



UNIVERSITÀ DEGLI STUDI DI PADOVA

DIPARTIMENTO DI INGEGNERIA INDUSTRIALE DII

Corso di Laurea Magistrale in Ingegneria Aerospaziale

**MULTI-OBJECTIVE AERODYNAMIC OPTIMISATION
OF A MORPHING WING**

Relatore

Prof. Ernesto Benini

Laureando

Andrea Magrini

mat. 1105705

ANNO ACCADEMICO 2017/2018

Abstract

The thesis presents an aerodynamic optimisation of a morphing leading edge airfoil. The NASA GA(W)-1 general aviation airfoil is allowed to morph between 0 and 25% of its chord and is represented by the CST parameterization technique, employing an original dedicated procedure to keep the curve arc length constant in the deformable part, in order to limit the axial stress. Three validated aerodynamic tools are used: XFOIL, the RANS Spalart-Allmaras model and the RANS transitional model $\gamma - Re_\theta$ or Transition SST. An optimisation loop based on genetic algorithms is applied to both single-objective and multi-objective problems and is extended also to a three-dimensional morphing leading edge wing. Results at varying the formulation, the optimiser settings and the aerodynamic solver are presented and discussed. Finally, a meta-model based optimisation is enquired, providing an approximation of the aerodynamic data obtained with the three models. A meta-model assisted optimisation loop is also used to solve a previously known single-objective problem to compare the convergence performance.

La tesi presenta l'ottimizzazione aerodinamica di un profilo con bordo d'attacco deformabile (morphing). Un profilo subcritico di tipo NASA GA(W)-1 è stato parametrizzato con la tecnica CST, impiegando un'apposita procedura originale per mantenerne costante la lunghezza nella zona deformabile, estesa tra 0 e 25% della corda. Tre modelli aerodinamici validati sono stati impiegati: XFOIL, il modello RANS di Spalart e Allmaras ed il modello di transizione RANS $\gamma - Re_\theta$ o Transition SST. Un ciclo di ottimizzazione basato su algoritmi genetici è stato applicato a problemi singolo- o multi-obiettivo, estendendo anche ad una configurazione tridimensionale, con un'ala a bordo d'attacco deformabile. Per ciascuna formulazione proposta sono stati presentati e discussi i risultati ottenuti con diverse impostazioni dell'ottimizzatore e differenti solutori aerodinamici. Infine si è indagato l'uso di modelli surrogati all'interno del ciclo di ottimizzazione, che forniscono un'approssimazione dei dati aerodinamici precedentemente ottenuti con i tre modelli. Tramite un'ottimizzazione assistita da un meta-modello si è risolto un problema singolo-obiettivo già noto, al fine di comparare l'efficacia nella convergenza per i due casi.

Ringraziamenti

Questo lavoro è il frutto delle competenze acquisite durante l'intero ciclo degli studi universitari. Desidero ringraziare quanti hanno reso possibile questo percorso, sostenendomi materialmente e moralmente: i miei genitori e familiari ed i compagni di corso con i quali ho condiviso l'impegno e le soddisfazioni. Un ringraziamento specifico per il lavoro svolto va al Prof. Benini, per la sua costante disponibilità e gentilezza. Ringrazio altresì il Dott. Dal Monte per il fondamentale aiuto nell'utilizzo del cluster di calcolo.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 1 |
| 1.1 | Objectives of the present thesis | 2 |
| 1.2 | Summary of the thesis | 3 |
| 2 | Morphing Systems: an Overview | 5 |
| 2.1 | Introduction | 5 |
| 2.2 | Morphing concepts | 5 |
| 2.3 | Morphing airfoils | 8 |
| 3 | Morphing Leading Edge Devices | 11 |
| 3.1 | Introduction to morphing leading edges | 11 |
| 3.2 | Morphing leading edge concepts | 13 |
| 3.3 | Morphing leading edge design | 15 |
| 3.3.1 | Aerodynamic design | 15 |
| 3.3.2 | Structural design | 17 |
| 3.3.3 | Actuation Mechanism design | 19 |
| 4 | Aerodynamic Optimisation of a Morphing Wing | 23 |
| 4.1 | Problem statement | 23 |
| 4.1.1 | Reference configurations | 24 |
| 4.2 | Standard optimisation framework | 25 |
| 4.3 | Synthetic summary | 26 |
| 5 | Shape Parameterization | 29 |
| 5.1 | Available parameterization methods | 30 |
| 5.1.1 | Discrete methods | 30 |
| 5.1.2 | Analytical methods | 30 |
| 5.1.3 | Polynomials and Splines | 31 |
| 5.1.4 | Other methods | 32 |
| 5.2 | Parameterization methods suitable for morphing | 32 |
| 5.2.1 | CST formulation | 33 |
| 5.3 | Airfoil representation through CST | 35 |
| 5.3.1 | Inverse fitting via CST | 38 |

| | | |
|----------|--|------------|
| 5.3.2 | Application of CST to morphing airfoils | 39 |
| 5.4 | Constant length curve parameterization with CST | 43 |
| 6 | Aerodynamic Analysis Tool Validation | 47 |
| 6.1 | CFD simulation set-up | 48 |
| 6.1.1 | Grid Generation | 48 |
| 6.1.2 | Material Properties and Reference Quantities | 49 |
| 6.1.3 | Flow equations | 49 |
| 6.2 | Test case validation | 53 |
| 6.3 | Comparison between CFD and XFOIL results | 56 |
| 6.4 | Further validation activity | 59 |
| 6.4.1 | Transition prediction | 60 |
| 6.5 | Comparison of results for various turbulence models | 65 |
| 7 | Two-dimensional Optimisation of a Morphing Leading Edge Airfoil | 71 |
| 7.1 | Mathematical formulation of the optimisation problem | 71 |
| 7.2 | Optimisation algorithm | 72 |
| 7.2.1 | Genetic algorithms | 72 |
| 7.2.2 | NSGA-II algorithm | 75 |
| 7.3 | Design variables | 77 |
| 7.4 | Single objective optimisation using XFOIL | 77 |
| 7.4.1 | Single objective optimisation at DP1 | 79 |
| 7.5 | Multi-objective optimisation using XFOIL | 83 |
| 7.5.1 | Application of a MOOP to a practical case | 85 |
| 7.6 | Single objective optimisation using CFD | 90 |
| 7.6.1 | Single objective optimisation based on the Spalart-Allmaras model | 91 |
| 7.6.2 | Single objective optimisation based on the Transition SST model | 95 |
| 8 | Three-dimensional Optimisation of a Wing with a Morphing Leading Edge | 99 |
| 8.1 | Reference geometry | 99 |
| 8.2 | Geometry handling | 100 |
| 8.3 | CFD solver overview | 101 |
| 8.4 | Single objective optimisation | 102 |
| 8.5 | Multi-objective optimisation | 104 |
| 9 | Meta-model Based Optimisation | 109 |
| 9.1 | Application of surrogate models | 110 |
| 9.2 | Artificial neural networks | 113 |
| 9.3 | Two-dimensional optimisation with complete fitness replacement by ANN | 115 |
| 9.3.1 | ANN training | 116 |
| 9.3.2 | Results of the optimisation sessions | 119 |
| 9.4 | Approximation of the CFD data by a regression model | 120 |
| 9.4.1 | Surrogate model for the Spalart-Allmaras dataset | 121 |
| 9.4.2 | Surrogate model for the Transition-SST dataset | 122 |

| | | |
|-----------|--|------------|
| 9.5 | Meta-model assisted optimisation | 123 |
| 10 | Conclusions | 125 |
| | Bibliography | 127 |

List of Figures

| | | |
|------|--|----|
| 1.1 | An example of morphing in eagle’s flight | 2 |
| 2.1 | Possible morphing mechanisms | 6 |
| 2.2 | An overview of aircraft with morphing devices [7] | 6 |
| 2.3 | Four examples of airfoil morphing in literature | 9 |
| 3.1 | Two different patents for droop nose devices | 14 |
| 3.2 | Airbus gapless high lift leading edge system of A350 [72] | 14 |
| 3.3 | Two different mechanisms for morphing leading edges | 15 |
| 3.4 | SADE Smart Leading Edge Concept for the ground test [37] | 18 |
| 3.5 | Eccentric beam actuator in neutral and deployed position [53] | 18 |
| 3.6 | SARISTU drive chain actuation mechanism [40] | 21 |
| 4.1 | NASA GA(W)-1 general aviation purpose aerofoil | 24 |
| 4.2 | Three-dimensional view of the reference wing | 25 |
| 4.3 | Synthetic scheme of the basic optimisation framework employed | 27 |
| 5.1 | Representation of the <i>partition of unit</i> property for the Bernstein polynomials [41] | 35 |
| 5.2 | Geometric parameters of the CST representation | 36 |
| 5.3 | RMS error versus BPO n for NACA4412 airfoil | 39 |
| 5.4 | Parameterization results for a NACA4412 with BPO = 10 | 40 |
| 5.5 | Morphing LE example with CST parameterization for NASA HSNLF(1)-0213 profile | 41 |
| 5.6 | Residual error between perturbed and original CST parametric NASA HSNLF(1)-0213 profile in the fixed region | 42 |
| 5.7 | An example of family of morphed LE for a NASA GA(W)-1 airfoil | 43 |
| 5.8 | Relative variation of the length on the morphed region of an airfoil as a function of the variation of the third curvature coefficient A_2 | 44 |
| 5.9 | Scheme of the algorithm employed to obtain a new curve with imposed length variation | 45 |
| 5.10 | Scheme of the algorithm employed to obtain a constant length curve when some coefficients are imposed | 45 |
| 5.11 | An example of morphing LE airfoil with constant length | 46 |

| | | |
|------|--|----|
| 6.1 | Details of the hybrid grid used | 49 |
| 6.2 | Details of the hybrid grid used | 49 |
| 6.3 | $C_l - \alpha$ validation results for $M = 0.20, Re = 6.2 \times 10^6$ | 53 |
| 6.4 | Lift-Drag polar validation results for $M = 0.20, Re = 6.2 \times 10^6$ | 54 |
| 6.5 | Variation of RMS residuals with the number of grid elements, computed for $M = 0.20, Re = 6.2 \times 10^6$ | 55 |
| 6.6 | $C_l - \alpha$ validation results for $M = 0.20, Re = 6.2 \times 10^6$ | 57 |
| 6.7 | $C_d - C_l$ polar validation results for $M = 0.20, Re = 6.2 \times 10^6$ | 57 |
| 6.8 | $C_m - \alpha$ validation results for $M = 0.20, Re = 6.2 \times 10^6$ | 58 |
| 6.9 | C_p validation results for for CFD and XFOIL at $M = 0.15, Re = 6.3 \times 10^6, \alpha = 4.17$ [deg] | 58 |
| 6.10 | Details of a structured grid | 60 |
| 6.11 | Scheme of the natural transition phenomenon over a flat plate | 62 |
| 6.12 | $C_l - \alpha$ polar comparison between various turbulence models, computed at $M = 0.20, Re = 6.2 \times 10^6$ | 66 |
| 6.13 | Lift-Drag polar comparison between various turbulence models, computed at $M = 0.20, Re = 6.2 \times 10^6$ | 67 |
| 6.14 | $C_m - \alpha$ polar comparison between various turbulence models, computed at $M = 0.20, Re = 6.2 \times 10^6$ | 67 |
| 6.15 | $C_l - \alpha$ polar comparison for different meshes, employing the $\gamma - Re_\theta$ model, computed at $M = 0.20, Re = 6.2 \times 10^6$ | 69 |
| 6.16 | Lift-Drag polar comparison for different meshes, employing the $\gamma - Re_\theta$ model, computed at $M = 0.20, Re = 6.2 \times 10^6$ | 70 |
| 6.17 | $C_m - \alpha$ polar comparison for different meshes, employing the $\gamma - Re_\theta$ model, computed at $M = 0.20, Re = 6.2 \times 10^6$ | 70 |
| 7.1 | A general scheme for GA | 73 |
| 7.2 | An example of Pareto surface | 74 |
| 7.3 | Result for a SOO for minimum C_d at DP1, with only bounds on the DVs as constraints | 78 |
| 7.4 | Evolution of mean fitness with different mutation operators | 80 |
| 7.5 | Influence of crossover fraction | 81 |
| 7.6 | Optimised airfoil at DP1 with XFOIL and C_p distribution (dashed = baseline) | 82 |
| 7.7 | Three Pareto Fronts for a different <i>Pareto Fraction</i> | 84 |
| 7.8 | Spread of the DVs along the Pareto Fronts for three PF | 84 |
| 7.9 | Pareto Fronts for different population sizes | 85 |
| 7.10 | Pareto Fronts for different population sizes | 88 |
| 7.11 | Convergence history for an optimisation run with PS=100 after 100 generations | 88 |
| 7.12 | Decision variables scatter along the Pareto front for two sessions with PS=100 | 89 |
| 7.13 | Off-design post-check conditions attainment for points in two Pareto fronts | 90 |

| | | |
|------|--|-----|
| 7.14 | Aerofoils at the bounds of the Pareto front attaining the post-check conditions | 91 |
| 7.15 | Baseline (red) and optimised airfoil with CFD and SA at DP1 | 92 |
| 7.16 | Baseline and optimised pressure coefficient distribution for SOO at DP1 using CFD with SA | 93 |
| 7.17 | Contours of Mach number for the baseline (lhs) and the optimised (rhs) airfoil for SOO at DP1 using CFD with SA | 93 |
| 7.18 | Baseline and optimised airfoil by CFD with $\gamma - Re_\theta$ at DP1 | 95 |
| 7.19 | Baseline and optimised pressure coefficient distribution for SOO at DP1 using CFD with $\gamma - Re_\theta$ | 96 |
| 7.20 | Contours of Mach number for the baseline (lhs) and the optimised (rhs) airfoil for SOO at DP1 using CFD with $\gamma - Re_\theta$ | 96 |
| 7.21 | Contours of turbulent kinetic energy k for the baseline (lhs) and the optimised (rhs) airfoil for SOO at DP1 using CFD with $\gamma - Re_\theta$ | 97 |
| 8.1 | Planar view of the reference wing, with the reference sections indicated | 100 |
| 8.2 | Three-dimensional view of the reference wing | 101 |
| 8.3 | Illustration of the computational domain | 103 |
| 8.4 | Baseline and optimised airfoil obtained with the 2D MOOP employing the SA model | 104 |
| 8.5 | Pareto fronts obtained after the 3D SO problem (Initial Pareto) and after the 3D MO problem (Final Pareto) | 105 |
| 8.6 | Wing section at station 2 for the shape minimising the C_d in the MOOP, compared to the baseline geometry | 106 |
| 9.1 | Functional scheme of a single neuron | 113 |
| 9.2 | A feed-forward three-layer neural network | 114 |
| 9.3 | Bar plot on the relative error between test data and the prediction of NN 1 | 118 |
| 9.4 | Bar plot on the relative error between test data and the prediction of NN 1 | 119 |
| 9.5 | Comparison between the optimal solutions provided using ANN and XFOIL, in terms of decision variables | 120 |
| 9.6 | Comparison between the optimal solution provided using ANN and XFOIL, in terms of aerodynamic coefficients | 121 |
| 9.7 | Histogram of relative error distribution for drag prediction of the SA using different models | 122 |
| 9.8 | Histogram of relative error distribution for lift and moment prediction of the SA using different models | 122 |
| 9.9 | Scheme of the surrogate model assisted optimisation loop | 123 |
| 9.10 | Baseline and optimised airfoil found with the meta-model assisted loop | 124 |

List of Tables

| | | |
|-----|---|-----|
| 4.1 | Design Points (DPs) considered | 25 |
| 5.1 | Represented class of geometries with various Class function exponents . | 34 |
| 5.2 | Geometric variables in CST model | 37 |
| 6.1 | Summary of CFD settings for the Spalart-Allmaras model | 56 |
| 6.2 | Summary of CFD settings for the Transition SST model | 69 |
| 7.1 | Comparison between baseline and optimised aerodynamic coefficients for the first SOO at DP1 | 78 |
| 7.2 | Comparison between the aerodynamic coefficients of the baseline and op- timised airfoil with XFOIL at DP1 | 82 |
| 7.3 | Design Points considered in the MOOP | 87 |
| 7.4 | Comparison between baseline and CFD-SA optimised airfoil aerodynamic coefficients at DP1 | 92 |
| 7.5 | Comparison between the best solution obtained by XFOIL, calculated also with the CFD and SA model | 94 |
| 7.6 | Comparison between the best solution obtained by CFD with SA, calcu- lated also with XFOIL | 94 |
| 7.7 | Comparison between baseline and CFD with $\gamma - Re_\theta$ optimised airfoil aerodynamic coefficients at DP1 | 97 |
| 8.1 | Comparison between the aerodynamic indexes of the baseline and opti- mised airfoil for the 2D multi-objective problem discussed | 104 |
| 8.2 | Comparison between the aerodynamic coefficients of the baseline and opti- mised airfoil for the 3D multi-objective problem discussed | 106 |
| 9.1 | Features of the three ANNs presented | 117 |
| 9.2 | Comparison between the aerodynamic coefficients of the baseline and the optimised airfoil, found with the metamodel assisted loop | 124 |

Abbreviations

- AEC** Adaptive Evolution Control.
- ANN** Artificial Neural Network.
- AoA** Angle of Attack.
- BL** Boundary Layer.
- BPO** Bernstein's Polynomial Order.
- CAD** Computer-Aided Drafting.
- CF** Crossover Fraction.
- CFD** Computational Fluid Dynamics.
- CST** Class/Shape Transformation.
- DFR** Direct Fitness Replacement.
- DNS** Direct Numerical Simulation.
- DoE** Design of Experiment.
- DP** Design Point.
- DV** Decision or Design Variable.
- EA** Evolutionary Algorithm.
- FEC** Fixed Evolution Control.
- FEM** Finite Element Method.
- GA** Genetic Algorithm.
- GB** Gradient-Based.

GLR Generalized Linear Regression.

GPR Gaussian Process Regression (Kriging).

IFR Indirect Fitness Replacement.

LDG Landing.

LE Leading Edge.

LHS Latin Hypercube Sampling.

LSE Least Square Error.

MM Meta-model.

MOOP Multi-Objective Optimisation Problem.

MSE Mean Square Error.

NEC No Evolution Control.

NN Neural Network.

PF Pareto Fraction.

PS Population Size.

RANS Reynolds Averaged Navier-Stokes.

RBF Radial Basis Function.

RMS Root Mean Square.

SA Spalart-Allmaras.

SDS Selectively Deformable Structures.

SIMP Simple Isotropic Material with Penalization.

SM Surrogate Model.

SOO Single-Objective Optimisation.

SOOP Single-Objective Optimisation Problem.

SVR Support Vector Regression.

TE Trailing Edge.

TO Take-Off.

TO Topology Optimisation.

UAV Unmanned Aerial Vehicle.

Nomenclature

A_i Curvature correction coefficient in CST.

C_d Drag coefficient.

C_l Lift coefficient.

C_m Moment coefficient.

C_p Pressure coefficient.

K Mean kinetic energy.

L Length.

M Mach number.

Re Reynolds number.

α angle of attack.

γ Intermittency.

μ Dynamic viscosity.

ν Kinematic viscosity.

ψ Normalized abscissa x/c .

ζ Normalized ordinate y/c or z/c .

c chord.

k Turbulent kinetic energy.

Chapter 1

Introduction

The interest of men for flight has very ancient roots. The possibility of soaring into the air and moving without limitations has always fascinated men and birds appear in cave pictures or in bird-shaped neolithic statues. The mysterious geometries described by birds in flight were associated to doom and prophecies and the ability of decoding the significance of the animals' behaviour was an art reserved to few people. One of the first scientific observation of birds' flight can be attributed to Leonardo Da Vinci, who also proposed a flying machine on the basis of his understanding of the flight mechanisms that make possible for birds to soaring into the air. It is not surprising, then, that the natural world was of main inspiration for the birth of modern aeronautics. The first heavier-than-air aircraft with an engine, the Wright flyer, used cables to vary the twist of wings in order to control roll, as observed in birds. The capability of natural organisms to adapt their shape to different flight conditions, in order to exhibit the best response to the environmental state, oriented the early years of aeronautic research towards biomimetic strategies, based on flexible and highly deformable structures.

The development of more powerful engines and the increasing in flight speed and payload capacity led to more rigid frames, resisting the higher aerodynamic and inertial loads, which, indeed, behave optimally only for a phase of the flight. In particular, each aircraft is optimised to operate at best for a given flight condition, such us high speed cruise for transportation aircraft or low speed loitering for surveillance UAVs. However, the starting and final phases of flight, the take-off and landing, require a different behaviour of the aircraft, in order to guarantee high lift at low speed with minimum drag at take-off and sufficient drag and high lift at landing. The demanding of such a variation in the performance of airplanes motivates the use of slats and flaps to vary the camber of the airfoils composing the wing. Thus, changing the shape of the wing section is the most common way to adapt the performance of a system optimized for a range of conditions, when operating outside this range.

Coming back to nature, the variation in wing shape during the flight is very accentuated in birds and allows them to move with unsuspected agility in the air, executing manoeuvres not possible for a human artefact and landing precisely in a very short space. This is illustrated in Figure 1.1. The ability of manipulating certain characteristics of a

vehicle in order to better match the vehicle's state to the environment and increase its performance is referred to as *morphing* [29]. Actually, there is not a complete agreement about what shape change must be involved in order to speak about morphing, but it can be stated that it always requires a modification of the geometry to improve the behaviour and the performance in a given state. In more recent times, a new interest to the biomimetic paradigm has risen and the natural world can still have an influence on the technology development, not only in aeronautics but also in many other disciplines, starting for example from nanotechnology.



Figure 1.1: An example of morphing in eagle's flight

Adapting the shape of the wings according to the flight phase is a strategy to investigate, in order to extend the optimal range of operation and achieve a better efficiency also in off-design conditions. The Advisor Council for Aeronautics Research in Europe (ACARE) is promoting the development of novel systems to meet the requirements of VISION 2020 of emission per passenger kilometre $CO_2 < 50\%$, $NO_x < 80\%$ and noise $< 50\%$ [1]. Future aircraft will benefit of new solutions for fuel consumption reduction such as laminar flows, active flow control and low noise emission and in this path morphing structures have been investigated in many European projects, e.g. the New Aircraft Concepts Research (NACRE) project, the Smart High Lift Devices for Next Generation Wings (SADE), the Smart Intelligent Aircraft Structures (SARISTU) project, the Novel Air Vehicle Configurations (NOVEMOR) project and the CHANGE project.

1.1 Objectives of the present thesis

This thesis is intended to set up a proper procedure to optimise the aerodynamic response of a morphing leading edge device. Among several morphing concepts, this

latter kind of adaptation is interesting thanks to the advantages that can potentially show with respect to traditional high-lift systems. The discrete surface variation used to increment the lift during take-off and landing, though effective in this task, leads to a remarkable drag rise and has been identified as one of the main airframe noise source. A gapless and smooth variation of the airfoil shape with the employment of deformable skins can be helpful in eliminating this problem and enforcing a laminar flow at high angles of attack, reducing the drag and improving the aerodynamic efficiency. The objective of the thesis, therefore, is to identify an optimisation framework able to deal with a morphing leading edge geometry. Each functional block of the optimisation loop is analysed: the parameterization technique, the aerodynamic analysis tool and the optimisation engine. A suitable parameterization method should be able to allow for a local control of the shape and the computation of geometrical or structural quantities, so as to limit the deformation occurring with the morphing. In order to evaluate the aerodynamic performance of each candidate solution, validated aerodynamic tools are required and their influence on the optimal shapes provided is object of investigation. Moreover, the purpose of the study is also to check the ability of the optimisation algorithm to handle different problems, both single and multi-objective, in presence of non-linear constraints. After assessing the convergence at varying the objectives, the constraints and the boundary conditions, an interest is also present in inquiring the possible advantages that a meta-model assisted optimisation can have, with respect to a standard loop.

1.2 Summary of the thesis

In this work a morphing leading edge concept was analysed. The aim of the study was to set-up a procedure to efficiently parameterize a morphing shape, calculate its aerodynamic performance and optimise it using a genetic algorithm. In a first step, morphing systems were analysed referring to published researches, thus doing a literature review to understand what are the possible configurations available, their actual readiness, how to model them and simulate within an optimisation framework. After assessing the past studies on the topic, the attention focussed on the the aerodynamic calculation and in particular to the validation of a model, able to correctly replicate experimental data collected on wind tunnel tests for a given airfoil. Three formulations were individuated as the best behaving, a first one based on the panel solver XFOIL with the eN transition model, a second using Computational Fluid Dynamics (CFD) with the fully turbulent Spalart-Allmaras model and a third using the Transition SST or $\gamma - Re_\theta$ model.

With all these available, the following task was to individuate a parameterization technique suitable to deal with morphing. From a bibliographic research and by comparing several methods, the Class/Shape Transformation (CST) by Kulfan and Bussoletti was chosen, because of its ability to prescribe smooth changes over a partial region of an airfoil. In order to make it easier to deal with structural constraints, which appear to be the real issue of a morphing architecture, a dedicated procedure to enforce the

parameterization to produce a constant arc length, whenever was possible, was developed.

The optimisation framework built was based on a genetic algorithm as optimisation engine, in the standard form present in the Matlab® Global Optimisation Toolbox. This optimiser was linked with the parameterization block, employing the CST technique with constant arc length, and then the aerodynamic evaluation of an individual was carried out by one of the aforementioned tool. In particular, XFOIL was used in an initial survey, when the influence of the optimiser settings was widely analysed performing several runs with different initial populations and genetic operators. In a second phase, when the behaviour of the optimiser appeared sufficiently known for the actual case treated, ANSYS® Fluent was used for the aerodynamic analysis, with the software Gambit v. 2.4.6 to mesh the domain. This surveying activity was performed for both a single and a multi-objective problem.

Having individuated the set of parameters controlling the optimisation algorithm which was able to produce a smooth convergence, a variety of problems was faced, starting from single-objective and then switching to multi-objective. Specifically, a SOOP aiming at minimizing the C_d was solved with all the three aerodynamic models, comparing their results in terms of shapes and performance. Also three MOOPs were analysed, with different formulations employing either XFOIL or the CFD. One of them replicated a past activity aiming at optimising a multi-element high lift system based on slats and flaps, according to some airworthiness specifications, but adapted to a morphing leading edge airfoil.

The step following an extensive analysis of a two-dimensional geometry was the extension of the knowledge acquired to a three-dimensional case, involving a 3D wing, a portion of which could undergo morphing at its leading edge. The optimisation framework was adapted to this new situation, employing the same 2D parameterization method for some wing sections and Gambit with Fluent to perform the mesh generation and the CFD computation.

Finally, the exploitation of meta-models into the optimisation was investigated. For each aerodynamic model a regression method was built and the same single-objective optimisation as before was performed with two approaches: a complete fitness replacement for the XFOIL computation using an ANN and a meta-model assisted loop with an ANN too for the Transition SST model, in which at each generation the approximation was used in combination with a gradient based solver to locally refine the population.

Chapter 2

Morphing Systems: an Overview

2.1 Introduction

The path of evolution of the aeronautic science has led us from lighter-than air uncontrolled ships to an intensive use of air vehicles for civil and military purposes in a little over one century. Impressive results have been achieved in aeronautics, passing from wood and canvas first aeroplanes to supersonic composite planes of our days. The need for higher cruise speed and heavier payloads favoured the establishment of a certain paradigm in aircraft design, attaining maximum stiffness with lightweight structures to perform best in a certain regime of flight. This leads to an inevitably compromise, in which performance in various flight conditions is sub-optimal.

An attempt to match variable requirements in different phases is represented by morphing aircraft. As previously reported in Chapter 1, according to a general definitions morphing is the ability of manipulating certain characteristics of a vehicle to better match the vehicle's state to the environment and increase its performance [7, 81]. This definition indicates as morphing all the systems aiming at improving the performance and the efficiency by a modification of the vehicle, such us retractable gear and deployable slats and flaps. In the scientific literature, morphing is often referred to as the capability of a structure to adapt to the operational state by modifying its shape. In this sense there are many examples of morphing aircraft, as reported thoroughly in [7, 76, 81].

2.2 Morphing concepts

The scheme in Figure 2.1 illustrates the possible mechanisms of wing morphing [7]. Large modifications comprise variations in wing plane, like wing span morphing or sweep changing, and out of plane, like twisting or dihedral angle variation. Small changes include airfoil modification, like variation in camber or thickness distribution. Figure 2.2 reports an overview of built aircraft with morphing elements. Excluding ever-present traditional slats and flaps for camber variation, most known examples of morphing are the variable sweep wings of the F-14 Tomcat or the Panavia Tornado, the

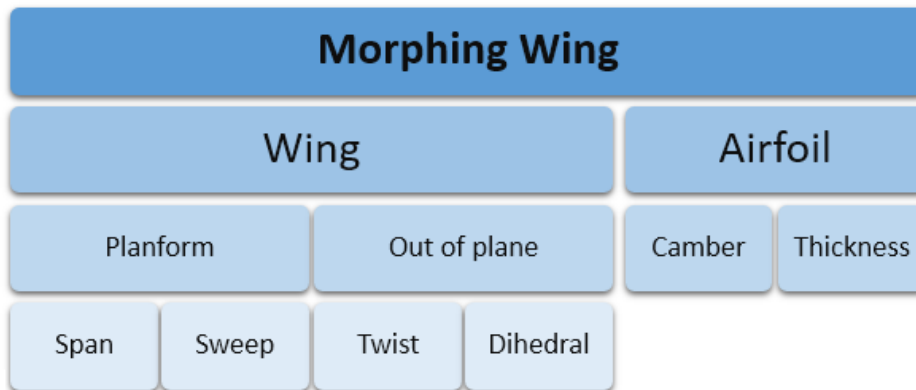


Figure 2.1: Possible morphing mechanisms



























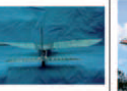








| | | | | | | |
|---|---|---|---|--|---|---|
| 1903 | 1931 | 1931 | 1932 | 1937 | 1947 | 1951 |
|  |  |  |  |  |  |  |
| Wright Flyer <i>Twist</i> | Pterodactyl IV <i>Sweep</i> | MAK-10 <i>Span</i> | IS-1 <i>Bi-to monoplane</i> | LIG-7 <i>Chord</i> | MAK-123 <i>Span</i> | X 5 <i>Sweep</i> |
| 1952 | 1964 | 1964 | 1966 | 1967 | 1967 | 1969 |
|  |  |  |  |  |  |  |
| XF10F <i>Sweep</i> | F 111 <i>Sweep</i> | XB 70 <i>Span bending</i> | Su 17 IG <i>Sweep</i> | MIG 23 <i>Sweep</i> | SU 24 <i>Sweep</i> | Tu 22 M <i>Sweep</i> |
| 1970 | 1972 | 1974 | 1974 | 1979 | 1981 | 1985 |
|  |  |  |  |  |  |  |
| F 14 <i>Sweep</i> | FS 29 <i>Span</i> | B 1 <i>Sweep</i> | Tornado <i>Sweep</i> | AD 1 <i>Obliquing</i> | Tu 160 <i>Sweep</i> | AFTI/F 111 <i>M.A.W.</i> |
| 1993 | 1994 | 2001 | 2002 | 2003 | 2004 | 2005 |
|  |  |  |  |  |  |  |
| FLYRT <i>Span</i> | MOTHRA <i>Camber</i> | AAL <i>Pitch</i> | F/A 18 A.A.W. | Virginia Tech <i>Span</i> | Univ. of Florida <i>Twist</i> | Univ. of Florida <i>Gull</i> |
| 2006 | 2006 | 2007 | 2007 | 2007 | 2008 | 2010 |
|  |  |  |  |  |  |  |
| MFX 1 <i>Sweep & Span</i> | Univ. of Florida <i>Sweep</i> | Virginia Tech <i>Camber</i> | Univ. of Florida <i>Folding</i> | MFX 2 <i>Sweep & span</i> | Delft Univ. <i>Sweep</i> | Virginia tech <i>Camber</i> |

Figure 2.2: An overview of aircraft with morphing devices [7]

last one still operative. Referring to this morphing mechanism, it can be pointed out that its employment has had a limited spread and more recent military aircraft achieve equivalent or better performance at subsonic and supersonic regimes even without variable sweep wings. This example shows in a qualitative way how the use of morphing devices needs to be carefully assessed. The adoption of morphing technology should be demonstrated to bring to performance benefits, considering not only the aerodynamic point of view, but the entire system in terms of weight estimation for both structures and actuators and energy requirement. Establishing whether morphing could be a valid solution depends also on the mission profile and given the variety of mission type and regime, such wide is the possibility in morphing strategy. Historically, morphing solutions have met with an unfavourable cost/benefits ratio, due to an increase in weight, complexity, cost and failure probability [81].

With current efforts to produce green and efficient aircraft, novel technologies such as morphing become interesting in that they can overcome the drawbacks of using standard rigid configuration with deployable attachments, to discretely vary the wing section shape for additional lift or flight control. Moreover, morphing becomes more attractive also for the new class of Unmanned Aerial Vehicles (UAVs). Adaptive wings can allow to perform multi role missions expanding the flight envelope of a single aircraft, which can remove the need for various vehicles. In order to match different requirements, large shape variations in wings are required. Planform morphing can improve the aircraft performance over a wide range of conditions and its benefits are more apparent for UAV type missions, with several conditions of flight speed and payloads [81]. Large spans allow an extended range and low drag, but limit manoeuvrability and cruise speed. A variable span could guarantee high efficiency during cruise and good manoeuvrability in targeting, while asymmetrical span can be used for roll control in MALE UAVs [2]. Most span morphing concepts are based on a telescopic mechanism with sliding skins, like the Adaptive Aspect Ratio AdAR wing by Swansea University [86]. Other concepts of wing morphing involves twist and dihedral angle variations. A computational and experimental study of a wing with two outboard morphing partitions that are variable in twist and dihedral angles is presented in [69]. An extensive review of research in morphing systems can be found in [7].

Adapting the wing shape to flight conditions requires several aspects to be examined. Deformable structures allow to match the desired configuration without hinges and sharp edges. Current high lift systems inducing discrete camber morphing are responsible for a penalty in performance, due to local source of turbulence which causes noise and drag increase [66]. The adoption of morphing airfoils capable of modifying their camber requires adaptive structures with highly deformable skins. These should be able not only deform in the correct way without plasticity, exhibiting high deformability in the deployed configuration, but also have to bear aerodynamic loading without any significant aeroelastic behaviour and provide stiffness to span bending. The actuation device can be the penalizing subsystem in morphing structures for the weight impact on the system. The optimisation of the mechanism to minimize the actuation energy can lower the power requirement and thus the size and the weight of the actuator,

but the mechanism must provide sufficient locking when folded, in order to guarantee a minimum deformation under pressure loads. Smart materials with embedded or distributed actuation can limit the additional weight with respect to traditional systems. Fitting morphing devices to existing wings can shrink the potential of the technology, giving results largely dependent on the system considered. According to [3], in order to extend the design space of morphing devices and entirely exploit their benefits a new design philosophy should be applied, including the device in the first step of the analysis, to remove existing constraints and enlarge the optimisation space.

