties. The classical choice in uncertainty-based optimization is the first strategy, because the second solution is intrusive, requiring an intensive modification of the system equation. As a result, the cost of solving the equations and therefore the optimization problem can be very high in most engineering applications. A major part of the thesis work has attempted to address this issue, and the methods developed to improve the efficiency are presented in the next chapter.

In the remaining part of the section, multi-objective optimization concepts and algorithms are described and uncertainty modelling and UQ methods are presented to better understand the elements involved in uncertainty-based optimization and appreciate the causes of the increase in cost.

2.2 Methods for optimization

A key element in optimization is the optimization algorithm. Because robust optimization is inherently multi-objective [62], optimization methods for multi-objective problems are considered in this work. Multi-objective optimization permits to optimize several conflicting criteria without resorting on scalarization, i.e. weighting several objectives, which inevitably affects the optimization results. It is a very important tool in the early stage of design to explore the design space following many directions of improvement. It also provides the decision makers with a set of optimal solutions, from which the most interesting ones can be selected and developed throughout the design process.

Furthermore, the attention is focused on problems in which the functional dependency of the objective function cannot be exploited, either because it is absent or it is unreliable. In this kind of problems, derivative-free optimization methods can be used to solve the optimization problem. These methods are very versatile, using the model of the system as a "black box". This is advantageous because it does not require an intrusive modification of the model. For instance, in the application to fluid dynamics, it is not necessary to modify the CFD solver. It also permits to easily change the objective function of the optimization. On the other hand, they may be very expensive from a computational point of view, because of their rate of convergence. Examples of such approaches are the Nelder-Mead method and its modified versions, the pattern search methods, or genetic algorithms. In this work, the Non-dominated Sorting Genetic Algorithm and the Bi-objective Mesh Adaptive Direct Search (BIMADS) algorithm are chosen, because they are suited to tackle the multi-objective case. A brief description of these methods is given after a brief overview of some key concepts of multi-objective optimization.

2.2.1 Multi-objective optimization and Pareto dominance

In many applications, conflicting criteria drive the design process. In this case the optimization problem has a multi-objective formulation. Indeed, a general uncertainty-based optimization problem (Eq. (2.4)) is inherently multi-objective: the mean value is used in conjunction with the variance in robust design, or with the probability of the occurrence of failure events in the application to reliability-based design.

The solutions of multi-objective problems ($m > 1$) is a set of solutions (unless the objectives are concurrent). These solutions represent a trade-off between the objectives considered in the optimization and they are selected according to the Pareto

dominance relation.

Definition 1. *For two vectors* $\mathbf{a}, \mathbf{b} \in \mathcal{X}$ *with* $\mathcal{X}$ *the design variables domain*

$$\mathbf{a} \succ \mathbf{b} \text{ (}\mathbf{a} \text{ dominates } \mathbf{b}\text{) iff } \beta_j f_j(\mathbf{a}) \leq \beta_j f_j(\mathbf{b}) \quad \forall j = 1, ..., m \text{ and}$$
$$\exists j \text{ such that } \beta_j f_j(\mathbf{a}) < \beta_j f_j(\mathbf{b})$$

$$\mathbf{a} \sim \mathbf{b} \text{ (}\mathbf{a} \text{ is indifferent to } \mathbf{b}\text{) iff } \mathbf{a} \text{ does not dominate } \mathbf{b} \text{ and } \mathbf{b} \text{ does not dominate } \mathbf{a}$$

Coefficient β_j *indicates the goal of the optimization of function* f_j*:*

$$\beta_j = \begin{cases} 1 & \text{if } f_j \text{ is to be minimized} \\ -1 & \text{if } f_j \text{ is to be maximized} \end{cases}$$

Definition 2. *A decision vector* $\boldsymbol{x} \in \mathcal{X}$ *is said to be* non-dominated *regarding a set* $\mathcal{A} \subseteq \mathcal{X}$ *if and only if*

$$\nexists \mathbf{a} \in \mathcal{A} : \mathbf{a} \succ \mathbf{x}$$

Moreover, $\boldsymbol{x}$ *is said to be* Pareto-optimal *iff* $\boldsymbol{x}$ *is non-dominated regarding* $\mathcal{X}$*.*

The solution vectors that are non-dominated regarding $\mathcal{X}$ are the Pareto optimal solutions, and they belong to the set denoted as $\mathcal{P}$. The image of the set $\mathcal{P}$ is called Pareto front and it defined as $\mathcal{T}_{\mathcal{P}} = \{\mathbf{f}(\boldsymbol{x}) \mid \boldsymbol{x} \in \mathcal{P}\}$.

The Pareto dominance is used in multi-objective optimization problems to rank the design vectors and to define the optimal solutions. Please note that single-objective optimization ($m = 1$) can also be treated with this criterion. In this case, the third option in Definition 1 cannot be satisfied, since only one function is considered: either vector $\mathbf{a}$ dominates vector $\mathbf{b}$ or viceversa.

2.2.2 Genetic Algorithm

Evolutionary algorithms have been employed to tackle many engineering optimization problems. Applications to aerodynamic optimization problems can be found for instance in Refs. [65, 25, 66]. Main advantages of such approaches are the possibility to tackle multi-objective problems without scalarization and to exhaustively explore the design space. On the other hand, they are very expensive from a computational point of view because they require many computations of the objective function, even in region of the domain where designs with poorer performance are present.

Genetic algorithms use operators inspired by natural evolution to get to the optimal solution. In such a framework, each design vector is considered as an individual,

and the design variables in the vector represent the chromosomes. A group of individuals is called a generation. Each individual is assigned a fitness function, that is a measure of the probability that this individual will be a parent of the individuals of the next generation. The fitness in the optimization algorithms is associated to the objective function: the better the objective function, the higher the fitness value. In the single objective case, the fitness function is a measure of the objective function. In the multi-objective case, the fitness function is assigned based on the position of the individual in the objective space with respect to the current Pareto front. With the fitness value, the selection step is performed to determine the parents of the new generation; mutation and crossover operators are then used to mantain genetic diversity and to obtain different children from a single set of parents.

The genetic algorithm performs a loop over the generations (see Fig. 2.5a). The starting point is an initial generation of individuals that are typically randomly chosen. For each iteration, a new generation of individuals is computed using selection, crossover and mutation operators. Then, the objective functions are evaluated for each individual, and each design vector is assigned a fitness function. The evaluation of the individual in the case of design under uncertainty require the computation of the statistics by means of a UQ method. From the fitness of these individuals, a new generation is computed and the loop goes on for a sufficient number of generations.

In this work, the Non-Dominated Sorting Genetic Algorithm is used [67]. Typical values for the crossover and mutation probabilities $pc = 0.9$, $pm = 0.1$ are chosen and the sharing parameter is set using a formula based on the number of design variables.

2.2.3 BiMads

The Bi-objective Mesh Adaptive Direct Search (BiMads) algorithm [68] is the extension of the Mads (Mesh Adaptive Direct Search) algorithm, to bi-objective optimization problems [69].

The Mads algorithm is a direct search algorithm for single-objective optimization. The underlying idea is to search an improved solution by evaluating the objective function for all the points of a mesh defined around the current design vector in the design variables space. The procedure consists on a two-step iteration based on a "search" and a "poll" step. The search step evaluates the objective function in the mesh points around the current design vector; if a better design vector is not found, a poll step is performed, which consists in a local exploration of the design variables space. If an improved solution is found, by either the search or the poll step, this

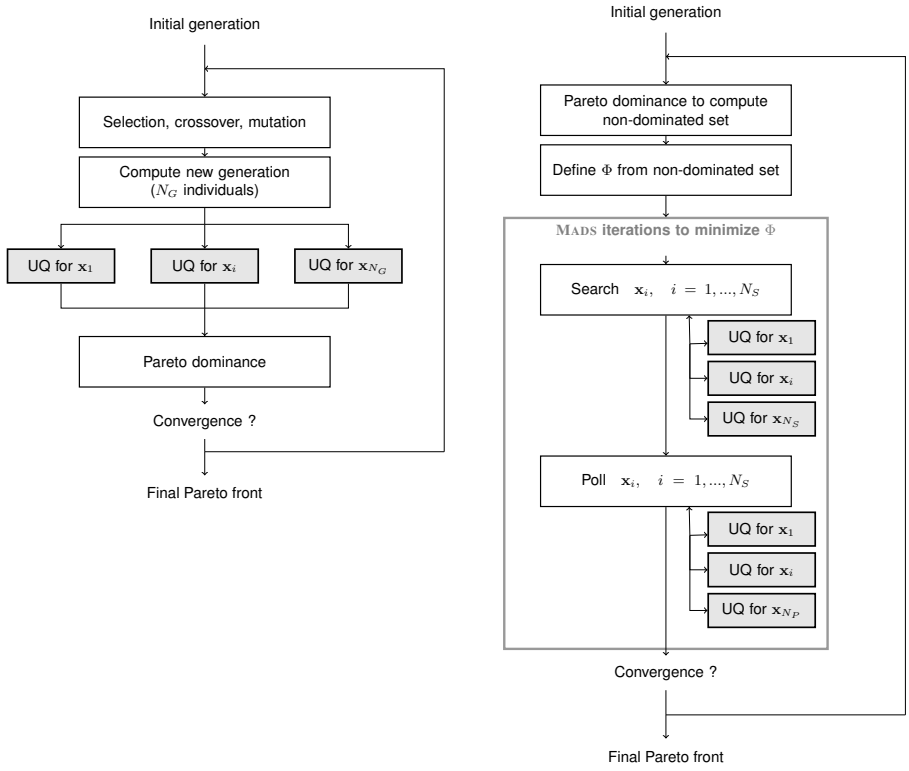

(a) Genetic algorithm. **(b)** BIMADS algorithm.

Figure 2.5: Scheme of optimization algorithms used in the work.

one is set as the new design vector around which a new search mesh is built. Otherwise, the mesh is refined and the iteration is repeated. The optimization loop stops either when a fixed number of function evaluations is reached or when the mesh size reaches a given tolerance.

For bi-objective problems, an approximation of the Pareto front is constructed using the BIMADS method, which launches successive runs of MADS. In the initialization of the method, two separate single-objective optimization problems (one for each objective) are performed by means of MADS. From the evaluations of the objective functions computed in the initialization, an initial list of non-dominated solutions is defined. The following loop is a repetition of three steps (see Fig. 2.5b). The first is the computation of a reference point in the objective space from the list of non-dominated solutions. This reference point is used to define a scalarization function Φ of the objective functions (for instance, the function is a measure of the distance

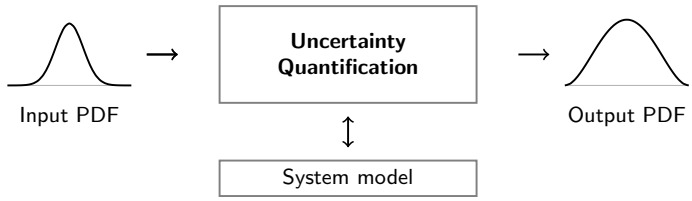

Figure 2.6: Scheme of uncertainty propagation from the input of the system model to the output.

from the reference point in the objective space). Function Φ is defined in such a way that it attains lower values in the non-dominated region of the objective space. Secondly, a MADS run is performed to minimize this scalarized function. Then, a new list of non-dominated solutions is determined from the evaluations obtained in the preceding step. Finally, the repetition of the three step is terminated according to a fixed number of iterations. A version of this algorithm exists for optimization problem with more than two objective functions; it is not employed here because only bi-objective cases are considered.

The BIMADS method is very efficient for low-dimensional problems, but its rate of convergence degrades as the number of design variables increases. Similarly to genetic algorithms, the BIMADS method treats the evaluation of the function for each design vector as a black-box. Therefore, it is very suited to coupling with the UQ inner loop. Concerning the implementation, the NOMAD software [70, 71] developed by the authors of the method is coupled with the UQ method.

2.3 Uncertainty quantification

The goal of uncertainty quantification is the quantitative description of the likelihood of a quantity of interest given its dependence on the uncertainty. It can only be obtained by combining the knowledge of the sources of uncertainties and the behavior of the quantity of interest functions with respect to the uncertain variables [72]. In a probabilistic framework, where the uncertain variables are described by a Probability Distribution Function (PDF), the goal of uncertainty quantification becomes the computation of the PDF of the quantity of interest. Thus, an uncertainty quantification methods propagates the uncertainty of the input variables into the uncertainty of the output variables by leveraging a model of the input uncertainty (e.g. the input PDF) and the system model connecting the input to the output variables (see Fig. 2.6).

Uncertainty quantification essentially involves two steps: (i) the definition and modelling the uncertainty, and (ii) the analysis of the uncertainty propagation through the simulation. These subsequent steps are discussed in the following sections.

2.3.1 Modelling uncertainty

It is clear from Fig. 2.1 that many sources of uncertainty exist in the lifecycle of an aircraft. Sources of uncertainty are typically grouped into the following categories:

- *Model input uncertainty and model parameter uncertainty*: this type of uncertainty includes the variability of the input variables of the system model (geometric variables, initial and boundary conditions, environmental parameters) and the variability of the parameters used in the model (physical constants, model coefficients).

- *Model form uncertainty*: also referred to as model inadequacy or model discrepancy, this is the uncertainty involved in model selection and it is due to a lack of confidence in selecting the best model [73]. This lack of confidence typically stems from a lack of knowledge of the underlying true physics, or by the impossibility to use the complete, accurate model of the true physics.

- *Model error*: this type is associated with the numerical simulation of the system model and it may include discretization errors of the equation describing the system model, round-off errors due to finite digits number representation, or iteration residual errors. It is important to note that although model errors are ascribed to the different sources of uncertainty, they differ from the other uncertainty because they are defiencies of the numerical implementation of the model and they can be analysed and reduced by a verification process. Uncertainty, on the other hand, are associated to the choice of the physical models and to the specification of the input parameters required for performing the analysis [74].

In this work, the focus is drawn on model input uncertainty, as also highlighted in Fig. 2.6. With regard to model form uncertainty and model error, these types can be investigated by the process of model Verification and Validation [75]. The verification step is aimed at determining whether a model implementation accurately represents the developer's conceptual description of the model. It basically answers the question: "are we solving the equations correctly?" [74]. So, for instance, the discretization and round-off errors can be estimated by comparison between numerical results of the programmed simulation model and analytical results of the exact

mathematical model. The validation process attempts to determine how accurate is a model in representing the real world. It answers the question: "are we solving the correct equations?". Model form uncertainty can be characterized by Bayesian approaches [73, 76, 77], or through model accuracy assessment by comparison between simulation results and experimental measurements [78].

Classification of uncertainty

The most popular uncertainty classification divides uncertain variables into two types: aleatory and epistemic. *Aleatory uncertainties* are irreducible uncertainties arising from the inherent variation associated with the system [57]. This type is also called *stochastic* to highlight the random stochastic nature associated with these variables. Aleatory uncertainties are typically environmental conditions, measurement data, manufacturing tolerances. On the other hand, *epistemic uncertainties* represent a lack of knowledge associated with the modeling process that are reducible through the introduction of additional information [79]. These variables are generally parameters of the system model.

For example, in the field of fluid dynamics, aleatory uncertainties are the uncertainties on the geometric parameters [80, 81, 82] and on the operating conditions [60, 25]. Other works focus on epistemic uncertainties, or mixed aleatory-epistemic uncertainties [83]. For example, in Ref. [84], a combination of an aleatory uncertainty on the freestream Mach number with an epistemic uncertain variable, i.e. the kinematic eddy viscosity of the the Spalart-Allmaras model [85], is considered.

Theoretical models for uncertainty

Different approaches are used to provide an appropriate representation and modelling of uncertainty according to its specific features and the amount of information available about it. Traditionally, probability theory has provided the mathematics for the representation of uncertainty in non-deterministic design, thanks to its long history and sound theoretical foundation [52, 57]. Recently, other mathematical theories for the representation of uncertainty have been introduced, including evidence theory, possibility theory and interval analysis. The terms *non-probabilistic* or *imprecise* probability approaches are used to refer to these mathematical models. As probability theory has deep root in the research of non-deterministic design and it is the theory used throughout the thesis, a more detailed discussion is given on it.

Probability theory Probability theory is a mathematical construction to describe phenomena whose outcomes occur non-deterministically. Few definitions to describe probability theory are given in this section, which are derived by the exposition of Ref. [86, Appendix A]. A complete discussion on probability theory may be found for instance in Refs. [87, 88].

Formally, the probability space is a triple $(\Theta, \mathcal{F}, P)$ where

(i) Θ is the sample space, containing all the possible outcomes or everything that could occur in the particular universe under consideration. An element of Θ is denoted by ϑ.

(ii) $\mathcal{F}$, is the collection of all the possible events, where an event is a set of outcomes of Θ. The collection of events $\mathcal{F}$ is called a σ–field if

- $\emptyset \in \mathcal{F}$
- $A \in \mathcal{F} \Rightarrow \bar{A} \in \mathcal{F}$, where A is an event and $\bar{A}$ is the complement of A;
- $A_1, A_2, ..., A_N \in \mathcal{F} \Rightarrow \left(\cup_{i=1}^{N} A_i\right) \in \mathcal{F}$;

The previous properties indicate that the field contains the empty set (event with none of the outcome), for each event included in the field it contains the complement of such an event and it contains the event containing all outcomes.

(iii) P is a probability measure of each event in $\mathcal{F}$.

An $\mathbb{R}$-valued random variable ξ is a map from the sample space into $\mathbb{R}$, i.e. $\xi : \Theta \mapsto \mathbb{R}$. The previous definition of a random variable can be extended to $\mathbb{R}^{n_\xi}$-valued functions $\boldsymbol{\xi}$, and if all the n_ξ coordinates of $\boldsymbol{\xi}$ are random variables, then $\boldsymbol{\xi}$ is called a random vector.

The integral of ξ with respect to P over the event $A \in \mathcal{F}$ is given by

$$\int_{\Theta} I_A(\vartheta)\xi(\vartheta)\, dP(\vartheta) = \int_{A} \xi(\vartheta)\, dP(\vartheta) \tag{2.6}$$

where I_A is the indicator function of A (i.e. a function that is equal to 1 if $\vartheta \in A$, and 0 elsewhere). If the previous integral exists and it is finite, ξ is said to be P-integrable over A. A random vector $\boldsymbol{\xi}$ is P-integrable over A, if each of the elements are P-integrable over A. When the integration domain is the complete sample space Θ, the integral corresponds to the expectation operator

$$\mathbb{E}[\xi] = \int_{\Theta} \xi(\vartheta)\, dP(\vartheta). \tag{2.7}$$

Let us denote $L_\beta(\Theta, \mathcal{F}, P)$, with β an integer number greater or equal to 1, the collection of $\mathbb{R}$-valued random variables ξ defined on $(\Theta, \mathcal{F}, P)$ such that

$$\mathbb{E}[|\xi|^\beta] < \infty. \tag{2.8}$$

The case $\beta = 2$ gives the L_2-space which possesses some essential properties. The expectation $\mathbb{E}[\xi\ \lambda]$ of two random variables ξ, λ defines an inner product on L_2

$$\mathbb{E}[\xi\ \lambda] = \langle \xi, \lambda \rangle \tag{2.9}$$

and the associated L_2 norm reads

$$||\xi||_{L_2} = \sqrt{(\mathbb{E}[\xi\ \xi])}. \tag{2.10}$$

This aspect is important since most uncertainty modelling in probabilistic framework consider second-order random vectors, i.e. random vectors belonging to L_2, assuring the existence of a finite value of the L_2 norm.

From the probability measure P it is possible to define the cumulative distribution function and the density function of a random variable. The Cumulative Distribution Function (CDF) of a random variable ξ defined on a probability space $(\Theta, \mathcal{F}, P)$ is given by

$$F_\xi(\tilde{\xi}) = P\left(\{\vartheta : \xi(\vartheta) \leq \tilde{\xi}\}\right) = P\left(\xi \leq \tilde{\xi}\right) \tag{2.11}$$

where $\tilde{\xi} \in \mathbb{R}$ is the realization of ξ: $\tilde{\xi} = \xi(\vartheta)$. The following properties hold true for the cumulative distribution function

$$\lim_{\tilde{\xi} \to +\infty} F_\xi(\tilde{\xi}) = 1 \qquad \lim_{\tilde{\xi} \to -\infty} F_\xi(\tilde{\xi}) = 0. \tag{2.12}$$

The distribution function is increasing, with range [0, 1]. If the cumulative distribution function F_ξ is continous in $\mathbb{R}$, there is an integrable function f_ξ such that

$$F_\xi(b) - F_\xi(a) = \int_a^b f_\xi(\tilde{\xi}) d\tilde{\xi}, \quad a \leq b \tag{2.13}$$

where $a, b \in \mathbb{R}$. Function f_ξ is called Probability Density Function (PDF) and it has

the following properties:

$$f_\xi(\tilde{\xi}) = \frac{dF_\xi}{d\tilde{\xi}} \qquad \int_{-\infty}^{\tilde{\xi}} f_\xi(z)dz = F_\xi(\tilde{\xi})$$
$$f_\xi(\tilde{\xi}) \geq 0 \qquad \int_{-\infty}^{+\infty} f_\xi(x)dx = 1. \tag{2.14}$$

Some important properties apply to transformation of random variables. Let us consider a random variable ξ defined on a probability space $(\Theta, \mathcal{F}, P)$ and a function g such that $\lambda = g(\xi)$ is a random variable defined on $(\Theta, \mathcal{F}, P)$. The expectation of λ can be computed as follows

$$\begin{aligned}\mathbb{E}[\lambda] &= \int_\Theta \lambda(\vartheta)\, dP_\lambda(\vartheta) = \int_{\mathbb{R}} \tilde{\lambda}\, dF_\lambda(\tilde{\lambda}) = \int_{\mathbb{R}} g(\tilde{\xi}) dF_\xi(\tilde{\xi}) \\ &= \int_{\mathbb{R}} g(\tilde{\xi}) f_\xi(\tilde{\xi}) d\tilde{\xi},\end{aligned} \tag{2.15}$$

or in other words the expectation of λ can be computed from the probability density function of ξ. The previous equation is sometimes calld the "Law of the Unconscious Statistician" [88, p. 200] and it is particularly important to UQ. In fact, if ξ is the input of the system under consideration and λ is the output, equation (2.15) clearly states that to reconstruct the output PDF it is necessary to combine the input-output function (i.e. the system model) and the PDF of the input (as presented in Fig. 2.6). In the trivial case $\lambda = g(\xi) = \xi$ the previous equation reduces to the classical definition of expectation

$$\mathbb{E}[\xi] = \int_{\mathbb{R}} \tilde{\xi} f_\xi(\tilde{\xi}) d\tilde{\xi}. \tag{2.16}$$

In addition, the moment of order r of a random variable ξ defined on the probability space $(\Theta, \mathcal{F}, P)$ is given by

$$\mu^r(\xi) = \mathbb{E}[\xi^r] = \int_{-\infty}^{+\infty} \tilde{\xi}^r f_\xi(\tilde{\xi}) d\tilde{\xi}. \tag{2.17}$$

If $\xi \in L_r(\Theta, \mathcal{F}, P)$, then $\mu^r(\xi)$ exists and it is finite. The central moment of order r is obtained by subtracting the first-order moment $\mu^1 = \mu$, callded the expected value or mean value, as follows

$$\mu^r(\xi) = \int_{-\infty}^{+\infty} \left(\tilde{\xi} - \mu\right)^r f_\xi(\tilde{\xi}) d\tilde{\xi}. \tag{2.18}$$

The second-order central moment is called variance: $\mu^2(\xi) = \sigma^2(\xi)$.

Finally, the joint distribution function of a random vector $\boldsymbol{\xi} = \{\xi_1, ..., \xi_{n_\xi}\}$ of n_ξ random variables is the direct extension of the univariate CDF

$$F_{\boldsymbol{\xi}}(\boldsymbol{\xi}) = P\left(\bigcap_{i=1}^{n_\xi}\{\xi_i \leq \tilde{\xi}_i\}\right) \tag{2.19}$$

and the joint density function of a random vector $\boldsymbol{\xi}$ is given by

$$f_{\boldsymbol{\xi}}(\boldsymbol{\xi}) = \frac{\partial F_{\boldsymbol{\xi}}}{\partial \xi_1 ... \partial \xi_{n_\xi}}. \tag{2.20}$$

The random variables of vector $\boldsymbol{\xi}$ are independent if and only if

$$P(\xi_i \leq \tilde{\xi}_i, i \in J) = \prod_{i \in J} P(\xi_i \leq \tilde{\xi}_i), \qquad \forall J \subset I, \;\; \tilde{\xi}_i \in \mathbb{R} \tag{2.21}$$

where $I = \{1, ..., n_\xi\}$. This also implies that the joint distribution and joint density functions can be computed as follows

$$\begin{aligned} F_{\boldsymbol{\xi}}(\boldsymbol{\xi}) &= \prod_{i=1}^{n_\xi} F_{\xi_i}(\xi_i) \\ f_{\boldsymbol{\xi}}(\boldsymbol{\xi}) &= \prod_{i=1}^{n_\xi} f_{\xi_i}(\xi_i). \end{aligned} \tag{2.22}$$

To sum up, modelling uncertainty in a probabilistic framework means determining the PDF of the random variable (or the joint density function for a random vector).

Non-probabilistic approaches Non-probabilistic approaches are typicallly used when poor knowledge is available about the uncertainty under study. In these cases, assigning a precise probability distribution for a proposition may be too strong a solution in most applications.

A common non-probabilistic method used to describe epistemic uncertainty is **evidence theory**. Evidence theory, also called the Dempster–Shafer theory, is a mathematical construct to measure uncertainty involving two measures of likelihood: belief, i.e. lower limit of likelihood and plausibility, i.e. upper limit of likelihood. Belief and plausibility are computed with a function which is assigned to a specific interval of the possible realizations based on the information on the system. This function can be thought of as the amount of likelihood that is associated with the

specific interval but without any specification of how this likelihood might be apportioned over the inverval [57]. Thus, evidence theory does not assume any particular value within the interval and nor does it assign a specific distribution to the interval. The information or evidence used to assign this likelihood function comes from a wide range of sources, e.g. experimental data, theoretical evidence, experts' opinions concerning belief in the occurrence of an event. This theory has proved effective in representing epistemic uncertainty in case of uncertainty information coming from multiple sources with application in fluid dynamics to the treatment of uncertainty on the closure coefficients used in the Spalar–Allmaras turbulence model [89].

Similarly to evidence theory, **possibility theory** involves the definition of two functions: possibility, i.e. the largest value of the likelihood assigned to an event, and necessity, i.e. the complement likelihood of the possibility of the complement event. For both evidence and possibility theories, an upper bound and lower bound of the cumulative distribution function can be defined in order to describe the uncertainty and propagate it to the output [57].

Instead, in **interval analysis** the value of a variable is replaced by a pair of numbers representing the maximum and minimum values that the variable is expected to take. Applications based on interval analysis do not attempt to infer an uncertainty structure of the output of the system based on the input uncertainty, so interval analysis does not provide an uncertainty structure. Interval arithmetic rules are then employed to perform mathematical operations with the interval numbers, so as to propagate the interval bounds through the computational model and compute the bounds of the output variables.

2.3.2 Propagating uncertainty

As previously mentioned, in this work the framework of probability theory is used to model uncertainty, because it provides a solid and comprehensive theoretical foundation and offers the most versatile statistical tools. Within a probabilistic framework, the uncertain variables vector $\boldsymbol{\xi}$ is defined as a random vector of second-order, independent random variables (see Sec. 2.3.1 for details) that can have any arbitrary probability density $f_{\boldsymbol{\xi}}(\boldsymbol{\xi})$.

The goal of uncertainty quantification is the following: given the probability density function of $\boldsymbol{\xi}$, calculate the probability density function of the quantity of interest u, i.e. the performance. Specifically, in the case of uncertainty quantification applied to optimization the actual goal is to compute one or more statistics of the performance u; for instance, in robust design the objective functions $\mathbf{f}$ are given by

low-order moments such as the mean value and the variance

$$\mathbf{f} = \{\mu_u(\boldsymbol{x}), \sigma_u^2(\boldsymbol{x})\}, \tag{2.23}$$

where $\boldsymbol{x}$ is the vector representing a specific design, for instance the i-th design to be evaluated inside the optimization loop: $\boldsymbol{x} = \boldsymbol{x}_i$ (recall Fig. 2.4). The statistical moment $\mu_u^r(\boldsymbol{x})$ of order r of the output of interest u is computed as

$$\mu_u^r(\boldsymbol{x}) = \mathbb{E}[u^r](\boldsymbol{x}) = \int_{\Xi} u(\mathbf{x}, \boldsymbol{\xi})^r f_{\boldsymbol{\xi}}(\boldsymbol{\xi}) d\boldsymbol{\xi}, \tag{2.24}$$

where $\Xi \subseteq \mathbb{R}^{n_\xi}$ is the support of the density function $f_{\boldsymbol{\xi}}$ which is equal to $f_{\boldsymbol{\xi}} = \prod_{i=1}^{n_\xi} f_{\xi_i}$ for independent random variables (see Eq. (2.22)). Several methods exist to reconstruct the variability of u and to compute Eq. (2.24). Common approaches in uncertainty-based optimization are Monte Carlo methods, Polynomial Chaos methods [90, 91] and collocation-based methods [92, 93]. A fundamental distinction among these approaches is the classification into intrusive and non-intrusive methods [86].

Intrusive methods involve the reformulation of the system governing equations and consequently the modification of the simulation codes, so as to incorporate uncertainty directly into the system. Typical example of this type is the intrusive Polynomial Chaos Expansion [86, p. 28], which consists in representing the stochastic process by means of an expansion of orthogonal polynomials. In other words, the solution $\phi(\boldsymbol{x}, \mathbf{y}, t, \boldsymbol{\xi})$ of Eq. (2.3) is replaced by an expansion of orthogonal polynomials in the space of the stochastic variables $\boldsymbol{\xi}$. By truncating the expansion up to a desired order and inserting the expansion in the system governing equations, a system of coupled deterministic equations is found instead of the stochastic equation (2.3). The coefficients of the polynomials become the unknowns of this augmented system model, which can be solved by editing the existing analysis codes. It is clear that this approach requires a significant modification of the deterministic simulation code.

On the other hand, *non-intrusive* approaches treat numerical simulation code as a black-box model and do not need modification to the existing deterministic simulation codes. With a proabilistic approach, a non-intrusive method typically draws n_s samples of the input PDF and for each sample solves the system model equation (Eq. (2.3)) by means of a suited simulation code (see Fig. 2.7). From each simulation, a sample of the output u is obtained, and the application of specific operators to the samples of u enable the reconstruction of the statistics. As a result, these methods are very versatile and they can be easily integrated inside an optimization loop. In

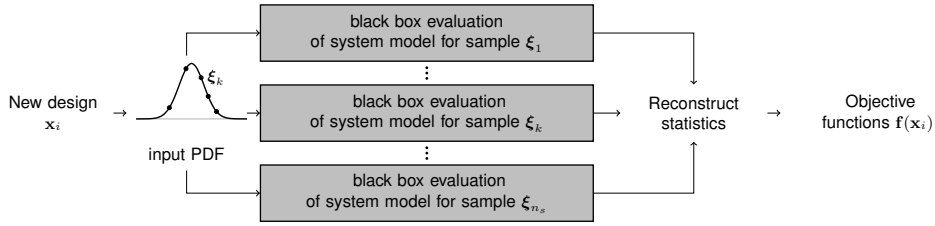

Figure 2.7: Block diagram of non-intrusive UQ methods.

addition, they can be distributed and performed in parallel. In the following section the most common approaches are discussed: Monte Carlo methods, non-intrusive Polynomial Chaos and Stochastic Collocation.

Monte Carlo approaches

In Monte Carlo approaches, a sequence $\boldsymbol{\xi}_1, \boldsymbol{\xi}_2, ..., \boldsymbol{\xi}_{n_s}$ of n_s samples of the random vector is sampled according to the probability distribution of $\boldsymbol{\xi}$, which is obtained from the uncertainty modelling. In this case, repeated simulations of Eq. (2.3) are required for each sample in the stochastic space to compute the realizations $\tilde{u}_k = u(\phi(\boldsymbol{x}, \mathbf{y}, t, \boldsymbol{\xi}_k)),\ k = 1, ..., n_s$. The integral in Eq. (2.24) is then replaced by a summation over the samples

$$\mu_u(\boldsymbol{x}) \approx \frac{1}{n_s} \sum_{k=1}^{n_s} \gamma_k \tilde{u}_k, \tag{2.25}$$

where γ_k is a weighting coefficient associated with each realization ($\gamma_k = 1\ \forall k$ for non-biased sampling).

In most application Quasi-Monte Carlo (QMC) methods are used, which differ from the usual Monte Carlo simulation in that they use quasi-random sequences instead of pseudo random numbers. The quasi-random sequences are the low discrepancy sequences, which are deterministic and fill the space more uniformly than pseudo random sequences. Different methods can be used to compute these sequences [86, p. 49]: in Fig. 2.8 Sobol sequences are used for the QMC case. Another improved sampling technique is Latin Hypercube Sampling (LHS), which draws samples within equiprobable bins in the parameter range.

This type of methods have several advantages: they are simple, because it is sufficient to resolve the deterministic model several times, and they are robust, because

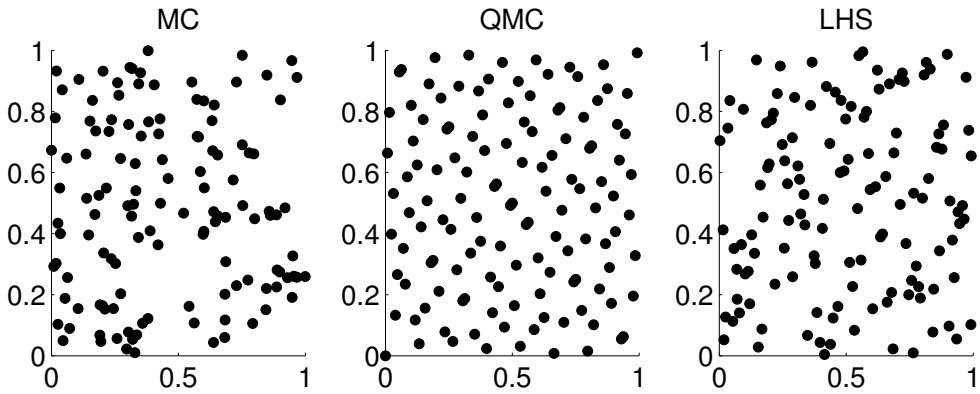

Figure 2.8: Samples sequence of two uniformly distributed variables with Monte Carlo (left), Quasi-Monte Carlo with Sobol sequence (center) and Latin Hypercube Sampling (right).

they do not rely on the regularity of the output of Eq. (2.3) with respect to the uncertain input [86, p. 8-9]. In addition, the rate of convergence scales with $n_s^{-1/2}$ with n_s the number of samples. Thus, the convergence rate does not depend on the dimensionality of the random vector, which is a clear advantage with respect to other methods which suffers of the curse of dimensionality. However, a large number of realizations are required to achieve a sufficiently good estimate of the statistical moments. This issue is of paramount importance for fluid dynamics applications or in any applications in which the computation of Eq. (2.3) is very expensive.

Non-Intrusive Spectral Projection with Polynomial Chaos

The employment of the Non-Intrusive Spectral Projection (NISP) based on the Polynomial Chaos (PC) expansion can reduce the computational cost of the uncertainty propagation, especially when a small number of uncertain variables are considered.

In fact, under specific conditions, a stochastic process can be expressed as a spectral expansion based on suitable orthogonal polynomials, with weights associated to a particular probability density function. The first study in this field is the Wiener process that was later developed in Ref. [91]. The basic idea is to project the variables of the problem onto a stochastic space spanned by a complete set of orthogonal polynomials Ψ that are functions of the random variables in vector $\boldsymbol{\xi}$. For example, the solution of Eq. (2.3) ϕ has the following spectral representation:

$$\phi\left(\boldsymbol{x}, t, \mathbf{y}, \boldsymbol{\xi}\right) = \sum_{k=0}^{\infty} \nu_k\left(\boldsymbol{x}, t, \mathbf{y}\right) \Psi_k\left(\boldsymbol{\xi}\right), \tag{2.26}$$

where Ψ_k are the PC orthogonal polynomials and ν_k the coefficients of the expansion, which depend on the remaining variables (design, spatial and time variables). In practice, the series in Eq. (2.26) has to be truncated to a finite number of terms

$$\phi\left(\boldsymbol{x}, t, \mathbf{y}, \boldsymbol{\xi}\right) \approx \sum_{k=0}^{\mathrm{N}} \nu_k\left(\boldsymbol{x}, t, \mathbf{y},\right) \Psi_k\left(\boldsymbol{\xi}\right), \tag{2.27}$$

where the number N is determined by

$$\mathrm{N}+1=\frac{\left(n_\xi+n_o\right)!}{n_\xi!\; n_o!}.$$

In the previous equation, n_ξ is the dimension of the random vector $\boldsymbol{\xi}$ and n_o is the order of the univariate polynomial expansion from which the multivariate polynomials $\Psi_k(\boldsymbol{\xi})$ are obtained via tensorization. In other words,

$$\Psi_k\left(\boldsymbol{\xi}\right)=\prod_{i=1}^{n_\xi} \varphi^{o_{i,k}}(\xi_i). \tag{2.28}$$

where $\varphi^{o_{i,k}}$ is the univariate polynomial in the variable ξ_i of order $o_{i,k}$, which varies from 0 to n_o depending on both k and i. An example for a two-dimensional case with $n_o = 2$ is given in Tab. 2.1. The polynomial basis $\varphi^{o_{i,k}}(\xi_i)$ is chosen according to the Wiener-Askey scheme [91] in order to select orthogonal polynomials with respect to the probability density function $f_{\xi_i}(\xi_i)$ of the specific input. For instance, if a uniform distribution is considered, Legendre polynomials are employed, while Hermite polynomials are used for a normal distribution (see Fig. 2.9). The corresponding polynomials for a group of commonly used PDF is given in Tab. 2.2.

The orthogonality property can be advantageously used to compute the PC coefficients of the expansion ν_k in a non-intrusive PC framework. In fact, from the

k	0	1	2	3	4	5
$o_{1,k}$	0	1	0	2	1	0
$o_{2,k}$	0	0	1	0	1	2
Psi_k	1	ξ_1	ξ_2	ξ_1^2	$\xi_1\xi_2$	ξ_2^2

Table 2.1: Terms in the tensorization of Eq. (2.28) for $n_\xi = 2$ and $n_o = 2$.

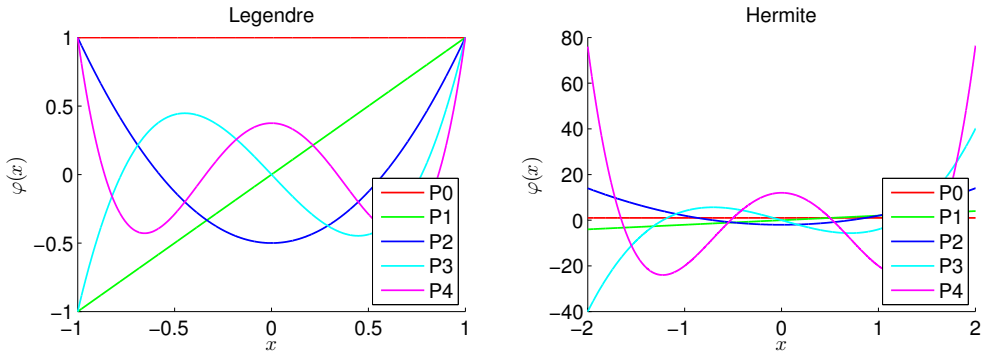

Figure 2.9: Orthogonal polynomial used in PC: Legendre basis (left) and Hermite basis (right).

Distribution	Polynomials $\varphi_i(\xi_i)$	Support
Uniform	Legendre	$[a, b]$
Gaussian	Hermite	$(-\infty, +\infty)$
Gamma	Laguerre	$[0, \infty)$
Beta	Jacobi	$[a, b]$

Table 2.2: Probability distribution and associated PC orthogonal polynomials.

orthogonality property it directly follows that

$$\nu_k = \frac{\langle \phi(\boldsymbol{x}, t, \mathbf{y}, \boldsymbol{\xi}), \Psi_k(\boldsymbol{\xi}) \rangle}{\langle \Psi_k(\boldsymbol{\xi}), \Psi_k(\boldsymbol{\xi}) \rangle} \quad \forall k. \tag{2.29}$$

where the operator $\langle \cdot, \cdot \rangle$ represents the inner product with respect to the probability density function, i.e.

$$\langle \phi(\boldsymbol{x}, t, \mathbf{y}, \boldsymbol{\xi}), \Psi_k(\boldsymbol{\xi}) \rangle = \int_{\Xi} \phi(\boldsymbol{x}, t, \mathbf{y}, \boldsymbol{\xi}) \; \Psi_k(\boldsymbol{\xi}) \; f_{\boldsymbol{\xi}}(\boldsymbol{\xi}) \; d\boldsymbol{\xi}. \tag{2.30}$$

Equation (2.29) states that the coefficients ν_k can be computed, up to the normalization factor $\langle \Psi_k, \Psi_k \rangle$, with an n_ξ-dimensional integral over Ξ. In other words,

$$\nu_k = \frac{\displaystyle\int_{\Xi} \phi(\boldsymbol{x}, t, \mathbf{y}, \boldsymbol{\xi}) \; \Psi_k(\boldsymbol{\xi}) \; f_{\boldsymbol{\xi}}(\boldsymbol{\xi}) \; d\boldsymbol{\xi}}{\langle \Psi_k(\boldsymbol{\xi}), \Psi_k(\boldsymbol{\xi}) \rangle} \quad \forall k. \tag{2.31}$$

The computation of the right-hand side of Eq. (2.31) can be estimated with several

approaches:

- *Monte Carlo method for NISP*: in this case the integral is replaced by a summation over a sequence of n_s samples

$$\langle \phi(\boldsymbol{x}, t, \mathbf{y}, \boldsymbol{\xi}), \Psi_k(\boldsymbol{\xi}) \rangle = \frac{1}{n_s} \sum_{i=1}^{n_s} \gamma_k \phi(\boldsymbol{x}, t, \mathbf{y}, \boldsymbol{\xi}_i) \Psi_k(\boldsymbol{\xi}_i) \tag{2.32}$$

  Similarly to the case presented in Sec. 2.3.2, improved sampling technique may be used to compute the sample sequence and reduce the sampling error.

- *Quadrature methods*: the integral is computed by a discrete sum

$$\langle \phi(\boldsymbol{x}, \mathbf{y}, t, \boldsymbol{\xi}), \Psi_k(\boldsymbol{\xi}) \rangle \approx \sum_{i=1}^{n_{QP}} g(\boldsymbol{\xi}_i) W(\boldsymbol{\xi}_i) \tag{2.33}$$

  where g is a function defined as $g = \phi(\boldsymbol{x}, \mathbf{y}, t, \boldsymbol{\xi}) \Psi_k(\boldsymbol{\xi})$, $\boldsymbol{\xi}_i$ are the n_{QP} quadrature nodes and $W(\boldsymbol{\xi}_i)$ are the weights associated with each node and defined according to the chosen quadrature formula [94]. The most common approach is using Gauss quadrature formulas, which are high-order approximations of integrals and closely related to the orthogonal polynomial family. In fact, the nodes $\boldsymbol{\xi}_i$ of the quadrature rule are the zeros of the orthogonal polynomial with respect to the density function $f_{\boldsymbol{\xi}}$. For instance, in the case of uniformly distributed variables Gauss-Legendre quadrature can be used, and Gauss-Hermite quadrature for normally distributed variables. The number of nodes and thereby samples n_s for a full quadrature rule is computed based on the order of the polynomials and the number of uncertain variables

$$n_s = (n_o + 1)^{n_\xi} . \tag{2.34}$$

  The Gauss quadrature rules require the computation of a new set of nodes for each desired order of the polynomial. To circumvent this problem, nested quadrature formulas are used, where formulas with different accuracy share nodes. This is the case for example of the Clenshaw-Curtis quadrature [86]. A major problem of quadrature formula is the exponential increase in the number of nodes with increasing dimensionality of the random vector $\boldsymbol{\xi}$. This problem, also known as "curse of dimensionality", may limit the application of NISP to situations where the stochastic dimension is small or moderate.

- *Sparse grid quadrature methods*: The sparse tensorization of quadrature formu-

las addresses the issue of the curse of dimensionality of full quadrature rules. In this case the multi-dimensional product is constructed as a linear combination of tensor product interpolants, each with relatively few points in its respective point set [74]. This greatly reduces the total number of abscissas, improving the computational efficiency of NISP and enlarged its domain of application [86].

Among these approaches, Gaussian quadrature formulas are used in this work [65], because they ensure higher accuracy and the saving provided by sparse grid methods is significant for a dimension of the random vector higher than the one considered in this study.

For any method used to solve the integral in Eq. (2.31), it is clear that the solution of the deterministic problem in Eq. (2.3) is required for each node of the quadrature rule or of the sample in the Monte Carlo sequence.

Once the polynomial chaos basis is defined and the associated ν_k coefficients are computed, the mean value and the variance of the stochastic variable $\phi\left(\mathbf{x}, t, \mathbf{y}, \boldsymbol{\xi}\right)$ are obtained by leveraging the orthogonality of the polynomials Ψ_k with respect to the probability function $f_{\boldsymbol{\xi}}$:

$$\begin{aligned}
\mu_\phi|_{PC} &= \left\langle \sum_{k=0}^{N} \nu_k\left(\boldsymbol{x}, t, \mathbf{y}\right) \Psi_k\left(\boldsymbol{\xi}\right) \right\rangle = \nu_0\left(\boldsymbol{x}, t, \mathbf{y}\right) \\
\sigma_\phi^2|_{PC} &= \left\langle \left(\sum_{k=0}^{N} \nu_k\left(\boldsymbol{x}, t, \mathbf{y}\right) \Psi_k\left(\boldsymbol{\xi}\right) \right)^2 \right\rangle - \nu_0^2\left(\boldsymbol{x}, t, \mathbf{y}\right) \\
&= \sum_{k=1}^{N} \nu_k^2\left(\boldsymbol{x}, t, \mathbf{y}\right) \langle \Psi_k^2 \rangle.
\end{aligned}$$

Please note that the uncertainty quantification methods presented so far are applied to the solution ϕ of Eq. (2.3) in order to make the discussion more general. However, the uncertainty quantification method in the robust optimization problem is applied to the quantity of interest u, which can be a function of $\boldsymbol{x}$ only (for instance, it is computed with an integral of the solution ϕ in both $\mathbf{y}$ and t). Thus, the discussion on the methods is still valid, only the dependence on the spatial coordinates $\mathbf{x}$ and time t drops out.

Polynomial chaos approaches rely on the regularity of the function ϕ in the stochastic space, so that the response ϕ may be approximated with accuracy by means of a polynomial basis. In addition, quadrature formula are generally global in that they spread nodes over the entire stochastic domain. However, in some applications

function ϕ may present large gradients only in small areas or be discontinous over Ξ. In these cases, the convergence of PC methods may be very slow and the method could be inefficient in practice. In these cases, a stochastic sampling may be preferable or a local adaptivity may advantageous to concentrate samples/nodes close to critical areas.

Multi-element collocation methods

When ϕ is discontinuous, the convergence of a global approximation, such as the one presented in the preceding section, deteriorates. It may then be advantageous to use a multi-element method, where the impacts of the discontinuity are confined to a limited number of small elements surrounding the discontinuity [95, 96].

In local UQ methods, such as the multi-element stochastic collocation approach, the weighted integrals in Eq. (2.24) defined over the parameter space Ξ are computed as a summation of integrals over n_e disjoint subdomains $\Xi = \bigcup_{j=1}^{n_e} \Xi_j$

$$\mu_u^r(\mathbf{x}) \approx \sum_{j=1}^{n_e} \int_{\Xi_j} u(\mathbf{x}, \boldsymbol{\xi})^r \; f_{\boldsymbol{\xi}}(\boldsymbol{\xi}) \; d\boldsymbol{\xi}. \tag{2.35}$$

To some extent the subdomains represent a numerical grid of elements in the space of the uncertain variables. An interesting approach is that of refining the element discretization with a step-by-step approach according to the quality of the current computation. The advantage of an adaptive strategy is using only as many samples as required for a sufficiently accurate computation of the statistics. In this work, the Simplex Stochastic Collocation (SSC) is chosen.

Simplex Stochastic Collocation The SSC method is a method implementing a simplex elements discretization of the probability space (see Fig. 2.10) with higher order interpolation and randomized sampling [97]. In the SSC approach [93, 98], the integrals in the simplex elements Ξ_j are computed by approximating the response surface $u(\boldsymbol{\xi})$ by an interpolation $w(\boldsymbol{\xi})$ of n_s samples $\mathbf{v} = \{v_1, \ldots, v_{n_s}\}$. Each sample v_k of the function $u(\boldsymbol{x}, \boldsymbol{\xi})$ is computed with the solution $\phi_k(\mathbf{x}, t, \mathbf{y}, \boldsymbol{\xi}_k)$ obtained by solving Eq. (2.3) for the k-th sample $\boldsymbol{\xi}_k$ of the random parameter vector $\boldsymbol{\xi}$

$$\mathcal{L}(\mathbf{x}, t, \mathbf{y}, \boldsymbol{\xi}_k; \phi_k(\mathbf{x}, t, \mathbf{y}, \boldsymbol{\xi}_k)) = 0 \quad \forall k = 1, \ldots, n_s. \tag{2.36}$$

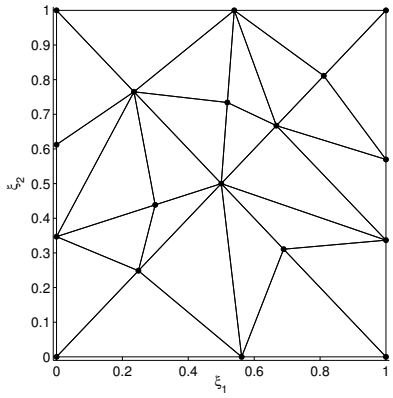

Figure 2.10: Simplex elements (triangles in two-dimensional stochastic space).

The interpolation of the samples $w(\boldsymbol{\xi})$ consists of a piecewise polynomial function

$$w(\boldsymbol{\xi}) = w_j(\boldsymbol{\xi}) \qquad \text{for } \boldsymbol{\xi} \in \Xi_j, \tag{2.37}$$

with $w_j(\boldsymbol{\xi})$ a polynomial interpolation of degree p of the samples $\mathbf{v}_j = \{v_{k_{j,1}}, \ldots, v_{k_{j,n_{se}}}\}$ at the sampling points $\{\boldsymbol{\xi}_{k_{j,1}}, \ldots, \boldsymbol{\xi}_{k_{j,n_{se}}}\}$ in element Ξ_j, where $k_{j,l} \in \{1, \ldots, n_s\}$ for $j = 1, \ldots, n_e$ and $l = 1, \ldots, n_{se}$, with n_{se} the number of samples in the simplex.

The main features of the SSC method [93, 98] are: (i) the possibility to compute an error estimate, (ii) the adaptive h-refinement and (iii) p-refinement capabilities, (iv) the treatment of non–hypercube parameter spaces, and (v) the treatment of discontinuities in the stochastic space. Properties (i), (ii), (iii) are discussed here, because they are the key elements used in the method proposed in the thesis. Please refer to Refs. [93, 98] for the other features.

With regard to the first property, the error ε of the interpolation is estimated using hierarchical surpluses ε_j, computed as the difference between interpolated and exact value of the k-th sample in each element Ξ_j

$$\varepsilon_j = w(\boldsymbol{\xi}_{j,k}) - v(\boldsymbol{\xi}_{j,k}). \tag{2.38}$$

In this way, both local and global error in each simplex can be computed, thus identifying the simplex element, where the error is largest.

Concerning h-refinement, SSC has two components: simplex selection (based on the previous error estimator), and simplex splitting, designed to avoid clustering. In the first step, the simplex element where the error ε is largest is selected for refinement. Then, a region of this element is selected for node insertion and the new sample

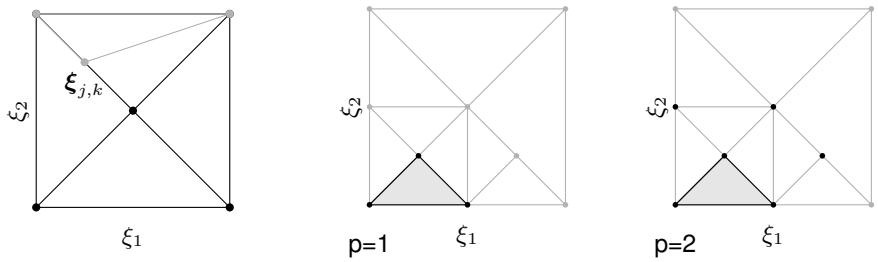

(a) h-refinement.

(b) p-refinement: samples higlighted with black circles are used in the interpolation for $p = 1$ (left) and $p = 2$ (right).

