Politecnico di Torino

von KARMAN INSTITUTE FOR FLUID DYNAMICS

Politecnico di Torino

Master's Degree in Aerospace Engineering

Academic year 2025/2026

Low-Reynolds Airfoil Characterization for BEMT-Based Propeller Modeling

March 2026

Supervisors:
Prof. Miguel Alfonso Mendez
Prof. Andrea Ferrero

Candidate:
Flavio Bordin - 333049

Abstract

Small unmanned aerial vehicles (sUAVs), and especially multi-rotor crafts, have seen an acceleration in technological development in recent years, as adoption in civilian contexts increases. Improving their performance and controllability requires effective control algorithms running onboard, which in turn depend on light and accurate models for the drones' propellers' aerodynamic characteristics. Small drones' propellers operate at low Reynolds numbers ($10^4 < Re < 10^5$), where the physics is very sensitive to geometry and inflow conditions. Furthermore, rapid maneuvers and wind gusts can introduce unsteady aerodynamic effects relevant for the whole drone's flight dynamics. High-fidelity methods like computational fluid dynamics can capture these complex phenomena, but their high computational cost makes them unsuitable for real-time analysis. Thus, although limited by its simplifying assumptions, a BEMT-based reduced model of the propeller is a promising candidate for low-cost prediction for drone control purposes. The present work explores the limitations of a drone's airfoils' numeric characterization and their influence on prediction accuracy when used in a BEMT-based solver in low Reynolds conditions, assuming steady, uniform inflow. Both XFOIL and Computational Fluid Dynamics (CFD) simulations were employed to compute the polars of the T1045-II propeller's airfoils; these were then employed in a BEMT solver specifically implemented for this work. In order to assess the tools' reliability in challenging conditions, the results from the full methodology (i.e. BEMT coupled with numerically computed polars) are compared to experimental data. This study shows that BEMT is a computationally efficient model for predicting propeller performance, but its accuracy is strongly influenced by the model's assumptions and by the quality of the airfoils' polars, whose accurate calculation is particularly problematic for small propellers. Both methods yielded satisfactory polar predictions on a documented test case, but revealed unexpected discrepancies on the propeller's own sections, particularly at low incidence, where differing resolution of a laminar separation bubble on the pressure side drives a significant offset in lift and, to a lesser extent, drag. At the propeller level, thrust

predictions agreed well with experiments regardless of the polar source, while torque was systematically underestimated, pointing to the role of three-dimensional rotational effects and wake modeling limitations beyond the reach of the present framework.

Acknowledgements

I would like to express my sincere gratitude to Prof. Miguel Alfonso Mendez for his guidance throughout this thesis, and for providing me with a competent and available team at the von Karman Institute for Fluid Dynamics, without whom this work would not have been possible. I am grateful to Prof. Andrea Ferrero for his valuable pointers on the CFD methodology, which proved essential in orienting the numerical approach. A special thank you goes to Dr. Manuel Ratz for his insight and guidance on the geometry extraction procedure. I would also like to thank Dr. Romain Poletti for his patient help and expertise with CFD and OpenFOAM throughout the work. PhD candidate Yannick Lecomte deserves my sincere gratitude for his suggestions on the BEMT implementation, his careful reading of the manuscript, and his general guidance over the course of this project. Last but not least, a thank you to Dr. Alessandro Grava, whose expertise on propeller aerodynamics proved greatly helpful in this work.

Table of Contents

List of Figures

Chapter 1

Introduction

If you asked someone to imagine the future of humanity, chances are they would describe the landscape of a city populated by vertiginous skyscrapers, inhabited by people and humanoid robots alike, infested by flying vehicles of all sizes breezing through the buildings at high speed like bees in a hive. This, or the extinction of humankind.
In the event we manage to avoid the latter, we will need to overcome many technological challenges. One of these is the development of autonomous drones, in particular small unmanned aerial vehicles (sUAV), which promise unprecedented versatility in a variety of applications.
In recent years, the sUAV sector has attracted significant investments worldwide as it transitions from its hobbyist or defense-related roots to commercial and industrial use. Advancements in various fields have enabled drones technology to reach the accelerated pace of development witnessed recently, however many challenges still stand between drones and their widespread adoption in our society.

Among these challenges, flight endurance and operation in adverse weather conditions are still limiting widespread adoption of drones. In terms of updates to current drone platforms, new solutions are being explored at the Von Karman Institute for Fluid Dynamics (VKI). These need light and reliable computational models of the aerodynamics of the drones' propellers to operate efficiently. In particular, the models need to be capable of capturing the complex phenomena related to the low-Reynolds regime and unsteadiness that dictate propeller performance.

1.1 Context

The term "small unmanned aerial vehicle" generally refers to an unmanned aircraft with a maximum takeoff mass below 25 kg. These vehicles can be divided into groups based on their architecture (Fig. 1.1): fixed-wing sUAVs have the architecture of a conventional airplane and use the propellers to generate thrust, they require a conventional takeoff to operate; hybrid fixed-wing sUAVs have again a base conventional structure but use added propellers to generate the lift needed for vertical takeoff; single rotor architectures use the same structure of helicopters; finally, multi-rotor sUAVs use multiple vertically-oriented rotors for maneuvers.

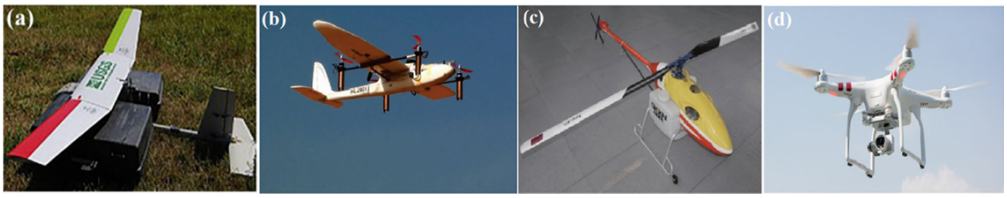

Figure 1.1: sUAV architectures: (a) fixed-wing, (b) hybrid fixed-wing, (c) single rotor, (d) multi-rotor (Credit: [1]).

The sector of drones is attracting high investments worldwide. A 2022 market report [2] estimated the European quadcopter market at 7.2 billion euros in 2023, making up a third of the global market, and forecasts a 14.6% average annual growth rate for the following decade. Although drones have a long history of military use, they are increasingly used in civilian applications [3].

This expansion has been possible thanks to developments in various enabling technologies. Lithium-ion batteries provide the energy density and discharge rates required to complete most operations, while cheap and miniaturized MEMS sensors enable efficient data acquisition from the drone platform. Together with progress in control algorithms, positioning systems, and communication technologies, these developments have made reliable remote operation possible.

The main civilian applications of drones described in the literature [3, 1] include monitoring and inspection, logistics and photography and image acquisition. In particular, monitoring and inspection operations, such as in infrastructure, precision agriculture, environmental monitoring and disaster response, benefit from the ability of sUAVs to access remote or hazardous locations and acquire data once there.

Despite their fast development, sUAVs still face important challenges. Flight endurance is still limited by battery capacity and energy density [1], while adverse weather conditions such as wind and rain impede the drones' flight. These limitations still affect sUAVs integration and thus require technological developments to be overcome.

1.2 Motivation

Small unmanned aerial vehicles still face important technological challenges. Among these, operation in adverse weather conditions — wind gusts, turbulent inflow, rapid maneuvers — remains a central limitation, as these disturbances directly degrade flight performance and stability. Overcoming them requires effective onboard control algorithms capable of responding in real time. In the case of propeller-driven platforms, these algorithms depend on fast and accurate models of propeller aerodynamics: the model must be light enough to run continuously onboard, yet faithful enough to be useful.

Accurate prediction of small propeller performance is, however, a genuinely difficult problem. The blade sections of small propellers operate at low Reynolds numbers ($Re < 10^5$), a regime in which the flow is dominated by complex viscous phenomena — laminar separation, transition, recirculation bubbles — that are highly sensitive to geometry, inflow conditions, and turbulence levels [4, 5]. Small changes in operating conditions can produce large changes in aerodynamic behaviour, making reliable prediction challenging even with high-fidelity tools. This difficulty is further compounded by the computational constraints of onboard operation: full three-dimensional CFD and other high-fidelity approaches can capture these phenomena accurately, but at a cost that is orders of magnitude beyond what real-time use allows. Intermediate-fidelity methods exist but often remain either too expensive.

Blade Element Momentum Theory (BEMT) emerges as a promising candidate for onboard propeller modeling [6]. By decomposing the blade into independent radial elements and describing each through two-dimensional airfoil aerodynamics, BEMT reduces the propeller characterization to a lightweight calculation that can be evaluated at negligible cost once the blade's sections aerodynamic input

data, hereafter also referred to as polar maps, have been precomputed offline. This separation of offline preparation and online evaluation makes BEMT well suited to the computational constraints of onboard control. However, the model carries its own limitations: its simplifying assumptions restrict its physical fidelity, and its predictions are only as reliable as the polar maps it is fed. In the low-Reynolds regime, where airfoil aerodynamics are particularly difficult to characterize, the quality of these polar maps is far from guaranteed.

The objective of the present work is to investigate these limitations on a real small propeller. Specifically, it asks how well low-Reynolds airfoil polar maps can be computed numerically, and how much the accuracy of that data influences the propeller performance predicted by BEMT. This study treats the data's reliability as the central question, assessing the consequences of polar uncertainty at the propeller level through comparison with experimental measurements.

The relevance of this question extends beyond the specific case studied here. Variable-pitch propellers (VPPs), currently being explored as a means of enhancing responsiveness and control authority in small drones, introduce additional aerodynamic complexity and place even greater demands on the underlying propeller model. More broadly, at the Von Karman Institute for Fluid Dynamics (VKI), the RE-TWIST project [7] is developing model-based control agents for engineering systems [8] that require lightweight, reliable propeller models such as the BEMT. The present work contributes a necessary step toward both of these directions by establishing what current numerical approaches can and cannot deliver in the low-Reynolds propeller aerodynamics context.

Chapter 2

Theoretical Background

This chapter provides the theoretical background underlying both the physics of the problem and the numerical methods employed in this work.

The first part addresses the aerodynamics of airfoils at low Reynolds numbers. The boundary layer and its chordwise evolution are discussed first, as they govern the onset of the complex viscous phenomena characteristic of this regime. Particular attention is then given to laminar separation bubbles — their structure, dynamics, and dependence on Reynolds number — as these are the dominant flow feature affecting airfoil performance in the range of interest, and the primary source of difficulty for numerical prediction.

The second part presents the two numerical methods used for airfoil characterization. XFOIL, based on a viscous-inviscid interaction formulation with an integral boundary layer model, is introduced first as the computationally efficient baseline approach. The unsteady Reynolds-Averaged Navier-Stokes (URANS) framework, closed with the k-ω SST turbulence model and the γ-Re_θ transition model, is then presented as the higher-fidelity alternative. For both methods, the theoretical formulation is outlined alongside its known limitations in the low-Reynolds regime, providing the basis for interpreting their respective behaviors in the results.

The third part introduces the Blade Element Momentum Theory (BEMT). The model's formulation is developed from its two constituent theories, namely Momentum Theory and Blade Element Theory, and the corrective models adopted to extend its range of validity are described. Emphasis is placed on the assumptions

underlying the theory, which enable the great simplification of propeller aerodynamics operated by the model, and on its reliance on the polar maps of the blades' sections.

2.1 Low-Reynolds airfoil aerodynamics

Airfoil performance in a freestream current is strongly dependent on the Reynolds number, which quantifies the ratio of inertia to viscosity, here defined on the basis of chord length c as

$$Re = \frac{U_\infty c}{\nu} \,. \tag{2.1}$$

The blade sections of small propellers usually operate at $10^4 < Re < 10^5$, which is regarded as low-Reynolds regime for airfoils. In these conditions, disturbances in the boundary layer that develops on the airfoil are often damped by viscosity, impeding turbulent transition. This dynamic gives rise to a rich variety of phenomena related to the laminar boundary layer and its chordwise evolution, which fundamentally alter the flow on the airfoil and thus its performance. A qualitative schematic is provided in Fig. 2.1.

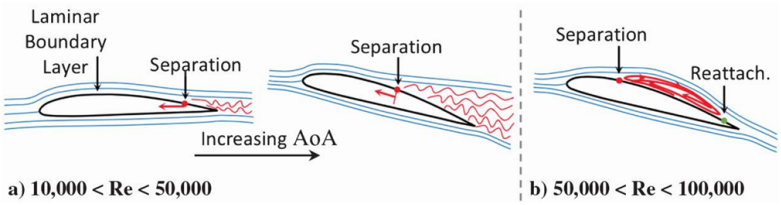

Figure 2.1: Qualitative illustration of flow over an airfoil in low-Reynolds conditions (Credit: [4]).