2.3 Morphing airfoils

Since this work deals with a morphing airfoil, a general sight into this topic is here proposed. As previously stated, morphing can involve a partial or total change in camber or a variation in thickness distribution. Camber represents the effective curvature of an airfoil and its control has been performed from the origin of flight. Common aerodynamic surfaces like slats and flaps are a way to vary the camber in low speed phases of take-off and landing, where additional lift is required. Flaps are more spread because changing the trailing edge results easier, since it deals with deflecting a thin surface with no other roles except control. Conversely, leading edge slats must stand higher aerodynamics loads and potential impacts, confer stiffness to the wing, can host anti-icing devices based on fluids. Discrete surfaces are however used in almost all vehicles and slats are more spread in transportation aircraft. Research showed that a gapless and seamless continuous control of camber can lead to high L/D ratios in subsonic regimes [7]. Compliant structures with total camber variation are more suitable for small or medium aircraft like UAVs, since in other cases wings have not only aerodynamic functions and host various other accessory subsystems.

Figure 2.3 shows four different concepts of morphing airfoil. Figure 2.3a depicts a variable camber compliant wing (VCCW) by the Us Air Force Research Laboratory described in [35]. Up to 6 % camber modification along span under aerodynamic load was demonstrated using a bench top model. Variable camber wing allows to maximize fuel reduction by adapting wing contour to a specific flight condition or configuration, obtaining higher efficiency with respect to conventional flapped wing due to a lower drag. The fundamental design requirement for the compliant mechanism was to use a single actuation force for both the leading and trailing edge to deform the camber by 6%, modifying the initial NACA 2410 profile to a NACA 8410. The wing deformation was produced using a series of miniature electronic actuators and compliant mechanisms, which acted independently to produce either uniform or nonuniform camber variation along the span.

Figure 2.3b illustrates the final prototype realized of a morphing camber wing with a rigid leading edge and a compliant mechanism with corrugated skin for trailing edge deformation [51]. The model was the result of a multidisciplinary optimisation activity and was based on distributed compliance and smart piezoelectric actuators. It was tested on wind tunnel, to validate the design tools and assess aero-servo-elastic effects,

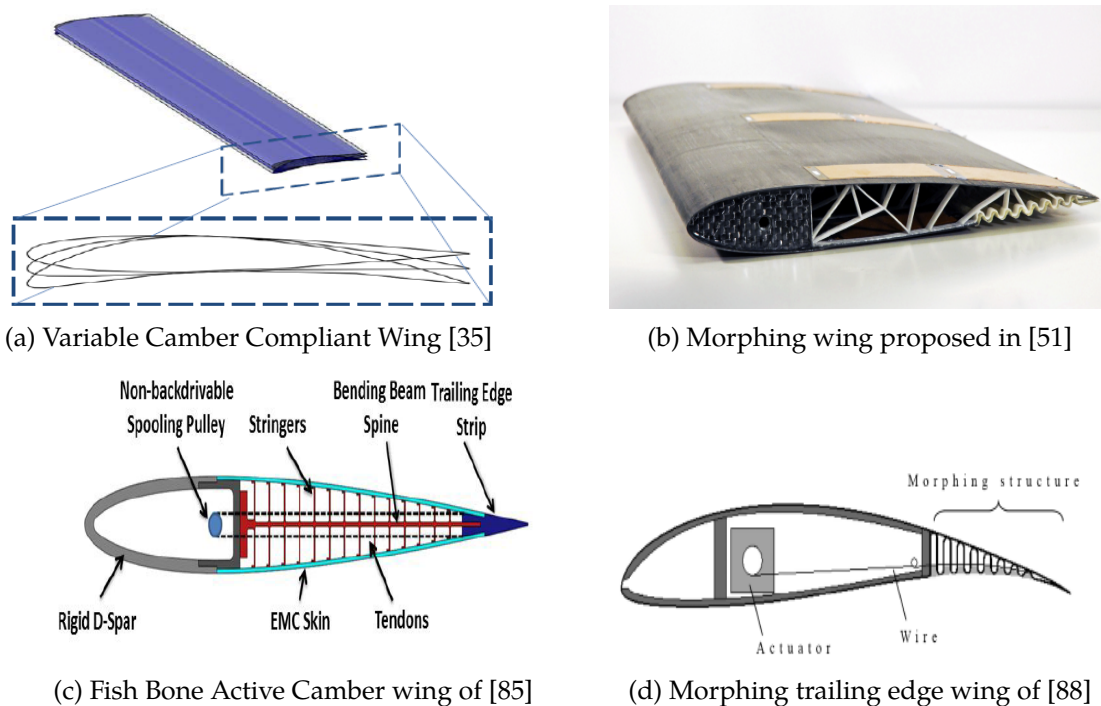


Figure 2.3: Four examples of airfoil morphing in literature

and also in flight tests, showing a potential application for small and medium UAVs.

Figure 2.3c reports the FishBAC concept of Woods and Friswell [85]. The Fish Bone Active Camber (FishBAC) morphing airfoil was an alternative morphing camber design, biologically inspired to achieve large and continuous changes of the profile. The structure consisted of a thin chordwise bending beam spine with stringers branching off to connect it to a pre-tensioned elastomeric matrix composite skin surface. The Core and the skin were designed to have almost zero Poisson's ratio in spanwise direction and their deformation was controlled via an antagonistic tendon system with a non backdrivable pulley. The tendon system allowed a compliant deformation of the structure and pretensioning provided out of plane stiffness and eliminated buckling in the lower skin when morphing. Locking was guaranteed by the non backdrivable pulley without energy input, even for chordwise bending stiffness under aerodynamic loads. The system was prototyped and wind tunnel tested, showing an increase in L/D ratio of 20-25 % in the same regimes of use of conventional flapped airfoils [83].

Figure 2.3d shows the morphing trailing edge concept developed in [88]. It was based on corrugated structures, which have high load bearing capability in the direction perpendicular to corrugation, whereas flexible deformation is possible in the corrugation direction. The deflection was controlled by wire tension and was possible only in downward direction, imitating flaps. Due to the high anisotropy of the structure, a

large deformation could be achieved with the use of only one actuator, like in the Fish-BAC concept. Wind tunnel tests of a prototype proved the effectiveness of the morphing system and a comparison with a standard reference wing with discrete flap showed an extended linearity of C_l .

Another option available is a camber modification by a morphing leading edge. Since it is the topic of this work, it will be treated separately in the next chapter.

Chapter 3

Morphing Leading Edge Devices

3.1 Introduction to morphing leading edges

As previously seen in chapter 2, camber variation with compliant structures can be used to continuously adapt the profile geometry in order to optimise the performance among different flight conditions and various options are available to realize such an adaptation. Morphing leading edge concepts occupy an important area of research in that they could help overcoming some of the drawbacks of standard high lift slat systems.

High lift systems have been investigated starting from the first quarter of the 20th century [32]. The evolution in aerodynamics and engine performance has progressively led to faster cruise speed, representing the reference condition for the design, which is thus optimised for high speed. In take off and landing the performance of the vehicle can be penalised, requiring longer runways. In order to reduce the take off and landing speed, higher lift coefficients are required and trailing edge flaps helped to achieve this, while leading edge slats were found to delay the stall at higher angles of attack. In the second half of the last century, the development in high lift systems led to very complex solutions with double or triple slotted flaps, like in the Boeing 747. With the enlargement of commercial aviation, new environmental issues like air pollution and noise rose, inducing a new attention towards environment respect.

Extensive research in engine technology led to the actual state-of-the-art, considerably reducing the impact of the propulsive system in air pollution and noise. In recent years the adoption of high by-pass ratio engines has effectively limited the noise generation [66], focusing on other sources of disturbance. Leading edge slats and deployed landing gear have been identified as the main source of noise during approach [13]. As reported in [58], there are basically two methods for noise reduction: the first is to increase the distance between the noise source and the observer, the second is reducing the acoustic emission of the source. As regards this second option, an aeroacoustic driven design of high lift systems can collide with the aerodynamic performance required, leading to a penalisation of the system also in terms of complexity or weight. Research in this topic was performed within the DLR projects LEISA [32, 82] and SLED

[58], in which a single slotted slat optimisation showed a noise reduction up to 4 dB for a 3D configuration. A comparison of conventional and morphing high lift devices in noise terms can be found in [66], where numerical investigations were performed to evaluate potential re-arrangements of slats with various types of fillers and also a droop nose.

Traditional discrete leading edge systems have two main drawbacks. The first one regards the deployed configuration, with the slat forming a gap with the wing remainder, which is source of noise and turbulence, inducing turbulent boundary layer and drag increase [38]. The second problem occurs during the retracted configuration, due to the presence of a dent in the location where the slat meets the remainder of the wing, which still induces transition to turbulent boundary layer. The risen sensitivity towards environmental issues, which leads to recommendations for civil aviation, like those of ACARE [1], induces the need for novel solutions to efficiency improvement and noise impact. Although modification of the current technology of slats could reduce their noise by re-arranging the slat gap and increasing the chord, the slots remaining in the airfoil do not comply with high quality surfaces required for future aircraft. For the purpose of drag reduction, most experts agree that laminarization is the only technology with potential for step changes in drag reduction within a suitable timeframe [38].

In this context, seamless and gapless morphing leading edge devices represent a viable solution to overcome the drawbacks of actual high lift systems. Due to their promising performance increase and environmental benefits, these devices received a fair amount of research in the past and are currently being investigated in various projects around the world. Future wings will probably benefit of higher aerodynamic efficiency with high aspect ratio, slim profiles and smooth surface for laminar flow, like the FNG wing produced by German ProHMS project, cited in [33] and used in various other projects. The reduced space in the leading edge and the need for high quality surfaces for laminar flow indicate in seamless compliant structures a valid alternative to be investigated.

A morphing leading edge requires a flexible skin which can adapt to the desired shape and an actuation system which drives the skin into the correct position. Typical technological challenges in realizing a morphing structure are: [67]

- Highly elastic skin capable of assuming the desired configuration without degradation and withstanding air loads;
- Appropriate mechanism to deform the skin while transferring loads from it to the carrying structure;
- Connection of mechanism to the elastic skin;
- Actuation systems integration in lumped or distributed configuration.

Designing a morphing leading edge is a multidisciplinary activity involving various phases. The aerodynamic optimisation of the drooped nose should take into account not only the aerodynamic indexes, but also structural constraints on the deformation


capability of the skin and the reversibility of the process. Various approaches are available for the design and appropriate tools for optimisation were developed in the context of European projects and academic research. After a suitable aerodynamic shape has been individuated, a structural optimisation of the skin has to be done, in order to guarantee that the desired configuration can be reached without plasticity and with minimum actuation energy. A third phase in the design process is the choice of the mechanism realizing the deformation, which can be a kinematic chain based on rigid members pivoted together and hinged to the structure, or a compliant mechanism made up of a single deformable structure without linkages. The optimisation of the actuation system can be performed with different strategies and constraints and locking capability and integration issues must be considered in a post process phase.

3.2 Morphing leading edge concepts


Even though research interest in smart Leading Edge (LE) devices has risen in last two decades, conceptual solutions have been proposed and patented in last century, inspiring future works of more recent times. [52] reports four patents of variable camber airfoil from 1916 to 1980. One of the major problems in deformable airfoils is the design of the skin. A downward rotation of the nose, indeed, causes the upper skin to stretch while the lower to compress, a requirement contrasting the characteristic of an elastic skin. Attempts to overcome this problem in the past led to slotted configuration, with the lower part of the skin consisting of a fixed part in the aft, interfering with the rotating nose skin, like visible in the Boeing patent of 1982 [14] in Figure 3.1a. Although this solution solves the problem of the skin in the lower part, it collides with the need for a high quality surface for laminar flow. A redistribution of the skin material at the leading edge is possible, for example in the form of bulging of the nose, increasing the tip radius. This is a natural reaction of the structure, seeking for a minimum energy deformation, and also a good strategy to control the pressure suction peak at high angles of attack [31].

A well-known example of gapless and seamless mechanism is the Dornier patent of 1979 in Figure 3.1b [91]. The concept consisted of three flexible panels and an internal framework with discrete struts and force introduction points, for transmitting the actuator force into the flexible panels, for shape adaptation and the aerodynamic loads into the front spar. It represented a quite advanced formulation, since it presented a continuous end-to-end skin, led in the deformation by guiding profiles moved by a single lever-arm actuated by a rotary actuator. The thin skin in the tip transforms the compression load in the lower panel to an increase in tip radius. Nevertheless, the presence of rigid parts (9 and 10 in the figure) and the guided deformation induce high stresses and strains in the skin, which instead should be designed to adapt naturally to the desired configuration without forcing it.

Airbus developed a gapless slat for the high lift system of A350 XWB-900, based on a rigid slat hinged at the bottom of the nose and able to rotate circularly, keeping a contact with the upper part, as illustrated in Figure 3.2. This design, adopted for the inboard



(a) US patent 4553722, Boeing Company, 1982 [14]



(b) patent DE 2907912 A1, Dornier Company, 1979 [91]

Figure 3.1: Two different patents for droop nose devices

wing, was found to be a good choice from an aerodynamic and integration point of view and was driven by the requirement of attaining a maximum L/D ratio in the second segment climb at take-off, providing sufficient protection against flow separation also in landing [72]. This solution is implemented in A380 too and represents a first step towards operational employment of smart systems overcoming slotted slats drawbacks, though still not based on morphing.




Figure 3.2: Airbus gapless high lift leading edge system of A350 [72]

As regards the choice of the mechanism, various solutions were proposed, based both on rigid kinematic chains like those on Figure 3.1 and compliant mechanisms with no linkages or pivots. The former are classical mechanisms based on rigid members, linked together in order to move describing a certain path, connected to the skin by appropriate joints to induce the desired deformation. An example of such a device is reported in Figure 3.3a, illustrating the LeaTop prototype which makes use of a conventional kinematic chain actuated by a linear actuator. The latter are a different class of mechanisms based on compliance. According to [6], a compliant mechanism is a single-piece structure designed to transmit motion and force mechanically relying solely upon elastic deformation of its constituent elements. This kind of mechanism has to comply between sufficient stiffness to support external loads and enough flexibility to deform

under acceptable actuating inputs. The connection between the mechanism and the skin can be realized by stringers, in a way that a continuous deformation towards the correct configuration is achieved with minimal errors. A save in weight and complexity is obtainable through compliance [60], and an example of a morphing leading edge with a compliant mechanism found in [79] is reported in Figure 3.3b.




Figure 3.3: Two different mechanisms for morphing leading edges

3.3 Morphing leading edge design

In this section the design process of morphing leading edge devices will be analysed basing on a literature review of past and current research. Aeronautical morphing structures design is a multidisciplinary problem, which involves basically three major areas:

- Aerodynamic design of the drooped configuration;
- Structural design and optimisation of the skin;
- Design of the actuating mechanism.

Each of these areas can be treated separately in the design process, starting from the aerodynamic optimisation which produces a certain shape and then optimising the thickness distribution of the skin for that shape and the mechanism, or they can be considered in contemporary, identifying the relationship among them and including in the optimisation some constraint, deriving from those relationships. In a literature survey, the second approach appeared to belong to well structured European or national projects aiming at developing a morphing prototype, while the first appeared to be more related to independent academic research on specific topics of morphing, where an integral vision of the problem is not available.

3.3.1 Aerodynamic design

One of the main reasons about employing morphing wings is the potential aerodynamic performance increase, due to the capability of adapting to changing external

conditions and providing more than one optimal operative points. The aerodynamic design of the gapless droop nose device must therefore be led by aerodynamic parameters. Typical objectives of the design are the maximization of the L/D ratio at constant C_l [27, 39, 85] or with C_l as a constraint in the optimisation [19, 73] and also maximize C_l at a constant condition [45]. While in most works the aerodynamic shape is found via an optimisation process, in other cases it is found parametrizing the airfoil and seeking for the curve which satisfies certain conditions, such as a smooth link with the fixed part of the profile and a constant girth [74] or the downward deflection on the tip radius in the drooped configuration [53].

Another issue related to the aerodynamic design is including structural constraints in the optimisation loop. A more traditional approach considers the aerodynamics separately from the structure, producing a target shape optimised for some flight conditions and then a morphing system is produced that can deform the structure to match that shape. The accuracy of the result will depend on the number of degrees of freedom of the mechanism, which will match in an approximate way the target if too simple, but could become excessively complex and heavy if stricter requirements to the target shape are required, vanishing the benefit of an increased efficiency. In [27], for example, a more integral approach was considered in which rather than performing the aerodynamic optimisation and then designing the morphing system to obtain the required shape, the morphing system was included in the loop, passing from a multiple single-objective optimisation to a single multi-objective optimisation and obtaining a reduction in the Pareto front. A similar philosophy appeared in [17, 19], where a two level optimisation was used firstly to get the best aerodynamic shape with constraints on skin stress and energy deformation, and secondly to design the compliant mechanism to deform the airfoil via topology optimisation.

As regards the parametrization technique, several methods were found in literature. In [74] a fourth order polynomial was used for the camber line expression and the thickness. Ten coefficients needed were determined on the basis of some assumptions, such as a constant tip radius and girth and smoothness of the link with the fixed part of the airfoil. [73] used a linear superposition of a baseline function with type functions based on the Hicks-Henne's. In [25] the CST method was chosen to parametrize the entire morphing airfoil, while for the same aim Radial Basis Functions are used in [27]. CST were also used in [17, 19] where leading and trailing edges were morphing. In this case the parametrization is based on four terms: a shape function, a class function and two additional terms related to the aerofoil leading edge and trailing edge shapes. The calculation of first and second order derivatives is very fast with this technique and can be used to compute geometrical quantities, useful also for structural modelling.

Finally, Genetic Algorithms appear quite spread as optimisation method and were used e.g. in [19, 27, 73]. SUPLEX method is found instead in [39], where no structural constraint is reported during the aerodynamic analysis.

3.3.2 Structural design

Structural design is the second phase of the design process of a morphing device. Even though, as previously seen, structural constraints may already appear in the aerodynamic optimisation, it is here assumed that a certain target shape has been found and is known. A morphing structure must exhibit conflicting properties: it should be able to reach the deformed configuration with minimum actuating force and without damaging, while it should also support aerodynamic loads with no relevant aeroelastic deformation and confer torsional and bending stiffness to the wing, in this case. These requirements pose serious questions about what structural configurations can satisfy them, what materials are eligible for the task, what kind of actuation is needed and what it is a suitable design process. It is evident that these questions have to be answered in a global approach: the structure arrangement varies with the material chosen and different sizing and optimisation techniques are required for different classes of materials. The interaction between the actuation system and the morphing structure can change the structural behaviour if forcing the deformation, so it should also be considered in the loop. Moreover, the aeroelastic frame cannot be neglected and aerodynamic loads appear in the structural analysis varying either the deformed shape or the actuation force required. From these qualitative considerations, it appears how this second phase is the core of the design chain; while the aerodynamic optimisation can always lead to an optimal shape for given conditions, it is not direct that shape is achievable with a certain structure. The high number of degrees of freedom in the choice of materials, structural arrangements and actuation makes this phase of the design critical in that a result can be obtained that vanishes the benefits of the morphing technology due to an increased weight and complexity.

A review of morphing skins can be found in [75]. Although several concepts were analysed in literature for adaptive structures in general, in the case of morphing leading edge, realized prototypes show a preference for composite materials, due to their potential high anisotropy which is required for the task. Since stiffness, elasticity and strength are tensor properties in composites, it is possible to combine the required properties in the need direction obtaining selective deformability. Selectively Deformable Structures (SDS) were investigated for application in the European project SADE, as reported in [4]. Connecting properly elementary cells with high anisotropic behaviour, it is possible to arrange a structure exhibiting high deformability and strain capacity in one plane, while stiffness and strength in other directions. Two configurations were proposed, one with the SDS filled with an elastomer to create a quite smooth surface and support the cells, and one with the SDS arrangement covered by a flexible skin. This concept is similar to those reported in [59], where a honeycomb was filled with polyurethane elastomer and covered with flexible skins. The solution, indeed, seems more suitable for trailing edge devices, where the structure can occupy the entire space and there are less problems of integration and supplementary resistance needed for ice and impacts. In the SADE European project, in fact, a SDS was chosen for the trailing edge, while a glass fibre prepreg tailored skin was optimised at the leading edge [39]. The structural design process reported was based on a pre-sizing of the skin and an optimisation to determine the thick-

ness distribution, the number and position of the support stations along the skin, the kinematical hinge points and actuation angles at the support stations. A first thickness distributions was chosen on the basis of the strain derived from the curvature difference between the baseline and drooped configuration. Constraints for manufacturing like minimum skin thickness and distance for the stacking sequence of layers were applied. Once the optimisation, using a SIMPLEX scheme, had converged, a simplified kinematic mechanism was designed basing on the position of the supports station and an assumption of a circular motion of the hinge points. After completing the design of both skin and actuator, a complete Finite Element Method (FEM) model was used to assess the global behaviour also in presence of aerodynamic loads. The project ended with a wind tunnel test campaign of a 5 m span prototype [37], like that in Figure 3.4. In the same project SADE other concepts for structural design were investigated. In [54]




Figure 3.4: SADE Smart Leading Edge Concept for the ground test [37]

the same leading edge was designed with a glass fiber composite skin, reinforced by eight I-shape stringers on both upper and lower part. An eccentric beam was proposed as actuation system and is visible in Figure 3.5. In [76], instead, a variable stiffness com-




Figure 3.5: Eccentric beam actuator in neutral and deployed position [53]

posite skin was investigated within the same European project. Traditional laminate composites are made up of a stacking sequence of unidirectional laminae, in which fibres angles are the same within the same lamina. A spatial stiffness variation can be

achieved changing the fibres orientation into a single lamina, thus enlarging the design space and exploiting better the huge potential of composites, which allow for a large number of degrees of freedom. Three possibilities were accounted for stiffness variation: changing the fibre angle with a thickness variation, changing the fibre angle without a thickness variation and changing thickness keeping a constant fibre angle. The skin parametrization by lamination parameters allowed to consider these three options and in order to reduce the design variables a symmetric layup was chosen. Skin stiffness optimisation was made in an aeroelastic framework, in which aerodynamic loads were computed with an inviscid panel method and the fluid-structure static interaction was included in the loop. The actuation system was then optimised through topology optimisation and several solutions were proposed. The peculiarity of this work is that static aeroelasticity is present in the optimisation process, where convergence is sought between deforming pressure loads and internal elastic forces of the structure.

Interesting observation about applicative problems for morphing skins are reported in [30], in which additional functions required to the skin in an operative environment, such as impact resistance, de-icing, lightning protection and abrasion shielding, are examined. The starting point was the morphing leading edge developed by the DLR in the European project SARISTU. In order to avoid plastic bending of the skin, its thickness needs to be as low as possible. There again, the aforementioned functions required for airworthiness could require additional layers in the outer part, possibly leading to exceed the strain capability. For this reason a hybrid laminate was investigated, with a thin elastomeric layer achieving all the functions integrated in the GFRP laminate. Experimental tests on erosion-shielding and de-icing showed a 100% thick increase when the required functions are integrated as additional layers, with a tremendous impact on the strain in the outer surface. These results demonstrate how further steps in the morphing structures have to be done, before reaching the level needed for applicative implementations of the technology, in which accessory aspects can lead to reformulate a whole assembly concept.

3.3.3 Actuation Mechanism design

After completing the aero-structural design, a proper actuation mechanism to drive the structure in the desired deformed configuration is to be determined. As previously seen in section 3.2, actuation systems can be loosely divided into two classes: kinematic chains and compliant mechanisms, an example of which can be found in Figure 3.3. The former typology is based on rigid members pivoted together and linked to the skin, with the task of transmitting forces or displacements from an actuation input point to output points, thanks to the degrees of freedom given by hinges. In this case members are modelled in first assumption like rigid bodies and their deformation is undesirable. Compliant mechanisms, instead, exploit structural flexibility to transmit forces or deliver motion and are monolithic and jointless structures in which deformation is necessary, though a certain stiffness is also required to support external loads. Some advantages of compliance are a reduced number of manufactured parts, the absence of

linkages associated to high stress and possible misalignments, the lower weight compared to kinematic chains, the absence of relative moving parts needing lubrication.

A spread tool for designing actuation mechanisms is Topology Optimisation (TO). The basic idea of a topology optimisation is to find the optimal region where material should be present to achieve a certain target objective. Examples of use of this tool to design compliant mechanisms for leading edge devices are [62, 73, 79, 80, 87], while conventional kinematic chains for the same purpose are reported in [39, 40, 52, 70, 76]. A systematic analysis of optimisation tools and processes for compliant mechanisms for morphing structures can be found in [46].

The objective of the optimisation is basically the minimization of the displacement error between the deformed configuration driven by the actuation system and the target configuration, often expressed in term of Least Square Error (LSE), with some structural constraints for example in the strain level. In [87] TO with the Simple Isotropic Material with Penalization (SIMP) method was used considering different fiber ply-orientation for the mechanism. SIMP method varies the local density of the discrete element of the design space, leading to a material distribution with denser area where the mechanism should be and areas of void. A drawback of this method is the necessity of a post-processing analysis because of “grey” areas. Filtering techniques are available to lead to a more solid-void solution but an engineer’s view is often required to translate the result into a feasible structure. [62] reports a comparison between two TO method, specifically load path and density based SIMP. The first one is based on an initial parent lattice of beam members to be determined a priori as a choice variable. The algorithm evaluates the load paths made up of a connection of the beams which lead to the convergence criteria on the objective function. Even though this method reduces the need for post processing analysis, the solution is depending on the initial parent lattice grid. In the reference cited a method to reduce the parameters for the description of the network of beams was proposed, because an high number of nodes necessary to better match the objective can lead to excessive degrees of freedom. A similar comparison between the two aforementioned methods can be found in [73].

In the work reported in [79], as part of the European project NOVEMOR, a wingtip morphing leading edge was designed following a three steps process: 1) the design of the skin via Simplex based optimization; 2) the design of the compliant mechanism via continuum gradient-based topology optimization and 3) the design of the support for the compliant mechanism also via continuum gradient-based topology optimization. The material selected for the compliant mechanism was aluminium 7075 alloy with a thickness of 5 mm, while a linear actuator was chosen. LSE minimisation with constraints on the amount of material used was performed within a SIMP model for TO and the MMA method as update scheme. The result of the mechanism optimisation was then used to determine the boundary conditions for the support, which in turn was optimised solving a minimum mean compliance problem. The final prototype deriving from this work is shown in Figure 3.3b.

TO was used in an aeroelastic framework in [76], where several solutions for the morphing leading edge of the SADE project were proposed. The concept implemented in

the prototype for ground tests was based on a kinematic chain made up of a main lever from which four branches pivoted on it depart and are hinged to four omega stringers reinforcing the skin. A bell crank lever was used for actuation. After the skin optimisation, the kinematic chain was optimised varying the positions of the levers with their bearing support, in order to minimize the actuation energy. The resulting system was finally incorporated into the global FEM model of both skin and actuator to verify the operational behaviour under aerodynamic loads. Results of displacements test also in wind tunnel can be found in [37, 39], while a figure of the prototype with the kinematic system visible is Figure 3.4. Another concept evaluated during the project is an eccentric beam actuator (EBAM) [54], already proposed in Figure 3.5. The system was based on an aluminium beam with five disks attached to the skin, shaped in order to convert a rotational motion of the beam in a vertical displacement of the skin.




Figure 3.6: SARISTU drive chain actuation mechanism [40]

Starting from the results of the SADE project, a similar mechanism was designed for the SARISTU project within the activity stream “AS01 Enhanced adaptive droop nose for a morphing wing” under the lead of DLR [40]. In this case a 3D wing section was designed and the aim of the work was considering also integration problems like additional protection against impacts, ice and lightnings, manufacturing and industrialization. Dealing with a 3D wing section, the kinematics had to be optimised to reach different shapes along the wing span, which actually are only slightly differing from each other. The final choice was a drive chain in which the linear motion of the actuator was translated via two hinges and one lever into a rotary motion. A single linear actuator was used to move a modular drive chain, to which at each wing station a main lever was connected by a cross link. In this way a low weight solution was achieved, with only one actuation point and a distributed system of hinges to convert the main bar translation into the rotation necessary for drooping the leading edge. Figure 3.6 illustrates the concept presented.

Another European project aimed at designing, manufacturing and testing an actuation system that could deform a morphing leading edge seamlessly was the LeaTop project [16]. The actuation system designed via TO was a statically determinate kine-

matic chain of rigid members which drove a thin aluminium skin into the desired deformed configuration. A proper design of the actuation system was required to make sure that strains in the skin are low and that aerodynamic loads are transmitted to the front spar of the wing through it, and not through the skin itself to avoid excessive strain in the region near the fixed part of the airfoil. Experimental test conducted proved the ability of the system to reach the desired deformation while providing sufficient stiffness to support aerodynamic loads and the locking capability [70]. A technology readiness level (TRL) of 3 was achieved and the system has been already shown in Figure 3.3a.

Chapter 4

Aerodynamic Optimisation of a Morphing Wing

As explained in previous sections, morphing airfoils are reported to be an interesting technology to improve aircraft performance, particularly in take-off and landing phases, where conventional discrete control surfaces are cause of a drag penalty and one of the main airframe noise source. In the present thesis an aerodynamic optimisation of morphing airfoil is performed. The current section presents the problem and gives an overview on the structure of the thesis.

4.1 Problem statement

The problem faced in this work is part of a more generic procedure aimed at designing a morphing wing. The nature of the problem requires a global perspective to be pursued in order to guarantee that the technology considered leads to effective advantages to the system it belongs to. The design is mainly led by the aerodynamic behaviour of the device, but without considering its systemic role, the solution can be impracticable to realize and integrate within the main system.

The presence of different requirements and trends when including considerations of various areas, such as aerodynamics, aeroacoustic, structural mechanics, integration and so on, suggests the need for an optimum design which, taking into account the main objective and the principal constraints, can effectively represent the best compromise against conflicting requirements. For example, considering simply the aerodynamic, the optimal morphed leading edge could involve such a shape change to be achievable only with a very strict structural optimisation, a large amount of material and a robust mechanism which, in turn, vanishes the efficiency improvement of the aerodynamic due to an increased complexity and weight. Starting from this basic consideration, the design procedure of the morphing leading edge appears to be preferably optimisation led, in order to take into account of the aforementioned problems. Of course the degree of complexity of the model used in the optimisation framework is limited and only some

constraints can be invoked; many other aspects will be assessed in successive iterations of the design.

The purpose of this work was to set the first level optimisation chain which can produce an optimal aerodynamic shape, subdued to aerodynamic, geometric and structural constraints. The aim of the study was to identify a suitable parameterization method, able to handle the morphing geometry, assess the influence of the aerodynamic model on the optimal shape, analyse the influence of the optimisation strategy and propose a framework able to produce optimised shapes either for single or multi-objective formulations. Several cases were analysed, with an increasing level of complexity. In a first phase, a two-dimensional geometry was considered, in order to evaluate the performance of the optimisation under different conditions and to set-up properly the optimisation loop. In a second phase, the work focused on a more complex three-dimensional shape, extending the previous tools to optimise a morphing leading edge wing. The final phase dealt with possible improvements of the optimisation, by employing more sophisticated algorithms and including meta-modelling as an optimisation strategy.

4.1.1 Reference configurations

In the initial phase of the study, a 2D geometry was chosen in order to set up the optimisation loop and test its subsystems. The geometry consisted in a general aviation airfoil designed by NASA, the GA(W)-1, for which experimental aerodynamic characteristics are available in literature. The airfoil is depicted in Figure 4.1, in which the morphing part is indicated in green and extends from 0 to 25% of the chord length. For the 2D optimisation, the reference conditions considered are reported in next table 4.1. Further specifications regarding particular problems treated will be reported in the proper section.




Figure 4.1: NASA GA(W)-1 general aviation purpose aerofoil

After assessing the optimisation in a 2D fashion, the framework was extended to deal with a 3D case. A portion of a reference wing was optimised. The geometry is depicted in Figure 4.2. Being c the root chord, the wing had a wingspan of $2c$ and a chord of $c/2$. The leading edge of the tip section was at the same abscissa of the trailing edge of the root section, providing the wing a sweep angle of 26.565 [deg] for the leading

| Name | Reynolds | Mach | AoA [deg] |
|------|----------|------|-----------|
| DP1 | 6E+06 | 0.20 | 15 |
| DP2 | 6E+06 | 0.20 | 5 |

Table 4.1: DPs considered

edge. Also a sweep was introduced from the hub to the tip for at most 1 [deg]. The three-dimensional component of the flow field around the body appeared relevant, due to the low aspect ratio, the sweep, the taper and the twist.




Figure 4.2: Three-dimensional view of the reference wing

4.2 Standard optimisation framework

The basic optimisation framework employed consisted in a standard optimisation loop, written and executed in a Matlab environment. The chain was composed by three main blocks: the geometry block, the aerodynamic block, the constraint evaluation block and the optimisation engine block. Figure 4.3 synthetically illustrates a scheme of the framework. The geometry block, surrounded by a dashed purple line in the figure, allowed to generate the geometry, using the information derived from the optimisation engine. It communicated with both the second and third block providing a parametric constant arc length geometry representation. The aerodynamic block, represented with a dashed red line, received as input the geometry and computed the aerodynamic indexes for the given boundary conditions. The meshing software Gambit generated the grid, while an high fidelity CFD commercial software, ANSYS Fluent, was used to

solve the flow field. The information computed was blended in the constraint evaluation blocks, surrounded by a green dashed line, in which both geometric and aerodynamic constraints were calculated, producing a penalised fitness value. This entered the optimisation engine block, orange in the figure, which manipulated the decision variables and tried to improve the solution. The algorithm employed was a Genetic Algorithm (GA). The loop was then repeated, unless some stopping criteria are satisfied. During all the computations, in order to save time and polling the design space, a database containing the decision variables already investigated and the related fitness was created. This allowed to avoid to launch the CFD code if a design vector had been already computed and was useful for a further approximation of the fitness landscape.

In the initial phase of the work, a panel solver was used instead of a CFD tool, for the great benefit in terms of computational cost. This will be discussed more in detail in chapter 6, but the reader should consider to replace the two sub-blocks in the figure, inside the aerodynamic red-dashed block, with a single subsystem consisting in the panel solver XFOIL.