Figure 2.11: Scheme of the h- and p-refinement of the SSC in two-dimensional stochastic space.

$\boldsymbol{\xi}_{j,k}$ is inserted (see Fig. 2.11a). A threshold τ is assigned to determine whether to stop the refinement: when $\varepsilon < \tau$, the adaptive refinement is stopped.

When using SSC, the error convergence of first degree SSC decreases for an increasing n_ξ. To cure this issue, a p-refinement criterion for higher degree SSC has been conceived. In practice, the idea is to use higher degree polynomial interpolation stencil based on Newton-Cotes quadrature. Figure 2.11b shows for a two-dimensional stochastic space how more than three samples can be used to build a higher degree polynomial interpolation by enlarging the stencil. The choice of the optimal polynomial degree p is selected based on the number of uncertainties n_ξ and the required order of convergence. A super-linear convergence can be obtained for smooth responses (see [93] for more details).

The inner loop of the SSC is constituted by the following operations:

- Initial discretization of $n_{\mathrm{s_{init}}}$ samples (2^{n_ξ} vertexes of the hypercube enclosing the probability space Ξ and one sampling point in the interior);
- Computation of v_k computed in the initial samples by solving $n_{\mathrm{s_{init}}}$ deterministic problems (Eq. (2.36)) for the parameter values corresponding to the initial sampling points $\boldsymbol{\xi}_k$ located in Ξ only;
- Polynomial approximation (Eq. (2.37)) and p-refinement;
- Error ε estimate and h-refinement according to the threshold τ.

Finally, the probability distribution function and the statistical moments μ_u^r of $u(\boldsymbol{\xi})$ given by Eq. (2.35) are then approximated by the probability distribution and the

moments μ_w^r of $w(\boldsymbol{\xi})$

$$\mu_u^r(\mathbf{x}) \approx \mu_w^r(\mathbf{x}) = \sum_{j=1}^{n_e} \int_{\Xi_j} w_j(\mathbf{x}, \boldsymbol{\xi})^r \; f_{\boldsymbol{\xi}}(\boldsymbol{\xi}) \; d\boldsymbol{\xi}, \tag{2.39}$$

in which the multi-dimensional integrals are evaluated using a weighted Monte Carlo integration of the response surface approximation $w(\boldsymbol{\xi})$ with $n_{mc} \gg n_s$ integration points. This is a fast operation, since it only involves integration of piecewise polynomial function $w(\boldsymbol{\xi})$ given by (2.37) and does not require additional evaluations of the exact response $u(\boldsymbol{x}, \boldsymbol{\xi})$.

CHAPTER 3

Efficient methods for robust optimization

From the discussion of the previous chapter it is evident that the coupling of the optimization loop with the uncertainty quantification analysis results in a substantial increase of the computational cost of the optimization procedure. This increase is due to the higher number of evaluations of the solution of Eq. (2.3) required by the uncertainty propagation technique. It is therefore proportional to the cost of solving the system equations, which can be very high in most engineering applications.

For instance, let us consider an optimization problem in aerodynamics. A typical goal of aerodynamic optimization is the maximization of the aerodynamic efficiency, i.e. the ratio of aerodynamic lift over drag, by means of a modification of the wing shape. In this case, the mathematical model describing the aerodynamic system is represented by the governing equations of fluid dynamics. The numerical solution of this equation is achieved by Computational Fluid Dynamics (CFD) tools that require significant CPU time. The uncertainty-based optimization procedure requires several CFD computation for each sample in each optimization iteration (see Fig. 3.1). Thus, the optimization procedure requires a considerable computational effort.

Two different strategies to reduce the computational cost of the optimization under uncertainty are developed in the thesis. In the first strategy, the loop presented in Fig. 3.1 is not substantially modified, but the samples drawn to reconstruct the statistics of interest are evaluated by models of the system of different fidelities. In this case, the cost reduction is obtained by combining the employment of a high-fidelity, expensive simulation and a lower fidelity, cheaper simulation. The combined

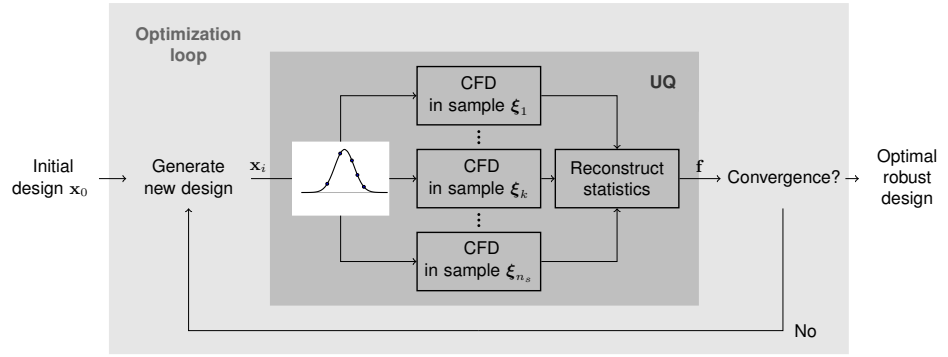

Figure 3.1: Block diagram of the uncertainty-based optimization loop.

use of models of different fidelity and level of approximation is called *multi-fidelity optimization* and it is used to balance computational cost and accuracy.

The second method leverages the employment of an adaptive UQ method, that improves the accuracy of the objective estimates with increasing number of samples. The rationale of the novel strategy is to compute with accuracy only the interesting, non-dominated solutions, while a rough, cheap estimate of the objective functions is sufficient for clearly dominated solutions. To understand whether higher accuracy is needed or not for the current individual of the optimization iteration, an exchange of information is necessary between the outer loop based on the optimization algorithm and the inner UQ loop. As a result, this is a *coupled method* for optimization under uncertainty to the extent that the inner UQ loop depends on the current status of the optimization loop.

The first section of the chapter is devoted to the description of the multi-fidelity strategy, while the second section discusses the adaptive method. For both methods, preliminary results based on algebraic test cases are presented.

3.1 Multi-fidelity uncertainty-based optimization method

In aerodynamic applications, to alleviate the cost of optimization procedures, the numerical model of the aerodynamics provided by the accurate, yet expensive, CFD solvers could be replaced with a lower-fidelity model. Examples of this type of model are potential flow models, which have been employed in the past both in fixed wing [99] and rotor airfoil design [22, 100], as well as surrogate models of

the performance, such as response surfaces [23]. These models may require a small amount of CPU time, but their range of applicability is limited and accuracy may be poor. The employment of inappropriate aerodynamic models inevitably affects the estimate of the objective function: inaccurate estimates, such as those obtained with a low-fidelity model, may drive the optimization procedure to false optima.

An alternative approach is based on the combined use of models of different fidelity and level of approximation in order to try to balance computational cost with the accuracy required. In the literature this solution is usually referred to as *multi-fidelity* or *variable-fidelity* optimization and it has been already applied in deterministic optimization [101, 102, 103, 104]. In the last years, some applications have considered the problem of uncertainty-based airfoil optimization, as well. Among these, a multi-fidelity method is developed in Refs. [105, 106] based on a hierarchical kriging approximation which exploits an estimate of the solution obtained by the same CFD solver, applied to fine grids for obtaining high fidelity results, and to coarser meshes for low fidelity ones. Another approach presented in Ref. [107] is based on the application of the uncertainty quantification technique to the system response obtained by the low-fidelity model and to the correction function, i.e. the discrepancy between the high-fidelity model and the low-fidelity model. Other works [83] builds a corrected low fidelity method with a space mapping from a CFD solution computed on a coarse mesh and uses this surrogate within a NISP algorithm.

In the following, a general framework is built for a multi-fidelity strategy inside the uncertainty-based optimization loop. The method is based on a simple switching logic between a low-fidelity, cheap simulation and a high-fidelity, expensive simulation. The underlying assumption of the method is that the discrepancy between high fidelity method and the low fidelity method is not uniform throughout the stochastic space, but it varies from region to region. Thus, there may be areas in the stochastic domain where the low fidelity model is sufficiently accurate, and areas where the error between the low- and high-fidelity methods is very high. This behavior was observed during a set of first experiments on the aerodynamic models considered for the aerodynamic design of airfoils in the hovering conditions. Thus, the main idea for the multi-fidelity strategy stems from the application itself. In this section, a generic description is given and an algebraic test case is tackled, whereas the application to the aerodynamic case is discussed in Chapter 4.

3.1.1 Description of the multi-fidelity strategy

The starting point is the definition of two models of different fidelity for the system under consideration: an accurate, expensive High-Fidelity (HF) model and a cheap, Low-Fidelity (LF) model. An initial characterization of the discrepancy between the estimates of the models must be performed. To do this, the quantity of interest u for a reference design is computed with both models in each point of the sampling set Ξ_s in the stochastic space. The set Ξ_s is defined by the chosen UQ non-intrusive method: it is the set containing the quadrature points if the coefficient of the NISP method are computed by quadrature, for example. The reference design is either the initial design or a design representative of the class of possible solutions defined by the chosen parameterization.

Then, a set Ξ_c of *control points* is defined which represents a subset of the samples set Ξ_s in the stochastic space. This subset includes some representative samples of the difference between the two models: for example, the sample corresponding to the nominal condition, or the samples with highest and lowest discrepancy in each dimension of the stochastic space may fall into this subset. Examples for a one-dimensional and a two-dimensional case are given in Fig. 3.2a. The control points are the points in which the *control function*, i.e. the discrepancy between the HF and LF models, is going to be computed. In this way, a measure of the difference between the models is estimated. The sample points will be evaluated with either the low or high-fidelity model according to a switching logic based on the value of the discrepancy of the closest control point.

Thus, for each individual inside the optimization loop at a generation, the following steps are performed (summarized in Fig. 3.2b):

1. compute the performance using the LF model u_{LF} and using the HF model u_{HF} for the control points and evaluate the control function $\mathcal{C}$ for each control point
$$\mathcal{C}_j = |(u_{HF,j} - u_{LF,j})/u_{HF,j}| \quad \forall j \in \Xi_c,$$
which represents the error of the low-fidelity estimate;

2. for each sample point $k \in \Xi_s$, the closest control point $j \in \Xi_c$ is found and if the control function $\mathcal{C}_j$ for such control point exceeds an assigned threshold value $\overline{\mathcal{C}}$, the simulation of the system model (Eq. (2.3)) is performed with HF; otherwise, the LF estimate is employed;

3. the mean value and variance are computed by means of the chosen UQ method using the samples in Ξ_s.

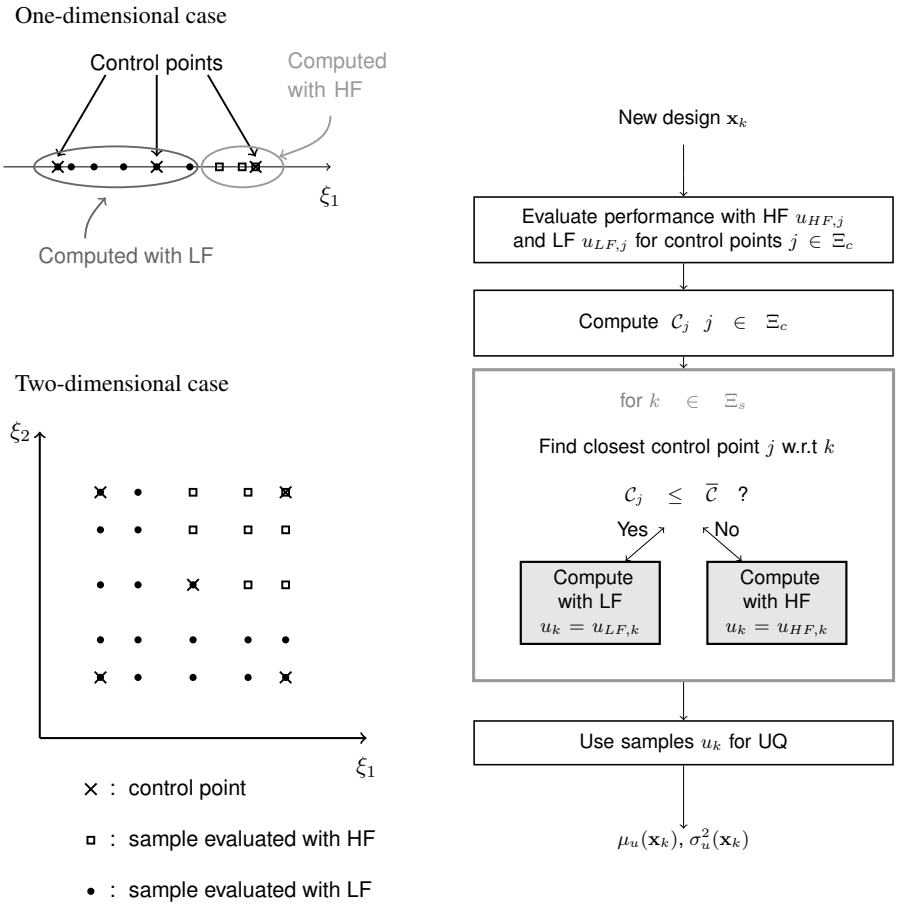

(a) Examples of control points.

(b) Flowchart of the multi-fidelity selection strategy.

Figure 3.2: Scheme of the multi-fidelity method.

As a result, some samples in the stochastic space are evaluated with the low-fidelity model, and others with the high-fidelity model (see Fig. 3.2).

The method is independent of the type of sampling in the stochastic space, and of the input PDF as well. However, it is strongly influenced by the choice of control points. In the cases presented in the thesis, three control points are chosen for a problem with one uncertain variables and five control points for a two-dimensional problem in the stochastic space. Additional intermediate control points can be used to improve the resolution of the control function. This may be a good strategy in the

case of control functions that have many local minima/maxima. It could reduce the number of HF samples for a given design, but it would increase the number of HF computations required by the switching logic. A trade-off number should be pursued, which may also take into account the available resources. As it will be discussed in the next chapter, a physical knowledge of the problem under consideration and of the models employed in the optimization would support the definition of the control points in the multi-fidelity strategy.

To some extent, this multi-fidelity method could be regarded as an extension of an Evolution Control strategy [108] in the case of uncertainty-based optimization. As a matter of fact, Evolution Control is a multi-fidelity technique for evolutionary algorithms in which either a certain number of individuals within a generation or a certain number of generations within an evolution are evaluated with the high-fidelity fitness function. The "control" consists in switching to the higher fidelity method for those selected individuals, thereby adding high-fidelity information about the fitness function of the population. In the application to robust optimization, the proposed method performs a control of the evolution in the stochastic space rather than in the design space. In other words, the switching logic acts on each individual of every generation, but only for a subset of samples in the stochastic space.

3.1.2 Preliminary results on algebraic test case

In this section, the multi-fidelity method is applied to an algebraic test case. A bi-objective ($m = 2$) optimization is considered, with one, uniformly distributed uncertain variable ($n_\xi = 1$) and two design variables ($n = 2$). The optimization problem reads

$$\begin{aligned} \text{minimize:} \quad & \mathbf{f}(\boldsymbol{x}) = \{f_1, f_2\} = \{\mu_u, \sigma_u^2\} \\ \text{by changing:} \quad & \boldsymbol{x} \in \mathcal{X} = [1,2] \times [1,2] \\ \text{with:} \quad & \xi \sim \mathcal{U}(0,1) \end{aligned} \tag{3.1}$$

where the function representing the performance is the following

$$u(\boldsymbol{x}, \xi) = \xi - x_1 \xi^5 + \cos(2\pi x_2 \xi) + \kappa, \qquad \kappa = 5. \tag{3.2}$$

The mean value and variance of function u are represented in the design space in Fig. 3.3. To test the multi-fidelity strategy, a High-Fidelity (HF) model and a Low-Fidelity (LF) model of function u are considered. In the former case, the HF function

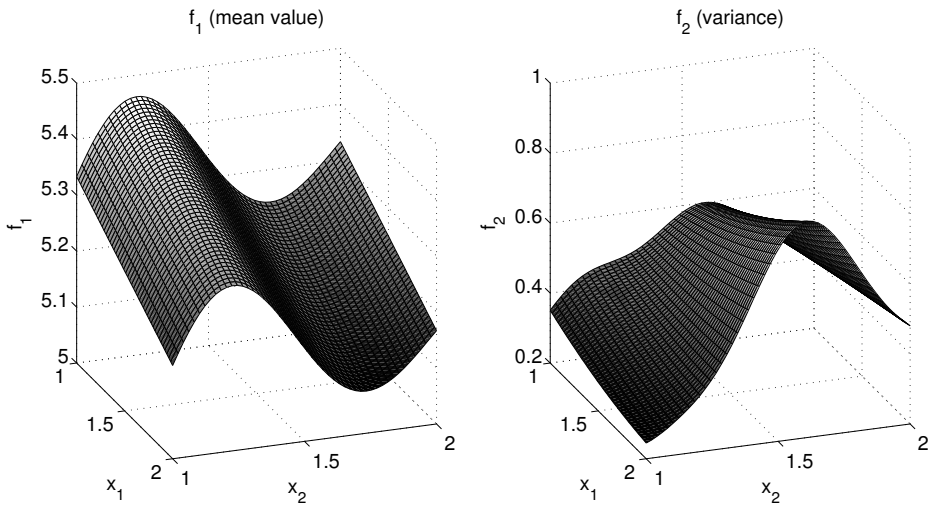

Figure 3.3: Objective functions for problem (3.1).

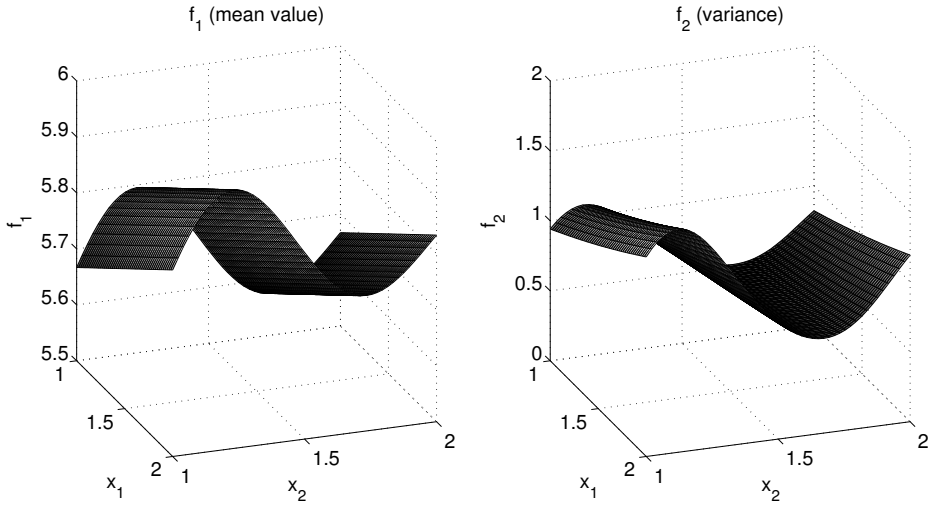

Figure 3.4: Objective functions for problem (3.1) with the LF model u_{LF} (Eq. (3.3)).

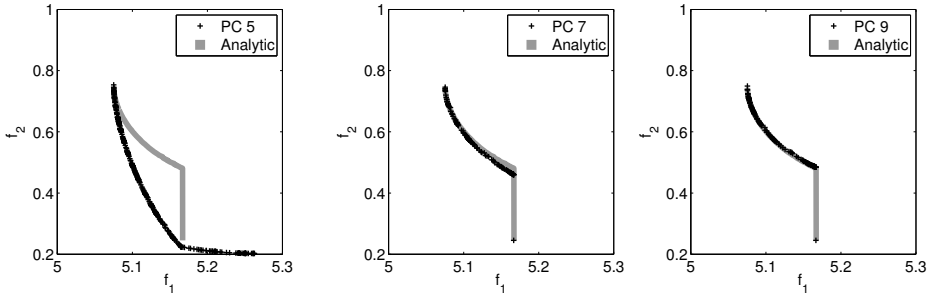

(a) 5-th order PC. **(b)** 7-th order PC. **(c)** 9-th order PC.

Figure 3.5: Pareto fronts obtained with different orders of PC expansion (HF model) and analytic function.

is the exact function $u_{HF} = u$, whereas the LF model is given by

$$u_{LF}(\boldsymbol{x}, \xi) = \xi + x_1 \xi^5 + \cos(2\pi x_2 \xi) + \kappa, \quad \kappa = 5. \tag{3.3}$$

The LF model of function u is compared to the HF model in Fig. 3.6a. The LF function u_{LF} is very similar to the accurate function for $\xi < 0.5$, while a higher discrepancy exists for increasing values of the stochastic variable. The difference between the two models strongly influences the resulting mean value and variance, which are reported in Fig. 3.4. This difference no doubt affects the results of the optimization problem.

The optimization under uncertainty stated in Eq. (3.1) is tackled by coupling the NISP method with a full Gauss quadrature method (please refer to Sec. 2.3.2) with the BIMADS optimization algorithm (presented in Sec. 2.2.3). The optimization is run three different times using different models to compute the samples of the NISP method: (i) using only HF samples, (ii) using only LF samples and (iii) with the multi-fidelity strategy which combines HF and LF samples. In the three cases the optimization is performed for 22 iterations and a total of 1000 evaluations of the objective functions. The NISP method with a PC expansion of Legendre 9-th order is used. Legendre polynomials is chosen to represent function u because the uncertain variables are uniformly distributed. The 9-th order is chosen after a convergence analysis on the result of the bi-objective optimization with different PC orders. As presented in Fig. 3.5, the overlap of the Pareto front obtained with the PC expansion and the analytic front improves with increasing order and an accurate solution is obtained for a 9-th order polynomial basis. Figure 3.6b shows that the PC method with HF samples is very accurate in representing the analytic function, while the PC

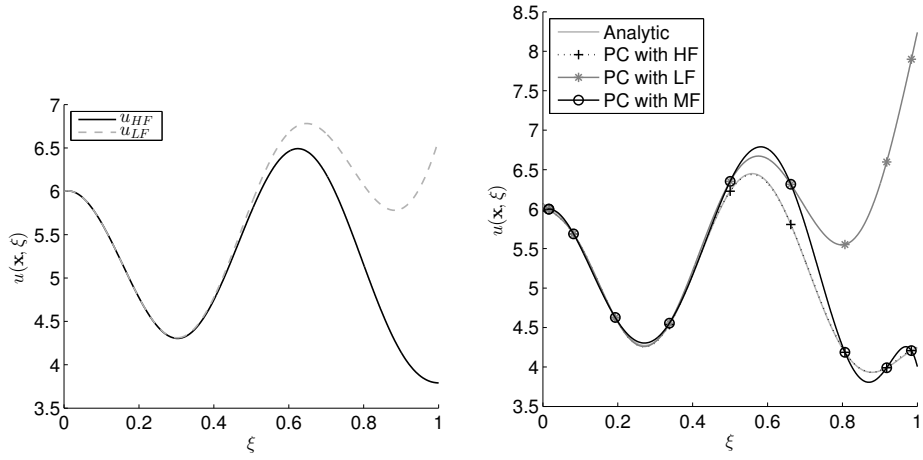

(a) HF and LF models for $\boldsymbol{x} = \{1.4, 1.6\}$.

(b) PC expansion using HF, LF and MF samples for $\boldsymbol{x} = \{1.99, 1.78\}$.

Figure 3.6: Representation of function $u(\boldsymbol{x}, \xi)$.

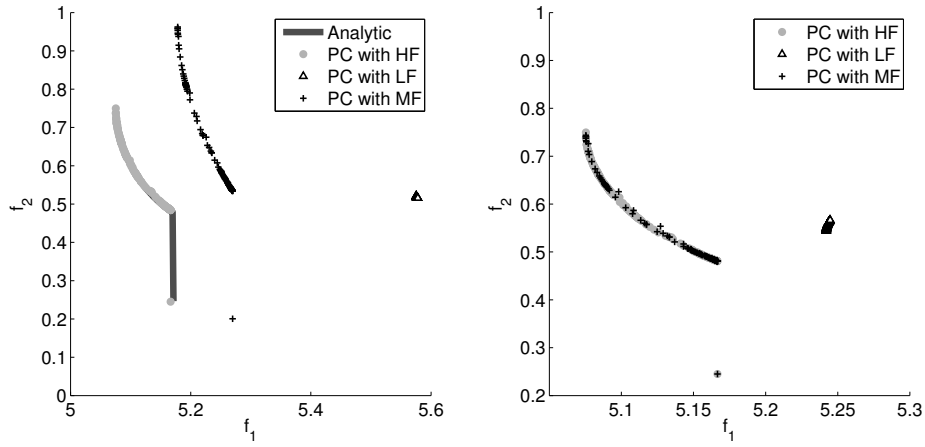

(a) Original Pareto fronts.

(b) Post-processed Pareto fronts.

Figure 3.7: Pareto fronts obtained with HF, LF and MF.

expansion based on the LF samples exhibits significant errors for $\xi > 0.5$. The MF method provides a trade-off between these two results: despite a small error at the maximum, the overall trend over the uncertain variable domain is preserved, owing to the employment of the HF samples in the last three quadrature points.

The results of the optimization are presented in Fig. 3.7. In particular, Figure 3.7a shows the analytic front (i.e. the front obtained by computing the analytic mean and

	HF	LF	MF
HF function evaluations	9000	–	5000
LF function evaluations	–	9000	7000
Distance to analytic front	7.79e-04	1.03e-1	8.69e-04

Table 3.1: Cost and distance to analytic front of HF, LF, MF approaches.

variance) and the approximate fronts in the objective space, where the objectives are computed according to the specific method (HF, LF or MF). Instead, in Fig. 3.7b the fronts are post-processed with the analytical function u to compare the results of the different optimization runs with the same metrics. The MF results overlap with the HF front, while the LF front is very far from both fronts. The accuracy of the result obtained with the MF method is also demonstrated by computing the distance to the analytic Pareto front. This distance is computed by taking the arithmetic mean of the distance of each point of the Pareto front to the closest point of the analytical Pareto front, as suggested in Ref. [109]. In mathematical terms, the distance d_{A-B} of the Pareto set $\mathcal{P}_A$ to the Pareto set $\mathcal{P}_B$ reads

$$d_{A-B} = \frac{1}{|\mathcal{P}_A|} \sum_{i=0}^{|\mathcal{P}_A|} \min_{\mathbf{f} \in \mathcal{P}_B} ||\mathbf{f}(\mathbf{x}_i) - \mathbf{f}(\mathbf{x}_{\mathcal{P}_B})||, \tag{3.4}$$

where symbol $|\cdot|$ indicates the cardinality (i.e. number of elements) of the set, symbol $||\cdot||$ is the Euclidean distance, $\mathbf{f}$ and $\mathbf{x}$ are respectively the objective functions and the design variables of the solutions in the Pareto set. The distance is reported in Tab. 3.1, where it is possible to note that the distance of the HF and MF method with respect to the analytic result are very similar. Furthermore, Table 3.1 presents the cost of the three strategies: the MF method requires a total number of evaluations (12000) higher than the HF and LF method, because both the HF and LF models have to be evaluated in the control points. However, the total number of HF evaluations is lower: if the LF simulation is 10 times faster than the HF evaluation, the cost of the HF optimization has been reduced by 37%. Thanks to the MF strategy, the accuracy of the Pareto front is preserved and the global cost is reduced.

3.2 Adaptive method with error bounding boxes

Another way to alleviate the cost of uncertainty-based optimization is to create a stronger coupling between the inner UQ loop and the outer optimization loop. Specifically, the two loops could share information, and the accuracy in the stochastic space may be increased as the algorithm gets closer to the optimum. Only few attempts to develop such an idea are given in the literature. For instance, in Ref. [110] a method has been presented, which improves the coupling between a Monte Carlo simulation and the stochastic annealing algorithm. To this end, the number of samples used to compute the statistics of the performance in the Monte Carlo simulation are increased as the annealing temperature decreases. This strategy has been extended to genetic algorithms in Ref. [111] to deal with multi-objective problems. In this case, the number of samples increases after a certain number of generations. Another example is the hierarchical stochastic optimization presented in Ref. [112], where the resolution in the stochastic space is increased over the iterations. A more recent work [113] illustrates the Simplex^2 method, a method which efficiently couples the Nelder-Mead algorithm for the optimization loop and the Simplex Stochastic Collocation [93, 98] for the uncertainty quantification part. This method only adresses the problem of single-objective optimization.

In such a context, the novel method presented in the following is based on an adaptive strategy in the UQ loop to reduce the number of samples in the stochastic space for those designs that are not interesting for optimization purposes. The main idea is to use the error bounds on the estimates of the objective functions provided by the UQ method and to define a bounding box in the objective space around the objective function estimate. A comparison between the bounding box and the current Pareto front (or current best design in the case of single objective optimization) during the optimization loop determines whether a refinement is needed in the stochastic space. This method applies to both single- and multi-objective optimization problems. Unlike the methods presented in the literature, the adaptation is not set a priori, but instead it depends on the convergence history of the optimization procedure and the performance of each new design. The strategy is here developed with either a Non-dominated Sorting Genetic Algorithm or the BiMADS algorithm for the optimization loop, and the Simplex Stochastic Colocation for the UQ method. However, the strategy may apply to different algorithms in both the optimization and the stochastic part, as long as it is possible to define an adaptivity with the available UQ method.

3.2.1 Pareto dominance with error bounding boxes

In Sec. 2.2.1, the traditional concept of Pareto dominance is presented, i.e. the relation to define the concept of "optimality" in a multi-dimensional objective space. Here, the Pareto dominance relation is adapted for the application to the proposed adaptive method.

Let us assume that each estimate of the objective functions $f_j, j = 1, ..., m$ is affected by an error $\varepsilon_{f_j}, j = 1, ..., m$. The error is associated to the approximate method used in the UQ loop to reconstruct the statistics of the performance u, i.e. the objective functions. Because of the error in the estimation of the statistics, each objective is a function f'_j that could take any value in the interval

$$f'_j(\boldsymbol{x}) \in [\underline{f_j}, ..., \bar{f}_j] \quad \forall j = 1, ..., m, \tag{3.5}$$

where the lower bound $\underline{f_j}$ and upper bound $\bar{f}_j$ are obtained respectively by subtracting and adding the error ε_{f_j} to the function estimate f_j. For instance, assuming the first objective f_1 is the mean value μ, the lower bound $\underline{f_1}$ and and the upper bound $\bar{f}_1$ read

$$\underline{f_1} = \mu - \varepsilon_\mu \qquad \bar{f}_1 = \mu + \varepsilon_\mu,$$

where ε_μ is the error affecting the mean value. Inside the adaptive UQ loop, the error ε_{f_j} decreases with increasing number of samples. Typically, a tolerance τ is used, and when the error ε_{f_j} is lower than the tolerance the loop breaks. Thus, for a given design vector $\mathbf{a}$ the amplitude of the errors is a function of the tolerance τ applied in the adaptive uncertainty quantification method. This holds true also for the objective function f'_j. Thus, both the objective function f'_j and the associated error ε_{f_j} depends not only on the design variables $\boldsymbol{x}$, but also on the tolerance value τ:

$$f'_j = f'_j(\boldsymbol{x}, \tau) \qquad \varepsilon_{f_j} = \varepsilon_{f_j}(\boldsymbol{x}, \tau) \qquad \forall j = 1, ..., m. \tag{3.6}$$

Once the error is lower than the limiting tolerance $\hat{\tau}$, the objective function passed on to the optimization algorithm is set equal to the estimate f_j given by the UQ method.

Typically, the same value of the tolerance is set for each design in the optimization loop, and the optimization and the uncertainty quantification are completely decoupled. Instead, in the proposed strategy, the value of the tolerance is updated inside the inner UQ loop, according to a comparison between the current Pareto front and the m-dimensional hypercube composed of the upper and lower bounds of each objective function (see Fig. 3.8). The hypercube is referred to as *error bounding box* and it is a

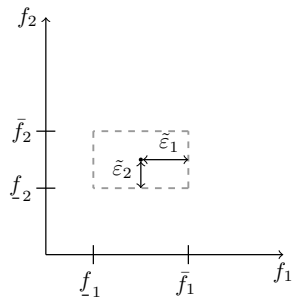

Figure 3.8: Error bounding box in the objective space.

function of the design vector and the tolerance:

$$\mathcal{B}(\boldsymbol{x}, \tau) = [\underline{f}_1, \bar{f}_1] \times ... \times [\underline{f}_m, \bar{f}_m]. \tag{3.7}$$

The error bounding box has also been used in Refs. [114, 115] to consider a confidence interval around the objective function. To compare the box and the Pareto front, a modified definition of the Pareto dominance criterion is used, which takes into account the error associated to the objective function of the current design. Let us consider a single iteration inside the optimization loop: the comparison is carried out between the current design $\mathbf{a}$ and one of the non-dominated solution vector $\mathbf{b}$ belonging to the current Pareto set.

Definition 3. *Consider two vectors $\mathbf{a}, \mathbf{b} \in \mathcal{X}$. The objective functions $\mathbf{f}(\mathbf{a}, \tau)$ of the first vector depend on the current tolerance τ (and thereby current refinement). The objective functions $\mathbf{f}(\mathbf{b})$ of the second vector have been computed in a previous step up to the limiting tolerance $\hat{\tau}$ and do not depend on τ. It is possible to state that:*

$$\mathbf{a} \succ \mathbf{b} \text{ ($\mathbf{a}$ dominates $\mathbf{b}$) iff } \beta_j \left[f_j(\mathbf{a}, \tau) - \beta_j \varepsilon_{f_j}(\mathbf{a}, \tau) \right] \leq \beta_j f_j(\mathbf{b}) \quad \forall j = 1, ..., m$$
$$\text{and } \exists j \text{ such that } \beta_j \left[f_j(\mathbf{a}, \tau) - \beta_j \varepsilon_{f_j}(\mathbf{a}, \tau) \right] < \beta_j f_j(\mathbf{b})$$
$$\mathbf{a} \sim \mathbf{b} \text{ ($\mathbf{a}$ is indifferent to $\mathbf{b}$) iff } \mathbf{a} \text{ does not dominate } \mathbf{b} \text{ and } \mathbf{b} \text{ does not dominate } \mathbf{a}$$

where coefficient β_j indicates the goal of the optimization of function f_j:

$$\beta_j = \begin{cases} 1 & \text{if } f_j \text{ is to be minimized} \\ -1 & \text{if } f_j \text{ is to be maximized} \end{cases}$$

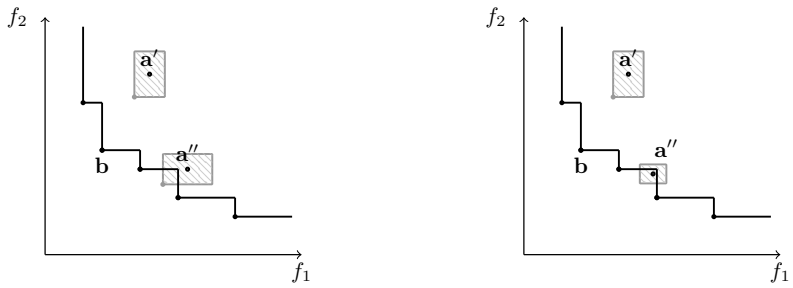

(a) Pareto dominance comparison: vector $\mathbf{a}'$ is clearly dominated, while the objectives of vector $\mathbf{a}''$ has to be computed with further refinement.

(b) Boxes after refinement of design $\mathbf{a}''$: the refinement proves that vector $\mathbf{a}''$ is non-dominated.

Figure 3.9: Scheme of Pareto dominance with error bounding boxes for a minimization problem.

With the use of the coefficient β_j, the definition considers the lower bound of the i-th function for a problem of minimization of f_j, and analogously the upper bound for maximization problem. Let us consider the example presented in Fig. 3.9 where a minimization of two objectives $\mathbf{f} = \{f_1, f_2\}$ is performed. In this case, to compare the grey shaded box with the current front in the objective space, the point corresponding to the lower bounds $(f_1 - \varepsilon_{f_1}, f_2 - \varepsilon_{f_2})$ is taken (the point is marked with a grey circle in the figure). Design $\mathbf{a}'$ is clearly dominated, whereas the box associated with design $\mathbf{a}''$ crosses the front. In the latter case, neglecting the error on the objective functions would result in excluding the design from the non-dominated set. However, a refinement in the stochastic space could improve the estimate of the objective functions and it could prove that the design is actually non-dominated (Fig. 3.9b).

In practice, if an adaptive UQ loop is considered, an estimate of the actual error is provided for each objective function: $\tilde{\varepsilon}_{f_j}$. The bounding box is built with the estimate of the error, rather than the error itself, and it is important that the estimate is conservative.

Definition 4. *In a minimization problem of m objectives affected by an error ε_{f_j}, $j = 1, ..., m$, it is possible to state with certainty that*

$$\mathbf{b} \succ \mathbf{a} \ (\mathbf{b} \text{ dominates } \mathbf{a})$$

if and only if

$$f_j(\mathbf{a},\tau) - f_j(\mathbf{b},\tau) > \varepsilon_{f_j}(\mathbf{a},\tau) + \varepsilon_{f_j}(\mathbf{b},\tau) \quad \forall j = 1,...,m. \tag{3.8}$$

When the objective functions are not affected by error, the previous definition corresponds to the inequality of the Pareto dominance relation (i.e. Definition 1).

Theorem 1. *Given a conservative estimate $\tilde{\varepsilon}_{f_j}$ of the error ε_{f_j} affecting the m objectives of a minimization problem, if*

$$f_j(\mathbf{a},\tau) - f_j(\mathbf{b},\tau) > \tilde{\varepsilon}_{f_j}(\mathbf{a},\tau) + \tilde{\varepsilon}_{f_j}(\mathbf{b},\tau) \quad \forall j = 1,...,m \tag{3.9}$$

then with certainty

$$\mathbf{b} \succ \mathbf{a} \ (\mathbf{b} \text{ dominates } \mathbf{a}).$$

Proof.
If the estimate is conservative, then

$$\tilde{\varepsilon}_{f_j} > \varepsilon_{f_j} \quad \forall j = 1,...,m \tag{3.10}$$

thus

$$\begin{aligned} \varepsilon_{f_j}(\mathbf{a},\tau) + \varepsilon_{f_j}(\mathbf{b},\tau) &< \tilde{\varepsilon}_{f_j}(\mathbf{a},\tau) + \tilde{\varepsilon}_{f_j}(\mathbf{b},\tau) \\ &< f_j(\mathbf{a},\tau) - f_j(\mathbf{b},\tau) \quad \forall j = 1,...,m. \end{aligned} \tag{3.11}$$

The latter relation is Definition 4. Thus, the dominance criterion with a conservative estimate of the error guarantees the detection of truly non-dominated solutions.

3.2.2 Algorithm

Leveraging the idea of the bounding box, an adaptive refinement on the error affecting the objective function is developed. The novel strategy proposes to modify only the inner UQ loop (Fig. 3.1), which provides the objective function $\mathbf{f}$ for a given design.

The operations of the adaptive strategy are presented in Algorithm 1. For the i-th design, the current design $\boldsymbol{x}_i$ and the set $\mathcal{Y}_{i-1}$ of the objective functions of the old design vectors are passed to the function. The set is used to compute the Pareto front according to the classical Pareto dominance relation, i.e. Definition 1. The tolerance τ is initialized to a coarse, moderately high value $\bar{\tau}$. Then, the objective function is evaluated up to the current tolerance with the UQ method of choice. With the adaptive

Algorithm 1 Scheme of error bounding box refinement

procedure EVALUATE($\boldsymbol{x}_i$, $\mathcal{Y}_{i-1}$)
  compute Pareto front from $\mathcal{Y}_{i-1}$ with Definition 1
  $\tau = \bar{\tau}$
  $\tilde{\boldsymbol{\varepsilon}} = BIG$
  while ($\tau \geq \underline{\tau}$) $\wedge$ ($\tilde{\boldsymbol{\varepsilon}} > \tau$) **do**
    evaluate $\mathbf{f}(\boldsymbol{x}_i)$ and $\tilde{\boldsymbol{\varepsilon}}(\boldsymbol{x}_i)$ with current τ
    if $\tilde{\boldsymbol{\varepsilon}} < \tau$ **then**
      compare current box $\mathcal{B}(\boldsymbol{x}_i)$ with front according to Definition 3
      if $\mathcal{B}(\boldsymbol{x}_i)$ is non-dominated **then**
        $\tau = \tau/10$
        $n_s = n_s + 1$
        continue
      else
        exit while
      end if
    else
      $n_s = n_s + 1$
    end if
  end while
  return $\mathbf{f}(\boldsymbol{x}_i)$
end procedure

UQ method, the error typically decreases by increasing the number of samples n_s, so the number of samples is increased if the error is higher than the tolerance. When $\tilde{\varepsilon}_{f_j} < \tau \; \forall j = 1, ..., m$, the error bounding box is compared with the current Pareto front: if the box is non-dominated, the tolerance is reduced by an order of magnitude and another iteration is performed. If the box is dominated according to Definition 3, the while loop is broken and the next individual is computed. The loop stops in any cases when the error is lower than the tight tolerance $\underline{\tau}$.

The choice on the method for reducing the tolerance is arbitrary and it is chosen based on the value of the limiting tolerance $\bar{\tau}$ and $\underline{\tau}$. Here, a factor equal to 1/10 is used because tolerance $\bar{\tau}$ and tolerance $\underline{\tau}$ differ of two or three orders of magnitude in the examples of the following sections. It should be noted that a smaller reduction results in a higher number of updates and number of comparison between the current front and the bounding box, but it could potentially reveal dominated solutions with a fewer number of samples.

The algorithm is formulated in such a way that the accuracy of the objective

estimates is variable throughout the design space. In fact, the limiting tolerance $\hat{\tau}$ is different for each design. Design vectors far from the Pareto front are computed with the coarse tolerance $\hat{\tau} = \bar{\tau}$, the Pareto solutions are computed with the tolerance $\hat{\tau} = \underline{\tau}$ providing the highest accuracy (and highest number of samples), and the intermediate solutions are computed with an intermediate tolerance on a case-by-case basis in the loop described in Algorithm 1.

As a final remark, the refinement proposed in this strategy is based on a comparison of the current design with the Pareto front. Thus, it is strictly associated to the performance of each design and to the importance of such a design in the optimization history. This is a key difference with respect to other approaches in the literature. For instance, the number of samples increases as the annealing temperature decreases in the stochastic annealing [110], or as the number of generations increases in stochastic genetic algorithm [111]. However in these cases, there is not a real connection with the values of the performance: due to the poor accuracy at the beginning of the optimization history, a good design may be discarded and, viceversa, a poor design may be included in the non-dominated set. In the proposed strategy, the non-dominated solutions are always computed with the highest accuracy (corresponding to the lowest tolerance $\underline{\tau}$). This requires a higher computational effort also for those individuals that may temporarily appear as Pareto optimal during the optimization process, and will be later dominated. However, the increase in the cost for these individuals is a necessary price to pay to guarantee that the optimization would end up to the accurate Pareto front. In future works, this condition may be relaxed to further increase the computational efficiency.

3.2.3 Convergence analysis

In this section, a proof of convergence is presented to prove that with a bounding box approach an optimization algorithm able to converge to the correct solution will converge also in the presence of an error up to the desired tolerance $\hat{\tau}$, if the error estimates are conservative.

From the property of the adaptive UQ method, the estimate of the objective function of a design vector $\mathbf{a}$ satisfies the following

$$\lim_{\tau \to 0} \varepsilon_{f_j}(\mathbf{a}, \tau) = 0 \tag{3.12}$$

Regarding the order of convergence, the UQ method used in this case is the Simplex Stochastic Collocation (described in Sec. 2.3.2), whose convergence is linear for first-order piecewise polynomials in a one-dimensional stochastic space, and superlinear

if a p-refinement is used. In the proposed strategy, a very low tolerance $\underline{\tau}$ is used for the solutions belonging to the Pareto front $\mathcal{P}$. In other words, by construction of the method, if $\mathbf{b} \in \mathcal{P}$, then $\varepsilon_{f_j}(\mathbf{b}, \tau) \to \underline{\tau}$.

Definition 5. *Let us consider the i-th iteration of the proposed optimization method, for a minimization problem of m objectives affected by an error $\varepsilon_{f_j},\ j = 1, ..., m$. Considering the i-th design vector $\boldsymbol{x}_i$, the estimate of the error satisfies*

$$\lim_{i \to \infty} \tilde{\varepsilon}_{f_j}(\boldsymbol{x}_i) = \underline{\tau} \qquad \forall j = 1, ..., m \tag{3.13}$$

Let us suppose to have an optimization method to solve the minimization problem of function f, whose solution is f^*. The quantity $f(\boldsymbol{x}_i)$ is the solution at the i-th iteration, which corresponds to the statistical moment of a given quantity for the design $\boldsymbol{x}_i$. Note that this quantity is considered as the exact statistical moment associated to the design $\boldsymbol{x}_i$. Instead, the approximated function computed with a tolerance and affected by an error is denoted with $\widetilde{f}_{j,i}(\boldsymbol{x}_i)$.

Definition 6. *A conservative estimate $\tilde{\varepsilon}_{f_j}$ of the error ε_{f_j} that affects the estimation $\widetilde{f}_j(\boldsymbol{x}_i)$ of the function $f_j(\boldsymbol{x}_i)$ is defined as follows*

$$\left\| f_j(\boldsymbol{x}_i) - \widetilde{f}_j(\boldsymbol{x}_i) \right\| \triangleq \varepsilon_{f_j} < \tilde{\varepsilon}_{f_j} \quad \forall j = 1, ..., m \tag{3.14}$$

Note that this conservative error estimation can be applied on the error of the mean, of the variance, or of a linear combination of some statistics if this is the objective function.

Theorem 2. *Let us suppose that the method converges to the solution of the optimization problem, i.e.*

$$\lim_{i \to +\infty} \left\| f_{j,i}(\boldsymbol{x}_i) - f_j^* \right\| = 0 \quad \forall j = 1, ..., m \tag{3.15}$$

or

$$\lim_{i \to +\infty} \|\delta_{j,i}\| = 0, \quad \forall j = 1, ..., m \tag{3.16}$$

Then, by considering the adaptive bounding-box approach, the convergence of the optimization method is preserved up to the tight tolerance $\underline{\tau}$ used to compute the objective functions, provided that a conservative error estimation is computed, i.e.

$$\lim_{i \to +\infty} \left\| \widetilde{f}_{j,i}(\boldsymbol{x}_i) - f_j^* \right\| = \underline{\tau} \tag{3.17}$$

Proof.
The proof stems directly from the definition of the conservative error estimation. In particular, for the j-th objective function $\left\|\widetilde{f}_{j,i} - f_j^*\right\|$ can be written as follows:

$$\begin{aligned}\left\|\widetilde{f}_{j,i} - f_j^*\right\| &= \left\|\widetilde{f}_{j,i} - f_j^* - f_j + f_j\right\| \\ &= \left\|(\widetilde{f}_{j,i} - f_j) + (f_j - f_j^*)\right\| \end{aligned} \tag{3.18}$$

Using the triangle inequality, $\|a + b\| \leq \|a\| + \|b\|$, we have

$$\begin{aligned}\left\|\widetilde{f}_{j,i} - f_j^*\right\| &= \left\|(\widetilde{f}_{j,i} - f_j) + (f_j - f_j^*)\right\| \\ &\leq \left\|\widetilde{f}_{j,i} - f_j\right\| + \left\|f_j - f_j^*\right\| \\ &\leq \varepsilon_{f_j} + \delta_{j,i}. \end{aligned} \tag{3.19}$$

By recalling the definition of conservative error estimation, the inequality also gives that

$$\left\|\widetilde{f}_{j,i} - f_j^*\right\| < \tilde{\varepsilon}_{f_j,i} + \delta_{j,i}. \tag{3.20}$$

By assumption $\delta_{j,i} \to 0$ as $i \to +\infty$. Since by construction of the algorithm, the conservative estimate $\tilde{\varepsilon}_{f_j,i}$ goes to the tight tolerance $\underline{\tau}$ (recall Definition 5), taking the limit yields

$$\lim_{i \to +\infty} \tilde{\varepsilon}_{f_j,i} + \delta_{j,i} = \underline{\tau}. \tag{3.21}$$

The previous result shows that the optimal solutions are known up to the tight tolerance $\underline{\tau}$. If a decoupled strategy and a loose tolerance $\bar{\tau}$ are used, the difference from the computed optimal solutions and the exact solutions may be significant. Instead, if a very tight tolerance is used, the discrepancy may be negligible. In addition, in the limit of $\tau \to 0$, Eq. (3.21) reads

$$\lim_{i \to +\infty} \lim_{\tau \to 0} \tilde{\varepsilon}_{f_j,i} + \delta_{j,i} = 0. \tag{3.22}$$

This means that if the tolerance in the estimation of the objective function is equal to zero, the computed optimal solutions converge to the exact solutions.

Finally, the order of convergence of the proposed strategy strictly depends on the convergence of the UQ method (the limit $\tau \to 0$) and the order of the optimization algorithm (the limit $i \to +\infty$).

3.2.4 Estimation of error bounds in the uncertainty quantification method

The Simplex Stochastic Collocation (SSC) is chosen for the UQ loop, because it relies on the adaptive refinement of the sampling in the stochastic space [93]. Other examples of methods in which an adaptive refinement could be used are Monte Carlo simulation methods, kriging-based or Polynomial-based approaches with nested quadratures.

As presented in the description of the SSC (see Sec. 2.3.2), there is a straightforward, quantitative way to relate the cost of the uncertainty propagation procedure and the accuracy of the result: the lower the tolerance τ on the error between function w and function v, the higher the cost and the better the estimates. On the other hand, a rough, cheap estimate on the statistics of the output v can be obtained with a coarser tolerance and fewer samples. To control the accuracy of the result, a measure of the error of the estimate is required. This measure of the error is the quantity that should be compared to the desired tolerance.

Different errors are computed in the SSC method, which are quantitative estimates of the error of the statistics (i.e. mean value and variance) due to the polynomial approximation w. In the following, three groups of errors are presented: Monte Carlo errors, estimated errors and conservative estimated errors. The first set serves as a reference for the assessment of the other estimates. Note that the errors on mean value and variance are considered here, but analogous estimates can be derived also for higher order statistics.

The first set is composed of the *Monte Carlo (MC) errors* , which are based only on the value of function w, and are obtained by means of random Monte Carlo sampling. The mean MC error ε_μ reads

$$\varepsilon_\mu = ||\mu_{mc} - \mu_w||, \tag{3.23}$$

where the mean μ_{mc} based on the exact function evaluation and the mean based on the approximation μ_w are computed by means of an MC integration over n_{mc}, as follows:

$$\begin{aligned} \mu_{mc} &= \frac{1}{n_{mc}} \sum_{i=1}^{n_{mc}} v(\boldsymbol{\xi}_i) \\ \mu_w &= \frac{1}{n_{mc}} \sum_{i=1}^{n_{mc}} w(\boldsymbol{\xi}_i). \end{aligned} \tag{3.24}$$

Similarly, the variance MC error is the following

$$\begin{aligned}\varepsilon_{\sigma^2} &= ||\sigma^2_{mc} - \sigma^2_w|| \\ \sigma^2_{mc} &= \frac{1}{n_{mc}} \sum_{i=1}^{n_{mc}} u^2(\boldsymbol{\xi}_i) - \mu^2_{mc} \\ \sigma^2_w &= \frac{1}{n_{mc}} \sum_{i=1}^{n_{mc}} w^2(\boldsymbol{\xi}_i) - \mu^2_w. \end{aligned} \tag{3.25}$$

These errors are useful to compare the function w and v with the same sampling technique, by ruling out the error associated with the MC method. As a result, these errors represent the accuracy of the polynomial approximation w. However they rely on the computation of the exact function in many random samples, thus they are not employed in practical cases where it is desirable to compute the exact function only for a small number of samples.

Therefore, the SSC provides a set of *estimated errors* that are based on the hierarchical surplus and the order of the polynomial employed in each element of the grid. In this case, the estimation of the error only relies on the (small) set of samples obtained in each vertex of the SSC grid. Let us define the hierarchical surplus for the new k-th sampling point $\boldsymbol{\xi}_{j,k}$ belonging to the j-th element (see Fig. 2.11a), as follows

$$\varepsilon_j = w(\boldsymbol{\xi}_{j,k}) - v(\boldsymbol{\xi}_{j,k})$$

where w is the interpolation before refinement and v is the exact function evaluation. Based on the hierarchical surplus, we define the estimated mean error $\hat{\varepsilon}_\mu$ as follows

$$\begin{aligned}\hat{\varepsilon}_\mu &= \sum_{j=1}^{n_e} \Omega_j \tilde{\varepsilon_j} \\ \tilde{\varepsilon}_j &= \frac{\varepsilon_j}{2^{\frac{p_j+1}{n_\xi}}}, \end{aligned} \tag{3.26}$$

where Ω_j is the probability contained in j-th simplex, p_j is the order of the polynomial approximation in the j-th simplex, and n_ξ is the number of the uncertain variables. The estimated variance error $\hat{\varepsilon}_{\sigma^2}$ then reads

$$\hat{\varepsilon}_{\sigma^2} = \left|\left| \sigma^2_w - \left[\frac{1}{n_{mc}} \sum_{i=1}^{n_{mc}} (w(\boldsymbol{\xi}_i) - \tilde{\varepsilon}_{j,i})^2 - \left(\frac{1}{n_{mc}} \sum_{i=1}^{n_{mc}} (w(\boldsymbol{\xi}_i) - \tilde{\varepsilon}_{j,i}) \right)^2 \right] \right|\right|,$$

where $\tilde{\varepsilon}_{j,i}$ is the corrected hierarchical surplus of the element which point i belongs to. This estimate is found by assuming that one could approximate the function $v(\boldsymbol{\xi}_i)$ by $w(\boldsymbol{\xi}_i) - \varepsilon_{j,i}$; this choice is conservative since typically the vertex based on which the surplus is computed is the one with the greatest error in each simplex. This error estimate works quite well for smooth functions when the interpolation reaches its theoretical convergence rate. However, in some cases it may be useful to adopt a more conservative estimate which is only based on the hierarchical surplus.