2.1.1 Boundary layer

The concept of boundary layer was first introduced by Prandtl in 1904. Specializing the Navier-Stokes equations for the boundary layer of a high-Reynolds, incompressible, steady, two-dimensional current flowing over a flat plate, Prandtl [9] arrived

at the following equations for mass conservation and momentum balance

$$\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} = 0 \ , \tag{2.2}$$

$$u\frac{\partial u}{\partial x} + v\frac{\partial u}{\partial y} = -\frac{1}{\rho}\frac{dp_e}{dx} + \frac{1}{\rho}\frac{\partial \tau_{xy}}{\partial y} \ , \tag{2.3}$$

where $p_e(x)$ is the pressure in the outer flow and τ_{xy} is the fluid's viscous shear stress term. An important conclusion is that pressure in the boundary layer is constant along the direction normal to the wall (y) and it is thus imposed by the outer inviscid flow as $p_e(x)$. The external inviscid flow drives the boundary layer dynamics through a pressure gradient dp_e/dx, which is a parameter of fundamental importance [10].

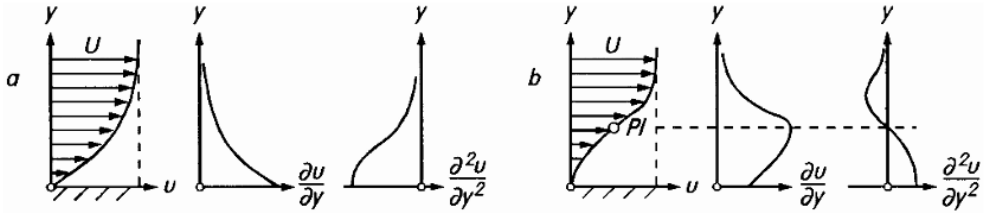

Figure 2.2: Velocity profile and derivatives in the boundary layer: (a) pressure drop, (b) pressure increase (Credit [10]).

Indeed, in the case of a steady current of a Newtonian fluid, the curvature of the velocity profile at the wall is directly tied to the outer pressure gradient through viscosity

$$\mu\frac{\partial^2 u}{\partial y^2}(x,0) = \frac{dp_e}{dx} \ . \tag{2.4}$$

The shape of the velocity profile, and thus the onset of boundary layer separation, are directly controlled by the outer pressure gradient. Separation is of central importance in low-Reynolds airfoil aerodynamics, as it is the mechanism through which the complex phenomena discussed in the following arise.

At the wall, viscosity leeches momentum from the fluid, and velocity profiles assume shapes like the one in Fig. 2.2. Momentum is ultimately extracted from the outer flow. When the pressure gradient is favorable ($dp_e/dx < 0$), the fluid near the wall gains momentum, this balances viscosity and leads to a stable configuration

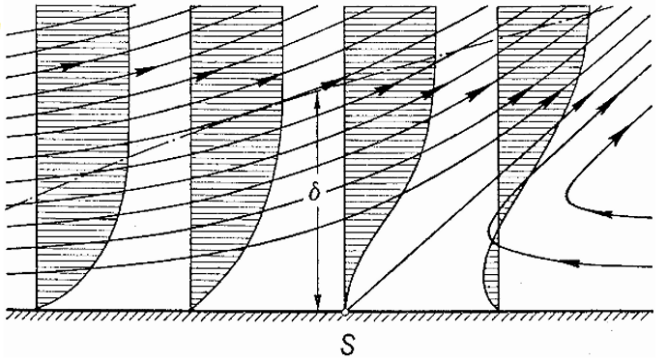

Figure 2.3: Schematic of boundary layer separation (Credit: [10]).

(Fig. 2.2(a)). However, with an adverse pressure gradient ($dp_e/dx > 0$) the fluid at the wall is slowed down further and takes the shape presented in Fig. 2.2(b). Having lost momentum to viscosity, the fluid can come to a stop and begin flowing in the opposite direction, leading to separation of the boundary layer (Fig. 2.3).

$$\text{Separation:} \qquad \frac{\partial u}{\partial y}(x,0) = 0 \ , \tag{2.5}$$

$$\text{Reverse flow:} \qquad \frac{\partial u}{\partial y}(x,0) < 0 \ . \tag{2.6}$$

These considerations are valid for both laminar and turbulent boundary layers, although there are relevant differences. In laminar boundary layers, the only mechanism that operates momentum transport from the outer flow is viscosity, which acts at short range. The fluid at the wall is thus slower and, although this grants less viscous drag, it is prone to separation; indeed laminar boundary layers can only withstand very small adverse pressure gradients before separating. On the other hand, in turbulent boundary layers the mixing operated by turbulent motion contributes greatly to transferring momentum to fluid close to the wall. The flow is thus more energized with respect to the laminar regime, delaying separation.

Thus, the laminar-turbulent transition of the boundary layer is of fundamental importance in its dynamics. This phenomenon can be modeled with linear stability theory: a small disturbance is superimposed on a known laminar base flow, like the

boundary layer, and its evolution in time is studied by linearizing the Navier-Stokes equations; when the disturbance gets damped the base flow is stable and stays laminar, when it is amplified the base flow ends up transitioning to turbulent. Tollmien [11] applied this theory to Blasius' solution of the laminar boundary layer on a flat plate, the resulting stability map is depicted in Fig. 2.4, which details which disturbances are amplified or damped as a function of the momentum-thickness Reynolds number. This parameter is defined as

$$Re_\theta = \frac{U_\infty \theta}{\nu} \,, \tag{2.7}$$

where θ is the momentum thickness of the boundary layer, defined later in this section. The instability modes he had obtained are now deemed Tollmien-Schlichting (T-S) waves, which are of viscous nature and are the dominant cause of instability in low-disturbance external flows.

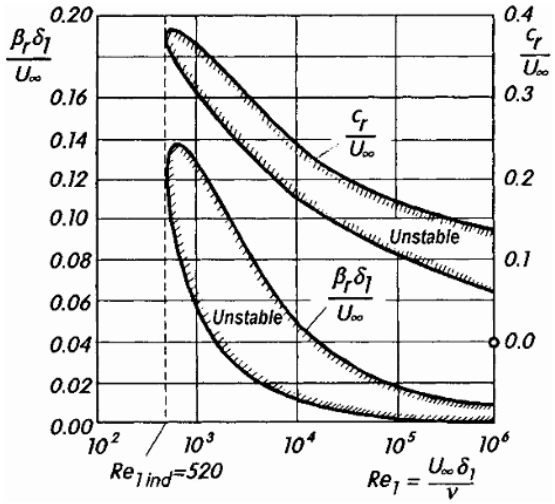

Figure 2.4: Stability map for Blasius' boundary layer (Credit: [10]).

We see that for $Re_\theta < Re_{\theta,ind} \approx 520$ small disturbances are damped, the boundary layer stays laminar. A general interpretation, valid for every kind of flow, is that at low Reynolds numbers, viscosity dominates over inertial effects, suppressing the growth of Tollmien–Schlichting waves as well as that of other small perturbations. This is the reason boundary layers at low Reynolds tend to stay laminar.

The gradient of the outer pressure dp_e/dx is very important for laminar-turbulent transition too: stability studies by Pretsch [12] on the Falkner-Skan velocity profile family[1] showed that adverse pressure gradients have a strong destabilizing effect, while favorable gradients stabilize the flow.

Another important parameter for stability is the level of disturbance in the outer flow, quantified by the turbulence intensity

$$Tu = \frac{\sqrt{\frac{2}{3}k}}{U_\infty} \,, \tag{2.8}$$

where k is the turbulent kinetic energy per unit mass. Experimental measurements on spheres show that Re_{crit} increases greatly with decreasing Tu [10].

The Reynolds number at which transition occurs is always greater than the indifference Reynolds number, $Re_{crit} > Re_{ind}$, as the disturbances take time to be amplified enough to provoke transition. Smith and Gamberoni [13] and van Ingen [14] independently proposed a simple rule to determine the point of transition: the e^N method. For a given velocity profile, the spatial amplification rate $-\alpha_i$ of the most unstable T-S wave is integrated along the chordwise direction from the point of neutral stability x_{ind} to obtain the total amplification factor

$$N(x) = \int_{x_{ind}}^{x} (-\alpha_i)\, dx' \,. \tag{2.9}$$

The method predicts transition when N reaches a critical threshold N_{crit}, which depends on the external disturbance conditions. Mack [15] established an empirical correlation between N_{crit} and the freestream turbulence intensity

$$N_{crit} = -8.43 - 2.4 \ln(Tu) \,, \tag{2.10}$$

so that high-turbulence conditions lead to earlier transition, that is, lower N_{crit}.

[1] This family of profiles arises as the analytical solution for the flow over a wedge, which is the theoretical canonical problem to study the pressure gradient's influence on boundary layer dynamics.

Integral quantities and the von Kármán equation

Sometimes it is useful to formulate the boundary layer equations in integral form, that is, employing integral quantities. The displacement thickness of the boundary layer measures the mass flux deficit of the boundary layer with respect to the outer flow. It is defined as

$$\delta^* = \int_0^\infty \left(1 - \frac{u}{U_\infty}\right) dy \, . \tag{2.11}$$

With the same reasoning, the momentum thickness of the boundary layer measures the momentum flux deficit, defined as

$$\theta = \int_0^\infty \frac{u}{U_\infty}\left(1 - \frac{u}{U_\infty}\right) dy \, , \tag{2.12}$$

which is directly related to viscous drag. The momentum balance equation (2.3) can be reformulated[2] in light of these definitions, yielding the von Kármán integral equation [16]

$$\frac{d\theta}{dx} + (2+H)\frac{\theta}{U_e}\frac{dU_e}{dx} = \frac{C_f}{2} \, . \tag{2.13}$$

where $C_f = \tau_w/(\frac{1}{2}\rho U_\infty^2)$ is the skin friction coefficient, and $H = \delta^*/\theta$ is the shape factor of the boundary layer. A similar integral equation can be obtained from the energy equation in compressible boundary layers. An important note is that these integral equations are valid for both laminar and turbulent boundary layers. They are generally easier to solve than the regular boundary layer equations and, although the solution doesn't provide the full field, the coefficients obtained themselves are often the objective of the analysis in the first place. One drawback is that these equations form an underdetermined system, which has to be closed with a parametric model of the velocity profile $u(y)$ to get an approximate solution. This approximate solution method is called *integral method*, it was first proposed by Pohlhausen in 1921 [17] but is still relevant and applied nowadays.

[2] In the derivation, the Bernoulli equation $-\frac{1}{\rho}\frac{dp_e}{dx} = U_e\frac{dU_e}{dx}$ is used to link the outer pressure to the outer velocity U_e.

2.1.2 Laminar separation bubble

When a laminar boundary layer separates, it generates a laminar shear layer at a small distance from the wall. If this shear layer then reattaches, a recirculation bubble is formed underneath: this structure is called a *laminar separation bubble* (LSB) (Fig. 2.5). Two main regions can be distinguished in the scheme in Fig. 2.5, divided by a streamline which begins at the separation point and ends at reattachment, and below which the net mass flow is zero everywhere. The underlying region is the bubble itself, and comprises slow recirculating fluid. The region above the streamline and limited above by the outer flow is the free shear layer resulting from the boundary layer separation.

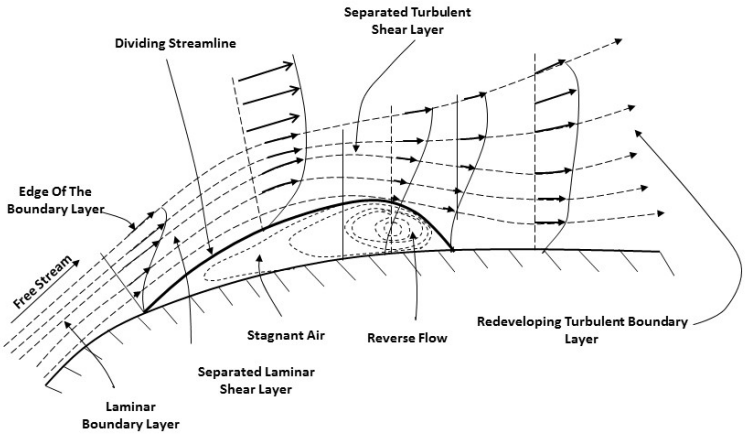

Figure 2.5: Schematic structure of a laminar separation bubble (Credit: [18]).

The fundamental phenomenon that enables reattachment, and thus the bubble formation, is the turbulent transition in the laminar shear layer; turbulence strongly increases momentum transport to fluid close to the wall and eliminates reverse flow. Viscous T-S waves originating in the laminar boundary layer have been found to be the primary instability cause in the shear layer, which is also sensitive to inviscid Kelvin-Helmholtz (K-H) instabilities (Fig. 2.6).