4.3 Synthetic summary

A brief summary of next chapters is here reported:

- **Chapter 5: Shape Parametrisation:** an overview of shape parametrisation techniques for 2D shapes is reported; the employed CST method by Kulfan is explained and its performance in inverse fitting and airfoil morphing tested;
- **Chapter 6: Aerodynamic analysis tool validation:** the tools used for aerodynamic analysis are presented and their results compared with experimental data from a test case, in order to validate the model;
- **Chapter 7: Two-dimensional Optimisation of a Morphing Leading Edge Airfoil:** the optimisation procedure is applied to a 2D geometry; results for different problems and found with different models are discussed, also with regard to the optimiser settings;
- **Chapter 8: Three-dimensional Optimisation of a Morphing Leading Edge Wing:** the optimisation procedure is applied to a 3D wing geometry;
- **Chapter 9: Meta-model Based Optimisation:** the optimisation framework is enriched with a meta-model, in order to reduce the fitness evaluation time; results are discussed at varying the regression methods as well as its application into the loop;

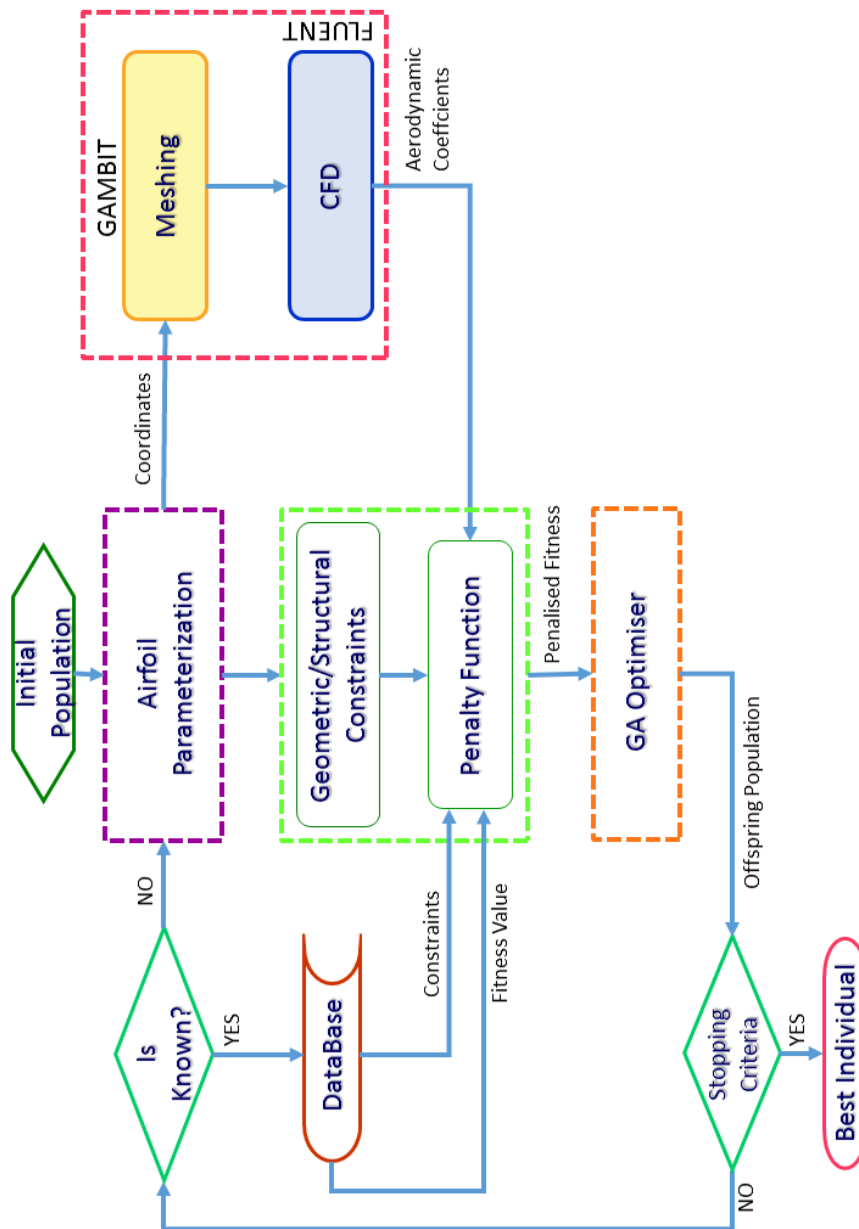


Figure 4.3: Synthetic scheme of the basic optimisation framework employed

Chapter 5

Shape Parameterization

Performing a shape optimisation problem requires the exploration of the design space where to search for the set of optimal solutions. The physical domain is explored varying the basic shape within it, while satisfying some constraints. In order to produce a family of shapes it is necessary to have a representation of the model, which can be controlled by altering the design variables. A parameterization consists in this representation through a set of parameters, whose variation allows to explore the physical domain and to find the best solutions inside it. The choice of the parameterization method has a large impact on the obtainable geometries and on the design space exploration and it should be assessed also with respect to the optimisation strategy, before performing it.

A valuable parameterization technique should have the following desirable properties [41, 61, 90]:

1. to provide high flexibility to cover the entire design space;
2. to have a small number of design variables as possible;
3. to produce smoothness and realisability of the shapes;
4. to have intuitiveness on the design parameters for geometrical and physical understanding by the designer in exploring the design space and setting up optimization constraints.

It is unlikely that a parameterization method possesses all this properties and then a selection of the most suitable strategy for the given problem should be made, considering potentials and drawbacks of each method. For example, a high number of design variables is able to provide high flexibility in exploring the design space, at the cost of an higher complexity of the optimisation loop, computational time and effort by the optimizer. Intuitive parameters allow for a well understanding of the design space and a simpler setting of the constraints. With non intuitive parameterization, extra functions can be needed to calculate the constraints, with an extra computational time and a possible reduced efficiency of the optimiser if they are not differentiable, such us in gradient

based methods. Smoothness is not just related to the mathematical regularity of the shapes, but affects also the general performance of the system in that it can influence for example both the aerodynamic and structural performance.

5.1 Available parameterization methods

Existing parameterization methods can be grouped in eight classes [42, 61]:

- Basis Vector
- Domain Element
- Partial Differential Equations (PDE)
- Discrete points (mesh methods)
- Polynomial and Spline
- Analytical methods
- CAD-based
- Free Form Deformation (FFD)

5.1.1 Discrete methods

Discrete methods are based on using the mesh point as the design variables. They are able to provide a wide design space and could possibly represent any shape. Unfortunately, since each point moves individually it is difficult to obtain smooth shapes, unless using some kind of filtering to dynamically couple the points and obtain smooth shapes and gradients. Although a high flexibility is achievable, the large number of design variables required to achieve it generates a very high dimensional design space, which can reduce the efficiency of the optimiser in searching for global optima at an unaffordable computational cost. Moreover, intuitive parameters are difficult to set and the result can be impracticable to manufacture.

5.1.2 Analytical methods

So called Analytical methods are based on a mathematical formulation to describe a certain shape. In general they rely on a superposition of a baseline geometry $y_b(x)$ with a linear combination of shape functions $f_i(x)$, whose combination coefficients AKA participating coefficients λ_i are the design variables:

$$y(x, \lambda) = y_b(x) + \sum_{i=1}^n \lambda_i f_i(x)$$

Shape functions can be of various type, e.g. Hicks-Henne, Wagner, Legendre, Bernstein or NACA series. A common analytical method was proposed by Hicks and Henne and employs a series of bump basis shape functions in the form

$$f_i(x) = [\sin(\pi x^{\beta_i})]^w$$

with $\beta_i = \ln(0.5)/\ln(x_{Mi})$ and x_{Mi} defines the position of the peak of the i -th function and w controls the width of the bump function. The technique has proved to be valid in parametrising airfoils, but is not trivial to extend it to 3-dimensional shapes.

5.1.3 Polynomials and Splines

Polynomials allow to describe a curve in a very compact way, with a limited number of control variables. The description of a shape in terms of elements of a power basis of polynomials is common to many polynomial-based approach. Among these there is the Bézier representation.

Bézier curves

A n -th degree Bézier curve can be expressed by

$$\vec{R}_g(u) = \sum_{i=0}^n \vec{P}_i B_{i,n}(u), \quad 0 \leq u \leq 1$$

with \vec{P}_i the control point, $B_{i,n}$ Bernstein polynomials of order n and $u \in [0; 1]$ the normalized abscissa. The control points \vec{P}_i describe a polygon approximating the curve and containing it and are seen as the design variables. The first and last points are positioned exactly at the beginning and the end of the curve and the convex hull of Bézier control polygon contains the curve, providing a way to set geometric constraints. The Bézier form is effective and accurate for shape optimisation of simple curves, but complex shapes require an high order p . Since the basis of polynomials is in power form and its matrix representation is Vandermonde like, inverse fitting is accompanied with numerical errors, which can badly affect the result.

B-spline

To limit the required degree, it is possible to use Bézier representation over a certain number of intervals composing the curve. This approach leads to B-splines, which are curves described as a sequence of Bézier curves, joined together at so called knot, using piecewise polynomial interpolation. The mathematical formulation of a p degree B-spline is in the form:

$$\vec{R}(u) = \sum_{i=0}^n \vec{P}_i N_{i,p}(u), \quad 0 \leq u \leq 1$$

where \vec{P}_i is the control point and $N_{i,p}(u)$ is the B-spline basis function. This basis function can be defined using the de Boor recursive algorithm:

$$N_{i,0}(u) = \begin{cases} 1 & \text{if } u_i \leq u < u_{i+1} \\ 0 & \text{otherwise} \end{cases}$$

$$N_{i,p}(u) = \frac{u - u_i}{u_{i+p} - u_i} N_{i,p-1}(u) + \frac{u_{i+p+1} - u}{u_{i+p+1} - u_{i+1}} N_{i+1,p-1}(u)$$

It is clear from the above formula that a set of u_i is required in order to compute the basis function. In fact, for a p order B-spline with n control points \vec{P}_i , $m = n + p + 1$ knots u_i are necessary. B-splines can provide a high flexibility in representing various kinds of shapes and can be used also for local control, since the control points have only a local effect on the curve.

A limit of polynomial based representation is the impossibility of accurately parametrise implicit conic shapes, like circles, ellipsis and hyperbolas. To extend the employability of these tools, rational formulation based on splines exist, namely Non-Uniform Rational B-Spline (NURBS). They are used to represent both curves and surfaces and are present in Computer-Aided Drafting (CAD) software, thanks to their flexibility and wide range of applicability.

5.1.4 Other methods

CAD-based methods uses various representation to build completely mathematically defined solids. They can handle complex and detailed geometries and have become a standard design tool in many areas, from aerodynamic to mechanics, providing connectors for multiple purposes, significant for multi-disciplinary optimisation (MDO). Intuitive parameters are also available to help in defining boundaries and constraints of the design space. Nevertheless, incorporating CAD-based models into an optimisation loop can be challenging and requires a well structured experience and competence.

PDE methods view the surface to represent as a boundary-value problem and produce the shape as the solution of PDE. Using PDE methods to parametrise existing complex model is time consuming and costly and only a surface geometry can be represented. Together with Basis Vector and Domain Element techniques, the applicability is suitable for problems involving a single discipline and simple shape changes.

5.2 Parameterization methods suitable for morphing

When facing a morphing problem, properties like local control on geometry and smoothness of produced shape should be considered as binding when choosing the parameterization technique to be used for the optimisation. Only a portion of the initial body must be varied, while maintaining a smooth surface and a certain regularity in the joint, to avoid the vanishing of the benefits of morphing because of geometric discontinuities. Having to include in the optimisation also structural constraints, even in

a first level survey in order to limit the possible solutions to those effectively manufacturable, a technique allowing to calculate derivatives and geometric properties, such as arc length, curvature or section area, would make it easier to set these constraints.

Among the aforementioned classes, polynomial based methods seem to present these features and appear as natural candidates for the purpose of the work. B-Splines, for example, have a local control property and the computation of derivatives can be made analytically. As a drawback, because of their generality they do not provide intuitive parameters. A more intuitive method to parametrise airfoil is the PARSEC, which combines a polynomial representation with a set of eleven intuitive parameters. A local control of the shape is, nevertheless, not immediate. Another polynomial based strategy is the CST, combining the flexibility of Bézier curves with a small number of intuitive parameters, which can be anyway sufficient for the purpose of LE morphing.

Class/Shape function Transformation (CST) is an application of generalized Bezier curve definition to adapt to any aircraft geometries, proposed by Kulfan and Bussoletti [42]. The CST technique is based on merging two fundamental factors, plus other additional terms to better represent specific shapes. The two factors are a Class function $C(\psi)$, which defines a certain class of geometries, and a Shape function $S(\psi)$, individuating unique geometrical properties of the parameterized body. It is applicable to a variety of problems, both in 2D and 3D, and appears a powerful tool to be used in optimisation problems. For the purpose of morphing, the characteristics of the CST method, such as the $C(2)$ regularity of the parameterization, the smooth filtering due to the shape function decomposition, the possibility of varying only certain portions of the shape and the analytical availability of lower order derivatives, appear suitable.

5.2.1 CST formulation

The CST parameterization in two dimensions is generally defined as follows:

$$\zeta(\psi) = C(\psi)_{N2}^{N1} S(\psi) \quad (5.1)$$

where $\psi = x/c$ and $\zeta = z/c$ are the normalized abscissa and ordinate of the body, respectively.

The Class function is expressed by:

$$C(\psi)_{N2}^{N1} = \psi^{N1} (1 - \psi)^{N2} \quad (5.2)$$

where $N1$ and $N2$ are two exponents that can be chosen to represent a certain class of geometries. Table 5.1 shows different classes of geometries represented by various couples of $N1$ and $N2$ exponents.

The Shape function is introduced to define unique properties of the represented body by modulating the Class function, representative of a generic shape in the geometric class. In Kulfan proposal the Shape function is expressed by a Bernstein polynomial of order n , in which its components are weighted by so called curvature correction







| N1 | N2 | shape | Figure |
|-------|-------|-----------------------------|---|
| 0.5 | 1 | round nose, pointed aft end |  |
| 0.5 | 0.5 | elliptic airfoil |  |
| 1 | 1 | biconvex airfoil |  |
| 0.75 | 0.25 | low drag projectile |  |
| 1 | 0.001 | cone or edge airfoil |  |
| 0.001 | 0.001 | rectangle or circular rod |  |

Table 5.1: Represented class of geometries with various Class function exponents

coefficients A_i . The formulation of the Shape function is, then:

$$S(\psi) = \sum_{i=0}^n A_i K_{i,n} \psi^i (1 - \psi)^{n-i} \quad (5.3)$$

where

$$K_{i,n} = \binom{n}{i} = \frac{n!}{n!(n-i)!}$$

is the binomial coefficient and A_i , $i = 0, n$ are $n+1$ coefficients that determines the particular shape described, scaling each term of the Bernstein polynomial. Without this weighting the Shape function would be a unit, because of the property of any Bernstein polynomial of order n for which the sum of its $n + 1$ terms over the interval $[0;1]$ is one. Figure 5.1 illustrates this property: with a unit Shape function a single parameterization exists for a given class function. Varying the order n of the Bernstein polynomials representing the Shape function, without scaling each term with the curvature correction coefficient A_i , produces the same unit $S(\psi)$ and thus the same geometrical shape, because the sum of the $n + 1$ terms composing the Bernstein polynomials is always one in the normalized interval $[0;1]$ (*partition of unity*).

The location of the peaks of the component Shape functions, each of the $n + 1$ terms in the sum of equation (5.3), are equally spaced along the chord and defined by

$$(\psi)_{S_{max,i}} = \frac{i}{n}, \quad i = 0, n \quad (5.4)$$

The corresponding locations of the peak of the component airfoil are also equally spaced and expressed by

$$(\psi)_{Z_{max}} = \frac{N1 + i}{N1 + N2 + n}, \quad i = 0, n \quad (5.5)$$

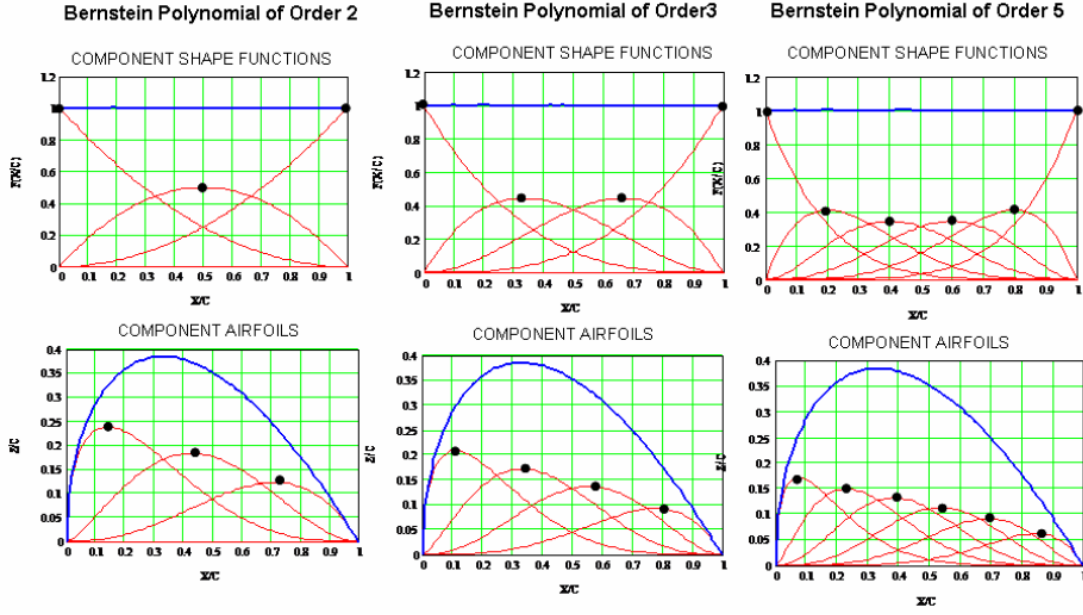


Figure 5.1: Representation of the *partition of unit* property for the Bernstein polynomials [41]

5.3 Airfoil representation through CST

For the parameterization of airfoils the $N1$ and $N2$ exponents of the Class function must be set to 0.5 and 1 respectively (see table 5.1). With this choice, the decomposition of the unit Shape function into components function and their scaling via the curvature correction coefficients A_i , which are the weight in the linear combination of the $n + 1$ terms, allows to represent a generic airfoil. It is visible by the expression of the of the Bernstein polynomials

$$B(\psi)_n = \sum_{i=0}^n K_{i,n} \psi^i (1 - \psi)^{n-i}$$

that the components of order 0 and n assume the value of 1 at $\psi = 1$ and $\psi = 0$ respectively, and the value of zero at $\psi = 1$ and $\psi = 0$, respectively. Thus, the coefficients A_0 and A_n affect only locally the shape of the airfoil at its boundaries, while the other $n - 1$ coefficients affect the entire chord. This allows to control the features of the airfoil at both leading and Trailing Edge (TE). Specifically,

$$A_0 = \sqrt{\frac{2R_{LE}}{c}} \quad (5.6a)$$

$$A_n = \tan(\beta) + Z_{TE}/c - Z_{LE}/c \quad (5.6b)$$

The formulation of the CST parameterization of an airfoil proposed in [42] is:

$$\zeta(\psi) = C(\psi)_1^{0.5} S(\psi) + \psi \zeta_{TE} \quad (5.7)$$

with $\zeta_{TE} = z_{TE}/c$ an extra term used to include the trailing edge thickness and $C(\psi), S(\psi)$ according to equations (5.2) and (5.3) respectively. In this way it is possible to list all the parameter controlling the shape of the airfoil because related to the curvature correction coefficients. Before doing this, an alternative formulation is used, as suggested in [19], in order to allow also for a leading edge deflection, which will be useful for the purpose of morphing:

$$\zeta(\psi) = C(\psi)_1^{0.5} S(\psi) + \psi \zeta_{TE} + (1 - \psi) \zeta_{LE} \quad (5.8)$$

with $\zeta_{LE} = z_{LE}/c$.

Figure 5.2 reports the geometric parameters of the airfoil as used in [19]. These parameters can be related to the curvature correction coefficients and to the constants ζ_{LE} and ζ_{TE} .




Figure 5.2: Geometric parameters of the CST representation

In order to parametrise a non symmetric airfoil, the upper and lower curves are treated separately and joined in the leading and trailing edge, thus equation (5.8) has to be applied for both upper and lower curve. It is here assumed, because of the 2-dimensional nature of the problem, that the LE is positioned at $\zeta = 0$, by posing x_{LE} in the figure equal to zero. Table 5.2 summarizes the model variables of the CST method in application to 2D airfoils.

$\zeta_{LE} = z_{LE}/c$, as already mentioned, allows to vary the vertical position of the LE and it is obviously the same for the upper and lower surface.

$$\zeta_{TE} = \begin{cases} \frac{z_{TE}}{c} + \frac{\Delta \zeta_{TE}}{2c} & \text{upper curve} \\ \frac{z_{TE}}{c} - \frac{\Delta \zeta_{TE}}{2c} & \text{lower curve} \end{cases}$$

includes the trailing edge position and thickness $\Delta \zeta_{TE}$.

R_{LE} is the leading edge nose radius. Normally it should be equal for both upper and lower curve; if different, a discontinuity will appear at LE.

C_{FS} and C_{RS} are the normalized abscissae of the fixed front and rear spar of the airfoil, delimiting the fixed portions when morphing. The leading and trailing edge deflection can be computed by $\delta_{LE} = \text{atan}(z_{LE}/C_{FS})$ and $\delta_{TE} = \text{atan}(z_{TE}/(c - C_{FS}))$.

Vectors of $n + 1$ coefficients A_l and A_u are necessary to uniquely individuating the shape described which, for the case of airfoils, is parameterized like the union of two curves.

The order n of the Bernstein polynomial determines the number of weighting coefficients of the component Shape functions and thus is related to the accuracy of the geometric reconstruction of a given body. This problem is to be treated in the next section.

| Parameter | Meaning |
|--------------------|---|
| c | airfoil chord |
| ζ_{LE} | normalized LE ordinate |
| ζ_{TE} | normalized TE ordinate |
| $R_{LE,u}$ | upper LE radius |
| $R_{LE,l}$ | lower LE radius |
| $\Delta\zeta_{TE}$ | trailing edge thickness |
| β | angle of inclination of the TE |
| δ | angle between the upper and lower curve at TE |
| C_{FS} | normalized front spar abscissa |
| C_{RS} | normalized rear spar abscissa |
| n | Bernstein polynomials order |
| A_u | upper curve curvature coefficient |
| A_l | lower curve curvature coefficient |

Table 5.2: Geometric variables in CST model

5.3.1 Inverse fitting via CST

Up to now, the CST formulation for an airfoil parameterization has been exposed, presenting the equation that defines an airfoil and the control variables, adjustable to represent different shapes. Now the problem of inverse fitting, i.e. the parameterization of an existing body expressed as a vector of $l + 1$ discrete point coordinates is treated. While in the previous section the normalized abscissa ψ has been considered the independent variable and the ordinate ζ a function of it, $\zeta = \zeta(\psi)$, now the $l + 1$ points belonging to the body are known, and thus a vector of $l + 1$ normalized coordinates $x/c = \psi$, $z/c = \zeta$ is available. The problem consists in determining the parameters in equation (5.8) in order to parametrise the shape. As regards ζ_{LE} and ζ_{TE} , they can be easily computed looking at the least and the maximum abscissa of the vector of coordinates and taking the respective ordinate. Once computed the two coefficients, in order to determine the $n + 1$ A_i coefficient for each side of the airfoil the problem must be written in a matrix form. The order n must be chosen. Starting from eq. (5.8), since the unknown terms are in the right side, it is possible to rewrite the equation like:

$$[C(\psi)_1^{0.5} B_n] \vec{A} = \zeta(\psi) - \psi \zeta_{TE} - (1 - \psi) \zeta_{LE} \quad (5.9)$$

Let $\vec{f}(\psi)$ be the right side of eq. (5.9), a column vector of $l + 1$ terms and $[C(\psi)_1^{0.5} B_n]$ an $(n + 1) * (l + 1)$ matrix $[T_n]$ such that:

$$[T_n]_{i,j} = \psi_j^{0.5} (1 - \psi_j) \binom{n}{i} (1 - \psi_j)^{n-i} \psi_j^i, \quad 1 \leq i \leq n + 1, \quad 1 \leq j \leq l + 1 \quad (5.10)$$

and \vec{A} a column vector of $n + 1$ unknown coefficients. Rearranging the linear equation in

$$[T_n] \vec{A} = \vec{f} \quad (5.11)$$

The above system is solvable with pseudo inversion, by multiplying each term for $[T_n]^T$, giving:

$$[T_n]^T [T_n] \vec{A} = [T_n]^T \vec{f} \quad (5.12)$$

This equals solving for the least square error from the analytic shape and the given one.

The system matrix $[T_n]^T [T_n]$ suffers from ill-conditioning and the solution can be affected by numerical errors. The QR decomposition method is therefore suggested so as to limit this problem, which becomes more important at increasing the order n . With the vector \vec{A} of $n + 1$ components all the curvature correction coefficients are determined and the parameterization concluded.

The CST technique shows the following properties [42]:

1. any smooth airfoil can be represented with the chosen degree of accuracy;
2. the entire space of smooth airfoils can be captured;
3. every airfoil can be derived from the unit shape function;

4. every airfoil is derivable from any other airfoil.

A survey on shape reconstruction with CST can be found in [12, 41, 42, 90]. As regards the influence of BPO, at increasing n the parameterization accuracy improves, until numerical errors highly affect the result and the system matrix becomes close to singular. This means that for a given set of points more than one parameterization exists, or vice versa two different parametrisations refer to practically the same shape. In fact, the solution of the system is the one that minimizes the Least Square Error (LSE), because the system matrix is not square; for this reason, when stating the passage through control points, this must be intended in the LSE sense. Figure 5.3 illustrates the RMS error between the parameterized airfoil and the starting one for a NACA 4412 as a function of BPO order n . Ordinate axis is logarithmic.




Figure 5.3: RMS error versus BPO n for NACA4412 airfoil

Following Figure 5.4a shows the parameterized NACA 4412 profile with control points marked by circles. Figure 5.4b reports the error difference between the parameterization and the initial discrete shape, calculated at the control points. Note that the control points used here are not the same of Figure 5.4a. A better fitting could also be obtained with an LSE optimisation for N1 and N2 Class function exponents, as suggested in [43].

5.3.2 Application of CST to morphing airfoils

One of the advantages of the CST method is that it can handle morphing geometries derived from a native one, as expressed by the aforementioned property 4. The problem can be formulated in this way: given a parameterized airfoil with BPO n , some m over $n + 1$ curvature correction coefficients A_i are perturbed, while maintaining a subset of $l + 1$ control points in the normalized domain $[0;1]$. In order to force the airfoil to conserve these control points, the other $n - m + 1$ coefficients must be scaled, in that each component Shape function, different from the lowest and highest one ($i = 0$ and $i = n$), affects the final shape in the entire interval. Perturbing m coefficients means

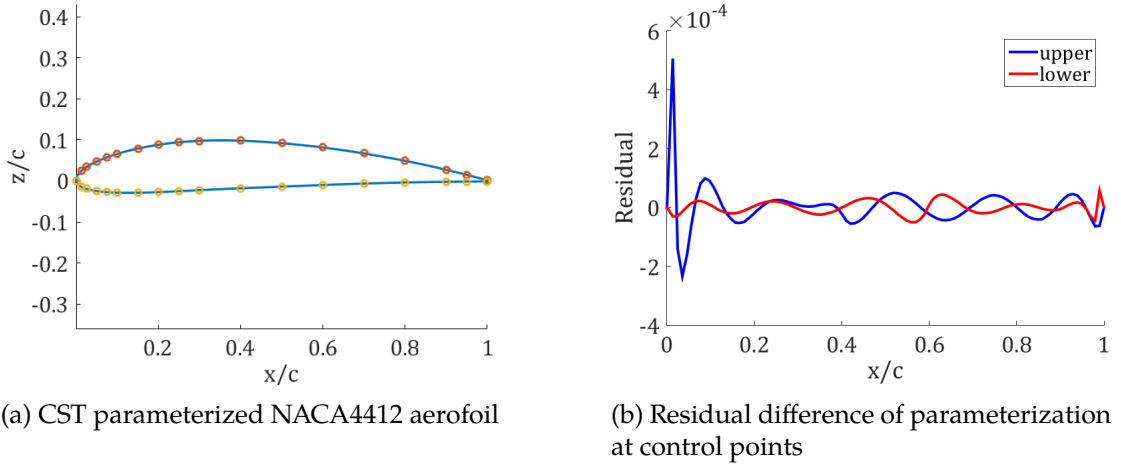


Figure 5.4: Parameterization results for a NACA4412 with BPO = 10

assigning them a value and then the remaining unknowns are the $n + 1 - m$ coefficients to be determined, while imposing the passage through the control points. This means to solve a system similar to that of equation (5.11), in which the system matrix is extracted from the original one and has dimensions $(n - m + 1) * (l + 1)$. Operatively it is derived by deleting the columns corresponding to the perturbed m coefficients; this columns multiplied per the scalar coefficient must then be subtracted from the right member. Formally,

$$[T_{n-m}]_{i,j} = \psi_j^{0.5} (1 - \psi_j) \binom{n}{i} (1 - \psi_j)^{n-i} \psi_j^i, \quad 1 \leq i \leq n - m + 1, \quad 1 \leq j \leq l + 1 \quad (5.13)$$

and

$$f_j = -\psi_j \zeta_{TE} - (1 - \psi_j) \zeta_{LE} - C(\psi_j)_1^{0.5} \sum_{i=0}^m A_i B_i(\psi_j), \quad j = 1, \dots, l + 1 \quad (5.14)$$

The linear system is finally:

$$[T_{n-m}] \vec{A} = \vec{f} \quad (5.15)$$

and can be solved with pseudo-inversion and QR method as before.

The first and last curvature coefficient A_0 and A_n are related to the LE radius and the TE position and the deflection angle (eq. (5.6)). The other coefficients express their influence over the entire domain, but mostly near the peak of their component airfoil (eq. (5.5)) and then the coefficients to be perturbed can be selected relying on the position of the peak with respect to the fixed subdomain, the latter being individuated by C_{FS} and C_{RS} in Figure 5.2.

Figure 5.5 shows an example of morphing using the CST parameterization. The original airfoil is a NASA HSNLF(1)-0213 (grey in the figure). The morphed profile, green in the figure, results from a perturbation of the original curvature coefficients for both




Figure 5.5: Morphing LE example with CST parameterization for NASA HSNLF(1)-0213 profile

the upper and lower side and a variation of ζ_{LE} . The fixed portion is between $[0.25;1]$ and the reduced system was solved to scale the free curvature coefficient and maintain the passage through the control points, marked with dots in the figure. The BPO was 8; the reconstruction error in the fixed portions depends on the coefficients perturbed. One can select which coefficients to perturb basing on the position of the component peak (eq. (5.5)) with respect to the subdomain. When introducing a variation in some coefficient, the remaining ones scale in order to maintain the fixed portion unchanged. Actually, since the inverse fitting is a Least Square Error approximation, the parameterized aerofoil will not pass exactly through the control points. This leads to a small variation of the perturbed profile in the fixed region, with respect to the original parameterized profile. The difference between the unperturbed parametric airfoil and the perturbed one in the fixed portion is presented in Figure 5.6, together with the error between the control points and the parametric unperturbed airfoil. It can be noted that the maximum error in the latter set (dashed line), where no perturbation has occurred, is below $2.5E-4$, while in the former (continuous line) it raises up to $1E-03$. Data have been obtained with a BPO of 10. Increasing the order of Bernstein polynomials can reduce the RMS value of the error, but results in a high conditioning number of the inverse fitting system matrix, so $n=10$ appears to be the trade-off value. Through the perturbation of the first three curvature coefficients A_0, A_1, A_2 for $n = 10$, whose component peaks are located at 0.0435, 0.1304, 0.2174 respectively, it is possible to obtain different morphed shapes of the airfoil in the morphing subdomain $[0;0.25]$. Starting from an unperturbed airfoil, the variation of the coefficients by multiplying each of them for a different scale factor produces a family of curves exhibiting $C(2)$ regularity. Obviously, not all possible sets of three perturbed coefficients are admitted to preserve an airfoil like shape. Some constraints can be introduced to discard those sets which violate them, for example to limit the maximum error between the fixed control points and the perturbed airfoil or the variation of area in the deformed region. Figure 5.7 presents the results for the aforementioned procedure applied to a NASA GA(W)-1 airfoil, in which a LE deflection of 15 [deg] has been imposed. The original airfoil is drawn in grey and the fixed spar is




Figure 5.6: Residual error between perturbed and original CST parametric NASA HSNLF(1)-0213 profile in the fixed region

dashed at $\psi=0.25$.

The former figure shows the results of the perturbation, when only geometrical constraints are considered, but also structural ones should be introduced, in order to limit strain or stresses in the morphed configuration. For this purpose it can be convenient to pass from a parametric representation of the airfoil in term of normalized abscissa ψ to an arc length representation.

Each side of the CST airfoil is expressed as a function

$$f : [0; 1] \rightarrow \mathbb{R}, \psi \mapsto f(\psi) = \zeta$$

To turn the parametric representation into a normalized arc length one, it is necessary to map the curve $f(\psi)$ into the curve $g(s)$, such that

$$g : [0; 1] \rightarrow \mathbb{R}^2, s \mapsto (x(s), y(s))$$

and s is the arc length, who can be expressed as a function of t :

$$s(t) = \int_0^t \sqrt{1 + \left(\frac{d\zeta}{d\psi}\right)^2} d\psi \quad (5.16)$$

The above relationship is explicit when passing from the adimensional parameter t to the arc length s , but being non linear an inverse explicit formulation is not available and recursive methods can be used to calculate the value of ψ corresponding to a given value of s . Anyway, since in the rebuilding of a CST parameterized aerofoil a discrete




Figure 5.7: An example of family of morphed LE for a NASA GA(W)-1 airfoil

number of points has to be used, it is possible to directly couple the ordinate, function of a known ψ , to the corresponding value of s calculated for the same ψ using eq.(5.16). In this way, since the number of points used are sufficiently high, no inverse computation of ψ for a given s is required and the mapping can be made straightforward, simply associating to a tabulated ψ its corresponding value of s .

5.4 Constant length curve parameterization with CST

The peculiarity of the problem treated requires a parameterization technique able to ensure that the constraints present in the optimisation, which will be both geometrical and structural, can be easily respected within the optimisation loop, thus avoiding for a long search in the design variable space for feasible solutions. Constraints treatment is, indeed, a big topic when dealing with optimisation and providing the shape generation algorithm the capability of automatically satisfy some of them improves the efficiency of the search for optimal feasible solutions. In this context, where morphed shapes are limited by the deformation the skin is able to bear, the generation, by the parameterization subsystem, of a family of curves whose length, in the morphed region, remains constant can make it easier to respect structural constraints. For this aim, a specific procedure has been individuated and tested.

The generation of a family of curves with constant length is not straightforward, because of the nonlinear relationship between the length and the control variables of the parameterization. Starting from the work of Bettadapura et al. [9] for Bézier curves, a semi-analytic algorithm to obtain a constant length family has been set. It is based on the observation of how the length of a CST parameterized airfoil varies at varying one of

its curvature correction coefficients, while recalculating the others to keep the geometry fixed in a region (as described in section 5.3.2). This relationship appears parabolic-like, as shown in Figure 5.8. The graph refers to the relative variation in length against the variation of the third curvature coefficient of the CST parameterization, when the reduced system is solved imposing the first three coefficients A_0, A_1, A_2 , with the former two equal to those of the baseline shape and the third variable, $A_2 = A_2 + \Delta$.




Figure 5.8: Relative variation of the length on the morphed region of an airfoil as a function of the variation of the third curvature coefficient A_2

In a first step, like in [9], an algorithm aimed at obtaining a curve whose length differs from the original one for a given quantity was developed. It is based on a parabolic approximation of the curve shown in Figure 5.8. Starting from a zero value for Δ , the relationship is locally approximated at second order by computing the length at three points of the curve, $[\Delta - \epsilon, \Delta, \Delta + \epsilon]$ and representing it as a parabola. Having a local approximation of the function $L(\Delta)$, the new value of Δ to obtain a length equal to the original one plus the imposed variation δL is computed. This procedure is recursively repeated until the final length reached equals the desired one. In this way a curvature coefficient is varied of the quantity Δ until the curve length equals the value sought. A qualitative scheme of the algorithm is shown in Figure 5.9. One coefficient must be chosen to be varied in order to attain the desired length change.