The *conservative estimated errors* are defined with a very similar structure with respect to the estimated errors, but they employ directly the hierarchical surplus ε_j, without using extrapolation based on h- and p-refinement. The conservative estimated mean error $\tilde{\varepsilon}_\mu$ is defined as

$$\tilde{\varepsilon}_\mu = \sum_{j=1}^{n_e} \Omega_j \varepsilon_j,$$

whereas the conservative estimated variance error $\hat{\varepsilon}_{\sigma^2}$ is

$$\tilde{\varepsilon}_{\sigma^2} = \left|\left| \sigma_w^2 - \left[\frac{1}{n_{mc}} \sum_{i=1}^{n_{mc}} (w(\boldsymbol{\xi}_i) - \varepsilon_{j,i})^2 - \left(\frac{1}{n_{mc}} \sum_{i=1}^{n_{mc}} (w(\boldsymbol{\xi}_i) - \varepsilon_{j,i}) \right)^2 \right] \right|\right|,$$

where $\varepsilon_{j,i}$ is the hierarchical surplus of the element which point i belongs to. To assess these estimates a comparison is presented for some analytical cases in the following subsection.

Analytic testcases for error estimation

To begin with, a one-dimensional sine function (see Fig. 3.10a) is considered:

$$u(\xi_1) = 3 + \sin(2\pi\xi_1) \tag{3.27}$$

where the uncertain variable ξ_1 is uniformly distributed in the domain $[0, 1]$. In Fig. 3.11a the behavior of the errors and estimators is presented, which confirms the effectiveness of the estimates. In fact, the estimated error and the conservative estimated errors are larger than the MC error, proving that they are conservative estimates of the actual error. In other words, the exact value of the statistics lies in the error bar defined around the SSC estimate and whose width can be described either by the estimated or the conservative error. In the case of the sine function, the error on the variance provides more information on the accuracy of the response approxi-

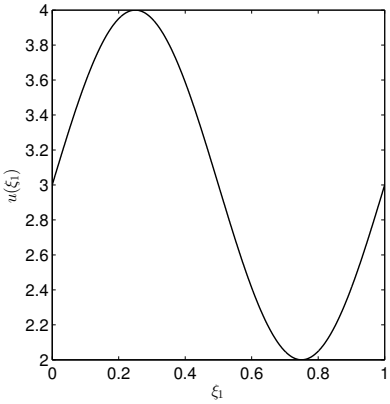

(a) One-dimensional sine function.

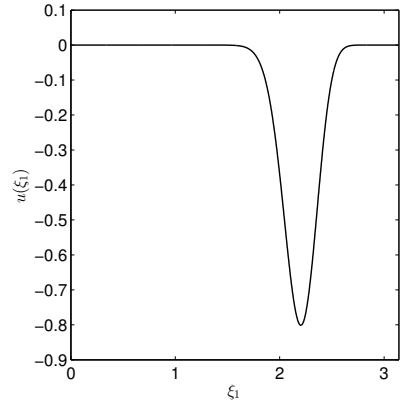

(b) One-dimensional Michalewicz function.

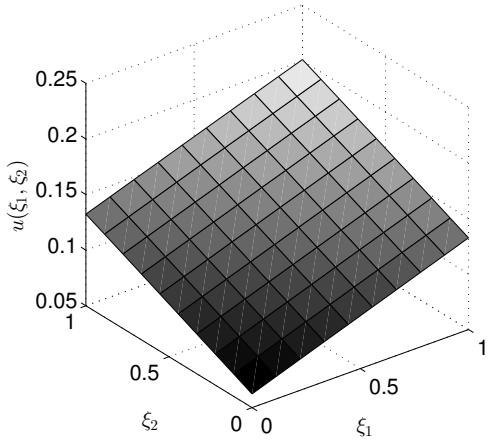

(c) Two-dimensional inverse tangent function.

Figure 3.10: Functions used to test SSC error estimates.

mation rather than the mean value, which is equal to the constant term in Eq. (3.27) and it is therefore affected by a small error even in the first iteration.

Then, the method is applied to a more complex function, the one-dimensional Michalewicz function (Fig. 3.10b), that is a very challenging testcase in global optimization problems

$$u(\xi_1) = -\sin(\xi_1)\ \sin^{20}\left(\frac{\xi_1^2}{\pi}\right). \tag{3.28}$$

The uncertain variable ξ_1 is again uniformly distributed in the domain $[0, \pi]$. Figure 3.11b shows the result of the application of SSC and proves that in this case the

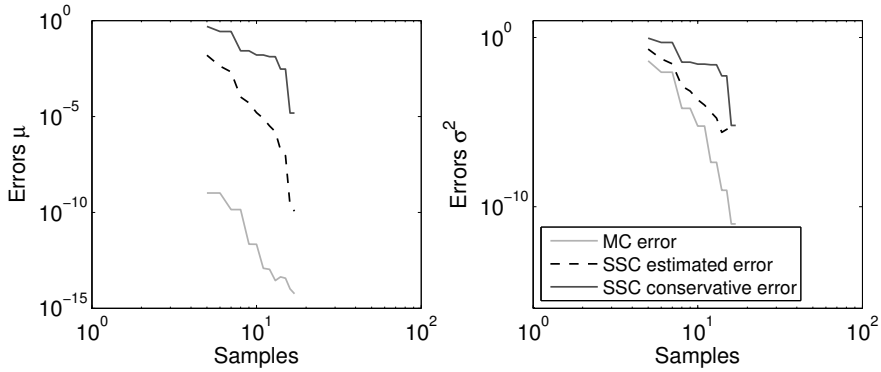

(a) One-dimensional sine function in Eq. (3.27).

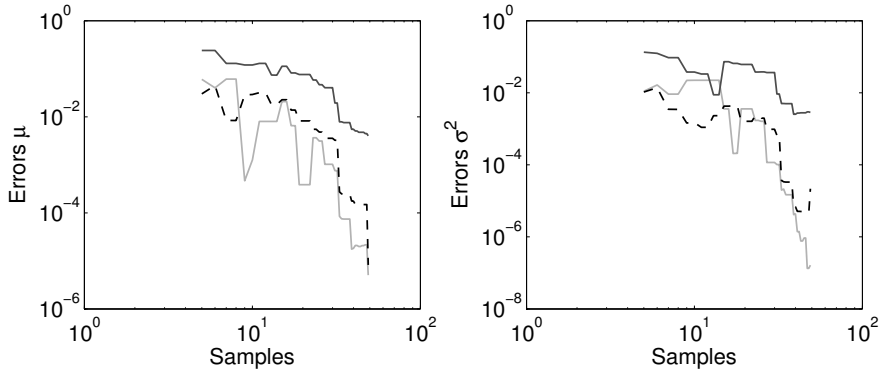

(b) One-dimensional Michalewicz function in Eq. (3.28).

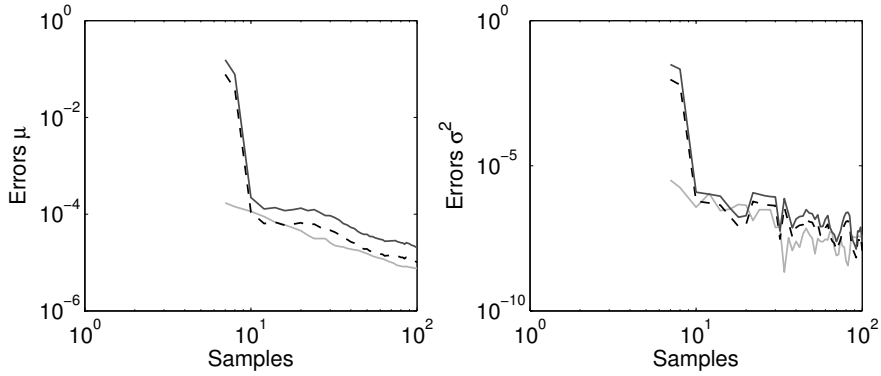

(c) Two-dimensional inverse tangent function in Eq. (3.29).

Figure 3.11: Monte Carlo, estimated and conservative errors for three analytic functions.

conservative error is a safer estimate of the error.

Finally, a two-dimensional function is considered to explore multi-dimensional problems. The inverse tangent presented in Ref. [93] is analysed (Fig. 3.10c), which reads

$$u(\boldsymbol{\xi}) = \text{atan}\left(\boldsymbol{\xi} \cdot \boldsymbol{\xi}^* + \xi_1^{*2}\right), \tag{3.29}$$

with $\boldsymbol{\xi}^* \in [0,1]^2$ is a two-dimensional vector of arbitrary values. In this case, the errors of the mean value have a smoother behavior with respect to the errors of the variance (see Fig. 3.11c). However, even in the latter case the convergence trend is clear and it appears that, despite the oscillations, the estimated error and the conservative error are consistent with the MC error.

As a final remark, these estimates are associated only to the error due to the polynomial approximation w. Although the method developed in this work considers a generic error on the estimate of the objective functions, regardless of the sources of that error, it is worth noting that these estimates are obtained by means of MC sampling, which is affected by an error itself. However, this error would decrease anyway when increasing the number of MC sampling points. Adding a bias to the estimates would be a possible strategy to include this error as well, without affecting the proposed optimization method.

3.2.5 Uncertainty-based optimization of algebraic test cases

This section presents some results obtained with the proposed adaptive strategy for two algebraic test cases, where there is a simple algebraic relationship between the performance u and the design and uncertain variables. In the comparison with the standard non-adaptive method, the proposed strategy is referred to as *coupled* strategy, and the classical approach is called *decoupled*.

Single-objective algebraic problem

In this section a single-objective ($m = 1$) optimization is considered, with two uncertain variables ($n_\xi = 2$) and two design variables ($n = 2$). The optimization problem reads

$$\begin{aligned} \text{minimize:} \quad & f(\boldsymbol{x}) = \mu(u) + \sigma(u) \\ \text{by changing:} \quad & \boldsymbol{x} \in \mathcal{X} = [0,6] \times [0,6] \\ \text{with:} \quad & \xi \sim \mathcal{U}(0,1) \end{aligned} \tag{3.30}$$

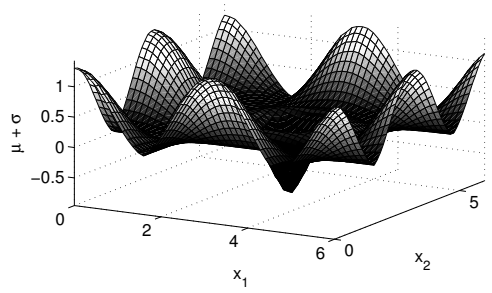

Figure 3.12: Objective of the optimization based on Yang function.

where the function representing the performance is the following

$$u(\boldsymbol{x},\xi) = \left[e^{-\sum_{i=1}^{n}(x_i/\kappa)^{2\gamma}} - 2e^{-\sum_{i=1}^{n}\sum_{j=1}^{n_\xi}\delta_{ij}\xi_j(x_i-\pi)^2}\right] \cdot \prod_{i=1}^{n}\cos^2 x_i \qquad (3.31)$$

with $\kappa = 15$, $\gamma = 5$ and δ_{ij} is the Kronecker delta ($\delta_{ij} = 1$ if $i = j$, $\delta_{ij} = 0$ if $i \neq j$). This function was presented in Ref. [116] and it is a modification to the classical test case for deterministic optimization. In this case, the uncertain variables are present in the exponential terms and they are as many as the design variables. The deterministic function has one global minimum $u(\boldsymbol{x}_{opt}) = -1$ in $\boldsymbol{x}_{opt} = \{\pi, ..., \pi\}$. When uncertain variables are considered the function changes its shape, but the global minimum remains the same for any realization of the uncertain variables. The mean value of u evaluated in the minimum is equal to -1 and the variance is equal to 0. Thus the solution of the robust case in Eq. (3.30) is equal to the deterministic global minimum $f(\boldsymbol{x}_{opt}) = -1$ in $\boldsymbol{x}_{opt} = \{\pi, ..., \pi\}$. The solution is close to the mid-point of the design variables domain, as represented in Fig. 3.12.

The solution of the optimization problem in Eq. (3.30) is obtained by means of the genetic algorithm described in Section. 2.2.2 and the SSC method (Section. 2.3.2). The coupling is obtained in three different ways: (i) with a decoupled approach in

Strategy	Percentual error on $f(\boldsymbol{x}^*)$	Distance from $\boldsymbol{x}^*$	N. function evaluations
Decoupled (1e-3)	1.0e-6	6e-4	24956
Decoupled (1e-1)	1.2e-5	2e-3	16820
Coupled	1.7e-7	4e-4	18875

Table 3.2: Comparison of results and cost between decoupled and coupled approaches for problem in Eq. (3.30)

which the tolerance is set to 1e-3, (ii) a decoupled approach in which the tolerance is set to 1e-1, (ii) the coupled strategy. In the latter case, the upper value $\bar{\tau}$ of the tolerance is set to 1e-1 and the lower value $\underline{\tau}$ is 1e-3. The three approaches reach the minimum value of the function in the correct minimum point, in spite of a slightly higher error for the decoupled approach with the loose tolerance (see Tab. 3.2). The computational cost of the decoupled strategy with tight tolerance is reduced by 24% with the coupled approach. The cost of the coupled method is comparable with that of the decoupled approach with loose tolerance, and the estimate of the minimum is very accurate. However, in this case a very good estimate of the optimal design is obtained even with a loose decoupled strategy. In the next section, a test case is presented which shows the real benefit of the coupled approach.

Multi-objective algebraic problem with NSGA

In this section the same bi-objective optimization problem tackled with the multi-fidelity strategy in Sec. 3.1.2 is solved (see Eq. (3.1)) by coupling the SSC method with the genetic algorithm.

Figure 3.13a presents the Pareto fronts of the analytic function and the fronts obtained by means of the decoupled method with tight tolerance, the decoupled method

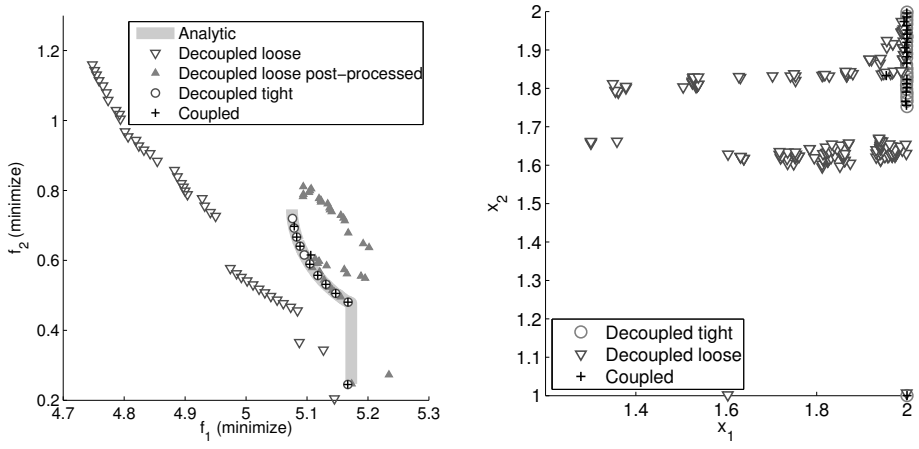

(a) Pareto fronts in the objective space. **(b)** Pareto fronts in the design space.

Figure 3.13: Results of multi-objective algebraic problem (Eq. (3.1)) with: (i) decoupled strategy with tight tolerance (1e-4), (ii) decoupled strategy with loose tolerance (1e-1) and (iii) coupled strategy. The loose tolerance result has been computed after the optimization with the analytical function for comparison.

with loose tolerance and the coupled strategy. The sets of solution obtained with the tight tolerance and the coupled approach compare well, while the loose tolerance results in a different branch of the front. This different branch obtained with the loose tolerance is a fictitious front, obtained only by the fact that the use of the loose tolerance does not provide an accurate estimate of the objective functions. Because of this, the front obtained with loose tolerance has been re-computed with the analytic function and plotted in the objective space for comparison. The solutions of the three fronts are plotted in the design space in Fig. 3.13b, where it is possible to note that the set obtained with the coupled approach overlaps the set obtained with the decoupled strategy with tight tolerance, while the method based on the loose tolerance end up on a different set of solutions that are not optimal.

The distance to the analytic Pareto front is presented in Fig. 3.14 and it is computed by means of Eq. (3.4). The comparison of the difference computed at each generation proves that the cheap optimization obtained with the loose tolerance results in an inaccurate front, while the coupled approach yields a front consistent with the accurate result. In addition, the approximation $w(\xi)$ obtained with the SSC method is presented in Fig. 3.15 for a representative design of the inaccurate Pareto front. In particular, design vector $\tilde{\boldsymbol{x}}$ is selected, which is closest to the mid-point between the two extremes of the Pareto front obtained with the loose decoupled approach. The SSC with the tight tolerance provides the best fit of the analytical function, whereas a good approximation is obtained with the coupled method. Instead, the number of samples used with the loose tolerance is not sufficient to correctly capture the behavior of the function. Thus, for this test case a tight tolerance is crucial to obtain the

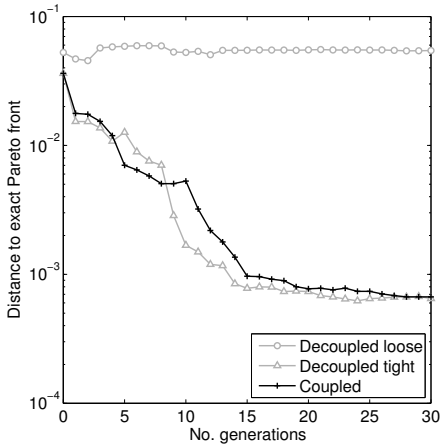

Figure 3.14: Distance to the Pareto front of the analytical function.

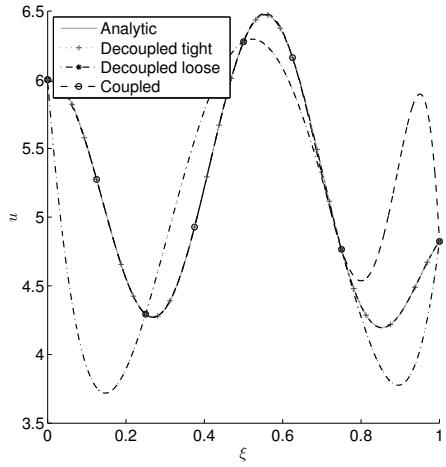

Figure 3.15: Approximations of the function u for the design vector $\tilde{\boldsymbol{x}}$ in the Pareto set obtained with the decoupled loose ($\bar{\tau}$ = 1e-1) method.

Strategy	Number of function evaluations		
	Initial generation	Evolution	Global
Decoupled (1e-1)	1077	5609	6686
Decoupled (1e-4)	4011	24482	28493
Coupled	1849	20368	22217

Table 3.3: Comparison of total number of function evaluations in the NSGA for the multi-objective algebraic optimization.

correct solutions of the optimization problem.

With regard to the computational cost of the coupled method, the number of function evaluations for each generation is shown in Fig. 3.16a and reported in Tab. 3.3. In the very first generation, the coupled approach has a significant reduction of the cost with respect to the tight decoupled approach: the reduction is close to 54%, thus the number of evaluations has been reduced to less than half the reference value. With increasing number of generations, the cost of the coupled approach increases and it almost converges to the cost of the tight decoupled method. In fact, with increasing number of generations during an evolution, each generation typically becomes more and more rich of individuals close to the current Pareto front, in other words design vectors with good objectives. For these "good" individuals the coupled approach has to compute the objective functions with a tighter tolerance, and the cost becomes higher. This is also shown in Fig. 3.16b, where the ratio of potentially non-

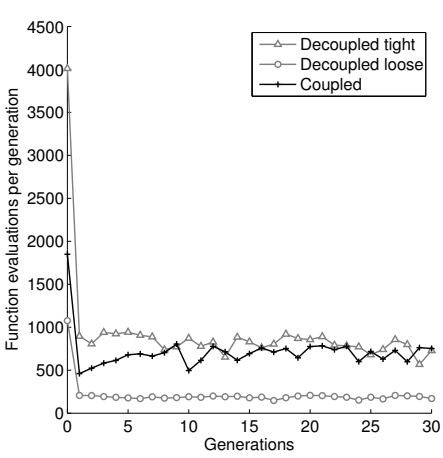

(a) Number of function evaluations.

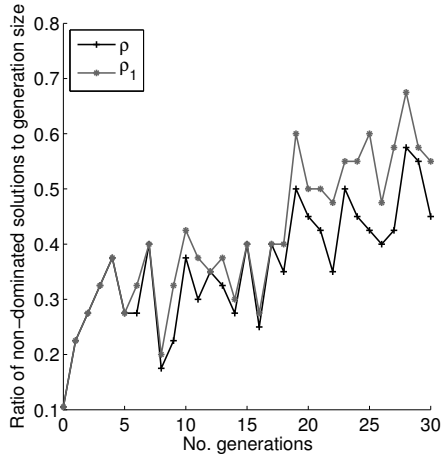

(b) Ratio of potentially non-dominated solutions to generation size.

Figure 3.16: Cost of the multi-objective algebraic optimization (Eq. (3.1)) with NSGA.

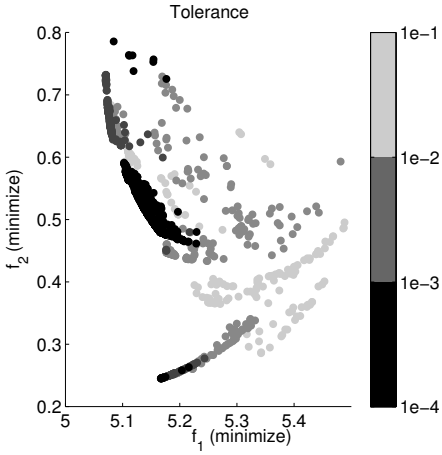

(a) Tolerance.

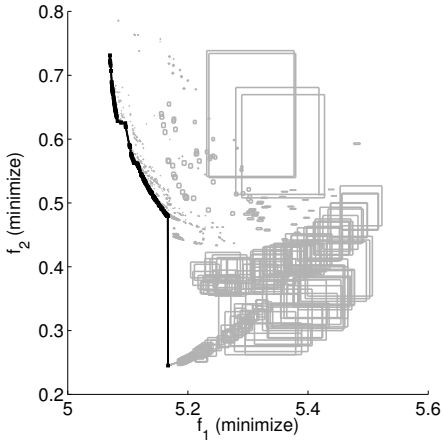

(b) Bounding boxes.

Figure 3.17: Tolerance and bounding boxes of the coupled approach for the multi-objective algebraic test case (Eq. (3.1)) with NSGA.

dominated solution over the generation size is presented. A solution is defined as potentially non-dominated if the objective vector is non-dominated with respect to the individuals computed up to that point. In mathematical terms, given the number of individuals in a generation N_G the ratio ρ is expressed as follows

$$\rho = \frac{\sum_{i=1}^{N_G} c_{k,i}}{N_G} \tag{3.32}$$

where the coefficient $c_{k,i}$ indicates whether the design vector $\boldsymbol{x}_{i,l}$ occurring at the l-th iteration in the current generation and i-th iteration in the global population is non-dominated regarding the population at the i-th iteration $\mathcal{X}_i$

$$c_{i,l} = \begin{cases} 1 & \text{if } \boldsymbol{x}_{i,l} \text{ is non-dominated regarding } \mathcal{X}_i \\ 0 & \text{otherwise.} \end{cases}$$

It appears that when the ratio exceeds 0.3, the cost of the coupled approach becomes similar to the cost of the tight decoupled approach. Nevertheless, the cost remains lower and the global cost is reduced by 22 % (see Tab. 3.3). Please note that the ratio ρ is a lower bound of the cost of the decoupled strategy, because it is computed in a post-processing stage with the non-dominated solutions based on the final estimate. Therefore, it does not include the refinement required for dominated solutions close to the Pareto front. To highlight this aspect, Figure 3.16b presents the modified ratio ρ_1, computed taken into account not only the non-dominated solutions, but also the solutions within a distance to the current Pareto front of $\underline{\tau}$ in each direction of the objective space. In this case the two ratios have similar trend.

Figures 3.17a presents the tolerance τ used for each design vector in the objective space. During the optimization, the algorithm uses the tighter tolerance $\underline{\tau}$ = 1e-4 only in the region close to the actual Pareto front. The loose tolerance $\bar{\tau}$ = 1e-1 is used for higher values of the mean value f_1. As a matter of fact, the resulting error bounding boxes are wider in that area, as presented in Fig. 3.17b.

Discussion on error estimates During the post-processing stage of the optimization, the error with increasing number of samples has been computed for a group of design vectors. For instance, in Fig. 3.18 the MC errors, estimated and conservative errors are plotted for the design $\tilde{\boldsymbol{x}}$, a trade-off design belonging to the Pareto front obtained with the decoupled loose method. The figure shows that in the first iterations of the SSC method the estimated errors is not conservative with respect to the MC error (which represents the reference value, as explained in Sec. 3.2.4).

In light of this result, the robust optimization has been performed again considering the conservative estimated errors affecting the objective functions. The resulting Pareto fronts are presented in Fig. 3.19a, which shows a situation similar to the one presented in the preceding section. The coupled approach results in a front overlap-

ping the result of the decoupled method with tight tolerance. Once again, the loose tolerance does not provide accurate results. The error bounding boxes computed in the coupled approach are larger in this case, due to the higher order of magnitude of the conservative estimated errors. The computational cost of the coupled approach in this case is higher and it approaches the result of the tight tolerance (see Fig. 3.20a).

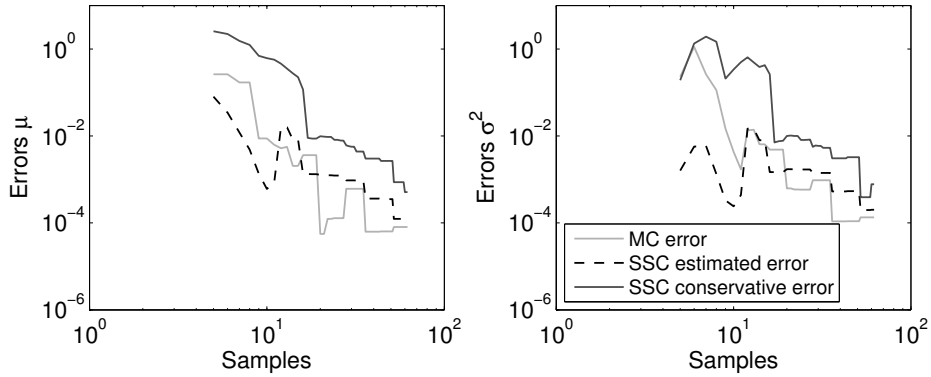

Figure 3.18: Errors affecting the mean value and the variance of function u for a representative design vector of the decoupled loose Pareto front.

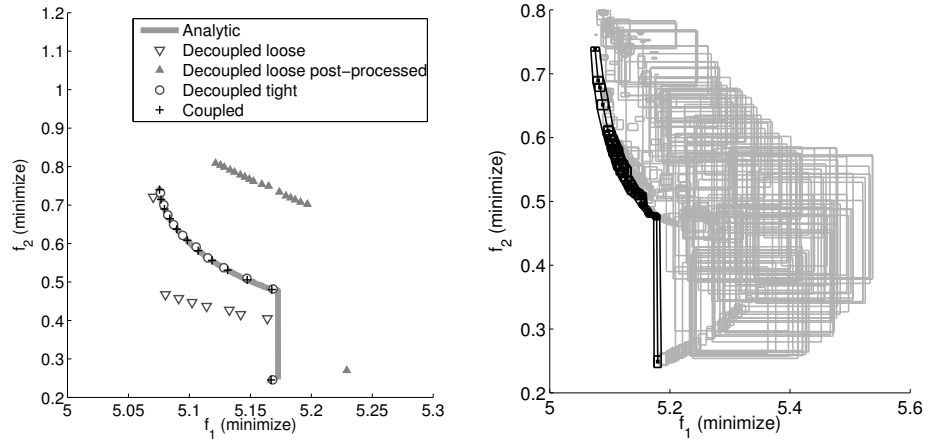

(a) Pareto optimal fronts: the loose tolerance result has been computed after the optimization with the analytical function for comparison.

(b) Error bounding boxes in the coupled method.

Figure 3.19: Results of multi-objective algebraic problem (Eq. (3.1)) with conservative estimates.

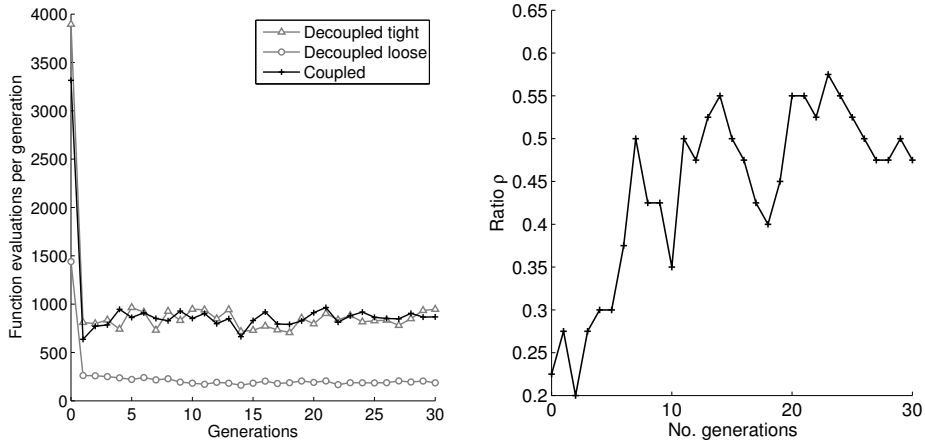

(a) Number of function evaluations.

(b) Ratio of potentially non-dominated solutions to generation size.

Figure 3.20: Cost of the multi-objective algebraic optimization (Eq. (3.1)) with NSGA and conservative estimates.

This is due to two main reasons. First, the ratio of potentially non-dominated solutions over generation size ρ is greater in this case as presented in Fig. 3.20b. Second, since the conservative estimates overestimate the actual errors, the bounding box is greater than the box that could be defined with the actual errors (see the boxes in Fig. 3.19b). As a result, it is more likely for the bounding box to dominate the Pareto front, thereby requiring a higher number of samples even for the design vectors that are far from the Pareto front.

The results obtained with the estimated and conservative estimated errors highlight two important aspects. First, even if the error estimates in the first iterations of the SSC are not strictly conservative, the proposed strategy is able to attain the accurate Pareto front as presented in the preceding section. In fact, for the non-dominated solutions the method always uses the lowest tolerance $\underline{\tau}$, overcoming the problems that may arise with a non-conservative error estimates in the first iterations. Thus, the method is robust and it is able to assess whether a solution that may appear non-dominated due to a poor accuracy of the objective function is actually dominated. Nevertheless, conservative estimates are important in the opposite situation, that is to avoid discarding potentially non-dominated solutions. Secondly, not only it is crucial to have an error estimate that is conservative, but also the estimate should be consistent with the actual error. A too conservative estimate may jeopardize the efficiency of the adaptive strategy.

With regard to the estimate provided by the SSC, a theoretical proof is given in Ref. [93]. Based on the experience gained with different functions considered throughout the thesis the robustness of the error estimates turns out to be very problem-dependent. The theoretical convergence is reliable as long as the function is smooth and as long as the initial sampling is sufficiently capturing the function variability. The variability of the function in an optimization problem not only depends on the uncertain variables, but also on the design variables, so tweaking the initial sampling is problem-specific and the effect of the sampling should be assessed for representative design before the optimization run.

Multi-objective algebraic problem coupled method with BIMADS

The multi-objective robust optimization problem presented in Sec. 3.1.2 (Eq. (3.1)) is tackled also with the BIMADS algorithm coupled to the SSC method. The results are presented in Fig. 3.21a, where the front obtained with the decoupled and coupled approaches are compared to the result of the analytic function. Similarly to the case presented in the preceding subsection, the distance computed with respect to the analytic result in Fig. 3.21b proves that the decoupled method with tight tolerance and the coupled method reaches the analytical result. Furthermore, the convergence is

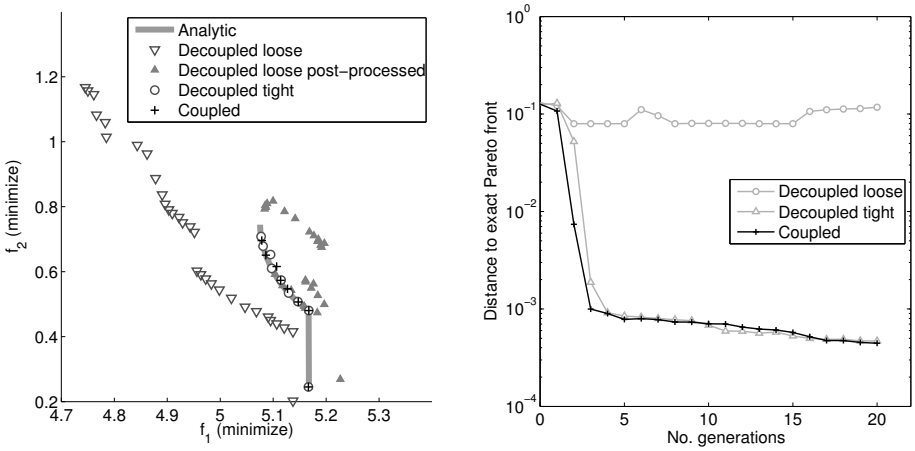

(a) Pareto optimal fronts: the loose tolerance result has been computed after the optimization with the analytical function for comparison.

(b) Distance to Pareto front of the analytical function.

Figure 3.21: Results of the multi-objective algebraic test case (Eq. (3.1)) with BIMADS.

Strategy	Decoupled (1e-1)	Decoupled (1e-4)	Coupled
Number of function evaluations	5331	26387	23894

Table 3.4: Comparison of total number of function evaluations in the BIMADS optimization loop.

reached in almost five iterations of the algorithm.

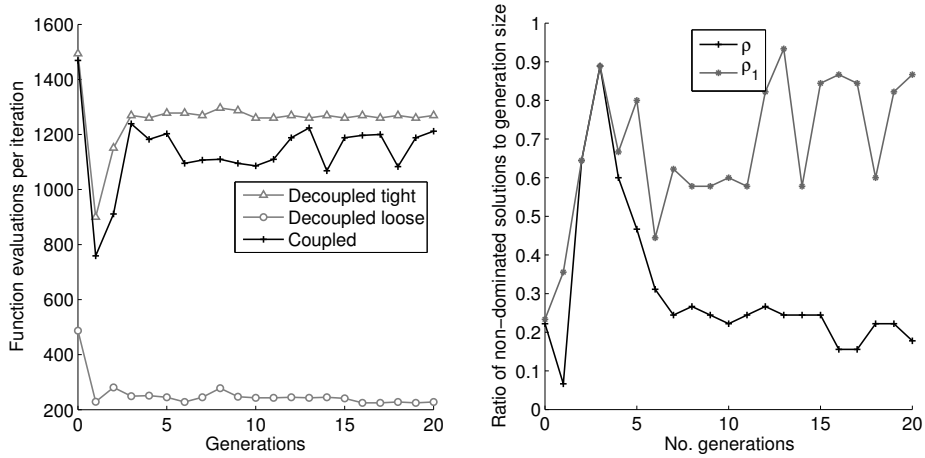

(a) Number of function evaluations.

(b) Ratio of potentially non-dominated solutions to generation size.

Figure 3.22: Cost of the multi-objective optimization with BIMADS.

The cost of the decoupled approach with tight tolerance is reduced by 30% with the coupled approach in the first iteration (see Fig. 3.22a). In the following iterations the coupled approach becomes more expensive and the cost gets closer to the cost of the tight decoupled method, although it remains lower. The reduction of the global cost is about 10% in this case. In fact, in the first iterations many design vectors close to the Pareto front are computed: as reported in Fig. 3.22b over 70% of the design vectors are potentially non-dominated solutions. In addition, during the remaining iterations, the design vectors tested are very close to the Pareto front, so a tighter tolerance is required in those cases. On average, the coupled approach in the last iterations uses 25 evaluations per design vector, whereas the decoupled tight approach uses 28 function evaluations.

With respect to the case performed with the NSGA, the coupled approach employs more function evaluations with the solutions computed in the BiMADS loop.

In fact, among the design vectors tested in the NSGA loop, the number of solutions which are "good", i.e. closer to the Pareto front, is greater in the BiMADS case. This is shown by the ratio ρ_1 in Fig. 3.22b, which clearly presents that the solutions within a distance to the current Pareto front equal equal to the tight tolerance in each direction of the objective space are dominant in the BiMADS iteration. This explains the higher cost of the coupled strategy with BiMADS.

3.3 Remarks on the methods

In the present Chapter two methods to improve the efficiency of robust optimization have been presented. In the algebraic test cases presented in the previous sections, both the multi-fidelity and adaptive methods provide a significant gain with respect to the traditional robust optimization methods. However, the gain is obtained in very different ways. The first approach leverages models with different fidelity (and cost) to improve the efficiency of the robust optimization loop. Its gain is somewhat associated with the difference in the cost of the chosen high- and low-fidelity model. Furthermore, the availability of two methods to model the system under consideration, which share (almost) the same accuracy in a part of the stochastic domain is a key element. On the other hand, the adaptive strategy can be used with a single model and it uses the information shared by the external optimization loop and the internal UQ loop.

A natural follow-up study would consist in the integration of the two methods. In fact, these act on two different levels: the multi-fidelity method plays with the models, the adaptive strategy with the tolerance values. The resulting adaptive, multi-fidelity strategy could provide even higher computational gain. This study is left for future activities and it was not investigated in the present thesis because time had to be devoted to numerous and expensive application of helicopter rotor airfoils. The hovering condition is considered in the next chapter.

CHAPTER 4

Robust optimization of fixed airfoil in hovering condition

The robust optimization techniques presented in the preceding chapter are applied to the optimal selection of the airfoil for a helicopter rotor blade in hover.

4.1 Optimization problem

Ideally the hover condition is a perfectly symmetric and steady flow condition when an isolated rotor is considered. However, when a rotor is installed on the helicopter (see Fig. 1.1) the interference with the other parts and the necessity to trim the entire aircraft requires the introduction of some cyclic pitch commands that lead to variable flow conditions met by the blade section in different azimuth positions. As a consequence, instead of considering a deterministic, representative operating condition for hover, it is assumed that the freestream condition of the airfoil section is affected by some level of uncertainty. In this thesis we deal with sources of uncertainty in the physical modelling employed at the design stage. This uncertainty could arise from errors in the model representation (e.g. the blade inflow model). Additionally, other physical sources of uncertainty may be considered, such as those related to blade structural flexibility, that may modify the local pitch angle. Because in the design stage the blade pitch angle θ and the induced velocity V_i dictate the operating condition of the airfoil (Fig. 4.1), uncertainties on these two system parameters are considered in the present analysis. The induced velocity is also a freestream boundary condition that permits consideration of the effect of the wake vorticity of the blade

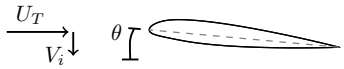

Figure 4.1: Scheme of operating conditions at a blade section: the uncertain variables are the pitch angle θ and the induced velocity V_i. The velocity component U_T is the velocity tangential to the rotor disk.

on the airfoil in a simple yet effective way [9].

The objective of the aerodynamic optimization is to maximize the lift-to-drag ratio E, which is a measure of the aerodynamic efficiency of the airfoil. This measure is chosen because improving the lift-to-drag ratio in the airfoil selection benefits the rotor performance [11]. The objective of this analysis is mainly to find a robust optimal airfoil shape. In other words, the optimal design must maximize the mean value μ_E and minimize the variance σ_E^2 of the objective function, computed by taking into account the uncertainties of the operating conditions. Unlike single-objective optimization where there is only one global optimum, in the case of multi-objective optimization, as the one under consideration here, there is a set of optimal solutions (the so-called Pareto front) corresponding to various trade-offs between the mean value and the variance.

The objective function E is a function of the output of the aerodynamic system, that is the solution of the differential equation modelling the aerodynamics:

$$\mathcal{L}\left(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi}; \phi(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi})\right) = 0, \tag{4.1}$$

where $\mathcal{L}$ is a nonlinear spatial differential operator describing the problem. The solution of the stochastic equation (4.1) is the vector $\phi(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi})$, which contains the unknown variables, e.g. density, momentum and total energy. The solution vector ϕ is a function of the two-dimensional space variable $\mathbf{y} \in \mathbb{R}^2$, the design variables $\mathbf{x}$ belonging to the design space $\mathcal{X}$, and the uncertain input parameters $\boldsymbol{\xi}$ belonging to the stochastic space Ξ.

The differential equation (4.1) represents a constraint for the optimization problem. Additional constraints are related to manufacturing feasibility and costs. The geometry of the airfoil is constrained in such a way that only airfoils having at most one inflection point of the camber line are allowed, and the pitching moment coefficient computed at the quarter-chord location is constrained to prevent excessive torque on the blade structure. The latter constraint is applied when the moment coefficient about the quarter-chord location exceeds 0.03. Despite the importance of

minimizing the pitching moment, such a value is chosen to explore the design space with sufficient freedom.

In mathematical terms, the resulting optimization problem can be stated as:

$$\begin{aligned} \text{maximize:} \quad & \mu_E\left(\phi\left(\mathbf{x},\mathbf{y},\boldsymbol{\xi}\right)\right) \\ \text{and minimize:} \quad & \sigma_E^2\left(\phi\left(\mathbf{x},\mathbf{y},\boldsymbol{\xi}\right)\right) \\ \text{subject to:} \quad & \mathcal{L}\left(\mathbf{x},\mathbf{y},\boldsymbol{\xi};\phi(\mathbf{x},\mathbf{y},\boldsymbol{\xi})\right) = 0 \\ & \mathbf{g}\left(\mathbf{x},\mathbf{y},\boldsymbol{\xi};\phi(\mathbf{x},\mathbf{y},\boldsymbol{\xi})\right) \leq \mathbf{0} \\ \text{by changing:} \quad & \mathbf{x} \end{aligned} \tag{4.2}$$

with uncertain input parameters $\boldsymbol{\xi} = \{\theta, V_i\}$ and under the constraints previously described and collected in vector $\mathbf{g}$. The robust optimization problem requires several ingredients (Fig: 1.10): (i) a model of the aerodynamic system which includes the aerodynamic model ($\mathcal{L} = 0$) which describes the flowfield around the airfoil, the freestream boundary conditions and the airfoil shape, (ii) an uncertainty quantification method that propagates the uncertainty affecting the freestream operating conditions to the uncertainty on the output, i.e. the aerodynamic loads computed by means of the aerodynamic model , and (iii) a shape optimization method that seeks to find the best design according to the objectives and contraints set in Eq. (4.2).

4.2 Numerical ingredients of robust optimization method

The different elements of robust optimization have been discussed in Chapter 2. In this section, the specific methods used for the robust aerodynamic optimization presented in the preceding section are outlined.

4.2.1 Uncertain operating conditions

In this work, a reference, or nominal, condition is first chosen which represents a typical condition for a section of a blade in hover. Then, an estimate of the uncertainty on each variable is computed. For the nominal condition, data from Ref. [117] are employed, considering two blade sections respectively at 40% and 80% of the span (see Tab. 4.1). The blade pitch angle and induced velocity are computed by means of the Blade Element Momentum Theory (BEMT) starting from data of the Bo105 rotor and assuming a radial distribution of inflow on the blades [118, p. 127]. Regarding the uncertain parameters, a variation of ± 2 deg is estimated for the blade pitch angle due to longitudinal cyclic and lateral cyclic pitch, and a value of ± 0.5 deg is added

in order to model the effect of blade flexibility with an uncertainty in the torsional stiffness. This level of uncertainty has also been found on the pitch measures taken during the HART II experimental campaign [117]. For the induced velocity, a variation of $\pm 10\%$ of the nominal value is considered due to the modeling of physical parameters of the system; such a level of uncertainty has been assessed in Ref. [35] as a result of uncertainty on both aerodynamic and structural uncertainties such as blade chord, air density, and rotor angular velocity. Furthermore, a $\pm 5\%$ is added based on empirical considerations to approximate the effects of the losses due to the wakes trailed from the blades and other three-dimensional effects of the rotor which are not yet known in the design stage of the blade.

In addition, the uncertain variables $\boldsymbol{\xi} = \{\theta, V_i\}$ are modelled in a probabilistic framework, meaning a probability density function is assigned to each variable. In particular, a uniform PDF is adopted (see Fig. 4.2a) in order to consider the worst-case scenario among the probabilistic descriptions. That is the one where each event in the uncertainty range is just as probable as the nominal condition (see Tab. 4.2). Another modelling assumption is that the uncertain variables are independent, which may not be true to some extent due to the connection between the rotor induced flow and the blade control. However, because the complete aero-servo-elastic system is not under examination but only the aerodynamic subsystem, the two variables are considered as separate and independent.

The uncertainties of the induced velocity V_i and the blade pitch angle θ affect the conditions of the flowfield around the airfoil. In particular, its influence on the airfoil angle of attack is the most significant, while minor effects are noticed on the Mach and Reynolds numbers. The level of uncertainty considered for each variable is consistent with those typically employed in the literature of uncertainty-based analysis in fluid dynamics (e.g. Refs.[119, 120, 121]). Their combined effect results in a large variation of the angle of attack once it is propagated through the BEMT equations. Such a variation is up to 60%, as shown in Fig. 4.2b for the section at 80% of the radius. The angle of attack has a fundamental role on the onset of bound-

r [-]	V_i [m/s]	θ [deg]	α [deg]	M [-]	Re [-]
0.4	10.52	10.11	6.0	0.258	1.6e6
0.5	12.1	11.7	5.9	0.32	2e6
0.8	13.82	9.71	5.2	0.513	3.2e6

Table 4.1: Nominal operating conditions.

Variable	Uncertainty	PDF type
V_i	$\pm$ 15%	Uniform
θ	$\pm$ 2.5 deg	Uniform

Table 4.2: Uncertainty on blade pitch angle θ and induced velocity V_i.

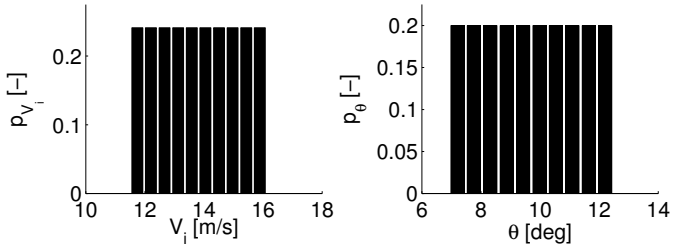

(a) Input uncertainty on induced velocity V_i and blade pitch angle θ.

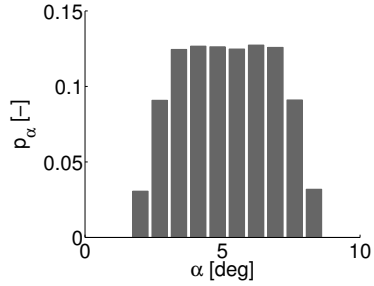

(b) Resulting uncertainty on the angle of attack α .

Figure 4.2: Uncertainty on the operating conditions (section at $r = 0.8$).

ary layer instability, the occurrence of large flow separation and the development of complex compressibility effects. In this case, taking uncertainty into account yields large variations of the flowfield characteristics, and the most significant effect is associated with compressibility. To show the above in a straightforward manner, the critical Mach number is computed at different angles of attack for the NACA 0012 airfoil by combining the isentropic one-dimensional flow equation for pressure with either the Prandtl-Glauert or the Karman-Tsien correction rule for the pressure coefficient [122, p. 325]. The result is plotted in Fig. 4.3a, where the variablity range of the flowfield parameters, angle of attack and Mach number, is also presented for different blade sections. It appears that for sections located at a distance from the root

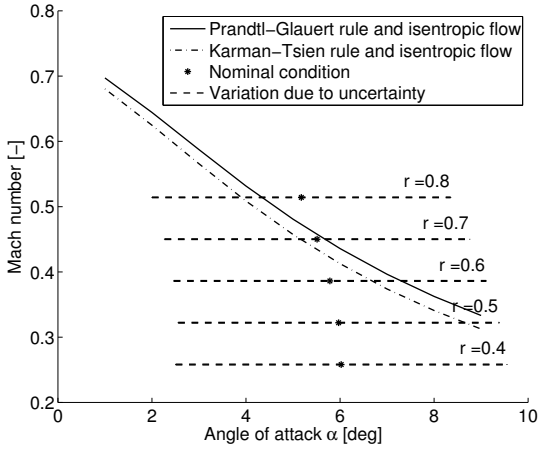

(a) Critical Mach number and uncertain operating conditions.

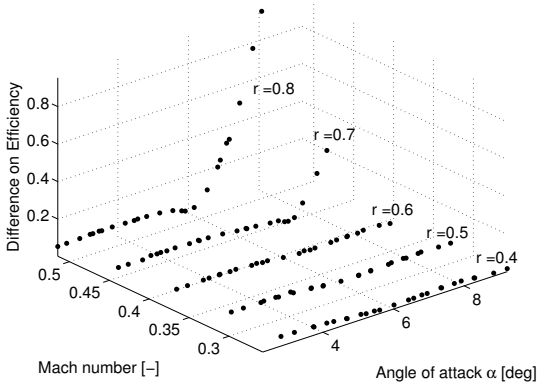

(b) Relative difference in the estimates of the lift-to-drag ratio between high and low fidelity models.

Figure 4.3: Analysis of the effect of uncertain operating conditions on the critical Mach number (a) and on the estimates of the airfoil lift-to-drag ratio (b) fro the NACA 0012 airfoil.

greater than 60% of the radius, the variability range crosses the critical Mach curve, which represents the onset of significant compressibility effects.

4.2.2 Aerodynamic models

As presented in the preceding section, computations both in the subsonic and low transonic range are required to reconstruct the mean and the variance of the performance. In transonic conditions, only aerodynamic solvers dealing with nonlinear compressible effects yield meaningful estimates. However, the computational cost

increases when using compressible aerodynamic models. Under the critical Mach curve, instead, cheaper computations can be performed with a subsonic, linear model. As a result, two aerodynamic models are considered, which corresponds to different level of fidelity: a low-fidelity model suited for flow past an airfoil in subsonic flow, and a high-fidelity model for (low) transonic conditions.

Low-fidelity model

The low-fidelity solver is XFOIL, an aerodynamic solver which couples panel and integral boundary layer methods developed for the analysis of subsonic, isolated airfoils [99]. The outer flowfield is modelled with an inviscid linear vorticity panel-based method, which is coupled to the viscous layers represented by the two-equation lagged dissipation integral method in order to correct the potential flow with the viscous boundary layer. Here, the transition of the boundary layer is predicted by applying the e^n criterion, while compressible effects for low Mach numbers are taken into account by the Karman-Tsien correction. This model is well suited for the simulation of flows around an airfoil at moderate values of the angle of attack and Mach number, and it permits a rapid evaluation of the aerodynamic coefficients.

High-fidelity model

The high-fidelity solver is MSES [123], a coupled Euler/integral boundary layer code for the analysis and design of multi-element or single-element airfoils. It couples the solution of the Euler equations implemented with a streamline-based discretization for the outer field with a two-equation integral boundary layer formulation for the thin region close to the airfoil. The coupling is performed by means of the displacement thickness and the system equations are solved simultaneously by a full Newton method. The prediction of the boundary layer transition from laminar to turbulent flow is carried out either with the full e^n method [124] or with the simplified envelope method [125]. Thus, the MSES solver is capable of representing the nonlinear compressibility effects arising in the transonic flow regime.

Comparison of aerodynamic models

Although the two models share a similar modelling of the boundary layer, the high-fidelity solver is capable of representing the nonlinear compressibility effects arising in the low transonic flow regime. In fact, once the critical Mach number is exceeded, the wave contribution to the global drag coefficient is no longer negligible. As a

result, the difference in the quantity of interest, i.e. the aerodynamic efficiency, becomes important. This is presented in Fig. 4.3b, where the relative error between the estimates of the two solvers is plotted in the range of Mach number and angle of attack dictated by the uncertainty on the induced velocity and blade pitch angle. Finally, the computational cost associated with the low-fidelity model is lower than the cost of the high-fidelity model: on average, the XFOIL solver is ten times faster than MSES.