Laminar separation bubbles can greatly alter the entire flow around an airfoil, influencing performance, particularly in low Reynolds conditions. Short LSBs measure few percentage of chord length, they tend to develop at higher Reynolds numbers ($10^5 < Re < 5 \cdot 10^5$) and typically have a lesser influence on flow dynamics

Figure 2.6: Atmospheric Kelvin-Helmholtz instability (Ville-en-Hesbaye, Liège, Belgium; 7th February 2026, 12:11 CET).

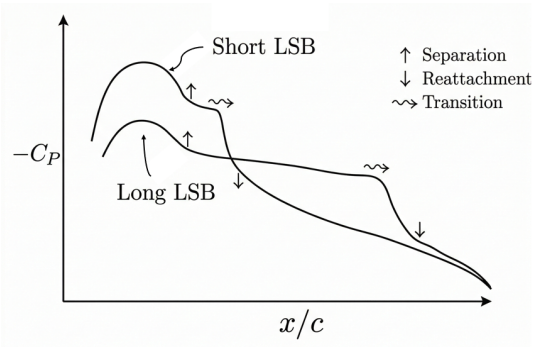

Figure 2.7: Scheme of pressure coefficient distribution on an airfoil suction side, as altered by a LSB.

and performance. Long LSBs on the other hand develop at lower Reynolds numbers ($10^4 < Re < 10^5$) and can cover a large portion of the chord, penalizing pressure drag. This dead air zone in the bubble fundamentally alters the pressure recovery of the outer flow above it, which experiences almost null deceleration and thus sees a constant pressure. The pressure is imposed on the slow-moving fluid below, we thus see an almost constant pressure distribution across the whole bubble, and, as a consequence, at the wall too (Fig. 2.7) [19].

As the Reynolds number decreases ($5 \cdot 10^4 < Re < 10^5$), viscosity damps instabilities more effectively, delaying separation, transition and reattachment, increasing the size of the LSB [4, 20]. In these conditions, an increase in incidence

causes the reattachment point to move downstream, enlarging the bubble.

In very low Reynolds number regimes ($10^4 < Re < 5 \cdot 10^4$), the separation point moves downstream because of the increased stability of the laminar boundary layer, but for the same reason the resulting shear layer is unable to transition and reattach, the airfoil is thus in trailing edge stall for most of its operational range. As the angle of attack increases, the separation point moves upstream, leading to a bigger stall and worse performance.

Across both regimes, high freestream turbulence intensity can directly cause transition in the shear layer and thus create shorter LSBs and milder, delayed stall, improving performance [21, 20]. However, at low turbulence intensity levels ($Tu < 1\%$) the instability mechanisms previously discussed dominate, and disturbances may merely act as a trigger for them.

In these low Reynolds conditions, unsteady dynamics are also observed [20]. The reattachment point can begin oscillating back and forth at low frequency, with shedding of vortical structures in the boundary layer. This phenomenon is deemed bubble flapping. Vortex shedding at high incidence, which causes strong dynamic loading, has been linked to K-H instabilities in the shear layer; related entrainment mechanisms can also increase the likelihood of mean reattachment due to the increased mixing. Static hysteresis phenomena have also been observed [5] as the flow alternates between stall and reattachment.

2.1.3 Airfoil performance

In the range $10^4 < Re < 5 \cdot 10^5$, the airfoil flow and performance are very sensitive to variations in the Reynolds number; as Re decreases we assist to a sharp decline in airfoil efficiency (Fig. 2.8, Fig. 2.9a). On the suction side, the laminar boundary layer is basically guaranteed to separate before turbulent transition, leading to either the formation of a LSB or stall, as previously discussed.

At the higher end of the range ($10^5 < Re < 5 \cdot 10^5$), the flow is qualitatively unaltered by the short LSB, as the shear layer readily reattaches as a turbulent boundary layer [4]. The coefficient of lift C_L exhibits a non-linear dependency in α, as the separation point moves upstream with increasing incidence. The decrease in performance relative to high Reynolds regimes is to be attributed to the higher viscosity and earlier separation.

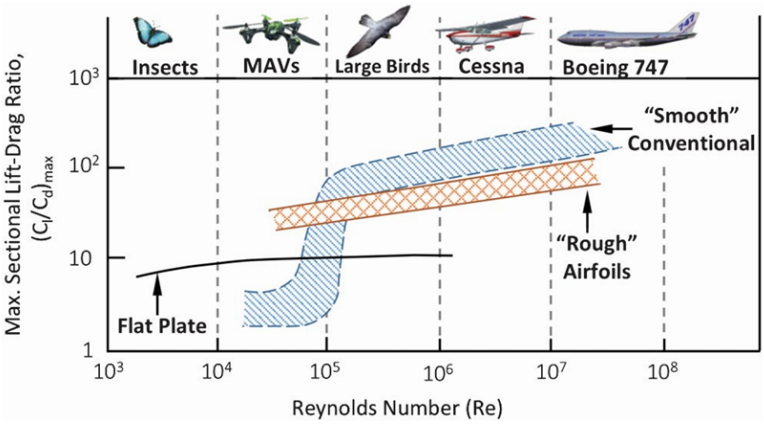

Figure 2.8: Airfoil efficiency with respect to Reynolds number (Credit: [4]).

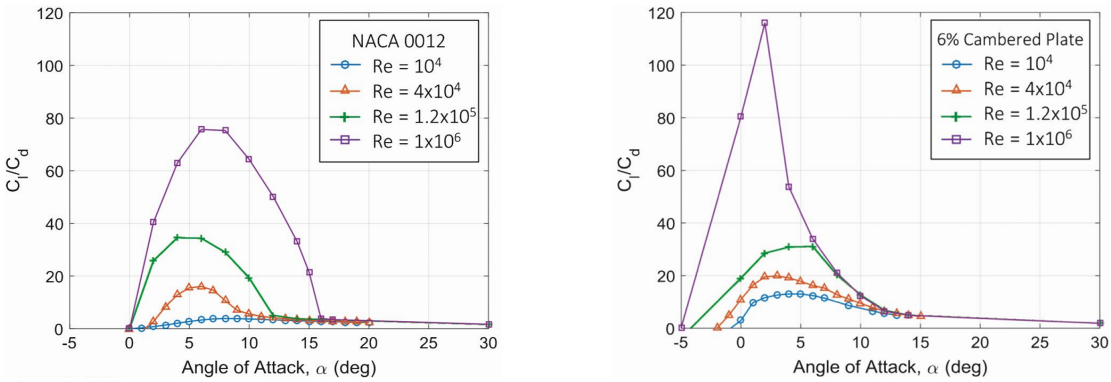

(a) NACA0012. **(b)** Cambered plate.

Figure 2.9: Aerodynamic efficiency against α, Re (Credit: [4]).

In the range $5 \cdot 10^4 < Re < 10^5$ efficiency further decreases, the flow dynamics become extremely sensitive to the Reynolds number [4]. The rich phenomenology of the LSB governs the flow, causing highly non-linear behaviour in $C_L(\alpha)$ and often unsteady loading, due to bubble flapping and vortex shedding. Stall is also anticipated, again hindering aerodynamic efficiency, as shown in Fig. 2.9a for the NACA0012.

Finally, with $10^4 < Re < 5 \cdot 10^4$ the airfoil is in trailing edge stall for most of its operation; efficiency again worsens.

A note is to be made on very thin airfoils, which exhibit better performance than conventional airfoils in low Reynolds conditions (Fig. 2.8) [4]. The geometry forces

a strong separation at the sharp leading edge on the suction side; early separation promotes transition and reattachment, which improves performance and stabilizes the flow (Fig. 2.9b).

2.2 Numerical modeling

Characterizing airfoil performance at low Reynolds numbers ($10^4 < Re < 10^5$) numerically represents a difficult challenge; as established previously, the flow dynamics in this regime are heavily dictated by complex viscous phenomena. In this work, two different numerical methodologies were used to obtain the numerical polar maps: the first is XFOIL, which is based on a Viscous-Inviscid Interaction (VII) formulation; the second approach is CFD, in particular the Unsteady Reynolds-Averaged Navier Stokes equations (URANS) with the $k-\omega$ SST turbulence model, coupled with the $\gamma - Re_\theta$ transition model.

2.2.1 Viscous-Inviscid Interaction in XFOIL

The phenomenology of the laminar boundary layer in low-Reynolds conditions fundamentally alters the shape of the airfoil as perceived by the outer, inviscid flow, which in turn drives the boundary layer by imposing the pressure gradient and outer velocity. Methods based on Viscous-Inviscid Interaction (VII) aim at modeling airfoil flow by coupling the inviscid flow and the boundary layer in a single formulation, while using specific tools to resolve these two regions separately. XFOIL is one such method; developed by Drela in 1989 [22, 23], it was conceived as a rapid analysis and design tool for low Reynolds airfoils specifically. Over the decades, it has proven to be a computationally efficient and reliable tool, remaining a standard for many aerodynamic applications today.

In the XFOIL framework, the outer inviscid flow is governed by potential flow theory. The flow field is mathematically constructed through the linear superposition of a uniform freestream U_∞ with incidence α, a vortex sheet γ distributed along the airfoil surface, and a source sheet σ applied over both the

surface and the wake. The resulting streamfunction reads

$$\Psi(x,y) = U_\infty \cos(\alpha) y - U_\infty \sin(\alpha) x + \frac{1}{2\pi}\int \gamma(s) \ln r(x,y,s) ds + \frac{1}{2\pi}\int \sigma(s)\theta(x,y,s) ds \ , \tag{2.14}$$

where s is the coordinate along the airfoil surface, and r, θ are the radius and angle between the point s and the flow field point (x, y). To solve this potential flow problem, a panel method is employed to discretize the airfoil (N_a panels) and the wake (N_w panels), leading to the formulation of a system of linear equations with the panel's vortex intensities $\boldsymbol{\gamma} \in \mathbb{R}^{N_a}$ as unknowns

$$\mathbf{A}(x,y,s)\boldsymbol{\gamma} = f(\boldsymbol{\sigma}, U_\infty, \alpha) \ . \tag{2.15}$$

where $\mathbf{A}$ is the induction matrix, and depends solely on the geometry of the problem. The right hand side has been compressed for simplicity; the system depends on the panels' sources intensities $\boldsymbol{\sigma} \in \mathbb{R}^{N_a+N_w}$, which are unknown at this point and have to be related to the boundary layer displacement thickness.

The viscous analysis uses an integral formulation of the boundary layer's equations, namely the compressible integral momentum and kinetic energy shape parameter equations

$$\frac{d\theta}{d\xi} + (2 + H - M_e^2)\frac{\theta}{u_e}\frac{du_e}{d\xi} = \frac{C_f}{2} \ , \tag{2.16}$$

$$\theta\frac{dH^*}{d\xi} + (2H^{**} + H^*(1-H))\frac{\theta}{u_e}\frac{du_e}{d\xi} = 2C_d - H^*\frac{C_f}{2} \ . \tag{2.17}$$

where H, H^*, H^{**} are shape factors of the boundary layer, here taking the role of closure quantities, and C_d is the mechanical energy dissipation coefficient. Following the integral method, in XFOIL the system is closed by prescribing a parametric boundary layer profile family. In the laminar regime, the code uses the Falkner-Skan family, while for turbulent boundary layers it uses Swafford's. To detect transition between the two regimes, the e^N method described in the previous section is employed. These equations again are discretized on the airfoil's paneling. Still, the boundary layer's equation system needs data on the outer flow driving the dynamics to produce a solution.

Finally, the two formulations are linked through the outer velocity u_e which drives the boundary layer, obtained by the inviscid method, and by the sources' intensities $\boldsymbol{\sigma}$ which alters the outer flow, and is related to the displacement thickness. With Drela's formulation, the two methods can be combined together in a single system of non-linear equations, greatly improving computation time with respect to other VII formulations at the time. Solving the complete system thus gives a representation of both boundary layer and outer flow field at the same time, fully characterizing the airfoil in a specific operative condition.

Over the decades since its introduction, XFOIL has proven to be a reliable and computationally efficient tool for airfoil analysis in a wide range of conditions, becoming a standard reference in low-Reynolds airfoil design and assessment. Its speed makes it particularly well suited to applications requiring large numbers of evaluations, such as the generation of polar maps across broad ranges of Reynolds number and angle of attack, which is precisely the use case required by BEMT-based propeller modeling. However, XFOIL also presents some limitations originating from its formulation. As noted in its official documentation [24], the solver often struggles to converge or loses accuracy in cases of massive boundary layer separation or highly unsteady phenomena, like the ones discussed earlier. Additionally, literature indicates that XFOIL's drag predictions become less reliable when large laminar separation bubbles are present [25]. In these situations, the basic assumptions of the viscous-inviscid model break down, requiring the use of higher-fidelity methods.