If the procedure is applied when some coefficients of the parameterization are imposed and one is chosen to vary in order to keep the length constant, it is possible to have a parameterization able to generate curves with constant length, given that this can only occur if the imposed coefficients belong to a certain region. This is, indeed, the second step. For example, the first two curvature correction coefficient are imposed. The new curve length is computed and the third coefficient is changed, according to the first step, in order to keep the length constant. The algorithm is summarized in Figure 5.10. From a practical point of view, the imposed coefficients are set by the optimizer, while the one to be varied in order to keep the length constant is chosen at the beginning by the user.

An example of a family of morphing LE airfoil with constant length is illustrated in Figure 5.11, for a NASA GA(W)-1 airfoil, with fixed spar at $0.25c$ and the third curvature

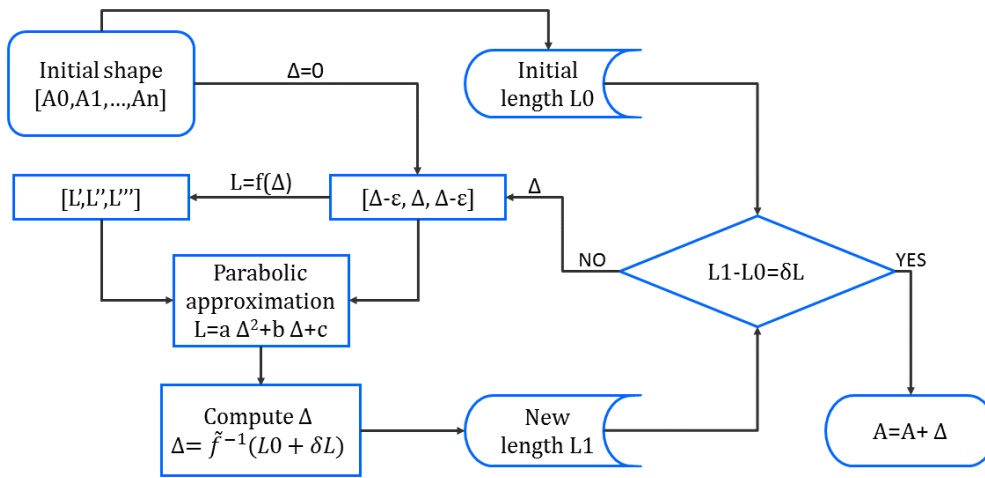


Figure 5.9: Scheme of the algorithm employed to obtain a new curve with imposed length variation

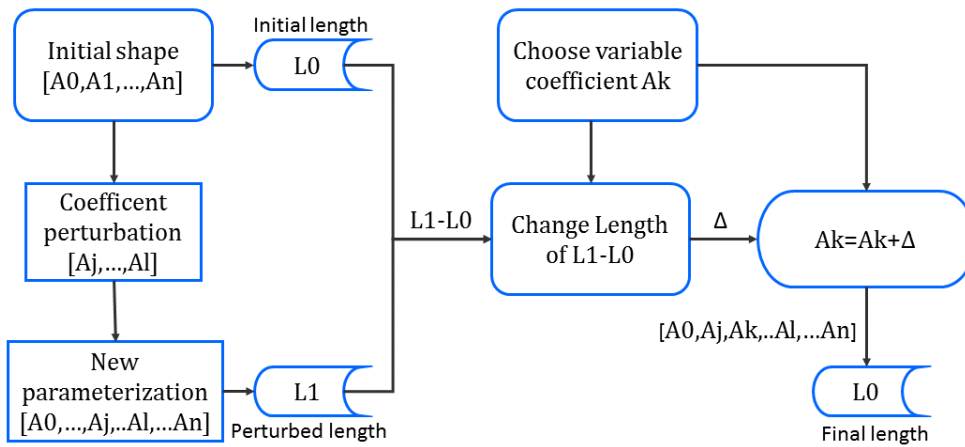


Figure 5.10: Scheme of the algorithm employed to obtain a constant length curve when some coefficients are imposed




Figure 5.11: An example of morphing LE airfoil with constant length

correction coefficient chosen to keep a fixed length, while imposing A_0 , A_1 and Z_{LE} .

Chapter 6

Aerodynamic Analysis Tool Validation

In the computational loop, the evaluation of the performance of a geometry in terms of aerodynamics has to be done using a numerical analysis tool. CFD offers a wide range of instruments to model and simulate a variety of fluid dynamics problems with possibly an high level of accuracy. On the other hand, a CFD analysis can be very expensive in terms of computational resources and its employment in an optimisation loop can result into an unacceptably long time, especially when stochastic methods, broad-spectrum but low efficiency, lead the search.

A powerful analysis tool in external aerodynamics problems is represented by panel methods, based on velocity potential flows theory. When the hypothesis of the linearisation are valid and the velocity field can be expressed by the gradient of a scalar potential, panel methods produce reliable results for the pressure and velocity distribution along an aerofoil, at a meaningless computational cost. Linear methods are applicable for thin bodies, at moderate angles of attack, when no extended separation is present and the regime is subsonic. A popular code based on linear flow, combining a viscous BL interaction to allow fixed or free transition prediction, moderate trailing edge separation and lift and drag prediction up to $C_{l,max}$ is Drela's XFOIL [24]. This code is widespread for external subsonic aerodynamic analysis, where its reliability and very low computational time makes it very competitive with CFD, given that the linear model remains valid. For a wider list of theoretical references about XFOIL, see [23]. Many optimisations with morphing report the use of the open source tool [27],[85],[22] and in [84] a comparison between XFOIL and a CFD code shows a bit difference between them for the FishBAC concept already presented (fig.2.3c).

For the purpose of the work, the use of an aerodynamic analysis tool with validated results is needed. Since *a priori* it is not possible to discard linearised flow methods, whose employment would greatly reduce the computational resources, both a so called high-fidelity CFD code and the low-fidelity XFOIL were tested and their results compared, together and with an experimental test case chosen. In the next section the CFD code set-up is reported and then outcomes obtained with both the tools are presented.

6.1 CFD simulation set-up

The set-up of a CFD simulation requires a wide choice of parameters and models, as well as a discrete representation of the domain. The models employed and the domain mesh are here described.

6.1.1 Grid Generation

For the nature of the finite volume approximation a domain discretization is needed. In general, the result can be largely affected by the mesh employed, because of both physical and numerical issues (e.g. a geometric scale unable to capture fluid gradients or a poor mesh elements quality). A grid generation for a CFD problem can be performed using loosely three methods: structured, unstructured and hybrid grids. The first ones are generated relying on some user defined background structures or scheme to distribute the elements within the domain. There are some advantages in such an approach, because it is quite controllable the orientation of the layers of elements and their density. The nodes can be easily ordered in a connection matrix and their position uniquely identified, so that a scheme can be easily replicated. Normally, elements used are quadrilateral in 2D and hexahedrals in 3D.

Unstructured grids are generated by some algorithm to fill the domain with elements without using a user defined control structure. This results in the possibility of meshing complex geometries in a quite fast way, with good quality and minimal user post-processing. A control in element size and distribution can be achieved through some function forcing the algorithm, but of course the exact orientation and density is less controllable than the structured grids. Typically triangular elements are used in 2D domain and tetrahedra in 3D.

A third option is hybrid grids, combining layers of structured elements connected to zones with unstructured ones. In this way it is possible to properly resolve critical areas by using structured elements, while filling the remaining parts with a less ordered elements distribution. A typical example is the adoption of a structured grid in the boundary layer, where the need for orthogonality and a precise control of cell height and stretch factor are desired. The connection between two grids can be a critical area if an excessive stretching and skewness are required. A proper setting of size functions to control the unstructured grid can limit the post-process adjustment by the user.

For the purpose of the work, both structured C-grid and hybrid O-grid with structured BL were tested. The former are quite easy to build thanks to the domain regularity but has shown to loosely match experimental data in CFD computation, unless increasing the number of elements to a value which has been judged too expensive for the computation and not leading to advantages with respect to the latter kind, actually chosen. After several tests, the boundary layer was filled with 50 layers of quadrilateral elements and the first cell height chosen to have an $y^+ < 1$, in order to properly solve the boundary layer without using wall functions. The software used for meshing was ANSYS GAMBIT v.2.4.6. An overall view of the hybrid grid generated is depicted in Figure 6.1, while in Figure 6.2 are illustrated some details of the mesh, at the leading

and at the trailing edge of the aerofoil.




Figure 6.1: Details of the hybrid grid used

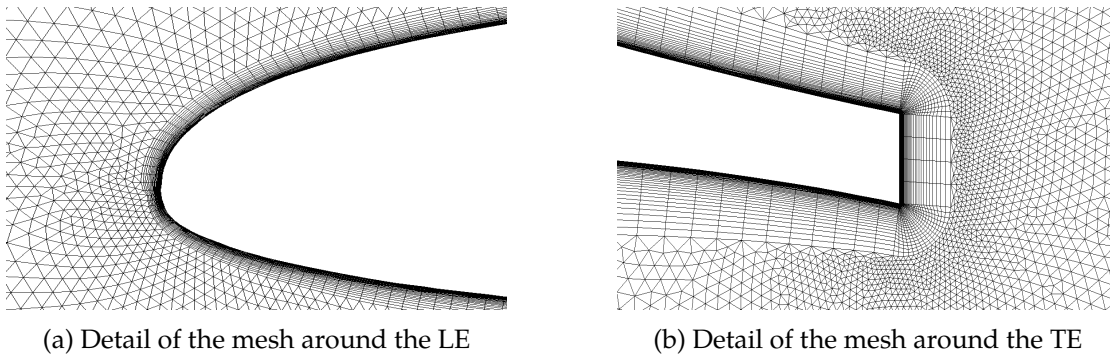


Figure 6.2: Details of the hybrid grid used

6.1.2 Material Properties and Reference Quantities

The problem involved air as a fluid, with a free stream Mach number of 0.20 and a Reynolds number of $6E+06$. It was modelled with an ideal flow with molecular weight $M_m = 28.966 \text{ kg/kmol}$, particular gas constant (\bar{R}) = 287.058 J/kgK and constant specific heats ratio $k = 1.4$. The viscosity model used was Shuterland's law with reference temperature $T_0 = 273.11 \text{ K}$ and reference viscosity $\mu_0 = 1.716 \cdot 10^{-5} \text{ kg/ms}^{-1}$.

6.1.3 Flow equations

The CFD code is employed to solve the flow equations in the domain, by typically using a discrete approximation of the quantities in order to numerically solve for them. The equations describing the physical behaviour of a Newtonian fluid are the Navier-Stokes equations, which are the meaningful and exact physical law for the motion of

a fluid. Since no general integrated form of them is available, a numerical scheme is adopted to solve them in practical cases. Moreover, the high non-linearity associated with them gives rise to strange phenomena, in presence of certain boundary conditions, one of which is turbulence. Turbulent flows exhibit rapid local fluctuations of the mechanical and thermodynamic variables over a multi-scale level, resulting in an instantaneous field not respecting the symmetry or the time invariance of the boundary conditions. In order to properly solve a turbulent flow using Navier-Stokes exact equations, the computational grid should be able to capture the very steep gradients present in the domain, requiring an element size of the order of the Kolmogorov scale $\eta = (\nu^3/\epsilon)^{1/4}$, where ν is the kinematic viscosity and ϵ the turbulent dissipation, proportional to the specific power introduced into the fluid. Furthermore, the flow must be evolved also with respect to the time domain, requiring a time-step discretization. The direct solution of the Navier-Stokes equations is referred to as Direct Numerical Simulation (DNS). Being L_0 a characteristic length of the material domain used as a reference, the computational cost of a DNS is $N_P \times N_\tau \propto \text{Re}^3$, with N_P the number of points over a 3-D domain and N_τ the number of time-steps. This is due to the need of capturing all the turbulent structures arising from the Kolmogorov scale, which is of an order of magnitude of the molecular size, to a macro-scale which is of the order of magnitude of the physical material domain. The outcome is a computational time of hundreds of years for a high Reynolds flow over a large domain.

In an engineering and industrial approach, where computational resources are limited, as well as the time available, DNS are not employed. Instead, a different formulation of the Navier-Stokes equations is adopted, based on the so called *Reynolds average* or *ensemble average*. This way of averaging the flow quantities consists virtually in repeating many times an experiment, without changing the boundary conditions, and registering point by point all the physical quantities. After that, an average over the number of the experiments is performed point by point, giving rise to a mean flow, exhibiting the symmetry and time-invariance of the boundary conditions. In this way each quantity can be decomposed as the sum of its mean value $\langle u \rangle = U$ (in a Reynolds sense) and a turbulent fluctuation, proper of the turbulence motion u' : $u = \langle u \rangle + u'$.

With this decomposition, the original Navier-Stokes equations become the so-called RANS (Reynolds Averaged Navier-Stokes) equations:

$$\vec{\nabla} \cdot \vec{U} = 0, \quad \vec{\nabla} \cdot \vec{u}' = 0 \quad (6.1a)$$

$$\frac{\partial \vec{U}}{\partial t} + \vec{U} \cdot \vec{\nabla} \vec{U} = -\frac{1}{\rho} \vec{\nabla} P + 2\nu \vec{\nabla} \cdot \vec{\bar{E}} - \vec{\nabla} \langle \vec{u}' \vec{u}' \rangle \quad (6.1b)$$

$$\frac{\partial K}{\partial t} + \vec{U} \cdot \vec{\nabla} K = \vec{\nabla} \cdot \left\{ -\frac{P}{\rho} \vec{U} + 2\nu \vec{\bar{E}} \cdot \vec{U} - \vec{U} \langle \vec{u}' \vec{u}' \rangle \right\} - \underbrace{2\nu \vec{\bar{E}} : \vec{\bar{E}}}_D + \underbrace{\langle \vec{u}' \vec{u}' \rangle : \vec{\nabla} \vec{U}}_{-P} \quad (6.1c)$$

$$\frac{\partial k}{\partial t} + \vec{U} \cdot \vec{\nabla} k = \vec{\nabla} \cdot \left\{ -\frac{\langle p' \vec{u}' \rangle}{\rho} \vec{U} + 2\nu \vec{\bar{e}}' \cdot \vec{u}' - \left\langle \frac{\vec{u}' \cdot \vec{u}'}{2} \vec{u}' \right\rangle \right\} - \underbrace{2\nu \vec{\bar{e}}' : \vec{\bar{e}}'}_\epsilon - \underbrace{\langle \vec{u}' \vec{u}' \rangle : \vec{\nabla} \vec{U}}_P \quad (6.1d)$$

where upper case denotes mean values and prime or lower case denotes turbulent quantities.

The first equation expresses the conservation of the mass and is the continuity equation. The second equation is the balance of the momentum. The term $\vec{\bar{E}}$ is the symmetric part of the gradient of the mean velocity U . The tensor $\langle \vec{u}'\vec{u}' \rangle$ is the Reynolds stress tensor, made up of the turbulent velocity components. The third equation is the energy balance for the mean kinetic energy $K = \frac{U_i U_i}{2}$, while the fourth is the energy balance for the turbulent kinetic energy $k = \frac{\langle u'_i u'_i \rangle}{2}$. The underlined terms represent meaningful quantities for the RANS approach. D is the energy dissipation associated to the mean flow and it is a quantity always lower than zero. It dissipates kinetic energy in the mean flow through viscous stresses. The term P is the production of turbulent kinetic energy and is a quantity statistically positive. In the K equation it compares with a minus sign, while in the k equation it has a plus sign. This can be interpreted as the fact that for some conditions, the dissipation term activates, transferring energy from the mean flow to the turbulent one, and thus creating the turbulence. The energy from the mean flow is converted into turbulent energy and in this regime the term ϵ , which is negative, dissipates it, creating turbulent structures over a multi-scale level, until the scale length is sufficient to ensure that viscous forces can eliminate the energy introduced. This quite strange and complex mechanism is due to the high non-linearity of the Navier-Stokes equations, but in the RANS approach it is not simulated, it is modelled according to various formulations. Popular turbulence closures are the $k - \epsilon$ or the $k - \omega$, but a variety of models exists, always with a certain degree of arbitrary elements. The next subsection presents the turbulence closure employed for the 2-D simulations, chosen over the models available in the CFD code employed.

Turbulence modelling

Turbulence closure for RANS equations is one of the main variable when solving turbulent flows, because the result can be largely affected by the capability of the model to replicate physical phenomena. The software used for the CFD calculation, ANSYS Fluent, allows for a quite wide choice of turbulence closures. For aeronautical application a common one is the Spalart-Allmaras (SA) model.

Spalart-Allmaras turbulence model [71] is a Boussinesq model, based on one transport equation modelling the kinematic eddy viscosity ν_t . In the Boussinesq approach, the Reynolds stress tensor is modelled in a similar fashion to the stress tensor due to molecular viscosity transport and is expressed by:

$$-\langle u'_i u'_j \rangle = -\frac{2}{3}k\delta_{ij} + 2\nu_t E_{ij} \quad (6.2)$$

where curl brackets denotes ensemble average, prime quantities denote fluctuations variables, k is the turbulent kinetic energy and E the symmetric part of mean velocity deformation tensor.

In SA model, the transport equation for ν_t is

$$\frac{\partial}{\partial t}(\rho\nu_t) + \frac{\partial}{\partial x_i}(\rho\nu_t u_i) = G_\nu + \frac{1}{\sigma_{\nu_t}} \left[\frac{\partial}{\partial x_j} \left\{ (\mu + \rho\nu_t) \frac{\partial \nu_t}{\partial x_j} \right\} + C_{b2}\rho \left(\frac{\partial \nu_t}{\partial x_j} \right)^2 \right] - Y_\nu + S_{\nu_t} \quad (6.3)$$

where G_ν is the production of turbulent viscosity, and Y_ν is the destruction of turbulent viscosity that occurs in the near-wall region due to wall blocking and viscous damping. σ_2 and C_{b2} are the constants and ν is the molecular kinematic viscosity. $+S_{\nu_t}$ is a user-defined source term. Since the turbulent kinetic energy is not computed in the model, in the previous Boussinesq formulation the symmetric part of the Reynolds stress depending upon it is simply ignored.

The eddy viscosity is modelled as:

$$\mu_t = \rho\nu_t f_{\nu 1} \quad (6.4)$$

with $f_{\nu 1}$ a viscous damping term given by:

$$f_{\nu 1} = \frac{\chi^3}{\chi^3 + c_{\nu 1}^3} \quad (6.5a)$$

$$\chi = \frac{\nu_t}{\nu} \quad (6.5b)$$

The production term is modelled as:

$$G_\nu = C_{b1}\rho S_T \nu_t \quad (6.6)$$

with:

$$S_T = S + \frac{\nu_t}{\tilde{k}^2 d^2} f_{\nu 2} \quad (6.7a)$$

$$f_{\nu 2} = 1 - \frac{\chi}{1 + \chi f_{\nu 1}} \quad (6.7b)$$

C_{b1} , $C_{\nu 1}$ and \tilde{k} are constants and d is the distance from the wall. S is expressed by

$$S = \sqrt{\Omega_{ij}\Omega_{ij}} \quad (6.8)$$

with Ω_{ij} the deviatoric part of the mean velocity deformation tensor.

The destruction term is expressed by:

$$Y_\nu = C_{w1}\rho f_w \left(\frac{\nu_t}{d} \right)^2 \quad (6.9)$$

where:

$$f_w = g \left[\frac{1 + C_{w3}^3}{gC + C_{w3}^6} \right]^{16} \quad (6.10a)$$

$$g = r + C_{w2}(r^6 - r) \quad (6.10b)$$

$$r = \frac{\nu_t}{S_T \tilde{k}^2 d^2} \quad (6.10c)$$

Reference values of the eight model constants $C_{b1}, C_{b2}, \sigma_\nu, C_{v1}, C_{w1}, C_{w2}, C_{w3}, \tilde{k}$ are provided in [71] and preloaded into Fluent. The model, designed specifically for aerospace applications, has shown to provide good results for wall-bounded flows and adverse pressure gradients boundary layers.

6.2 Test case validation

The calibration of the CFD settings was performed in order to obtain a result as close as possible to an experimental test case, chosen to validate the numerical simulation. The test case was a general aviation airfoil section, developed by NASA and described in [49]. In a first step, the grid influence was evaluated, changing the element density and the structured BL grid, until a good agreement with experimental data has been achieved. The final hybrid O-grid employed, already shown in Figure 6.1, had about 60000 cells. Simulations were run to compute the lift coefficient variation with the angle of attack, the Lift-Drag polar and the variation of the moment coefficient with the angle of attack. Test conditions were a Reynolds number close to $6E06$ with a Mach number of 0.15, 0.20 and 0.28. The comparison of numerical and experimental data over this range of Mach number allowed for a better assessment of the validity of the models embedded in the simulation. For the sake of brevity, only one set of validation results is here reported in following Figure 6.3 and Figure 6.4.




Figure 6.3: $C_l - \alpha$ validation results for $M = 0.20$, $Re = 6.2 \times 10^6$

The process of validation is, in general, time consuming, requiring an iterative procedure in which many variables can be chosen, starting from the grid, the turbulence model and its coefficients and coming to the discretization scheme and advection method.




Figure 6.4: Lift-Drag polar validation results for $M = 0.20$, $Re = 6.2 \times 10^6$

For this work, many grids were tested, improving by time regions of the mesh which appeared still too coarse to capture the fluid gradients. The evaluation of the grid performance was carried out by computing the entire polar for C_l , C_d and C_m and comparing to the experimental data available. Once the trend of those data proved to be replicated correctly, the target was trying to get as close as possible to the experimental charts. In fact, though there was an accordance between the trends, as it is visible in the figures, in particular for the drag the CFD plot was still biased with respect to the experimental one, while for the coefficient of moment at low angles of attack the graphs were quite different. In order to fill this gap, a huge effort was required, by refining the mesh, changing turbulence closure, using different advection schemes and discretization order. Notwithstanding this, a gap still persisted between the Lift-Drag polars and also for a portion of the AoA-Moment polar, as explained later. Partial results obtained are not reported in this text, but qualitatively speaking, the convergence to the actual outcomes was rapid, while in many attempts of getting closer to the experimental data only a slight change was obtained.

In order to deal also with a quantitative index, once individuated a good mesh a successive refinement was performed, calculating the three polars each time and computing the RMS error between the CFD chart and the experimental one:

$$RMS_{error} = \frac{1}{NoP} \sum_{i=1}^{NoP} \sqrt{(y_{CFD} - y_{Exp})^2}$$

where y_{CFD} is the data obtained by the CFD simulation, while y_{Exp} is the experimental data. To compare the values in the same abscissa, since the starting data were sampled, a cubic spline interpolation was used. The trends in the RMS errors versus the number of elements is illustrated in Figure 6.5. The boundary conditions and turbulence modelling were the same for all the cases. The number of elements changed from about 30'000 to

120'000, refining each time the mesh around the airfoil and reducing the growth factor starting from their side. No remarkable variation appeared; specifically, the AoA-Lift residuals and the Lift-Drag residuals showed an opposite behaviour, with the maximum of the first associated with a minimum of the second. For the AoA-Moment errors, instead, the variation was negligible and the curve was very flat. In examining this charts, let remember that they still have an indicative value, since the absolute RMS error is by itself not meaningful for comparison. Nevertheless, such a report can help choosing the best grid overall, considering the opposite effect of the refinement on all the quantities enquired.




Figure 6.5: Variation of RMS residuals with the number of grid elements, computed for $M = 0.20$, $Re = 6.2 \times 10^6$

Moreover, since the aim of this analysis stage was the practical implementation of the grid model for the optimisation process, a variable taken into account was the time of convergence. The absence of remarkable modifications of the polars at increasing the number of cells suggested to pick the grid with a number of elements of 40'000, providing practically the same outcomes of the other finer meshes, in order to speed up the expensive optimisation session.

The analysis dealt also with changing the turbulence model. For this 2D case the aforementioned Spalart-Allmaras model appeared a good candidate and simulations were performed using it and the $k - \omega SST$, which is also, generally speaking, well behaving even for more complex flows. The tests revealed a better performance of the Spalart-Allmaras, which was then chosen and employed, also for the results reported above, about the influence of the grid resolution.

Furthermore, different flow solver strategies, schemes and discretizations were tested. It was found that convergence speed could be increased with non-segregated solver, where all the equations are solved in a coupled way. The next Table 6.1 summarizes the characteristics of the final hybrid grid chosen and the settings of the CFD solver Fluent.

| | |
|-----------------------------------|--------------------------------------|
| Cells | 40194 |
| Boundary Conditions | Far-field (Mach, Pressure, Velocity) |
| Boundary Layer Rows | 50 with growth ratio < 1.12 |
| Turbulence Model | Spalart-Allmaras |
| Pressure-Velocity Coupling | Non-segregated solver |
| Gradient approximation | Third order MUSCL |

Table 6.1: Summary of CFD settings for the Spalart-Allmaras model

6.3 Comparison between CFD and XFOIL results

As previously stated, the employment of a panel method is preferable for its incomparably smaller computational time with respect to the CFD. The problem treated deals with a subsonic flow around a slender aerofoil, at moderate angle of attack. Since no extended stall was developed, the viscid-inviscid formulation of XFOIL could be able to predict the velocity and pressure field with an acceptable accuracy. Indeed, from Figure 6.6 it is visible a slight difference between CFD with SA model and XFOIL C_l - α curve. Figure 6.7, instead, shows a similar trend for the CFD and XFOIL Drag-Lift polar. Both curves were shifted with respect to the experimental one, positively for the CFD and negatively for XFOIL. The latter appeared able to capture better the experimental trend in the laminar bucket. As regards the moment coefficient, it is known that its representation by a numerical code is even more difficult than that of the drag coefficient. The actual grid employed was the best result obtained for all the three polars. Specifically, for the $C_m - \alpha$ graph in Figure 6.8 a very good agreement with experimental data was achieved for $\alpha = 10 - 18$ [deg] by the SA model, while for lower angles of attack XFOIL seemed to better predict experimental data. These strange deviation of the CFD for lower angles resulted low sensitive to the grid. A further effort was put in trying to improve the CFD prediction, firstly by changing the mesh, but unfortunately the discordance appeared to be little responsive to the mesh kind itself. Finer discretizations with a more filled boundary layer produced quite the same result in this regime, but exhibited an increased deviation for higher angles, where the grid presented, instead, worked quite well. Given that, the actual grid was the one which proved to lead to the best result overall, in a Pareto fashion, for the combination of models and boundary conditions reported in Table 6.1.

As a further check on the validity of the tools, also a simulation of a local distribution was performed. Specifically, since among the experimental data also some pressure coefficients distribution were available, this parameter was computed by both the solvers and the results are illustrated in Figure 6.9. A good agreement with the test points was reached in both the cases, with the CFD giving a best result for the pressure side and the suction side up to 0.15 c. From this point on, XFOIL had a slightly better agreement.




Figure 6.6: $C_l - \alpha$ validation results for $M = 0.20$, $Re = 6.2 \times 10^6$




Figure 6.7: $C_d - C_l$ polar validation results for $M = 0.20$, $Re = 6.2 \times 10^6$




Figure 6.8: $C_m - \alpha$ validation results for $M = 0.20$, $Re = 6.2 \times 10^6$




Figure 6.9: C_p validation results for for CFD and XFOIL at $M = 0.15$, $Re = 6.3 \times 10^6$, $\alpha = 4.17$ [deg]

Potential flow solvers like XFOIL appears an efficient and inexpensive tool useful for aerodynamic forces calculation and also transition analysis, but provides valid results only if the basic hypothesis allowing linearisation are satisfied. Due to the moderate angle of attack and the variety of shapes tested in the current optimisation causing potentially non linear representable effects, such as extended stall and separation, a CFD finite-volume based calculation was preferable. Notwithstanding some models contained in the code are quite arbitrary, e.g. in turbulence closure, various phenomena can be represented with a moderate level of uncertainty and if the code is validated using experimental data, the results produced can be an acceptable approximation of the physics involved.

The slight difference between the two approaches shown for the test case suggested the possibility of exploiting the speed of XFOIL in the first phase of the optimisation loop set-up, when many runs were necessary to check the robustness of the loop and tune the optimiser. This allowed also to inspect the design space and check the influence of the constraints in the final result, in order to avoid non-linear constraints too strict or not consistent. In a second phase, when the optimisation framework appeared robust and efficient, the CFD was used in place of XFOIL for the aerodynamic indexes computation.

The outcome from this second phase showed an unexpected difference between the results produced by XFOIL and those of the CFD. This will be discussed in details in 7.6. As a consequence of that discordance, several further attempts were made in changing the CFD settings and verify the accuracy of the predictions, with respect to the experimental data. The next section reports the activity performed for this aim, which was carried out after comparing the optimised shape for a Single Objective Optimisation Problem (SOOP) obtained with both XFOIL and Fluent, and it is here reported in continuity with the topic of the chapter.

6.4 Further validation activity

The lack of accordance between the CFD and XFOIL in producing optimised solutions for a SOOP at DP1, notwithstanding the similar predictions for the baseline geometry, and the presence of a gap for the CFD data with respect to the experimental ones suggested the opportunity to further push on the validation, by employing different strategies. In the previous phase a wide exploration of the possibilities had been already performed, testing many kinds of grids with different levels of refinement, varying the discretization order, the advection scheme and checking most promising turbulence models, the Spalart-Allmaras (SA) and the $\kappa - \omega$ SST.

Since a gap was still persisting and possibly affecting the 2-D optimisation results, the process was continued, taking into consideration the use of structured grids. These meshes are known to behave well, since the orientation of the elements is more regular and controlled and the user can handle the grid generation by setting in a very precise way the node distribution. An example of O-shaped structured grid is depicted in Figure 6.10. It is visible the regularity of the mesh, as well the presence of local refinements in the wake region and in some parts of the airfoil, in order to keep the orthogonality of

the elements detaching from the walls. Also in this case, the cells number was varied from 43920 to 360220. The results in terms of adherence to the experimental data, instead, showed no remarkable variation with respect to the hybrid mesh, employing the same boundary conditions and models. The RMS residuals were always dominated by those reported in Figure 6.5 for hybrid meshes. This datum indicated from one point the goodness of the grid produced in the first validation process, which appeared to be the best overall as regards the accordance obtained and the numerical behaviour. Actually, structured grids suffered from difficult convergence at high angles of attack, where the layers are no more orthogonal to the stream direction.

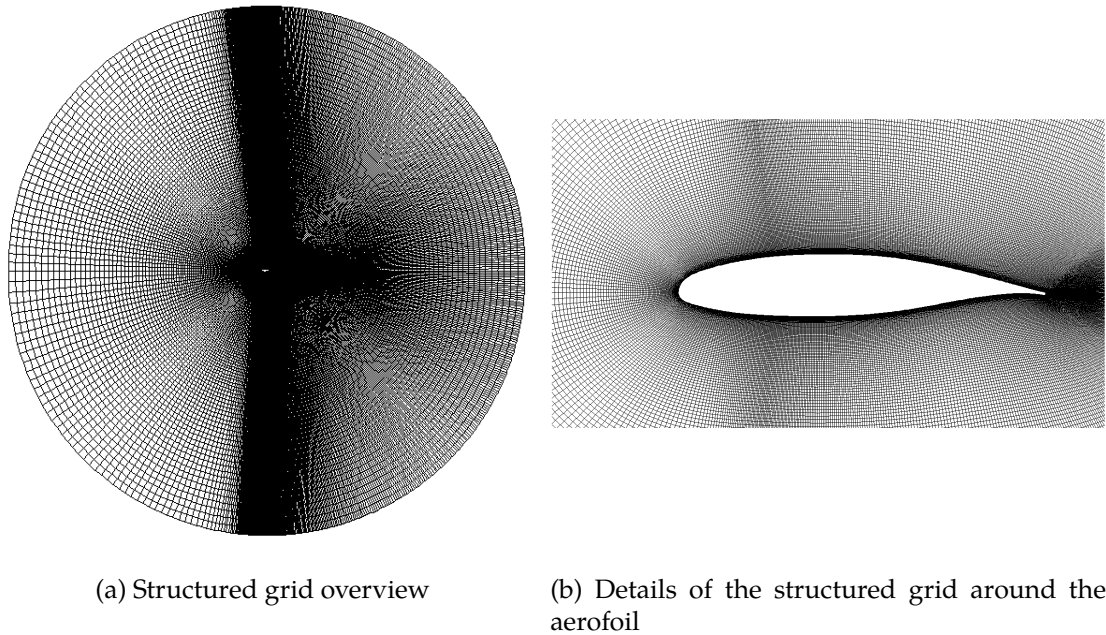


Figure 6.10: Details of a structured grid

From the other point, having excluded a grid issue, the influence of the turbulence model was checked in a final attempt, employing not only the Spalart-Allmaras and the $\kappa - \omega SST$, the former behaving better also for structured grids, but also transition models, accounting for the passage from laminar to turbulent boundary layer. The good behaviour of XFOIL can be, in fact, associated with this capability, lacking in most of the common used turbulence models which consider the turbulence already developed at the inlet.

6.4.1 Transition prediction

Transition from laminar to turbulent boundary layer is a quite complex research field, making use of theoretical backgrounds coupled with empirical correlations to predict the development of the boundary layer and the transition onset. The topic has been

widely studied from the beginning of the modern aerodynamics, but it is still an intriguing issue, getting importance as the development of laminar flow airfoil for both external aerodynamics and turbomachinery is spreading out. An overview of the laminar to turbulence transition may possibly require a whole thesis and of course this goes beyond the scope of the present work. A brief report of the phenomenological aspects of the transition will be here presented, in order to properly deal with the models available in ANSYS Fluent.

Transition phenomenology

The transition itself is an instability phenomenon, in which an initially laminar boundary layer evolves through a turbulent condition, because of some external forcing causing the flow field to jump to a different behaviour. There are many reasons why, from a physical point of view, a laminar flow cannot be any more and a turbulent transition occurs.

The first studied transition is the so-called natural transition, e.g. observed in the classical example of the flat plate in a zero pressure gradient flow. The early research was based on inviscid stability theory, suggesting that a boundary layer becomes unstable in presence of a flexural point in the velocity profile. Indeed, a laminar boundary layer is proved to undergo transition because of the presence of viscous instability waves, also referred to as Tollmien-Schlichting waves. When the freestream turbulence level is low ($<1\%$) instability occurs beyond a critical Reynolds number, at which viscosity destabilizes the waves and they start growing very slowly. The phenomenon evolves very gradually, such that it can be completed also 20 times farther downstream the point where initially transition developed. When the waves become non-linear and an inviscid mechanism has come into play, turbulence spots are born and they enlarge until eventually coalescing into a turbulent boundary layer. A schematic of the natural transition process is shown in Figure 6.11. Even though this first kind of phenomenon sets spontaneously, as a consequence of the propagation of the flow downstream from the leading edge, it is still rather complex.