4.2.3 Shape parameterization

The shape of the airfoil is the design output of the optimization problem. To define a finite, sufficiently small set of design variables, a parameterization is required. In this work, the Class/Shape function Transformation (CST) [126] is employed to describe the airfoil shape.

Class/Shape function Transformation

The parameterization is well-defined by specifying two functions: a geometry class function C and a shape function S that defines the particular shape of the geometry. The former is defined as follows

$$C(\chi) = \chi^{N_1}(1-\chi)^{N_2}, \tag{4.3}$$

where χ is the non-dimensional chordwise coordinate ($\chi = x/c$) and N_1 and N_2 are appropriate coefficients that determine the class of the geometry; in the case of an airfoil with a rounded nose and a sharp trailing edge: $N_1 = 0.5$ and $N_2 = 1$ (see Ref. [126] for further details). The shape function $S(\chi)$ is given by a Bernstein polynomial of order N. The CST is applied to the camber mean-line ζ_c and to the thickness distribution ζ_t of the airfoil (both obtained by projection in the direction orthogonal to the x-axis) and it reads

$$\begin{aligned}
\zeta_c(\chi) &= C(\chi)\ S(\chi) + \chi\ \zeta_{TE_c} = C(\chi) \cdot \sum_{i=0}^{n} A_i^c S_{N,i}(\chi) + \chi\ \zeta_{TE_c} \\
\zeta_t(\chi) &= C(\chi)\ S(\chi) + \chi\ \zeta_{TE_t} = C(\chi) \cdot \sum_{i=0}^{n} A_i^t S_{N,i}(\chi)
\end{aligned} \tag{4.4}$$

where the coefficients A_i^c, A_i^t define the Bernstein polynomial coefficients of the camber line and thickness distributions respectively, $S_{n,i}$ is the i-th term of the Bernstein

polynomial of order N defined as:

$$S_{N,i} = \frac{N!}{i!(N-i)!}\,\chi^i \cdot (1-\chi)^{N-i}. \tag{4.5}$$

The coordinate of the trailing edge of the mean camberline is ζ_{TE_c}, whereas the coordinate of the trailing edge of the thickness distribution ζ_{TE_t} is set to zero in order to consider closed trailing edge airfoils. Finally, from the meanline and camber distribution the upper surface ζ_u and lower surface ζ_l are computed as follows

$$\zeta_u(\chi) = \zeta_c(\chi) + \frac{1}{2}\zeta_t(\chi), \qquad \zeta_l(\chi) = \zeta_c(\chi) - \frac{1}{2}\zeta_t(\chi). \tag{4.6}$$

Convergence analysis of the CST

The order of the polynomial is chosen based on a convergence analysis of the CST. To do this, the parameterization has been applied with an increasing polynomial order to represent a group of helicopter rotor airfoils (e.g. NACA23012, SC1095, SC1094, HH02). A least squares fitting is employed to determine the coefficients of the CST which best represents each airfoil. The mean squared error MSE is computed as follows

$$MSE = \frac{1}{n_p}\sum_{i=1}^{n_p}\left(\tilde{\zeta}_i - \zeta_{CST_i}\right)^2,$$

where n_p is the number of points on the airfoil surface, $\tilde{\zeta}_i$ is the exact airfoil and ζ_{CST_i} is the CST representation, both evaluated at the i-th point. The analysis proves that with a polynomial of 7-th order the mean quadratic error of the CST reaches a value equal to or lower than a desired threshold set to 10^{-7}. A representative set of results is shown for the airfoil SC1094 R8 in Fig. 4.4. Thus, the order of the polynomial is set to seven and the resulting design variables are 17: $\mathbf{y} = \{A_i^c, A_i^t, \zeta_{TE_c}\}$ with $i = 0, ..., 7$.

4.2.4 Optimization method

In the present work, the optimizer is based on the Nondominated Sorting Genetic Algorithm (NSGA) [67]. Main tuning parameters of the algorithm are the population size, the number of generations, the crossover and mutation probabilities pc, pm and the so-called sharing parameter r used to take into account the relative isolation of an individual along a dominance front. Typical values for pc, pm are, respectively, 0.9

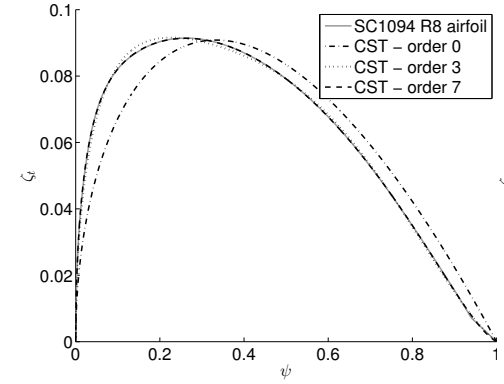

(a) Thickness distribution ζ_t.

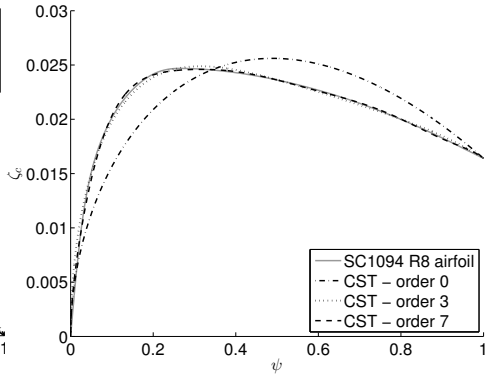

(b) Camber distribution ζ_c.

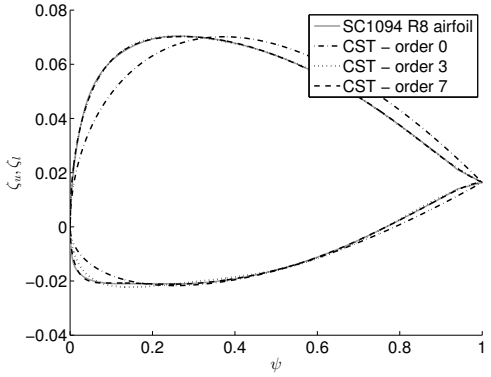

(c) Resulting upper surface ζ_u and lower ζ_l surface.

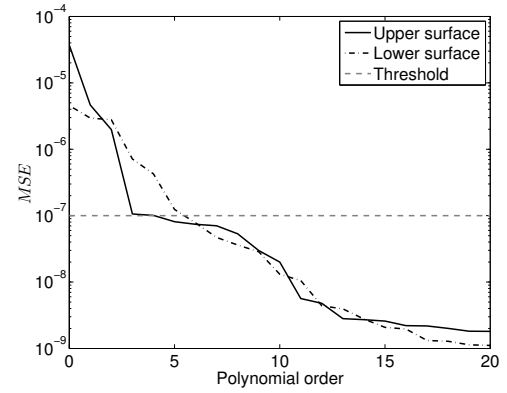

(d) Mean squared error MSE for upper and lower surfaces with increasing polynomial order.

Figure 4.4: Results of the application of the CST for the SC1094 R8 airfoil: each dash-dotted line in (a), (b) and (c) corresponds to a given polynomial order.

and 0.1; values of r are defined following a formula given in Ref. [67] that takes into account the population size and the number of objectives.

The constraint on the pitching moment is obtained by means of a penalty function that is applied to the lift-to-drag ratio when the moment coefficient exceeds the maximum value reported in Section 4.1. The penalty function is static, with a fixed value of the penalty function set to a couple of order of magnitude greater than any feasible lift-to-drag ratio. In the same way, another geometric constraint is added in order to disregard geometries which are selected by the evolution of the generations, but possess many inflections points. This strategy is used to avoid computing a CFD solution on a design that is not feasible from a geometric point of view.

With regard to the convergence of the optimization, several empirical approaches exist for evolutionary algorithms. For single-objective optimization, a possible criterion to assess convergence of the solution is tracking the difference between the mean value of the objective function in the population and the minimum value (for a minimization problem) of the objective function. In fact, when the minimal and the mean value of the fitness function in the population converge to the same asymptotic value, the population is considered to converge to a set of almost identical (optimal) individuals. For multi-dimensional cases, measures for assessing convergence include the global number of dominated solutions per generation, the diversity or spread of the individuals in the design space and the domination rate [127, 128]. The latter measure represents the percentage of non-dominated individuals that become dominated in the following generation, and should then decrease as the population evolves.

Finally, the optimization based on genetic algorithms typically starts with a first generation that constitutes a Design Of Experiment (DOE), that is the evaluation of the objective function by sampling in the design space to capture the variation of the function with respect to the independent variables. The designs involved in the initial DOE have been selected in order to satisfy the geometric constraint; this step is of paramount importance to avoid the excessive occurrence of undesirable designs in the following generations. Because of the structure of genetic algorithms, it is convenient to implement the algorithm in parallel mode, so that each CFD computation of an individual in a generation is distributed to the available processor. In this case, the algorithm is implemented using MPI libraries.

The optimization loop based on the NSGA and the CST parameterization has been tested for a benchmark case proposed by the AIAA Aerodynamic Design Optimization Discussion Group [129]. The results are presented in Ref. [130] and reported here in Appendix A.

4.2.5 Uncertainty quantification method

Before running the robust optimization, a preliminary convergence study in the stochastic space is performed to determine the performance of the UQ methods and the number of deterministic computations required for each design run. Three methods for non-intrusive UQ are selected for comparison: (i) QMC with Sobol sequences, (ii) non-intrusive PC (NISP) and (iii) SSC. Please refer to Sec. 2.3.2 for detailed explanation of these methods. The comparison is based on the capability to reconstruct the mean and variance of the lift-to-drag ratio for the NACA 0012 airfoil. For this comparison, the XFOIL code for the section at 40% of the blade radius is considered. The

PC order n_o	2	3	4	5	6	7	8	9	10	20	30
Number of samples	9	16	25	36	49	64	81	100	121	441	961

Table 4.3: PC order and number of samples ($(n_o+1)^{n_\xi}$, with n_ξ =2) for the NACA 0012 test.

computations with the QMC are performed for 50, 100, 500, 1000 and 5000 samples. The order of the PC expansion and the corresponding number of samples are reported in Tab. 4.3 (please note the exponential increase of the number of samples, according to Eq. (2.34)). Results with the SSC are obtained in two different ways: (i) using only first-order Lagrange polynomials and (ii) using the p-refinement capability. The maximum order of the polynomial p is set to 5, which in principle provides an order of convergence $\mathcal{O}$ equal to three according to the relation $\mathcal{O} = \frac{p+1}{n_\xi}$ [93]. A uniform h-refinement is considered because it should avoid clustering when increasing the number of samples.

Figure 4.5 shows that the PC, QMC and SSC results converge to the same value with an increasing PC order and number of samples. Taking the QMC result obtained with the higher number of evaluations (5000) as a reference, the error on the mean μ_E and the variance σ_E^2 are computed and plotted in Fig. 4.6.

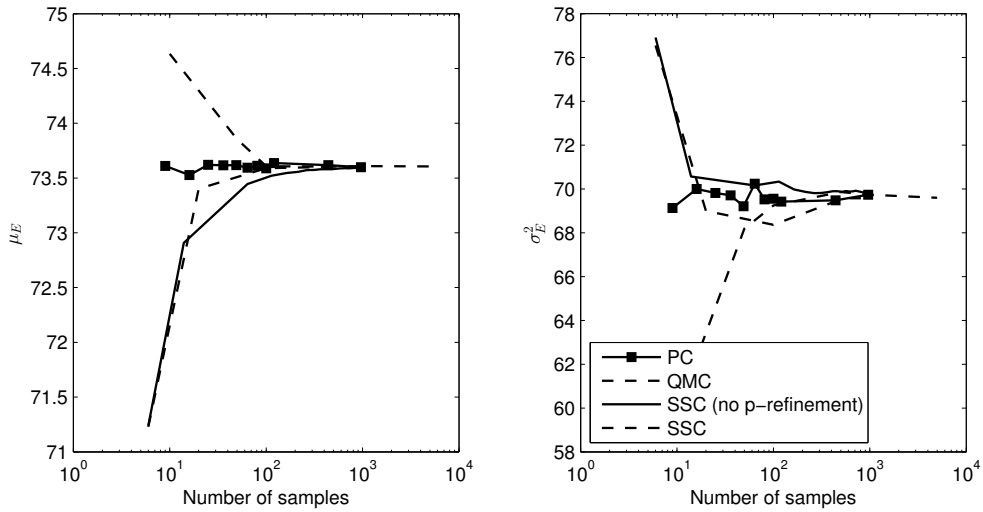

Figure 4.5: Estimate of the mean value (left) and variance (right) of the lift-to-drag ratio E obtained with QMC, PC and SSC methods.

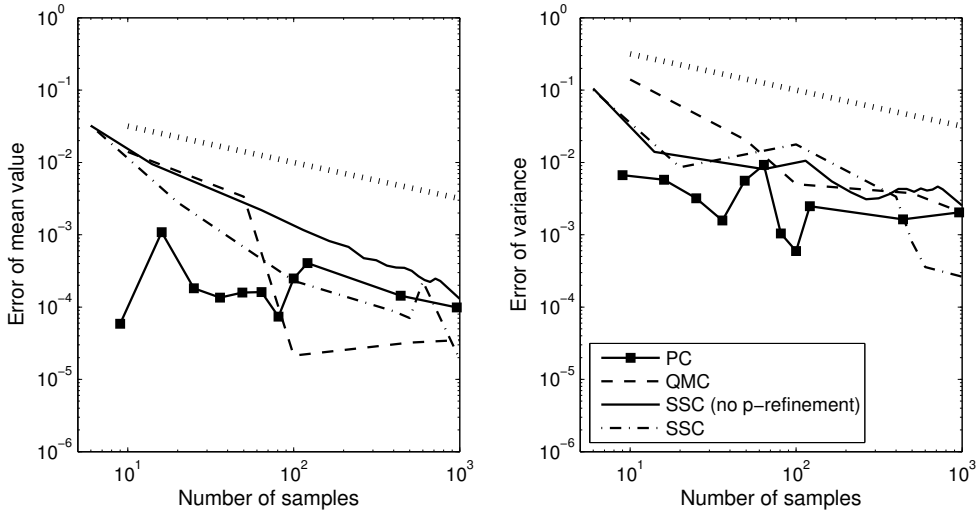

Figure 4.6: Error on the estimated mean value (left) and variance (right) of the lift-to-drag ratio E obtained with PC, QMC and SSC methods (dotted line is the theoretical MC convergence).

First, the error on the estimated mean values are lower than those for the variance. In fact, a higher number of samples are generally required to compute with accuracy a higher order statistical moment (i.e. the variance), whereas the mean value is well estimated with fewer terms. In addition, the error on the mean value computed by the PC and SSC methods is very small even for a small number of samples. With increasing number of samples, the QMC bridges the gap with the other methods, but a number samples higher than 100 is prohibitive for a CFD application in a two-dimensional space. This proves the capability of the PC and SSC approaches to provide an accurate estimate of the statistics with a much lower number of samples with respect to the Monte Carlo method for the problem of interest. It is noted that although the estimate on the mean obtained with the SSC with linear polynomials are quite similar to the estimates of the QMC, the SSC variance in the first iterations has already a lower error.

The PC-based method appears to be accurate right from very low orders of the PC expansion: as a matter of fact, by plotting the response surface of the quantity of interest E in the stochastic space V_i, θ (see Fig. 4.7), it is possible to note that the response is smooth and continuous. Indeed, a global meta-model based on a polynomial basis is well suited to approximate this type of functions. Given these results, the fourth-order PC, which corresponds to 25 samples per design, is chosen for the

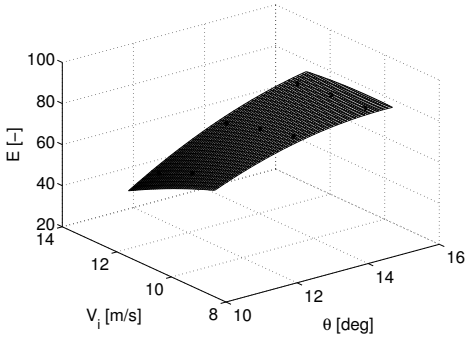

(a) PC order $n_o = 2$.

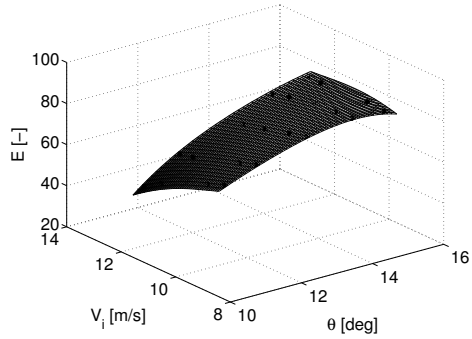

(b) PC order $n_o = 4$.

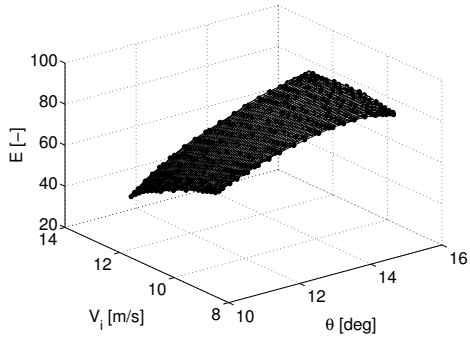

(c) PC order $n_o = 20$.

Figure 4.7: PC meta-model of the lift-to-drag ratio E with increasing order.

optimization, as it represents a good trade-off between accuracy and computational cost.

With regard to the SSC, Figure 4.6 shows how the order of convergence can be improved by using the p-refinement which improves the first-order convergence of the first-order Lagrange basis [93]. It is also worth noting that the SSC error is affected by the MC error used in this implementation to compute the statistical moments from the interpolation, as given in Eq. (2.39). Improvement of the error of this type of representation can be obtained using quadrature rules to solve the integrals based on the approximating function w.

In Fig. 4.8 the error of the SSC with respect to the QMC statistics is compared to the estimates of the error obtained from the hierarchical surpluses (as explained in Sec. 3.2.4). These estimates are used by the SSC to stop the adaptation loop. The stopping criterion is based both on a maximum number of samples and on the comparison between the estimated errors and a tolerance τ. The percentage error

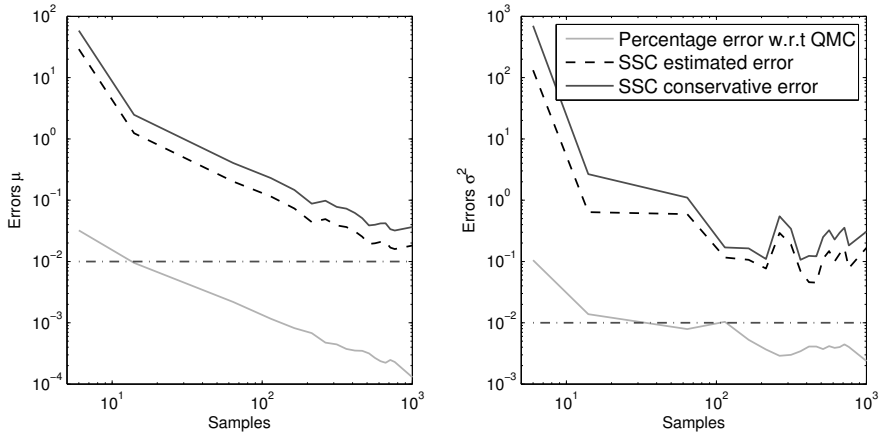

Figure 4.8: Errors of the SSC estimate of the mean value (left) and variance (right)

in Fig. 4.8 falls under 10^{-2} when the SSC estimated errors is 5e-1. Based on these results, the stopping criteria of the SSC adaptivity are set with a maximum number of samples equal to 50 and a tolerance on the estimated error of each statistical moment equal to 5e-1.

It is noted that if SSC were applied to an outer section of the blade section, where compressibility effects are significant (see Fig. 4.3), a similar behavior of the statistics would be expected. As a matter of fact, despite the onset of compressibility effects the response of the lift-to-drag ratio (obtained from integral values) is smooth (this will be presented in Fig. 4.10b). Thus, the convergence of the SSC should be smooth and similar to what presented for the 40% case.

4.3 Multi-fidelity method results

In this chapter, the multi-fidelity method presented in Chapter 3 is built with respect to the different degree of accuracy of the physical models used to describe the flowfield. In particular, two aerodynamic models are considered in the multi-fidelity strategy: (i) the compressible coupled Euler equations/integral boundary layer solver MSES and (ii) the aerodynamic solver coupling a panel method and integral boundary layer model XFOIL. The former is more expensive than the latter from a computational point of view. However, it is more accurate in those operating conditions that exhibit important compressible effects. The proposed multi-fidelity method switches from the low-fidelity model to the high-fidelity model only in those operating conditions. Unlike most multi-fidelity methods in aerodynamic optimization, the switching logic

here is based on physical considerations, so the multi-fidelity approach can be labeled as *physics-based*. The low-fidelity physical model does not require training or training data at each optimization loop, as surrogate models do. An example of a similar approach in the literature for deterministic aerodynamic optimization is presented in Ref. [131].

4.3.1 Application of the multi-fidelity strategy

The method described in Sec. 3.1.1 is applied to the aerodynamic optimization problem in hovering condition. The switching logic is set leveraging physical considerations based on compressibility effects. In Sec. 4.2.2, the estimates of the lift-to-drag ratio obtained with XFOIL and MSES solvers for the NACA 0012 are compared. The XFOIL code provides sufficiently accurate results below the critical Mach number curve (see Fig. 4.3). When compressibility is not negligible, the high-fidelity model computes a more accurate estimate of the objective function. Thus, the multi-fidelity method should switch from the low-fidelity model to the high-fidelity model only when significant compressibility effects are present. In this way, the high-fidelity model would be used only for a small number of samples, namely the samples in which the freestream Mach number exceeds the critical value as shown in Fig. 4.3a. Because the operating conditions for the onset of significant compressibility effects depend on the shape of the airfoil, they should be evaluated for each particular design in the optimization loop. Therefore, a set of control points are necessary, in which a suited control function will be computed.

With regard to the control points, the subset Ξ_c includes the limiting operating conditions in the stochastic space. Because the physical problem considered here is dominated by the effect of the angle of attack and the Mach number, the samples $\boldsymbol{\xi}_i$ are mapped into the (α, M)-plane. The subset of control points is chosen as the set of conditions including (i) the nominal condition (α_n, M_n), and the conditions corresponding to the bounds of the complete samples set: (ii) minimum Mach number M_{min} and angle of attack at the minimum Mach number $\alpha_{M_{min}}$, (iii) maximum angle of attack α_{max} and minimum Mach number M_{min}, (iv) maximum Mach number M_{max} and minimum angle of attack α_{min}, (v) maximum Mach number M_{max} and angle of attack at maximum Mach number $\alpha_{M_{max}}$ (see the crosses in Fig. 4.9).

Thanks to the monotonicity of the discrepancy function (which is represented by points in Fig. 4.3b, the five selected control points should be sufficient for the multi-fidelity method applied to the specific problem. In a more general case, in addition to the extreme points and the center point of the domain, intermediate points should

be considered to control a non-smooth or non-monotonic discrepancy function. For instance, control points could be placed in points with steep gradients or change in the sign of the gradient of the discrepancy function.

Once the control points are defined, a measure of the compressibility effects is required as a control function. The discrepancy between the low-fidelity model and the high-fidelity model is chosen, because this difference is strictly connected to the increase of the wave drag, as presented in Section 4.2.2. For each individual, the operations detailed in Sec. 3.1.1 are performed (see Fig. 3.2b).

4.3.2 Preliminary results

The multi-fidelity strategy is applied to the NACA 0012 as a reference configuration. The PC meta-model which represents the system output $E(\boldsymbol{x}, \boldsymbol{\xi})$ as a function of the uncertain variables (i.e. induced velocity V_i and blade pitch angle θ), is obtained by means of the high-fidelity (HF) model, the low-fidelity (LF) model and finally with the proposed multi-fidelity (MF) strategy. The result is presented in Fig. 4.10, where the samples of the PC method are marked as squares when the HF estimate is employed and dots for the LF computations. As expected, the LF model does not account for the loss in the lift-to-drag ratio encountered at higher values of pitch angles and low values of induced velocity (which correspond to higher angles of attack, since the induced velocity acts as a downwash on the airfoil flowfield). Instead, the correction of the MF method, which is mainly applied in that area, yields the

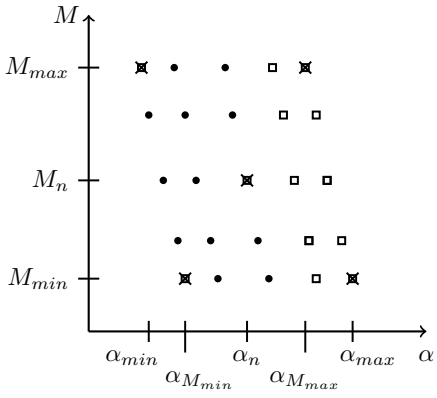

	α	M
i	α_n	M_n
ii	$\alpha_{M_{min}}$	M_{min}
iii	α_{max}	M_{min}
iv	α_{min}	M_{max}
v	$\alpha_{M_{max}}$	M_{max}

Figure 4.9: Points used in multi-fidelity strategy in (α, M)-plane: control points (cross), points evaluated by low-fidelity model (filled dots) and points evaluated with the high-fidelity model (empty squares).

quantitative behavior of the response obtained with the HF model.

The statistics resulting from the PC expansion are then collected in Table 4.4, which presents the mean value, variance and probability of the lift-to-drag ratio E obtained with the HF method and the relative error obtained with the LF model and MF strategy. The error on the MF estimates of the function employed in the optimization loop are reduced with respect to the LF case. Figure 4.11 shows the Probability Density Function (PDF) of the lift-to-drag ratio obtained by means of Monte Carlo sampling of the PC meta-model. It appears that the multi-fidelity strategy well captures the tail of the distribution and it shares the same support with respect to the

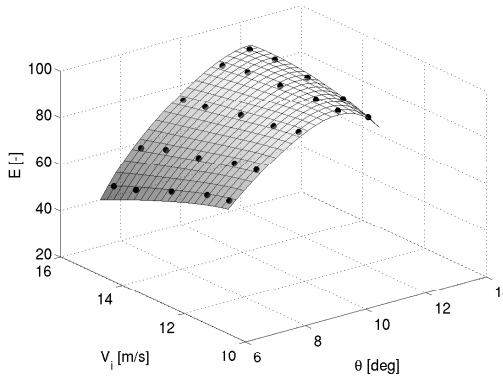

(a) Lift-to-drag ratio meta-model obtained with LF samples.

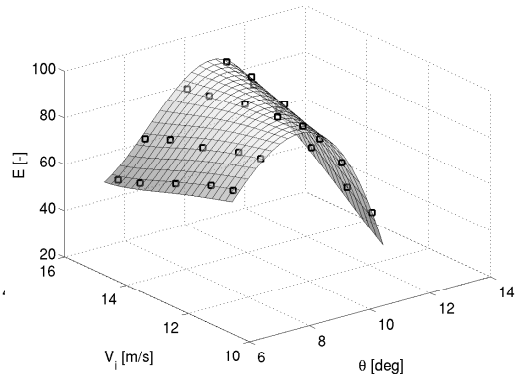

(b) Lift-to-drag ratio meta-model obtained with HF samples.

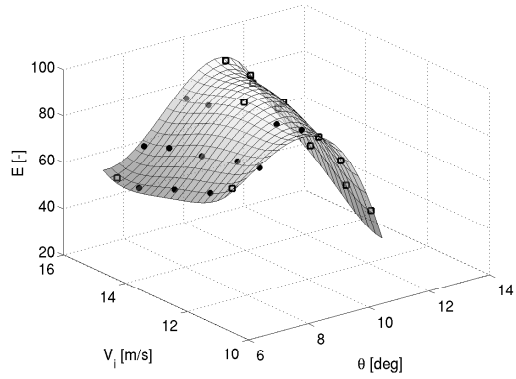

(c) Lift-to-drag ratio meta-model obtained with MF samples.

Figure 4.10: Lift-to-drag ratio evaluation on the PC sampling points using the multi-fidelity approach (empty squares for HF samples, filled dots for LF samples) at $r = 0.8$.

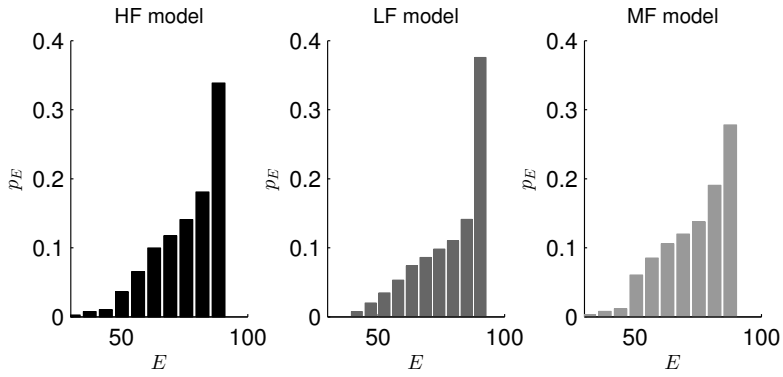

Figure 4.11: Comparison of uncertainty output prediction of the high fidelity (HF) method, low fidelity (LF) method and multi-fidelity (MF) method at $r = 0.8$.

high-fidelity PDF, although both the low-fidelity and the multi-fidelity models lack accuracy in predicting the PDF peak.

Furthermore, the estimates of the mean value and variance of the lift-to-drag ratio obtained with the HF, LF and MF optimization strategies are compared for the set of 200 individuals belonging to the first generation of the genetic algorithm for the blade section at 80% of the radius. The estimates are non-dimensionalised and plotted in Fig. 4.12, where the dispersion with respect to the solid line represents the error between the method reported on the ordinate and the HF result on the abscissa. The MF strategy greatly reduces the dispersion of the LF results, demonstrating the effectiveness of the switching logic.

4.3.3 Optimization results

Two blade sections are considered: an inboard section at 40% of the blade radius and an outboard section at 80% (the corresponding operating conditions are reported in

Method	μ	σ^2	σ/μ	$p(E < 50)$	$p(50 \leq E \leq 80)$	$p(E > 80)$
HF	74.89	128.99	0.15	0.03	0.54	0.43
	err (μ)	err (σ^2)	err (σ/μ)	err $(p(E < 50))$	err $(p(50 \leq E \leq 80))$	err $(p(E > 80))$
LF	3.92 %	6.77 %	0.56 %	20.97 %	17.28 %	22.93 %
MF	3.02 %	0.68 %	3.46 %	4.12 %	13.08 %	16.64 %

Table 4.4: Reconstruction of statistics with low-fidelity, high-fidelity and multi-fidelity models at $r = 0.8$.

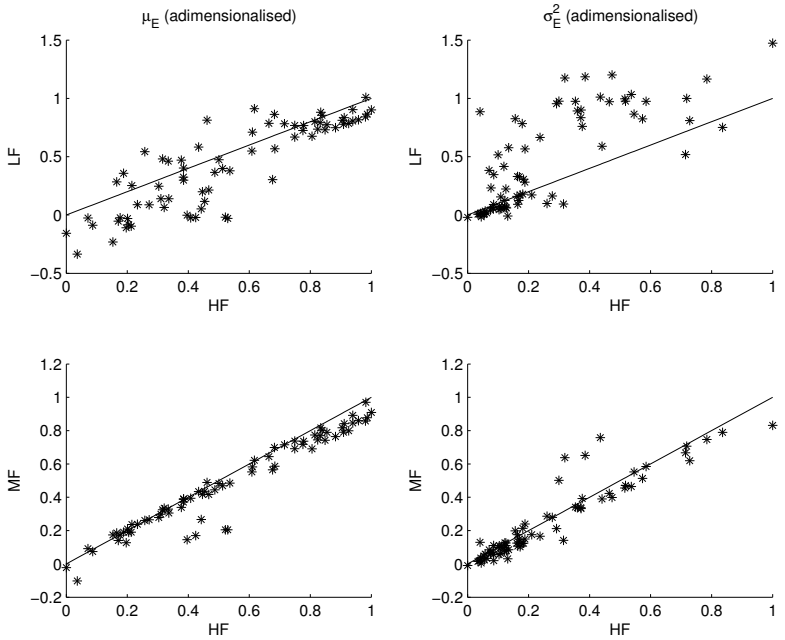

Figure 4.12: Comparison of objective estimates of the HF, LF and MF methods for the first generation of the robust genetic optimization obtained at $r = 0.8$.

Table 4.1). The operating Mach number of the blade section in this case is mainly affected by the tangential velocity due to the rotor rotation, which grows linearly going outboard. Thus, the inboard section features subsonic flow, whereas the outboard section undergoes low transonic flow. The results for the inboard section are presented to assess the effectiveness of the robust optimization technique by comparing the results with other classical optimization approaches, such as the deterministic single-point and multi-point optimization problems. The second set of results are aimed at presenting the performance of the multi-fidelity algorithm.

4.3.4 Inboard section

In this section, the results of the optimization of the inboard section airfoil obtained with the low-fidelity model are presented in order to examine the effectiveness of the robust optimization approach. In addition, because the operating conditions at 40% blade span do not involve significant compressibility effects, the low-fidelity model in this condition provides good estimates of the quantity of interest at a very low computational cost.

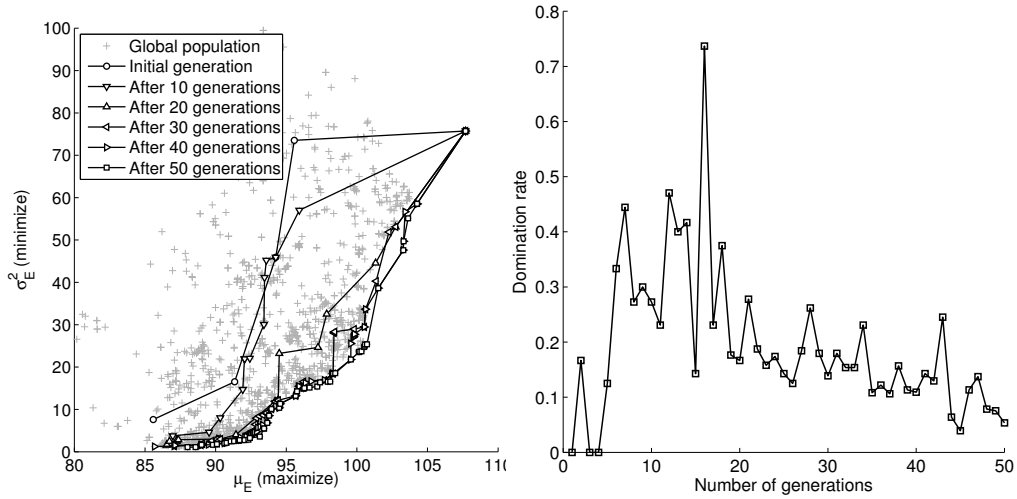

(a) Pareto front evolution. **(b)** Domination rate.

Figure 4.13: Convergence of the Pareto front (LF method).

Convergence analysis of optimizer

To begin with, some properties of the optimization algorithms are discussed. The population size and the number of generations for the convergence of the genetic algorithm should be chosen according to the number of parameters and objectives of the particular optimization problem. However, in practice the population size and number of generations are set by also taking into account the constraints related to the global amount of CPU time devoted to the computation. It must be recalled that taking too small a population may rapidly lead to a local optimum from which the iteration process on the generation number will not evolve. On the other hand, a large population would permit the computation of only a few generations, resulting in a poorly converged solution. Here, after performing an initial generation of 200 individuals to explore the design space, 40 individuals are then retained and evolved. The evolution of the Pareto front is presented in Fig. 4.13a, where the objective functions of the individuals are presented at different generations. It appears that by increasing the number of generations the front moves to higher values of the mean and lower values of the variance, with a decaying rate.

To assess the convergence of the genetic optimization, the domination rate is used, which represents the percentage of non-dominated individuals that become dominated in the following generation, and should then decrease as the population evolves. The domination rate is plotted in Fig. 4.13b. Despite the presence of oscil-

lations due to the stochastic nature of the genetic algorithm, the convergence trend is still clear.

Analysis of Pareto fronts

The final Pareto front is shown in Fig. 4.14a, where the optimal mean airfoil and the optimal variance airfoil are highlighted with different symbols. Estimates of the mean and variance for selected reference airfoils are compared against the Pareto front in the same figure. The NACA 23012 has been here analysed, since it is a reference airfoil for rotorcraft applications. The airfoil labelled Single-Point is the optimal airfoil obtained with a single-objective deterministic optimization in the nominal operating condition. Likewise, the Multi-Point airfoil is the result of a multi-point single-objective optimization, where the cost function is the average sum of the lift-to-drag ratio over three operating conditions (nominal, higher angle of attack and lower angle of attack conditions). For these airfoils, the statistics have been evaluated in a post-processing phase by computing the lift-to-drag ratio in the same uncertain conditions employed for the robust approach. In the objective space, the reference airfoils have either a high mean value (NACA 23012, Single-Point and Multi-Point) or a low variance (NACA 0012). The non-dominated solutions represent a trade-off between the two requirements: in fact, they have a high mean value comparable to those of the Single-Point airfoil, but they also have a variance lower than the value of the NACA0012. In this sense, they outperform the reference airfoils. This is especially true in terms of robustness (i.e. reduction of the variance), even for the cases where the maximum mean values of the lift-to-drag ratio are obtained. As a consequence, the robust approach appears to be a better design method, even when compared with the multi-point optimization. This is because it results in the enhancement of the performance indicators. The airfoils with higher mean and lower variance and a representative trade-off airfoil lying close to the midpoint of the Pareto front are plotted in Fig. 4.14b.

The results obtained with the high-fidelity and multi-fidelity methods are presented in Figs. 4.15a and 4.15b. In order to compare the three fronts, the airfoils belonging to the LF and MF fronts have been evaluated with the HF model after optimization. Figure 4.16a shows the comparison of the Pareto fronts. The results obtained by means of both the low-fidelity and multi-fidelity methods are in very good agreement with those of the high-fidelity model. In particular, the optimal mean airfoil is the same in the three cases, and the optimal variance airfoils are very similar (Fig. 4.16b). In fact, the basic physics that dominates the flow in this condition is cap-

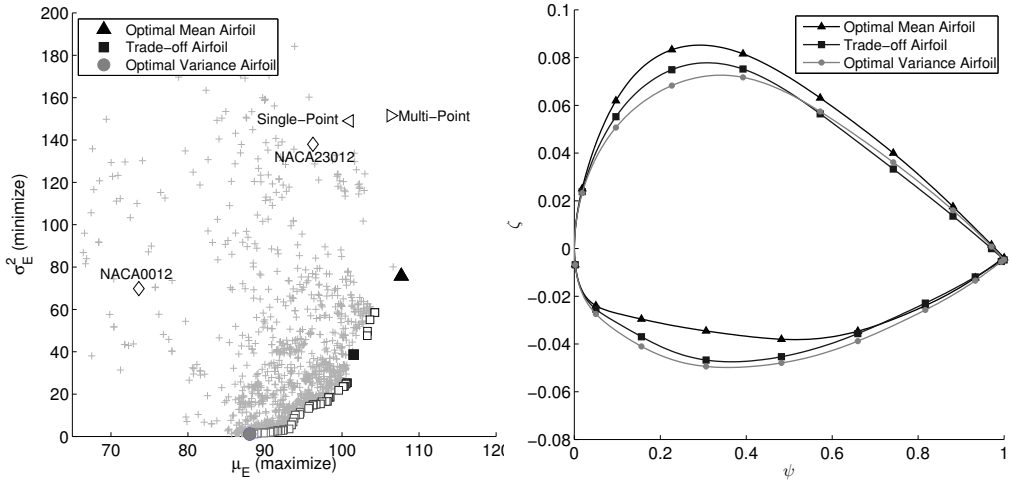

(a) LF Pareto front. **(b)** LF Airfoils.

Figure 4.14: Result of the optimization at $r = 0.4$ for the LF methods.

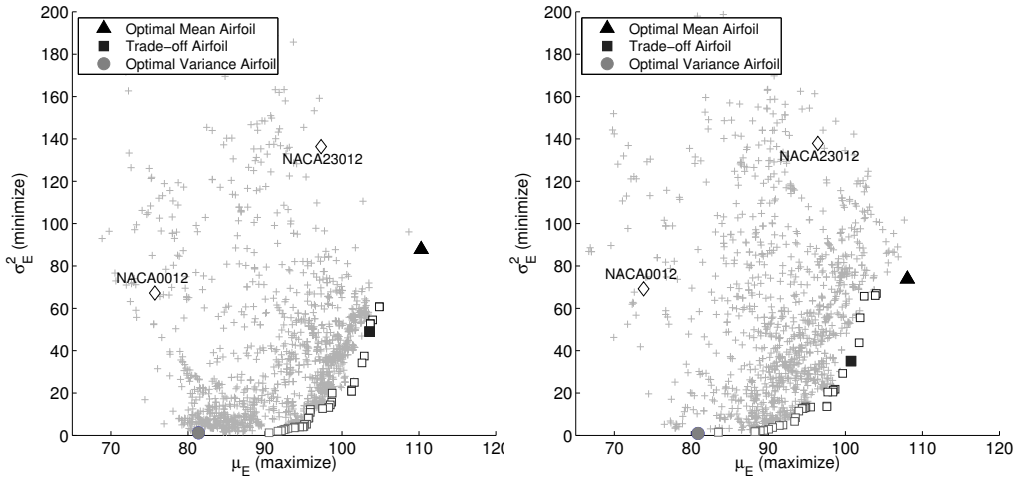

(a) HF Pareto front. **(b)** MF Pareto front.

Figure 4.15: Result of the optimization at $r = 0.4$ for the HF and MF methods.

tured in a similar way by the two numerical models. Nevertheless, the multi-fidelity result is closer to the low-fidelity one. Since the Mach number in this case does not exceed 0.258, the control of the multi-fidelity strategy acts only on a relatively small number of points, as defined in Section 4.3.1.

A quantitative measure of this difference is computed by using Eq. (3.4). The

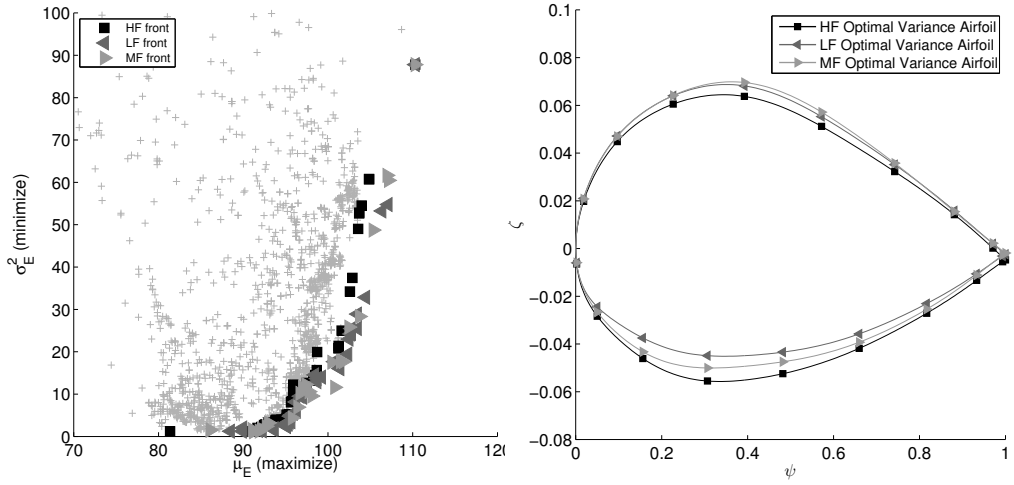

(a) Pareto fronts post-processed with the HF model.

(b) LF, HF, MF Optimal Variance Airfoils.

Figure 4.16: Comparison of LF, HF and MF results at $r = 0.4$.

	LF	HF	MF
LF evaluations	32100	-	31625
HF evaluations	-	32450	6325
Cost	3210	32450	9487

Table 4.5: Comparison of computational cost of different robust optimization strategies: cost is computed considering 1 cost unit for each HF evaluation and 0.1 cost unit for each LF evaluation.

distance d_{MF-HF} of the MF set to the HF set is equal to 44.1, while the distance d_{MF-LF} of the MF set to the LF set is 4.4, proving that the MF set is closer to the LF set rather than the HF set. With regard to the computational cost, the number of function evaluations are computed to compare the three strategies. In particular, the cost is computed by approximating the ratio of the cost of a single LF run over the cost of a HF run to 0.1 and the results are presented in Table 4.5. It is evident that the low-fidelity and the multi-fidelity strategy permit a great reduction of the computational cost, although the optimization result is still consistent with the high fidelity-based method.

Analysis of optimal airfoils

In the following, the optimal mean airfoil and the optimal variance airfoil obtained with the high-fidelity model are analysed in more detail. First, the meta-models representing the lift-to-drag ratio in the stochastic space are compared in Fig. 4.17. As expected from the optimization objectives, the lift-to-drag ratio of the optimal variance airfoil is almost constant over the domain at the expense of a poorer performance. The output PDF is then obtained by Monte Carlo sampling of the meta-model and it is presented in Fig. 4.18. The PDF of the optimal mean airfoil has a spread support which extends to higher values of the lift-to-drag ratio, whereas the PDF of the

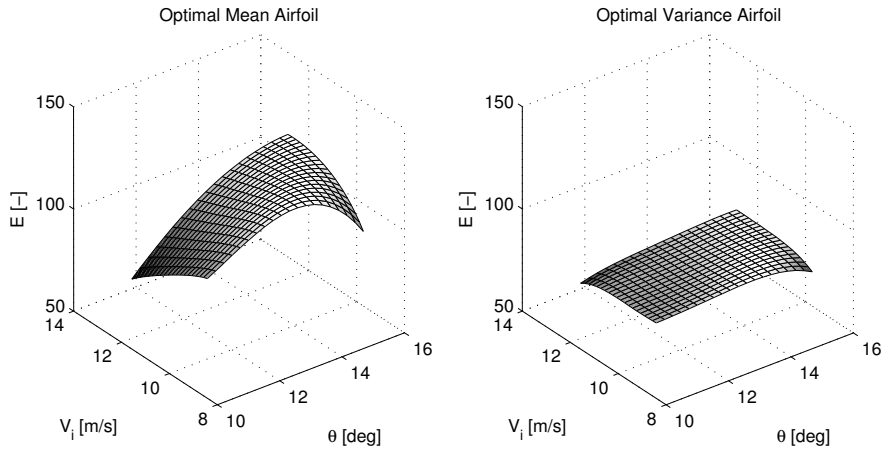

Figure 4.17: Meta-model of lift-to-drag ratio of the optimal mean and variance airfoils (HF).

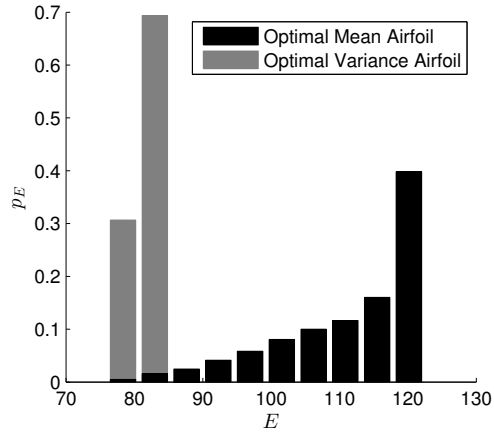

Figure 4.18: Sampled PDF of lift-to-drag ratio of optimal mean and variance airfoils (HF).

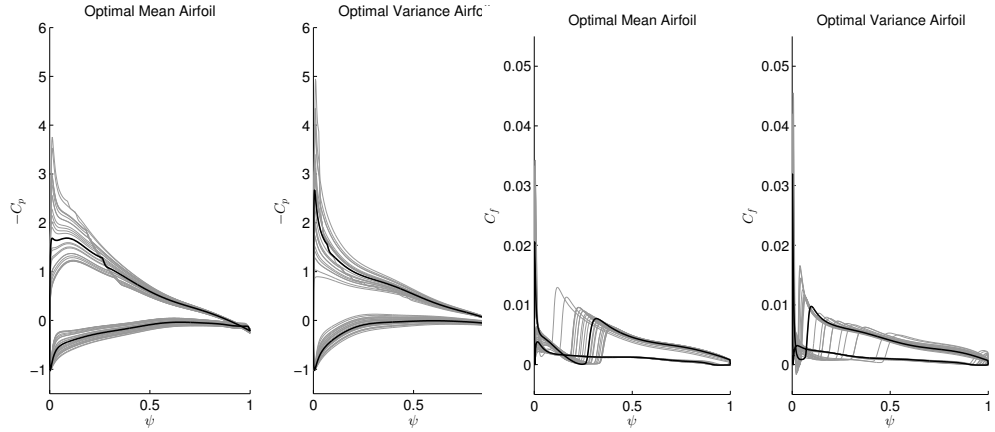

(a) Pressure coefficient. **(b)** Friction coefficient.

Figure 4.19: Comparison of the flowfield of the optimal airfoils (HF): a thick black line is employed for the nominal condition and thin grey lines for the remaining samples.

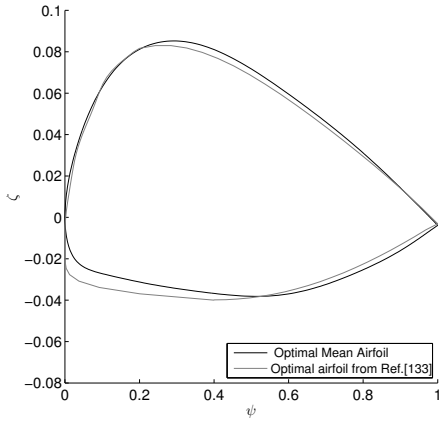

Figure 4.20: Comparison of optimal mean airfoil and optimal airfoil from Ref. [132].

optimal variance airfoil is narrower.

The preceding figures have shown that the optimal variance airfoil appears to achieve the target of minimizing the variance, by reducing the peak of the lift-to-drag ratio with respect to the angle of attack. In order to do so, the optimization process leads to an airfoil shape that reduces the aerodynamic load, by unloading the region after the suction peak on the upper surface and by increasing the local curvature on the lower surface with respect to the optimal mean airfoil (see Fig. 4.19). The reduction of the upper surface thickness results in a sharper leading edge. The fast

transition to a turbulent boundary layer induced by the highly loaded nose contributes to the reduction in the lift-to-drag ratio due to the increase in friction drag. On the other hand, the optimal mean airfoil has higher camber with a more loaded upper surface and a nearly flat region on the lower surface to reduce the flow acceleration. In addition, the pressure peak at the leading edge is less severe and more gradual with respect to the optimal variance airfoil, moving towards a laminar flow airfoil. This results in a lower drag coefficient, at the cost of a higher sensitivity of the performance with respect to changes in the operating conditions. As a final remark, the optimal mean airfoil obtained in this optimization appears to be very similar to the result obtained in Ref. [132] for the subsonic lift optimization of an airfoil with constrained drag (see Fig. 4.20). However, the leading edge in that case is more drooped owing to the higher freestream angle of attack.

4.3.5 Outboard section

In the case of the outboard section at 80% of the blade radius, the freestream condition is low transonic and the compressibility effects are expected to result in non-negligible differences among the three optimization strategies.

Analysis of Pareto fronts

The Pareto front obtained with the proposed multi-fidelity strategy is presented in Fig. 4.21a. The MF Pareto front is closer to the front obtained with the HF method (Fig. 4.21b), even though the starting point of the method is the full LF method which results in a very different front (Fig. 4.21c). To better appreciate the improvement of the MF method with respect to the LF optimization, the LF and MF fronts are re-computed with the HF model, and compared with the HF results in Fig. 4.21d. The multi-fidelity result is closer to the HF result with respect to the LF front: using the definition in Eq. (3.4), the distance d_{MF-HF} of the MF set to the HF set is equal to 9.5, while the distance d_{LF-HF} of the LF set to the HF set is 22.8 (see Tab. 4.6). Also, the MF method defines a Pareto front on the objective space, whereas the LF results are more sparse and do not present a clear trend.

d_{MF-HF}	d_{LF-HF}
9.5	22.8

Table 4.6: Distance of MF and LF fronts to HF front for $r = 0.8$.