2.2.2 Computational Fluid Dynamics

An alternative to VII calculations is to simulate the flow over an airfoil through Computational Fluid Dynamics (CFD). By directly solving the discretized Navier-Stokes equations, CFD is better suited for capturing the complex and unsteady aerodynamic effects that dictate performance near and beyond the stall angle, albeit at a significantly higher computational cost than XFOIL. Given what was discussed earlier regarding flow unsteadiness and laminar phenomenology in low-Reynolds airfoil flows, an approach based on the incompressible, two-dimensional Unsteady Reynolds-Averaged Navier-Stokes (URANS) equations with the $k-\omega$

SST turbulence model and the $\gamma - Re_\theta$ transition model was chosen.

URANS equations

In the vast majority of applications, directly solving the Navier-Stokes equations would be too demanding computationally; thus, one needs to model part of the flow physics. In the URANS equations, the turbulent fluctuations in time are separated from the mean flow and then modeled via a turbulence model. Indicating with $\bar{\cdot}$ the ensemble average of a quantity, the Reynolds decomposition splits the velocity field into a mean $\overline{\mathbf{u}}(\mathbf{x}, t)$ and a fluctuating $\mathbf{u}'(\mathbf{x}, t)$ component:

$$\mathbf{u}(\mathbf{x}, t) = \overline{\mathbf{u}}(\mathbf{x}, t) + \mathbf{u}'(\mathbf{x}, t) \tag{2.18}$$

Introducing this decomposition into the incompressible Navier-Stokes equations leads to the equivalent URANS equations:

$$\nabla \cdot \bar{\mathbf{u}} = 0 \,, \tag{2.19}$$

$$\frac{\partial \bar{\mathbf{u}}}{\partial t} + (\bar{\mathbf{u}} \cdot \nabla)\bar{\mathbf{u}} = -\frac{1}{\rho}\nabla \bar{p} + \nu \nabla^2 \bar{\mathbf{u}} + \frac{1}{\rho}\nabla \cdot \boldsymbol{\tau}_R \,, \tag{2.20}$$

where $\boldsymbol{\tau}_R$ is the Reynolds stress tensor, defined as[3]

$$\boldsymbol{\tau}_R = -\rho\overline{(\mathbf{u}' \otimes \mathbf{u}')} \,. \tag{2.21}$$

The Reynolds stresses emerge from the averaging of the non-linear convective term in the momentum equation, and represent the effect of turbulence on the mean flow. However, because the turbulent fluctuations are not resolved in order to reduce computational cost, the term $\boldsymbol{\tau}_R$ must be expressed as a function of the mean field through a turbulence model to close the system.

The $k - \omega$ SST turbulence model

Menter's $k-\omega$ Shear Stress Transport (SST) turbulence model (1994) [26] belongs to the family of turbulent-viscosity models. These assume that turbulence transports

[3]The symbol "$\otimes$" indicates the outer product of two vectors.

momentum in the mean flow in a diffusive fashion, behaving analogously to the viscosity in a Newtonian fluid, as

$$\boldsymbol{\tau}_R = -2\rho\nu_T\overline{\mathbf{S}} - \frac{2}{3}\rho k\mathbf{I} \,, \tag{2.22}$$

where ν_T is the turbulent kinematic viscosity and $\overline{\mathbf{S}}$ is the strain rate tensor of the mean velocity, defined as

$$\overline{\mathbf{S}} = \frac{1}{2}(\nabla\overline{\mathbf{u}} + \nabla\overline{\mathbf{u}}^{\mathrm{T}}) \,. \tag{2.23}$$

Each model belonging to this family uses a different method to calculate the turbulent viscosity. The $k-\omega$ SST in particular is a two-equation model: it calculates the turbulent viscosity using the turbulent kinetic energy k and the specific turbulent dissipation rate ω, which are obtained by solving the following transport equations coupled to the core URANS equations (2.19), (2.20)

$$\frac{\partial k}{\partial t} + \overline{\mathbf{u}}\cdot\nabla k = P_k - \beta^* k\omega + \nabla\cdot[(\nu + \sigma_k\nu_T)\nabla k] \,, \tag{2.24}$$

$$\frac{\partial \omega}{\partial t} + \overline{\mathbf{u}}\cdot\nabla\omega = \frac{\alpha}{\nu_T}P_k - \beta\omega^2 + \nabla\cdot[(\nu + \sigma_\omega\nu_T)\nabla\omega] + 2(1-F_1)\sigma_{\omega 2}\frac{1}{\omega}\nabla k\cdot\nabla\omega \,, \tag{2.25}$$

while the expression used for the turbulent viscosity is:

$$\nu_T = \frac{a_1 k}{\max(a_1\omega, SF_2)} \,. \tag{2.26}$$

The complete formulation of the blending functions F_1 and F_2, as well as the closure constants, are omitted here for brevity but are detailed in Menter's original formulation [26].

The SST model employs a zonal blending approach, combining two classical turbulence models to leverage the best aspects of both. The model utilizes blending functions to seamlessly transition between formulations based on the distance from the wall:

1. In the near-wall region, it operates as Wilcox's standard $k-\omega$ model. This

formulation is highly robust and accurately predicts flow behavior under adverse pressure gradients without requiring the complex near-wall damping functions used by other models.

2. In the outer flow region, the model switches to a $k-\varepsilon$ formulation, recast in terms of ω. This eliminates the standard $k-\omega$ model's sensitivity to freestream turbulence properties, ensuring reliable far-field boundary conditions.

The SST formulation is particularly well suited for low-Reynolds applications, where solution accuracy at the wall is crucial. Indeed, the model directly solves the equations down to the viscous sublayer without the need for empirical wall functions, capturing flow separation much more accurately. Furthermore, the $k-\omega$ SST model serves as the foundational turbulent framework for the $\gamma - Re_\theta$ transition model, which is currently the industry standard.

The $\gamma - Re_\theta$ transition model

A primary objective in modern Computational Fluid Dynamics (CFD) is to capture transition phenomenology with general-purpose solvers in order to improve prediction accuracy. However, traditional transition models relied on non-local, geometry-dependent formulations, which required knowledge of a greater portion of the fluid domain. In unstructured, parallelized CFD codes, this long-range cell dependence would be computationally costly and difficult to implement.

To overcome these limitations, Langtry and Menter [27, 28] developed the $\gamma - Re_\theta$ transition model based on the Local Correlation-Based Transition Modeling (LCTM) framework. The fundamental strength of this model is its strictly local and universal formulation; by evaluating transition criteria using only local cell variables and resolving the flow through regular transport equations, it eliminates the need for geometry-specific operations. This PDE-based approach makes the model perfectly compatible with modern CFD codes.

The model introduces two additional transport equations that are solved together with the URANS and $k-\omega$ SST system. The first equation, governing the spatial and temporal evolution of the intermittency γ, reads:

$$\frac{\partial \gamma}{\partial t} + \overline{\mathbf{u}} \cdot \nabla \gamma = P_\gamma - E_\gamma + \nabla \cdot \left[\left(\nu + \frac{\nu_T}{\sigma_\gamma} \right) \nabla \gamma \right] , \tag{2.27}$$

where P_γ and E_γ are the production and destruction terms driving the transition, and σ_γ is a closure constant. The intermittency modulates the turbulence production in the underlying $k-\omega$ SST model; in purely laminar regions, $\gamma = 0$, which effectively deactivates turbulent mixing.

The second equation governs the transport of the transition onset momentum thickness Reynolds number, $\tilde{Re}_{\theta t}$:

$$\frac{\partial \tilde{Re}_{\theta t}}{\partial t} + \overline{\mathbf{u}} \cdot \nabla \tilde{Re}_{\theta t} = P_{\theta t} + \nabla \cdot \left[\sigma_{\theta t}(\nu + \nu_T) \nabla \tilde{Re}_{\theta t} \right] , \tag{2.28}$$

where $P_{\theta t}$ is the source term and $\sigma_{\theta t}$ controls the diffusion rate.

The fundamental insight that enables this local formulation relies on the vorticity Reynolds number, defined as

$$Re_v = \frac{\rho y^2 \Omega}{\mu} , \tag{2.29}$$

where y is the absolute distance to the nearest wall and Ω is the local vorticity magnitude. Langtry and Menter demonstrated a direct correspondence between the maximum vorticity Reynolds number and the momentum thickness Reynolds number for a Blasius velocity profile, given by

$$Re_\theta = \frac{\max(Re_v)}{2.193} . \tag{2.30}$$

Instead of integrating the velocity profile across the boundary layer, the momentum thickness Reynolds number can thus be calculated locally, and transition can be assessed utilizing empirical correlations.

The logic behind this transition model is as follows: the intermittency γ modulates turbulence generation in the URANS equations and is itself governed by a transport equation. The source term P_γ is activated when the local Re_θ, estimated via Equation (2.30), exceeds a critical threshold value $Re_{\theta c}$. This critical value is mathematically linked to the transition variable $\tilde{Re}_{\theta t}$. This transported variable is forced to adhere to empirical correlations in the freestream and subsequently diffuses its value toward the wall; the boundary layer thus receives the necessary information regarding the outer flow conditions driving its dynamics without requiring non-local integration.

The $k-\omega$ SST turbulence model with $\gamma - Re_\theta$ transition equations represents a higher-fidelity framework for the low-Reynolds aerodynamic regime analyzed in this work. In airfoil flow, the transition model ensures that the boundary layer can stay laminar, accurately separate under adverse pressure gradients, and naturally transition within the separated shear layer. This resolves the LSBs and improves prediction of stall and thus of airfoil performance with respect to regular, fully turbulent URANS simulations.

2.3 Blade Element Momentum Theory

As anticipated, the Blade Element Momentum Theory (BEMT) [29] is a low order model used to efficiently compute the aerodynamic coefficients for propellers and wind turbines in a defined operative condition (Fig. 2.10a).

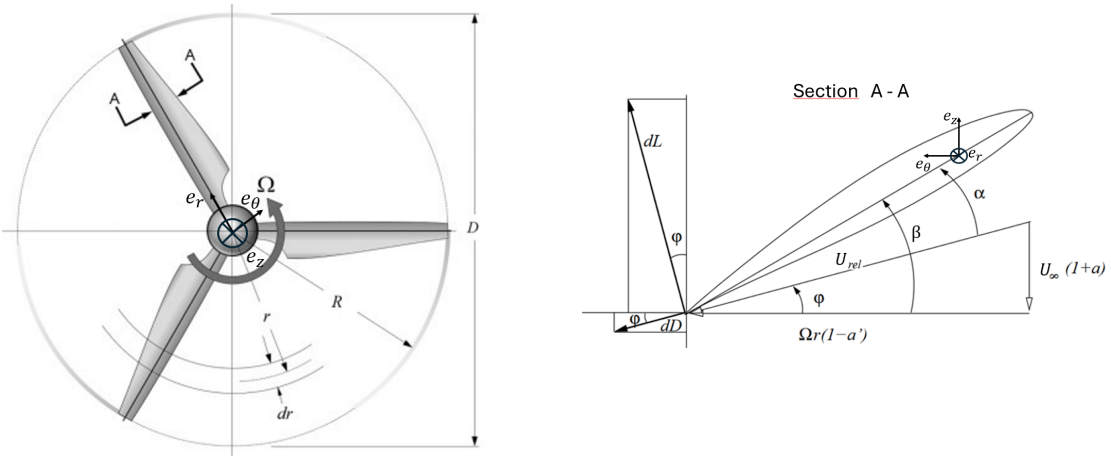

(a) Propeller. **(b)** Blade element (Adapted from: [30].

Figure 2.10: Propeller schematics.

The setting for the basic formulation of this theory is a steady propeller under uniform, axially symmetrical inflow. The far-field velocity U_∞ and the propeller's rotational velocity Ω fully define its operation, together with the fluid's properties; these are distilled into two non-dimensional parameters, namely the advance ratio

J and the propeller's Reynolds number Re_p, defined as

$$J = \frac{U_\infty}{\Omega R} \,, \tag{2.31}$$

$$Re = \frac{U_{0.75} c_{0.75}}{\nu} \,, \tag{2.32}$$

where R is the propeller's radius, while Re is taken as the chord-based Reynolds numbe of the blades' section at 75% of total radius length. The propeller's performance is given by the aerodynamic coefficients of thrust T and torque Q; the following convention has been adopted

$$C_T(J, Re_p) = \frac{T}{\frac{1}{2}\rho(\Omega R)^2 \pi R^2} \,,$$

$$C_Q(J, Re_p) = \frac{Q}{\frac{1}{2}\rho(\Omega R)^2 \pi R^3} \,,$$

$$C_P(J, Re_p) = \frac{Q\Omega}{\frac{1}{2}\rho(\Omega R)^2 \pi R^3 \Omega} = C_Q(J, Re_p) \,.$$

The propulsive efficiency is readily obtained as

$$\eta(J, Re_p) = \frac{T U_\infty}{Q\Omega} = J \frac{C_T}{C_Q} \,, \tag{2.33}$$

which quantifies how much of the power is spent in forward motion.