For the case of a swept wing, Tollmien-Schlichting waves are not the only source of disturbances that can cause transition. When large sweep angles are present a 3D boundary layer develops near the wall, having a significant component in the sweep direction. This is referred to as cross-flow instability, which can cause transition to occur much earlier than would otherwise be the case of a pure Tollmien-Schlichting wave growth.

At higher freestream turbulence levels ($>1\%$), the first and possibly the further stages of the natural transition are bypassed, in a way that turbulent spots are produced within the boundary layer by the influence of the freestream disturbances. This case is known as the by-pass transition, which can actually occur also for turbulence injection into the boundary layer (as in cooling holes for turbine blades) or in presence of rough walls. The instability arises when the skin friction deviates from the laminar value and turbulent spots appear. Upstream, the fluctuations already present in the freestream are thought to cause pressure oscillations, inducing laminar fluctuations that grow similarly to T-S




Figure 6.11: Scheme of the natural transition phenomenon over a flat plate

waves and eventually lead to the formation of spots.

Another kind of transition is referred to as the separation-induced and occurs when a laminar boundary layer separates. In this situation, transition may occur in the shear layer of the separated flow as a consequence of the inviscid instability mechanism. The enhanced mixing caused by turbulence may also induce the flow to reattach, forming a laminar-separation / turbulent-reattachment bubble on the surface. Separation-induced transition happens behind boundary layer trip wires or in presence of strong adverse pressure gradients. The bubble length depends on the transition process within the shear layer and may involve all of the stages listed for natural transition. The freestream turbulence level plays a large role in determining the length of the separation bubble. Traditionally, separation bubbles have been classified as long or short, based on their effect on the pressure distribution around an airfoil. Short bubbles reattach shortly after separation and only have a local effect on the pressure distribution. Long bubbles can completely alter the pressure distribution around an airfoil. Small changes in the Reynolds number can have a huge impact on the length of a bubble, thus complicating the already challenging prediction of their size. Separation induced transition can also occur around the leading edge of an airfoil, if the leading edge radius is small enough. The size of the leading edge separation bubble is a strong function of the freestream turbulence intensity, the leading edge geometry, the angle of attack and to a lesser extent the Reynolds number. The size of the leading edge, moreover, can alter the pressure distribution downstream, causing an adverse pressure gradient which is responsible for a

lack of performance of many airfoils in off-design conditions.

Approaches for predicting transition

The important effects of laminar-turbulent transition is not included in most today's engineering CFD simulations. This is because of many reasons, the first of which is that the phenomenon can arise in several ways, as previously described, and so modelling does not offer the same wide spectrum of CFD-compatible model formulations, that are currently available for turbulent flows. Furthermore, RANS equations do not lend themselves to an easily description of transitional flows, where both linear and non-linear effects are relevant, since Reynolds averaging filters the waves growth. Methods based on linear stability, such us the well-known e^n method by van Ingen [77] and Smith and Gamberoni [68] are not compatible with general-purpose CFD codes. In fact, they require an aprioristic knowledge of the geometry and the grid topology, involving in addition numerous non-local operations that are difficult to implement in highly parallelised solvers, such us computing integral boundary layer quantities or tracking along streamlines.

The e^N method represents, anyway, a widespread tool used for years in engineering to predict the transition with a good confidence. It is based on the local linear stability theory and the assumption of a parallel flow, in order to calculate the growth of the amplitude of the perturbations from the neutral point of the boundary layer to the transition location. The solution requires three steps. The first step is to calculate the laminar velocity and temperature profiles along the walls. This second step consists in computing the local growth rates of the unstable waves for each profile. Finally, in the third step these local growth ratio are integrated for each streamline in order to determine the N factor. Once the disturbance amplitude ratio e^N exceeds the critical factor N the transition is assumed to develop. For isolated airfoils the method has shown very good agreement with experimental data, and it has been shown a certain accuracy also for suction boundary layers [78]. The main problem with it is that the N factor represents the amplification factor of an initial unknown amplitude in the boundary layer, related to the external environment through a generally unknown interaction. For this reason, N is not universal and must be determined from wind tunnel calibration. The method is thus semi-empirical, which is very frequent when dealing with transitional flows.

Some attempts have been made to couple empirical correlations with the RANS decomposition, using the Lagrangian approach to obtain local values of the interest quantities and follow their evolution along the boundary layer. Notwithstanding the wide effort in this area, similar strategies are based on structured grids and are available only for certain geometries, while nowadays the CFD codes use an unstructured and parallel approach, requiring a different solution.

Available models for general turbulence prediction in RANS are the Low-Reynolds models, the Intermittency based models and the Laminar Kinetic Energy models [44]. The concept of intermittency has been derived by observing the oscillatory behaviour of the flow in the transition region, where turbulent spots appear and convect downstream. It is defined as the probability that a certain point is located inside the turbulent region

or the fraction of time in which the boundary layer is turbulent and it is commonly indicated by γ . Within this concept, two approaches are possible: so called Algebraic intermittency models, for which the intermittency is propagated streamwise according to some prescribed function, or intermittency transport models, using a transport model to describe the evolution of this quantity. Whenever γ is invoked, the transition onset location must be established via another model, typically an empirical correlation.

The algebraic intermittency is usually implemented by multiplying the eddy viscosity by the intermittency itself. Upstream of transition γ is zero, so is the eddy viscosity and the boundary layer is laminar. Once the transition onset occurs, the intermittency is slowly increased from zero to one, until the transition is completed. Other models base the intermittency on the ratio of the local boundary layer momentum thickness θ^* Reynolds number Re_θ to the starting and ending transition Reynolds number, derived e.g via empirical correlations.

Another approach relies on a transport equation for γ , where the source terms are modelled to mimic the behaviour of some algebraic equations. In this way the process is modelled also across the boundary layer, having more chance to predict a realistic transition. Such a strategy is used in the $\gamma - Re_\theta$ model by Menter and Langtry [50].

So called Low-Re turbulence models, instead, employ damping functions that are designed to predict the viscous sublayer behaviour and rely on the diffusion of turbulence from freestream to the boundary layer and on the interaction of the former with the model source terms.

Finally, there are the aforementioned Laminar Kinetic Energy Transport models, based on the idea of modelling the laminar fluctuations all along the boundary layer evolution. For example, the transition can occur when the laminar fluctuations reach a certain level. Using a locally formulated transport equation for the laminar kinetic energy, representing the amplitude of non-turbulent streamwise fluctuations in the pre-transitional boundary layer, and coupling this laminar kinetic energy equation to a turbulent kinetic energy (κ) and a turbulent eddy frequency (ω) equation, it is possible to predict the onset of the transition, without any intervention by the user and relying only on local variables. Such an approach is employed in the $k - kl - \omega$ Transition by Walters and Cokljat [36].

Fluent models for transitional flows

In ANSYS Fluent, the available transition models are the so called $k - kl - \omega$ Transition by Walters and Cokljat, the Transition SST or $\gamma - Re_\theta$ by Langtry and Menter, and the γ transition model. Each of the above models is based on different formulations to describe the turbulent field and how it affects the internal solution. For example, the first and the third are three-equation models, while the second is a four-equations one. All of them rely on the $k - \omega$ interpretation of the RANS equation, with the Shear Stress Transport switching between $\kappa - \omega$ and $\kappa - \epsilon$ for the near-wall and far-field, respectively.

$k - kl - \omega$ Transition model by Walters and Cokljat [36] is a three-equation model, based on the $\kappa - \omega$ framework and invoking a further transport equation to predict

the magnitude of low-frequency velocity fluctuations in the pre-transitional boundary, that have been identified as the precursors to transition. The transition process itself is represented by an energy transfer from the laminar kinetic energy κ_L to the turbulent kinetic energy κ . Conceptually, the latter is assumed to represent the magnitude of the fluctuations that display the characteristics of fully turbulent flow, such as strong three-dimensionality, multiple length and time-scales, energy cascading and significant viscous dissipation. The former, instead, accounts for the low-frequency fluctuations promoting the transition upstream its onset. The initiation of the transition process in the model is based on local single-point flow conditions.

The Transition SST or $\gamma - \text{Re}_\theta$ by Langtry and Menter [50] is a four equation model, in which a transport equation for the intermittency and for the momentum thickness Reynolds number are coupled with the $\kappa - \omega$ framework. The intermittency is used to turn on the production term of the turbulent kinetic energy, downstream of the transition point in the boundary layer. This differs from the typical usage of intermittency, where it is often used to modify the eddy viscosity. The authors argue that, from a modeling standpoint, that approach has the advantage of capturing the effect of large free-stream turbulence levels on laminar boundary layers and the related increase in the laminar skin friction and heat transfer. The fourth transport equation for Re_θ is necessary to capture the non-local influence of the turbulent intensity, which changes due to the decay of the turbulent kinetic energy in the freestream, as well as in changes in the freestream velocity outside the boundary layer. This additional equation is an essential part of the model, in that it links the empirical correlation to the onset criteria in the intermittency equation. This additional transport equation is an essential part of the model as it ties the empirical correlation to the onset criteria in the intermittency equation, giving more chance to incorporate several formulations for different boundary conditions, in order to obtain a general purpose and even customizable tool. The coupling with the $\kappa - \omega$ SST is performed by modifying the production and destruction terms of the κ equation of the original model.

The γ transition model is a proprietary formulation of ANSYS, representing a possible enhancement of the the $\gamma - \text{Re}_\theta$. Unlike the latter, it uses only a transport equation for γ , avoiding the need for a fourth equation, reducing the computational effort, providing a capability for cross-flow instability. It is again a strictly local variables based model, in which the source terms for the intermittency equation are based on empirical correlations to trigger the transition onset.

6.5 Comparison of results for various turbulence models

After briefly introducing the intriguing and wide issue of transition, in this section the results obtained employing different turbulence models embedding transition prediction are presented, comparing them to those derived using more standard fully turbulent models, like Spalart-Allmaras and $\kappa - \omega$. As previously explained, in the first stage of the validation activity a study on the mesh sensitivity and on the turbulence modelling suggested that the Spalart-Allmaras turbulence model solved in an hy-

brid mesh, with a structured boundary layer and an unstructured external domain, can provide the best outcome overall, compared to other solutions employing structured meshes or $\kappa - \omega / \kappa - \epsilon$ formulations. Despite the acceptable level of accuracy between experimental data and predicted ones, a very different optimised individual was found by the optimiser for a SOOP, minimising the drag coefficient at an angle of attack of 15 [deg], using XFOIL, which enables transition prediction via eN model, and fully turbulent CFD. The detailed description of this activity is documented in chapter 7. As a consequence of such a discordance, the validation process for the CFD was continued, by examining different turbulence closures, able to deal with transitional flows. It is then interesting to compare the predictions of all the models tested, to see if it is possible to fill the gap between CFD data and the experimental data, present in Figure 6.3, Figure 6.4, Figure 6.8. All the plots were derived using the same hybrid grid, whose details are in Figure 6.1.




Figure 6.12: $C_l - \alpha$ polar comparison between various turbulence models, computed at $M = 0.20$, $Re = 6.2 \times 10^6$

Figures 6.12, 6.13, 6.14 illustrate the polars computed with several models.

The Lift versus the angle of attack is depicted in Figure 6.12. It is immediately noticeable how XFOIL and the $\kappa - \omega$ SST overpredicts and underpredicts the lift, respectively. For the other models, the points were closer and the difference can be appreciated in the zoomed figure, for $\alpha = 10 \div 18$ [deg]. The Transition SST or $\gamma - Re_\theta$ was closer to the experimental C_l , except at 18 [deg], where Spalart-Allmaras behaves better. A wider gap instead was present for the $\kappa - \kappa_l - \omega$.

The Lift-Drag polar is shown in Figure 6.13. Here a more various behaviour can be observed among the models. In general, it was not found a formulation able to capture with excellent accuracy the experimental chart over the whole range of angles of attack.




Figure 6.13: Lift-Drag polar comparison between various turbulence models, computed at $M = 0.20$, $Re = 6.2 \times 10^6$

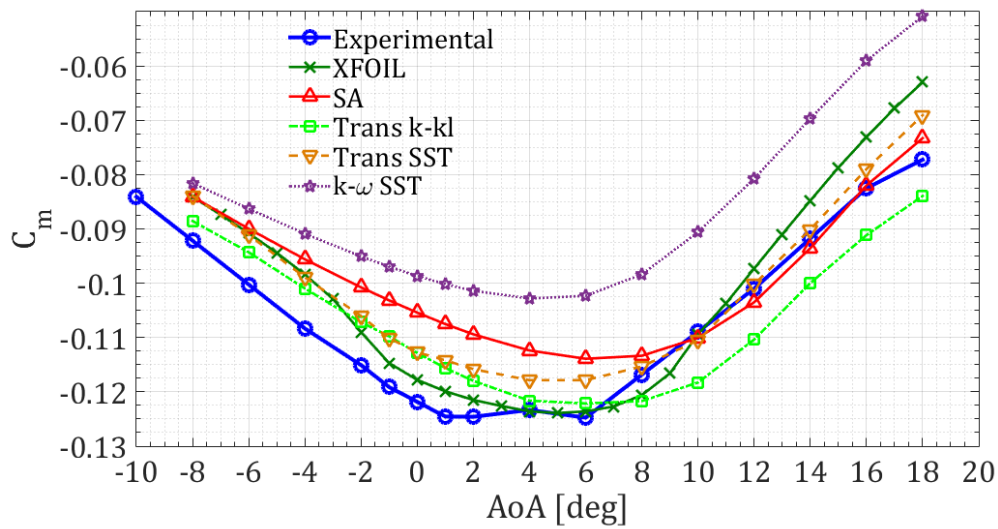


Figure 6.14: $C_m - \alpha$ polar comparison between various turbulence models, computed at $M = 0.20$, $Re = 6.2 \times 10^6$

The Transition $\kappa - \kappa_l$ was very good between -4 and 4 [deg], but lost accordance for higher α . The Menter's Transition SST provided a better fidelity for $\alpha = 4 \div 10$ [deg], while near the maximum lift condition, occurring at 18 [deg], all the models tended to overpredict the drag, except for XFOIL. It is worth remarking how transition models were all able to capture the laminar bucket, where the drag shows a little dependency by the lift coefficient, which was not true for fully turbulent formulations.

An increased level of difficulty is dwelling in predicting the moment coefficient. Particularly, it resulted easier to match the experimental chart at higher angles of attack, where the C_m versus α relationship is almost linear, as visible in Figure 6.14. Here, SA confirmed to behave very good, and such was the Transition SST. For the other models, especially for the standard SST, the accuracy was lower. Instead, for a zero incidence condition XFOIL and the $\kappa - \kappa_l$ behaved better.

The comparison of all the models tested led to some synthetic conclusions:

- No one model is able to provide an optimal accuracy over the entire extent of the charts;
- The lift coefficient is well captured by all the models, with the Transition SST having a better match and XFOIL and $\kappa - \omega SST$ a worse;
- The drag coefficient can be correctly predicted, by choosing the appropriate model depending on the boundary conditions (AoA);
- The moment coefficient is captured with better accuracy at high angles of attack, particularly by SA and Transition SST;
- The standard $\kappa - \omega SST$ performs worse, with respect of the other models;
- The Transition SST and the Transition $\kappa - \kappa_l$ are the most promising models.

On the basis of these observations, it was further investigated the $\gamma - Re_\theta$ model, which appeared to be more accurate at higher angles of attack, with which the optimisation should deal. Specifically, the influence of the grid size and type and the boundary conditions was assessed, in order to find the best configuration to employ in an optimisation loop. Some of the outcomes are reported in Figure 6.15, Figure 6.16, Figure 6.17; the same considerations as before can be made here, with none of the meshes able to guarantee a full replication of the charts. Furthermore, the response of the model to a variation in the cell size and number was different in different portions of the polars, this resulting in the choice of a range to be better captured, penalising the prediction outside that range. Again, since the interest was focused toward high incidence angles, the final choice was the mesh named 'Hyb-finer', with about 60'000 elements. For details of this mesh, refer to Figure 6.1 and Figure 6.2, since there is no remarkable difference, except for a slightly denser distribution of points along the profile sides. Structured grids showed a worse accuracy at an higher computational cost and so were discarded.

The boundary conditions employed for this second model are summarized in Table 6.2

| | |
|------------------------------------|--|
| Cells | 6E+04 |
| Boundary Conditions | Far-field (Mach, Pressure, Velocity) |
| Inlet turbulence quantities | $\gamma = 0.5$, Intensity=0.5%, $\mu/\mu_t = 1$ |
| Boundary Layer Rows | 50 with growth ratio < 1.11 |
| Turbulence Model | Transition SST |
| Pressure-Velocity Coupling | Non-segregated solver |
| Gradient approximation | Second Order Upwind |

Table 6.2: Summary of CFD settings for the Transition SST model




Figure 6.15: $C_l - \alpha$ polar comparison for different meshes, employing the $\gamma - Re_\theta$ model, computed at $M = 0.20$, $Re = 6.2 \times 10^6$




Figure 6.16: Lift-Drage polar comparison for different meshes, employing the $\gamma - Re_\theta$ model, computed at $M = 0.20$, $Re = 6.2 \times 10^6$




Figure 6.17: $C_m - \alpha$ polar comparison for different meshes, employing the $\gamma - Re_\theta$ model, computed at $M = 0.20$, $Re = 6.2 \times 10^6$

Chapter 7

Two-dimensional Optimisation of a Morphing Leading Edge Airfoil

In a qualitative way the optimisation procedure consists in a series of operations allowing to change a configuration in order to attain a desired performance, bound to the respect of some kind of constraints. In engineering application the procedure is typically recursive, since some explicit models of the problem are missing and a straightforward formalisation is not available. The procedure is then configured as a loop, invoked several times, within which are present various “subsystems”, each one with a specific function. The optimisation properly said is a part of this loop and is a mathematical problem, whose formalisation will be given later. Translating an operative challenge into a mathematical model is a necessary step for the optimisation and its result will be obviously greatly affected by the hypothesis and approximations embedded in them. High accuracy models can result in impracticably expensive computational resources and a balanced formulation with a moderate number of assumptions and theories validated for the case in exam is always desirable.

7.1 Mathematical formulation of the optimisation problem

A general constrained multi-objective optimisation problem (MOOP) consists in:

$$\begin{aligned} &\text{minimize} && f_i(\mathbf{x}), i = 1, \dots, n, \mathbf{x} \in \Omega \\ &\text{such that:} && G_j(\mathbf{x}) = 0, j = 1, \dots, l \\ &&& H_k(\mathbf{x}) \leq 0, k = 1, \dots, m \end{aligned}$$

where:

- $\mathbf{x} = (x_1, \dots, x_n)$ is the vector of the decision variables;
- $f_i : \Omega \subseteq \mathbb{R}^n \rightarrow \mathbb{R}$ is the i -th objective (or fitness) function;
- $G_j(\mathbf{x})$ is the j -th equality constraint function;

- $H_k(\mathbf{x})$ is the k-th inequality constraint function;

If $i = 1$ the problem is single objective. Note that looking for the minimum of f is not restrictive, since each maximisation problem can turn into a minimisation one by inverting the objective.

7.2 Optimisation algorithm

A variety of optimisation strategies are available in literature to solve the minimization problem. The choice depends upon various aspects, *in primis* the level of knowledge of the objective function and the fitness landscape. In the problem faced neither the relationship between the decision variables and the objectives is known with an explicit mathematical formulation, nor the feasible multidimensional region, where local minima are likely to be present. Furthermore, non-linearity is present both in the objective function, due to aerodynamic effects such as stall and separation, and in the constraints, including geometrical and structural non linear quantities e.g. curvature. These factors complicate the use of gradient-based methods and suggest the employment of a stochastic search of the design space.

7.2.1 Genetic algorithms

Genetic algorithms (GA) are direct, parallel, stochastic methods for global search and optimization, which imitate the evolution of the living beings. Unlike other methods, GAs treat a set of candidate solutions, forming a so called population. Following bio-inspired evolution mechanism, the population is evolved many times until a convergence criterion is reached. A brief summary of terminology used is here reported.

- *gene*: a representation of a scalar decision variable according to an encoding scheme (e.g binary);
- *chromosome*: an encoding of a vector in the domain Ω consisting in a set of genes and individuating a possible solution; it is a synonym of individual;
- *population*: a set of N chromosomes which will be evolved by the algorithm;
- *generation*: each iteration of the algorithm involving the evolution of a population;
- *selection*: the process of individuating n best candidates among a population, whose information will be transmitted to the next generation;
- *recombination*: a process where couples of chromosomes (parents) are recombined to form new ones (children) whose genes are the same of the parents but in a different order; AKA crossover;
- *mutation*: a random change of some genes of a chromosome, due to transcript error in nature, used to promote diversity in the population.

A general scheme of a GA is shown in Figure 7.1. The process of selection, crossover and mutation occurs sequentially and there are different methods to perform each of them.

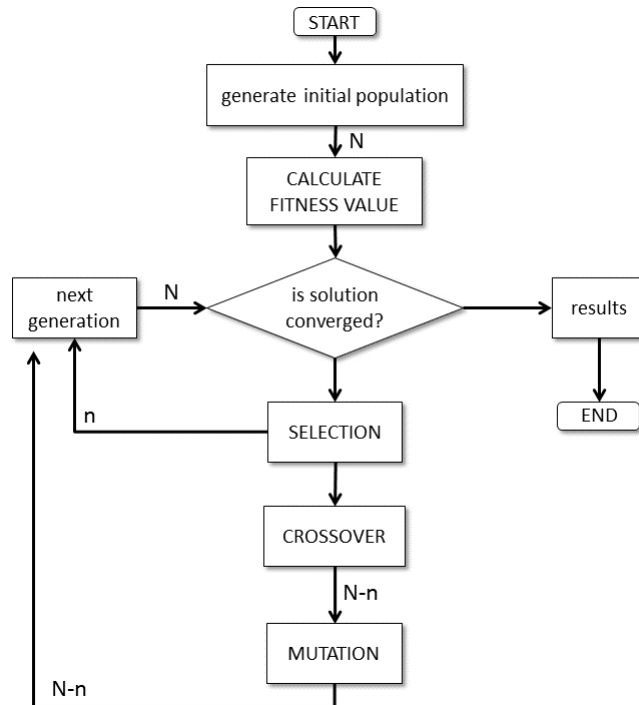


Figure 7.1: A general scheme for GA

Selection

As stated above, selection is a process in which the individuals, which will create the offspring population are chosen, according to some criterion. The selection has two main purposes:

1. to choose the most perspective individuals, which will take part in the generation of next population or will be directly copied (elitism);
2. to give an opportunity to individuals with comparatively bad value of the fitness functions to take part in the creation of the next generation. This allows us to preserve the global character of the search process and not to allow a single individual to dominate the population, bringing to local extrema.

For a single objective optimisation, a common method of selection is the Roulette-wheel. Qualitatively, individuals with an higher performance index, directly related to their fitness value, which indeed has to be minimised, are given a larger probability of being selected on a spin of the roulette wheel. Couples of individuals are chosen to form children in the next recombination process.

In a multi-objective problem, the concept of a dominating solution and Pareto front must be invoked.

Definition 7.1. A vector \mathbf{v} is said to dominate a vector \mathbf{u} if

$$\forall i = 1, \dots, n \ f_i(\mathbf{v}) \leq f_i(\mathbf{u}) \quad \text{and} \quad \exists j = 1, \dots, n \ | \ f_j(\mathbf{v}) < f_j(\mathbf{u})$$

In other words, \mathbf{v} is giving one smaller value than \mathbf{u} for at least one fitness function.

Definition 7.2. A vector $\mathbf{x} \in \Omega$ is a Pareto optimal solution if and only if there is no vector $\mathbf{y} \in \Omega$ for which $(f_1(\mathbf{y}), \dots, f_n(\mathbf{y}))$ dominates $(f_1(\mathbf{x}), \dots, f_n(\mathbf{x}))$.

The set of Pareto optimal solutions forms a Pareto front. In Figure 7.2 a Pareto front for two objective functions f_1 and f_2 is indicated by P , while Y denotes the feasible design space, f_1^* and f_2^* are the optimal solutions for the corresponding fitness function. In a multi-objective optimisation the main goal is not only to find a set of solutions as close as possible to the Pareto optimal front, but also to find a set of solutions as diverse as possible, to cover the entire front.




Figure 7.2: An example of Pareto surface

When selecting individuals among a population, a non-domination criterion can be applied. In non-dominated sorting methods the population is sorted according to an ascending level of non-domination. Once the first set of non-dominated solution is found, it receives a rank 1 and is temporarily discarded. The remaining individuals are sorted again and the non-dominant receives a rank 2 and are discarded. The process is repeated until all the individuals are ranked. At this point a further parameter can be assigned before operating selection itself, to take into account the distance between the individuals belonging to the same non-dominated set. The selection operator is then applied basing on both the non-domination rank and the distance. Roulette wheel or tournament selection are popular. The latter consists in choosing at random pairs of individuals and extracting a random number $R^+ \in [0; 1]$. if $R^+ > r$, $0.5 < r \leq 1$ than the fitter of the two is selected, otherwise the less fit. r is a user-set number.

When elitism is employed, some of the fittest individuals within a population are selected to propagate in the next generation. This ensures that the “good genes” are conserved and the solution proceeds monotonically towards a minimum. The rate of elite chromosomes preserved in the generation can be set according to some indexes, but care must be taken to avoid excessive polarisation of the population by the fittest members, which can lead to a local optimum.

Recombination

The recombination allows to propagate the generation and evolve towards the optimum. As previously mentioned, recombination is known as crossover, because of the mechanism of reproduction of a pair of chromosome. In a single point crossover, a position is chosen for the recombination, then two children are produced, by cutting their parents chromosomes in the position assigned and reassembling the individuals using pieces of each parent. In scattered crossover a random binary string of the same chromosome length is generated and then a gene from the first chromosome is chosen if the corresponding bit in the random string is one, or a gene from the second chromosome is chosen if the bit is 0. This method is valid only in case of non linear constraints. In elitist methods the fraction of the new generation, other than elite children, produced by crossover can be selected.

Mutation

Mutation occurs in nature because of transcription errors in the chromosomes replication. Although it can lead to devastating effects and produce meaningless genes, it is part of a survive strategy to promote diversity among a species. Similarly, in GA mutation is used to avoid polarisation in a population and guarantee a wide search of the design space. Mutation simply consists in a bit flip when a binary encoding is used. In constrained problems other strategies are used to ensure not to violate the bounds and the constraints.

7.2.2 NSGA-II algorithm

The Multi-Objective Optimisation Problem (MOOP) has been solved in Matlab using the Optimisation Toolbox, which provides a quite wide range of optimisation tools, for both Single Objective Optimisation Problems (SOOP) and MOOP, with linear or non-linear constraints and bounds. In particular, for multi-objective optimisation the algorithm named NSGA-II by Deb [21] . It is a non-dominated sorting elitist MO genetic algorithm, whose principles and basic steps will be presented below.

NSGA-II uses the concept of non-dominated sorting previously reported, together with an explicit diversity-preserving mechanism, based on crowding distance. First of all, an offspring population Q_t is created starting from the parent population P_t . Then the two populations are combined together to form an R_t population of size $2N$, being N the population size of P_t and Q_t . A non dominated sorting is thus used to classify the

entire population and the new one is formed picking individuals from different non-dominated fronts, until N new elements are chosen. Individuals picked from the last non-dominated front contributing to form the new population are chosen according to a niching strategy, in order to select elements residing in the least crowded region of that front. The synthetic NSGA-II procedure reported in [21] is:

- Step 1** Combine parent and offspring populations and create $R_t = P_t \cup R_t$. Perform non-dominated sorting and form fronts $F_i, i = 1, 2, \dots$
- Step 2** Set new population $P_{t+1} = \emptyset$ and counter $i = 0$. While $\text{size}(P_{t+1}) + \text{size}(F_i) < N$, do $P_{t+1} = P_{t+1} \cup F_i$ and $i = i + 1$.
- Step 3** Apply the crowding-sort procedure to the current front and include the most widely spaced $N - \text{size}(F_i)$ solutions by using crowding distance values in the sorted F_i to P_{t+1} .
- Step 4** Create offspring population Q_{t+1} from P_{t+1} by using the crowded tournament selection, crossover and mutation operators.

The last mentioned *crowded tournament operator* is a pair tournament operator according to which the solution i wins over solution j if i has a better rank, $r_i < r_j$ or, if $r_i = r_j$, if $d_i > d_j$.

Crowding distance is a parameter used to measure the distance of the individuals in the fitness space, in order to promote the ones less crowded in the offspring population, and to ensure a wide coverage of the Pareto frontier. The crowding distance assignment procedure consists in:

- Step 1** Assign $d_i = 0$ for each solution $i = 1, \dots, l$ in F .
- Step 2** For each objective function $f_m, m = 1, \dots, M$, sort the set in descending order of f_m .
- Step 3** For $m = 1, \dots, M$, assign $d_{I_1^m} = d_{I_l^m} = \infty$ and for all other individuals with $j = 2, \dots, l - 1$ assign:

$$d_{I_j^m} = d_{I_j^m} + \frac{f_{I_{j+1}^m} - f_{I_{j-1}^m}}{f_m^{\max} - f_m^{\min}}$$

with I_j^m denoting the index of the j -th member in the list sorted for f_m .

The crowding distance, in other words, measures the perimeter of the cuboid whose opposite vertices are the two nearest solutions to the considered one. An advantage of this approach is that solutions compete each other also in terms of how dense they are in the fitness space, thus diversity is explicitly considered when forming the offspring population. Furthermore, elitism is ensured by non-dominated sorting of both parents and offspring and inclusion of non-dominated fronts in the new set P_{t+1} . A disadvantage underlined by its author dwells in the possibility of a convergence failure if the number of non-dominated solutions in the first ranked set is less than the population size N .

The current elitist algorithm is implemented in Matlab with a slightly different formulation, which responds to the problem of controlling elitism. This problem is related to an important issue which always appears when employing GAs, the concept of exploration versus exploitation. In NSGA-II elite solutions are emphasized by the crowded tournament selection operator and when individuals from the next non-dominated front are picked to form the offspring population. Although in this procedure diversity along the Pareto frontier is promoted by the crowded tournament, lateral diversity, i.e. the diversity along non-dominated fronts of different ranks, is rapidly lost. In order to avoid this, a strategy could be to limit the number of individuals belonging to the first non-dominated front and try to distribute individuals through higher ranked fronts according to a defined distribution. The number of individuals in the best non-dominated front is a user-defined, problem dependent parameter called *Pareto Fraction*. Different values of it give rise to different Pareto frontiers, There should exist a value for which the final front converges as closest as possible to the true, seldom known, front.

For more issues related to NSGA-II the reader is referred to [21] for theoretical aspects and to Matlab Documentations or also [65] for the implementation procedure.

7.3 Design variables

As a consequence of the parametrisation method employed, 5 or 7 decision variables (DV) were individuated. It has to be kept in mind how the CST method describes the airfoil, in order to properly understand what are the degrees of freedom available (see chapter 5). Each component shape function expresses its influence mostly over a region near to the shape component peak; this is because of the graph of Bernstein polynomials. Since for $n = 10$, the chosen degree of the representation, the first three component shape functions, controlled by coefficients A_0, A_1, A_2 , have their peak within the interval $[0; C_{fs} = 0.25]$ these are the available degrees of freedom, to which the ordinate of the LE Z_{LE} can be added. Given that, being each side of the airfoil represented by an explicit function, the total number of degrees of freedom is 7: $[A_0, A_1, A_2]_{up}$, $[A_0, A_1, A_2]_{low}$ and Z_{LE} . By varying them it is possible to obtain a family of curves, all maintaining the same shape in the fixed region, as depicted for example in Figure 5.7.

In order to obtain a constant length parametrisation, instead, as explained in section 5.4, it is possible to reduce the number of degrees of freedom from 7 to 5, being A_2 determined in order to keep the length of the airfoil section in the morphing region constant, for both upper and lower side. In this way the total number of DV is 5: $[A_0, A_1]_{up}$, $[A_0, A_1]_{low}$ and Z_{LE} . Except for the first unconstrained run, in all other cases the constant length parametrisation was employed and thus the DVs were those 5 here presented.

7.4 Single objective optimisation using XFOIL

A Single Objective Optimisation (SOO) was firstly performed to enquire the decision variables space, the fitness landscape, the effects of constraints and genetic oper-

ators. Furthermore, an initial SOO allows to check the consistence and correctness of the problem formulation and to verify how the optimiser set-up afflicts the results. For these purposes, a first optimisation with 7 DVs (no constant length parametrisation) was carried out referring to the reference 2D configuration at DP1 (see Table 4.1). XFOIL was used to evaluate the fitness C_d . The optimiser was Matlab[®] *ga* function with default options and a population size of 50 individuals. Constraints were only bounds on DVs to avoid very distorted shape slowing down the run. The outcome after 50 generations is depicted in Figure 7.3 and the quantitative indexes are reported in following Table 7.1. Despite such a shape variation results infeasible in reality, the run was intended to check the correct operation of all subsystems in the loop and give a sort of benchmark on the best allowable result for DP1.




Figure 7.3: Result for a SOO for minimum C_d at DP1, with only bounds on the DVs as constraints

| Index | Baseline | Optimised | Relative variation [%] |
|-------|----------|-----------|------------------------|
| C_d | 0.03122 | 0.01401 | -55.12 |
| C_l | 1.9416 | 2.0967 | 7.99 |
| L/D | 62.191 | 149.657 | 140.64 |

Table 7.1: Comparison between baseline and optimised aerodynamic coefficients for the first SOO at DP1

7.4.1 Single objective optimisation at DP1

In a second step, constraints were added and constant length parametrisation employed. The problem can be formalized as:

$$\begin{aligned} &\text{minimize} && C_d(\mathbf{x}), \mathbf{x} \in \Omega \\ &\text{such that:} && |L_m - L_0|/L_0 \leq 1e - 6 \\ &&& |A_m - A_0|/A_0 \leq 0.05 \\ &&& C_{l,b} - C_l \leq 0 \\ &&& (C_{m,b} - C_m)/C_{m,b} \leq 0.1 \\ &&& \Omega = \mathbf{x} \in \mathbb{R}^5 | \mathbf{Lb} \leq \mathbf{x} \leq \mathbf{Ub} \end{aligned}$$

Upper and lower bounds on DVs were set in order to avoid the generation of very distorted shapes, with no possibility of respecting the constraints. A bounding box on the morphing region was built and all the geometries exceeding that were simply discarded and no aerodynamic computation performed. Finally, the maximum error between fixed control points, from which the shape should pass outside the morphing region, was monitored and included as a constraint, with a tolerance of 5e-04. This constraint showed to be violated only for geometries infeasible also at least with respect to the other geometrical constraints, so it did not appear to be strict and represents only a further indication on how good the shape generated is.

Influence of the genetic operators

Within the above formulation, several runs were performed still using XFOIL in order to investigate the presence of local optima and to evaluate the effect of different initial populations and genetic operators. Convergence to the same point was observed regardless the choice of operators and initial population; this allows to state that a global optimum is likely to exist and was reached by the optimisation. Qualitatively, it appeared a clear trend in the search direction who led to the aforementioned point. The results of this analysis stage indicated that some genetic operators can produce a rapid vanishing of the population diversity, if the initial one was low. The initialization of the population was carried out by firstly generating a randomly distributed set within the domain, delimited only by bounds on DVs. After that, two approaches were followed: retaining only points without constraints violation or using a wider tolerance for constraints in order to have a wider initial population. The former approach, indeed, showed that the feasible domain is quite tight, thus the absolute average distance between individuals is low. The latter, instead, allows to have a wider distributed population including also non-feasible individuals. Different initial populations can be obtained if the number of the points produced by the above process, whatever the constraint tolerance is, is greater than the population size. By this method, the influence of the starting set on the ultimate result was analysed, showing a global dependency between initial diversity and

genetic operators used. Though for certain combinations the diversity rapidly vanished, leading to a complete dominance of few individuals, the same optimum was reached.