The airfoils obtained by the optimization are presented in Fig. 4.22. With respect to the mean optimal airfoils, although the MF and the HF results have a different maximum thickness position, they present similar curvature in the region close to the leading edge. Note that the LF optimal mean airfoil shows a greater leading edge radius, and a distribution of the lower side which resembles the airfoil obtained for the inboard section. As expected, the LF result does not take into account the change

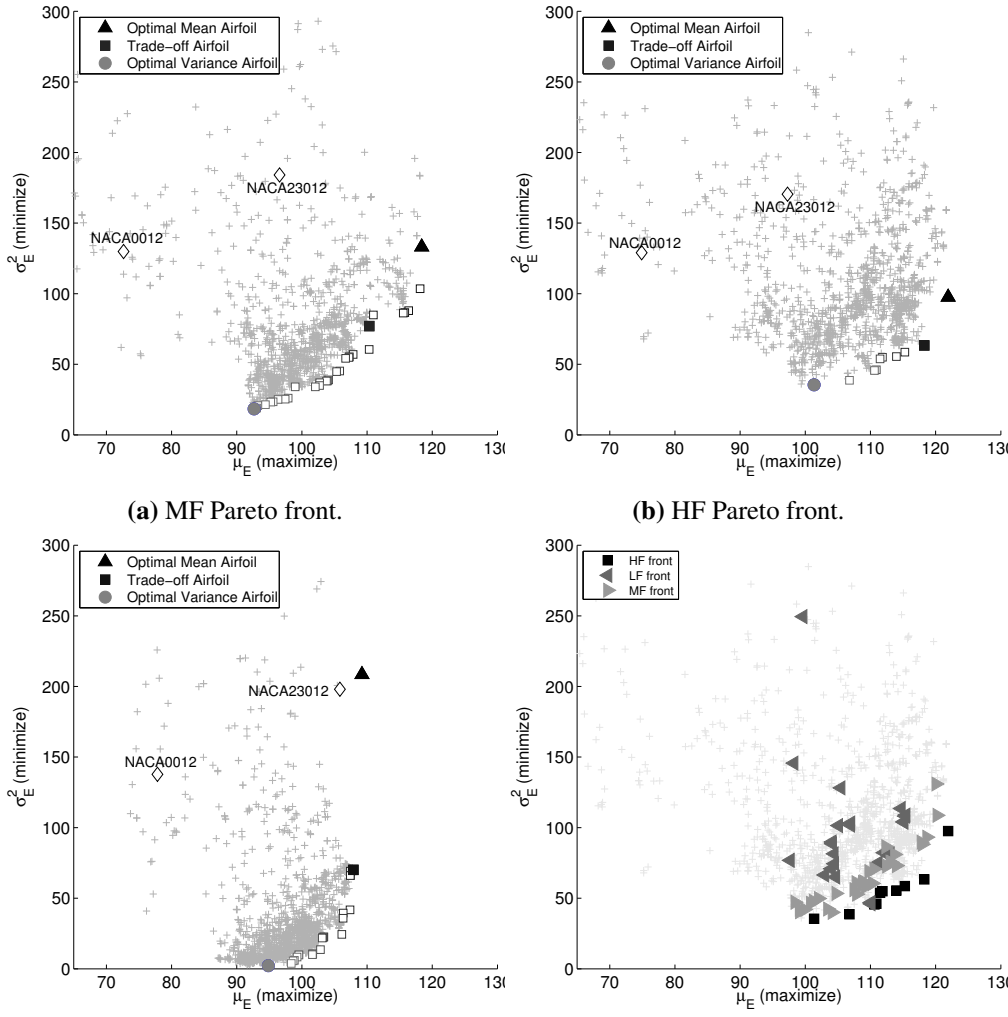

(a) MF Pareto front.

(b) HF Pareto front.

(c) LF Pareto front.

(d) Pareto fronts post-processed with the HF model.

Figure 4.21: Result of the optimization at $r = 0.8$.

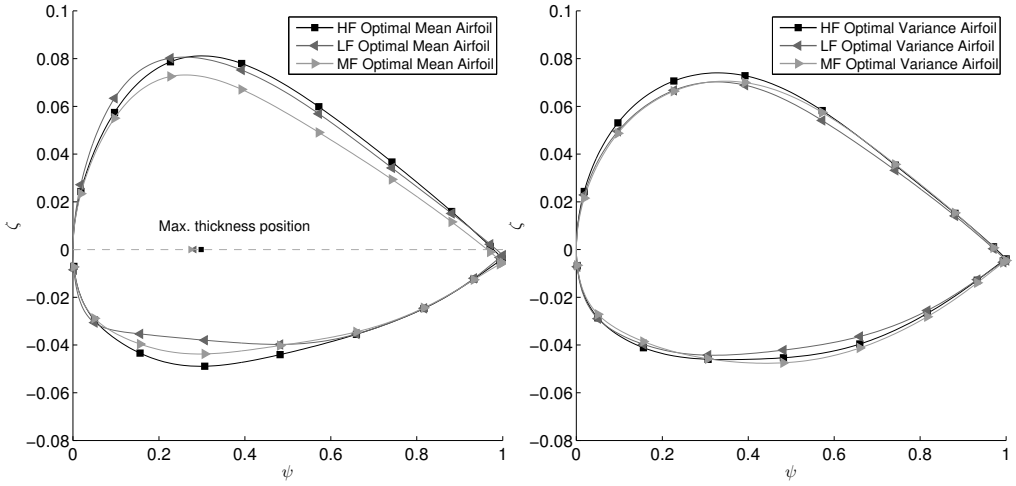

(a) LF, HF, MF Optimal Mean Airfoils. **(b)** LF, HF, MF Optimal Variance Airfoils.

Figure 4.22: Comparison of LF, HF and MF results.

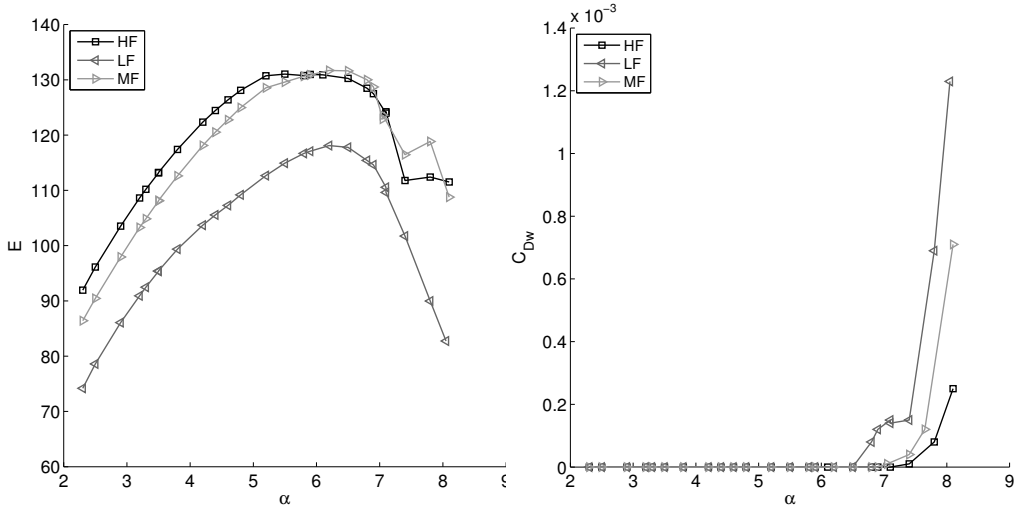

(a) Lift-to-drag ratio. **(b)** Wave drag coefficient.

Figure 4.23: Comparison of the performance and of the optimal mean airfoils resulting from LF, HF and MF methods: results of LF and MF optimizations are re-computed with HF model.

in the operating conditions. The difference on the optimal variance airfoils is instead less pronounced.

The performance of the mean optimal airfoils are compared to appreciate the

	LF	HF	MF
LF evaluations	32600	-	29166
HF evaluations	-	33475	10324
Cost	3260	33475	13241

Table 4.7: Comparison of computational cost of different robust optimization strategies at $r = 0.8$: cost is computed considering 1 cost unit for each HF evaluation and 0.1 cost unit for each LF evaluation.

difference in the results. The lift-to-drag ratio computed for all the three optimization strategies with the HF model is presented in Fig. 4.23a, which shows that the HF and MF results present similar trend in terms of lift-to-drag ratio. Both results exhibit an overall lift-to-drag ratio which is higher than the one obtained by the LF optimization; this is explained in Fig. 4.23b by the behavior of the wave drag coefficient C_{D_W}, which is computed as the integral of the momentum defect over the streamtubes in the region outside the boundary layer [133]. The LF airfoil suffers from a larger drag penalty for higher values of the angle of attack. In fact, the effect of compressibility is greater at higher angle of attack with respect to the HF result. Additionally, the increase in the drag coefficient due to compressibility appears at a lower value of the angle of attack. The result of the HF optimization does not suffer from a harsh drop off after the maximum of the lift-to-drag ratio is reached, thereby postponing the drag rise due to compressibility effects. In this respect, the MF result represents a trade-off solution between the two results.

With regard to computational cost, the number of evaluations of the solvers is presented in Table 4.7. The comparison proves that the computational effort of the HF method is reduced by 60% with the MF strategy. Also, the MF strategy employs one third of the HF evaluations required by the full HF optimization. This proves the main benefit of the proposed MF strategy, which generates solutions very close to the HF ones, but with a strong reduction of the computational cost. The MF strategy is thus capable of indicating regions of improvement in the design space, whereas a detailed analysis of the optimal solutions is left to higher fidelity methods and/or wind tunnel testing.

Analysis of optimal airfoils

To understand the validity of the robust approach in the outboard case as well, the HF results are discussed in more detail. With regard to the Pareto front airfoils presented

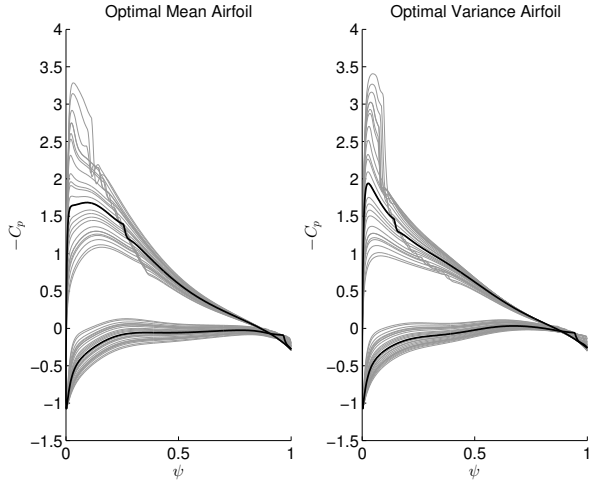

Figure 4.24: Comparison of the pressure coefficient of the optimal airfoils (HF) at $r = 0.8$: a thick black line is employed for the nominal condition and thin grey lines for the remaining samples.

in Fig. 4.22, it appears that the relationship between the optimal mean airfoils and the optimal variance airfoils found in the inboard section case, still holds. The optimal mean airfoils tend to be those with a mild curvature at the leading edge and more rear loaded. The optimal variance airfoils present a higher curvature at the leading edge. Indeed, the analysis of the pressure coefficient of the two airfoils, presented in Fig. 4.24, shows that the optimizer generates an optimal mean airfoil with an almost flat region of pressure coefficient behind the leading edge and reduces the number of operating conditions in the considered range for which a weak shock wave occurs. The optimal variance instead counteracts the growth of the lift coefficient resulting from the increase in the angle of attack, by producing a shock wave behind the peak close to the leading edge in a larger range of operating conditions. This balance results in a lower variance, at the expense of lower lift-to-drag ratio.

4.4 Adaptive method results

The adaptive method based on error bounding boxes (presented in Sec. 3.2) is applied to the optimal selection of the airfoil for a helicopter rotor blade in hover (see Eq. (4.2)). The following set of results are obtained for the mid-span section ($r = 0.5$), see Tab. 4.1 for the operating conditions. For this blade section, the low fidelity XFOIL code [99] provides sufficiently accurate estimates, because compress-

ibility effects are not present, as shown in Fig. 4.3b. A different section of the blade is used with the intent to explore how the optimal shapes change along the blade radius. Because time had to be devoted to the forward flight case (presented in the next chapters), the analysis of the optimal load distributions along the blade radius is left for future activities.

The robust aerodynamic optimization has been performed by a decoupled approach with tight tolerance ($\underline{\tau}$ = 5e-1) and a coupled approach with $\underline{\tau}$ = 5e-1 and $\bar{\tau}$ = 5. The values of the tolerance are set by applying the same convergence analysis presented in Sec. 4.2.5 for the considered blade radius (see Ref. [134]). In addition, the decoupled approach with loose tolerance is not considered here, because it is deemed that in real-world applications due to the substantial cost of the robust optimization procedure it may not be possible to assess the effect of the tolerance. If a single run has to be performed, it is best to choose a tight tolerance that guarantees good results. For the optimization loop the NSGA has been used, because the efficiency of the BIMADS method degrades with such a high number of design variables.

The Pareto fronts obtained with the decoupled and coupled methods are presented in Fig. 4.25a and Fig. 4.25b, respectively. In the figures, the solution with higher mean value, higher variance and a trade-off solution are highlighted. It is clear that the coupled method compares well to the decoupled approach. The coupled front is presented in Fig. 4.26 with the result of the decoupled method. The coupled front al-

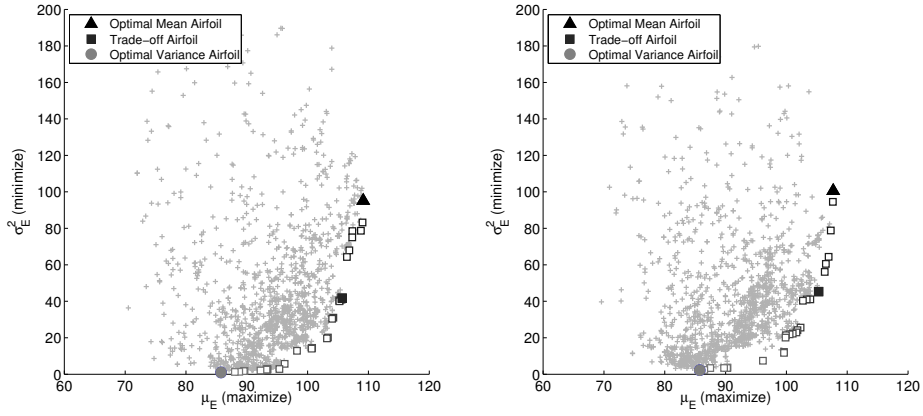

(a) Pareto front of decoupled tight approach.

(b) Pareto front of coupled method.

Figure 4.25: Pareto fronts of decoupled and coupled methods.

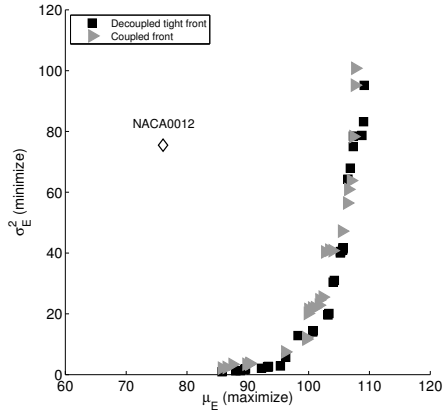

Figure 4.26: Comparison of Pareto fronts (front obtained with coupled approach has been post-processed with tight tolerance for comparison).

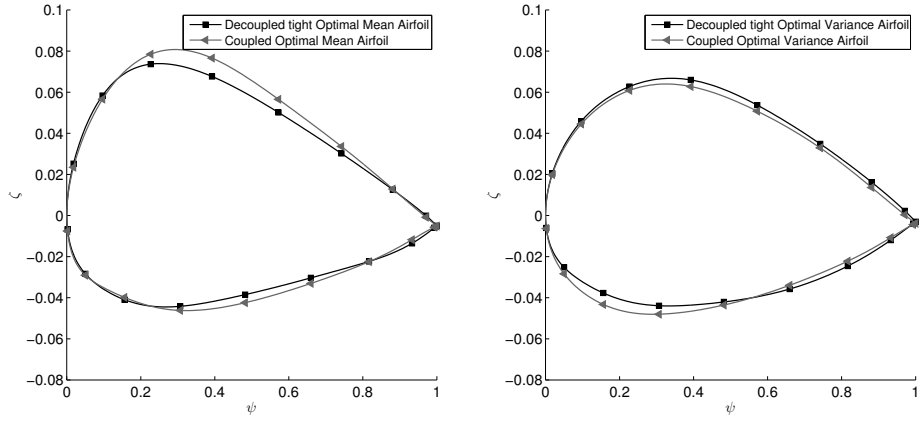

(a) Optimal mean airfoils. **(b)** Optimal variance airfoils.

Figure 4.27: Optimal airfoils of decoupled and coupled methods.

most overlaps the decoupled results and the distance between the fronts computed by means of Eq. (3.4) is equal to 2.25. The solutions belonging to the front improves the mean value of the lift-to-drag ratio of the reference airfoil, i.e. the NACA 0012, and a subset has also reduced variance. To compare the optimal shapes, the airfoils with higher mean value and higher variance are presented in Fig. 4.27a and Fig. 4.27b. Despite minor differences, the airfoils of the fronts obtained with the two methods share similar trends.

With regard to cost, the coupled method greatly reduces the computational effort of the decoupled approach (see Tab. 4.8): in the first generation the cost is reduced

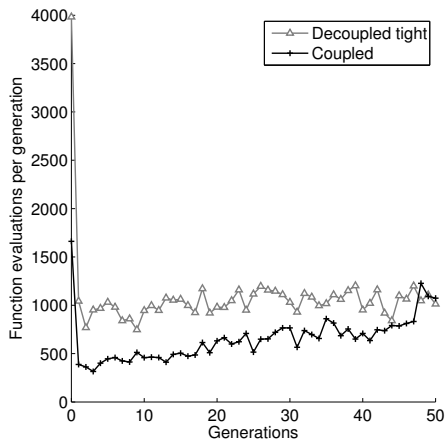

(a) Number of function evaluations.

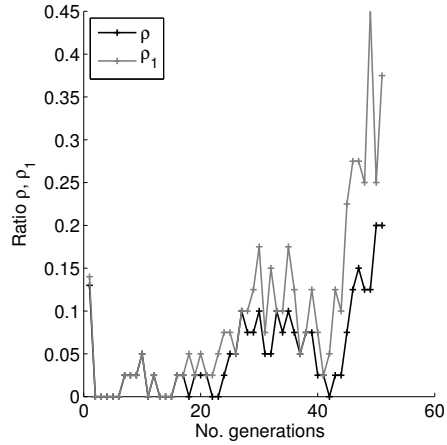

(b) Ratio of potentially non-dominated solutions to generation size.

Figure 4.28: Cost of the decoupled and coupled methods for the robust aerodynamic optimization.

Strategy	Number of function evaluations		
	Initial generation	Evolution	Global
Decoupled tight	3982	51258	55240
Coupled	1661	31760	33421

Table 4.8: Comparison of total number of function evaluations for the robust aerodynamic optimization.

by 58% and the gain on the global cost is equal to 40%. The great reduction is due to the fact that the optimization problem is high-dimensional and nonlinear. Thus, the optimization algorithm samples many design vectors in the design space with poor objectives, especially in the first iterations. In Fig. 4.28b the ratio ρ of potentially non-dominated solution over the generation size is presented, which is computed from Eq. (3.32). As presented in Sec. 3.2.5, a solution is defined as potentially non-dominated if the objective vector is non-dominated with respect to the individuals computed up to that point. In this case, ratio ρ is lower than 0.2 (see Fig. 4.28b) with a growing trend after 20 generations. This means that the generation becomes rich of non-dominated individuals only after a few generations of the algorithm. As pointed out in Sec.3.2.5, ratio ρ is a lower bound of the actual cost, since it does not take into account the refinement needed for dominated solutions close to the Pareto front.

To highlight this aspect, Figure 4.28b presents the modified ratio ρ_1, computed taken into account not only the non-dominated solutions, but also the solutions within a distance to the current Pareto front of $\underline{\tau}$ in each direction of the objective space. The second ratio has a trend closer to the one of the actual cost of the coupled strategy.

To conclude, the coupled approach gives an accurate estimate of the Pareto front with a greatly reduced computational effort.

4.5 Final remarks on the methods

The methods outlined in Chapter 3 have been applied in the preceding sections to the case of the aerodynamic optimization of the blade airfoil in hovering condition.

When considering uncertainties in the operating conditions, different flow regimes could be encountered in the domain defined by the stochastic variables, which require specific physical models and numerical solvers. The multi-fidelity strategy proposed in the work deals with two different models of the aerodynamics which have different accuracy and computational cost. In particular, a rapid estimation is obtained by means of a panel/integral boundary layer method for the stochastic sample associated with subsonic flow conditions, while a coupled Euler equations/integral boundary layer solver is employed for those operating conditions in the stochastic variables domain that exhibit important compressible effects. A switching logic is used for each design to determine which samples should be computed by the rapid low-fidelity or by the high-fidelity codes. The method proved effective in finding airfoils that improve the lift-to-drag ratio, by postponing the boundary layer transition to turbulence in the subsonic case, and by reducing the drag due to compressibility effects in the low transonic case, in the range of considered operating conditions. The multi-fidelity strategy reduces the computational cost by 60% with respect to the high-fidelity strategy, and results in Pareto airfoils consistent to the high-fidelity results. Thus, the proposed method is well-suited to indicate regions of improvement in the design space, which can be later analyzed in detail by means of a high-fidelity model. In addition, the multi-fidelity strategy could be easily applied to other models at one's disposal, and it could be extended to other flowfield characteristics and physical models.

The gain of the multi-fidelity strategy is close to 60% on the global population, whereas the cost reduction with the adaptive method is close to 40%. However, the counterpart of the lower gain of the adaptive strategy is the greater accuracy in computing the Pareto front: the distance between the reference front and the front obtained with the new strategy is 9.5 in the case of the multi-fidelity strategy and

2.25 for the adaptive strategy. As pointed out in Sec. 3.2.2, the method pays the price of computing the non-dominated individuals with the tightest tolerance to preserve the accuracy of the final front.

Another important aspect when comparing these methods is related to the possibility of using parallel computation. The adaptive strategy in the current form does not allow for a parallel mode (except for the very first iteration). This is a crucial difference if parallel computing resources are available. A possible solution to cure this issue in the adaptive strategy could be to perform the convergence check and the tolerance update only after a certain number of iterations, rather than at each iteration.

Although a comparison of the methods is proper at this stage, it is also important to note that the two methods could be easily integrated, as was previously stated in Sec. 3.3. The resulting adaptive, multi-fidelity strategy could provide higher computational gain and it is left for future activities.

Finally, robust optimization proved to be a reliable and effective design tool when compared with two deterministic approaches. Furthermore, extremely low values of the variance in the Pareto front comes at the expense of poor performance, while airfoils with higher mean values possess lower values of the variance of the lift-to-drag ratio with respect to deterministic results. The reduction of the variance could lead for instance to a reduction of the required rotor shaft torque in variable operating conditions.

CHAPTER 5

Robust optimization of airfoils in forward flight

Forward flight is one of the most peculiar flight condition for helicopters. As discussed in Chapter 1, the flow encountered by the blade in such a condition is periodically changing over the azimuth angle and the extremes of this cycle are very different from each other. In the advancing side the rotor blade element encounters transonic flow and low values of the angle of attack. In the retreating side, the angle of attack is typically high and the local Mach number is close to the incompressible limit. For this reason, aerodynamic design should take into account the variability of the flow encountered by the blade.

Several approaches exist to consider the variation of the performance over the azimuthal angle. A full unsteady optimization may be considered, by inserting unsteady models in the optimization loop and directly evaluating the cycle of lift and drag coefficients for each design. However, such an approach is very expensive from a computational point of view and the accurate prediction of the time-varying aerodynamic load of pitching/plunging airfoil is still a challenge. In most cases, one can resort on a multi-point approach, where the evaluation of the aerodynamic performance is obtained by means of steady models. This approach is the one pursued in the thesis. In particular, a two-point optimization problem is tackled, considering the two most representative conditions of forward flight: the retreating and advancing blade conditions (see Fig. 1.4). Because the optimization only considers two positions over the azimuth, the evaluation of the performance of the optimal shapes over the azimuth angle is performed in the post-processing phase with steady models.

To motivate the choice of steady models, the beginning of the chapter discusses the comparison between unsteady CFD model predictions and steady computations.

In the following sections, the optimization problem of improving the aerodynamic performance of fixed airfoils is tackled from both deterministic and robust points of view. In the deterministic case presented in Sec. 5.2, the objective of the optimization is to minimize the drag while maintaining a level of lift that guarantees satisfaction of trimming condition. In the case of robust design, a single condition is not considered and the trimming condition cannot be applied to a number of values of the angle of attack. Thus, the idea is to optimize the lift-to-drag ratio as a measure of the performance of the airfoil, impose inequality constraint on the lift coefficient to guarantee a sufficient level of lift and then check after optimization that the trimming condition can be satisfied. This approach is used in Sec. 5.3 and results coming from deterministic and robust optimization are presented and compared in Sec. 5.4. Because in both cases only two points over the azimuthal angle are considered in the optimization loop, an *a posteriori* assessment of the performance over the azimuthal angle is presented in Sec. 5.5.

5.1 Unsteady CFD simulation

The unsteady effects encountered on a helicopter rotor blade affect the prediction of both local aerodynamics along the blade radius and the overall performance of the rotor. A proper simulation of the complex rotor blade environment requires considering simultaneous velocity and incidence fluctuations of the incoming airflow, compressible effects, three-dimensional effects and accurate viscous separation modelling and coupling of the aerodynamic loads with the blade dynamics and aeroelasticity. Nevertheless, most of the studies regarding the aerodynamics of blades have tackled the problem by separating the influence of these combined effects.

In particular, when analysing the performance of a blade element, most of the studies have simulated the case of a two-dimensional unsteady flow around a pitching or plunging airfoil. However, in forward flight, the combination of the rotational velocity of the rotor and the advancing velocity causes the blade section to encounter a time-varying incident velocity. This is clearly presented in Fig. 5.10, taken from Ref. [9], where the operating conditions of a blade element are presented for different blade radii. It is clear that both the local angle of attack and the freestream Mach number undergo a periodic variation over the rotor rotation period.

In the literature only few works have considered a translation parallel to the freestream direction which is necessary to simulate the speed fluctuation encountered

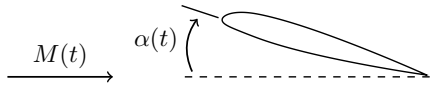

Figure 5.1: Airfoil undergoing time-varying freestream Mach number and angle of attack.

by the blade section [135]. Among these, References [136, 137] have been devoted to the experimental investigation of the effects of coupled fluctuations in angle of attack and freestream speed, as it occurs in the real flow around the rotor blade. These studies have tackled conditions below stall (with maximum angle of attack amplitude lower than the static stall angle) and through stall as well. A few theoretical aerodynamic models based on the extension of Theodorsen method [1, p.431] are presented in Ref. [138] for an airfoil undergoing a combination of harmonic pitching, plunging and fore-aft motion in incompressible flow. Results based on these theories have been used for validation in recent years for the assessment of a surrogate model of the unsteady aerodynamics of a pitching airfoil with time-varying freestream speed in Ref. [139]. A compressible formulation of the theories presented in Ref. [138] has been developed and assessed in Ref. [140] against a CFD code based on the Euler equations. A numerical analysis of dynamic stall caused by pitching motion and freestream fluctuations has been performed by means of a CFD tool based on the SST $k-\omega$ turbulence model in Ref. [141]. The shortage of data is likely to be due to the fact that most of these works focus on both dynamic stall and time-varying incident velocity: because dynamic stall simulation is complex on its own, considering a time-varying freestream speed in addition to a time-varying angle of attack may increase the complexity and understanding of a phenomenon that is not fully understood and is still the subject of much research [1, p. 526]. Another explanation for the lack of experimental data is the challenge associated with producing high-amplitude unsteady oscillations in the wind tunnel [142].

In such a context, an investigation of the unsteady aerodynamic loads acting on the reference airfoils, i.e. the NACA 23012 and the NACA 0012 airfoils, is performed. The investigation consists of the simulation of combined pitching motion and time-varying Mach number as depicted in Fig. 5.1.

5.1.1 Unsteady CFD set-up

The simulation is performed by means of the SU^2 software suite [143]. In particular, the unsteady Reynolds-Averaged Navier-Stokes equations together with the Spalart–

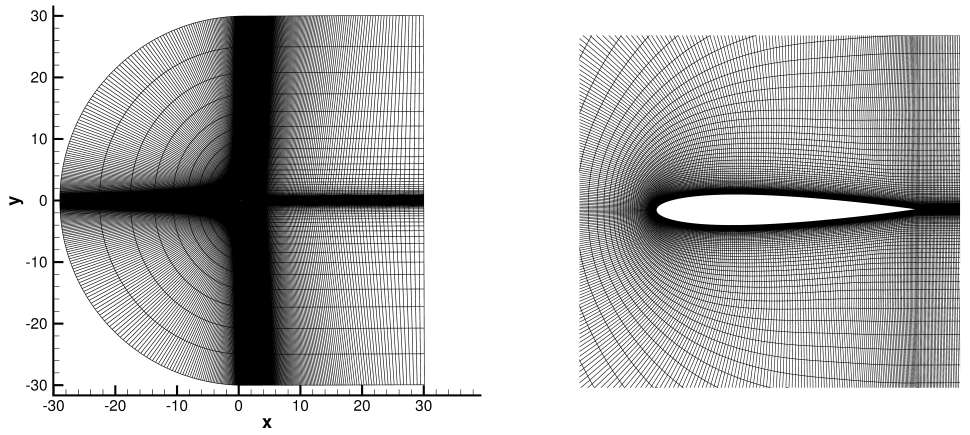

(a) C-mesh. **(b)** Close-up on the airfoil.

Figure 5.2: Mesh for unsteady CFD computations.

Allmaras model [85] is used to simulate the flowfield around the airfoil. The numerical grid is a C-type, structured mesh (Fig. 5.2) with 77792 rectangular cells. The grid extends $\pm$ 30 chords around a unit chord airfoil (Fig. 5.2b) and the first cell is equal to y^+ since wall law is not used in the simulation.

The time-varying freestream Mach number and angle of attack have the following time dependency

$$\begin{aligned} V(t) &= V_\infty \left(1 + \lambda \sin(\omega t)\right) \\ \alpha(t) &= \alpha_0 + \Delta\alpha \sin(\omega t + \overline{\psi}) \end{aligned} \tag{5.1}$$

In order to apply a sinusoidal variation to the freestream velocity, a time-varying translation of the airfoil is performed with the following time dependency

$$h(t) = -\frac{V_\infty \lambda}{\omega} \cos(\omega t). \tag{5.2}$$

The translation is applied in the direction of the mean angle of attack α_0. The time-varying translation and pitching are implemented with a volumetric rigid displacement applied to each cell of the mesh in each iteration of the physical time.

5.1.2 Validation

The CFD setting presented in the preceding section is validated against experimental and computational measurements of the effects of combined angle of attack and velocity fluctuations below stall. the fluctuations in the angle of attack $\alpha(t)$ and incident velocity amplitude $V(t)$ are given by the following

$$V(t) = V_\infty \left(1 + \lambda \sin(\omega t)\right) \tag{5.3}$$

$$\alpha(t) = \alpha_0 + \Delta\alpha \sin(\omega t + \overline{\psi}) \tag{5.4}$$

where V_∞ is the mean freestream speed, λ is the percentage of variation in the freestream speed, ω is the angular velocity of the fluctuations, α_0 is the mean angle of attack and $\Delta\alpha$ is the amplitude of the oscillation in angle of attack and $\overline{\psi}$ is the phase angle. Please note that the freestream velocity is directed according to the mean angle of attack α_0.

In this type of simulation is of utmost importance to fulfill the requirement of geometric, kinematic and dynamic similarity between the numerical simulation and the experimental reference data. The first requirement consists in setting up the geometry of the numerical simulation scaling the geometry of the experimental set up while maintaining its proportion. Second, two measurements (either numerical simulation or wind tunnel testing) are dynamically similar when the forces at corresponding points are similar. In the application to the simulation of an oscillating airfoil, the equivalence of the Reynolds number and the Mach number guarantees the fulfilment of this requirement [144, p.262]. Finally, with regard to kinematic similarity, the velocity and accelerations of the numerical simulation should be consistent with those of the reference set up. In particular, for a pitching airfoil with varying freestream, the amplitude and frequency of the oscillation, the type of motion and the freestream must match the reference values. For the frequency to be consistent with the experimental data, it is necessary to match the reduced frequency k, an indicator of the degree of unsteadiness in the problem that appears when non-dimensionalizing the Navier-Stokes equations [1, p.427]. For an airfoil of chord c undergoing a periodic variation in angle of attack or a freestream fluctuation with angular frequency ω, the reduced frequency is defined as

$$k = \frac{\omega c}{2V_\infty} \tag{5.5}$$

where V_∞ is the freestream speed (or mean value, if a periodic variation in velocity amplitude is considered).

For the validation, different test cases are considered. In the first test case, ex-

α_0 [deg]	$\Delta\alpha$ [deg]	V_∞ [m/s]	Re	M_∞	λ	k	$\overline{\psi}$ [deg]
6.	0.	25	$4{\cdot}10^5$	0.07	0.114	0.101	0

Table 5.1: Operating conditions for validation with experimental data from Ref. [136].

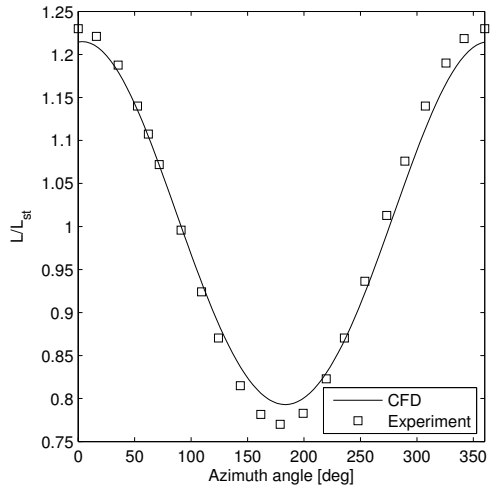

Figure 5.3: Lift ratio L/L_{st} as a function of the azimuth angle $\psi = \omega t$ for an airfoil undergoing freestream Mach number variation (Eq. (5.3)).

perimental data are taken from Ref. [136] for the case of the NACA 0012 airfoil undergoing the speed fluctuation in Eq. (5.3). The corresponding operating conditions are reported in Tab. 5.1. The computations have been performed considering 100 intervals for each period $T = 2\pi/\omega$, and 800 steps in the dual-time iteration. The number of steps in dual time are set to 800 to guarantee that the residuals of the equations are always lower than 1e-4 for every time step. It is noted that the simulation with time-varying freestream Mach number requires a streamwise translation of the airfoil with amplitudes depending on the chosen variation of speed, but also on the reduced frequency. In some cases, the amplitude of this motion may exceed 5 chords, therefore it has been found necessary to both reduce the time steps (to apply a small translation at each outer iteration) and increase the number of dual time steps (to guarantee time-accuracy). Figure 5.3 presents the lift ratio L/L_{st} between the unsteady lift L and the steady value L_{st} plotted against the azimuth angle $\psi = \omega t$. A very good agreement is found between the CFD result obtained with the SU^2 software and the experimental data, with an error affecting the amplitude of the lift ratio by 6%.

	α_0 [deg]	$\Delta\alpha$ [deg]	M_∞ [-]	λ [-]	k [-]	$\overline{\psi}$ [deg]
Time varying α	2	2	0.3	0.	0.2	0
Time varying M	1	0	0.5	0.4	0.2	0
Time varying AoA and M	1	1	0.5	0.4	0.2	180

Table 5.2: Operating conditions for validation with numerical data from Ref. [140].

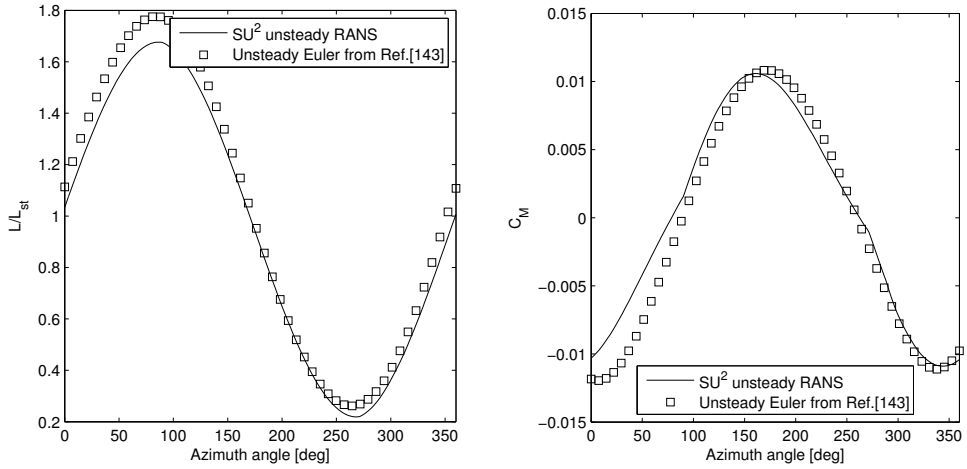

(a) Lift coefficient. **(b)** Moment coefficient.

Figure 5.4: Lift and moment coefficient for $\alpha = 2 + 2\sin(\omega t)$ deg, $k = 0.2$.

Secondly, the numerical data presented in Ref. [140] for the NACA 0006 airfoil have been compared to the results obtained with the unsteady CFD set-up presented in the preceding section. The data of this reference are obtained with an Euler equations model and many combinations of pitching and fluctuating velocity field have been considered. Table 5.2 shows the different combinations considered in the validation process. Although Reference [140] uses a different aerodynamic model than the one used in the present thesis, this data set is used to build confidence in the unsteady RANS set-up because it is rich in data for a wide range of conditions.

The resulting aerodynamic coefficients are presented in Figs. 5.4, 5.5, 5.6. Figure 5.4 presents a case of pure pitching motion, where good agreement is found for both the lift and moment coefficients. In Fig. 5.5 the coefficients for the case of time-varying freestream are presented: in this test case, an Euler computation with SU^2 has also been performed, to demonstrate that the discrepancy in the result is ascribed

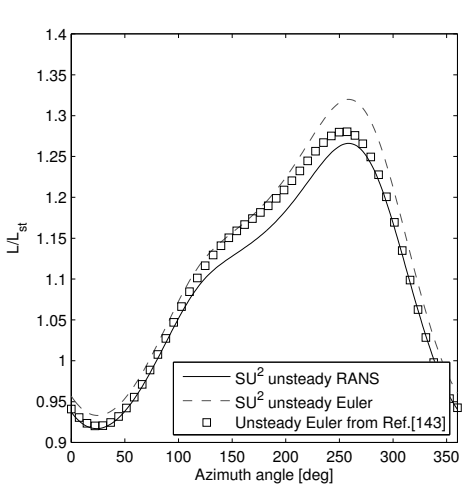

(a) Lift coefficient.

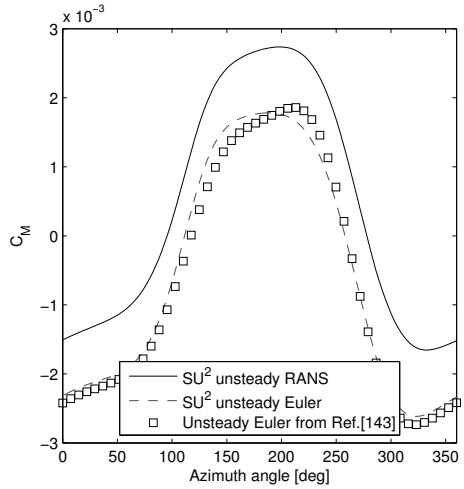

(b) Moment coefficient.

Figure 5.5: Lift and moment coefficient for $M = 0.5(1 + \lambda \sin(\omega t))$, $\alpha = 1$ deg, $k = 0.2$, $\lambda = 0.4$.

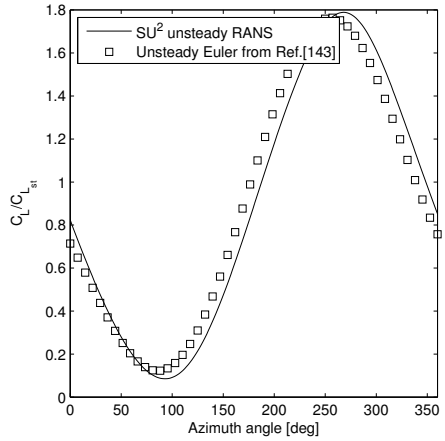

Figure 5.6: Lift coefficient for $M = 0.5(1+\lambda \sin(\omega t))$, $\alpha = 1+1\sin(\omega t+\overline{\psi})$ deg, $k = 0.2$, $\lambda = 0.4$, $\overline{\psi} = 180$ deg.

α_0	$\Delta\alpha$	M_∞	Re	λ	k	$\overline{\psi}$
2	2	0.04	3e5	0.5	0.097	0

Table 5.3: Operating conditions for validation with experimental data from Ref. [145].

to the different choice of aerodynamic model and not to the unsteady RANS set-up. Finally, a case of combined pitching and freestream fluctuations is considered. The corresponding lift coefficient is presented in Fig. 5.6, which shows a good accuracy of the CFD set-up.

Finally, an experimental case of combined pitching and freestream speed fluctuations is used as a validation case. The operating conditions are presented in Tab. 5.3 and the airfoil used in this case is a NACA 0018 airfoil. Figure 5.7 shows the lift coefficient and the excellent agreement between CFD results and wind tunnel results.

As a final remark, it is noted that experimental data on this kind of oscillating airfoils are scarce and that measurements of the drag coefficients are not assessed in the analysed literature.

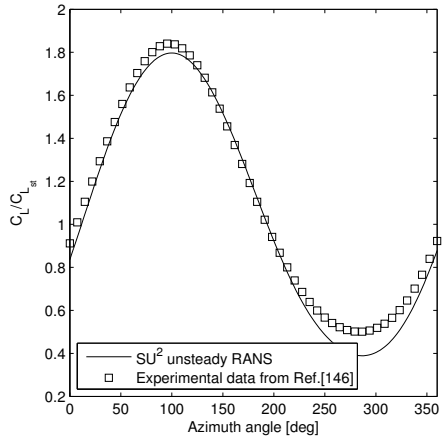

Figure 5.7: Lift coefficient for the pitching airfoil with fluctuating freestream experiment (conditions in Tab. 5.3).

5.1.3 Steady vs. unsteady simulation

In this final section the estimate of the lift and drag coefficients obtained with the CFD set-up are compared to the steady approximation used in the optimization problem. This comparison is obtained by computing the aerodynamic coefficients by means of

	NACA 0012			NACA 23012		
	Unsteady	Steady	Error	Unsteady	Steady	Error
$avg(C_L)$	0.591	0.588	0.5 %	0.756	0.769	1.8 %
$avg(C_D)$ [counts]	126.61	110.18	12 %	121.26	101.51	16 %

Table 5.4: Average drag coefficient and lift coefficient obtained from steady and unsteady computations for the NACA 23012 airfoil.

an unsteady RANS simulation presented in the preceding section and different steady models:

- for the advancing side, an Euler equations simulation with SU^2 is performed, with the set-up that will be later presented for the optimization problem (Sec. 5.2.2). Because the Euler equations are used, a correction associated with viscous effects is added to the inviscid SU^2 estimate combining the van Driest II method and a form-factor correction as presented in Ref. [148];
- for the retreating side, the XFOIL solver is adopted
- for the remaining intermediate points, the MSES code is used.

The results for the NACA 0012 airfoil and NACA 23012 airfoil are presented in Fig. 5.8 and Fig. 5.9, respectively. The average aerodynamic coefficients over the azimuth angle are presented in Tab. 5.4. An excellent agreement is observed with regard to the lift coefficient, both in terms of average value and amplitude oscillations. The difference between the unsteady CFD-based average and the steady value is less than 1 % for the NACA 0012 airfoil and less than 2 % for the NACA 23012 airfoil. In terms of drag coefficient, the discrepancy between the unsteady results and the steady approximation is more pronounced. The unsteady wake effects are completely neglected in the steady approximation and they greatly affect the drag coefficient, more so than in the case of the lift coefficient. This is most likely the reason for this reduction of accuracy. Nevertheless, the average values of the drag coefficient differ by 12 % and 16 % for the NACA 0012 and NACA 23012, respectively. The magnitude of the error for the steady approximation is encouraging and the different steady models are considered an acceptable solution to significantly reduce the computational cost of the optimization loop (both deterministic and robust).

In addition, the work presented in Ref. [24] backs up this conclusion. In fact, Reference [24] demonstrates that optimal solutions are not greatly affected by the use of either an unsteady or steady model, as long as the reduced frequency is moderate.

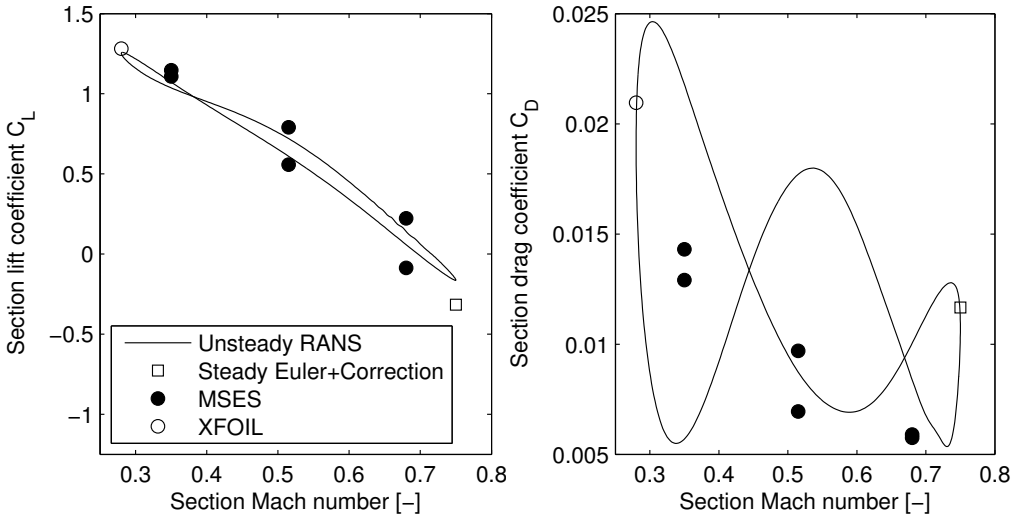

(a) Lift coefficient. **(b)** Drag coefficient.

Figure 5.8: Cycle of lift and drag coefficient of the NACA 0012 airfoil.

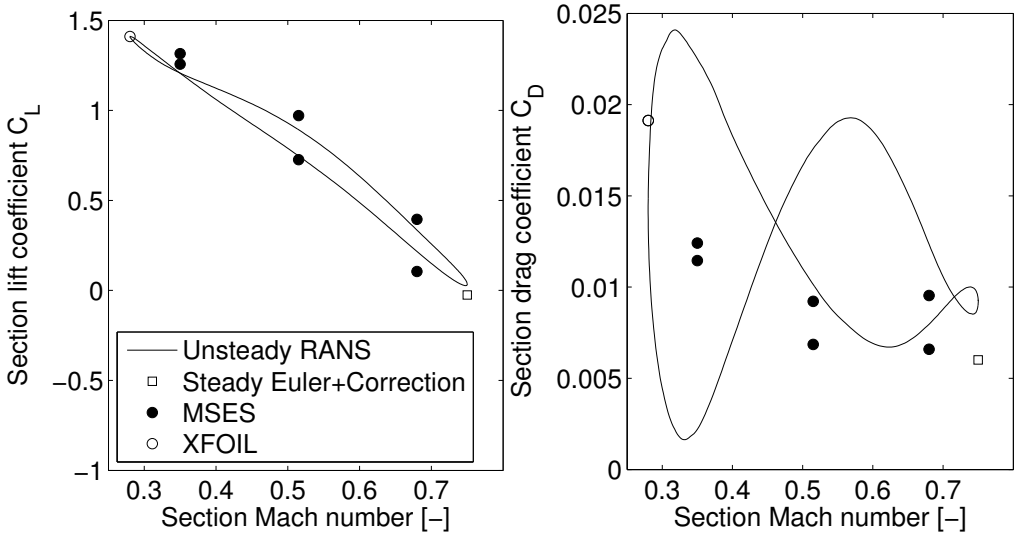

(a) Lift coefficient. **(b)** Drag coefficient.

Figure 5.9: Cycle of lift and drag coefficient of the NACA 23012 airfoil.

This is certainly the case, since the 1/rev frequency is equal to 0.034 and the highest frequency of the angle of attack reaches 0.12.

5.2 Deterministic forward flight optimization

The deterministic optimization of helicopter rotor airfoil in forward flight consists typically in the minimization of the drag coefficient C_D ensuring the satisfaction of constraints on the lift C_L and moment C_M coefficients. Examples of such an approach are given in Refs. [23, 20, 21]. The constraint on the lift coefficient is set to maintain the trim condition of the entire helicopter and reduce the interaction with the blade dynamics.

In this work, a target lift coefficient $\overline{C_L}$ is set to consider a specific trim condition and the moment coefficient is bounded with a given threshold $\overline{C_M}$ to avoid excessive loading on the blade structure. The objectives of the advancing and retreating side are conflicting with one another. In fact, an optimal shape for the advancing side would be a mildly cambered airfoil to postpone drag rise to higher Mach number values, whereas the adoption of a highly-cambered airfoil in the retreating side could better adapt to higher angle of attack, for instance by means of a nose-droop (see Fig. 1.7).

In mathematical terms, two single-objective optimization problems may be for-

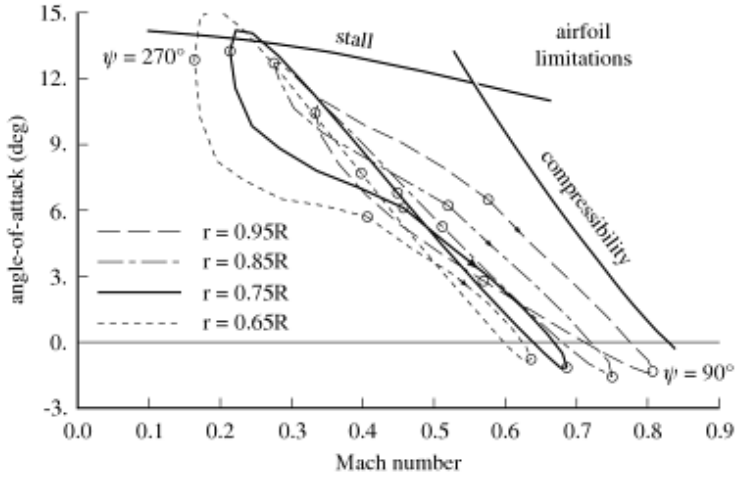

Figure 5.10: Operating conditions for blade sections (picture from Ref. [9]).

mulated as follows

$$\begin{aligned} &\text{minimize: } C_{D_i}(\boldsymbol{x}) \\ &\text{subject to: } C_{L_i}(\boldsymbol{x}) = \overline{C_{L_i}} \quad |C_{M_i}(\boldsymbol{x})| \leq \overline{C_{M_i}} \\ &\text{by changing: } \boldsymbol{x} \end{aligned} \tag{5.6}$$

where $\boldsymbol{x}$ are the design variables. The subscript i is equal to 1 in the advancing side and 2 in the retreating side. The blade section of the Bo105 rotor [117] at $r = 0.85$ is considered and the operating conditions for advancing and retreating sides are taken from the corresponding "figure-of-eight" in the (M, α)–plane from Ref. [9, p.296] presented in Fig. 5.10. Both operating Mach number and the corresponding constraints are reported in Tab. 5.6. The values of the lift constraint $\overline{C_{L_i}}$ are set for each side with the intent of producing a design that improves the performance of the NACA 23012 airfoil, a classical shape for helicopter rotors. In particular, the aerodynamic coefficients of the NACA 23012 airfoil in the operating conditions are presented in Tab. 5.5. It is possible to note that the NACA 23012 airfoil does not provide positive lift in the advancing side, thus a slightly positive value of $\overline{C_{L_1}}$ is set to generate lift. Furthermore, the constraints on the moment coefficient has been enlarged with respect to the value obtained for the NACA 23012 airfoil to give more freedom to the design optimization and explore design solutions which may have greater moment coefficient but better performance in the transonic condition.

To complete the scenario of deterministic optimization problems, a one-objective optimization with a trade-off performance is performed. The optimization problems in this case is defined as follows

$$\begin{aligned} &\text{minimize: } \frac{1}{2}\sum_{i=1}^{2} w_i C_{D_i}(\boldsymbol{x}) \\ &\text{subject to: } C_{L_i}(\boldsymbol{x}) = \overline{C_{L_i}} \quad |C_{M_i}(\boldsymbol{x})| \leq \overline{C_{M_i}} \;\; i = 1, 2 \\ &\text{by changing: } \boldsymbol{x}, \end{aligned} \tag{5.7}$$

where w_i are weights chosen from the drag coefficient of the optimal airfoils obtained in the one-point optimization case (see Tab. 5.6). The result of such an optimization is a compromise solution between advancing and retreating side that is used as a reference of a trade-off optimal solution.

Case	α[deg]	M[-]	Re [-]	C_L [-]	C_D [counts]	C_M [-]
1.	-1.5	0.75	4.6e6	-0.025	86.97	0.02
2.	12.5	0.28	1.7e6	1.411	191.3	0.03

Table 5.5: Aerodynamic coefficients of NACA 23012 airfoil.

Case	M[-]	Re [-]	$\overline{C_L}$	$\overline{C_M}$	w_i [1/counts]
1.	0.75	4.6e6	0.025	0.08	1/2
2.	0.28	1.7e6	1.411	0.04	1/160

Table 5.6: Operating conditions and aerodynamic constraints.