The Blade Element Momentum Theory combines two models for propellers, namely the Momentum Theory (MT) and the Blade Element Theory (BET), which by themselves cannot fully characterize the propeller. Indeed, the first provides a simplified model of the flow field given the loading on the propeller, while the second calculates the loading given the flow field at the rotor.

The first step that needs to be taken to join the two theories is to link the definition of the relative velocity at the disk needed by BET to the flow field determined by MT. In the reference system represented in Fig. 2.10, the induced velocity at the disk $\mathbf{u}_i(r, \theta, 0) = \mathbf{u}(r, \theta, 0) - U_\infty \mathbf{e}_z$ is expressed through the induction

coefficients[4]

$$a(r,\theta) = \frac{u_{i,z}(r,\theta,0) - U_\infty}{U_\infty} \ , \tag{2.34}$$

$$a'(r,\theta) = -\frac{u_{i,\theta}(r,\theta,0)}{\Omega r} \ . \tag{2.35}$$

Given this definition, in the rotor's frame of reference, the velocity $\mathbf{u}_{rel}$ at the rotor (Fig. 2.10b) has components

$$u_{rel,z}(r,\theta,0) = U_\infty[1 + a(r,\theta)] \ , \tag{2.36}$$

$$u_{rel,\theta}(r,\theta,0) = \Omega r[1 - a'(r,\theta)] \ , \tag{2.37}$$

the relative inflow angle is thus obtained as

$$\phi = \arctan\left[J\frac{1 + a(r,\theta)}{1 - a'(r,\theta)}\right] \tag{2.38}$$

The following and final step is to derive the thrust and torque expressions from the two theories and equate them to close the system and obtain a solution.

2.3.1 Momentum Theory

The Momentum Theory, and in particular the Streamtube Theory (STT), models the propeller as an actuator disk and analyzes it by radial discretization and subsequent application of the Euler equations to control volumes defined by the infinitesimal annuli that constitute the disk (Fig. 2.11) [31]. The theory provides a simplified flow field given the loading on the actuator disk, as anticipated earlier. This model is based on the following (main) hypotheses:

1. Stationary, axially symmetrical inflow with asymptotic velocity U_∞;
2. Radial independence of the annuli;
3. Axially symmetrical actuator loads concentrated on the disk;

[4] As the tangential velocity is discontinuous across the disk, at the disk it is taken as the average of the values before and after the disk: $u_{i,\theta}(r,\theta,0) = 1/2 \cdot [u_{i,\theta}(r,\theta,0^-) + u_{i,\theta}(r,\theta,0^+)]$.

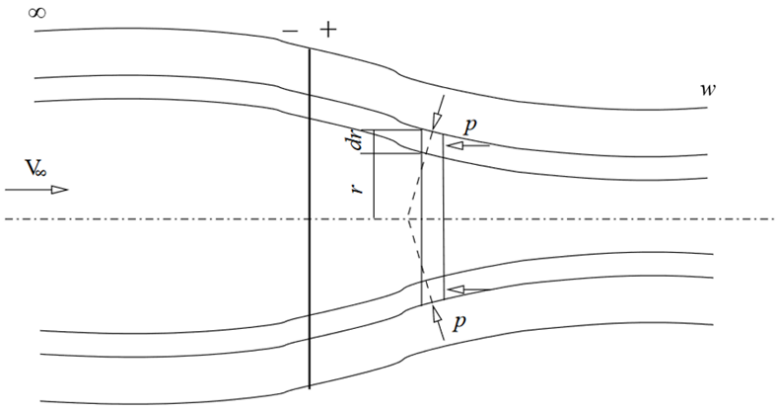

Figure 2.11: Schematic of an annulus' control volume (Credit: [30]).

4. Discontinuity of pressure and tangential velocity across the disk.

The result of the STT is to link the loading to the flow field, providing a mathematical relation between an annulus' local aerodynamic coefficients dC_T, dC_Q and the induction factors a, a' at that point in the flow field, as[5]

$$dC_T(r) = \frac{dT}{\frac{1}{2}\rho(\Omega R)^2\pi R^2} = 8J^2F(r,\phi)[1+a(r)]a(r)\frac{r}{R}d\left(\frac{r}{R}\right) \ , \tag{2.39}$$

$$dC_Q(r) = \frac{dQ}{\frac{1}{2}\rho(\Omega R)^2\pi R^3} = 8JF(r,\phi)[1+a(r)]a'(r)\left(\frac{r}{R}\right)^3 d\left(\frac{r}{R}\right) \ . \tag{2.40}$$

The function $F(r,\phi)$ derives from the tip-loss corrective model, explained later on.

2.3.2 Blade Element Theory

Given the flow field at the rotor, the Blade Element Theory provides the loading, determined using the blade elements' two-dimensional aerodynamic coefficients and the airfoil relative to the blades (Fig. 2.10b). This model relies on the following assumptions:

1. Radial independence of the blade elements;

2. Two-dimensional flow on the blade elements;

[5]Refer to Branlard [31] for proof.

The resulting equations[6] again link the local coefficients dC_T, dC_Q to the relative flow field

$$dC_T(r,\phi) = C_z(r,\phi) J^2 \frac{Bc(r)}{\pi R} \left(\frac{u_{rel}(r,\phi)}{U_\infty} \right)^2 d\left(\frac{r}{R}\right) , \tag{2.41}$$

$$dC_Q(r,\phi) = C_\theta(r,\phi) J^2 \frac{Bc(r)r}{\pi R^2} \left(\frac{u_{rel}(r,\phi)}{U_\infty} \right)^2 d\left(\frac{r}{R}\right) , \tag{2.42}$$

where B is the number of blades, while C_z and C_θ are the projections of the two-dimensional aerodynamic coefficients C_L, C_D along the axial and tangential directions (Fig. 2.10b), obtained as

$$C_z = C_L \cos\phi - C_D \sin\phi , \tag{2.43}$$

$$C_\theta = C_L \sin\phi + C_D \cos\phi . \tag{2.44}$$

The relative flow is described by the modulus u_{rel} and the incidence relative to the rotor plane ϕ. The expressions (2.41), (2.42) again depend solely on r, and can be integrated along the radius to obtain the global coefficients.

2.3.3 Blade Element Momentum Theory

By equating the MT's and the BET's expressions for the local coefficients, it is possible to obtain two equations for the induction factors

$$a(r,\phi) = \left(\frac{4F(r,\phi)\sin^2\phi}{\sigma(r) C_z(r,\phi)} - 1 \right)^{-1} , \tag{2.45}$$

$$a'(r,\phi) = \left(\frac{4F(r,\phi)\sin\phi\cos\phi}{\sigma(r) C_\theta(r,\phi)} + 1 \right)^{-1} , \tag{2.46}$$

[6] Refer to Branlard [31] for proof.

where $\sigma = Bc/(2\pi r)$ is the local solidity. Following Ning's procedure [32], a single equation for the inflow angle ϕ can be obtained

$$f(r,\phi) = \frac{\sin\phi}{1+a(r,\phi)} - J\frac{R}{r}\frac{\cos\phi}{1-a'(r,\phi)} = 0\ . \tag{2.47}$$

Ning's equation is the result of the BEMT, as obtaining the inflow angle on all the infinitesimal annuli making up the propeller is enough to fully characterize the propeller.

2.3.4 Prandtl's tip-loss corrective model

The various assumptions at the base of the BEMT have greatly simplified the problem, leading to a simple one-equation solution. However, these restrict the field of application of the BEMT. To overcome these limitations, corrective models have been formulated, through empirical and/or theoretical analyses. The present work only employed the Prandtl's tip-loss corrective model, although multiple can be applied at the same time.

The hypoteses (2-3) from MT and (1-2) from BET are valid only for a rotor with an infinite number of blades, producing a uniform vortex sheet in the wake. A real propeller with a finite number of blades sheds discrete helical tip vortices, which create a non-uniform induced velocity field and force the blade loading to vanish at the tip. To correct the BEMT, Glauert [29] modeled the semi-helical vortical wake as a succession of vortex sheets, which generate an induced velocity at the propeller, in the same fashion of Prandtl's lifting line model. The BEMT is corrected through a tip-loss factor included in the STT equations (2.39), (2.40), defined as

$$F(r,\phi) = \frac{2}{\pi}\left(e^{-f}\right)\ , \tag{2.48}$$

$$f = \frac{B}{2}\frac{R-r}{r\sin\phi}\ . \tag{2.49}$$

For an infinite number of blades $B \to \infty$, $F \to 1$ and the BEMT regresses to its original, uncorrected formulation.

Chapter 3

Numerical Methodology

3.1 The Modular Numerical Framework

The evaluation of propeller performance in this study relies on a modular numerical framework especially developed for it; a qualitative schematic is provided in Fig. 3.1. The primary objective of the pipeline is flexibility, as it must be applicable to different propeller geometries and operative conditions. Separating the workflow into distinct blocks allows one to easily switch between different configurations without requiring a fundamental restructuring of the code.

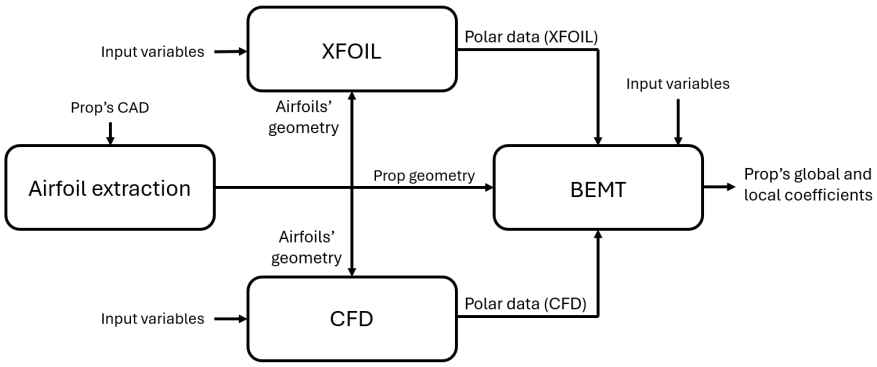

Figure 3.1: Pipeline scheme.

To guarantee consistency across the entire pipeline, all global variables are centralized within a single master configuration script, named `INPUT.py`. The

operational conditions and the spatial discretization depend heavily on the specific rotor under analysis; they are thus defined *a priori* in this central file. These parameters include the freestream velocity V_∞, the rotational speed Ω, the fluid properties ρ and ν, and the radial node distribution along the blade span.

Furthermore, the modular nature of the pipeline requires a clear data interface between consecutive blocks. The individual numerical modules do not share active memory variables during execution; they communicate exclusively through standardized output files. Specifically, the extracted geometric parameters and the airfoil coordinates are transferred via `geometry.csv` and `.dat` files, while the aerodynamic data is collected in dedicated `.csv` polar maps. This structure guarantees the complete independence of the calculation tools. The aerodynamic characterization of the blade sections can thus be performed utilizing either a Viscous–Inviscid Interaction solver (XFOIL) for rapid evaluations, or a high-fidelity CFD solver (OpenFOAM) for a more detailed transition and stall analysis. The choice of solver solely affects the intermediate polar generation; the resulting data seamlessly feeds the final Blade Element Momentum Theory (BEMT) integration step without altering the main integration logic.

The following sections describe each component of the pipeline in detail. The calculation of the airfoils' two-dimensional aerodynamic coefficients precedes logically the BEMT solver, but for the sake of clarity it is presented immediately after the airfoil extraction section and before the polar data calculation.

3.2 Geometry Extraction

In practical implementations of the BEMT, the continuous propeller blade must be divided into a finite number of radial elements; the domain is thus discretized using a succession of boundary nodes $\{r'_i, r'_{n+1}\}$, with $i = 1, \ldots, n$ (Fig. 3.2). These nodes define a series of n annuli, where the central radial coordinate r_i and the radial width Δr_i of each element are evaluated as

$$r_i = \frac{r'_i + r'_{i+1}}{2} \,, \tag{3.1}$$

$$\Delta r_i = r'_{i+1} - r'_i \,. \tag{3.2}$$

To properly resolve the aerodynamic forces, the local pitch angle β_i and the chord length c_i must be evaluated at each centre r_i, along with the global hub and tip radii R_{hub} and R.