For example, the following Figure 7.4 illustrates the variation of the mean fitness of the population for two different mutation operators, starting from the same initial population, whose elements are all feasible and initial distance between them is low, because of the tight feasible region individuated. All other operators and parameters were the same. The *Gaussian* mutation operator adds a random number taken from a Gaussian distribution with mean 0 to each entry of the parent vector. The standard deviation of this distribution is determined by the parameters *Scale* and *Shrink*. The former determines the standard deviation of the first population and was set as 1. The second controls how the standard deviation evolves through the optimisation:

$$\sigma_k = (\sigma_{k-1} - \textit{Shrink} \times k/\textit{Max Gen.})$$

and was set to be 0.9. The *Uniform* mutation operator is a two-step process. First, the algorithm selects a fraction of the vector entries of an individual for mutation, where each entry has a probability *Rate* of being mutated. In the second step, the algorithm replaces each selected entry by a random number selected uniformly from the range for that entry. *Rate* was set to be 0.01. With this latter operator the mean rapidly converged to the best fitness of the current population, meaning that diversity had vanished. It is worthwhile to underline that both operators are not recommended for problems with bounds or linear constraints, because they can produce individuals violating them. Notwithstanding, for the runs considered Uniform mutation showed not to be able to preserve diversity, leading indeed to a faster stop in 42 generations because of the stall in fitness variation over last 15 iterations. Gaussian Mutation kept a higher diversity and stopping criteria was that maximum generations (50) was reached. The final point was, indeed, almost the same regardless the choice of operators.




Figure 7.4: Evolution of mean fitness with different mutation operators

As regards the influence of the crossover operator, various tests were conducted and it has been observed that for all other conditions kept fixed, different crossover operators

gave almost the same trend in the evolution of mean and best fitness and average distance between individuals. Another parameter investigated was the crossover fraction, the number of children in a population generated by crossover. The crossover fraction used in those runs was 0.65:0.85. For higher crossover fractions the convergence was more rapid in the first iterations, while proceeding slowly towards the minimum; for lower crossover fraction convergence was more linear towards the same minimum, as expected. This is reported in Figure 7.5. For all the three cases the initial population was the same, with a wider tolerance on constraints violation and a good distance between individuals. The crossover operator was the intermediate, according to which children are created by taking a weighted average of the parents:

$$\text{child} = \text{parent1} + \text{rand} \times \text{Ratio} \times (\text{parent1} - \text{parent2})$$

rand is a random number between 0 and 1. If ratio is in the same range, children lies on the line between the two parents, while if *Ratio* is greater than one the child lies on the same line, but also outside the connecting segment. For the cases illustrated *Ratio* was set to be 1.5. The mutation operator chosen was the *Adaptive Feasible*, which randomly generates directions that are adaptive with respect to the last successful or unsuccessful generation. The mutation chooses a direction and a step length that satisfies bounds and linear constraints. Elite count, the number of elite children preserved through each generation, was set to 2 and the population size was 50.

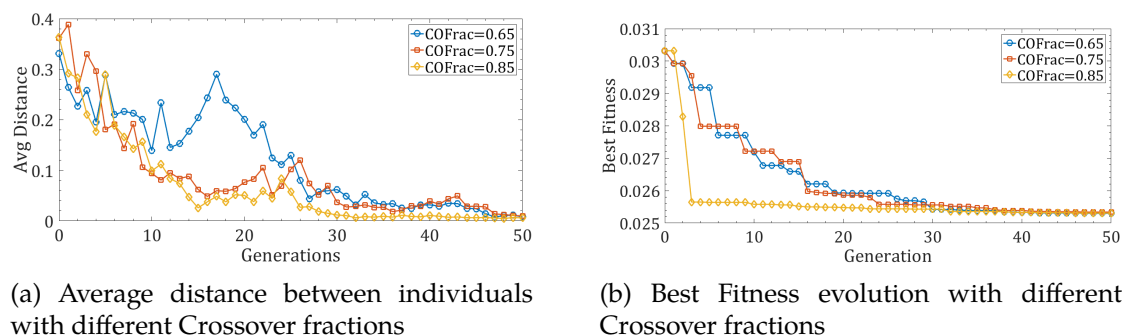


Figure 7.5: Influence of crossover fraction

Even for these cases, the same optimum was reached regardless the configuration adopted.

Best result for SOO at DP1 using XFOIL

After reporting some of the investigations performed in order to better understand how the genetic operators influence the evolution for the particular case examined, here it is presented the best result obtained. Again, it is recalled that practically all the runs produced the same final point. Figure 7.6 depicts the baseline and the optimised airfoil together with the calculated C_p distribution for the optimised profile, while Table 7.2 reports the variations in the aerodynamic coefficients.




Figure 7.6: Optimised airfoil at DP1 with XFOIL and C_p distribution (dashed = baseline)

| Index | Baseline | Optimised | Relative variation [%] |
|-------|----------|-----------|------------------------|
| C_d | 0.03122 | 0.02530 | -23.4 |
| C_l | 1.9416 | 2.0582 | 6.00 |
| L/D | 62.191 | 81.340 | 30.79 |

Table 7.2: Comparison between the aerodynamic coefficients of the baseline and optimised airfoil with XFOIL at DP1

As it can be noticed, the results for aerodynamic coefficients variations are less accentuated than those of the first unconstrained optimisation, though still promising.

7.5 Multi-objective optimisation using XFOIL

After successfully testing the algorithm for a SO case, the NSGA-II was applied to a multi-objective (multi-point) optimisation. In this phase too, thanks to the affordable computational cost of XFOIL, various runs were performed with the aim of evaluating the change in Pareto front with different settings. The first problem solved was:

$$\begin{aligned}
 &\text{minimize} && C_d(\mathbf{x})@DP1, \mathbf{x} \in \Omega \\
 &\text{minimize} && C_d(\mathbf{x})@DP2, \mathbf{x} \in \Omega \\
 &\text{such that:} && |L_m - L_0|/L_0 \leq 1e - 6 \\
 &&& |A_m - A_0|/A_0 \leq 0.05 \\
 &&& C_{l,b} - C_l \leq 0 \\
 &&& (C_{m,b} - C_m)/C_{m,b} \leq 0.1 \\
 &&& \Omega = \mathbf{x} \in \mathbb{R}^5 | \mathbf{Lb} \leq \mathbf{x} \leq \mathbf{Ub}
 \end{aligned}$$

It was a multi-point optimisation with the same objective of the SO case, minimised for two different conditions (see Table 4.1). Two population sizes were used, 50 and 100 individuals, both pre-initialized with a random search to obtain points within the feasible region. Another parameter investigated was the Pareto fraction and its effects on the final frontier.

Figure 7.7 illustrates three different Pareto fronts calculated for three different Pareto fractions. The mutation operator was the Adaptive Feasible, the crossover Intermediate with a Ratio of 1.5 and the population size 50. The Pareto Front slightly changed among the cases, and an higher difference in few points was noticeable for the first case with the lowest Pareto Fraction.

It is also interesting to plot how the design variables were spread through the Pareto fronts. This is helpful in discovering if a correlation exist and to verify that the front is thoroughly caught. Figure 7.8 shows in ordinate the value of the first objective function and in abscissa the component of the DV vector. It appears a possible correlation for A_{0up} and A_{1up} , with their values quite evenly distributed along the frontier forming a curve. For A_{0ow} and A_{1ow} , points are collocated in a narrow region, approximatively vertical; this is true particularly for A_{1ow} . The ordinate of the leading edge, instead, shows a wider variation and points are scattered along the front.

As regards the influence of population size (PS), Figure 7.9 depicts fronts for a population of 50 and 100 individuals. The Pareto Fraction (PF) varies in order to compare fronts with the same number of individuals. Here again, the major difference between the two sets occurs at one edge. As a general statement, the highest the population size, the better; in a practical approach, 50 individuals are quite many when the fitness evaluation is made through a CFD, also compared to the number of DVs. From the analysis




Figure 7.7: Three Pareto Fronts for a different *Pareto Fraction*

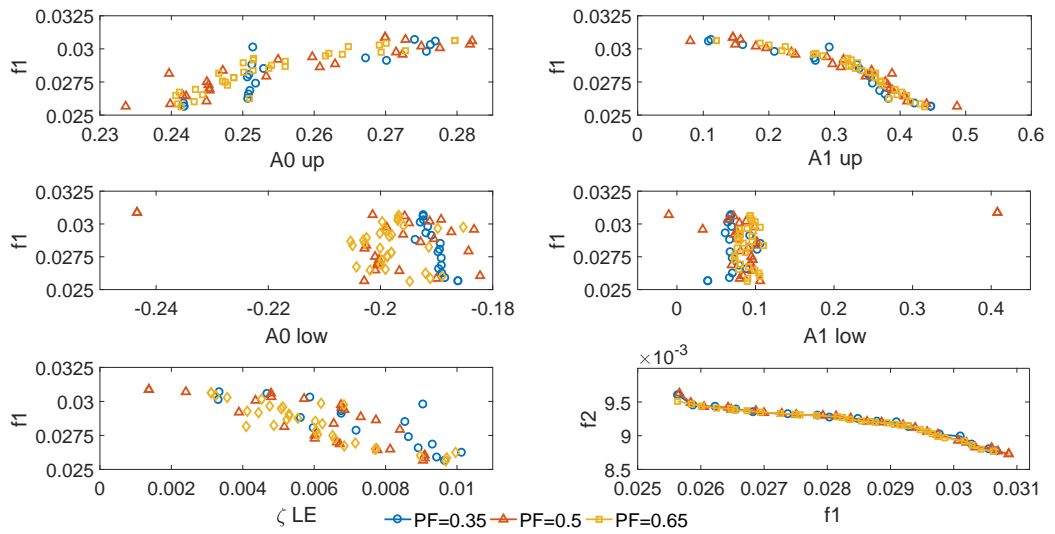


Figure 7.8: Spread of the DVs along the Pareto Fronts for three PF

reported, similar results were obtained by keeping the same number of individuals in the first non-dominated front. Obviously, for an higher PF the lateral diversity will be lower, with an higher risk of stalling against a local Pareto solution. Since all the fronts obtained are comparable, this occurrence appears improbable.




Figure 7.9: Pareto Fronts for different population sizes

A last observation on the graphs reported deals with the improvement brought by the optimisation, with respect to the baseline configuration. The original point in the fitness space is indicated as the baseline in the figures. The best improvement in the Pareto solutions for the first objective is -15.99%, giving a variation of the same trend of that of the single objective optimisation (refer to Table 7.2). For the second objective at DP2 the best improvement in the Pareto optimal solutions is -11.68%. The shapes corresponding to these points at the edges of the frontier are similar to that of Figure 7.6 and are omitted. Since at DP2 the angle of attack is low, the shape for which the objective at this DP is the minimum shows a slight change compared to the baseline one, while for the other edge solution, for which the objective at DP1 is the overall minimum, the variation is more evident, due to the higher angle of attack.

7.5.1 Application of a MOOP to a practical case

The algorithm previously described and implemented into a multi-point bi-objective optimisation was applied also to a practical case, specifically the optimisation of an high-lift system, based on both slats and flaps. The procedure followed is similar to that described in [48], which was adapted to a morphing LE aerofoil without flaps. In the paper, a multi-objective optimisation was used to increase the performance of a multi-element high-lift system of a regional aircraft. In the case study, the qualitative target of the procedure applied was the reduction of the time-to-climb in the second climb segment for the take-off (TO) and the reduction of the fuel burn in the final approach phase of the landing (LDG). According to certification issues and airworthiness, this resulted into defining six design points (DP) and minimizing a two objective function,

subject to some aerodynamic constraints.

Problem statement

The formal statement of the problem reported in the paper is:

$$F(x) = \begin{cases} F_{TO}(x) = - \left(\frac{0.271}{\sqrt{C_l}} - \frac{C_d}{C_l^{3/2}} \right) |_{@TO-DP1} \\ F_{LDG}(x) = C_d |_{@LDG-DP1} \end{cases}$$

subject to:

$$\begin{cases} C_l |_{@TO-DP1} \geq C_{l, \text{baseline}} |_{@TO-DP1} \\ C_l |_{@TO-DP2} \geq C_{l, \text{baseline}} |_{@TO-DP2} \\ C_l |_{@LDG-DP1} \geq C_{l, \text{baseline}} |_{@LDG-DP1} \\ C_l |_{@LDG-DP2} \geq C_{l, \text{baseline}} |_{@LDG-DP2} \\ -C_m \leq -1.15 C_{m, \text{baseline}} \end{cases}$$

with the additional off-design post-check condition:

$$\begin{cases} C_l |_{@TO-DP3} \geq C_{l, \text{baseline}} |_{@TO-DP3} \\ C_l |_{@LDG-DP3} \geq C_{l, \text{baseline}} |_{@LDG-DP3} \end{cases}$$

In order to replicate the above formulation with the morphing leading edge aerofoil so far treated, some assumptions were made to identify the analogous design point in this latter case. In the paper, the design points (DP) derived from the application of some airworthiness requirements fixing the maximum velocities for climb and descent phases and are the result of the reference aircraft considered and an experimental campaign. For the purpose of this work, none of this data were available, so the DP were chosen by comparing the polars of the multi-element airfoil and the morphing one and proportionally scaling the angles of attack. This simplifying assumption, strictly valid only if the polar would be perfect linear, appeared to be the only way of replicating the procedure reported. After all, the application of the optimisation loop built so far to a practical case was performed for the sake of testing whether the framework is able to handle a quite complex problem, which can hold in a practical implementation.

Table 7.3 indicates the design points at which the aerodynamic computation have been performed.

Results

Similarly to previous cases, different initial populations and Pareto fractions were tested among the optimisation sessions. Specifically, two couple of sets of 50 and 100 individuals were used to seed the optimiser, which operated with a Pareto Fraction of 0.5 or 0.35, an *Intermediate* crossover and a *Gaussian* mutation. Unlike in the aforementioned bi-objective problem with only the drag coefficient as objective, in this case the convergence resulted less straightforward and the Pareto fronts produced in each run showed

| Name | Reynolds | Mach | AoA [deg] |
|---------|----------|---------|-----------|
| TO-DP1 | 20E+06 | 0.17146 | 10.13 |
| LDG-DP1 | 20E+06 | 0.17146 | 5.26 |
| TO-DP2 | 20E+06 | 0.17146 | 15.15 |
| LDG-DP2 | 20E+06 | 0.17146 | 14 |
| TO-DP3 | 20E+06 | 0.17146 | 18.06 |
| LDG-DP3 | 20E+06 | 0.17146 | 16 |

Table 7.3: Design Points considered in the MOOP

a certain variation, indicating a difficulty of the optimiser to reach the true frontier. A report of four fronts is presented in Figure 7.10. The red front run for 100 generations, while the others for 50. It can be noticed how the front moves toward the left hand-side when increasing the PS and the generations, indicating that convergence was not reached with the 50 generations sessions, for which the front is non-convex. This instead is less evident for the red-marked run, which has been obtained after 100 generations: the front shows a tiny region of accentuated non-convexity, which is likely to be due to the need of a local refinement of the solution, still not occurring notwithstanding the quite long session. As regards the influence of the Pareto fraction (PF), two values were set, as indicated in the figure, with the light-blue marked run employing PF=0.35, while the remaining PF=0.5. A lower value was checked to see if promoting lateral diversity could accelerate the convergence through the frontier, since it was observed that the initial random population was quite scattered over a disperse region of the objective space, thus a certain number of iterations were needed to get close to the frontier and start filling all the slots in the first rank. The result depicted in the figure shows a little dependency of the final Pareto front on the PF, thus no particular improvement was noticed.

Observing the position of the Pareto fronts with respect to the baseline point, indicated in the same figure, it appears that the optimisation led mainly to a reduction of the F_{LDG} objective, which underwent a percent decreasing up to -11.61%, while the first objective F_{TO} reduced at least of -0.6057%, proving also in this case the ability of morphing of mostly diminishing the drag coefficient, while bringing a little improvement in the lift, at least for the geometry so considered so far.

Figure 7.11 shows the convergence history for the longest run of 100 generations. The Pareto front got gradually populated in the first fifteen generations and was close to the final one just before the half of the iterations. After that, many more iterations were needed to locally refine the front and eliminate most of the concavities, thus indicating a certain effort to reach a stable solution.

As regards the coverage of the Pareto in terms of decision variables spread, this is depicted by Figure 7.12. A correlation is quite evident for the first DV (related to the LE




Figure 7.10: Pareto Fronts for different population sizes




Figure 7.11: Convergence history for an optimisation run with PS=100 after 100 generations

radius of the suction side) and possibly also for the third (related to the LE radius of the pressure side), while for the other the points appear more scattered. Some lower dense areas for $x(1)$ and $x(3)$ can be associated with local non converged areas in the fronts, such as the one non convex. The plot of $x(1) = A0_{up} = \sqrt{2R_{LE}/chord}$ vs F_{TO} can be well represented by a 3-rd order polynomial regression with a resulting regression coefficient $R^2 = 0.9941$. A similar trend has been noticed while plotting the scatter of the DVs against the second objective F_{LDG} .

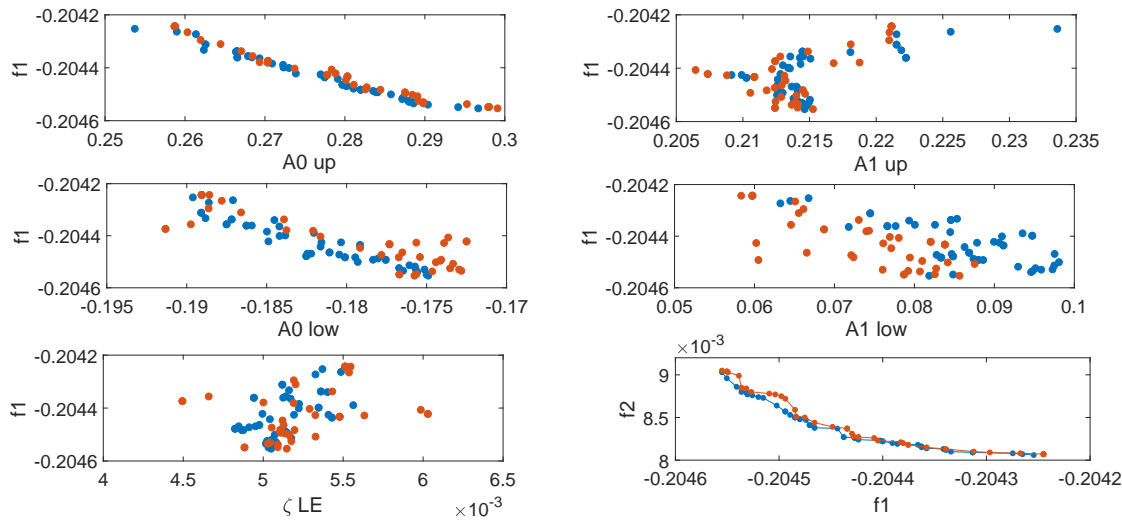


Figure 7.12: Decision variables scatter along the Pareto front for two sessions with PS=100

A final comment on the outcomes produced must deal with the off-design post-check conditions reported above. The satisfaction of these criteria was very partial for the points in the final Pareto fronts, as illustrated in next Figure 7.13, where the frontier are those of the third run (PS=100, PF=0.35, Generations=50) and of the fourth (PS=100, PF=0.5, Generations=100), already discussed. As it can be seen, only a few points in the fronts passed the post-check, evaluated at the end of the optimisation. Moreover, the conditions are satisfied without a precise pattern: there is a region in which Pareto solutions that passed the check are located, but there is no a continuity along the front and some points are isolated. This result could suggest the need of including the post-check conditions as additional constraints for this particular case, which is, actually, not thoroughly comparable with the reference case we are trying to adapt to a morphing design, since traditional high lift systems are much more efficient in improving lift than this morphing architecture, whose main benefit seems to be mainly a drag reduction, according to both multi-objective and single-objective optimisation performed so far.

Finally, two airfoils corresponding to the bounds of the post-check feasible solutions Pareto front for the longest run are reported in Figure 7.14. The original airfoil is depicted in magenta. Optimised profiles shows an upward leading edge deflection and a slightly different shape for the suction side, while the pressure side remains very similar




Figure 7.13: Off-design post-check conditions attainment for points in two Pareto fronts

in the two cases, with a flexural point typical of all the optimised shapes so far produced, mainly due to the constant arc length parameterization.

7.6 Single objective optimisation using CFD

After completing an analysis phase on both a single objective and a multi-objective optimisation using XFOIL to compute the aerodynamic performance, the following task to be performed was to repeat the optimisation embedding the CFD solver instead of the previous panel based. This stage was carried out in two steps, using two different models. Initially, a validated CFD model was available for the purpose and was based on the combination of an hybrid grid, with a structured boundary layer and triangular elements in the external field, and the Spalart-Allmaras (SA) turbulence model. The model was compared with the experimental data and with XFOIL when computing the lift, drag and moment coefficients over a variety of boundary conditions. This first validation activity is described in Section 6.2. The comparison showed a better accordance for the CFD in replicating the lift and the moment at high angles of attack, with a poorer prediction of the drag for both the solvers. The calculated and experimental charts presented a gap which was quite constant, with an overall trend among them appearing acceptable. On this basis, a single objective optimisation at DP1 was performed, firstly employing the CFD model based on the Spalart-Allmaras. The outcome of this is presented below.

By observing the optimal solution produced, there is a remarkable difference with that indicated by XFOIL and shown in Figure 7.6. Such a discrepancy suggested to go further with the validation activity, by examining other grids, specifically structured grids, and turbulence closures. A second model has been individuated, still based on a hybrid grid but embedding the transition prediction, the Transition SST [44]. The validation is proposed in Section 6.4. Below it is presented the solution obtained in this




Figure 7.14: Aerofoils at the bounds of the Pareto front attaining the post-check conditions

second run, showing a good agreement with that of XFOIL and proving that the transition prediction is responsible for the difference in the two optimised shapes, which is very relevant despite such a discordance did not appear in both global and local validation for the two models.

7.6.1 Single objective optimisation based on the Spalart-Allmaras model

By exploiting the enquiry on the optimiser settings presented above, when XFOIL was used to calculate the aerodynamic parameters, the single objective optimisation at DP1 as before was solved employing Fluent and the Spalart-Allmaras model, with a population size of 20 individuals, evolved for 40 generations. The final point was reached within the 32th generation and was kept until the end of the session. It is depicted in Figure 7.15, compared to the baseline geometry. It can be noticed that the shape is quite different from that of Figure 7.6, which was obtained by XFOIL. Specifically, while in the previous case the LE was morphed upward, producing a thickening of the suction side and a flexural point in the pressure side, in this case the shape is smoother and the LE is deflected downward. The pressure side shows a little hump because of the constant arc length.

Table 7.4 reports the improvements in aerodynamic indexes brought by the optimisation using CFD. The drag coefficient decreases of about 9%, while the overall efficiency Lift/ Drag improves of about 10%. Compare to the outcomes of Table 7.2, the variations are lower. This aspect is treated in more details in the next section.

In Figure 7.16 the pressure coefficient distribution is depicted for the baseline and the optimised shape. A lower suction peak occurs for the optimised airfoil and the pressure




Figure 7.15: Baseline (red) and optimised airfoil with CFD and SA at DP1

| Index | Baseline | Optimised | Relative variation [%] |
|-------|----------|-----------|------------------------|
| C_d | 0.040029 | 0.036433 | -8.98 |
| C_l | 1.88586 | 1.89545 | 0.51 |
| L/D | 47.112 | 52.026 | 10.43 |

Table 7.4: Comparison between baseline and CFD-SA optimised airfoil aerodynamic coefficients at DP1

recovery in the upper side is more gradual in the morphed region, up to 0.25 c . The pressure side shows no remarkable variation and the curves are overlapped.

The computed Mach number distribution for both the cases is depicted in Figure 7.17. The maximum Mach number is reached in the point where the C_p is minimum, of course, and it is visible from the picture that the optimised shape has a smaller area for the suction peak, corresponding to a smoother Mach number distribution. The flow remains fully subsonic, since the undisturbed Mach number is 0.2. Compressibility effects are limited to the portions of highest speed and for this reason the contours of other mechanical and thermodynamic quantities, such as velocity magnitude, static pressure, static temperature and density are qualitatively similar to those illustrated in Figure 7.17.

Comparison between XFOIL and CFD results

In order to try to understand the difference between XFOIL and the CFD, a simple cross check on the best solutions produced in the SOOP by both the tools was made. Table 7.5 presents the comparison between the best solution obtained by XFOIL and calculated also with the CFD. Table 7.6 presents the same comparison, considering the best solution obtained by the CFD. In the first table, the major difference occurs for the drag coefficient, motivating why the best geometry produced by XFOIL was discarded in the optimisation employing the CFD: the C_d in this second case is 65.86% higher than in the former. This is to look for in the different approximations embedded in the two tools. During the validation, the CFD showed an overestimation of the drag for the baseline




Figure 7.16: Baseline and optimised pressure coefficient distribution for SOO at DP1 using CFD with SA




Figure 7.17: Contours of Mach number for the baseline (lhs) and the optimised (rhs) airfoil for SOO at DP1 using CFD with SA

geometry, with respect to the experimental data, while XFOIL showed an underestimation. This happens also for the case presented in Table 7.5, but the difference is much higher. It is not straightforward to establish which one is wrong in computing the C_d , but according to validation process following this discrepancy, the different prediction may be caused by the absence of a transition modelling for the Spalart-Allmaras.

| Index | XFOIL | CFD | Relative variation [%] |
|-------|----------|----------|------------------------|
| C_d | 0.02530 | 0.041946 | 65.86 |
| C_l | 1.90165 | 1.90794 | 0.33 |
| C_m | -0.07629 | -0.07736 | -1.39 |

Table 7.5: Comparison between the best solution obtained by XFOIL, calculated also with the CFD and SA model

The comparison between the solvers for the best solution obtained by the CFD is reported in Table 7.6. Here the difference is less accentuated, but the discordance is present also for the C_m . The high-fidelity CFD should be able, in principle, to properly deal with high angles of attack, involving partial separation of the flow and adverse pressure gradients in the boundary layer, with a higher accuracy with respect to XFOIL. The result for this second case in terms of variations in the aerodynamic indexes is more similar to the trends emerging from the validation. The reason why the optimisation using the CFD chooses the shape of Figure 7.15 and the optimisation using XFOIL prefers that of Figure 7.6 is due to the lack of accordance in evaluating the drag in this latter case, indicating a diverging behaviour between the tools. A point in favour of the CFD solver, besides its more general validity for the less restrictive hypothesis, is that in literature the optimisation of morphing LE wings, under similar conditions, is reported to lead to downward deflections of the LE, which appears intuitive for an angle of attack as high as 15 deg.

| Index | XFOIL | CFD | Relative variation [%] |
|-------|---------|----------|------------------------|
| C_d | 0.03237 | 0.036433 | 12.55 |
| C_l | 1.8810 | 1.89548 | 0.77 |
| C_m | -0.0874 | -0.10164 | -16.29 |

Table 7.6: Comparison between the best solution obtained by CFD with SA, calculated also with XFOIL

Of course, being available the experimental data only for the baseline geometry, one must rely on the validation uniquely for this case. Even though it was carried out over a variety of Mach numbers and for the entire extent of the polar, as reported in 6.2, it is not strictly sure that the model remains valid also for morphed shape. Given this

discrepancy and the lack of a good match for the validation data, it was decided to further push on it, firstly by employing full structured grids, which could have a chance of better approximating the flow field, and on the other hand using turbulence models accounting for transition prediction. The results of this second validation activity are described in 6.4, while in the next section it is discussed the optimisation performed by using the Transition SST model.

7.6.2 Single objective optimisation based on the Transition SST model

The second SOOP solved was equal to the previous one, except for the aerodynamic model used, which was now based on the Transition SST or $\gamma - Re_\theta$ model by Menter and Langtry. A population of 20 individuals was evolved for 35 generation, leading to the optimal solution depicted in Figure 7.18.




Figure 7.18: Baseline and optimised airfoil by CFD with $\gamma - Re_\theta$ at DP1

It is here visible how the optimised airfoil is very similar to that produced by XFOIL in Figure 7.6. The leading edge deflection is positive, with a little concavity in the pressure side. The resulting C_p distribution is reported in Figure 7.19. A plateau appeared in the suction side near the leading edge and then the pressure recovery occurs quicker than the baseline geometry until 0.25 c. From that point on, even though the shape remains the same, because no morphing is allowed, the effect propagates downstream, bringing to a smoother acceleration between 0.4 and 0.8 c, in an area where the boundary is thickening and there is also a wide recirculation of the fluid, as visible in following contours. For the pressure side, the difference in the pressure distribution is less evident, although the stagnation point moves a little downstream.

It is interesting to plot the Mach number contours for the two airfoils, computed using the same CFD model, as shown in Figure 7.20. A larger area of high Mach number is visible near the leading edge of the optimised airfoil for the suction side, as inferable by the pressure distribution. This area, anyway, covers a portion of the airfoil where the boundary layer is thin and thus does not affect the drag in a relevant way. Instead, this leads to a higher pressure and more lift and the perturbation propagates downstream, in a way that the detachment and recirculation near the trailing edge is confined to a narrower region, resulting in a lower drag.

Finally, in order to visualise the recirculation near the trailing edge and the wake, the contour for turbulent kinetic energy is illustrated in Figure 7.21. The recirculation




Figure 7.19: Baseline and optimised pressure coefficient distribution for SOO at DP1 using CFD with $\gamma - Re_\theta$




Figure 7.20: Contours of Mach number for the baseline (lhs) and the optimised (rhs) airfoil for SOO at DP1 using CFD with $\gamma - Re_\theta$

is evident for the left hand side baseline profile, where the wake area is thick and a local peak of k starts from the trailing edge at the pressure side. For the right hand side optimal profile, the wake is thinner and its contours more defined. The peak is no more present and k appears smoother, showing that the aerodynamic behaviour is actually better.

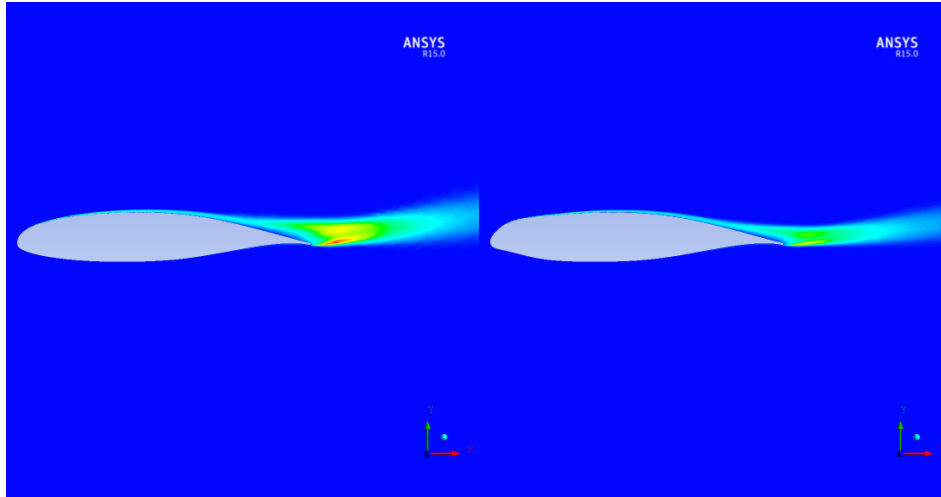


Figure 7.21: Contours of turbulent kinetic energy k for the baseline (lhs) and the optimised (rhs) airfoil for SOO at DP1 using CFD with $\gamma - Re_\theta$

The variation in the aerodynamic indexes for the two profiles is reported in Table 7.7. The drag coefficient, the objective of the optimisation, undergoes the major change, with a decreasing of 17.93% with respect of the baseline. The constraint on the moment coefficient was set to a maximum increasing of 10%, thus its increase of 8.47% within the tolerance imposed. The L/D ratio augments of 30.85%, confirming that morphing is capable of producing an improvement in the aerodynamic behaviour for the case analysed and that it mostly affects the drag, with a variation in the lift coefficient not comparable with that of multi-element high lift systems.

| Index | Baseline | Optimised | Relative variation [%] |
|-------|----------|-----------|------------------------|
| C_d | 0.03418 | 0.02805 | -17.93 |
| C_l | 1.8895 | 2.0288 | 7.35 |
| C_m | -0.08617 | -0.0935 | 8.47 |
| L/D | 55.281 | 71.087 | 30.85 |

Table 7.7: Comparison between baseline and CFD with $\gamma - Re_\theta$ optimised airfoil aerodynamic coefficients at DP1

Chapter 8

Three-dimensional Optimisation of a Wing with a Morphing Leading Edge

In order to test the optimisation procedure, a three-dimensional problem was also solved, extending the methodology already applied to the 2D geometry and presented in previous chapters. A wing model with a morphing leading edge was optimised to validate the process and give a first insight into the effect of morphing in a real body. In the next sections the geometry handling is reported, as well as the optimisation properly said and the outcomes obtained are discussed.

8.1 Reference geometry

The geometry analysed was a low aspect ratio wing, possibly replicating the shape of a tail plane of a regional aircraft. The wing span was double the length of the chord of the wing section at the wing root, while at the tip the section chord was one half of the previous quantity. The wing was swept by 20 [deg] and had a negative twist of 1 [deg] at maximum, from the root to the tip, which had a lower incidence with respect to the hub to avoid the tip to stall first. A planar view of the wing is shown in Figure 8.1. The geometry was described in terms of four stations, numbered from 1 to 4 as depicted. Sections 1 and 4 were fixed, while sections 2 and 3 could undergo morphing. This shape was chosen in order to deal with the simplest 3D case, in a small computational domain to reduce the computational time, but at the same time presenting an important level of three-dimensionality. The combination of a low aspect ratio, a sweep angle and a little twist, in fact, induces three-dimensional flows, which develop a velocity component in the spanwise direction, giving a certain generality to the example discussed.




Figure 8.1: Planar view of the reference wing, with the reference sections indicated

8.2 Geometry handling

The geometric representation of the 3D wing was obtained by specifying the four aforementioned sections and then demanding the meshing software Gambit to rebuild it, by lofting between the stations. This procedure allows a quite simple handling of the geometry, which was interpolated between the specified shapes, without recurring to a complete 3D parameterization, for example using NURBS, and without employing further CAD software. From a practical point of view, it was necessary to communicate Gambit the coordinates of the points of the sections from 1 to 4 and then instruct it to form a skin surface, using as reference edges the pressure and the suction sides of the four wing sections. While the first and the last were fixed in advance at the beginning of the simulations, sections 2 and 3 were described using the same CST parameterization employed for the 2D optimisation, since they lay on a plane. For each of them, 5 design variables were used, the same of 7.3, thus a total number of 10 DVs was needed to completely describe the 3D wing. Even though it is usual to employ at least five sections to represent 3D shapes, since the aim of this study was to set up an optimisation procedure for a morphing wing, and not to produce a ready-to-use result, it was chosen to take only two sections, at the beginning, and then eventually increase their number in a second step, when the behaviour of the optimisation loop was sufficiently known and robust.