5.2.1 Design variables

Particular attention is drawn to the design variables. First, to describe an airfoil shape with a finite set of variables, a shape parameterization is required. The same shape parameterization used in Sec. 4.2.3 is employed. However in this case, the airfoil shape is decomposed into the camber mean line and the normal thickness distribution. The difference lies in the definition of the thickness distribution which in this case is taken in perpendicular direction with respect to the camber mean line.

Thus, the camber line and thickness distribution are such that the upper surface ζ_u and lower surface ζ_l are obtained applying the thickness perpendicular to the camber line, as follows

$$\begin{aligned} \zeta_u &= \zeta_c + \zeta_t \cos(\epsilon) \\ \zeta_l &= \zeta_c - \zeta_t \cos(\epsilon), \end{aligned} \tag{5.8}$$

where $\epsilon = \arctan\left(\frac{\mathrm{d}\zeta_c}{\mathrm{d}\chi}\right)$. Please note that also the χ coordinate of the resulting airfoil will be affected by this addition of vectors.

With respect to the original formulation, the coordinate of the camber mean line at the trailing edge is set to zero to consider closed trailing edge and zero geometric angle of attack. In addition, the class functions for the camber line C_c and the thickness distribution C_t are defined as follows

$$C_c(\chi) = \chi(1-\chi) \quad C_t(\chi) = \sqrt{\chi}(1-\chi). \tag{5.9}$$

While the class function for the thickness distribution is the one suggested in Ref. [126] to define a rounded nose distribution close to the leading edge, the class function for

the camber presents a linear term to avoid vertical slope of the camber distribution at the leading edge. In addition, the linear term of the class function permits to directly relate the first coefficient of the shape function A_0^c with the slope of the camber line at the leading edge.

The shape function $S(\chi)$ is given by a Bernstein polynomial of order N, whose coefficients represent the design variables $\boldsymbol{x}_s$ of the optimization problem. A convergence study, suggested the employment of 4 design variables for each distributions that yields 10 design variables.

In addition to the coefficients describing the shape of the airfoil, the angle of attack can be regarded as a design variable. In fact, it is possible to solve the constrained optimization problem stated in Eq. (5.6) by acting on the airfoil design variables $\boldsymbol{x}_s$ to minimize the drag coefficient and on the angle of attack to track the target lift coefficient inside two nested loops. The underlying idea is an analogy with the blade collective pitch control: to increase trust it is necessary to act on the collective pitch, which changes the angle of attack of the blade section.

Thus, each airfoil tested in the optimization loop is obtained with a specific set $\boldsymbol{x}_s$, and its performance are evaluated computing the angle of attack α that provides the desired lift coefficient. The nested optimization loops used to implement this procedure are described in detail in Sec. 5.2.3. Thus, the resulting set of design variables is $\boldsymbol{x} = \{\boldsymbol{x}_s, \alpha\}$.

5.2.2 Aerodynamic models

A model describing the aerodynamic load acting on the airfoil is necessary to compute the objective function of the optimization problem. The aerodynamic models used in this optimization problem are different for each side of the blade in order to capture the specific features of the flow in such different operating conditions, while limiting the computational effort.

In the retreating side featuring subsonic condition below the static stall boundary, the XFOIL code [99] is adopted, that is presented in Sec. 4.2.2. This code is chosen because it provides a fast and sufficiently accurate estimation of the aerodynamic force coefficients for a two-dimensional section in the range of angle of attack considered in this problem. Please note that the highest value of the angle of attack in the range is below the static stall limit for most airfoils.

On the other hand, the MSES code used in the hovering case is not suitable for the conditions encountered in the advancing blade. In fact, the coupled Euler/boundary layer code is accurate and robust only with mild compressible effects. For fully

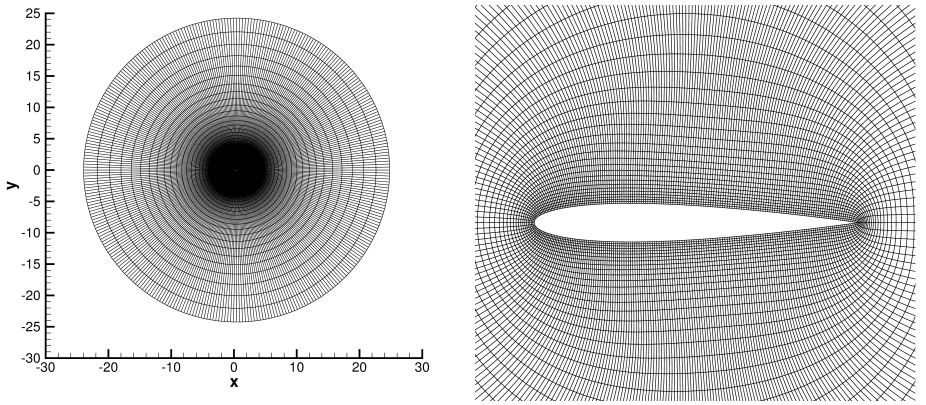

(a) O-mesh. **(b)** Zoom.

Figure 5.11: Coarse structured grid (20480 cells) for Euler computations inside the optimization loop.

developed transonic conditions, like the one experienced in forward flight, a finite volume-based numerical solver of the Euler equations is used. In this thesis the Stanford University Unstructured (SU^2) software suite [143, 146] is employed owing to the accessibility of the source codes, the extensive validation and verification test cases assessed in the literature and its capability to handle mesh deformation and unsteady grid movement. The SU^2 solver provides several space discretization schemes, among which the second-order Jameson-Schmidt-Turkel scheme is employed. An implicit Euler, local time-stepping is used to converge to the steady-state solution, and the GMRES method in conjunction with the LU SGS preconditioner is used to solve the resulting system [147]. This setting for the Euler equations is also used in the benchmark optimization test case presented in Appendix A.

For the computation inside the optimization loop, a two-dimensional coarse structured mesh of 20480 rectangular cells is used (see Fig. 5.11). The mesh extends ± 25 chords around the unitary-chord airfoil. For each airfoil in the optimization loop, the computation runs for 500 iterations starting from a baseline solution of the NACA0012 airfoil. The mesh is deformed by displacing the airfoil boundary cells according to the new geometry and by translating the displacement to the volume cells with a method based on the finite element method discretization of the linear elasticity equations [143]. The element stiffness is set inversely proportional to the wall distance and 500 smoothing iterations are applied to get a regularly deformed mesh.

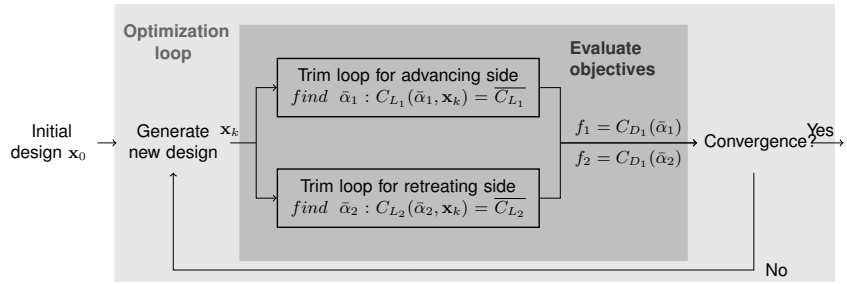

Figure 5.12: Outer NSGA loop for $\boldsymbol{x}_s$ and inner trim loop to determine $\overline{\alpha}$.

5.2.3 Optimization algorithm

To solve the optimization problem in Eq. (5.6) the Non-dominated Sorting Genetic Algorithm is chosen for its ability to explore the design space. Inside the genetic algorithm loop, an inner loop is nested to deal with the lift coefficient constraint (see Fig. 5.12). In particular, for each design vector $\boldsymbol{x}_{s,k}$ of the k-th iteration in the genetic loop, a secant method is used to find the angle of attack $\overline{\alpha}$ that guarantees $C_{L,k} = \overline{C_L}$. The i-th iteration of the secant loop reads

$$\alpha_i = \alpha_{l_i} + \frac{\alpha_{r_i} - \alpha_{l_i}}{C_L(\alpha_{b_i}) - C_L(\alpha_{a_i})} \left(\overline{C_L} - C_L(\alpha_{a_i})\right) \tag{5.10}$$

where α_{l_i} and α_{r_i} are the left and right boundaries of the interval in which the angle of attack is sought and they satisfy: $C_L(\alpha_{r_i}) > \overline{C_L}$ and $C_L(\alpha_{l_i}) < \overline{C_L}$. The left and right value are updated according to the value of the lift coefficient for the current α_i: if $C_L(\alpha_i) > \overline{C_L}$, then $\alpha_{r_{i+1}} = \alpha_i$; if $C_L(\alpha_i) < \overline{C_L}$, $\alpha_{l_{i+1}} = \alpha_i$, and if $C_L(\alpha_i) = \overline{C_L}$ the loop breaks.

The method converges in a few iterations, owing to the quasi-linearity of the lift coefficient for most airfoils in the vicinity of the considered values of the angle of attack. Lack of convergence within few iterations is not critical, because this happens for airfoils with poor values of the lift coefficient that should be discarded anyway. The result of the two nested loops is a set of optimal solutions defined by a design vector $\{\boldsymbol{x}_s, \alpha\}$. It is noted that the inner trim loop targets the equality constraint on the lift coefficient, but the design vector coming out of it may not satisfy the moment coefficient constraint.

5.3 Robust optimization problem

The objective of robust optimization is to design an airfoil that is minimally sensitive to the variation of the operating conditions. In this case, it is considered that the operating conditions at which the blade section will operate are affected by the uncertainty arising due to modelling assumptions of physical parameters necessary at the design stage.

Due to the uncertainty, the objective function f is no longer only a function of the design variables $\boldsymbol{x}$, but it also depends on the uncertain variables $\boldsymbol{\xi}$. Within this uncertain framework a minimization with a target lift coefficient in a specific design conditions loses its meaning. However, a robustly optimal airfoil with a satisfactory lift-to-drag ratio in a range of conditions could be used to trim the aircraft in a specific condition keeping always a low drag coefficient. In addition, if the robustness of the lift-to-drag ratio will translate into the drag coefficient at a specific target lift coefficient, a robust shape would ensure less variability of the required power due to aerodynamic drag.

In particular, for the problem under analysis, the following objectives are defined for each side of the blade.

1. On the advancing side, rotor blade encounters transonic flow, so airfoils should typically be designed to delay drag divergence to higher Mach numbers [1, 11]: however, since the airfoil should feature also a certain amount of lift, the lift-to-drag ratio is maximized for the first condition corresponding to the advancing blade at 90 deg (see Fig. 1.4), with the constraint of providing a lift coefficient equal to or greater than the value of the reference NACA 23012 airfoil. Thus, the lift-to-drag ratio C_L/C_D is the objective f_1 of the optimization: $f_1 = C_L/C_D$.

2. On the retreating side, the increase of the lift coefficient is typically sought [11]: for the blade at 180 deg the ratio $C_L^{3/2}/C_D$ is maximized, which is a measure related to the rotor figure of merit [1] and privileges the lift coefficient over the drag coefficient, that is $f_2 = C_L^{3/2}/C_D$.

The nominal operating conditions of the aerodynamic simulation are set using data from Ref. [9] (see Fig, 5.10). Finally, the constraints on the moment coefficient used in the deterministic case (and presented in Tab. 5.6) are still considered.

In mathematical terms, the resulting optimization problem can be stated as:

$$\begin{aligned} \text{maximize :} \quad & \mu\left(f_i\left(\boldsymbol{x}_s, \boldsymbol{\xi}\right)\right) \\ \text{and minimize :} \quad & \sigma^2\left(f_i\left(\boldsymbol{x}_s, \boldsymbol{\xi}\right)\right) \\ \text{subject to :} \quad & \mathbf{g}\left(\boldsymbol{x}_s, \boldsymbol{\xi}\right) \leq \mathbf{0} \\ \text{by changing :} \quad & \boldsymbol{x}_s, \end{aligned} \tag{5.11}$$

where the moment constraints are collected in vector $\mathbf{g}$. Similarly to the deterministic case, function f_i is equal to f_1 for the advancing side and equal to f_2 in the retreating case. The design variables are the variables defining the shape of the airfoil, while the angle of attack is considered as one of the uncertain variables in the robust design problem.

Indeed, the angle of attack α and the Mach number M encountered by the two-dimensional section of the blade are considered as uncertain, because they are affected by uncertainty on the modeling of the physical parameters considered in the design stage. For instance, both aerodynamic and structural uncertainties, such as blade chord, air density, and rotor angular velocity, may affect the value of the angle of attack and Mach number. Following a probabilistic framework, the uncertain variables are modeled as uniformly distributed random variables around a nominal value. The nominal conditions are reported in Tab. 5.5 and the uncertainty band is set to 5% for the Mach number and for the angle of attack.

The objective functions of Eq. (5.11) are the statistics of the performance f. These are computed by means of an uncertainty propagation technique, that is a method to propagate the uncertainty affecting the operating conditions into the quantity of interest f.

5.3.1 Uncertainty quantification

With regard to the method used to compute the statistics of function f, the uncertainty quantification is performed by the Polynomial Chaos (PC) expansion method [91] presented in Sec. 2.3.2 and assessed in Sec. 4.2.5. Because a uniform distribution is considered also in the forward flight case, Legendre polynomials are employed in the PC expansion. The order of the expansion is retained from the results of Sec. 2.3.2, because a similar convergence analysis has been performed and proved that a fourth-order polynomial is sufficient to accurately capture the statistics. The reconstruction of the statistics is based on the evaluation of the exact function for a set of samples $\boldsymbol{\xi}_k$ in the uncertain variables domain for each design vector $\overline{\boldsymbol{x}_s}$.

Case	Objective	Equality constraints	Inequality constraints	Design variables
DO1	$\min C_{D,1}$	$C_{L_1} = \overline{C_{L_1}}$	$\lvert C_{M_1}\rvert \leq \overline{C_{M_2}}$	$\boldsymbol{x}_s, \alpha$
DO2	$\min C_{D,2}$	$C_{L_2} = \overline{C_{L_2}}$	$\lvert C_{M_2}\rvert \leq \overline{C_{M_2}}$	$\boldsymbol{x}_s, \alpha$
DO0	$\min \sum_{i=1}^{2} w_i C_{D,i}$	$C_{L_1} = \overline{C_{L_1}}$	$\lvert C_{M_1}\rvert \leq \overline{C_{M_1}}$	$\boldsymbol{x}_s, \alpha$
		$C_{L_2} = \overline{C_{L_2}}$	$\lvert C_{M_2}\rvert \leq \overline{C_{M_2}}$	
RO1	$\max \mu(f_1)$ and $\min(\sigma^2(f_1)$		$\lvert C_{M_1}\rvert \leq \overline{C_{M_1}}$	$\boldsymbol{x}_s$
RO2	$\max \mu(f_2)$ and $\min \sigma^2(f_2)$		$\lvert C_{M_2}\rvert \leq \overline{C_{M_2}}$	$\boldsymbol{x}_s$

Table 5.7: Summary of optimization problems.

5.4 Results

As presented in the previous sections, different optimization problems are performed. A summary of the optimization cases presented in this chapter is given in Tab. 5.7. Letters D and R stand for the deterministic and robust cases, while numbers 1 and 2 refer to advancing and retreating side, respectively.

5.4.1 Deterministic optimization

The results of the deterministic optimization problems are: (i) airfoil DA1 with optimal drag coefficient in the advancing side, (ii) airfoil DA2 with optimal drag coefficient in the retreating side, and (iii) airfoil DA0 with optimal weighted drag coefficients. These airfoils are presented in Fig. 5.13 and the characteristics of their geometry is reported in Tab. 5.8. It is possible to note that the optimal airfoil for the retreating side is highly cambered to cope with the higher angle of attack and it has a lower thickness-to-chord ratio to reduce the drag coefficient. The optimal airfoil of the advancing side generates positive lift owing to a slight camber and greater thickness. The slope of the camber line at the leading edge is controlled by the first coefficient of the parameterization thanks to the modification of the CST discussed in Sec. 5.2.1. Both in the overall airfoil geometry and in each specific parameter presented in Tab. 5.8, it is possible to note that airfoil DA0 represents a trade-off solution between DA1 and DA2, although it appears more similar to the DA1 airfoil. This similarity is likely to be due to the choice of the weighting function.

With regard to performance, the drag coefficients of the optimal airfoils are presented in Tab. 5.9. The drag coefficient are expressed in drag counts, where one drag count is equal to 0.0001. The gain with respect to the reference NACA 23012 airfoil

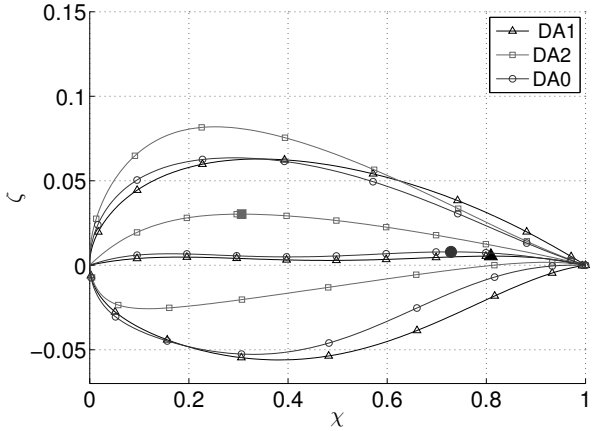

Figure 5.13: Optimal airfoils for optimization problem DA1, DA2 and DA0.

	DA1	DA2	DA0
Maximum thickness	0.1187	0.1049	0.1161
Position of maximum thickness	0.36	0.22	0.31
Maximum camber	0.0054	0.0303	0.0079
Position of max camber	0.81	0.31	0.73
First design variable (A_0^c)	0.065	0.289	0.107
Camber slope at LE [deg]	3.704	16.098	6.108

Table 5.8: Characteristics of the optimal airfoils (for unitary chord).

	C_{D_1} [counts]	$\alpha_1\|_{\overline{C_{L_1}}}$ [deg]	C_{M_1} [-]	C_{D_2} [counts]	$\alpha_2\|_{\overline{C_{L_2}}}$ [deg]	C_{M_1} [-]
DA0	0.8947	-1.059	0.0463	177.8	11.1	-0.0108
DA1	0.6927	-1.002	0.0459	–	–	–
DA2	191.99	-3.101	0.1047	126.0	9.27	-0.0393

Table 5.9: Performance of the optimal airfoils: performance of DA1 in the retreating side is not reported because DA1 airfoil cannot satisfy lift constraint (see Fig. 5.14b).

are presented in Tab. 5.10. The DA1 airfoil and the DA2 airfoil significantly improve the performance of the advancing and retreating side, respectively. Airfoil DA1 is an airfoil which reduces the drag coefficient by smoothing curvature variations of the upper and lower side, yielding a shock-free flowfield (see the Mach number contour in Fig. 5.14a, However, in the retreating side the DA1 airfoil cannot satisfy lift con-

	Percentage gain side 1	Percentage gain side 2	Relative gain side 1	Relative gain side 2
DA0	+99%	+7 %	1/97.2	1/1.08
DA1	+99%	–	1/125	–
DA2	–	+34%	–	1/1.52

Table 5.10: Gain of the optimal airfoils with respect to the NACA 23012 airfoil: gain of DA1 in the retreating side and gain of DA2 in the advancing side are not reported because these airfoils don't satisfy respectively the lift and moment constraint.

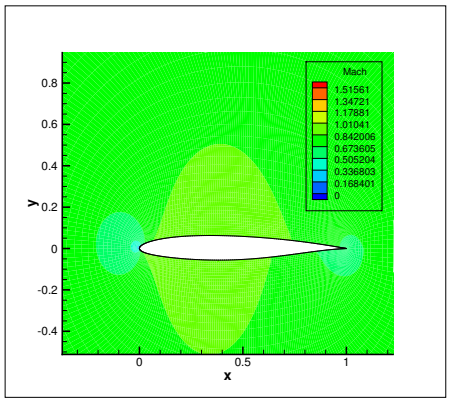

(a) Mach number contour in the advancing side for airfoil DA1.

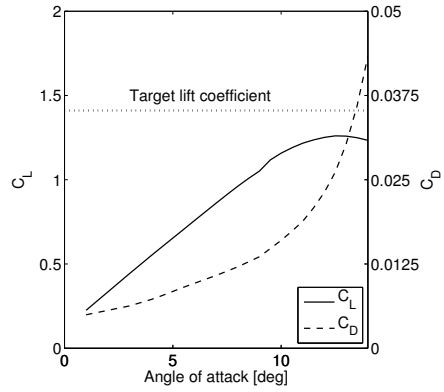

(b) Lift and drag coefficient in the retreating side for airfoil DA1.

Figure 5.14: Analysis of optimal DA1 airfoil.

straint (as presented in Fig. 5.14b), therefore its drag coefficient is not reported in Tab. 5.9 and its gain is not present in Tab. 5.10. Airfoil DA2 reduces the drag coefficient of the NACA 23012 airfoil by 34% by reducing the peak friction coefficient and by postponing transition to turbulent boundary layer, as presented in Fig. 5.15b. In the advancing side, the performance of such an airfoil is very poor and a strong shock wave is present on the lower side (see Fig. 5.15a). The trade-off airfoil has a significant gain in the advancing side, because it is more similar to the DA1 airfoil. As a matter of fact, a shock-free flowfield is obtained with this airfoil (Fig. 5.16a), whereas in the retreating side an earlier transition to a turbulent boundary layer and a higher friction coefficient are present (Fig. 5.16b).

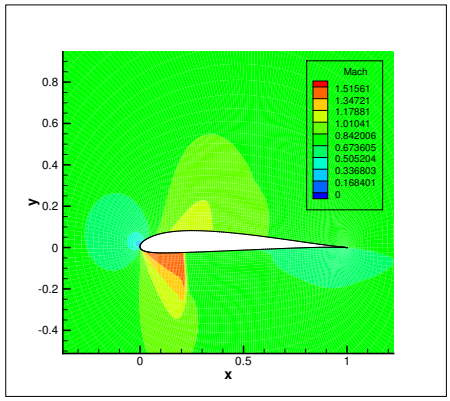

(a) Mach number contour in the advancing side for airfoil DA2.

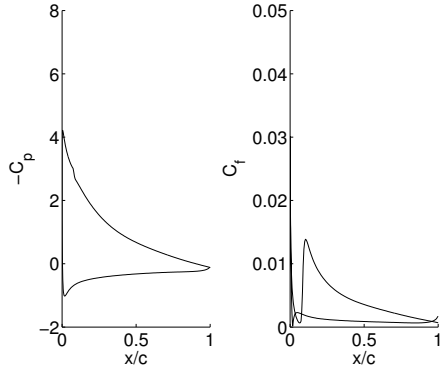

(b) Pressure and friction coefficient in the retreating side for airfoil DA2.

Figure 5.15: Analysis of optimal DA2 airfoil.

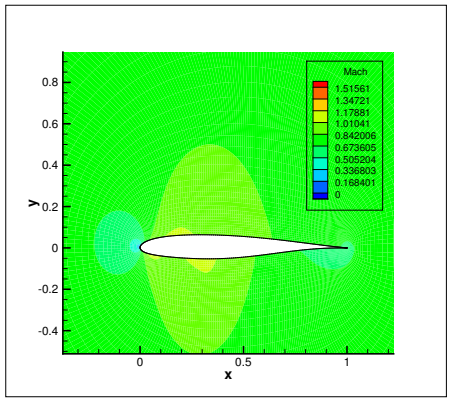

(a) Mach number contour in the advancing side for airfoil DA0.

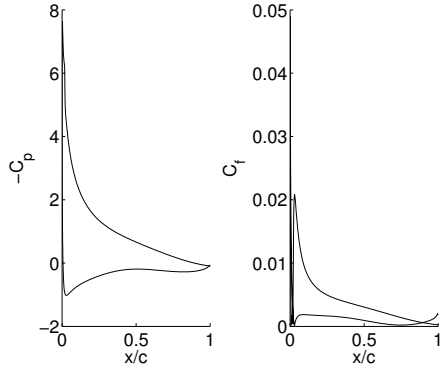

(b) Pressure and friction coefficient in the retreating side for airfoil DA0.

Figure 5.16: Analysis of optimal DA0 airfoil.

5.4.2 Comparison with robust optimization

Two different robust optimization problems are performed, the first one for the advancing side (problem RO1) and the second one for the retreating side (RO2). The Pareto front for the case of the advancing side is presented in Fig. 5.17a where the two objective functions are the mean value and the variance of the lift-to-drag ratio. Please note that in the figure the mean value and the variance are obtained with a scaling factor of 1/10 acting on f_1, because the value of $f_1 = C_L/C_D$ is very large due to the inviscid estimate of the drag coefficient which is very small (on the or-

Airfoil	$C_{D,1}\|_{\overline{C_{L_1}}}$ [counts]	$\alpha_1\|_{\overline{C_{L_1}}}$ [deg]	$\mu_{C_{D,1}}$ [counts]	$\sigma^2_{C_{D,1}}$ [counts^2]	$(\sigma/\mu)_{C_{D,1}}$ [-]
DA0	0.895	-1.059	5.26	65.75	1.54
DA1	0.693	-1.002	2.16	11.12	1.54
RA1	0.898	-1.510	2.03	5.668	1.17

Table 5.11: Drag coefficient of the deterministic airfoil minimizing C_{D_1} (DA1), the trade-off airfoil(DA0) and the airfoil selected from the robust front in the advancing side (RA1).

der of a drag count) for optimal airfoils. The front presents very robust solution that however comes at the expense of poor performance. In the higher part of the front, solutions with high mean value are found. Figure 5.17a does not present the results of the reference airfoil, i.e. the NACA 23012 airfoil, because this airfoil has a negative lift-to-drag ratio (due to the negative lift).

To compare the results obtained from the Pareto front in RO1 and the deterministic result DA1, a solution or a group of solutions from the front in Fig. 5.17a should be selected. To do this, the target lift coefficient $\overline{C_{L_1}}$ is compared to the lift coefficients obtained in the samples used for the reconstruction of the mean value and variance in the UQ loop. The lift coefficients for the solutions of the front are plotted in Fig. 5.17b. From this set airfoil RA1 is selected, which is the airfoil that exhibits a lift coefficient equal to the target value inside the uncertainty range considered and a

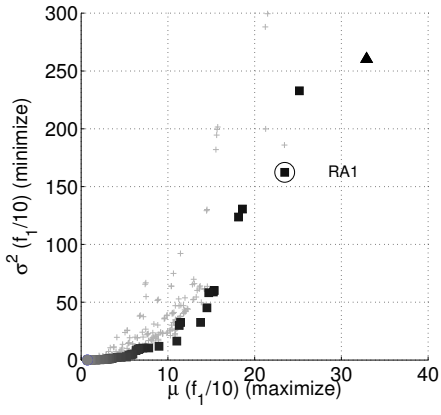

(a) Pareto front.

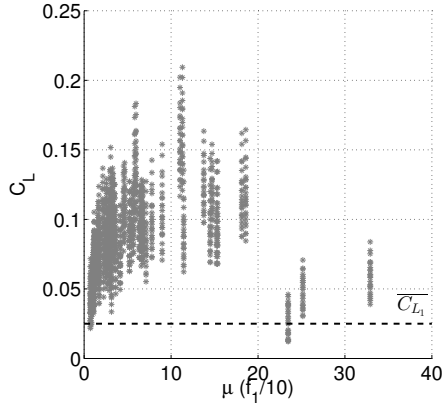

(b) Lift coefficient in the samples used in RO1 for the airfoils of the Pareto front.

Figure 5.17: Results of the robust optimization for the advancing side (RO1) .

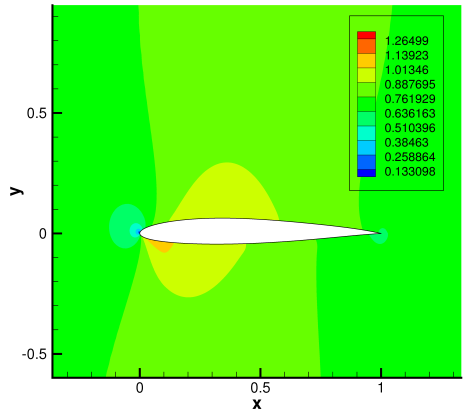

Figure 5.18: Mach number contour of RA1 airfoil at $\alpha_1|_{\overline{C_{L_1}}} = -1.51$ deg.

Airfoil	$\mu_{C_{L,1}}$	$\sigma^2_{C_{L,1}}$	$(\sigma/\mu)_{C_{L,1}}$	$\mu_{C_{M,1}}$	$\sigma^2_{C_{M,1}}$	$(\sigma/\mu)_{C_{M,1}}$
DA0	0.0257	4.52e-05	0.26132	0.0459	9.87e-07	0.0217
DA1	0.0211	4.04e-05	0.30151	0.0457	1.81e-06	0.0294
RA1	0.0253	8.20e-05	0.35816	0.0688	1.47e-06	0.0176

Table 5.12: Lift and moment coefficient of the deterministic airfoil minimizing C_{D_1} (DA1), the trade-off airfoil(DA0) and the airfoil selected from the robust front in the advancing side (RA1).

high mean value of the lift-to-drag ratio. Airfoil RA1 is also highlighted in Fig. 5.17a. The value of the angle of attack for which the lift coefficient of airfoil RA1 is equal to $\overline{C_{L_1}}$ is called $\alpha_1|_{\overline{C_{L_1}}}$ and it is equal to -1.51 degrees at the Mach number used for the deterministic optimization DO1. As presented in Tab. 5.11, in this condition the drag coefficient is higher than the one provided by airfoil DA1, but it is still lower than one drag count thanks to the fact that even this airfoil does not present extended region of supersonic flow and shock waves (see Fig. 5.18). This means that this airfoil is capable of satisfying the trim condition with a very small drag penalty with respect to the deterministic airfoil. However, if an uncertainty band is considered around the condition ensuring trim requirement, both the mean value and the variance of the drag coefficient for airfoil RA1 are smaller than the values of airfoil DA1 (see Tab. 5.11). The table also presents the coefficient of variation, defined as the ratio of the standard deviation σ to the mean value μ, to appreciate the dispersion with respect to the mean value.

To complete the analysis, Table 5.12 presents the mean value and variance of the lift and moment coefficients for airfoil RA1 and airfoil DA1. In the table, it is interesting to note that the mean value of airfoil RA1 is closer to the target value. It is noted that the mean value of the drag coefficient, lift coefficient and moment coefficient are obtained with the fourth-order PC expansion assessed in Sec. 4.2.5. Despite the presence of shock waves in this range of conditions, the integral values remain smooth throughout the stochastic space and a fourth-order polynomial provides sufficiently accurate estimates of the statistics.

The same analysis can be performed for the Pareto front of the retreating side. Figure 5.19a presents the optimal set of solutions obtained from the robust optimization and the result of the reference airfoil, the NACA 23012. As emerged from the results in the hovering case (Sec. 4.3.4) and in the advancing side, interesting solutions from the front are typically those with high mean performance. A possible criterion for decision making in the post-processing of the front is then the selection of non-dominated solutions with mean value higher than the value of the reference airfoil. This criterion helps removing the solutions with poor performance, but it still retains solutions with variance lower than the reference result thanks to the quality of the Pareto front.

For this subset of airfoils, the comparison between the robust front and the deterministic airfoil is performed in the same way presented in the advancing side. Figure 5.19b shows the lift coefficient of the samples used in RO2 with mean value

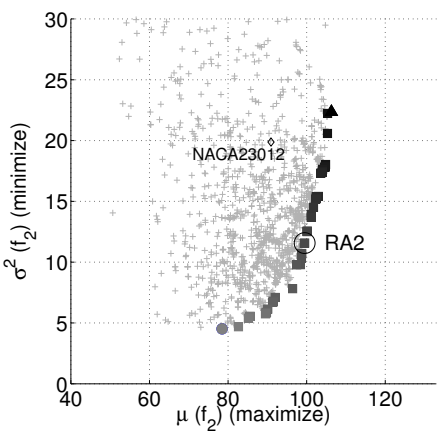

(a) Pareto front.

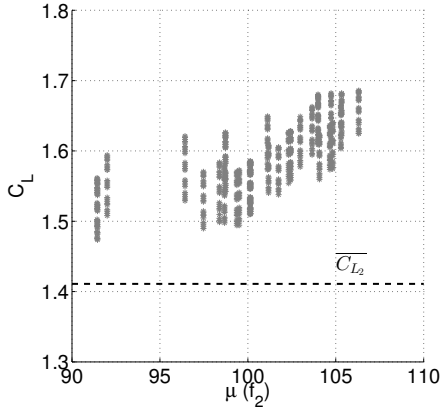

(b) Lift coefficient in the samples used in RO2 for the airfoils of the Pareto front.

Figure 5.19: Results of the robust optimization for the retreating side (RO2) .

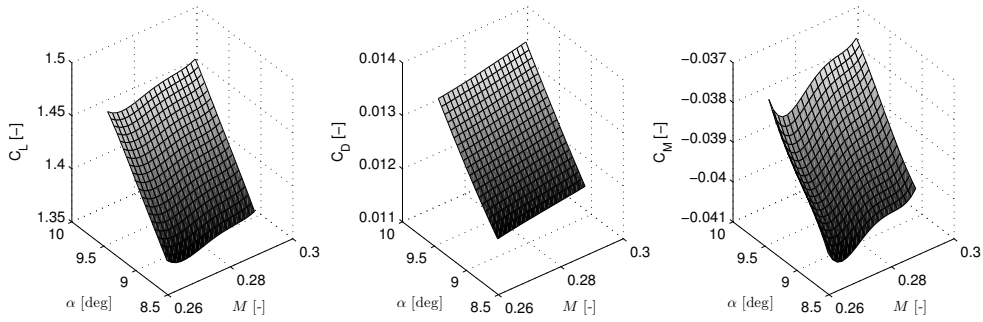

Figure 5.20: Load of DA2 in the uncertainty band around $\alpha_1|_{\overline{C_{L_1}}}$.

Airfoil	$C_{D,2}\|_{\overline{C_{L_2}}}$ [counts]	$\alpha_1\|_{\overline{C_{L_2}}}$ [deg]	$\mu_{C_{D,2}}$ [counts]	$\sigma^2_{C_{D,2}}$ [counts^2]	$(\sigma/\mu)_{C_{D,2}}$ [-]
DA0	177.8	11.13	178.4	91.32	0.0536
DA2	126.0	9.27	126.2	28.40	0.0422
RA2	148.1	10.66	148.4	33.83	0.0392

Table 5.13: Drag coefficient of the deterministic airfoil minimizing C_{D_2} (DA2), the trade-off airfoil(DA0) and the airfoil selected from the robust front in the retreating side (RA2).

higher than the reference value. In this case, every airfoil in the optimal set provides a lift coefficient that is greater than the target value $\overline{C_{L_2}}$. Thus, to select an airfoil from the front for comparison, airfoil RA2 has been chosen as a trade-off between the two objectives.

For this airfoil an angle of attack equal to 10.66 deg is needed to satisfy the trim condition. By applying the same uncertainty range used in the robust optimization case, it is possible to compute the mean value and the variance of the drag coefficient of RA2. These values are higher with respect to the value computed for airfoil DA2 (see Tab. 5.13. In fact, the optimal deterministic airfoil has a very low value of drag coefficient because it satisfies the lift constraint at a very low angle of attack (9.27 deg) thanks to the highly cambered mean line. However, airfoil DA2 cannot satisfy the moment coefficient constraint in the uncertainty range. As presented in Fig. 5.20, the moment coefficient exceeds the threshold value $\overline{C_{M_2}}$. On the contrary, the robust optimization formulation guarantees that the optimal airfoils would not incur moment penalties throughout the uncertainty range.

Finally, the discrepancy in the performance of airfoil DA2 and airfoil RA2 is due to the fact that the lift coefficient of the robust airfoils are significantly greater than

Airfoil	$\mu_{C_{L,2}}$	$\sigma^2_{C_{L,2}}$	$(\sigma/\mu)_{C_{L,2}}$	$\mu_{C_{M,2}}$	$\sigma^2_{C_{M,2}}$	$(\sigma/\mu)_{C_{M,2}}$
DA0	1.3998	8.37e-04	0.207	-1.069e-02	3.84e-06	-0.18
DA2	1.4113	8.67e-04	0.021	-3.924e-02	6.51e-07	-0.02
RA2	1.4016	6.24e-04	0.018	-1.383e-02	7.68e-06	-0.20

Table 5.14: Lift and moment coefficient of the deterministic airfoil minimizing C_{D_2} (DA2), the trade-off airfoil(DA0) and the airfoil selected from the robust front in the retreating side (RA2).

the target lift coefficient. Thus, the deterministic optimization and the robust optimization are essentially exploring different regions of the design space, according to the specific objectives. To some extent, a fair comparison between the two methods is meaningful if the two methods share the "same" objectives. This means, for example, that for this specific case the target lift coefficient should be higher for the comparison with the robust airfoils. Another possibility is to define the nominal condition for the robust case starting from the deterministic results. This would at least ensures the presence of the deterministic airfoil in the robust population (as long as it satisfies the contraints). This option appears more sound and it is the ones used in the final application of Chapter 6.

5.5 Validation with the azimuth angle

Up to this point, the performance of the optimal airfoils is considered only in the advancing and retreating side conditions. However, the blade section encounters different operating conditions over the period of rotation of the blade and an assessment of the performance with the azimuth angle is appropriate. To do this, the variation of the angle of attack, Mach number and Reynolds number with the azimuth angle must be computed. This computation is based on the varying freestream Mach number and angle of attack plotted in Fig. 5.10. The variation of angle of attack and Mach number presented in the figure are described by means of the following time law

$$\begin{aligned} M(t) &= M_\infty \left(1 + \lambda \sin(\Omega t)\right) \\ \alpha(t) &= \alpha_0 + \sum_{i=1}^{n_h} \Delta\alpha_i \sin(i\Omega t + \overline{\psi}_i), \quad n_h = 4 \end{aligned} \tag{5.12}$$

where Ω is the blade 1/rev frequency (such that the azimuth angle ψ satisfies $\psi = \Omega t$), M_∞ is the freestream Mach number and λ is the percentage variation of Mach

number. The angle of attack is described by means of a sine series, where α_0 is the mean angle of attack, $\Delta\alpha_i$ is the angle of attack variation of the i-th term in the series and $\overline{\psi}_i$ is the i-th phase angle. This structure of the variations has been obtained with a convergence analysis based on a sine series with increasing number of terms. Figure 5.21a shows the reference $\alpha - M$ cycle obtained from Ref. [9] (i.e. the curve corresponding to $r = 0.85$ in Fig. 5.10) with a black line and the samples obtained with the variation in angle of attack and Mach number described in Eq. (5.12) and presented in Fig. 5.21b. Please note that a variation in Reynolds number is associated with the variation in Mach number. The values of the variables in Eq. (5.12) are presented in Tab. 5.15. They directly correspond to the extrema used in the optimization loop: (i) a maximum value of the angle of attack equal to 12.5 deg at a Mach number of 0.28 and (ii) a minimum value equal to -1.5 deg at a Mach number equal to 0.75 (see Tab. 5.5).

For the baseline case (i.e. NACA 23012 airfoil), the variables in Eq. (5.12) are computed to represent the cycle depicted in Fig. 5.10. However, for the optimal airfoils new values of the amplitude of the sine terms describing the angle of attack must be computed, because the target lift coefficient is obtained in both sides at a

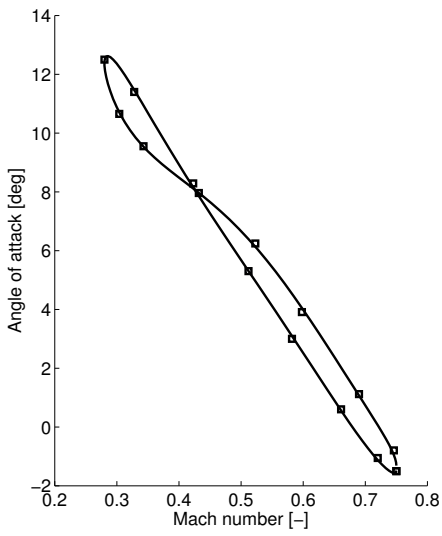

(a) Cycle $\alpha - M$: reference cycle from Fig. 5.10 (line) and samples from approximation in Eq. (5.12) (squares).

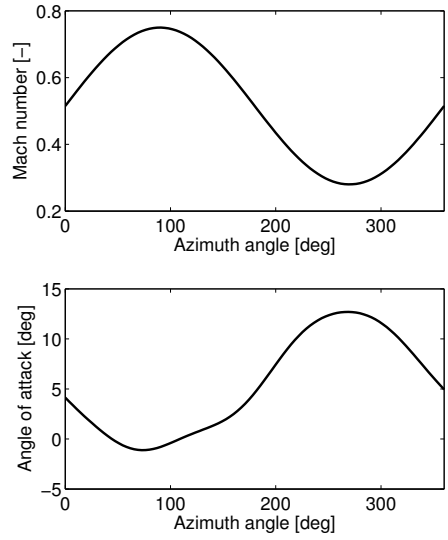

(b) Corresponding time varying Mach number (top) and angle of attack (bottom).

Figure 5.21: Time varying Mach number and angle of attack encountered by the blade element (Eq. (5.12)).

M_∞ [-]	λ [-]	α_0 [deg]	$\Delta\alpha_1$ [deg]	ϕ_1 [deg]	$\Delta\alpha_2$ [deg]	ϕ_2 [deg]	$\Delta\alpha_3$ [deg]	ϕ_3 [deg]	$\Delta\alpha_4$ [deg]	ϕ_4 [deg]
0.515	0.45	5.36	6.84	-173.8	0.67	-86.0	0.28	112.3	0.2	-22.5

Table 5.15: Variables of time-varying angle of attack and Mach number in Eq. (5.12).

	DA0	DA1	DA2	RA1	RA2
α_0 [deg]	5.04	5.75	3.09	5.49	4.58
Δ_1 [deg]	5.88	6.51	5.97	6.76	5.87

Table 5.16: Variables of time-varying angle of attack (Eq. (5.12)) for optimal shapes.

different angle of attack. In other words, the amplitude of the time-varying angle of attack presented in Fig. 5.21b must be corrected for each optimal shape. In general, the azimuth variation of the angle of attack is caused by the blade cyclic pitch control, the induced velocity of the rotor, the blade dynamics and aeroelasticity. Thus, as a first approximation it is possible to associate the 2/rev, 3/rev and 4/rev contributions in Eq. (5.12) to the nonlinear contributions depending on the aerodynamic effects of vortices, wakes and the blade dynamics. These contributions are essentially function of the lift produced by the blade, which remains unchanged for the optimal airfoils. However, the first sine term (1/rev) is directly related to the cyclic pitch control and can be adjusted to maintain the trim requirement.

After computing the angle of attack variation for each optimal shape (each term of Eq. (5.12) is presented in Tab. 5.16), the lift and drag coefficients are obtained at eight location over the azimuth angle. The steady model used to compute the coefficients in this validation over the azimuth angle are the following:

- for the advancing side, an Euler equations simulation with SU^2 is performed, with the set-up used in the optimization loop (Sec. 5.2.2). Because the Euler equations are used, a correction associated with viscous effects is added to the inviscid SU^2 estimate combining the van Driest II method and a form-factor correction as presented in Ref. [148];
- for the retreating side, the XFOIL solver is adopted
- for the remaining intermediate points, the MSES code is used

The aerodynamic coefficients for the deterministic optimal airfoils (DA0, DA1, DA2) are presented in Fig. 5.22. It is possible to note that the cycle of the lift co-

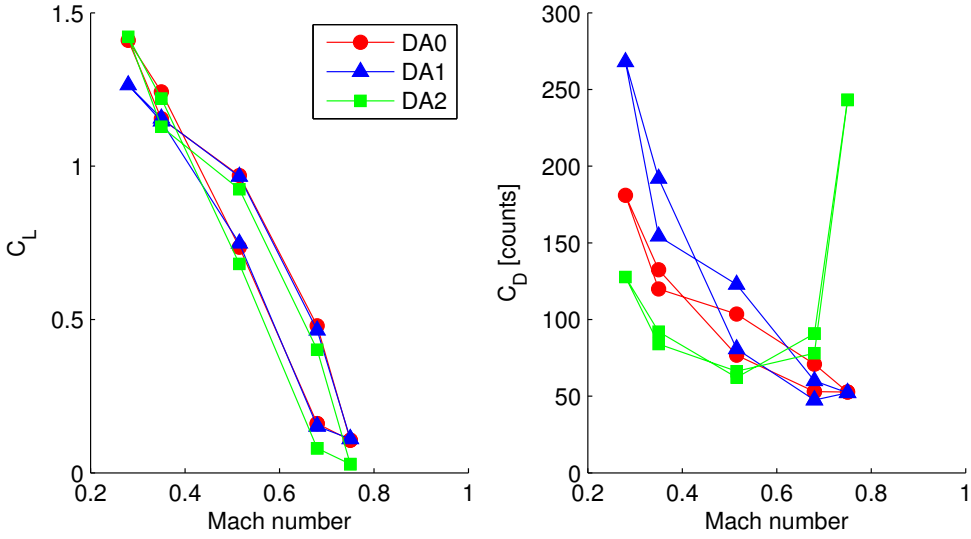

Figure 5.22: Lift and drag coefficient with azimuth for deterministic optimal airfoils.

efficient is similar for the three airfoils owing to the equality constraint on the lift coefficient used in the optimization loop. Please recall that airfoil DA1 is not capable of satisfying the trim constraint in the off-design condition (i.e. the retreating side), as emerged from Fig. 5.14b.

The average of the drag coefficient over the azimuth angle $avg(C_D)$ is then computed and the results are summarised in Tab. 5.17. The table also presents the results for the airfoils selected from the robust Pareto front. i.e. airfoil RA1 and airfoil RA2. The associated cycles of the aerodynamic coefficients is presented in Fig. 5.23. As expected, airfoil DA0 provides the best average drag coefficient over the azimuth angle. The penalty of airfoil DA1 in the off-design condition (i.e. the retreating side) is higher than the penalty suffered by airfoil DA2 in the off-design conditions. As a result, airfoil DA2 has a better average value. Airfoil RA2 has a small penalty with respect to its deterministic counterpart and it also presents a fairly symmetrical cycle. Airfoil RA1 on the other hand has a better average value than the corresponding deterministic airfoil.

Given these results, it could be argued that obtaining DA1 and DA2 from a combined two-point bi-objective problem could result in lower drag coefficient over the azimuth angle for DA1 and DA2. In fact, with such an approach the drag minimization and lift/moment constraints must be simultaneously satisfied for the retreating and advancing side conditions, resulting in a better average value over the azimuth.

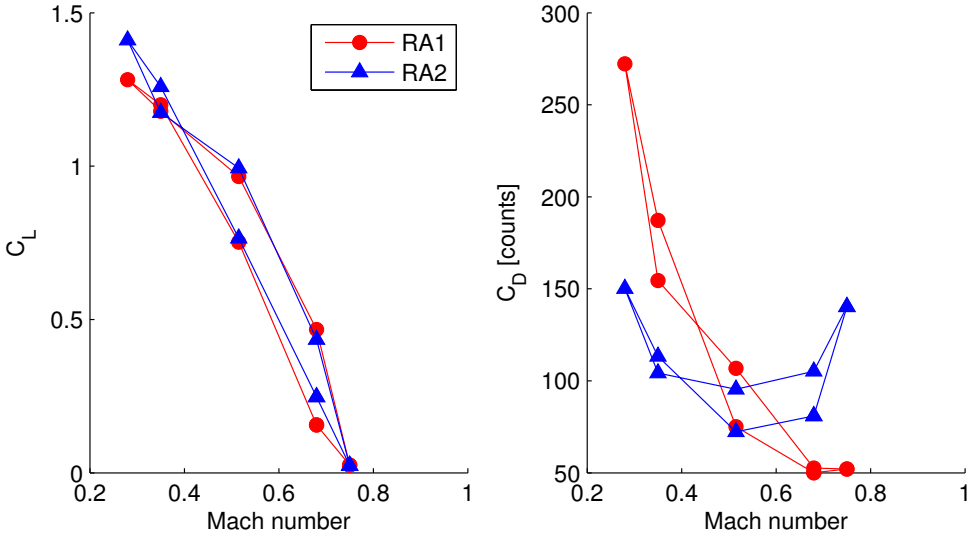

Figure 5.23: Lift and drag coefficient with azimuth for robust optimal airfoils.

Airfoil	DA0	DA1	DA2	RA1	RA2
$avg(C_D)$ [counts]	98.76	122.22	105.5	118.84	107.71

Table 5.17: Average drag coefficient for optimal airfoils in forward flight.

This solution was not pursued, because with a bi-objective approach the lift/moment constraints in the off-design condition limits the gain in each operating condition. The idea of the single-point deterministic optimization DO1 and DO2 is to define the "best", reference airfoils for the advancing and retreating side regardless of their performance in the off-design point. This choice is motivated by the use of these airfoils in the definition of the "ideal" morphing, discussed in the following chapter. Nonetheless, a bi-objective optimization will result in lower average drag coefficient of airfoils DA1 and DA2, so future activities may include a comparison with such an approach.

CHAPTER 6

Optimization of morphing airfoils

This chapter is devoted to the definition of optimal shapes for morphing airfoils to be applied to helicopter rotor blades. The morphing strategy considered in the work is conceived with the intent of changing its shape at the 1/rev frequency to enhance aerodynamic performance. Please note that the optimization of morphing airfoils presented here only takes into account the aerodynamic performance. As a matter of fact, optimization of morphing airfoils is typically performed separately from the morphing mechanism design [40] and a two-step procedure is typically pursued: (i) an aerodynamic optimization to set the optimal shape with respect to performance, and (ii) a topological optimization to find the best internal configuration to fit the best aerodynamic shape. The thesis deals with the first step, where technological aspects are inserted through geometrical constraints. In fact, the thesis presents an exploratory work on morphing geometries which aims at understanding the relationship between shape degree of freedom and actual aerodynamic gain. The continuous development of mechanisms used for morphing, such as cell-based structural topologies [149], may offer more flexibility with respect to the classical choices for morphing geometries. Thus, exploring and demonstrating the gain of the aerodynamic shape may drive the development of new mechanism for the realization of morphing structures.

In the following, a discussion on two different choices of morphing strategy is presented in Sec. 6.1 to assess the capability of morphing airfoils. In this first part, only deterministic evaluation are presented, because deterministic optimal shapes

(and corresponding operating conditions) represent the good starting point of robust approaches, as discussed in the preceding chapters. This discussion demonstrates how camber morphing is more effective to enhance the aerodynamic performance. Then, optimal shapes from deterministic and robust approaches are computed and presented in Sec. 6.2. Finally, the validation over the azimuthal angle is presented in Sec. 6.3.

6.1 Effects of camber and thickness morphing

In this section, the effect of modification of the thickness and the camber of the airfoils is discussed. An airfoil geometry can be characterized by the coordinates of the upper and lower surface. However, an airfoil can also be described by the construction of a thickness shape distributed around a camber line. This description dates back to the 1920s and 1930s, when the NACA 4-digit and 5-digit airfoils were generated by superimposing a simple meanline shape with a thickness distribution that was obtained by fitting a group of popular airfoils of the time.

The camber mean line mainly affects the chordwise load distribution, the angle of zero lift, the pitching-moment coefficient, the slope of the lift curve and the approximate position of the aerodynamic center [150, p.65]. Effects of increasing the camber line are generally a lower zero-lift angle of attack and a higher maximum lift coefficient. The counterpart of this is an increase in the drag coefficient, especially at low angle of attack, and an increase in the pitching moment coefficient. This latter problem, which is crucial in helicopter blade design, led to the almost universal use of symmetric airfoil sections on helicopter developed before the 1960s [1, p.360]. However, with the advent of new camber line distribution, such as the NACA 23012 airfoil which is designed to limit the moment coefficient, cambered airfoils were introduced also in rotorcraft blades. A clear example of the effect of camber on moment coefficient is given by reflexed airfoils, which have reduced camber over the afterward section producing less lift over this region, thereby less nose-down pitching moment. In this case the aft section is actually generating downforce and the moment coefficient at zero lift is positive.

Effects of thickness distribution includes the increase of maximum lift coefficient with increasing thickness, at least for thickness-to-chord ratio in the range used for aeronautical applications (i.e. around 16% [150, p.135] for the NACA 4- and 5-digit airfoils), above which the effect of thickness become detrimental for separation. Thickness distribution has a significant impact on profile drag and on wave drag. With regard to the latter, earlier drag rise and lift break are experienced for thicker

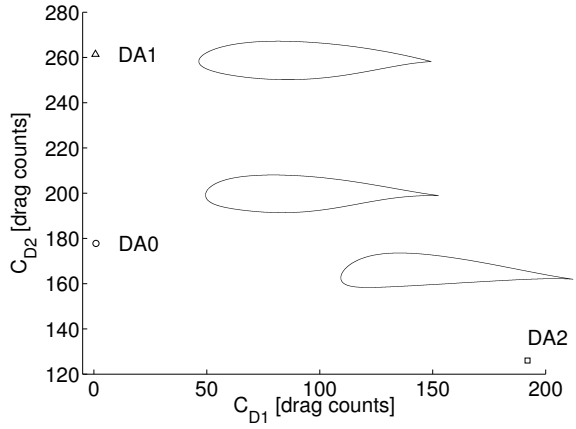

Figure 6.1: Optimal deterministic shapes and drag coefficient in the advancing and retreating sides

airfoil [150, p.274], although this is in general an effect of both thickness-to-chord ratio and distribution, and may vary from airfoil to airfoil.