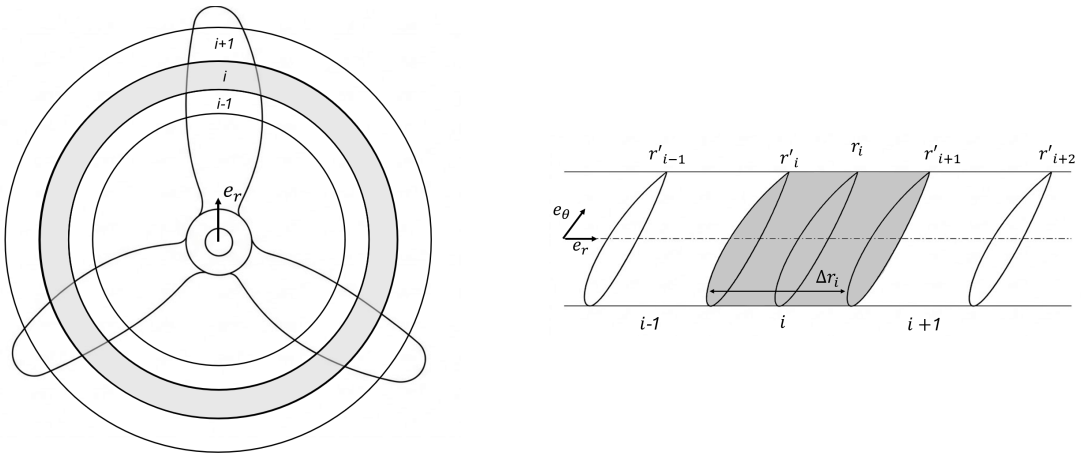

Figure 3.2: Radial discretization representation.

Furthermore, the BEMT solver requires a specific two-dimensional airfoil shape for each annulus to compute the necessary lift and drag coefficients. However, extracting and simulating a unique airfoil for every single radial element would introduce unnecessary computation; the geometric extraction is thus decoupled from the aerodynamic shape sampling. The procedure evaluates the macroscopic variables (c_i and β_i) at every BEMT node, but extracts the actual airfoil coordinates only at a few discrete spanwise locations. The script then automatically maps each radial element to the closest available extracted airfoil to properly assign the necessary polar data.

This entire procedure is automated by the `EXTRACTOR.py` script, which takes the three-dimensional CAD geometry in STL format and the radial node specifications from `INPUT.py` as its starting point (Fig. 3.3). Relying on the VKI library `batmanpy`, the script executes a precise sequence of geometric operations:

- It evaluates the global rotor limits R and R_{hub};
- It slices the blade at the requested radial centres r_i to extract three-dimensional geometric strips;

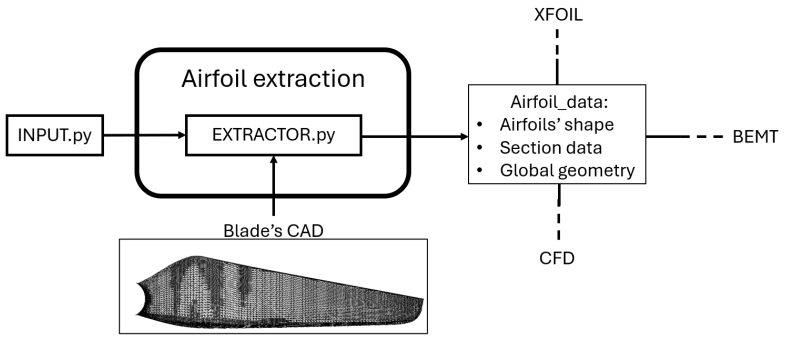

Figure 3.3: Geometry extraction block scheme.

- It unwraps these strips from their native cylindrical coordinates into a flat two-dimensional plane;

- It calculates the chord c_i and pitch β_i for each section;

- For the extracted airfoils, it translates the leading edge of the resulting sections to the origin, rotates them by the twist angle β to achieve a zero angle of attack, and finally normalizes the coordinates by dividing by the chord value.

All generated data is ultimately collected in the `airfoil_data` directory. This folder structurally separates the global variables (`geometry.csv`) from the arrays of local parameters mapped to their respective airfoils (`data.csv`). The raw normalized coordinate files for the discrete airfoils and their corresponding camberlines are stored in dedicated subdirectories.

```
airfoil_data/
├── camber/
│   ├── airfoil_r1_camber.dat
│   ├── airfoil_r2_camber.dat
│   └── ...
├── dat/
│   ├── airfoil_r1.dat
│   ├── airfoil_r2.dat
│   └── ...
├── geometry.csv
└── data.csv
```

Following this structural division, the coordinate files are passed to the aerodynamic solvers (XFOIL or CFD) to evaluate the fluid dynamics, while the macroscopic geometric parameters are passed directly to the BEMT solver.

3.3 Blade Element Momentum Theory (BEMT) Implementation

Following the radial discretization previously defined, the BEMT is utilized to evaluate the local thrust and torque coefficients for each annulus, which are subsequently summed to obtain the global propeller performance. This mathematical closure requires the airfoil aerodynamic coefficients to be known.

For the i-th annulus, Ning's equation (2.49) is solved to obtain the relative inflow angle ϕ_i. The local coefficients are then calculated by modifying equations (2.41) and (2.42) as follows:

$$\left(\frac{dC_T}{d(r/R)}\right)_i = C_z(r_i,\phi_i)\, J^2\, \frac{Bc_i}{\pi R}\left(\frac{U_{\text{rel}}(r_i,\phi_i)}{U_\infty}\right)^2 ,$$
$$\left(\frac{dC_Q}{d(r/R)}\right)_i = C_\theta(r_i,\phi_i)\, J^2\, \frac{Bc_i r_i}{\pi R^2}\left(\frac{U_{\text{rel}}(r_i,\phi_i)}{U_\infty}\right)^2 .$$

The propeller's global coefficients are then obtained by approximate integration along the blade span:

$$C_T = \int_{R_{\text{hub}}/R}^{1} \frac{dC_T}{d(r/R)}\, d\left(\frac{r}{R}\right) \approx \sum_{i=1}^{n}\left(\frac{dC_T}{d(r/R)}\right)_i \frac{\Delta r_i}{R} , \tag{3.3}$$
$$C_Q = \int_{R_{\text{hub}}/R}^{1} \frac{dC_Q}{d(r/R)}\, d\left(\frac{r}{R}\right) \approx \sum_{i=1}^{n}\left(\frac{dC_Q}{d(r/R)}\right)_i \frac{\Delta r_i}{R} . \tag{3.4}$$

This procedure is at the core of the present BEMT implementation, handled by `BEMT_run.py` and the associated functions in `BEMT_funcs.py` (Fig. 3.4).

The Python solver performs the following specific actions:

- It loads the geometry data and creates a two-dimensional grid interpolator for the airfoil data; to prevent evaluation failures outside the convex hull of the input data during numerical iterations, a hybrid approach is employed, falling

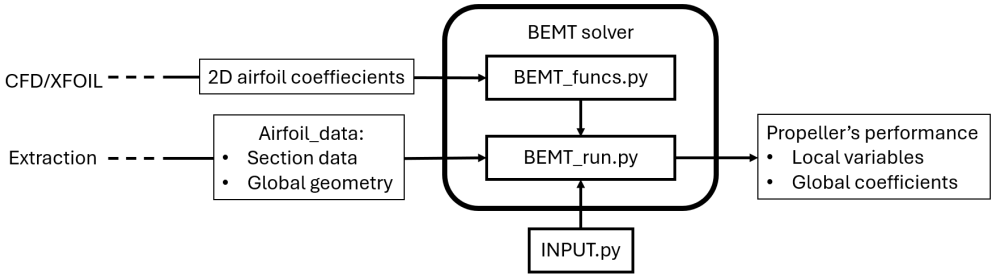

Figure 3.4: BEMT solver block scheme.

back from linear to nearest-neighbour interpolation;

- For each radial annulus, it evaluates the inflow angle ϕ_i using Brent's method to find the root of the momentum–blade-element residual (Ning's algorithm);
- The local Reynolds number depends on the evaluated induction factors; an iterative loop is therefore executed at each station to update Re_i until convergence is reached;
- It computes the local induction factors a_i and a'_i, the angle of attack α_i, and the load coefficients $C_{z,i}$ and $C_{\theta,i}$;
- It evaluates the incremental thrust and torque coefficients $[dC_T/d(r/R)]_i$ and $[dC_Q/d(r/R)]_i$ at each radial station;
- It integrates the local distributions along the blade span to obtain the global coefficients C_T, C_Q, C_P, and the propulsive efficiency η;
- It automatically saves the local and global results to dedicated CSV files in the `BEMT_results` directory.

The primary advantage of the BEMT solver is its computational efficiency. By reducing the propeller characterization to n independent nonlinear equations, the solution is virtually instantaneous; this makes the framework ideal for optimization loops and real-time control applications, where high-fidelity CFD would be prohibitively expensive. However, the modeling effort is heavily shifted offline: the airfoil polar data must be precomputed and tabulated across the relevant operative ranges. Once this database is available, the solver remains entirely lightweight and the aerodynamic performance can be evaluated at negligible cost.

3.4 XFOIL

When dealing with two-dimensional airfoil performance analysis, XFOIL is one of the most effective tools available. It can rapidly provide sufficiently accurate lift and drag data in low-Reynolds regimes, requiring only the airfoil geometry and operative conditions as input. It is important to note that XFOIL carries its own limitations, related to convergence issues and reduced accuracy in challenging conditions, as discussed in the previous chapter. XFOIL is distributed as a standalone executable and cannot be imported directly by Python scripts; however, it can be driven externally, which is consistent with the modular pipeline concept.

The polar data generation is automated by the script `xfoil_run.py` and the functions in `xfoil_funcs.py`, following the workflow in Fig. 3.5.

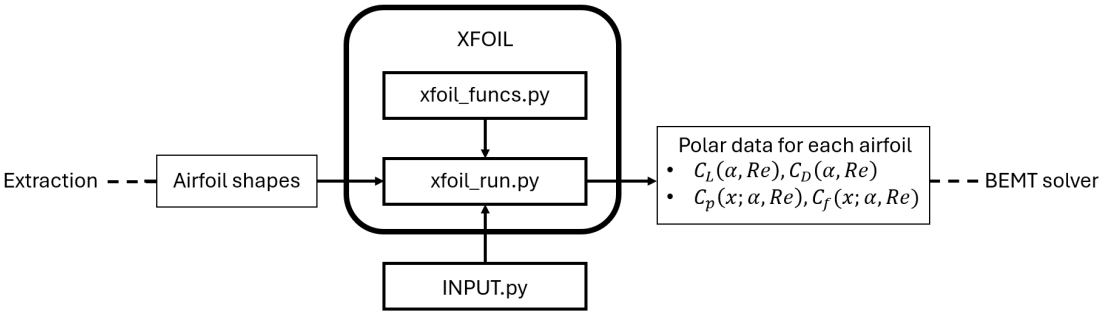

Figure 3.5: XFOIL automation block scheme.

The script requires an input grid of Reynolds numbers and incidence angles to compute the airfoil polar data. These ranges are dictated by the propeller geometry and operative conditions, and are therefore obtained directly from the BEMT pipeline. Furthermore, the transition parameter N_{crit} is linked to the freestream turbulence intensity Tu through Mack's correlation, in order to better replicate the operative environment.

To improve convergence, the script employs a specific initialization strategy. Rather than performing a single continuous sweep, the angle-of-attack progression is split into an ascending and a descending branch, both starting from $\alpha = 0°$. For each subsequent angle, the boundary layer state is initialized using the converged solution of the previous step; this localized restart procedure greatly improves

convergence. After each successful solution, the script saves the global aerodynamic coefficients together with the chordwise surface distributions of pressure C_p and skin friction C_f.

A fundamental requirement of the modular pipeline is uninterrupted execution. However, XFOIL is prone to hanging or diverging when massive flow separation occurs at high angles of attack or on particular airfoil geometries. The Python script therefore monitors the solver output in real time; if a fatal divergence is detected, such as NaN values appearing in the transition routine `TRCHEK2`, or if the process hangs entirely, execution is immediately terminated. Missing data points compromise the subsequent BEMT interpolation; an automated retry logic is therefore implemented to improve polar map data density. If the fraction of successfully converged angles for a given polar falls below a user-defined threshold, the script automatically retries the calculation at a slightly higher Reynolds number. This corrective loop is repeated up to a maximum number of attempts, improving dataset quality without requiring manual user intervention.

3.5 CFD

Computational fluid dynamics (CFD) offers a high-fidelity alternative to the viscous–inviscid interaction approach for airfoil characterization. As previously stated, the unsteady phenomena typical of low-Reynolds regimes are better captured by URANS simulations coupled with a transition model. This requirement, together with the convergence difficulties encountered by XFOIL on the specific T1045-II geometry, motivated the development of a parallel CFD module for polar data generation; the workflow is represented schematically in Fig. 3.6.

The propeller geometry and operative conditions dictate the required grid of incidence angles and Reynolds numbers, and thus define the exact simulation conditions. To handle the large volume of data required by the BEMT databases, the entire CFD procedure is highly automated: a dedicated meshing module generates the required spatial discretization, while a custom set of Bash functions manages the parametric execution within the OpenFOAM framework. The unsteady transient data is subsequently filtered and time-averaged, yielding the steady-equivalent polar coefficients necessary to close the BEMT equations.