Figure 8.2: Three-dimensional view of the reference wing

8.3 CFD solver overview

The aerodynamic performance of the wings generated during the optimisation procedure was computed by the CFD solver ANSYS Fluent, as well as in the 2D case reported in Section 7.6. In that situation, two CFD models were tested by comparing their predictions with the experimental data for the baseline geometry and then performing a SOOP. The results obtained employing both of them have been already discussed in Sections 6.2 and 6.4 as regards the validation and in Section 7.6, as regards the optimisation. Switching to a three-dimensional case, a choice must be done on which model to use. The Spalart-Allmaras turbulence model showed the best overall accuracy, with respect to other fully turbulent models, but still presented a not negligible gap with the experimental data and provided an optimal shape quite different from that of XFOIL, due to the lack of a transition analysis. On the other hand, the Transition SST model proved to give better accordance for the baseline polars and with the optimal shape of XFOIL. This was true in a 2D case, while for a 3D flow any test data were available. When using transition models in 3D domain, the presence of secondary fluid paths due to the tridimensionality, for example a velocity component in the sweep direction or the tip vortex, can cause the transition to occur in a different way from the corresponding 2D case; if the formulation employed does not have a robust crossflow transition prediction capability, a correct distribution of the flow field can be kept only by chance, but in general it is challenging to build a tool able to accurately predict the 3D development of the boundary layer in presence of pressure gradients and secondary flows. For this reason, since in principle and with the available data it was not possible to compare the two models, the Spalart-Allmaras and the $\gamma - Re_\theta$, in 3D, it was chosen to use the first

one. The SA is lighter in a computational sense, since it is a one equation model, while the other is a four equation, and it is more stable, since it is fully turbulent and does not deal with the computation of local quantities for assuming the transition onset.

With this in mind, now the results obtained in the 2D, for the validation and the optimisation, can be exploited to go for a 3D case. Gambit v. 2.4.6 was used to rebuild the geometry and generate the mesh. The grid was hybrid, with a structured O-shaped boundary layer of parallelepipeds and an unstructured tetrahedral grid for the external solution. This kind of meshes are able to provide good results without excessively increasing the number of cells; since 3D simulations are very time-consuming, dealing with a moderate number of elements can save days in the optimisation run. In this case no experimental data were available, to compare the CFD results with. The work done for the 2D problem could provide a suggestion on how to construct a mesh providing acceptable predictions, even though the need to limit the mesh size prevented from using a discretisation as fine as that employed in many two-dimensional grids tested. After various attempts, it was found that a good numerical stability could be obtained by employing 4 millions of cells, with a boundary layer of 30 rows and a growth ratio of 1.1. In this case the wall function approach was preferred, because reaching a $y^+ < 1$ would have required a larger amount of rows, in order to prevent from having high skewed cells in the transition region between structured and unstructured zone.

The domain is illustrated in Figure 8.3. Called c the wing root chord, the domain was a parallelepiped extending $7c$ upstream, $11c$ downstream, $8c$ in the wings span direction and $7c$ upward and downward of the wing. Again, even though in 2D it is common to collocate a pressure far-field condition at 20 chord lengths, here it was necessary to limit the extent of the space to avoid an excessive computational cost.

8.4 Single objective optimisation

The first problem solved was a single objective optimisation, in the same form of Section 7.6.1. An initial population of 20 individuals was evolved for 15 generations. In order to accelerate the computational time, each individual was simulated on a remote cluster, using 16 Intel Xeon E5-575 cores. An automatic procedure was developed, in order to perform the optimisation and the mesh in a local computer and linking Matlab with the Centos OS remote cluster, so as to execute there only the Fluent session. The results of this initial iterations showed that the aerodynamic constraint on the lift coefficient, desired to be at least as high as the baseline for each individual, was unlikely to be respected, causing the optimiser to produce infeasible solutions and lose a preferential direction in the search. Because of this, the baseline wing was the optimum for the formulation adopted, with other points exhibiting less drag but also a lower lift and, therefore, causing the penalty function to activate. By examining the scatter graph of the individuals simulated during the optimisation, it emerged a clear trend: the drag and the lift coefficient appeared conflicting and the points formed a Pareto front, depicted in blue in Figure 8.5, where the baseline point was also the one exhibiting the highest lift. As a consequence of this outcome, in order to assess the relationship between the drag




Figure 8.3: Illustration of the computational domain

and the lift coefficient, which appeared to be conflicting for the 3D case, since they both tended to decrease despite the lift was desired to have a non-negative variation, a multi-objective problem aiming at minimising the drag and maximising the lift was solved for the 2D case, where the optimal solution minimising the drag after 40 generations exhibited a very moderate lift increment. After evolving a population of 20 individuals for 40 generations again, it appeared that for the two-dimensional geometry the drag and the lift were not conflicting at all and thus a unique solution was found, that dominated the one of the previous single-objective problem. The airfoil obtained is depicted in Figure 8.4. The shape is more similar to other droop-nose airfoils found in literature, with the leading edge bent downward and a concavity in the pressure side. By comparing it with that of Figure 7.15, the major difference appears for the lower side. Table 8.1 reports the variation of the aerodynamic coefficients for the baseline and optimised airfoil obtained with the 2D multi-objective optimisation. Its solution dominated the one previously presented for the single-objective problem. Indeed, in that case the convergence stalled in the last 8 over 40 generations, the optimiser was not able to further refine the minimum and a better point, like that found in the multi-objective formulation, was not




Figure 8.4: Baseline and optimised airfoil obtained with the 2D MOOP employing the SA model

reached.

| Index | Baseline | Optimised | Relative variation [%] |
|-------|----------|-----------|------------------------|
| C_d | 0.04003 | 0.03616 | -9.67 |
| C_l | 1.8859 | 1.9276 | 2.21 |
| L/D | 47.112 | 53.31 | 13.16 |

Table 8.1: Comparison between the aerodynamic indexes of the baseline and optimised airfoil for the 2D multi-objective problem discussed

The conclusion of this 2D analysis revealed that in the single-objective problem of Section 7.6.1, the process got stuck into a solution nearby the minimum, but still not the best. Though in the multi-objective optimisation the final individual presented a little improvement in terms of aerodynamic indexes, the airfoil was quite different in the pressure side. This was likely to be due to the constant arc length parameterization, for which two shapes always exist, that keep the length constant and they have an opposite concavity. Which shape is chosen depends on the values of all the curvature correction coefficients and probably between the single and multi-objective case for the lower side a different concavity was produced. Coming back to the three-dimensional problem, a different behaviour with respect to the 2D case was found, since the drag and the lift resulted conflicting. The lack of agreement could be due to an insufficient model for the 3D case or to an effective different physics involved, which causes a response of the wing not assimilable to that of a single section of it.

8.5 Multi-objective optimisation

The Pareto front described by the set of individuals explored by the optimiser in the 3D SOOP brought to light the conflicting nature of the lift and the drag for the three-

dimensional case, suggesting to repeat the MO optimisation aimed at minimising the C_d while maximizing the C_l . The synthetic formulation of the MOOP was:

$$\begin{aligned}
 & \text{minimize} && C_d(\mathbf{x}), \mathbf{x} \in \Omega \\
 & \text{maximize} && C_l(\mathbf{x}), \mathbf{x} \in \Omega \\
 & \text{such that:} && |L_m - L_0|/L_0 \leq 1e - 6 \\
 & && |A_m - A_0|/A_0 \leq 0.05 \\
 & && (C_{m,b} - C_m)/C_{m,b} \leq 0.5 \\
 & && \Omega = \mathbf{x} \in \mathbb{R}^5 | \mathbf{Lb} \leq \mathbf{x} \leq \mathbf{Ub}
 \end{aligned}$$

The purpose was to further refine the Pareto described in the fitness function space by the set of individuals simulated during the SO problem. Figure 8.5 illustrates the scatter of the feasible individuals and highlights the initial Pareto, obtained with the single-objective optimisation and drawn in blue, and the final frontier obtained after evolving a population of 30 individuals for 10 generations in the MO approach. It is visible how a slight change occurred between the two frontiers, with the major variation for the left edge, whilst the baseline point still remained the one exhibiting the highest lift. Although some individuals with a higher lift were found by the optimiser, they did not respect the constant arc length constraint and thus they do not belong to the Pareto.




Figure 8.5: Pareto fronts obtained after the 3D SO problem (Initial Pareto) and after the 3D MO problem (Final Pareto)

The final shape obtained, minimising the drag coefficient, looked like that of Figure 7.15. In Figure 8.6 it is reported the airfoil lying on the section 2 of the wing. The deformation in the spanwise direction resulted very smooth and the airfoil at section 3 bore a striking similarity with that of section 2. In terms of aerodynamic indexes, the results are summarized in Table 8.2. The drag coefficient decreased of 3.73%, with




Figure 8.6: Wing section at station 2 for the shape minimising the C_d in the MOOP, compared to the baseline geometry

an overall L/D ratio increase of 1.174%. The variation was very modest, with respect to the bi-dimensional case. The outcome of this activity showed that for the geometry

| Index | Baseline | Optimised | Relative variation [%] |
|-------|----------|-----------|------------------------|
| C_d | 0.14078 | 0.13553 | -3.73 |
| C_l | 1.3976 | 1.3613 | -2.60 |
| L/D | 9.9275 | 10.044 | 1.174 |

Table 8.2: Comparison between the aerodynamic coefficients of the baseline and optimised airfoil for the 3D multi-objective problem discussed

considered the benefit produced by morphing was very limited. Actually, three aspects should be considered when discussing these data. The first one is that the aerodynamic model employed was not validated, due to the lack of experimental data in 3D, thus the accuracy of the result cannot be ascertained. Secondly, the MO problem ran only for 10 generations and further iterations might lead to a better improvement. Thirdly, the aim of this three-dimensional analysis was to test the procedure for an extended body and, though the wing presented a certain degree of three-dimensionality and the flow field as well, the geometry parameterization was kept very simple, with only two sections directly controlled by the optimiser. Providing more flexibility to the shapes generated, by using more stations to define the wing, may give a better chance to achieve an aerodynamic improvement closer to that of the 2D problems. Despite the modest increment brought by morphing, this does not necessarily prove its ineffectiveness, since the outcome was intended primarily to check the behaviour of the framework when dealing with a wing. Finally, the indications provided by this process are useful in that they can be used for a further development of the procedure. The algorithm written so as

to automatically link Matlab and the Linux-based environment, common to many computational servers, can be considered itself a step forward, allowing for a significant increase in the computational efficiency.

Chapter 9

Meta-model Based Optimisation

In previous chapters an optimisation framework has been presented, which was able to solve either single objective or multi-objective problems involving a morphing leading edge. In Chapter 7 it is reported the procedure applied to a bi-dimensional geometry, which was optimised to minimise the drag at one or two operating points. The evaluation of the performance of the aerofoil was performed by relying firstly on the panel solver XFOIL, which can calculate the aerodynamic of a 2D body in a very fast way and with good accuracy, and in a second moment on a CFD model built in Ansys Fluent. The employment of these aerodynamic tools, predicting in our case the lift, drag and moment coefficient, allows for the search of an optimal design, which is likely to be close to the actual solution, if the model is well calibrated and validated. In fact, the aerodynamic optimisation is dealing with a model and we optimised a shape within that model, not a physical object. It is only thanks to the validation described in Chapter 6 that our solutions have the chance to be close to the real best overall solution respecting the constraints imposed. The aerodynamic model built was simulating a very complicated physics, though the problem itself was not the most challenging, and embedded some assumptions and hypothesis which allows to simplify the computation and, as matter of fact, make it possible, because it is not even imaginable to solve directly the Navier-Stokes equations for a large domain at high Reynolds number. This preamble is meant to underline that we are dealing on a certain approximation of the reality in our aerodynamic optimisation and that the results presented are not the truth, but a hopefully good representation of a much more complex physics. In an engineering approach, this is the standard way of facing real world problems, for which the theory is just sufficient to lead to working solutions.

Notwithstanding the simplifications embedded in the aerodynamic model, which limit its applicability to a neighbourhood of the case treated, the computational effort spent to compute the performance of each individual was still remarkable, especially for the three-dimensional geometry described in Chapter 8. The optimisation procedure could be simplified more, using a faster tool to calculate the fitness, by building a further model that approximates the response of the CFD solver to the design variables. This procedure is referred to as meta-modelling or surrogate modelling, that is building a

model of the model. A surrogate model can help in individuating the optimal regions of the domain, without performing a large number of expensive simulation, which is typical of evolutionary algorithms. In this chapter it will be briefly presented some strategies used to build meta-models and some of them will be applied to the problem treated in this work.

9.1 Application of surrogate models

A surrogate model (SM) is an approximation of a simulation used to construct simpler and lower computational cost models [15]. It is built basing on the response of the simulations to a limited amount of intelligently chosen points to capture the relationship between the input and output variables. The original function $f(\vec{x})$ is then represented as $\tilde{f}(\vec{x}) = f(\vec{x}) + e(\vec{x})$, with $e(\vec{x})$ the error between the original and modelled function. $\tilde{f}(\vec{x})$ is the regression function, behind which is concealed the technique used: rational functions, radial basis functions, Kriging models, support vector machines, artificial neural networks, polynomial regression and splines are some of them.

Manríquez et al. [15] proposed a classification of the surrogate models application based on how they are employed in the optimisation loop. Substantially, two approaches are available, the so called Direct Fitness Replacement (DFR) or Indirect Fitness Replacement (IFR). In DFR the solutions are evaluated with the SM and are assumed to be comparable to those obtained by the real function. Of course, this is true only if the meta-model (MM) has been constructed properly, using a large set of evenly distributed data and assessing its prediction in other points, to verify that it behaves in the desired way. Otherwise, the evolutionary algorithm (EA) is likely to converge to a false optimum or a false Pareto front, vanishing the benefit of the simplification brought by the SM.

The evolution control, the way the SM deals with the original fitness function within the DFR approach can be: (i) No Evolution Control (NEC), (ii) Fixed Evolution Control (FEC) and (iii) Adaptive Evolution Control (AEC). If No Evolution Control is used, the SM simply replaces completely the original fitness function. The lack of a feedback can lead to misleading solutions if for some reasons the MM fails in predicting the response or is not enough accurate. Fixed Evolution Control foresees an alternation between the original and surrogate function. Only some individual or generations can be evaluated through the meta-model, and the remaining are calculated with the real function. In this way, it is possible to update the model during the optimisation, improving its accuracy since it is fed with points belonging to the search region. This approach has been successfully implemented in a variety of cases, leading to good results. The main issue is the switchback process, which should be set to provide a meaningful stream of information between the approximated and original model. The last options is the Adaptive Evolution Control, which automatically chooses whether to compute an individual with the surrogate or real function. according to some parameters. This implies a complex optimisation framework with nested loops, but can overcome the drawback of the previous instance, if the fixed exchange mechanism between the two models occurs in an misleading way.

In the IFR method, the original fitness function is used during the iterations, while one or more components are assessed in the surrogate model. A number of solutions evaluated using the SM are then selected and delivered to the real function, using the approximated fitness indirectly. Often the Evolutionary Algorithm (EA) is employed to explore the design space, while the MM intervenes in a local search, providing candidate solutions which are then assessed by the real function.

As already aforementioned, there exist a variety of regression techniques available to generate the approximation of a function. For example, in the simplest case of a $\mathbb{R} \rightarrow \mathbb{R}$ function, polynomial fitting is used to interpolate small sets of data, while splines allows for a good interpolation of large datasets. For multidimensional functions, these simple methods are not applicable straightforwardly any more and more sophisticated techniques have been developed. For example, starting from splines it is possible to derive the Radial Basis Functions (RBF), which are real valued functions depending on the distance between the input and their centre. Typical RBF are the linear, cubic, multi-quadratics or Gaussian. It is possible to build a network of RBF to interpolate data in a multidimensional environment, forming an artificial neural network (ANN). ANN are networks of elementary units which exchange information, generating an overall model which can be tuned to adapt to certain inputs. Kriging methods, AKA Gaussian Process Regression (GPR) are a group of geostatistical tools based on mean squared error to perform regression. They are employed in the DACE (Design and Analysis of Computer Experiment) model, which extends the potentiality of Kriging to manage three or more dimension. According to the application, the best method should be chosen, but it is difficult to have an *a priori* knowledge of what model best fits the data. Some of them are affordable only for small sets, other can handle a large number of samples. Statistical methods, such as Kriging, are powerful, because they embed the concept of uncertainty and give a better estimation of the error, but can be expensive in interpolating functions of many variables. Below a tool quite commonly used in evolutionary optimisation is presented, because it will be applied to the two-dimensional case geometry of Chapter 7, the artificial neural network.

As already aforementioned, there exist a variety of regression techniques available to generate an approximation model of a function. According to the cardinality of the problem, the dimension of the function as well as its nature (linear, non-linear, binary, noisy, etc.) and the purpose of the regression a tool can perform better than others and it is demanded to the analyst to interpret which one to choose, e.g. Polynomial Approximation, Linear or Non-Linear Regression, Generalized Linear Regression (GLR), Artificial Neural Networks (ANN), Radial Basis Functions (RBF), Gaussian Process Methods (GPR or Kriging), Support Vector Regression (SVR) etc. All of these methods can be embedded within an optimisation loop to construct a meta-model based optimisation.

Radial Basis Function are real-valued functions whose output depends on the distance between the input and the centre of the function itself. They can be used either for interpolation or for regression, also in a network formulation. The function is expressed by the model as a linear combination of a Radial Basis Function $\phi(\mathbf{x})$ plus a

trend function $P(\mathbf{x})$:

$$\hat{y}(\mathbf{x}) = \sum_{i=1}^N \omega_i \phi(\mathbf{x}) + P(\mathbf{x})$$

with ω_i N unknown weights to be determined and $\phi(\mathbf{x}) = \phi(\|\mathbf{x}^{(i)} - \mathbf{x}\|)$, depending on the distance metric between the observed point $\mathbf{x}^{(i)}$ and the untried one \mathbf{x} . The distance is usually computed in an Euclidean sense, while the RBF can be Gaussian, Thin Plate, Multiquadrics, Inverse-Multiquadrics and so on. The trend function is usually set to be constant. If N is the number of samples, it is possible to solve the problem by either imposing the exact interpolation at sample points or allowing for a smoothing effect, which can be positive in presence of noisy data.

Gaussian Process Regression (GPR) or Kriging methods can be considered a particular case of RBF, in that they still attribute the value of the model basing on a weighted sum of the points surrounding the input datum. The response is expressed by the sum of a trend function $\mathbf{f}^T(\mathbf{x})\boldsymbol{\beta}$ plus a stochastic process which is a Gaussian random function: $Z(\mathbf{x})$:

$$y(\mathbf{x}) = \mathbf{f}^T(\mathbf{x})\boldsymbol{\beta} + Z(\mathbf{x})$$

$\mathbf{f}^T(\mathbf{x})$ is a set of the regression basis function and $\boldsymbol{\beta}$ their coefficients. The trend function is usually taken as a low-order polynomial, even constant. $Z(\mathbf{x})$ is a stationary random process with zero mean, variance σ^2 and covariance $\text{Cov}[Z(\mathbf{x}), Z(\mathbf{x}')] = \sigma^2 R(\mathbf{x}, \mathbf{x}')$. $R(\mathbf{x}, \mathbf{x}')$ is the correlation function, depending on the distance between zany two sites \mathbf{x}, \mathbf{x}' in the domain and other parameters. It is usually taken as a Gaussian exponential function, such that $\text{Cov} = \text{Cov}(\|\mathbf{x} - \mathbf{x}'\|_2, \boldsymbol{\theta}, \mathbf{p})$, with $\boldsymbol{\theta}, \mathbf{p}$ the hyperparameters to be tuned. The regression process is performed to minimise the mean squared error and can bring also to an estimation of the uncertainty of the fit.

Support Vector Regression (SVR) is a tool derived from Support Vector Machines, a supervised statistical learning method originally developed in the 90s at the Bell Laboratories by Vapnik and others. In SVR, a non-linear function is learnt by a linear learning machine in a kernel-induced feature space, while the capacity of the system is controlled by a parameter which does not depend on the space dimensionality:

$$f(\mathbf{x}) = \langle \mathbf{w}, \mathbf{x} \rangle + b$$

The learning algorithm minimizes a convex functional and its solution is sparse, because the computed weights w_i do not depend above all the samples and all the training data close to the model prediction, within a threshold ϵ are ignored. SVR appears as a promising method that balances model simplicity and data fitting and is particularly suited of physical experiments, since the threshold can be tuned as the amount of experimental uncertainty.

Artificial Neural Networks are a general class of mathematical tools relying upon building an overall transfer function by combining several subunits, each one with a specific representation and depending on tunable parameters. They mimic the structure of brains, in that they are composed by layers of neurons, exchanging informations each

other. By selecting the network architecture and the neuron representation it is possible to use ANN in several application within machine learning. Since ANNs have been used to fit the aerodynamic data computed by XFOIL for the two-dimensional reference geometry, they will be presented in more details below.

9.2 Artificial neural networks

An artificial neural network (ANN) is a mathematical tool employed in computer science with many purposes: mimic the behaviour of physical systems, analyse of real world data by classification, pattern recognition, clustering or regression, assist machine learning, control dynamic systems. An ANN is made up of elementary units, the neurons, which are linked together and exchange information.

In Figure 9.1 a single neuron is depicted. Three different functional operations take place in the neuron. First, a scalar input p is multiplied by the scalar weight w , forming the product wp . Second, the product is added to the scalar bias b to form the net input n . This latter is then applied to the transfer function f , which produces the output a . The transfer function f of the neuron has to be set in advance and determines the kind of the neuron. Several possibilities are available, common options are linear functions or tangential sigmoid, but according to the problem one may use other sigmoids or inverse functions, as well as radial basis functions. The weight function w and the net input function, through the bias b in our case, represent the adjustable parameters of the neuron, producing a different response to the same input. The idea below the ANN is that each neuron can be set so as to produce a certain individual transfer function. By coupling many neurons and let them communicate, sharing inputs and outputs, it is possible to obtain an overall model which can be trained, i.e. changing each adjustable parameter of every neuron, in order to fit to the problem.




Figure 9.1: Functional scheme of a single neuron

In this work, ANNs were used for regression, in order to substitute to the time-consuming aerodynamic tools a function able to model their behaviour. Since we are modelling the response of a model, the CFD code, this technique is referred to as meta-modelling, that is a model of the model. When dealing with function approximation,

there are some categories of ANN which are known to be more suitable. As previously mentioned, an ANN is made up of a net of neurons, connected each other. Figure 9.2 shows a typical feed-forward network, composed by three layers of neurons. Each layer takes information from upstream, elaborates it and send it downstream. No interaction between the neurons of the same layer occurs, but each neuron of a layer communicates with all the neurons of the following one. Notice that unlike the model of Figure 9.1, here a neuron accepts an input vector. The weighting function, therefore, performs the matrix product $\mathbf{a} = [W]\mathbf{p}$. The bias vector \mathbf{b} is then added to \mathbf{a} and the sum processed by the transfer function \mathbf{f} . The relationship for a single layer is then $\mathbf{a} = \mathbf{f}([W]\mathbf{p} + \mathbf{b})$. Each output of a layer is the input of the following one. A layer receiving the first entry \mathbf{p} is called an *input layer*, while the last layer producing the output is called an *output layer*. Intermediate layers are called *hidden layers*. A feed-forward neural network used for function approximation has typically sigmoid hidden layers and linear output layers.

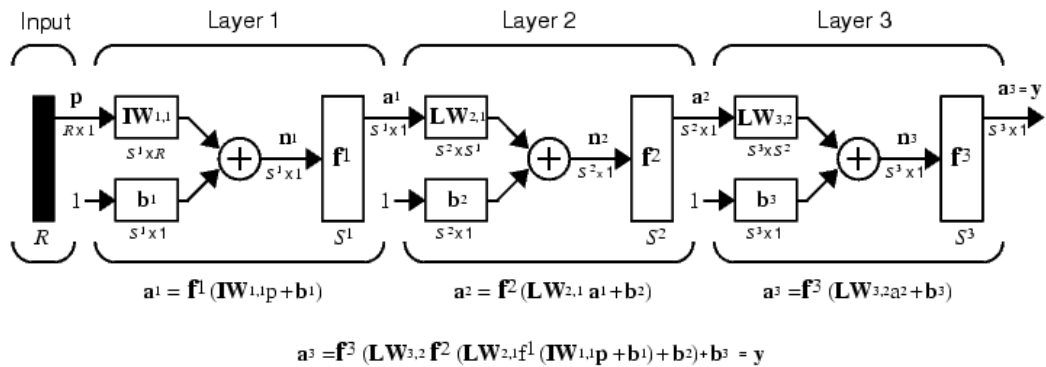


Figure 9.2: A feed-forward three-layer neural network

As previously highlighted, the weights and the biases represent the degree of freedom of an ANN, given a certain architecture and chosen the transfer function for each layer. The process of tuning these parameters in order to fit the network to a certain dataset is referred to as training. Training the network requires a set of sample points to feed the net with and the corresponding sample data. The network can deal with multi-dimensional data and can be trained to provide an approximation of the relationship between an $s \times r$ input matrix, being s the dimension of the vector and r the number of samples, and an $q \times r$ output matrix, with each row an output vector of dimension q . Generally this process is performed in three steps: first, the data are preprocessed and classified; second, the net is trained using a dedicated algorithm; third, the data are postprocessed. The aim of the training can be to minimise the mean squared error (MSE) between the targets t_i and the response a_i of the network:

$$MSE = \frac{1}{N} \sum_1^N (t_i - a_i)^2$$

, although other objectives are possible. In order to do that, the training algorithm up-

dates the weights and the biases and computes the response for a subset of the sample points. This subset is referred to as the training data. During the training, the MSE referred to the training data is expected to decrease. In order to evaluate the performance of the network, a two further subsets are used: a validation subset and a test subset. The MSE is computed for both of them too. Typically, the MSE for the former reaches a minimum during the training and the final parameters of the network are those corresponding to the iteration in which that minimum occurred. The test subset is used to further validate the network. Its MSE minimum should occur near to that of the validation subset, if data have been properly divided. The use of these three sets is needed to verify that a good response is available also for those points which are not directly involved in the training.

During the preprocessing, the data are divided into the three sets and eventually normalized. This scaling avoids a saturation for certain transfer functions and makes all the entries of a vector uniform and equally important. The division can be done in a random way or following a more structured procedure, for example filling the sets at the same time, picking the data and assigning each time to a different set, as in dealing a deck of cards. The proportion of the elements in each subset is a rule of thumb and generally the training data are the 70% of the total, while the remaining 30% is equally divided between the other two sets.

The simplest training algorithm is the gradient descent, which updates the parameters in the direction where the performance (e.g the MSE) decreases faster:

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \alpha_k \mathbf{g}_k$$

with \mathbf{x}_k the vector of current weights and biases, \mathbf{g}_k the gradient and α_k the learning rate. Several methods are available and provides different performances, according to the problem treated. The Levenberg-Marquardt algorithm is typically the first choice if the dataset is not too large, but the Bayesian Regularization can provide better accuracy for challenging problems. It is common to test many algorithms and then choose the best, since no *a priori* knowledge is often available, except for the experience of the analyst.

9.3 Two-dimensional optimisation with complete fitness replacement by ANN

With the purpose of testing the performance of a surrogate model for the prediction of the objective function in the optimisation loop, an artificial neural network was built to replace the aerodynamic calculation of XFOIL with a faster regression function. In order to check the ability of ANNs to approximate the relationship between the design variables and the fitness, the database filled during the several optimisation runs performed and containing for each individual its aerodynamic indexes, the drag, the lift and the moment coefficient, was used. This database had a large number of entries and allowed for a wide exploration of the design space. As previously underlined, when applying an ANN to fit samples, it is worthwhile to pre-process them in order to realize

a good coverage of the fitting space, discarding points which are too far from the subdomain of interest, and to form the three sets for training, validation and testing in a proper way. Having to deal with this, two approaches were used: the first one operates in the design variables space, the second in the fitness landscape.

The first approach is typical of the design of experiment (DoE) methodology, in which a poll of a function is performed, by choosing *a priori* some sample points and evaluating the function there. This way is somehow the only option when no information on the fitness landscape is available and the optimum position is completely unknown. The second approach exploits the large database build during the optimisations performed, which allows to look where the best points lie, and it comes after the first, because there is certain knowledge of the objective space. Both these methods have been employed and evaluated. Since in general the non-linear regression of a function produces bad predictions when extrapolating values, operating using samples produced by a previous optimisation is expected to produce better quality of the approximation, since we are operating nearby the subdomain which will be investigated by the optimiser when employing the surrogate model. In order to compare the results produced by several ANNs, trained using different datasets, a sample of about 800 points generated during a previous optimisation run was used to test the ANN and see whether it is able to accurately predict the response when applied to a real optimisation. In the following section the outcomes of this procedure will be presented and compared each other.

9.3.1 ANN training

The individuation of a good surrogate model to be used into an optimisation loop is, in general, an iterative process. When dealing with ANNs, the analyst has many degrees of freedom to play with, in order to improve his model: the neuron transfer function, the number of neurons in each layer, the network architecture. The presence of so many variables has led to the development of dedicated procedures to optimise the neural network itself. For the purpose of this work, such an advanced methodology has not been applied and a try and catch search has been performed, using general rules. A common architecture employed when using ANNs for function approximation is a feed-forward network, with one input layer containing as many neurons as the number of variables, e.g. five for our 2D case, one or more hidden layers and an output layer with the size of the output parameters, e.g. three for our case, C_d, C_l, C_m . This last one has linear neurons, while the other have tangential-sigmoid transfer function. Each layer operates in sequence and no feedback is present. Having the number of neurons in the input and output layer fixed and the transfer function for each layer, the only degree of freedom is the number of hidden layers and their size. As a general trend, the highest the number of neurons into an hidden layer, the highest the possibility of overfitting. The attention of the analyst should be focussed in reproducing the results with the best accuracy, using the minimum number of neurons and hidden layers, to avoid overfitting, which is one of the main drawbacks of ANNs. One hidden layer with ten neurons is the starting point used, with twenty neurons representing the upper bound. For each

ANN trained within this work, several attempts were made in order to obtain the best compromise. According to the dataset, the number of hidden layers was varied from three to five, with the neurons in each layer form six to sixteen.

In the processing of the samples, the standard division of 70%-15%-15% among the training, validation and test set was used. The building of the model was performed using the Neural Networks Toolbox™ of Matlab R2015a. As previously reported, two approaches were followed to get the samples: the first was based on a latin hypercube sampling (LHS) of the design variables space, the second was based on choosing some data among previous optimisation runs. The subdivision of the data among the three sets was based on a random process in the first case and on an interleaved method in the second.

It is here described the results for three neural networks, trained using different basis of data. Their features are described in the following Table 9.1. The network received as input the five DVs and gave back as output the values of the three aerodynamic indexes. For the sample data, they were obtained using XFOIL. The training algorithm was always the Levenberg-Marquardt, which produced better results compared to Bayesian Regularisation or Scaled Conjugate Gradients.

| NN ID | Hidden Layer Size | Samples | MSE validation | Regression coeff. |
|-------|-------------------|-------------------------|----------------|-------------------|
| NN 1 | [10 12 16] | 1000 from previous opt. | 0.0038 | 1 |
| NN 2 | [8 6 7 6] | 453 from LHS | 0.0092 | 0.99818 |
| NN 3 | [8 6 7 6 8] | 5574 from previous opt. | 0.0015 | 0.99996 |

Table 9.1: Features of the three ANNs presented

The first ANN was trained using 1000 samples picked from the database of the points explored in previous optimisation sessions. This points were selected by firstly sorting all the points available according to their fitness value, the drag coefficient, and then choosing one thousands among them to cover the whole interval. In general, the ANNs showed to hardly deal with points at the bounds of the domain, corresponding to very high drag values and low lift. The relative high value of the MSE for the validation subset in all the three networks was mostly due to these points, which the networks are not able to fit. The third network was trained using the same approach, but employing a larger number of points to better cover the fitness landscape. The second network, instead, was based on a latin hypercube sampling of the design space of 1000 samples. Among them, since XFOIL was not able to converge for some distorted shapes and because of the constant arc length parameterization there was a great sensitivity of the geometry to the value of the DVs, only converged points were retained and the other discarded, leading to the number of 453 reported in Table 9.1.

After training each network with the sets indicated, their accuracy was tested by ap-

plying them to a sample of 833 points, generated during a previous single optimisation run (here again the number is due to the discarding of bad points). In this way it was possible to see whether how the model built deals with a real optimisation in terms of accuracy, without actually performing it. A fast response parameter to measure the accuracy can be the number of samples between those of the test set, whose prediction is inside a certain bound on the relative error between the real value and that computed by the ANN. For example, in the first NN, NN 1, for the 31.93% of test points the drag was predicted with a relative error beyond 10%; for the third NN, NN 3, this value was 7.08%, but for the second, trained using LHS sampling, was 78.75%. It is clear how using previous optimisation points can lead to a very good accuracy in estimating the aerodynamic indexes for a real new optimisation run, while LHS, which is more generic, is susceptible of larger errors, because extrapolation is inefficient. As regards the lift and the moment prediction, the performance of the three networks was more even, with the worse result for the moment coefficient in NN 1, where almost 20% of the points violated the tolerance of 10% for the relative error but a better prediction for all the other case, with less than 10% of points exceeding the bound.

Figure 9.3 depicts the relative error between the real value and the value predicted by the NN 1 when testing it using the 833 samples from an optimisation session. Since these points were sorted according to their drag coefficient, it is evident from the bar plot that at the bounds of the domain, where the drag becomes very high, tens of time higher than the optimal, the lift decreases and the moment increases, the network was not able to fit the data. Besides, by comparing the distribution of the points used for training the network and those in which the test is performed, it is even more evident that outside the domain used of training the performance degrades.




Figure 9.3: Bar plot on the relative error between test data and the prediction of NN 1

The accuracy of a network can also be inspected by visualising the distribution of the relative error between for the response of the network to the training data. In Figure 9.4, for example, it is reported the distribution of the error for the training process of NN 2. The LHS helps in having a good coverage of the design space, but the neural network was not capable of fitting the data in a very accurate way, as proved by the spread of the histogram plot. The other two networks, instead, showed a much narrower graph and the value of the relative error does not exceed 6%.