6.1.1 Camber and thickness morphing from airfoil DA0

In the previous chapter, several optimal shapes have been presented and discussed. In particular, regarding the deterministic optimal shapes, three different shapes have been obtained: (i) DA0, the optimal trade-off airfoil, (ii) DA1, the optimal airfoil for the advancing side and (iii) DA2, the optimal airfoil for the retreating side. These airfoils are presented in Fig. 6.1 in the space of the objectives of the deterministic optimization (i.e. the drag coefficient in the advancing and retreating side). Airfoil DA0 represents the reference optimal solution for a non-morphing configuration. Instead, airfoil DA1 and DA2 give an indication of the best possible gain of a morphing strategy. In fact, the maximum, "ideal" gain of morphing airfoils consists of the gain of a morphing strategy that could morph the section of a blade from the shape of airfoil DA1 in the advancing side to the shape of airfoil DA2 in the retreating side. This is depicted in Fig. 6.2. This solution would have the lowest drag coefficient in the advancing side and at the same time the lowest drag coefficient on the retreating side. The morphing strategy is ideal in this case, because in practice any morphing mechanism would impose geometry constraints to the airfoil shape between the advancing side and retreating side that are not considered in this first instance. Indeed, state-of-the-art technology permits only local deformation of the camber line, as presented

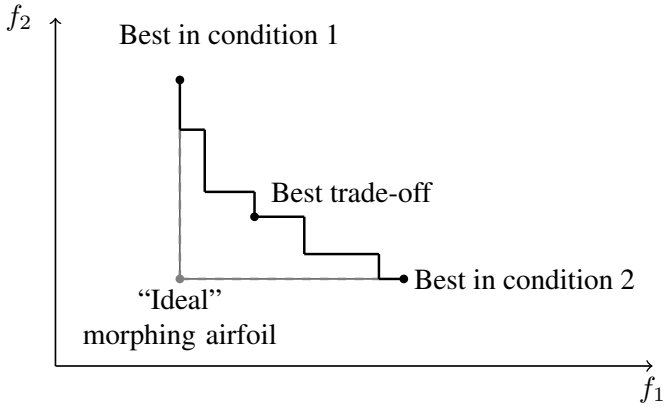

Figure 6.2: Scheme of two-point optimization and "ideal" morphing airfoil.

	DA1	DA2
Gain w.r.t DA0	22 %	29 %

Table 6.1: Gain of an "ideal" morphing with respect to the trade-off non-morphing solution DA0.

	Thickness	Camber
Thick DA0 + Camb DA1	DA0	DA1
Thick DA0 + Camb DA2	DA0	DA2
Thick DA1 + Camb DA0	DA1	DA0
Thick DA2 + Camb DA0	DA2	DA0

Table 6.2: Summary of airfoils obtained by imposition of either the thickness or camber line of DA1 and DA2 from DA0.

in Sec. 1.3, although an active field of research on materials is developing ways of increasing the degree of freedom of morphing geometries that could be adopted in future research activities. Nevertheless, the ideal gain is of the greatest significance, because it sets a reference gain. The gain is computed with respect to the optimal non-morphing geometry, which in this case is the trade-off airfoil DA0. As presented in Tab. 6.1, such a "perfectly" morphing airfoil could gain 22% of the drag coefficient in the advancing side and 29% in the retreating side, while maintaining the trim lift coefficient and moment constraint.

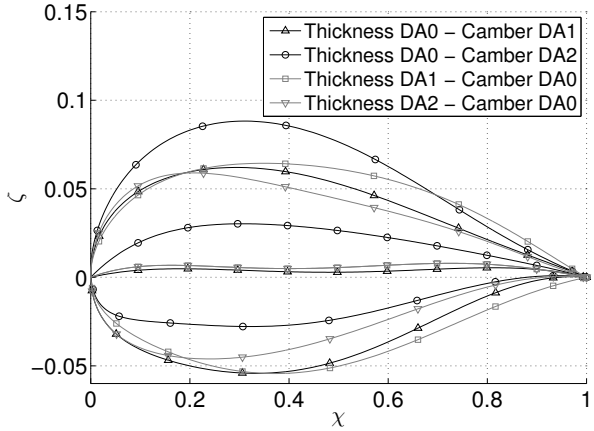

Figure 6.3: Airfoils obtained by changing either the thickness or camber line from airfoil DA0.

	C_{D_1} [counts]	$\alpha_1\|_{\overline{C_{L_1}}}$ [deg]	C_{D_2} [counts]	$\alpha_2\|_{\overline{C_{L_2}}}$ [deg]	Gain side 1	Gain side 2
DA0	0.8947	-1.059	177.8	11.1		
Thick DA0 + Cam DA1	0.805	-0.945	192.3	11.502	+10%	
Thick DA0 + Cam DA2	151.53	-3.316	141.6	9.55		+20%
Thick DA1 + Cam DA0	0.801	-1.128	–	–	+10.5%	
Thick DA2 + Cam DA0	11.26	-1.028	166.0	10.881		+6%

Table 6.3: Drag coefficient of airfoils obtained from deterministic optimal airfoils (see Tab. 6.2).

To further support the choice of camber line morph over thickness distribution morph, a comparison is made among airfoils generated from airfoil DA0 imposing either the thickness distribution of airfoil DA1 and DA2 or the camber line. In other words, four airfoils are constructed by fixing either the thickness or the camber line of airfoil DA0 and imposing the camber line or the thickness of airfoils DA1 and DA2. The airfoils are summarised in Tab. 6.2 and the resulting shapes are presented in Fig. 6.3. The first two solutions refers to a camber morphing in which the camber is modified and the thickness is held constant (and equal to the one of the best non-morphing solution). On the other hand, the last two solutions represent a thickness morphing, where the camber is held fixed to the one of airfoil DA0.

For these airfoils the angle of attack for which the lift coefficient satisfies the trim constraint is computed and the corresponding drag coefficients are evaluated. Results

$\alpha_1\|_{\overline{C_{L_1}}}$ [deg]	$C_{D_1}\|_{\overline{C_{L_1}}}$ [counts]	$C_{M_1}\|_{\overline{C_{L_1}}}$	$\alpha_2\|_{\overline{C_{L_2}}}$ [deg]	$C_{D_2}\|_{\overline{C_{L_2}}}$ [counts]	$C_{M_2}\|_{\overline{C_{L_2}}}$
-3.316	151.53	0.109153	9.55	0.01416	-0.0414

Table 6.4: Performance of Thick DA0 + Cam DA2 airfoil in both sides.

are presented in Tab. 6.3. The same model used in Chapter 5 are employed, thus in the advancing side the Euler equations-based CFD model is used and XFOIL in the retreating side (see Sec. 5.2.2). It is possible to note that the solutions obtained by holding the thickness fixed and changing the camber line (Thick DA0 + Cam DA1 and Thick DA0 + Cam DA2) improves the gain of the reference DA0 airfoil by 10% in the advancing side and 20% in the retreating side. With respect to the ideal gains presented in Tab. 6.1, this solution restores 50% of the ideal gain in the advancing side and 90% in the retreating side. On the other hand, solutions obtained by modifying the thickness distribution present a similar gain on the advancing side (10.5%), but a smaller gain in the retreating side (6%). These results show that the camber morphing gets closer to the ideal morphing gain and it is therefore more effective. Furthermore, it seems that the thickness and camber line equally contributes to the ideal gain in the advancing side, while in the retreating side the camber line plays a key role. As a matter of fact, in the advancing side transonic flow is experienced, and the minimization of the drag coefficient corresponds to the minimization of wave drag. This contribution equally depends on camber and thickness distributions over chordline. On the other hand, high angle of attack and subsonic flow are encountered in the retreating side, thus the gain is basically built upon the adaptation of the airfoil to the high angle of attack achieved through camber.

Finally, results for airfoil Thick DA1+Cam DA0 in the retreating side are not presented because this airfoil cannot meet the trim constrained, as it happened with airfoil DA1 in Chapter 5 (see Tab. 5.9 and Fig. 5.14b). In a similar way, airfoil with DA0 thickness and DA2 camber exceeds the moment coefficient in the advancing side as presented in Tab. 6.4 (similarly to what happened for DA2, please refer to Tab. 5.9). To some extent, these airfoils obtained by combination of thickness and camber of DA0, DA1 and DA2 inherit some of the characteristics of the original airfoils. However, it is noted that the moment coefficient of the airfoil Thick DA1+Cam DA2 also exceeds the threshold value $\overline{C_{M_2}}$ set in the optimization loop (see Tab. 6.4).

6.2 Optimal morphing airfoils

The optimization problems tackled in the case of morphing airfoils maintain the same structure of the cases presented in Chapter. 5. Thus, two optimization problems are performed, one for the advancing side and one for the retreating side, which seeks the design with minimal drag coefficient with an equality constraint on the lift coefficient and an inequality constraint on the moment coefficient. The same values of target lift coefficients, bounding moment coefficients and operating conditions are used (please refer to Tab. 5.6).

The difference in this case is that the optimization is performed by acting on the design variables describing the morphing geometry. The morphing design variables are described in the following section.

6.2.1 Morphing strategy and parameterization

The morphing strategy considered in this work is a conformable camber airfoil, which changes over the period of rotation of the blade. Several technological solutions can be employed to achieve this goal: for instance, the FishBone Active Camber [45], the controllable camber presented in Ref. [41] and chiral structures as the one developed in Ref. [47]. The basic idea in the definition of a camber morphing airfoil is that the thickness distribution is held fixed, while the camber is allowed to change its shape during flight.

The type of morphing introduced to some extent in the previous section permits modification of the camber over the entire chordwise coordinate. However, in general the camber line cannot be modified at any chordwise coordinate, but it is held fixed in particular areas to maintain the internal structure of the blade and to accommodate morphing mechanisms. In general, the region close to the leading edge is fixed, where the D-shaped spar used in the helicopter blade structure is found. Thus, two

(a) Camber line morphing. **(b)** Trailing edge morphing.

Figure 6.4: Types of morphing strategy: baseline geometry (black), camber morph (green) or trailing edge morph (blue) and resulting morphed geometry (red) .

types of camber morphing are considered. In the first case (Fig. 6.4a), complete freedom is given to the camber line. In the second type (Fig. 6.4b) the first part of the airfoil is fixed, while for the remaining part of the blade section, the camber can morph and work as a trailing-edge flap with a larger extension and continuous shape modification. In this way, an increasing level of constraints to the morphing strategy are given: (i) ideal morphing, with constraints neither on the thickness nor on the camber line, (ii) camber morphing, with a fixed thickness distribution, and (iii) trailing edge morphing, with a fixed thickness distribution and a partial constraint on the camber line.

A suited parameterization is required to describe and optimize a morphing camber airfoil. The parameterization introduced in Sec. 5.2.1 treats separately the camber line and the thickness distribution perpendicular to the local camber line. For the entire camber morphing, the design variables are the coefficients A_i^c of the shape function $S_c(\chi)$ used to describe the camber distribution as $\zeta_c(\chi) = C_c(\chi)S_c(\chi)$. The class function is the same function used in Chapter 5 and presented in Eq. (5.9). The order of the Bernstein polynomials is retained from the optimization in forward flight condition, but owing to the fact that the thickness distribution is now fixed, the number of design variables is reduced by half. It follows that the design variables for the case of camber morphing are $\boldsymbol{x}_{m,c} = \{A_i^c\},\ i = 0, ..., 4$.

For the trailing edge morphing modification, a piecewise cubic function is used. The cubic functions are defined to ensure continuity up to the second order. Following an approach similar to the one presented in Ref. [40], the morphing camber line is described as follows

$$\begin{aligned} \zeta_c(\chi) &= C_c(\chi)\ S_c(\chi) + H(\chi) \\ H(\chi) &= \sum_{i=1}^{n_c} h_i(\chi) \end{aligned} \tag{6.1}$$

and each cubic function $h_i(\chi)$ reads

$$h_i(\chi) = \begin{cases} 0, & \text{if } \chi \leq \psi_{0,i} \\ A_i^m \frac{(x-x_{0,i})^3}{(1-x_{0,i})^3}, & \text{if } \chi > \chi_{0,i}, \end{cases}$$

where a_i are the coefficients of the cubic functions and $x_{0,i}$ are the locations of the start of the morph. The number of cubic functions n_c is equal to two in order to represent also reflex cambered airfoil and the locations of the morphs are $x_{0,1} = 0.4$ and $x_{0,2} = 0.7$. The D-spar length $x_{0,1}$ is set to 40% of the chord, following the

FishBAC configuration [40]. The coefficients of the cubic functions are the design variables of the morphing airfoil $\boldsymbol{x}_{m,te} = \{A_1^m, A_2^m\}$.

6.2.2 Deterministic results

With two different operating conditions and two different morphing strategies, four optimization problems are performed, which are summarised in Tab. 6.5. Each optimization has been performed with an initial generation of 100 individuals and 10 generations of 55 individuals to widely explore the solution space. The results of these optimization are reported in Tab. 6.6, where the first and second columns indicate the morphing strategy and the objective, respectively. The last two columns present the gain with respect to the optimal non-morphing airfoil, i.e. airfoil DA0.

The camber morphing involving the entire camber line permits to gain 23% with respect to DA0 in the retreating side and 10% in the advancing side. In the latter condition, the optimal solution obtained by the optimization is indeed the airfoil obtained by summing the thickness distribution of DA0 to the camber line of airfoil DA1, i.e. Thick DA0+Cam DA1. In the retreating side, the optimization is capable of improving the drag coefficient of airfoil Thick DA0+Cam DA2 and at the same time overcoming the problem associated with the moment coefficient constraint.

On the other hand, the trailing edge morphing improves the starting airfoil of 5% in the retreating side, but it does not find better solutions than airfoil DA0 in the advancing side. A first explanation of this can be found in the fact that airfoil DA0 has a shape more akin to DA1 than DA2 (see Fig. 6.1) and it is an airfoil that already greatly reduces the wave drag coefficient of the reference airfoil (i.e. the NACA 23012 airfoil shown in Tab. 5.5). However, the modification that the camber needs to undergo to improve the performance in the advancing side (for example, a modification into the camber line of airfoil DA1) is not compatible with the trailing edge modification shown in Fig. 6.4b. This problem could be overcome only by changing the thick-

Case	Objective	Equality constraints	Inequality constraints	Design variables
DMO1	$\min C_{D,1}$	$C_{L_1} = \overline{C_{L_1}}$	$\lvert C_{M_1}\rvert \leq \overline{C_{M_2}}$	$\boldsymbol{x}_{m,c}, \alpha$
DMO2	$\min C_{D,2}$	$C_{L_2} = \overline{C_{L_2}}$	$\lvert C_{M_2}\rvert \leq \overline{C_{M_2}}$	$\boldsymbol{x}_{m,c}, \alpha$
DTO1	$\min C_{D,1}$	$C_{L_1} = \overline{C_{L_1}}$	$\lvert C_{M_1}\rvert \leq \overline{C_{M_2}}$	$\boldsymbol{x}_{m,te}, \alpha$
DTO2	$\min C_{D,2}$	$C_{L_2} = \overline{C_{L_2}}$	$\lvert C_{M_2}\rvert \leq \overline{C_{M_2}}$	$\boldsymbol{x}_{m,te}, \alpha$

Table 6.5: Summary of optimization problems.

Morphing	Objective	C_{D_1} [counts]	$\alpha_1\|_{\overline{C_{L_1}}}$ [deg]	C_{D_2} [counts]	$\alpha_2\|_{\overline{C_{L_2}}}$ [deg]	Gain side 1	Gain side 2
Camber	$\min f_1$	0.805	-0.945	192.3	11.502	+10%	
Camber	$\min f_2$	102.82	-2.5	135.7	9.62		+23%
TE morph	$\min f_1$	0.8947	-1.059	177.8	11.1	0%	
TE morph	$\min f_2$	100.50	-1.93	169.2	10.045		+5%

Table 6.6: Performance of the deterministic airfoils obtained with two strategies for camber morphing

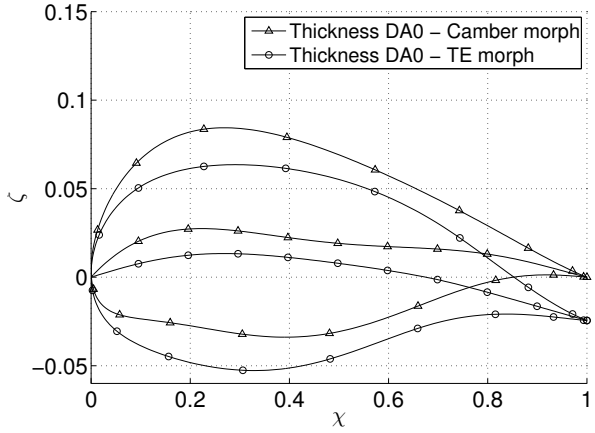

Figure 6.5: Airfoils obtained by morphing the camber mean line of airfoil DA0.

ness distribution together with the trailing edge modification. The best gain of such a strategy may be defined only by an optimization problem combining the coefficient of the thickness distribution A_i^t, the coefficient of the mean line A_i^c and the coefficient of the local trailing edge morph A_i^m.

To appreciate the difference in the morphing strategies, the optimal airfoils for the retreating side are presented in Fig. 6.5. It is worth noting how the camber morphing lifts up the camber line while maintaining the trailing edge location, and the trailing edge morph deflects the trailing edge to achieve an increase in lift coefficient.

6.2.3 Robust results

From the deterministic optimization, valuable information has been obtained. In fact, the preceding sections have shown how the camber morphing could potentially bridge the gap between a non-morphing airfoil and an "ideal" morphing airfoil and that the

Case	Objective	Inequality constraints	Design variables
RMO1	$\max(\mu(f_1))$ and $\min(\sigma^2(f_1))$	$\lvert C_M \rvert \leq \overline{C_M}$	$\boldsymbol{x}_{m,c}$
RMO2	$\max(\mu(f_2))$ and $\min(\sigma^2(f_2))$	$\lvert C_M \rvert \leq \overline{C_M}$	$\boldsymbol{x}_{m,c}$

Table 6.7: Summary of robust optimization problems.

morphing over the entire chordwise distribution of the camber line is more effective to this purpose. Starting from these results, the robust optimization problems are tackled. The same formulation used in Chapter. 5 is maintained (Eq. (5.11)), but the design variables used in this case are the design variable of the camber morphing $\boldsymbol{x}_{m,c}$, as this technique is more effective. Two optimization problems are then performed, which are summarised in Tab. 6.7.

Given the analysis of the previous chapter, the nominal operating conditions are set from the deterministic results. In other words, the nominal operating condition of the robust optimization in the advancing side is defined by the angle of attack $\overline{\alpha_1}$ of airfoil DMA1, while for the retreating side the nominal angle of attack is the angle $\overline{\alpha_2}$ of airfoil DMA2.

In the robust optimization problem (RMO2) the Pareto front shown in Fig. 6.6a is obtained. To select an interesting airfoil from the Pareto set the same procedure used in Sec. 5.4.2 is followed. Thus, the lift coefficient obtained in the samples over

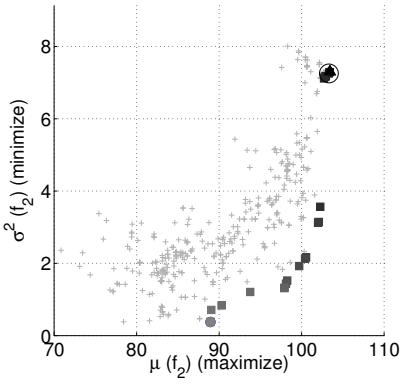

(a) Pareto front.

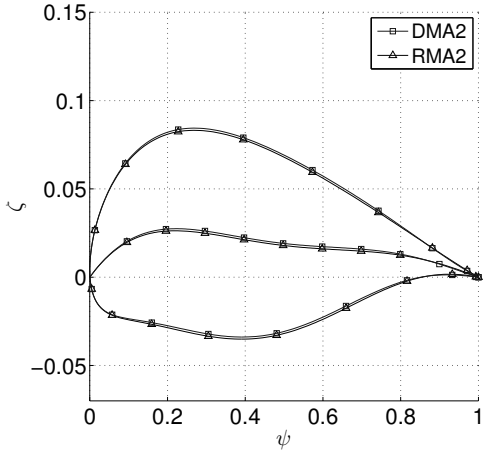

(b) Optimal deterministic and robust airfoils.

Figure 6.6: Results of the robust optimization for the retreating side (RMO2) .

Airfoil	$C_{D,2}\|_{\overline{C_{L_2}}}$ [counts]	$\alpha_1\|_{\overline{C_{L_2}}}$ [deg]	$\mu_{C_{D,2}}$ [counts]	$\sigma^2_{C_{D,2}}$ [counts^2]	$(\sigma/\mu)_{C_{D,2}}$ [-]
DMA2	135.6	9.62	135.9	39.99	0.047
RMA2	136.6	9.62	136.7	40.06	0.046

Table 6.8: Drag coefficient of the deterministic airfoil minimizing C_{D_2} (DMA2) and the airfoil selected from the robust front in the retreating side (RMA2).

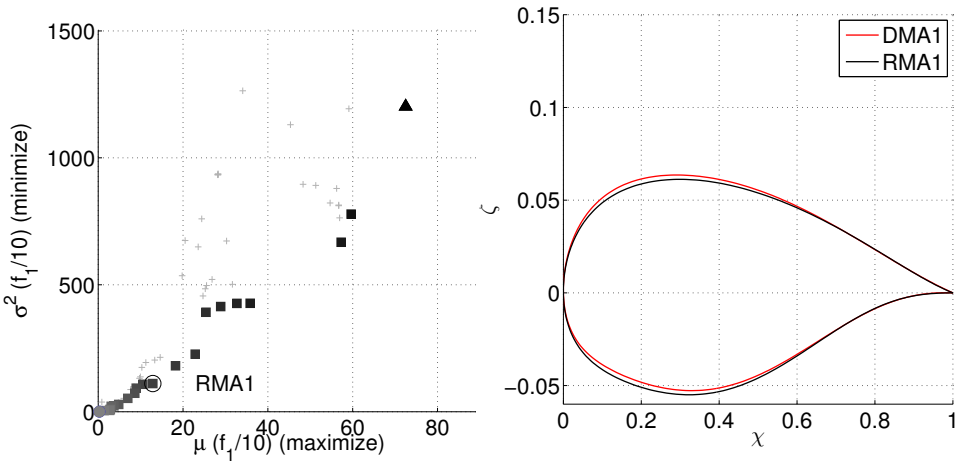

(a) Pareto front.

(b) Optimal deterministic and robust airfoils.

Figure 6.7: Results of the robust optimization for the advancing side (RMO1) .

the stochastic domain are plotted against the target lift coefficient and an airfoil is selected from the front. The selected airfoil is plotted in Fig. 6.6b together with the deterministic airfoil DMA2. This airfoil is referred to as RMA2 and its performance in the uncertainty range around the operating condition satisfying the lift constraint are reported in Tab. 6.8. It is possible to note that airfoils DMA2 and RMA2 are very similar and therefore they present similar performance. In this case, the result from the deterministic optimization is a robust airfoil in itself. However, it is important to note that this in general may not be the case and that robust optimization formulation represents a procedure to guarantee that the optimal airfoils would not incur in lift loss, drag increase or moment penalties throughout the uncertainty range.

The results of the robust optimization problem in the advancing side (RMO1) are shown in Fig. 6.7a. In the Pareto front, airfoil RMA1 is hihglighted which is chosen for its significantly low variance and average mean performance. The shape

Airfoil	$C_{D,1}\|_{\overline{C_{L_1}}}$ [counts]	$\alpha_1\|_{\overline{C_{L_1}}}$ [deg]	$\mu_{C_{D,1}}$ [counts]	$\sigma^2_{C_{D,1}}$ [counts^2]	$(\sigma/\mu)_{C_{D,1}}$ [-]
DMA1	0.805	-0.945	3.99	31.09	1.4
RMA1	0.827	-0.894	5.07	43.57	1.3

Table 6.9: Drag coefficient of the deterministic airfoil minimizing C_{D_1} (DMA1) and the airfoil selected from the robust front in the advancing side (RMA1) with camber morphing.

of airfoil RMA1 is presented in Fig. 6.7b along with the deterministic optimal airfoil. It is possible to note that even in this case the deterministic solution is very similar to a non-dominated robust airfoil. The deterministic airfoil is also very robust, as demonstrated by the values of the mean and variance presented in Tab. 6.9.

6.3 Validation with azimuth

After evaluating the performance in the advancing and retreating side, the drag coefficient of the optimal airfoils over the azimuth angle of the blade rotation is computed. In the case of a morphing airfoil, a smooth transition from the optimal shape in the advancing side to the optimal shape in the retreating side is required. The transition over the azimuth angle ψ is obtained as follows

$$\zeta = \frac{\zeta_1}{2}(1 + \sin(\psi)) + \frac{\zeta_2}{2}(1 - \sin(\psi)) \tag{6.2}$$

where ζ_1 is the shape in the advancing side ($\zeta(90^\circ) = \zeta_1$) and ζ_2 is the shape in the retreating side ($\zeta(180^\circ) = \zeta_2$). The transition in the case of the "ideal" morphing

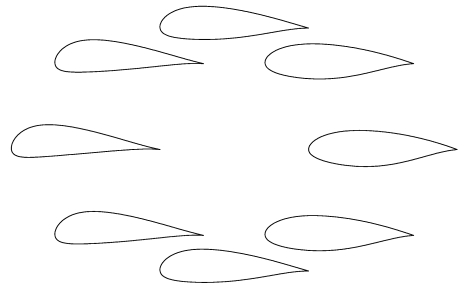

Figure 6.8: Shapes over the azimuth obtained by morphing DA1 into DA2.

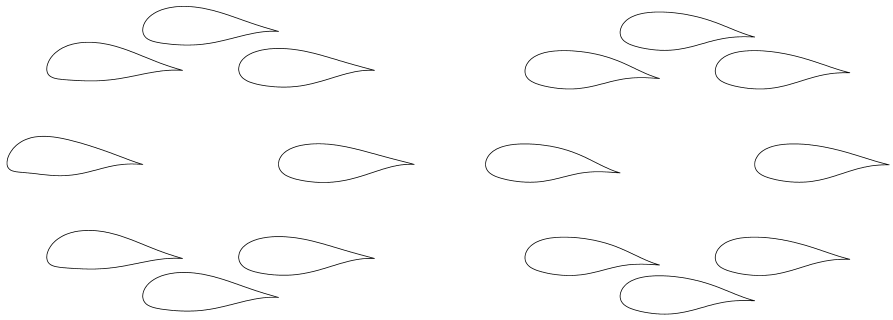

(a) Morphing DMA1 into DMA2. **(b)** Morphing DTA1 into DTA2.

Figure 6.9: Shapes over the azimuth.

DA1-DA2 is presented in Fig. 6.8, while the morphing shapes with geometrical constraints are shown in Fig. 6.9. The drag coefficient is then evaluated by following the procedure explained in Sec. 5.5.

The values of the mean angle of attack and the first harmonic amplitude are computed and reported in Tab. 6.10 for the following combinations: (i) DA1-DA2, i.e. airfoil DA1 morphing into DA2, (ii) DMA1-DMA2, i.e. airfoil DMA1 morphing into DMA2, and (iii) DTA1-DTA2, i.e. airfoil DTA1 morphing into DTA2. The lift and drag coefficients for these airfoils are presented in Fig. 6.10. It is possible to note that the greatest difference among these airfoils is the performance in the retreting side, where the lowest drag coefficient is obtained for the "ideal" morphing and the drag coefficient increases with increasing geometrical constraints. The performance in the advancing side is comparable among the different morphing strategies, because the viscous correction [148] is dominant with respect to the inviscid drag coefficient.

The average drag coefficient obtained with the "ideal" morphing and the camber morphing are presented in Tab. 6.11. The table also includes the gain with respect to the non-morphing optimal airfoil DA0 which is higher for the ideal case (22.6 %) and reduces to 17.1% and 3 % in the case of the camber morphing and trailing edge modification, respectively. It is clear that to obtain the maximum gain of the morph-

	DA1-DA2	DMA1-DMA2	DTA1-DTA2
α_0 [deg]	4.13	4.34	4.49
Δ_1 [deg]	4.95	5.10	5.36

Table 6.10: Variables of time-varying angle of attack (Eq. (5.12)) for optimal shapes.

	DA1-DA2	DMA1-DMA2	DTA1-DTA2
$avg(C_D)$ [counts]	76.44	81.84	95.14
Gain wrt DA0	22.6 %	17.1 %	3.6 %

Table 6.11: Average drag coefficient over the blade rotation period for morphing airfoils and associated gain.

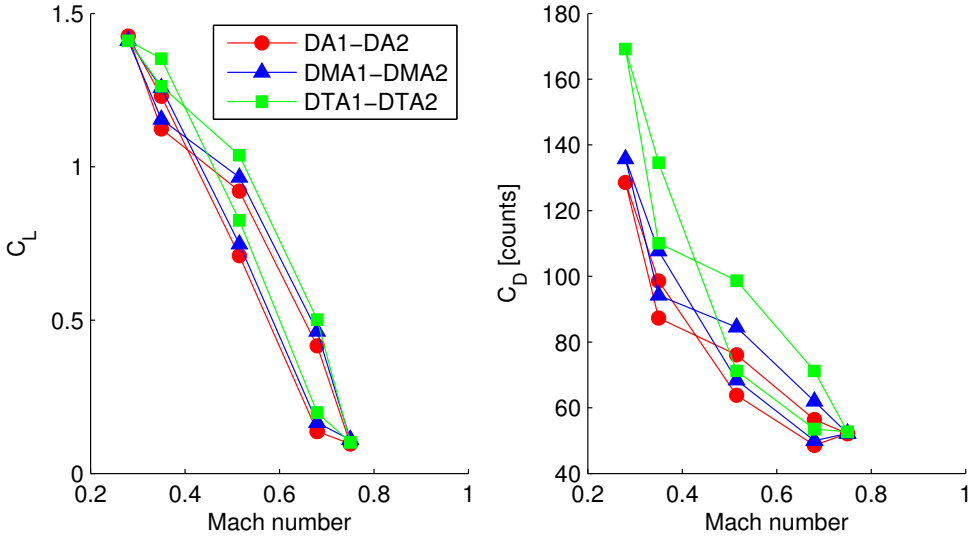

Figure 6.10: Lift and drag coefficient loop of DA1-DA2, DMA1-DMA2 and DTA1-DTA2.

ing strategies, an optimization which can model the thickness distribution as well is necessary.

CHAPTER 7

Conclusion

A methodology for robust optimization has been developed to tackle the problem of designing fixed and morphing airfoils of helicopter rotor blades. Robust optimization requires the coupling of the uncertainty quantification method and the optimization algorithm. This coupling increases the computational demand, especially in aerodynamic applications, where a single function evaluation may be very expensive. Two different methods have been developed to improve the numerical efficiency of the robust optimization coupled loop.

The first solution is the multi-fidelity strategy based on two different models of the aerodynamics which have different accuracy and computational cost. In the aerodynamic application, a rapid estimation is obtained by means of a panel/integral boundary layer method for the stochastic sample associated with subsonic flow conditions, while a coupled Euler equations/integral boundary layer solver is employed for those operating conditions in the stochastic variables domain that exhibit important compressible effects. A switching logic is used for each design to determine which samples should be computed by the rapid low-fidelity or by the high-fidelity codes. The adaptive strategy, on the other hand, reduced the number of function evaluations by computing the statistics with higher accuracy only for the solutions which are interesting for the optimization loop. This method uses a definition of the Pareto front based on the error bounding boxes, i.e. the error affecting the estimate of the objective function.

Both strategies have been applied to algebraic test cases and to the hovering op-

timization problem. In the latter application, the multi-fidelity strategy reduced the computational cost by 60% with respect to the high-fidelity strategy, and results in Pareto airfoils consistent to the high-fidelity results. Thus, the proposed method is well-suited to indicate regions of improvement in the design space, which can be later analysed in detail by means of a high-fidelity model. In addition, the multi-fidelity strategy could be easily applied to other models at one's disposal, and it could be extended to other flowfield characteristics and physical models. On the other hand, the adaptive strategy method has reduced the computational cost of the aerodynamic test case by 40%. However, the accuracy of the front in this case is even higher owing to the fact that the non-dominated solutions are computed with the tight tolerance.

In the hovering optimization problem the robust optimization method proved effective in finding airfoils that improve the aerodynamic lift-to-drag-ratio, by postponing the boundary layer transition to turbulence in the subsonic case, and by reducing the drag due to compressibility effects in the low transonic case, in the range of considered operating conditions.

In the forward flight case, deterministic results have proven how different the optimal airfoils for advancing blade is compared to the retreating case. Comparison with the robust approaches have shown similarities and differences between the use of the drag coefficient as a deterministic objective and the employment of the lift-to-drag ratio in a robust context.

Morphing airfoils have been considered as a possible solution to cope with the variable flow conditions encountered by blade element in forward flight. Deterministic estimates prove that an "ideal" morphing can increase the performance of an optimal non-morphing airfoils by 22.6 % on the average drag coefficient, and by 17 % when thickness constraint is considered. Results obtained with the trailing edge morphing are not satisfactory and it is required to include the thickness distribution in the design variables to assess the full potential of this geometry modification.

The employment of robust optimization is thus expected to be of great interest in the application of helicopter rotor airfoils, owing to the fact that even the airfoils with higher mean values achieve low values of the variance of the aerodynamic efficiency with respect to variations of the angle of attack. The reduction of the variance could lead for instance to a reduction of the required rotor shaft torque in variable operating conditions. Furthermore, in the different applications presented in the thesis, results have demonstrated how the robust optimization formulation represents a formal procedure to guarantee that the optimal airfoils would not incur in lift loss, drag increase or moment penalties throughout the uncertainty range.

Future works will be devoted to the unsteady assessment of the morphing ge-

ometries to ultimately verify the result of the steady optimization. In addition, an optimization including the thickness distribution as a design variable will be considered for the morphing cases. With regard to the robust technique developed in the thesis, the joint employment of the multi-fidelity method and coupled adaptive approach will be explored to further improve numerical efficiency. Furthermore, another important aspect when using robust optimization in real-world application is the possibility to handle arbitrary probability distributions. Thus, a feature perspective is to include moment-based PC or generalization of the stochastic collocation. Finally, the novel strategy reported in Appendix B for dealing with optimization under uncertainty is worth being further investigated for test problems of engineering interest, in particular in the framework of rotorcraft airfoil optimization.

Nomenclature

α	Angle of attack [deg]
β	Coefficient for multi-objective optimization problems
ϵ	Camber line slope angle
ε	Error of the objective function estimate
$\varepsilon_\mu, \varepsilon_{\sigma^2}$	MC errors on the mean value and variance obtained with SSC
$\hat{\varepsilon}_\mu, \hat{\varepsilon}_{\sigma^2}$	Estimated errors on the mean value and variance obtained with SSC
$\tilde{\varepsilon}_\mu, \tilde{\varepsilon}_{\sigma^2}$	Conservative estimated errors on the mean value and variance obtained with SSC
ζ	Non-dimensional coordinate of the airfoil
θ	Blade pitch angle [deg]
μ	Mean value
μ^r, μ_c^r	Statistical moment and centered moment
ν_k	Coefficient of the PC expansion
ξ	Uncertain variable
Ξ	Stochastic variables domain
Ξ_c	Set of control points in the stochastic variables domain
Ξ_s	Set of samples in the stochastic variables domain
ϕ	Solution of the system governing equations
φ	Univariate polynomial used in PC expansion
τ	Tolerance of the adaptive uncertainty quantification method
χ	Non-dimensional chordwise coordinate of the airfoil
ψ	Azimuth angle
Ψ	Multivariate polynomial used in PC expansion

ω	Angular frequency [rad/s]
Ω	Rotor angular frequency [rad/s]
$\mathbf{a}, \mathbf{b}$	Design vectors
A^m	Coefficient of shape parameterization for trailing edge morph
A^c, A^t	Coefficient of shape parameterization for camber and thickness
c	Airfoil chord
$\mathcal{B}$	Bounding box set
$\mathcal{C}$	Control function
$\overline{\mathcal{C}}$	Threshold value of control function
C_L, C_D, C_M	Lift, drag and moment coefficient
$\overline{C_L}$	Target lift coefficient
$\overline{C_M}$	Bounding moment coefficient
$\mathcal{D}$	Spatial domain
f	Objective function
$\mathbf{f}$	Objective functions vector
$\mathcal{F}$	σ-field of probability space
h	Cubic function for trailing edge morphing
H	Trailing edge morphing function
L	Lift
$\mathcal{L}$	Operator of the system model equations
k	Reduced frequency, $k = \omega\, c/2V_\infty$
m	Number of objective functions
M	Mach number
n	Number of design variables
n_ξ	Number of uncertain variables
n_e	Number of elements in SSC
n_{mc}	Number of Monte Carlo samples
n_s	Number of samples in the stochastic space
n_o	Order of the PC univariate polynomial
n_{QP}	Number of quadrature points
N	Number of terms in PC truncated expansion
N	Order of the CST Bernstein polynomials
N_1, N_2	Exponent of CST class function
$\mathcal{P}$	Pareto set
q	Number of constraints
r	Percentage of blade radius
R	Blade radius

$\mathcal{T}$	Temporal domain
t	Time variable
u	Performance or quantity of interest
V_∞	Freestream speed
w	Approximating function in SSC
x_0	Initial location of cubic function for trailing edge morph
$\boldsymbol{x}$	Design variables
$\boldsymbol{x}_{m,c}$	Design variables for camber morph
$\boldsymbol{x}_{m,te}$	Design variables for trailing edge morph
$\mathcal{X}$	Design variables domain
$\mathbf{y}$	Spatial variables
$\mathcal{Y}_\mathcal{P}$	Image of the Pareto set

List of abbreviations

BIMADS	Bi-objective Mesh Adaptive Search Algorithm
CFD	Computational Fluid Dynamics
DO	Deterministic Optimization
DMO	Deterministic Morphing Optimization
HF	High Fidelity
LF	Low Fidelity
NISP	Non-Intrusive Spectral Projection
NSGA	Non-dominated Sorting Genetic Algorithm
PC	Polynomial Chaos
RO	Robust Optimization
RMO	Robust Morphing Optimization
SSC	Simplex Stochastic Collocation
UQ	Uncertainty Quantification

APPENDIX A

Deterministic aerodynamic benchmark case

In this appendix, the Non-Dominated Sorting Genetic Algorithm used throughout the thesis is applied to the minimization of the drag coefficient of an isolated airfoil in transonic flow. This optimization problem is the first of a set of test cases proposed by the AIAA Aerodynamic Design Optimization Discussion Group to establish a series of increasingly complex set of benchmark problems suitable for exercising aerodynamic optimization methods in a constrained design space [129]. The first test case was taken by Ref. [151] and it is a constrained minimization problem.

Benchmark problem

The benchmark optimization problem is a minimization of the drag coefficient C_d of a modified NACA 0012 at a freestream Mach number of 0.85 at zero angle of attack with a constraint on the minimum thickness of the airfoil. In mathematical terms, it reads

$$\begin{aligned} &\text{minimize :} \quad C_d, \\ &\text{subject to :} \quad y \geq y_{baseline} \quad \forall x \in [0,1] \end{aligned} \tag{A.1}$$

with respect to the design variables $\mathbf{x}$. The baseline airfoil $y_{baseline}$ is the NACA 0012-*mod* airfoil, which is obtained by imposing zero-thickness at the trailing edge

and it is defined as follows

$$y_{baseline} = \pm 0.6 \left(0.2969\sqrt{x} - 0.1260x - 0.3516x^2 + 0.2843x^3 - 0.1036x^4\right).$$

The other constraint of the optimization problem is the equation of state of the system, that is the governing equation of the aerodynamics. The benchmark problem requires the aerodynamic model to be the Euler equations with a constant ratio of specific heat of 1.4.

Shape parameterization

In order to define a finite set of design variables $\mathbf{x} \in \mathbb{R}^n$, a parameterization is required. The airfoil is described by means of a parameterization called Class/Shape function Transformation (CST) [126], described in Sec. 4.2.3. The CST is first ap-

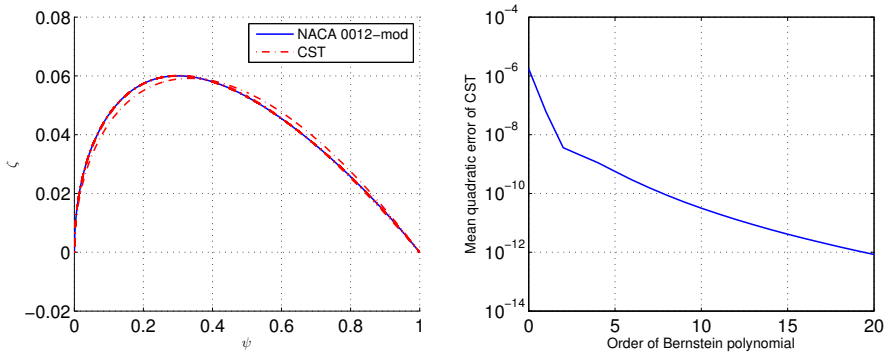

(a) CST representation with an increasing order of Bernstein polynomial.

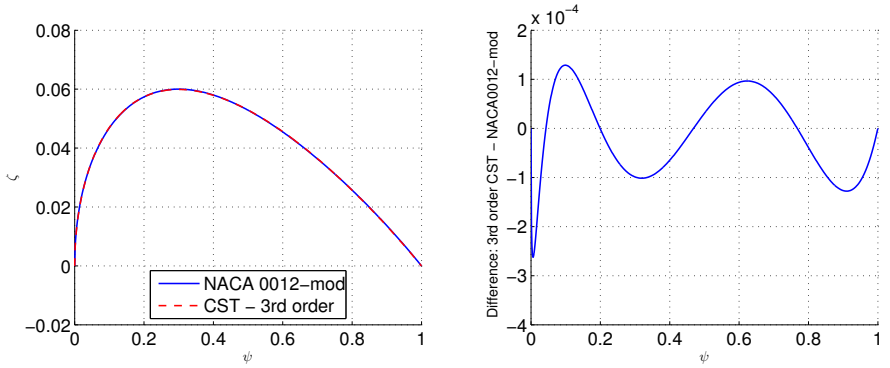

(b) Third-order CST results.

Figure A.1: Results of the shape parameterization for the NACA 0012-*mod* airfoil.

plied to represent the NACA0012-*mod* airfoil; the coefficients of the CST are obtained by means of a least squares approach for several orders of the polynomial. Figure A.1a shows that the CST is capable of representing the reference airfoil and that the error decreases with increasing polynomial order. In particular, with a third-order polynomial, the mean quadratic error of the CST attains a value lower than 10^{-8} and a very good representation of the reference airfoil is obtained (see Fig. A.1b). Note that a third-order CST corresponds to four design variables.

Aerodynamic model and mesh convergence

The evaluation of the objective function is performed by means of a Computational Fluid Dynamics (CFD) solver which provides the aerodynamic loads of the airfoil given its shape and operating conditions. The numerical solver of the Euler equations implemented in the Stanford University Unstructured (SU^2) software suite [143, 146] is employed. The solver provides several space discretization schemes, among which the second-order JST scheme is employed. An implicit Euler, local time-stepping is used to converge to the steady-state solution, and the GMRES method in conjunction with the LU SGS preconditioner is used to solve the resulting system.

With regard to the computational domain, half of a circular domain around the airfoil is considered. In fact, the solution of the optimization problem is expected to be symmetrical, because the angle of attack is set to zero and the optimization objective is the minimization of the drag of the airfoil. Thus half of a O-mesh extending

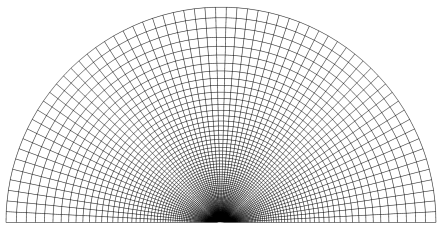

(a) 64×128 mesh.

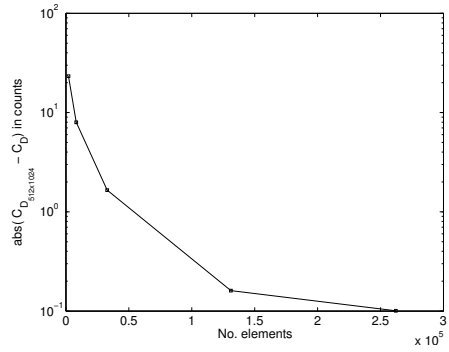

(b) Difference between the drag coefficient C_D and the drag coefficient $C_{D_{512\times1024}}$ of the finest mesh (512×1024 cells).

Figure A.2: Mesh convergence for the NACA 0012-*mod* airfoil.

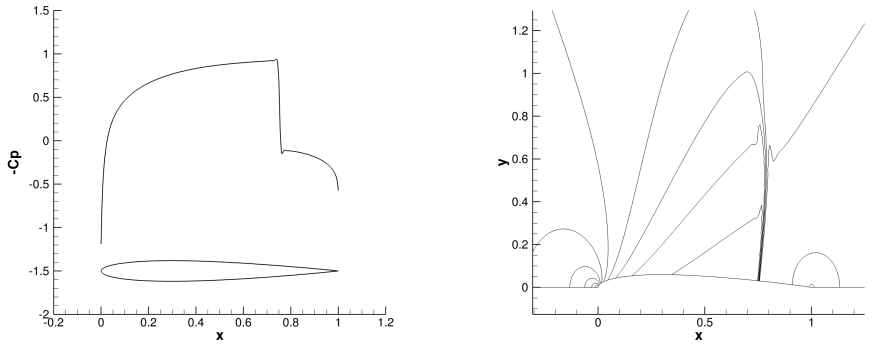

(a) Pressure coefficient on the airfoil. **(b)** Pressure coefficient contour.

Figure A.3: Pressure coefficient of the NACA 0012 *mod*(please note that the airfoil in the left figure has been scaled for visualization purpose).

$\pm$ 30 chords away of the airfoil is considered. Farfield boundary conditions are applied to the external half-circle, slip condition is applied to the airfoil boundary and symmetry conditions are applied to the lines connecting the leading edge and trailing edge of the airfoil to the farfield boundary. The benchmark problem required a mesh convergence study that is here considered only for the drag coefficient because of the employment of the symmetry property of the problem. For the baseline airfoil, several mesh grids are created by halving the number of cells in each direction: an example is shown in Fig. A.2a. In addition to this first set of mesh, a fine mesh based on 512$\times$512 cells is considered as a tradeoff between the 256$\times$512- and the 512$\times$1024-mesh. The values of the drag coefficient are presented in Tab. A.1, and the difference with respect to the drag coefficient of the finest grid $C_{D_{512\times1024}}$ is plotted against the size of the mesh in Fig. A.2b. The convergence trend is clear and based on the result of this analysis, the grid with 256$\times$512 cells has been chosen because it provides an error within 0.1 drag counts with respect to the subsequent refinement, and a difference of 0.16 drag counts with respect to the finest grid. Furthermore, it is worth noting that the computational model provides an estimate that compares well with the results in the benchmark reference [151]. Also the pressure distribution over the airfoil and in the domain (Fig. A.3) is in agreement with the result of the literature: due to the high Mach number the NACA 0012 airfoil presents a strong shock at about 75% of the chord and it suffers from a significant wave drag.

Once the mesh size is set up, the CFD solver is integrated in the optimization loop. Each computation of the flowfield of a new design is obtained starting from

	32×64	64×128	128×256	256×512	512×512	512×1024
C_D	0.0434738694	0.0465975412	0.0456345116	0.0457837056	0.045789711	0.04579979

Table A.1: Estimates of drag coefficient of the NACA 0012 *mod.*

the baseline computation and applying a suited displacement to the airfoil boundary of the mesh. The displacement is then propagated from the boundary to the closer mesh volumes by means of a mesh deformation method. The solver provides a mesh deformation tool based on the finite element method discretization of the linear elasticity equations; in this case the element stiffness is inversely proportional to the cell volume and 500 smoothing iterations are applied to get a regular deformed mesh.

Optimization algorithm

In the present work, the optimization strategy is based on the Non-dominated Sorting Genetic Algorithm (NSGA) [67]. Typical values for the crossover and mutation probabilities pc, pm are used, i.e. 0.9 and 0.1 respectively, and the geometric constraint is imposed by penalization of the fitness function. An initial Design of Experiment is performed with 100 individuals and the following generations are based on 60 design vectors. The algorithm was implemented in parallel mode, by using MPI libraries.

Results

Initially, four design variables have been considered to perform the optimization, based on the result discussed in Section 3.2.1. After the initial DOE and 10 generations, a reduction of 158 counts has been achieved (see Table A.2) that seemed in agreement with the results presented in the benchmark reference work [151] obtained with a a Bezier curve parametrization with 3 design variables. The contour of the pressure coefficient and the pressure distribution of this airfoil are presented in Fig. A.5. Nevertheless, the optimal result appeared quite far from the final goal reported by Ref. [151], and thus it was deemed that the number of design variables was not sufficiently high to describe an airfoil with higher performance.

Therefore, it has been decided to increase the number of design variables to release the level of freedom of the optimization problem. Thus, a 9-th order Bernstein polynomial has been chosen, which corresponds to ten design variables. In this case, the range of the coefficients has been reduced to avoid the regions in the design space that would yield unfeasible geometries (e.g. geometries with too many inflection points). Twenty generations after the original DOE have been computed which

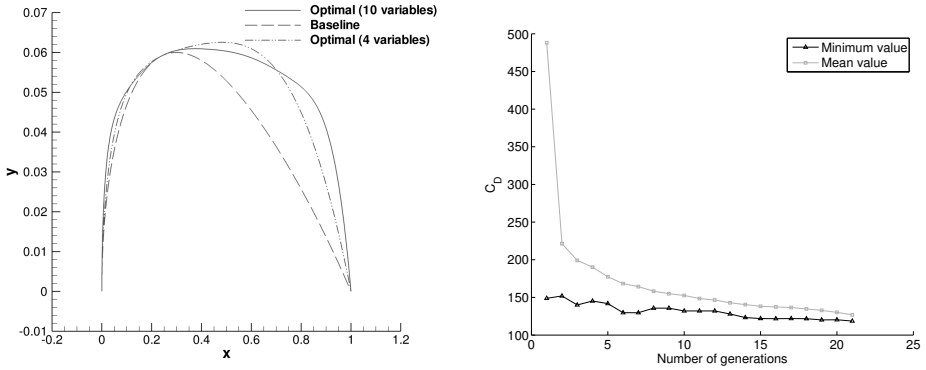

(a) Optimal airfoil. **(b)** Convergence history.

Figure A.4: Results of the aerodynamic optimization.

correspond to 1300 individuals; the optimal airfoil is presented in Fig. A.4a, with a comparison with the baseline airfoil. The figure shows that the geometric constraints is satisfied by the optimal solution. The evolution of the minimum and mean values of the objective function is shown in Fig. A.4b. It is noticeable that the convergence of the mean value is monotonic, and that a great improvement with respect to the reference airfoil is evident from the very first generation. As a matter of fact, genetic algorithms are very sensitive to the result of the initial DOE which has to be carefully prepared to get a rough yet comprehensive representation of the performance in the design space. The evolution is then capable of reducing the initial minimum by 20%, whereas globally the drag coefficient is reduced of about 340 counts with respect to the baseline airfoil. A similar reduction was reached in Ref. [151] using 24 design variables. The resulting optimal design presents a blunt trailing edge with a flat region that extends for the great majority of the chord length (Fig. A.6). Because of the shape of the airfoil the strong shock wave that would occur for the NACA 0012 airfoil is almost split into two contributions: one close to the leading edge and the second one close to the trailing edge. The first shock wave appears less steep with respect to the results of the literature owing to the lower local curvature of the airfoil that reaches the maximum thickness more gradually. The second shock instead appears stronger and results in an increased drag coefficient. Finally, a grid convergence analysis has been performed for the optimal design obtained with 10 design variables to assess the accuracy of the estimate. The result is reported in Table A.3 which proves the convergence of the estimates.

To sum up, the optimization method based on a genetic algorithm and the CST

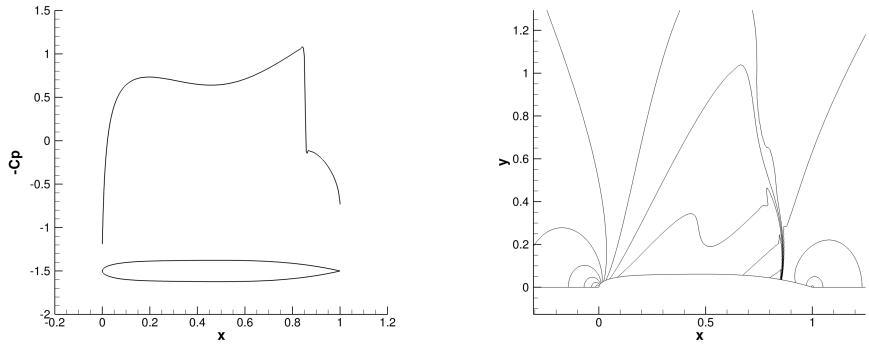

(a) Pressure coefficient on the airfoil. **(b)** Pressure coefficient contour.