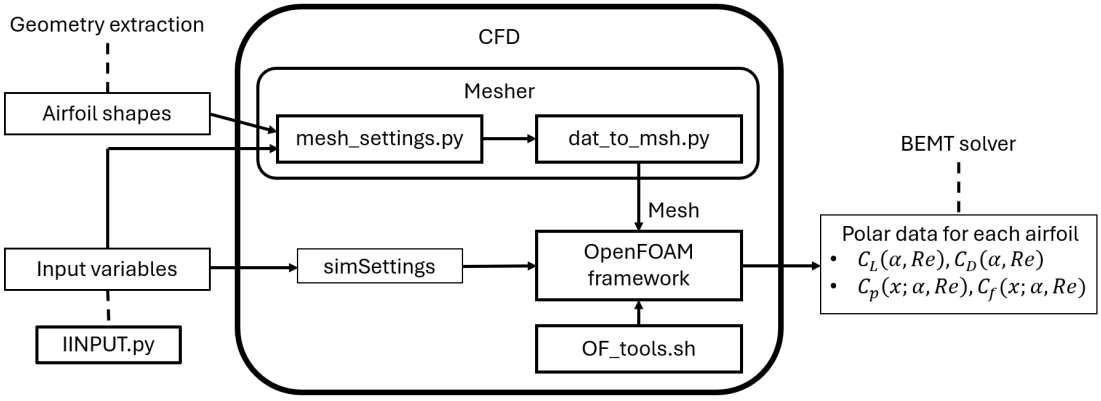

Figure 3.6: CFD block scheme.

3.5.1 Computational Domain and Meshing

The spatial discretization is generated by a dedicated mesher module written in Python using the Gmsh API. This programmatic approach allows complete parametrization of the geometry, enabling fast iterative mesh design and flexibility with respect to airfoil shape. The automation greatly accelerates the simulation workflow by minimizing user intervention.

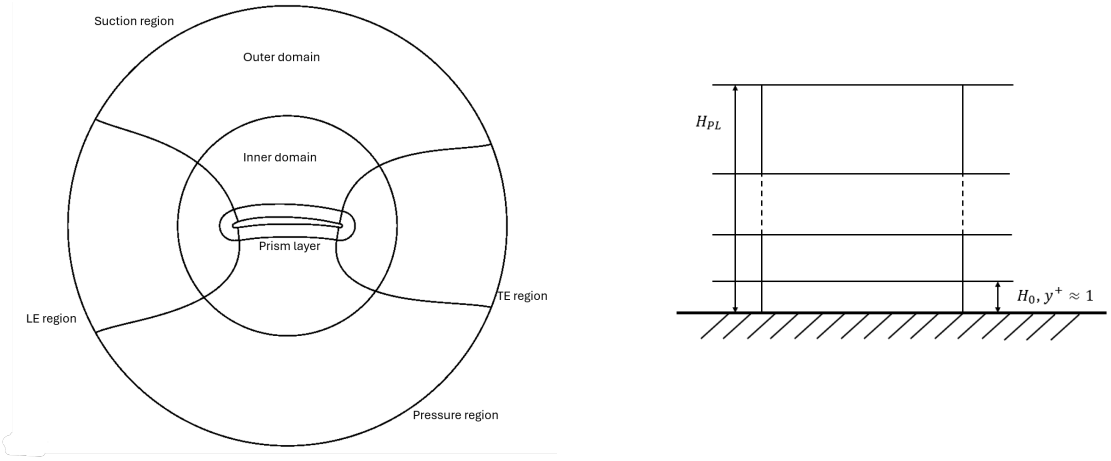

(a) Full domain scheme (not to scale).

(b) Prism layer detail.

Figure 3.7: Mesh topology schematics.

The fluid domain is discretized using a structured O-grid topology (Fig. 3.7a); a unitary extrusion is subsequently applied to yield the two-dimensional mesh

format required by OpenFOAM. To optimize computational cost, the domain is radially divided into three distinct zones: a prism layer to accurately resolve the boundary layer, an inner domain with sufficient refinement to capture the near wake and macroscopic flow structures, and an outer domain with progressively larger cells extending to the far-field boundary. Furthermore, the domain is tangentially subdivided into four regions: the leading edge, the trailing edge, the suction side, and the pressure side. Because the mesh is purely structured, this subdivision grants the user complete control over the discretization. The total cell count and their distribution within each block can be explicitly prescribed in both the radial and tangential directions; cells can thus be clustered precisely at the leading and trailing edges, and an extended tangential refinement is applied to properly resolve the shed wake structures in the inner and outer domains. Finally, an iterative smoothing algorithm is applied to the generated grid to reduce skewness and regularize cell shapes around geometrically challenging features, such as a sharp trailing edge. The script also automatically enforces cell-dimension matching at the block interfaces in the radial direction.

Accurate prediction of aerodynamic transition relies on the URANS k–ω SST model coupled with the γ–Re_θ transition model, which specifically requires proper resolution of the viscous sublayer. The height of the first cell H_0 must therefore guarantee a dimensionless wall distance $y^+ \lessapprox 1$, as illustrated in Fig. 3.7b. To size this first cell, the skin friction coefficient C_f and boundary layer thickness δ are derived from flat-plate correlations, with the script selecting the appropriate regime based on the local Reynolds number Re. For $Re < 5 \times 10^5$, laminar correlations are used:

$$C_f = \frac{1.328}{\sqrt{Re}} \,, \tag{3.5}$$

$$\delta = \frac{1.7208\, c}{\sqrt{Re}} \,. \tag{3.6}$$

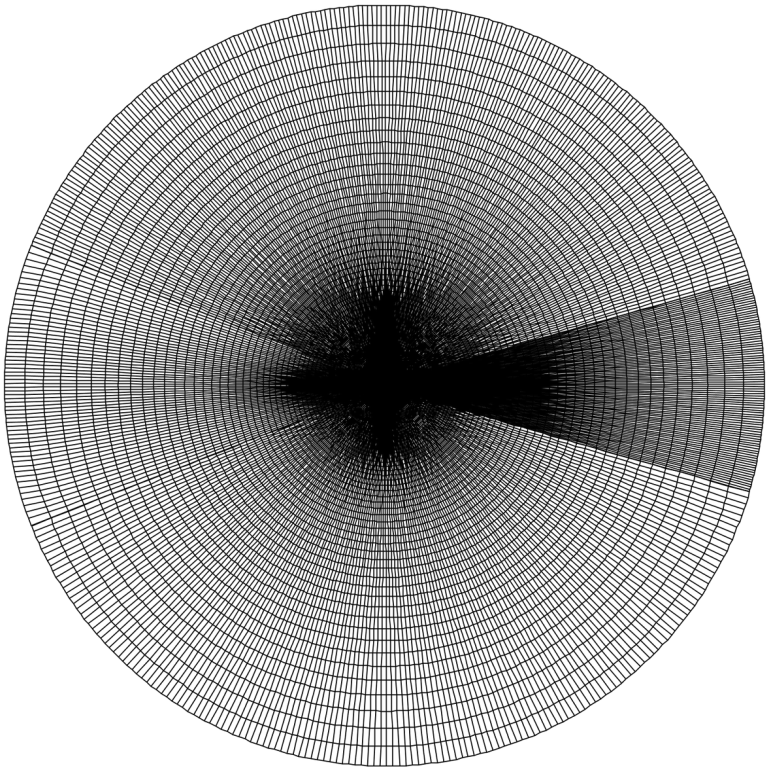

(a) Full domain.

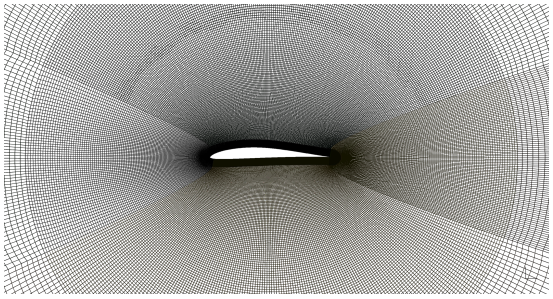

(b) Inner domain.

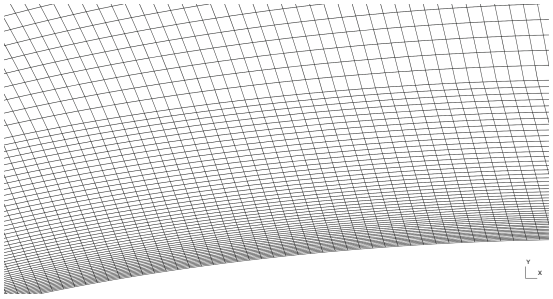

(c) Prism layer.

Figure 3.8: Sample mesh generated by the mesher module.

For $5 \times 10^5 \leq Re < 10^6$, the flow is in a transitional regime:

$$C_f = \frac{0.0592}{Re^{0.2}} \,, \tag{3.7}$$

$$\delta = \max\left(\frac{1.7208\,c}{\sqrt{Re}},\, \frac{0.046\,c}{Re^{0.2}}\right) \,. \tag{3.8}$$

For $Re \geq 10^6$, fully turbulent correlations are applied:

$$C_f = \frac{0.664}{\sqrt{Re}} \,, \tag{3.9}$$

$$\delta = \frac{0.046\, c}{Re^{0.2}} \,. \tag{3.10}$$

The first cell height is then computed to satisfy the near-wall requirement:

$$H_0 = y^+_{\text{target}} \, \frac{c}{Re \sqrt{C_f/2}} \,. \tag{3.11}$$

The total prism layer height H_{PL} is taken as a prescribed multiple of the estimated boundary layer thickness δ, ensuring the entire velocity profile is safely contained within the near-wall region.

An example of a mesh generated by the script is shown in Fig. 3.8.

3.5.2 The OpenFOAM Setup

The CFD simulations in this study were performed using OpenFOAM (Open Field Operation and Manipulation) [33]: a powerful open-source CFD framework running primarily on Linux. Unlike typical commercial software, OpenFOAM provides no graphical user interface; instead, all aspects of case setup and data management are handled through text-based dictionary files. While this requires familiarity with the file structure, it offers complete transparency, direct access to the underlying source code, and great flexibility. This text-based approach also lends itself naturally to automation, making it straightforward to script systematic analyses such as the generation of aerodynamic polar maps.

Physical Model

The governing equations solved in this work are the incompressible, two-dimensional URANS equations, closed with the k–ω SST turbulence model and the γ–Re_θ transition model, as detailed in Section 2.2.2. The γ–Re_θ formulation was selected primarily on the basis of software availability: it is the only transition model natively implemented in OpenFOAM, and therefore provides the most robust and accessible path to capturing boundary layer transition without requiring custom

solver development.

It is important to note that two-dimensional simulations cannot fully replicate the turbulent dynamics of physical three-dimensional flows. In three dimensions, vortex stretching drives a direct energy cascade, in which kinetic energy is transferred from large structures down to the small scales where viscosity dissipates it. Confining the flow to a plane eliminates this mechanism, forcing an inverse energy cascade in which small vortices merge into progressively larger, unphysical macro-structures. As a consequence, two-dimensional models tend to over-predict large-scale flow phenomena and cannot reproduce the turbulent decay observed in three-dimensional reality. Nevertheless, results from laminar Navier–Stokes simulations by Koning et al. [20] on the Eppler 387 airfoil demonstrate that aerodynamic coefficients obtained from two-dimensional simulations retain good agreement with experiments, with the main discrepancies appearing in post-stall and at lower Reynolds numbers. The significant reduction in computational cost therefore justifies the associated loss in accuracy, and the two-dimensional approach is adopted throughout this work.

Boundary Conditions

The governing PDEs require the following boundary conditions for two-dimensional flow over an airfoil. At the airfoil wall:

$$\mathbf{u}(x,y) = \mathbf{0} \ , \quad \frac{\partial p}{\partial n}(x,y) = 0 \ , \quad k(x,y) = k_{wf}(y_0, k_0) \ , \tag{3.12}$$

$$\omega(x,y) = \omega_{wf}(y) \ , \quad \frac{\partial \gamma}{\partial n}(x,y) = 0 \ , \quad \frac{\partial \widetilde{Re_\theta}}{\partial n}(x,y) = 0 \ , \tag{3.13}$$

where the subscript "0" denotes values associated with the first cell at the wall. The wall functions $k_{wf}(y_0, k_0)$ and $\omega_{wf}(y)$ model the inner boundary layer through empirical correlations, and are implemented in OpenFOAM via the `kLowReWallFunction` and `omegaWallFunction` commands, respectively; further details are given in the documentation [34, 35].