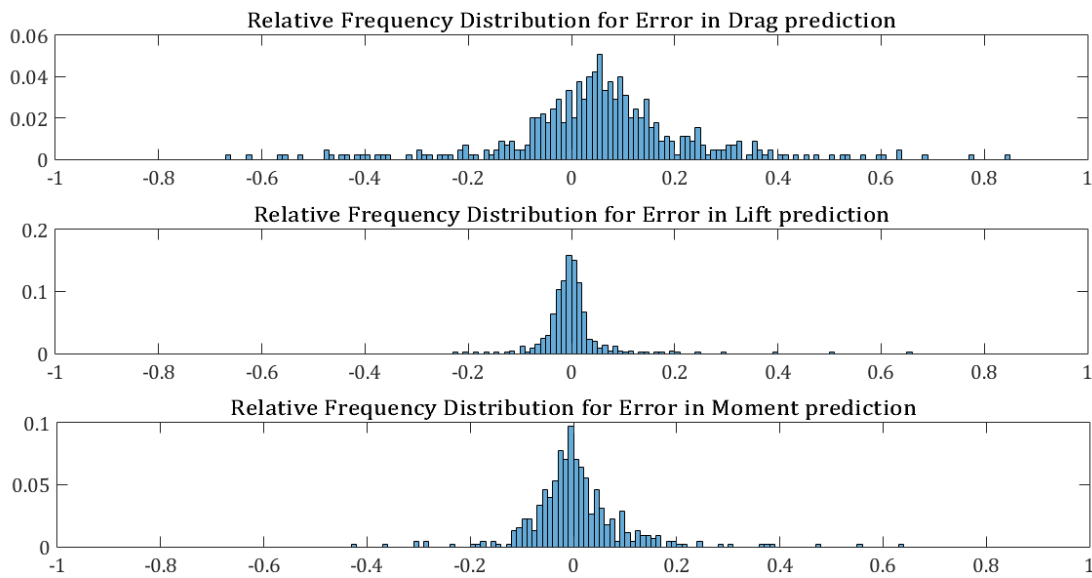


Figure 9.4: Bar plot on the relative error between test data and the prediction of NN 1

9.3.2 Results of the optimisation sessions

After training the networks, they were employed to solve a SOOP for a two-dimensional case. Specifically, the problem was the same of Chapter 7 (see Section 7.4), when the drag coefficient was minimised at DP1. Here, the NN had completely replaced XFOIL in evaluating the aerodynamic performance, following the simplest approach of No Evolution Control. For each neural network, three runs were launched. The initial population was the same and made up of 50 individuals and has been evolved for 50 generations in the first two runs and 100 generations in the third. At each run, the final optimal point produced by meta-modelling was compared with those indicated by XFOIL. The outcome of this procedure is summarised in the next two figures, Figure 9.5 and Figure. in Figure 9.5 it is reported the difference in terms of the values of the design variables. For the first, third and fifth design variable there seemed to be a good accordance with among the values individuated with the meta-model and those produced without it, employing directly XFOIL. For the second and the fourth, instead, an higher variation was present, especially for the NN 2, trained using LHS.




Figure 9.5: Comparison between the optimal solutions provided using ANN and XFOIL, in terms of decision variables

In Figure 9.6, instead, the comparison is presented in terms of the aerodynamic coefficients, which were approximated by the ANN. There was a good agreement between them, with a maximum error of about 2% for the objective C_d and 4% for the C_m occurring with the NN 2. In absolute terms, the third network, NN 3, reached the result closest to the global optimum of the standard loop with XFOIL. This can be attributed to the larger number of the training set and its derivation from previous optimisation searches.

In the procedure described, the complete replacement of the “real” objective function with a surrogate model based on an artificial neural network led to identify a neighbourhood of the optimum, proving that an ANN can approximate the relationship between the design variables and the aerodynamic indexes with a good level of accuracy, if the training set is adequate.

9.4 Approximation of the CFD data by a regression model

The complete replacement of XFOIL with an ANN showed that the original minimum was caught and thus the model was able to accurately represent the aerodynamic behaviour of the individuals. The approximation of the relationship between the decision variables and the aerodynamic indexes was then performed also for the Spalart-Allmaras and the Transition SST models. The results obtained are presented below. The metamodel built for the latter dataset was used within an optimisation loop, solving the same SOOP of the previous cases, to compare their solutions.

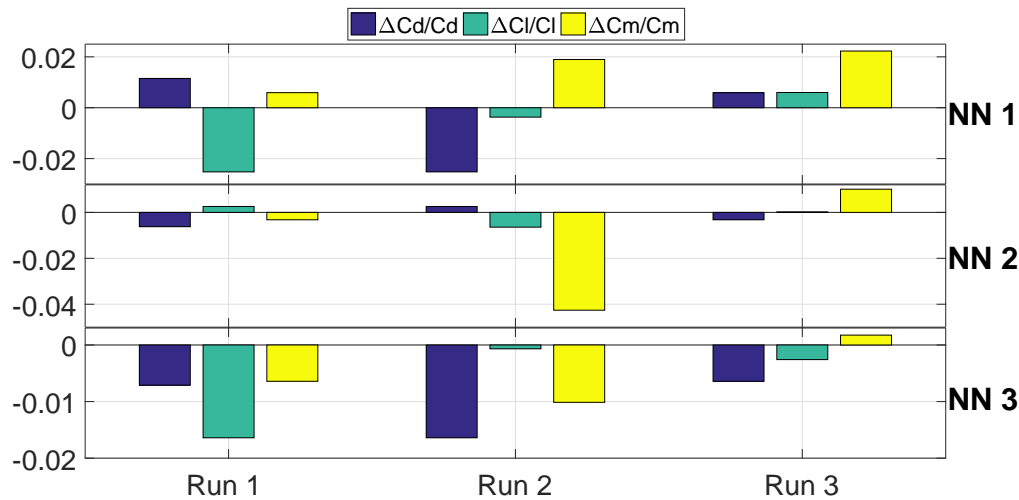


Figure 9.6: Comparison between the optimal solution provided using ANN and XFOIL, in terms of aerodynamic coefficients

9.4.1 Surrogate model for the Spalart-Allmaras dataset

The use of an ANN to fit the data coming from previous optimisations and computed by Fluent using the Spalart-Allmaras model revealed a poor accordance, inducing to try to several techniques in order to identify the one with more chance to represent the aerodynamic data. This activity can be summarized by the histograms of the relative error distribution between the predicted value and the real one, for each model tested. In Figure 9.7 this is illustrated for the drag value. Five models are present: an neural network (ANN), a radial basis function interpolation (RBFInterp) a radial basis function network (RBFNet), a support vector regression (SVR) and a Gaussian regression process (RGP). Among all, the RBF interpolation was the unique providing an exact fit for the training data and revealed to be the best, followed by the ANN, the SVR and the RBF network. RGP, instead, gave a very disperse result. Notice that only a portion of the interval where the relative errors felt are displayed, for the sake of clarity. Indeed, higher errors were obtained by all the models, though occurring for a very small amount of points in the dataset.

As regards the prediction of the lift and moment coefficient, the analogous histogram as before is shown in Figure 9.8. For the lift, the RBF network provided the best accordance, with the least dispersion, followed by the RBF interpolation and the RGP. The moment coefficient was replicated with a lower variability and most of the points felt between -2% and 2%, with only the RBF network error spread more evenly. Again, the RBF interpolation and the SVR gave a good response.




Figure 9.7: Histogram of relative error distribution for drag prediction of the SA using different models




Figure 9.8: Histogram of relative error distribution for lift and moment prediction of the SA using different models

9.4.2 Surrogate model for the Transition-SST dataset

The last set of data which was tried to be approximated by a regression model was that derived from the CFD computations with the Transition SST model. Here an ANN was used with 650 training points, since it worked well for XFOIL and the two solvers provided similar solutions for the SOOP. Other methods gave a poor accordance, leading to very disperse error distribution and unacceptable overfitting. Also in this case, an effort was dedicated into finding the best network architecture. It was noticed that in many cases the ANN output a constant response for some points at the bounds of the domain, for which either the drag was much higher and the lift was lower. Though normalization was performed both in pre and post processing, the networks appeared saturated most of the times, regardless the transfer function used. Each network was trained ten times, picking that giving the least MSE and number of relative error prediction out of a specified threshold of 10%. After many attempts, a network wa found

that was able to fit the data acceptably without exhibiting a flat response. It had 6 hidden layers of 8 neurons each. The model was then used into an hybrid optimisation loop, which will be presented below.

9.5 Meta-model assisted optimisation

A meta-model assisted optimisation loop was built in order to exploit the meta-model to perform a local refinement of the population. The loop had an hybrid architecture, coupling a GA as a local explorer of the design space based on external CFD computation and a surrogate model with a gradient based solver for a local improvement. The scheme of the loop is depicted in Figure 9.9. Generally, the model can be trained after performing some iterations with a standard global GA optimiser and exact fitness evaluation, in order to generate a dataset. After that, the model is built and at each generation the GA acts first, producing an offspring population. Each individual is locally refined by exploiting the surrogate model with a gradient based (GB) solver, e.g. the interior point algorithm. For improved robustness, the population is recomputed using the CFD and only those individual effectively improved substitute their original parent. This process is intended to use the GA for exploration, while exploiting its results for a local refinement based on gradient solvers, more suitable to deal with local minima.

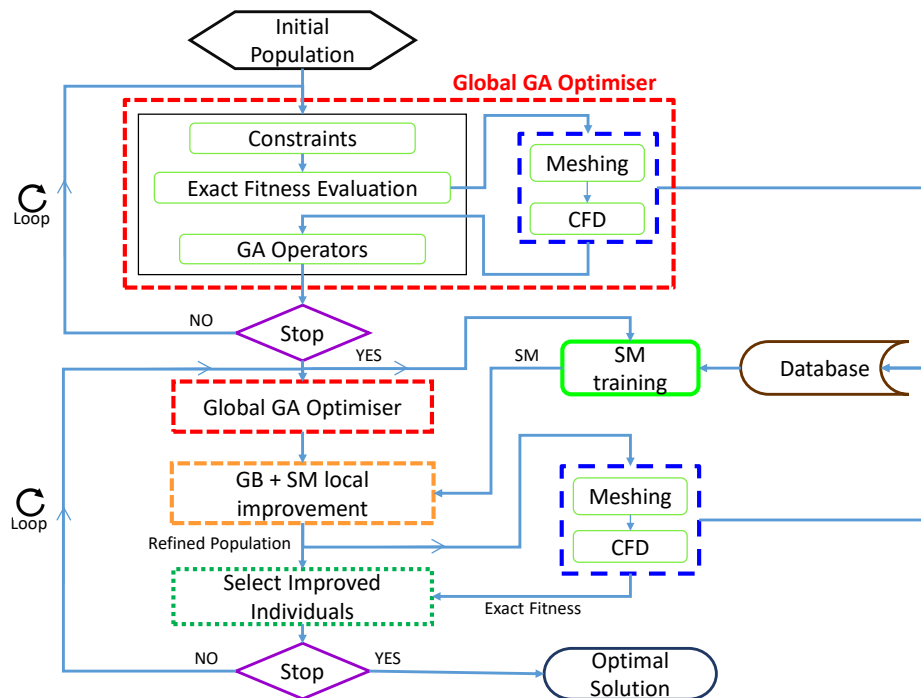


Figure 9.9: Scheme of the surrogate model assisted optimisation loop

The solution reached after 10 generations is presented in Figure 9.10. The shape was similar to that of Figure 7.18, with a more accentuated deformation. In terms of aerodynamic indexes, the results are reported in Table 9.2. The objective value was very close to the one originally found (refer to Table 7.7) and the same was also for the other indexes. Despite the higher number of CFD calls in each generation, since an individual refined by the gradient solver was re-evaluated with the exact fitness, the overall convergence time was lower, confirming the validity of the meta-model assisted procedure in leading the optimiser towards optimal points. The population diversity resulted preserved, because the refinement was performed in the neighbourhood of each initial point.




Figure 9.10: Baseline and optimised airfoil found with the meta-model assisted loop

| Index | Baseline | Optimised | Relative variation [%] |
|-------|----------|-----------|------------------------|
| C_d | 0.03418 | 0.02829 | -17.23 |
| C_l | 1.8895 | 2.0370 | 7.81 |
| C_m | -0.08617 | -0.0946 | 8.90 |
| L/D | 55.28 | 72.00 | 30.24 |

Table 9.2: Comparison between the aerodynamic coefficients of the baseline and the optimised airfoil, found with the metamodel assisted loop

Chapter 10

Conclusions

The study provided an insight into morphing leading edge aerofoils, highlighting some aspects which are required to be carefully inspected and proposing an optimisation framework which has been tested in a variety of situations to discover its drawbacks and validate the procedure. The parameterization procedure developed proved to be able to keep the constant arc length all over the optimisation, enforcing a constraint which otherwise would be hardly respected, since only a precise implicit relationship between the curvature correction coefficients must be respected, in order for it to happen, without the aforementioned procedure. Though imposing a zero length change can effectively limit the axial stress, the bending component due to the curvature variation too should be considered. Because of the absence of a reference structural configuration this was not included, although the availability of an explicit function describing each side of the aerofoil, for which derivatives can be computed analytically, allow for a quite easy implementation of a curvature constraint.

The influence of the aerodynamic model revealed itself to be crucial in determining the best solution and the inclusion of a transition prediction in the model, which happened for XFOIL and the Transition SST model, appeared to be the main cause of the difference in the shapes indicated by them and the Spalart-Allmaras model for the single-objective optimisation. In a 3D fashion, the minimisation of the drag coefficient was hampered by the constraint on the lift, required not to diminish with respect to the baseline configuration. Since this did not occur during the optimisation, the initial shape was already the best feasible solution found. In order to investigate this, a multi-objective problem was solved in 2D, aiming at minimising the drag and maximising the lift under the same boundary condition. The two objectives did not result conflicting, providing a unique solution. Coming back to the 3D case, the same MO problem was solved and the partial result obtained after 10 generations showed again that the baseline wing exhibited the highest lift, while the best drag decrease was of -3.73%. The modest improvement obtained with the 3D optimisation can be due to an effective different behaviour, for which a drag decrease is linked to a lift penalty, or to the inadequacy of the aerodynamic model in describing the physics involved. Despite this result, the activity was mainly intended to test the procedure in 3D and provided useful information that

can be exploited to further develop the process.

A topic studied in the thesis was the effectiveness of the optimiser and possible improvements of the loop. The engines natively present in Matlab R2015 b are general purposes tools, which are not thought specifically for engineering problems, though widely used for this purpose. In the single-objective optimisations solved, the convergence was more difficult when employing the CFD model. The global search capability of the algorithm was evident, since it led towards the same point in all the several runs performed, but a local refinement of the solution required quite an effort. A test with the Particle Swarm algorithm in combination with the Spalart-Allmaras model revealed a faster convergence and the use of another evolutionary paradigm is interesting to be studied. The NSGA-II for the multi-objective optimisations showed a certain difficulty in individuating the dominant solutions particularly for the case study of the multi-element high lift device optimisation. Also for this purpose more advanced tools are available, specifically intended to attain the best compromise between exploration and exploitation in a controlled way and the application of another engine to the problem solved deserves to be further examined. The use of meta-modelling was also enquired in the study, as a possible improvement to the framework. For the approximation of the aerodynamic performance computed by XFOIL and the Transition SST model an artificial neural network proved to be suitable and led to a solution close to that originally found without any surrogate model. In order to promote an easier refinement capability of the loop, a surrogate model assisted optimisation was performed, in which a gradient base solver exploited the surrogate model for a local refinement of each individual at each GA generation. This procedure was able to generate a solution close to that originally found with an exact fitness evaluation, but in a overall lower computational effort.

The results provided by this study demonstrates the effectiveness of a morphing leading edge aerofoil in substantially reducing the drag coefficient in the 2D problems solved. The technology appears, from an aerodynamic point of view, suitable for the purpose of improving the efficiency and overcoming the drawbacks of traditional high-lift systems, though the lift generation in the cases analysed resulted not comparable to those systems. Several still unsolved problems have, so far, limited the practical implementation of the technology to prototypes or demonstrators. The future research in this topic should keep a global perspective and focus towards the aero-structural modelling of morphing devices, in order to avoid for a useless effort in optimising one aspect and obtaining a solution incompatible from another point of view. The design of natively morphing architectures will definitively enlarge the optimisation space, removing some existing constraints and allowing for a better exploitation of their benefits.

Bibliography

- [1] ACARE. *2008 addendum to the strategic research agenda*. ACARE, 2008.
- [2] Rafic Ajaj et al. "An integrated conceptual design study using span morphing technology". In: *Journal of Intelligent Material Systems and Structures* 25.8 (2014), pp. 989–1008.
- [3] Rafic M. Ajaj, Christopher S. Beaverstock, and Michael I. Friswell. "Morphing aircraft: The need for a new design philosophy". In: *Aerospace Science and Technology* 49 (2016), pp. 154–166. ISSN: 1270-9638. DOI: <http://dx.doi.org/10.1016/j.ast.2015.11.039>. URL: <http://www.sciencedirect.com/science/article/pii/S1270963815003831>.
- [4] G. Amiryants. "Selectively deformable structures for design of adaptive wing smart elements". In: *ICAS 2010*. 2010. ISBN: 1-61782-049-0.
- [5] Andre Arruda and Alvaro Abdalla. "Numerical and experimental investigation and optimization of a morphing airfoil". In: *13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*. Fort Worth, Texas, 2010. ISBN: 1-60086-954-8.
- [6] David Baker and Michael Friswell. "The Design of Morphing Aerofoils using Compliant Mechanisms". In: *19th International Conference on Adaptive Structures and Technologies*. Ascona, Switzerland, 2008.
- [7] Silvestro Barbarino et al. "A review of morphing aircraft". In: *Journal of Intelligent Material Systems and Structures* 22.9 (2011), pp. 823–877.
- [8] M.H. Beale, M.T. Hagan, and H.B. Demuth. *Neural Network Toolbox™ User's Guide*. The MathWorks Inc., 2016.
- [9] *Length-Constrained Bezier Curve Smoothing*. Jan. 27, 2013.
- [10] Michellet Brigh et al. *Development of advanced high lift leading edge technology for laminar flow wings*. 2013. ISBN: 1-62410-181-X.
- [11] Marco Bumazzi and Rolf Radespiel. "Design and Analysis of a Droop Nose for Coanda Flap Applications". In: *Journal of Aircraft* 51.5 (2014), pp. 1567–1579.
- [12] *A Study of the ST Parametrisation Characteristics*. San Antonio, Texas, 2009.

- [13] L.C. Chow, Knut Mau, and Hugues Remy. "Landing Gears and High Lift Devices Airframe Noise Research". In: *8th AIAA/CEAS Aeroacoustics Conference and Exhibit*. Ed. by AIAA. 2002. DOI: 10.2514/6.2002-2408. URL: <http://dx.doi.org/10.2514/6.2002-2408>.
- [14] J.B. Cole. "Variable-camber airfoil". Pat. US 4553722 A. US Patent 4,553,722. 1985. URL: <https://www.google.com/patents/US4553722>.
- [15] Alan D'Ánaz-ManrÁnquez et al. "A Review of Surrogate Assisted Multiobjective Evolutionary Algorithms". In: *Computational Intelligence and Neuroscience 2016 (2016)*. DOI: 10.1155/2016/9420460.
- [16] Roeland De Breuker. *LEATOP (Leading Edge Actuation Topology Design and Demonstration)*. Final report summary. European Commission, 2014. URL: http://cordis.europa.eu/result/rcn/147524_en.html (visited on 07/30/2016).
- [17] Alessandro De Gaspari. "Design, Manufacturing and Wind Tunnel Validation of an Active Camber Morphing Wing Based on Compliant Structures". In: *ICAST 2014*. The Hague, The Netherlands, 2014. ISBN: 1-62410-396-0.
- [18] *A Two Levels Approach for the Optimal Design of Morphing Airfoils*. Forth Worth, Texas, 2010.
- [19] Alessandro De Gaspari and Ricci Sergio. "Knowledge-Based Shape Optimization of Morphing Wing for More Efficient Aircraft". In: *International Journal of Aerospace Engineering* 2015 (2015).
- [20] Alessandro De Gaspari, Ricci Sergio, and Luca Riccobene. "Design, Manufacturing and Wind Tunnel Validation of an Active Camber Morphing Wing Based on Compliant Structures". In: *ICAST2014*. The Hague, The Netherlands, 2014.
- [21] Kalyanmoy Deb. *Multi-Objective Optimization Using Evolutionary Algorithms*. New York, NY, USA: John Wiley & Sons, Inc., 2001. ISBN: 047187339X.
- [22] Andrea Di Renzo. "Sviluppo di metodologie per il progetto di ali con profili a curvatura variabile". MSc thesis. Politecnico di Milano, 2011.
- [23] Mark Drela, ed. *XFOIL 6.9 User Primer*. 2001. URL: http://web.mit.edu/drela/Public/web/xfoil/xfoil_doc.txt (visited on 10/01/2016).
- [24] *XFOIL: An Analysis and Design System for Low Reynolds Number Airfoils*. University of Notre Dame, June 1989, 1989.
- [25] Sha Du and Haisong An. "Design and feasibility analyses of morphing airfoil used to control flight attitude". In: *Strojnikovski vestnik* 58.1 (2012), pp. 46–55.
- [26] Laurene Fausett. *Fundamentals of Neural Networks. Architectures, Algorithms and Applications*. Florida Institute of Technology, 1994.
- [27] J.H.S. Fincham and M.I. Friswell. "Aerodynamic optimisation of a camber morphing aerofoil". In: *Aerospace Science and Technology* 43 (2015), pp. 245–255. ISSN: 1270-9638. DOI: <http://dx.doi.org/10.1016/j.ast.2015.02.023>. URL: <http://www.sciencedirect.com/science/article/pii/S1270963815000802>.

- [28] ANSYS Fluent. *Fluent 15.0 User's Guide*. ANSYS Inc., 2013.
- [29] Michael Friswell. "Morphing Aircrafts: An Improbable Dream?" In: *ASME 2014 Conference on Smart Materials, Adaptive Structures and Intelligent System*. Vol. 1. Newport, Rhode Island, USA: ASME, 2014. ISBN: 0-7918-4614-8.
- [30] Sebastian Geier et al. "Experimental and finite element analyses of multifunctional skins for morphing wing applications". In: *Proc. SPIE 9803, Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems 2016*. 2016. URL: <http://dx.doi.org/10.1117/12.2219357>.
- [31] Helen Heap and Bill Crowther. "A review of current leading edge device technology and options for innovation based on flow control". In: *RAes Conference*. London, 2003.
- [32] Andreas Hovelmann. "Aerodynamic Investigations of Noise-Reducing High-Lift Systems for Passenger Transport Aircraft". PhD thesis. KTH, 2011.
- [33] P. Iannelli et al. "Design of a High-Lift System for a Laminar Wing". In: *5TH EUROPEAN CONFERENCE FOR AERONAUTICS AND SPACE SCIENCES (EU-CASS)*. 2013.
- [34] Adam Jirasek and Olivier Amoigno. "Design of a High-Lift System with Droop Nose Device". In: *Journal of Aircraft* 46.2 (2009), pp. 731–735.
- [35] James J. Joo et al. "Variable Camber Compliant Wing - Design". In: *23rd AIAA/AHS Adaptive Structures Conference*. Ed. by AIAA. Kissimmee, Florida, 2015. doi: 10.2514/6.2015-1050. URL: <http://dx.doi.org/10.2514/6.2015-1050>.
- [36] D. Kelth Walters and D. Cokijat. "A Three-Equation Eddy-Viscosity Model for Reynolds-Averaged Navier–Stokes Simulations of Transitional Flow". In: *Journal of Fluids Engineering* 130 (2008). doi: 10.1115/1.2979230.
- [37] M. Kintscher et al. "Low Speed Wind Tunnel Test of a Morphing Leading Edge". In: *Italian Association of Aeronautics and Astronautics XXII Conference*. Napoli, Italia, 2013.
- [38] Markus Kintscher, Hans Monner, and Olaf Heintze. "Experimental testing of a smart leading edge high lift device for commercial transportation aircrafts". In: *ICAS 2010*. 2010. ISBN: 1-61782-049-0.
- [39] Markus Kintscher et al. "Design of a smart leading edge device for low speed wind tunnel tests in the European project SADE". In: *International Journal of Structural Integrity* 2.4 (2011), pp. 383–405.
- [40] Johannes Kirn and Stefan Storm. "Kinematic solution for a highly adaptive droop nose". In: *ICAST 2014*. The Hague, The Netherlands, 2014.
- [41] *CST Universal Parametric Geometry Representation Method with Application to Supersonic Aircraft*. Sendai, Japan, 2007.
- [42] Brenda M. Kulfan and John E. Bussoletti. "Fundamental Parametric Geometry Representations for Aircraft Component Shapes". In: *AIAA Paper 2006-6948* (2006), pp. 1–45.

- [43] *A Surface Parameterization Method for Airfoil Optimization and High Lift 2D Geometries Utilizing the CST Methodology*. Orlando, Florida, 2009. URL: <http://works.bepress.com/ddmarsha/21/>.
- [44] Robin Blair Langtry. "A Correlation-Based Transition Model using Local Variables for Unstructured Parallelized CFD codes". PhD thesis. Universität Stuttgart, 2006.
- [45] DongSeop Lee et al. "Multi-objective design optimization of morphing UAV aerofoil/wing using hybridised MOGA". In: *WCCI 2012 IEEE World Congress on Computational Intelligence*. Brisbane, Australia, 2012. ISBN: 1-4673-1509-5.
- [46] Kerr-Jia Lu. "Synthesis of Shape Morphing Compliant Mechanisms". PhD thesis. University of Michigan, 2004.
- [47] P. B. Martin et al. "Dynamic Stall Measurements and Computations for a VR-12 Airfoil with a Variable Droop Leading Edge". English. In: *American Helicopter Society 59th Annual Forum*. ID: 832030453. Phoenix, Arizona, 2003.
- [48] Andrea Massaro, Rita Ponza, and Ernesto Benini. *Multi-Objective and Multi-Point Optimization of a Multi-Element Airfoil Using a Surrogate-Assisted Evolutionary Algorithm*. Università degli Studi di Padova, 2009.
- [49] Robert J. McGhee et al. *Low-speed aerodynamic characteristics of a 17-percent-thick airfoil section designed for general aviation applications*. Technical Note NASA TN D-7428. NASA, 1973.
- [50] F.R. Menter, R. Langtry, and S. Völker. "Transition Modelling for General Purpose CFD Codes". In: *Flow Turbulence Combust* 77 (2006), pp. 277–303. DOI: 10.1007/s10494-006-9047-1.
- [51] Giulio Molinari. "Multidisciplinary Optimization of Morphing Wings with Distributed Compliance and Smart Actuation". Dr. Sc. ETH Zurich, 2016. Chap. 7.
- [52] H.P. Monner et al. "Design of a smart droop nose as leading edge high lift system for transportation aircrafts". In: *50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. Palm Springs, California, 2009. ISBN: 1-56347-973-7.
- [53] R. Morishima, S. Guo, and S. Ahmed. "A composite wing with a morphing leading edge". In: *51st AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. Orlando, Florida, 2010. ISBN: 1-60086-742-1.
- [54] Ryoko Morishima. "Analysis of composite wing structures with a morphing leading edge". PhDn thesis. Cranfield University, 2011.
- [55] P. Salunke Nilesh, Ahamad R.A. Juned, and S.A. Channiwala. "Morphing skins". In: *American Journal of Mechanical Engineering* 2.4 (2014). Airfoil Parameterization Techniques: A Review, pp. 99–102. DOI: 10.12691/ajme-2-4-1. URL: <http://pubs.sciepub.com/ajme/2/4/1>.
- [56] George M. Phillips. *Interpolation and Approximation by Polynomials*. 1st ed. Springer-Verlag New York, 2003. Chap. 7. ISBN: 978-1-4419-1810-9. DOI: 10.1007/b97417.

- [57] Les Piegl and Wayne Tiller. *The NURBS Book*. Springer-Verlag, 1996. ISBN: 3-540-61545-8.
- [58] Michael Pott-Pollenske, Jochen W. Wild, and Lothar Bertsch. "Aerodynamic and Acoustic Design of Silent Leading Edge Device". In: *20th AIAA/CEAS Aeroacoustics Conference*. Ed. by AIAA Aviation. 2014. DOI: 10.2514/6.2014-2076. URL: <http://dx.doi.org/10.2514/6.2014-2076>.
- [59] Jinhao Qiu et al. "Mechanics for the World: Proceedings of the 23rd International Congress of Theoretical and Applied Mechanics, ICTAM2012 Smart Skin and Actuators for Morphing Structure". In: *Procedia IUTAM* 10 (2014), pp. 427–441. ISSN: 2210-9838. DOI: <http://dx.doi.org/10.1016/j.piutam.2014.01.037>. URL: <http://www.sciencedirect.com/science/article/pii/S2210983814000388>.
- [60] Laxminarayana Saggere and Sridhar Kota. "Design of adaptive structures using compliant mechanisms". In: vol. 3329. 1998, pp. 672–676. DOI: 10.1117/12.316940. URL: <http://dx.doi.org/10.1117/12.316940>.
- [61] Jamshid A. Samareh. "A Survey of Shape Parametrisation Techniques". In: *NASA/CP-1999 209136* (1999), pp. 333–344.
- [62] M. Santer and S. Pellegrino. *Topological optimization of compliant adaptive wing structure*. ID: 700244697. 2009.
- [63] Hermann Schlichting. *Boundary-Layer Theory*. 7th ed. McGraw-Hill, 1979.
- [64] Gallart Marc Secanell. "Development of a design tool for aerodynamic shape optimization of airfoils". MASC thesis. University of Victoria, 2004.
- [65] Aravind Seshadri. *Multi-objective Optimization Using Evolutionary Algorithms (MOEA)*. Article.
- [66] A. Shmilovich. "Wing leading edge concepts for noise reduction". In: *27th International Congress of Aeronautical Science - ICAS 2010*. 2010.
- [67] Michael Sinapius et al. "DLR's Morphing Wing Activities within the European Network". In: *Procedia IUTAM* 10 (2014), pp. 416–426.
- [68] A.M.O. Smith and Nathalie Gamberoni. *Transition, Pressure Gradient and Stability Theory*. Report ES 26388. Douglas - Aircraft Division, 1956.
- [69] D.D. Smith et al. "Computational and experimental validation of the active morphing wing". In: *Journal of Aircraft* 51.3 (2014), pp. 925–937. URL: <http://dx.doi.org/10.2514/1.C032262>.
- [70] Jurij Sodja et al. "Experimental evaluation of the morphing leading edge concept". In: *23rd AIAA/AHS Adaptive Structures Conference*. Kissimmee, Florida, 2015. ISBN: 1-62410-346-4.
- [71] P. Spalart and S. Allmaras. *A One-Equation Turbulence Model for Aerodynamic Flows*. Technical Report AIAA-92-0439. AIAA, 1992.

- [72] H. Strüber. "The aerodynamic design of the A350 XWB-900 high lift system". In: *29th Congress of the International Council of the Aeronautical Sciences*. St. Petersburg, Russia, 2014.
- [73] Rujie Sun et al. "Multidisciplinary design optimization of adaptive wing leading edge". In: *Science China Technological Sciences* 56.7 (2013), pp. 1790–1797.
- [74] Hiroharu Suzuki, Kenichi Rinoie, and Asei Tezuka. "Laminar Airfoil Modification Attaining Optimum Drag Reduction by Use of Airfoil Morphing". In: *Journal of Aircraft* 47.4 (2010), pp. 1126–1132.
- [75] C. Thill et al. "Morphing skins". In: *The Aeronautical Journal* 112.1129 (Mar. 2008). Publisher: Royal Aeronautical Society, pp. 117–139. issn: 0001-9240.
- [76] G.A.A. Thuwis. "Stiffness and Layout Tailoring of a Morphing High-lift System with Aeroelastic Loads". PhD thesis. TU Delft University, 2012.
- [77] J.L. Van Ingen. *A suggested Semi-Empirical Method for the Calculation of the Boundary Layer Transition Region*. Report VTH-74. Technical University of Delft, 1956.
- [78] *The eN method for transition prediction. Historical review of work at TU Delft*. Seattle, Washington, 2008.
- [79] Srinivas Vasista, Johannes Riemenschneider, and Hans Monner. *Design and testing of a compliant mechanism-based demonstrator for a droop-nose morphing device*. AIAA, 2015. ISBN: 1-62410-346-4.
- [80] Srinivas Vasista et al. "Compliant structures-based wing and wingtip morphing devices". In: *Aircraft Engineering* 88.2 (2016), pp. 311–330.
- [81] Terrence A. Weisshaar. "Morphing Aircraft Technology - New Shapes for Aircraft Design". In: *RTO-MP-AVT-141*. Ed. by NATO. Neuilly-sur-Seine, France, 2006.
- [82] Joche Wild, Michael Pott Pollenske, and Bj Nagel. *An integrated design approach for low noise exposing high-lift devices*. AIAA, 2006. ISBN: 1-56347-813-7.
- [83] Benjamin K.S. Woods, Onur Bilgen, and Michael I. Friswell. "Wind tunnel testing of the Fish Bone Active Camber morphing concept". In: *Journal of Intelligent Material Systems and Structures* (2014). doi: 10.1177/1045389X14521700. eprint: <http://jim.sagepub.com/content/early/2014/02/04/1045389X14521700.full.pdf+html>. URL: <http://jim.sagepub.com/content/early/2014/02/04/1045389X14521700.abstract>.
- [84] *Aerodynamic Modelling of the Fish Bone Active Camber Morphing Wing Concept*. Bristol, UK, 2014.
- [85] Benjamin K.S. Woods and Michael I. Friswell. "Multi-objective geometry optimization of the Fish Bone Active Camber morphing airfoil". In: *Journal of Intelligent Material Systems and Structures* (2015). doi: 10.1177/1045389X15604231. eprint: <http://jim.sagepub.com/content/early/2015/09/15/1045389X15604231.full.pdf+html>. URL: <http://jim.sagepub.com/content/early/2015/09/15/1045389X15604231.abstract>.

- [86] Benjamin K.S. Woods and Michael I. Friswell. "The Adaptive Aspect Ratio morphing wing: Design concept and low fidelity skin optimization". In: *Aerospace Science and Technology* 42 (2015), pp. 209–217. ISSN: 1270-9638. DOI: <http://dx.doi.org/10.1016/j.ast.2015.01.012>. URL: <http://www.sciencedirect.com/science/article/pii/S1270963815000279>.
- [87] Tong Xinxing et al. "Topology optimization of compliant adaptive wing leading edge with composite materials". In: *Chinese Journal of Aeronautics* 27.6 (2014), pp. 1488–1494.
- [88] Tomohiro Yokozeki, Aya Sugiura, and Yoshiyasu Hirano. "Development of Variable Camber Morphing Airfoil Using Corrugated Structure". In: *Journal of Aircraft* 51.3 (2014), pp. 1023–1029. DOI: 10.2514/1.C032573. URL: <http://dx.doi.org/10.2514/1.C032573>.
- [89] Feng Zhu. "Geometry Parametrisation and Aerodynamic Shape Optimisation". PhD thesis. University of Sheffield.
- [90] Feng Zhu and Ning Qin. "Intuitive Class/Shape Function Parameterization for Airfoils". In: *AIAA Journal* 52.1 (2013), pp. 17–25. ISSN: 0001-1452. DOI: 10.2514/1.J052610. URL: <http://dx.doi.org/10.2514/1.J052610>.
- [91] H.D.I. Zimmer. "Quertriebskoerper mit veraenderbarer profilierung, insbesondere flugzeugtragfluegel". Pat. DE2907912 A1. DE Patent App. DE19,792,907,912. 1980. URL: <https://www.google.com/patents/DE2907912A1?cl=un>.