Figure A.5: Pressure coefficient of the resulting optimal airfoil when four design variables are considered (the airfoil in the left figure has been scaled for visualization purpose).

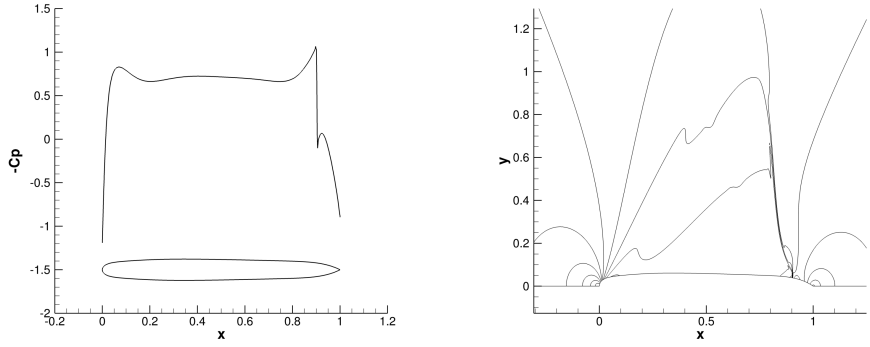

(a) Pressure coefficient on the airfoil. **(b)** Pressure coefficient contour.

Figure A.6: Pressure coefficient of the resulting optimal airfoil (the airfoil in the left figure has been scaled for visualization purpose).

NACA 0012	Optimal airfoil (4 design variables)	Optimal airfoil (10 design variables)
0.045784	0.029953	0.011862

Table A.2: Comparison of drag coefficient.

shape parameterization has proved capable of computing a new design with a performance improved by 70% with respect to the baseline airfoil, and in agreement with

	64×128	128×256	256×512	512×512	512×1024
C_D	0.011763	0.011901	0.011862	0.011829	0.011849

Table A.3: Estimates of drag coefficient for the optimal airfoil with increasing mesh size.

the results of the literature. The CST shape parametrization proved very effective, since it has been possible to achieve these results with only 10 design variables. In addition, the method has dealt with geometric constraints with success.

APPENDIX B

An alternative formulation for design under uncertainty

As presented in Sec. 2.1, a problem of uncertainty-based optimization can be formulated in several ways depending on the application and final goal. "Classical" formulations typically relies on the computation of the mean value, the variance, or a function of the two. These quantities however may not be capable of decribing in a comprehensive way the Probability Density Function (PDF) of the output, which represents the behavior of the system. In recent years efforts have been made to find formulations which exert a control on the PDF of the performance inside the optimization loop [152, 153]. Nevertheless, the computation of the PDF inside the optimization loop may increase the computational cost, and defining an attainable target PDF may often be difficult, because the PDF is a feature of the problem and it is not known a priori in most cases.

Another example is the work presented in Ref. [154], which considers the influence of high-order statistics during the optimization process. The aim here is to provide some useful indications for obtaining a good trade-off between the high-quality information given by high-order statistics and the feasibility of the whole optimization loop. In particular, a multi-objective optimization method taking into account high-order statistic moments, such as the third and fourth-order statistic moments, i.e. skewness and kurtosis, respectively, is considered. From this work it appears that the skewness may not be very useful in the definition of the optimization objectives to control the PDF, because it is difficult to handle in the optimization loop and it does not limit the tails of the PDF.

Here, an alternative formulation for design under uncertainty is explored, which is based on the computation of the mean value and the minimum of the function. The aim of the method is to exert a stronger control on the system output variability in the optimization loop at a moderate cost. This would reduce post-processing analysis of the PDF of the resulting optimal designs, by converging rapidly to the interesting individuals. In other words, in the set of designs resulting from the optimization, the new approach should be capable of discarding poor-performance design. Also, no a priori assumption on the optimal PDF is made.

Optimization problem

In an uncertainty-based optimization, the objectives are formulated to drive the variability of the performance u with respect to changes in the uncertain variables $\boldsymbol{\xi}$ to a desirable behavior, by acting on the design variables $\boldsymbol{x}$ [155]. Choosing the objectives in an uncertainty-based optimization basically means choosing the statistics of the performance that one wishes to optimize. In this work, a different formulation is presented with the goal of inserting comprehensive information about the PDF of the performance in the optimization loop at a moderate cost, and without conditioning the result with constraints or targets. The novel strategy is also explored to assess if a particular choice of the objectives could influence the global cost of the optimization problem, for instance by reducing the number of the final optimal designs. In the new approach, the mean value of the performance is accompanied by the minimum of the performance. By maximizing both the mean value and the minimum, the optimization addresses the problem of finding a design that guarantees a desirable minimum performance while maintaining high mean performances. In addition, this approach could be suited to tackle reliability-based design problems, where the goal of the optimization is to reduce the occurrence of undesired system response (e.g. failures). Although the current work only explores problems related to enhancement of the performances, reliability-based design is deemed a possible application of the novel formulation and it will be explored in future work.

In a probabilistic framework (see Sec. 2.3), the uncertain variables are treated as continous random variables in the stochastic space Ξ and they are assigned a probability density function p_ξ. The objectives of the optimization are measures of the response of the system under consideration to changes in the random variables. When the formulation based on the minimum $\min_u$ is considered the optimization problem

in Eq. (2.4) reads

$$\begin{aligned} \text{maximize:} \quad & \mu_u\left(\boldsymbol{x}, \boldsymbol{\xi}\right), \min_u\left(\boldsymbol{x}, \boldsymbol{\xi}\right) \\ \text{subject to:} \quad & \mathbf{g}\left(u\left(\boldsymbol{x}, \boldsymbol{\xi}\right)\right) \geq \mathbf{0} \\ \text{by changing:} \quad & \boldsymbol{x} \in \mathcal{X}. \end{aligned} \tag{B.1}$$

The performance u is typically a function of the solution of the equations describing the system under consideration which can be defined as follows

$$\mathcal{L}\left(\boldsymbol{x}, \boldsymbol{\xi}; \phi(\boldsymbol{x}, \boldsymbol{\xi})\right), \tag{B.2}$$

where $\mathcal{L}$ is the mathematical operator of the equations and φ is the solution vector and $u = u(\phi)$. In the next section, the computation of the minimum in the uncertainty quantification method is discussed.

Computation of the minimum

The computation of the minimum is performed in three different ways. The first, straightforward approach is the computation of the minimum value among the samples obtained for each point of a given quadrature formula (in this case, a Gaussian quadrature). This solution is referred to as Quadrature Point (QP) minimum computation and it is expressed as follows

$$\min_{\boldsymbol{\xi}_i \in \mathcal{N}_{QP}} v\left(\boldsymbol{\xi}_i\right), \tag{B.3}$$

where v is the output of the system, i.e. the quantity of interest, given a particular design and $\boldsymbol{\xi}_i$ are the samples drawn in the n_ξ-dimensional stochastic space belonging to the set $\mathcal{N}_{QP}$ of n_{QP} quadrature points.

The other two methods compute the minimum by leveraging the analytical approximations provided by two uncertainty quantification methods: (i) the Polynomial Chaos expansion, and (ii) the Simplex Stochastic Collocation, described in Sec. 2.3.2. The former provides a meta-model of function v over the entire stochastic space, starting from the set of n_{QP} samples $\mathbf{v}$ [91]. On the other hand, SSC computes a polynomial approximation either on the whole domain (when the global extremum diminishing condition is satisfied [93]) or on each element Ξ_j in which the grid is divided into.

From the analytical approximation w (either the truncated PC expansion of Eq. (2.27) or the piecewise polynomial of the SSC in Eq. (2.37)), the gradient of the response for

each design is computed and it is employed in a descent method based on Newton-Raphson iterations. From the location $\boldsymbol{\xi}_k$ at the k-th iteration, the new minimum location $\boldsymbol{\xi}_{k+1}$ is obtained using the descent step size γ and the descent direction $\mathbf{d}_k$ computed from the gradient $\nabla w(\boldsymbol{\xi}_k)$ and the Hessian $H(\boldsymbol{\xi}_k)$ of the analytical function in the current point:

$$\begin{aligned} \boldsymbol{\xi}_{k+1} &= \boldsymbol{\xi}_k + \gamma \mathbf{d}_k \\ H(\boldsymbol{\xi}_k)\mathbf{d}_k &= \nabla w(\boldsymbol{\xi}_k). \end{aligned} \tag{B.4}$$

If the Hessian is singular, the descent direction is taken proportional to the gradient. In the case of the PC expansion, the loop is started with the sample with lower value of function v, i.e. the minimum of the QP approach. With the SSC method, if a global interpolation is employed, the strategy used with PC still holds. Instead, when a local interpolation is employed, the descent method described in Eq. (B.4) is applied to each element Ξ_j in the stochastic grid and then the lower value of the local minima found on each element is computed.

It is worth noting that the underlying idea on building a meta-model of the function v is that the cost of sampling Eq. (B.2) is very high from a computational point of view. For instance, in the case of aerodynamic design, a Computational Fluid Dynamics (CFD) code is used to implement the equations of state, and a single CFD run can easily be 1000 times slower than the operations of a single optimization iteration. It follows that the cost of drawing a new sample is far higher than the cost of solving the iterative method in Eq. (B.4) to compute the minimum, because this procedure does not involve further sampling of Eq. (B.2).

Results

The novel optimization strategy is based on the mean and minimum values. While the mean value is a typical quantity in uncertainty quantification, an assessment of the computation of the minimum is required as this is a key feature of the new approach. Thus, the first section presents the result concerning the computation of the minimum. In the second part, the impact of the minimum on the optimization loop is assessed. Thus, the results first include the assessment of the techniques presented in the preceding section to compute the minimum. Secondly, the impact of the minimum on the optimization loop is discussed.

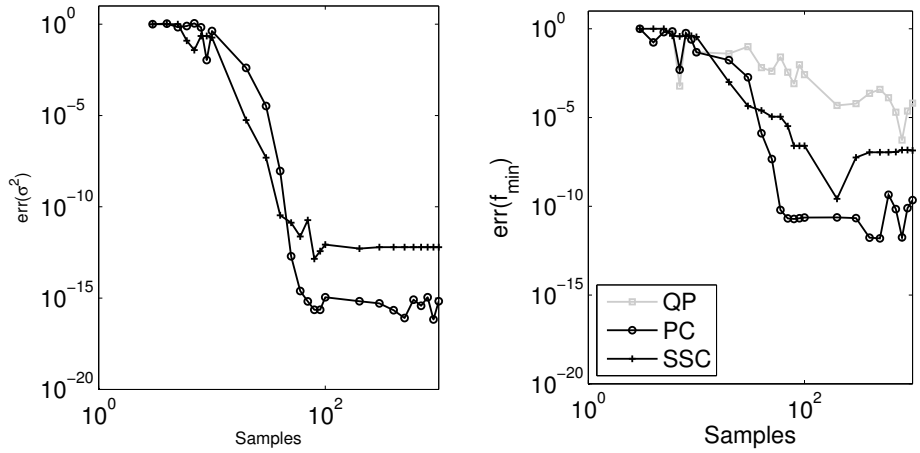

(a) Variance. **(b)** Minimum.

Figure B.1: Errors of the variance and the minimum for two-dimensional Michalewicz function computed with PC (grey circles) and QP (black squares).

Computation of the minimum

The computation of the minimum is applied to the Michalewicz function, a very challenging test case for optimization algorithms, which presents one global minimum, many local minima and wide, flat regions [116]. The function $u(\xi_1)$ is considered here in the one-dimensional stochastic space, and it reads

$$u(\xi_1) = -\sin(\xi_1)\ \sin^{20}\left(\frac{\xi_1^2}{\pi}\right) \tag{B.5}$$

where the uncertain variable has uniform distribution in the interval $[0, \pi]$. This test case is tackled by means of QP and PC approaches and the results of the errors for the variance and the minimum are presented in Fig. B.1a and Fig. B.1b, respectively. The QP estimate of the minimum presents some oscillations and the error remains higher with respect to the estimates obtained with the PC and SSC methods. In these cases, the error affecting the variance is slightly faster and it reaches a value close to the machine epsilon, whereas the error on the minimum remains a few order of magnitude larger. As a matter of fact, since the minimum is a point value, it is generally more difficult to be well captured with respect to an integral value. However, an error of magnitude -10 is acceptable for use in the optimization loop. Thus, the results prove that computing the minimum has only a slightly higher cost with respect to computing the variance and that it is better to resort on functional approximation

instead of using quadrature points.

Optimization with minimum

To assess the impact of the computation of the minimum a modification of the Michalewicz function [116] is considered. The function depends on two uniformly distributed uncertain variables in the interval $[0, \pi]$, and on two design variables: $\boldsymbol{x} \in \mathcal{X} = [0.1, 1]^2$. The modified function reads

$$\begin{aligned} u(\boldsymbol{x}, \boldsymbol{\xi}) &= M(\boldsymbol{x}, \boldsymbol{\xi}) + G(\boldsymbol{x}, \boldsymbol{\xi}) + P(\boldsymbol{x}, \boldsymbol{\xi}) \\ M(\boldsymbol{x}, \boldsymbol{\xi}) &= -\sum_{i=1}^{2} \sin(\xi_i + x_1)\ \sin^{20}\left(\frac{i\ (\xi_i + x_1)^2}{\pi}\right) \\ G(\boldsymbol{x}, \boldsymbol{\xi}) &= \frac{1}{2\pi x_2} \exp^{-\frac{(\xi_1 - \pi/2)^2}{2x_2} - \frac{(\xi_2 - 3\pi/4)^2}{2x_2}} \\ P(\boldsymbol{x}, \boldsymbol{\xi}) &= 0.2\ (\xi_1 - \pi/2)\ (\xi_2 - \pi/2) \end{aligned} \tag{B.6}$$

Term M is a modification of the Michalewicz function with a phase angle equal to the first design variable, term G represents a bell function employed to uncorrelate the mean and the minimum, and the function P is added to avoid the flat regions of the original function that would yield a high peak of the PDF at zero. The optimization problem seeks the set of design that maximize both the mean value and the minimum

$$\max_{\boldsymbol{x} \in \Sigma} \left(\mu_u, \min_u\right). \tag{B.7}$$

The optimization is performed by means of the Non-dominated Sorting Genetic Algorithm (NSGA) in the formulation by Deb [156] with a population of 40 individuals and a number of generations equal to 60.

Figure B.2a presents the Pareto fronts obtained using the minimum computation based on a Polynomial Chaos representation of 50-th order, the corresponding Quadrature Points, and a random sampling on the analytic function with 500000 samples which is taken as a reference. The SSC is not considered here for this explorative case. The solution found by the PC-based approach compares well with the reference result, while the approach based on the QP ends up on a different front and a different set of optimal designs, as presented in Fig. B.2b. Because the computation of the minimum based on the PC expansion proves robust and effective, it will be the method of choice in the following analysis.

The results of optimization (B.7) are now compared to the results obtained for

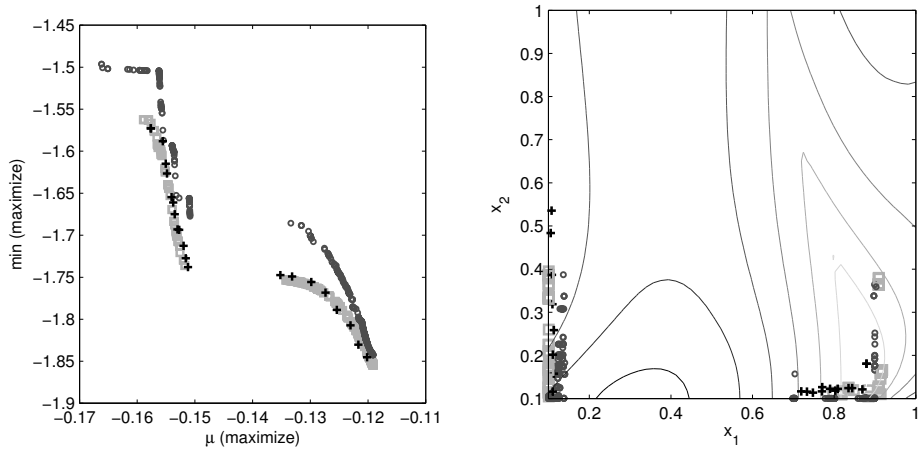

(a) Pareto fronts.

(b) Contour of the minimum with Pareto front designs and optimal Pareto designs.

Figure B.2: Optimization results for the two-dimensional modified Michalewicz function in Eq. (B.6) obtained with PC (light grey squares), QP (grey circles) and sampling on analytical function (black plus markers).

the same problem and with the formulation of robust design that maximizes the mean value and minimizes the variance σ^2:

$$\max_{\boldsymbol{x}\in\mathcal{X}} \ \mu_u \ \wedge \ \min_{\boldsymbol{x}\in\mathcal{X}} \sigma_u^2. \tag{B.8}$$

The Pareto fronts obtained with the two different formulations are presented in the $\mu, \min$-plane in Fig. B.3a and in the μ, σ^2-plane in Fig. B.3b. It appears that both strategies are capable of finding the sets of design with higher mean values, where both fronts completely overlap. However, the minimum-based optimization discards the branch composed of the designs with minimum variance, because they are dominated in the $\mu, \min$-plane. On the other hand, the front obtained by maximizing the mean and minimum values includes a branch with higher minimum values, which, in turn, the robust optimization cannot find. It is thus evident that the optimization based on the mean value and the minimum results in the subset with higher mean performance of the front obtained with the classical robust approach, but it also contains a branch with improved minimum values. Please note that in many applications, the solutions with very low variance comes at the expense of very poor performance values, to the extent that these solutions may be rejected in the decision making stage. In these cases, it is clear the benefit of a method that detects solutions with high mean

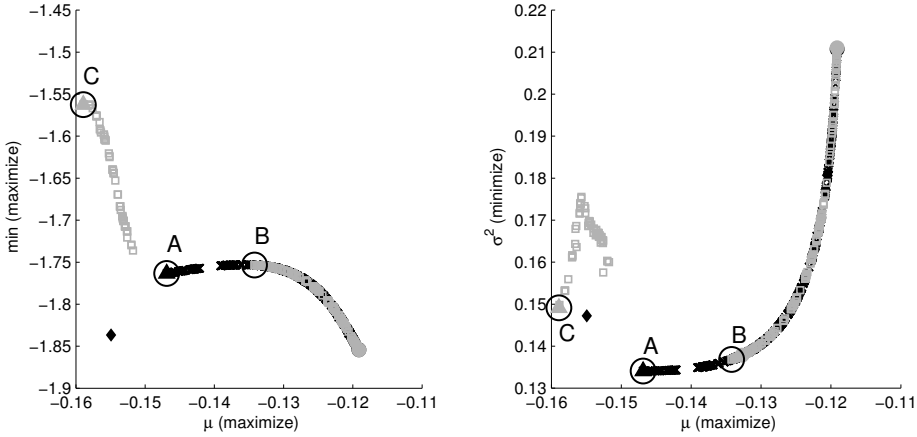

(a) Reference design (black diamond) and Pareto fronts in $\mu, \min$-plane.

(b) Reference design (black diamond) and Pareto fronts in μ, σ^2-plane.

Figure B.3: Optimization results with different formulations: minimum-based optimal solutions (light grey squares) and variance-based optimal solutions (black crosses)

value in the μ, σ^2-front.

Three optimal solutions from the non-overlapping branches are compared: from the optimization based on the mean value and variance, the design with optimal variance (A) is selected and, from the optimization based on the minimum, the designs with optimal variance (B) and optimal minimum value (C) are chosen. For these designs the PDF of function u is plotted in Fig. B.4. The functions share a similar shape, although the PDF of design B has higher mean value. It is interesting to note that for this particular case, despite a lower mean value of the performance, the design with greatest minimum may be an interesting solution for design purposes, because it possesses not only zero probability of low performances, but also the highest maximum value of the performance. This solution would never show up in the classical optimization based on the mean value and the variance. The optimal designs have been obtained without any assumptions on the final shape of the PDF, proving that the method is well suited for exploration of the optimization problem.

As a final remark, the size of the optimal set of designs obtained with the new formulation is comparable to that obtained with the classical approach. In fact, because the function of the test case in Eq. (B.6) is designed in order to uncorrelate the mean value, the result is a wide Pareto front. However, in the engineering case tackled in Ref. [154] the Pareto front obtained with the novel approach is much smaller with respect to the other; the difference in the size of the optimal set speeds up convergence

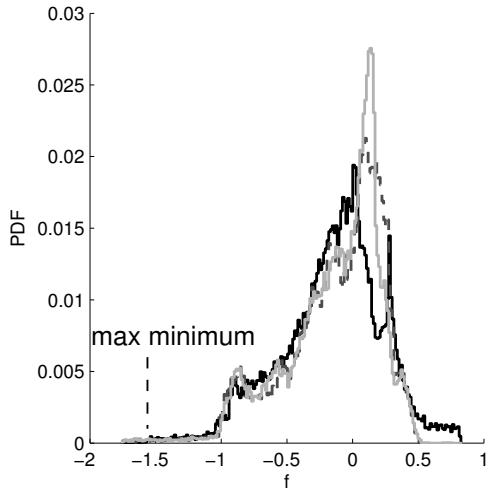

Figure B.4: PDF of selected optimal designs: design A (light grey solid line), design B (grey dashed line) and design C (black solid line).

	μ_u	σ_u^2	$\min_u$
A	-0.1469	0.1341	-1.763
B	-0.1342	0.1368	-1.753
C	-0.1589	0.1491	-1.562

Table B.1: Performance of the selected optimal designs.

of the genetic algorithm to the final results, reducing the computational effort of the novel formulation.

Conclusions

A novel formulation for design under uncertainty have been explored, which is aimed at maximizing both the mean value and the minimum of the performance. Three methods for computing the minimum are presented, which are based on the Polynomial Chaos expansion and the Simplex Stochastic Collocation. The method based on the quadrature points of the PC expansion appears less efficient with respect to the other strategies which directly exploit the continuous analytical approximations provided by these methods. The computation based on the PC expansion is effective and because it is the most straightforward and widely-employed method for uncer-

tainty quantification, it may be the method of choice for this new formulation. With regard to the impact on the optimization loop, the method has proven effective in selecting only the designs with higher mean performance from the front obtained with the classical robust approach, and in finding new interesting designs with improved minimum values. As a matter of fact, this exploratory study represents one of the first result in which a set of optimal designs is found which is discarded with a classical approach.

Future activities will be devoted to consolidate the computation of the minimum. In particular, the implementation of a refinement strategy based on the minimum in the SSC algorithm will be developed and assessed for different test cases with the aim of further reducing the computational cost of the new formulation. In addition, test cases of engineering interest will be solved with the new optimization formulation and compared to existing strategies. In this case, because reliability-based analysis seems to be a natural application of this formulation, test cases of this problem will be considered.

Bibliography

[1] J. G. Leishman, Principles of helicopter aerodynamics, Cambridge University Press, 2006.

[2] W. Johnson, CAMRAD-II.

[3] W. Johnson, Rotorcraft dynamics models for a comprehensive analysis, Tech. rep., Revised version of paper presented at the American Helicopter Society 54th Annual Forum (1998).

[4] W. Johnson, Rotorcraft aerodynamics models for a comprehensive analysis, Tech. rep., Revised version of paper presented at the American Helicopter Society 54th Annual Forum (1998).

[5] Mbdyn - free multibody dynamics simulation software.

[6] P. Masarati, Mbdyn theory and developer's manual, Tech. rep., Politecnico di Milano (2010).

[7] P. Masarati, D. J. Piatak, G. Quaranta, J. D. Singleton, J. Shen, Soft-inplane tiltrotor aeromechanics investigation using two comprehensive multibody solvers, Journal of the American Helicopter Society 53 (2) (2008) 179–192.

[8] A. R. S. Bramwell, G. Done, D. Balmford, Bramwell's Helicopter Dynamics-flight, 2nd Edition, Butterworth-Heinemann, 2001.

[9] W. Johnson, Rotorcraft Aeromechanics, Cambridge University Press, New York, 2013.

[10] J. Seddon, Basic Helicopter Aerodynamics, BSP Professional Books, Oxford, 1990.

[11] W. G. Bousman, Airfoil Design and Rotorcraft Performance., in: American Helicopter Society 58th Annual Forum, Montreal, Canada, 2002.

[12] W. G. Bousman, Aerodynamic Characteristics of SC1095 and SC1094 R8 Airfoils., Tech. Rep. TP-2003-212265, NASA (2003).

[13] W. McCroskey, L. Carr, K. McAlister, Dynamic stall experiments on oscillating airfoils, AIAA Journal 14 (1976) 57–63.

[14] T. S. Beddoes, A synthesis of unsteady aerodynamic effects including stall hysteresis, in: Vertica, 1976.

[15] J. G. Leishman, Challenges in modeling the unsteady aerodynamics of wind turbines, in: ASME 2002 Wind Energy Symposium Reno, Nevada, USA, January 14–17, 2002, pp. 141–167.

[16] D. W. Gross, F. D. Harris, Prediction of in-flight stalled airloads from oscillating airfoil data, in: 25th Annual National Forum of the American Helicopter Society, Washington D.C., May 14–16, 1969.

[17] R. E. Gormont, A mathematical model of unsteady aerodynamics and radial flow for application to helicopter rotors, Tech. Rep. 72, US ARMY Technical Report (1973).

[18] R. L. Bielawa, Synthesized unsteady airfoil data with applications to stall flutter calculations, in: 31st Annual National Forum of the American Helicopter Society, Washington D.C., 1975.

[19] Y. C. Fung, An Introduction to the Theory of Aeroelasticity, Dover, 2008.

[20] A. Morris, C. Allen, T. Rendall, Aerodynamic optimisation of hovering helicopter rotors using efficient and flexible shape parameterisation, in: 26th AIAA Applied Aerodynamics Conference, 2008.

[21] D. Fanjoy, W. A. Crossley, Aerodynamic shape design for rotor airfoils via genetic algorithm, Journal of the American Helicopter Society 43 (3) (1998) 263–270.

[22] B. R. Jones, W. A. Crossley, A. S. Lyrintzis, Aerodynamic and aeroacoustic optimization of rotorcraft airfoils via a parallel genetic algorithm, Journal of Aircraft 37 (6) (2000) 1088–1096.

[23] A. Massaro, E. Benini, Multi-Objective Optimization of Helicopter Airfoils Using Surrogate-Assisted Memetic Algorithms, Journal of Aircraft 49 (2) (2012) 375–383.

[24] S. K. Nadarajah, A. Jameson, Optimum Shape Design for Unsteady Flows with Time-Accurate Continuous and Discrete Adjoint Method, AIAA Journal 45 (7) (2007) 1478–1491.

[25] D. S. Lee, J. Periaux, E. Onate, L. F. Gonzalez, N. Qin, Active Transonic Aerofoil Design Optimization Using Robust Multiobjective Evolutionary Algorithms, Journal of Aircraft 48 (3) (2011) 1084–1094.

[26] C. A. Tatossian, S. K. Nadarajah, P. Castonguay, Aerodynamic shape optimization of hovering rotor blades using a non-linear frequency domain approach, Computers and Fluids 51 (1) (2001) 1–15.

[27] H. Farrokhfal, A. R. Pishevar, Aerodynamic shape optimization of hovering rotor blades using a coupled free wake–CFD and adjoint method, Aerospace Science and Technology 1 (2012) 1–10.

[28] S. Chae, K. Yee, C. Yang, T. Aoyama, S. Jeong, S. Obayashi, Helicopter Rotor Shape Optimization for the Improvement of Aeroacoustic Performance in Hover 47 (5).

[29] C. Allen, T. Rendall, Cfd-based optimization of hovering rotors using radial basis functions for shape parameterization and mesh deformation, Optimization and Engineering 14 (1) (2013) 97–118.

[30] C. B. Allen, T. C. S. Rendall, Computational-fluid-dynamics-based twist optimization of hovering rotors, Journal of Aircraft 47 (6) (2010) 2075–2085.

[31] A. Le Pape, P. Baumier, Numerical optimization of helicopter rotor aerodynamic performance in hover, Aerospace Science and Technology 9 (3) (2005) 191–201.

[32] D. Leusink, D. Alfano, P. Cinnella, J.-C. Robinet, Aerodynamic rotor blade optimization at Eurocopter - a new way of industrial rotor blade design, in:

51st AIAA Aerospace Sciences Meeting 7-10 January 2013, Grapevine, Texas, 2013, pp. 1–14.

[33] M. Imiela, High-fidelity optimization framework for helicopter rotors, Aerospace Science and Technology 1 (2011) 1–15.

[34] R. Celi, Recent Applications of Design Optimization to Rotorcraft - A Survey, Journal of Aircraft 36 (1).

[35] C. Siva, M. S. Murugan, R. Ganguli, Uncertainty Quantification in Helicopter Performance Using Monte Carlo Simulations, Journal of Aircraft 48 (5) (2011) 1503–1511.

[36] S. Murugan, R. Chowdhury, S. Adhikari, M. Friswell, Helicopter aeroelastic analysis with spatially uncertain rotor blade properties, Aerospace Science and Technology 16 (1) (2012) 29 – 39.

[37] S. Barbarino, O. Bilgen, R. M. Ajaj, M. I. Friswell, D. J. Inman, A Review of Morphing Aircraft, Journal of Intelligent Material Systems and Structures 22 (9) (2011) 823–877.

[38] J. Bailly, Y. Delrieux, Improvement of noise reduction and performance for a helicopter model rotor blade by active twist actuation, in: Proceedings of 35th European Rotorcraft Forum, 22–25 September, Hamburg, Germany, 2009.

[39] Q. Zhang, F. Hoffmann, Benefit studies for rotor with active twist control using weak fluid-structure coupling, in: Proceedings of 35th European Rotorcraft Forum, 22–25 September, Hamburg, Germany, 2009.

[40] J. Fincham, M. Friswell, Aerodynamic optimisation of a camber morphing aerofoil, Aerospace Science and Technology 43 (2015) 245 – 255.

[41] F. Gandhi, M. Frecker, A. Nissly, Design Optimization of a Controllable Camber Rotor Airfoil, AIAA Journal 46 (1) (2008) 142–153.

[42] Smart rotor.

[43] S. Kottapalli, Low speed and high speed correlation of smart active flap rotor loads, in: American Helicopter Society Specialists Conference on Aeromechanics, 2010.

[44] M. Potsdam, M. V. Fulton, A. Dimanlig, Multidisciplinary cfd/csd analysis of the smart active flap rotor, in: 66th Annual Forum of the American Helicopter Society, Phoenix, Arizona, May 11–13, 2010.

[45] M. S. Murugan, B. K. S. Woods, M. I. Friswell, Morphing helicopter rotor blade with curvilinear fiber composites, in: 38th European Rotorcraft Forum, 2012.

[46] A. De Gaspari, S. Ricci, A two-level approach for the optimal design of morphing wings based on compliant structures, Journal of Intelligent Material Systems and Structures 22 (10) (2011) 1091–1111.

[47] A. Airoldi, M. Crespi, G. Quaranta, G. Sala, Design of a Morphing Airfoil with Composite Chiral Structure, Journal of Aircraft 49 (4) (2012) 1008–1019.

[48] M. F. Kerho, Adaptive Airfoil Dynamic Stall Control, Journal of Aircraft 44 (4) (2007) 1350–1360.

[49] S. Murugan, B. Woods, M. Friswell, Hierarchical modeling and optimization of camber morphing airfoil, Aerospace Science and Technology 42 (2015) 31 – 38.

[50] P. B. Martin, K. McAlister, M.S.Chandrasekhara, W. Geissler, Dynamic stall measurements and computations for a vr-12 airfoil with a variable droop leading edge, in: 59th Annual Forum of the American Helicopter Society, Phoenix, Arizona, May 6–8, 2003.

[51] M. Trenker, Design concepts for adaptive airfoils with dynamic transonic flow control, Journal of Aircraft 40 (4) (2003) 734–740.

[52] W. Yao, X. Chen, W. Luo, M. van Tooren, J. Guo, Review of uncertainty-based multidisciplinary design optimization methods for aerospace vehicles, Progress in Aerospace Sciences 47 (6) (2011) 450–479.

[53] E. Obert, Aerodynamic Design of Transport Aircraft, IOS Press, Amsterdam, The Netherlands, 2009.

[54] M. Nikbay, M. N. Kuru, Reliability based multidisciplinary optimization of aeroelastic systems with structural and aerodynamic uncertainties, Journal of Aircraft 50 (3) (2013) 708–715.

[55] C. L. Pettit, Uncertainty Quantification in Aeroelasticity: Recent Results and Research Challenges, Journal of Aircraft 41 (5) (2004) 1217–1229.

[56] B. Mohammadi, O. Pironneau, Applied Shape Optimization for Fluids, Oxford University Press, Oxford, England, 2009.

[57] J. C. Helton, J. D. Johnson, W. L. Oberkampf, An exploration of alternative approaches to the representation of uncertainty in model predictions, Reliability Engineering & System Safety 85 (1–3) (2004) 39–71.

[58] Z. Tang, J. Périaux, Uncertainty based robust optimization method for drag minimization problems in aerodynamics, Computer Methods in Applied Mechanics and Engineering 217–220 (2012) 12–24.

[59] H.-G. Beyer, B. Sendhoff, Robust optimization – A comprehensive survey, Computer Methods in Applied Mechanics and Engineering 196 (33–34) (2007) 3190–3218.

[60] L. Huyse, S. L. Padula, R. M. Lewis, W. Li, Probabilistic Approach to Free-Form Airfoil Shape Optimization Under Uncertainty, AIAA Journal 40 (9) (2002) 1764–1772.

[61] G. Taguchi, M. Phadke, Quality engineering through design optimization, in: K. Dehnad (Ed.), Quality Control, Robust Design, and the Taguchi Method, Springer US, 1989, pp. 77–96.

[62] M. Trosset, Taguchi and robust optimization, Tech. rep., Rice University, Houston, TX (1996).

[63] A. Saltelli, M. Ratto, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, S. Tarantola, Global Sensitivity Analysis. The Primer, John Wiley & Sons Ltd, West Sussex, England, 2008.

[64] D. Lucor, C. Enaux, H. Jourden, P. Sagaut, Stochastic design optimization: Application to reacting flows, Computer Methods in Applied Mechanics and Engineering 196 (2007) 5047–5062.

[65] P. Congedo, C. Corre, J.-M. Martinez, Shape optimization of an airfoil in a BZT flow with multiple-source uncertainties, Computer Methods in Applied Mechanics and Engineering 200 (1-4) (2011) 216–232.

[66] R. A. E. Mäkinen, J. Periaux, J. Toivanen, Shape design optimization in 2d aerodynamics using genetic algorithms on parallel computers, in: Parallel Computational Fluid Dynamics, Elsevier Science, 1995, pp. 395–402.

[67] N. Srinivas, K. Deb, Multiobjective function optimization using nondominated sorting genetic algorithms, Evolutionary Computation 2 (3) (1994) 221–248.

[68] C. Audet, G. Savard, W. Zghal, Multiobjective optimization through a series of single–objective formulations, SIAM Journal on optimization 19 (1) (2008) 188–210.

[69] C. Audet, J. Dennis, Jr., Mesh adaptive direct search algorithms for constrained optimization, SIAM Journal on Optimization 17 (1) (2006) 188–217.

[70] S. Le Digabel, Algorithm 909: NOMAD: Nonlinear optimization with the MADS algorithm, ACM Transactions on Mathematical Software 37 (4) (2011) 1–15.

[71] M. Abramson, C. Audet, G. Couture, J. Dennis, Jr., S. Le Digabel, C. Tribes, The NOMAD project, Software available at `https://www.gerad.ca/nomad/`.

[72] Q. Wang, Uncertainty quantification for unsteady flow using adjoint-based approaches, Ph.D. thesis, Institute for Computational and Mathematical Engineering, Stanford University (2008).

[73] I. Park, R. V. Grandhi, A bayesian statistical method for quantifying model form uncertainty and two model combination methods, Reliability Engineering & System Safety 129 (2014) 46–56.

[74] G. Iaccarino, Quantification of uncertainty in flow simulations using probabilistic methods, in: VKI Lecture Series on Non-equilibrium gas dynamics, From physical models to hypersonic flight, 2009, pp. 1–29.

[75] W. L. Oberkampf, T. G. Trucano, Verification and validation in computational fluid dynamics, Progress in Aerospace Sciences 38 (3) (2002) 209–272.

[76] W. Edeling, P. Cinnella, R. Dwight, Predictive RANS simulations via bayesian model-scenario averaging, Journal of Computational Physics 275 (2014) 65–91.

[77] S. H. Cheung, T. A. Oliver, E. E. Prudencio, S. Prudhomme, R. D. Moser, Bayesian uncertainty analysis with applications to turbulence modeling, Reliability Engineering & System Safety 96 (9) (2011) 1137–1149, quantification of Margins and Uncertainties.

[78] C. J. Roy, W. L. Oberkampf, A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing, Computer Methods in Applied Mechanics and Engineering 200 (25-28) (2011) 2131–2144.

[79] J. Jakeman, M. Eldred, D. Xiu, Numerical approach for quantification of epistemic uncertainty, Journal of Computational Physics 229 (12) (2010) 4648–4663.

[80] M. P. Rumpfkeil, Optimizations Under Uncertainty Using Gradients, Hessians, and Surrogate Models, AIAA Journal 51 (2) (2013) 444–451.

[81] M. Padulo, M. S. Campobasso, M. D. Guenov, Novel Uncertainty Propagation Method for Robust Aerodynamic Design, AIAA Journal 49 (3) (2011) 530–543.

[82] M. Dodson, G. T. Parks, Robust Aerodynamic Design Optimization Using Polynomial Chaos, Journal of Aircraft 46 (2) (2009) 635–646.

[83] H. Shah, S. Hosder, L. Leifsson, S. Koziel, Y. A. Tesfahunegn, Multi-fidelity Robust Aerodynamic Design Optimization Under Mixed Uncertainty, in: 17th AIAA Non-Deterministic Approaches Conference, AIAA, Kissimmee, FL, 2015.

[84] Y. Zhang, S. Hosder, L. Leifsson, S. Koziel, Robust airfoil optimization under inherent and model-form uncertainties using stochastic expansions, in: 50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, AIAA, Nashville, TN, 2012.

[85] P. Spalart, S. Allmaras, A one-equation turbulence model for aerodynamic flows, in: 30th Aerospace Sciences Meeting and Exhibit, AIAA, Reno, NV, 1992.

[86] O. P. Le Maître, O. M. Knio, Spectral methods for uncertainty quantification. With Applications to Computational Fluid Dynamics., Springer, 2010.

[87] J. Jacod, P. E. Protter, Probability essentials, Springer Verlag, Berlin Heidelberg, 2004.

[88] J. K. Blitzstein, J. Hwang, Introduction to Probability, CRC Press, Boca Raton, FL, 2015.

[89] H. Shah, S. Hosder, T. Winter, A Mixed Uncertainty Quantification Approach with Evidence Theory and Stochastic Expansions, in: 16th AIAA Non-Deterministic Approaches Conference, AIAA, National Harbor, FL, 2014, pp. 1–25.

[90] R. Ghanem, P. Spanos, Spectral techniques for stochastic finite elements, Archives of Computational Methods in Engineering 4 (1) (1997) 63–100.

[91] D. Xiu, G. E. Karniadakis, The Wiener–Askey Polynomial Chaos for Stochastic Differential Equations, SIAM Journal on Scientific Computing 24 (2) (2002) 619–644.

[92] L. Mathelin, M. Hussaini, T. Zang, Stochastic approaches to uncertainty quantification in cfd simulations, Numerical Algorithms 38 (1-3) (2005) 209–236.

[93] J. A. S. Witteveen, G. Iaccarino, Refinement criteria for simplex stochastic collocation with local extremum diminishing robustness, SIAM Journal on Scientific Computing 34 (3) (2012) A1522–A1543.

[94] A. Quarteroni, R. Sacco, F. Saleri, Matematica Numerica, Springer, Milan, Italy, 2008.

[95] J. Foo, G. E. Karniadakis, Multi-element probabilistic collocation method in high dimensions, Journal of Computational Physics 229 (5) (2010) 1536 – 1557.

[96] X. Ma, N. Zabaras, An adaptive high-dimensional stochastic model representation technique for the solution of stochastic partial differential equations, Journal of Computational Physics 229 (10) (2010) 3884 – 3915.

[97] J. A. S. Witteveen, A. Doostan, R. Pecnik, G. Iaccarino, Uncertainty quantification of the transonic flow around the RAE 2822 airfoil, in: Annual Research Briefs 2009, 2009.

[98] J. A. S. Witteveen, G. Iaccarino, Simplex stochastic collocation with random sampling and extrapolation for nonhypercube probability spaces, SIAM Journal on Scientific Computing 34 (2) (2012) A814–A838.

[99] M. Drela, Xfoil: An analysis and design system for low reynolds number airfoils., Conference on Low Reynolds Number Airfoil Aerodynamics, University of Notre Dame 54 (1989) 1–12.

[100] N. A. Vu, J. W. Lee, J. I. Shu, Aerodynamic design optimization of helicopter rotor blades including airfoil shape for hover performance, Chinese Journal of Aeronautics 26 (1) (2013) 1–8.

[101] S. Choi, J. J. Alonso, I. M. Kroo, Two-level multifidelity design optimization studies for supersonic jets, Journal of Aircraft 46 (3) (2009) 776–790.

[102] T. D. Robinson, K. E. Willcox, E. M. S., R. Haimes, Multifidelity Optimization for Variable-Complexity Design, in: 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, AIAA, Portsmouth, VA, 2006.

[103] D. Leusink, D. Alfano, P. Cinnella, Multi-fidelity optimization strategy for the industrial aerodynamic design of helicopter rotor blades, Aerospace Science and Technology 42 (0) (2015) 136–147.

[104] K. Collins, L. Sankar, D. Mavris, Application of low- and high-fidelity simulation tools to helicopter rotor blade optimization, Journal of the American Helicopter Society 58 (4) (2013) 1–10.

[105] W. Yamazaki, Efficient robust design optimization by variable fidelity kriging model, in: 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, AIAA, Honolulu, HI, 2012.

[106] W. Yamazaki, D. J. Mavriplis, Derivative-enhanced variable fidelity surrogate modeling for aerodynamic functions, AIAA Journal 51 (1) (2103) 126–137.

[107] L. W.-T. Ng, M. Eldred, Multifidelity uncertainty quantification using non-intrusive polynomial chaos and stochastic collocation, in: 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials ConferenceAIAA, Honolulu, HI, 2012.

[108] Y. Jin, M. Olhofer, B. Sendhoff, A framework for evolutionary optimization with approximate fitness functions, IEEE Transactions on Evolutionary Computation 6 (2002) 481–494.

[109] H. Shen, E. McBean, Y. Wang, Sensor placement under nodal demand uncertainty for water distribution systems, in: R. M. Clark, S. Hakim (Eds.), Secur-

ing Water and Wastewater Systems, Vol. 2 of Protecting Critical Infrastructure, Springer International Publishing, 2014, pp. 123–133.

[110] L. Painton, U. Diwekar, Stochastic annealing for synthesis under uncertainty, European Journal of Operational Research 83 (3) (1995) 489 – 502.

[111] U. M. Diwekar, W. Xu, Improved genetic algorithms for deterministic optimization and optimization under uncertainty. part i. algorithms development, Industrial & Engineering Chemistry Research 44 (18) (2005) 7132–7137.

[112] S. Sankaran, C. Audet, A. L. Marsden, A method for stochastic constrained optimization using derivative-free surrogate pattern search and collocation, Journal of Computational Physics 229 (12) (2010) 4664 – 4682.

[113] P. Congedo, J. Witteveen, G. Iaccarino, A simplex-based numerical framework for simple and efficient robust design optimization, Computational Optimization and Applications 56 (1) (2013) 231–251.

[114] M. Hendriks, M. Geilen, T. Basten, Pareto analysis with uncertainty, in: Proceedings of the 2011 IFIP 9th International Conference on Embedded and Ubiquitous Computing, EUC '11, IEEE Computer Society, Washington, DC, USA, 2011, pp. 189–196.

[115] M. Mlakar, T. Tušar, B. Filipič, Comparing solutions under uncertainty in multiobjective optimization, Mathematical Problems in Engineering 2014 (2014) 1–10.

[116] X. S. Yang, S. Deb, Engineering optimization by cuckoo search, International Journal of Mathematical Modelling and Numerical Optimization 4 (2010) 330–343.

[117] B. G. van der Wall, 2nd HHC Aeroacoustic Rotor Test (HART II) - Part I: Test Documentation., Braunschweig, Germany, 2003.

[118] J. Leishman, Principles of helicopter aerodynamics. 2nd edition, Cambridge University Press, 2006.

[119] V. Pediroda, C. Poloni, A. Clarich, A fast and robust adaptive methodology for airfoil design under uncertainties based on game theory and self-organising-map theory, in: 44th AIAA Aerospace Sciences Meeting and Exhibit, AIAA, Reno, NV, 2006.

[120] F. Simon, P. Guillen, P. Sagaut, D. Lucor, A gpc-based approach to uncertain transonic aerodynamics, Computer Methods in Applied Mechanics and Engineering 199 (2010) 1091–1099.

[121] I. Abdallah, A. Natarajan, J. SÃ¸rensen, Impact of uncertainty in airfoil characteristics on wind turbine extreme loads, Renewable Energy 75 (2015) 283–300.

[122] A. H. Shapiro, The Dynamics and Thermodynamics of Compressible Fluid Flow, Vol. I, The Ronald Press Company, 1953.

[123] M. B. Giles, M. Drela, Two-dimensional transonic aerodynamic design method, AIAA Journal 25 (9) (1987) 1199–1206.

[124] M. Drela, Implicit implementation of the full e^n transition criterion, in: 21st Applied Aerodynamics Conference, Orlando, FL, 2003.

[125] M. B. Giles, M. Drela, Viscous-inviscid analysis of transonic and low reynolds number airfoils, AIAA Journal 25 (10) (1987) 1347–1355.

[126] B. M. Kulfan, J. E. Bussoletti, "Fundamental" Parametric Geometry Representations for Aircraft Component Shapes, in: 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Portsmouth, VA, 2006.

[127] L. Marti, J. García, An approach to stopping criteria for multi-objective optimization evolutionary algorithms: the MGBM criterion, Evolutionary Computation (2009) 1263–1270.

[128] T. Goel, N. Stander, A study on the convergence of multiobjective evolutionary algorithms, in: 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference, Forth Worth, TX, 2010.

[129] Aerodynamic design optimization discussion group.
URL `https://info.aiaa.org/tac/ASG/APATC/AeroDesignOpt-DG`

[130] F. Fusi, A. Guardone, G. Quaranta, P. M. Congedo, Drag minimization of an isolated airfoil in transonic inviscid flow by means of genetic algorithms, in: 17th AIAA Non-Deterministic Approaches Conference, AIAA, Kissimmee, FL, 2015.

[131] L. Leifsson, S. Koziel, Multi-fidelity design optimization of transonic airfoils using physics-based surrogate modeling and shape-preserving response prediction, Journal of Computational Science 1 (2) (2010) 98–106.

[132] A. M. Morris, C. B. Allen, T. C. S. Rendall, Development of generic cfd-based aerodynamic optimisation tools for helicopter rotor blades, in: 25th AIAA Applied Aerodynamics Conference, AIAA, Miami, FL, 2007.

[133] M. Drela, A User's Guide to MSES 3.05, MIT Department of Aeronautics and Astronautics, 2007.

[134] F. Fusi, P. M. Congedo, An adaptive strategy on the error of the objective functions for uncertainty-based derivative-free optimization, Journal of Computational Physics.

[135] D. J. Gosselin, 2d dynamic stall simulations with time-varying freestream representative of helicopter flight, Ph.D. thesis, Ottawa-Carleton Institute for Mechanical & Aerospace Engineering (2014).

[136] D. Favier, A. Agnes, C. Barbi, C. Maresca, A new simulation of airfoil dynamic stall due to velocity and incidence fluctuations, in: 19th AIAA, Fluid Dynamics, Plasma Dynamics and Lasers Conference, American Institute of Aeronautics and Astronautics, 1987.

[137] D. Favier, C. Maresca, C. Barbi, Unsteady aerodynamics of an airfoil in combined translation/pitch oscillations below and through stall, in: 1st National Fluid Dynamics Conference, American Institute of Aeronautics and Astronautics, 1988.

[138] B. G. van der Wall, J. G. Leishman, On the influence of time-varying flow velocity on unsteady aerodynamics, Journal of The American Helicopter Society 39 (4) (1994) 25 – 36.

[139] B. Glaz, L. Liu, P. Friedmann, J. Bain, L. Sankar, A surrogate based approach to reduced-order dynamic stall modeling, in: 51st AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, American Institute of Aeronautics and Astronautics, 2010.

[140] A. I. Jose, J. G. Leishman, J. D. Baeder, Unsteady aerodynamic modeling with time-varying free-stream mach numbers, Journal of the American Helicopter Society 51 (4) (2006) 299–318.

[141] K. Gharali, D. A. Johnson, Dynamic stall simulation of a pitching airfoil under unsteady freestream velocity, Journal of Fluids and Structures 42 (2013) 228 – 244.

[142] C. Strangfeld, H. Mueller-Vahl, D. Greenblatt, C. Nayeri, C. O. Paschereit, Airfoil subjected to high-amplitude free-stream oscillations: theory and experiments, in: 7th AIAA Theoretical Fluid Mechanics Conference, AIAA Aviation, American Institute of Aeronautics and Astronautics, 2014.

[143] F. Palacios, , M. R. Colonno, A. C. Aranake, A. Campos, S. R. Copeland, T. D. Economon, A. Lonkar, T. W. Lukaczyk, T. W. R. Taylor, J. J. Alonso, Stanford University Unstructured (SU2): An open-source integrated computational environment for multi-physics simulation and design, 51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition.

[144] J. D. J. Anderson, Introduction to flight, McGraw-Hill International Edition, 2005.

[145] C. Strangfeld, H. Mueller-Vahl, D. Greenblatt, B. G. van der Wall, C. Nayeri, C. O. Paschereit, Pitching airfoil subjected to high amplitude free stream oscillations, in: European Rotorcraft Forum 2015, 2015.

[146] F. Palacios, T. D. Economon, A. C. Aranake, S. R. Copeland, A. K. Lonkar, T. W. Lukaczyk, D. E. Manosalvas, K. R. Naik, A. S. Padron, B. Tracey, A. Variyar, J. J. Alonso, Stanford University Unstructured (SU2): Open-source Analysis and Design Technology for Turbulent Flows, in: 52nd Aerospace Sciences Meeting, 2014.

[147] C. Hirsch, Numerical Computation of Internal and External Flows, second edition Edition, Butterworth-Heinemann, Oxford, 2007.

[148] G. Kenway, J. R. R. A. Martins, Multipoint high-fidelity aerostructural optimization of a transport aircraft configuration, Journal of Aircraft 51 (1) (2014) 144–160.

[149] M. E. Pontecorvo, S. Barbarino, F. S. Gandhi, Cellular honeycomb-like structures with internal inclusions in the unit-cell, in: ASME 2012 Conference on Smart Materials, Adaptive Structures and Intelligent Systems, Georgia, USA, September 19–21, 2012, pp. 383–393.

[150] I. H. Abbott, A. E. von Doenhoff, Theory of wing section, Dover Publications, Inc., New York, 1949.

[151] J. C. Vassberg, N. Harrison, D. Roman, A. Jameson, A systematic study on the impact of dimensionality for a two-dimensional aerodynamic optimization model problem, in: 29th AIAA Applied Aerodynamics Conference, Honolulu, Hawaii, 2011.

[152] P. Seshadri, P. Constantine, G. Iaccarino, G. Parks, Aggressive Design Under Uncertainty, in: 16th AIAA Non-Deterministic Approaches Conference, AIAA SciTech, 2014. arXiv:arXiv:1409.7089v1.

[153] G. Petrone, G. Iaccarino, D. Quagliarella, Robustness Criteria in Optimization under Uncertainty, in: EUROGEN 2011, Capua, Italy 2011, 2011.

[154] P. M. Congedo, G. Geraci, G. Iaccarino, On the use of high-order statistics in robust design optimization, in: 5th European Conference on Computational Mechanics (ECCM V), July 20-25, 2014, Barcelona, Spain, 2014.

[155] G. Schuëller, H. Jensen, Computational methods in optimization considering uncertainties – An overview, Computer Methods in Applied Mechanics and Engineering 198 (1) (2008) 2–13.

[156] N. Srinivas, K. Deb, Multiobjective function optimization using nondominated sorting genetic algorithms, Evol. Comput. 2 (1995) 221–248.