At the outlet:

$$\frac{\partial \mathbf{u}}{\partial n}(x,y) = \mathbf{0}\ , \quad p(x,y) = 0\ , \quad \frac{\partial k}{\partial n}(x,y) = 0\ , \quad \frac{\partial \omega}{\partial n}(x,y) = 0\ , \tag{3.14}$$

$$\frac{\partial \gamma}{\partial n}(x,y) = 0\ , \quad \frac{\partial \widetilde{Re}_\theta}{\partial n}(x,y) = 0\ . \tag{3.15}$$

At the inlet:

$$\mathbf{u}(x,y) = \frac{\nu\, Re}{c}\Big(\mathbf{e}_x \cos\alpha + \mathbf{e}_y \sin\alpha\Big)\ , \quad \frac{\partial p}{\partial n}(x,y) = 0\ , \quad k(x,y) = k_\infty\ , \tag{3.16}$$

$$\omega(x,y) = \omega_\infty\ , \quad \gamma(x,y) = 1\ , \quad \widetilde{Re}_\theta(x,y) = Re_{\theta,fs}(Tu_\infty)\ , \tag{3.17}$$

where $Re_{\theta,fs}(Tu_\infty)$ is the freestream correlation for $\widetilde{Re}_\theta$ employed in the transition model [28].

Given the O-grid topology and the necessity of varying the angle of attack, the inlet and outlet boundary conditions are computed dynamically in OpenFOAM using `inletOutlet` for k, ω, γ, and $\widetilde{Re}_\theta$, and `freestreamVelocity` and `freestreamPressure` for the velocity and pressure fields. These conditions determine whether a given cell face acts as an inlet or outlet by evaluating the local velocity flux.

To accurately reproduce user-defined turbulence conditions at the airfoil leading edge, the inlet turbulence quantities are corrected to account for turbulent decay along the streamwise distance. Assuming an isotropic, homogeneous freestream of velocity U, the k–ω SST transport equations reduce to

$$\frac{Dk}{Dt} = -\beta^* k\omega\ , \tag{3.18}$$

$$\frac{D\omega}{Dt} = -\beta\omega^2\ , \tag{3.19}$$

where β and β^* are closure constants of the k–ω SST model. Estimating the convective travel time from the inlet to the leading edge as $t_d = (R_{\text{outer}} - c)/U$,

these equations admit the analytical solution

$$\omega_\infty = \frac{\omega_{\text{target}}}{1 - \beta\,\omega_{\text{target}}\,t_d}\,, \tag{3.20}$$

$$k_\infty = k_{\text{target}}\,(1 + \beta\,\omega_\infty\,t_d)^{\beta^*/\beta}\,, \tag{3.21}$$

which ties the inlet values to the target values at the leading edge. The user prescribes a target turbulence intensity Tu at the leading edge, and the simulation automatically sets the inlet conditions to satisfy this requirement. The corresponding target specific dissipation rate is obtained from

$$\omega_{\text{target}} = \frac{\sqrt{k_{\text{target}}}}{0.09^{1/4}\,l_t}\,, \tag{3.22}$$

where $l_t = C_{lt}\,c$ is the turbulent length scale and C_{lt} is a user-prescribed constant.

Numerical Schemes

The transport equation for the generic scalar variable ϕ is cast in OpenFOAM's discretization framework in its differential, conservative form:

$$\underbrace{\frac{\partial \phi}{\partial t}}_{\text{Increment}} + \underbrace{\nabla \cdot (\mathbf{u}\,\phi)}_{\text{Convection}} = \underbrace{\nabla \cdot (\Gamma_\phi\,\nabla\phi)}_{\text{Diffusion}} + \underbrace{S_\phi}_{\text{Source}}\,, \tag{3.23}$$

where Γ_ϕ denotes the diffusion coefficient of the variable. Following the Finite Volume Method (FVM), OpenFOAM discretizes the governing equations by integrating them over the control volumes defined by the mesh. The specific model equations, those of the URANS solver coupled with the k–ω SST and γ–Re_θ models, are omitted here for brevity, as the principal ones have already been presented in Section 2.2.2. The numerical schemes employed in the simulations are defined in the `fvSchemes` dictionary.

Time derivative. The transient term $\partial/\partial t$ is discretized with a blended Crank–Nicolson scheme (`CrankNicolson 0.9`), in which the second-order Crank–Nicolson (CN) method is blended with the first-order implicit Euler scheme to improve

stability. The blending parameter $\psi_{CN} \in [0,1]$ controls the relative weight of the two schemes: $\psi_{CN} = 1$ recovers pure CN, while $\psi_{CN} = 0$ recovers pure implicit Euler. The discretized time derivative of the generic variable ϕ reads

$$\frac{\phi_i^{n+1} - \phi_i^n}{\Delta t} = \psi_{CN} \cdot \frac{1}{2}\Big[F_i^{n+1}(\mathbf{x}, t, \phi) + F_i^n(\mathbf{x}, t, \phi)\Big] + (1 - \psi_{CN})\, F_i^n(\mathbf{x}, t, \phi) \ , \quad (3.24)$$

where $F_i(\mathbf{x}, t, \phi)$ denotes the right-hand side of the transport equation evaluated at cell i, and the superscript indicates the time-step level. For the present simulations $\psi_{CN} = 0.9$ is adopted.

Gradient scheme. The gradient interpolation scheme, required by the solver to compute higher-order discretizations and auxiliary quantities such as the strain-rate tensor, is set to `cellLimited Gauss linear 1` for all variables except velocity. This scheme evaluates the gradient of ϕ at cell i of volume Ω by applying Gauss's theorem:

$$\int_\Omega \nabla\phi \, d\Omega = \oint_{\partial\Omega} \phi \, \mathbf{n} \, d\sigma \ , \quad (3.25)$$

$$\nabla\phi_i \approx \frac{1}{\Omega} \sum_{f=1}^{N_f} \phi_f \, \mathbf{n}_f \, S_f \ , \quad (3.26)$$

where f indexes the faces of cell i, S_f and $\mathbf{n}_f$ are the face area and outward unit normal, and ϕ_f is the variable evaluated at the face centroid. Because the solver stores data at cell centres only, the `linear` keyword obtains ϕ_f by a distance-weighted interpolation between neighbouring cell-centre values.

While this volume integration is highly accurate, unconstrained gradients can produce unphysical values on complex meshes. The `cellLimited` modifier therefore scales the raw gradient by a scalar limiter $\alpha_i \in [0,1]$:

$$\nabla\phi_i^{\lim} = \alpha_i \, \nabla\phi_i \ . \quad (3.27)$$

To prevent extrapolated face values from overshooting the local solution range, the solver first identifies the extrema $\phi_{\max}$ and $\phi_{\min}$ over cell i and its immediate

neighbours, then computes a local limiter α_f for every face f:

$$\alpha_f = \begin{cases} \dfrac{\phi_{\max} - \phi_i}{\nabla\phi_i \cdot \mathbf{d}_{if}} & \text{if } \nabla\phi_i \cdot \mathbf{d}_{if} > 0 \ , \\ \dfrac{\phi_{\min} - \phi_i}{\nabla\phi_i \cdot \mathbf{d}_{if}} & \text{if } \nabla\phi_i \cdot \mathbf{d}_{if} < 0 \ , \end{cases} \tag{3.28}$$

where $\mathbf{d}_{if}$ is the vector from the centre of cell i to face centroid f. The cell limiter α_i is set to the minimum over all face limiters, so the gradient is reduced just enough to satisfy the most restrictive face constraint and suppress numerical oscillations.

It was found that applying this limiter to the velocity-gradient scheme introduces instability in the transition model calculations; the root cause requires further investigation. For the present work, an unlimited `Gauss linear` scheme is therefore adopted for the velocity gradient.

Divergence scheme. The convective fluxes are discretized with the second-order linear-upwind scheme (`linearUpwind` in OpenFOAM). Upwind methods preserve solution stability by reconstructing face values from the upstream cell, consistent with the hyperbolic nature of convection. Applying Gauss's theorem, the convective flux through the boundary of cell i is

$$\int_\Omega \nabla \cdot (\phi\,\mathbf{u})\, d\Omega = \oint_{\partial\Omega} \phi\,\mathbf{u} \cdot \mathbf{n}\, d\sigma \approx \sum_{f=1}^{N_f} \phi_f\,(\mathbf{u}_f \cdot \mathbf{n}_f)\, S_f \ . \tag{3.29}$$

The face values ϕ_f are obtained by linear reconstruction from the upwind cell centre. For two cells i and $i+1$ sharing face f, the reconstructed values are

$$\phi_{f,i} = \phi_i + \nabla\phi_i \cdot \mathbf{d}_{if} \ , \tag{3.30}$$

$$\phi_{f,i+1} = \phi_{i+1} + \nabla\phi_{i+1} \cdot \mathbf{d}_{(i+1)f} \ , \tag{3.31}$$

where $\mathbf{d}_{if}$ is the vector from the centre of cell i to face f, and $\nabla\phi$ is the cell-centre gradient. Defining $\mathbf{n}_f$ as the face normal pointing out of cell i, the upwind flux

selection reads

$$\phi_f \left(\mathbf{u}_f \cdot \mathbf{n}_f\right) S_f \approx \begin{cases} \phi_{f,i} \left(\mathbf{u}_f \cdot \mathbf{n}_f\right) S_f & \text{if } \mathbf{u}_f \cdot \mathbf{n}_f \geq 0 \ , \\ \phi_{f,i+1} \left(\mathbf{u}_f \cdot \mathbf{n}_f\right) S_f & \text{if } \mathbf{u}_f \cdot \mathbf{n}_f < 0 \ . \end{cases} \tag{3.32}$$

The advecting velocity $\mathbf{u}_f$ is pre-computed by the solver to satisfy global continuity. In addition, OpenFOAM applies a gradient-based limiter to the reconstructed face values, bounding them within the local solution extrema and preventing spurious oscillations in regions of strong gradients.

Case Management and Post-Processing

The structured case management required to build an airfoil performance dataset is implemented through custom Bash functions stored in `OF_tools.sh`. The function `derivedCases` generates a family of cases from a single base case by modifying the operative conditions according to user-defined variables and their values, stored in `parametersFile.sh`. The cases generated in this way are grouped in a dedicated folder and can be executed via `foamRun_rec`, which recursively applies a `foamRun` command to a given directory. To quickly modify simulation settings and control parameters across a group of cases, the custom function `castCase` provides this functionality recursively on the corresponding folder.

For post-processing, OpenFOAM function objects are used to extract the aerodynamic coefficients $C_L(t; \alpha, Re)$ and $C_D(t; \alpha, Re)$ as time signals, and the surface distributions $C_p(x; \alpha, Re)$ and $C_f(x; \alpha, Re)$ via surface sampling.

The custom function `extractPolarMapAvg` computes the time-averaged polar map entries $\overline{C}_L(\alpha, Re)$ and $\overline{C}_D(\alpha, Re)$ from a group of simulations. It first filters the initial transient and then evaluates the mean coefficients. For each case, the end of the transient is identified by monitoring the moving average of the drag coefficient, defined as

$$\overline{C}_{D,i} = \frac{1}{t_n - t_i} \sum_{j=i}^{n} C_{D,j} \left(\frac{t_{j+1} - t_{j-1}}{2}\right) \ , \quad i = 2, \ldots, n-1 \ , \tag{3.33}$$

where t_i and $C_{D,i}$ are the i-th time instant and the corresponding drag coefficient. The transient is considered ended once all moving averages within a specified time

window fall within a user-defined relative tolerance. The function `extractCpCfAvg` applies the same transient-filtering procedure to the drag coefficient and then averages the pressure and friction coefficient distributions over the remaining time steps to obtain the mean chordwise profiles. All output data is saved in a standardized file structure to facilitate readability and management.

Chapter 4

Airfoil characterization results

In this chapter the airfoil characterization results are presented, as obtained with the XFOIL and CFD methods discussed in the previous chapter. First, a test case on the Eppler 387 airfoil will be reported, to serve as proof of validity of the numerical simulations.

Before applying the numerical methods to the T1045-II propeller sections, it is necessary to establish their reliability in the low-Reynolds regime on a well documented test case. For this reason, the chapter opens with a validation study on the Eppler 387 airfoil, a geometry specifically designed for low-Reynolds applications for which extensive experimental data exists across multiple facilities. This validation serves a dual purpose: it assesses the accuracy of both XFOIL and CFD against ground truth, and it provides confidence in the mesh and numerical setup adopted for the CFD simulations, in lieu of a dedicated mesh convergence study.

The validation is followed by the characterization of a representative airfoil of the T1045-II propeller, which was selected as test case for the full numerical methodology as it has been used at the VKI for experimental studies on low-Reynolds aerodynamics. The polar maps generated are then fed to the BEMT for evaluation of the full methodology against experimental data in the following chapter.

4.1 Test case: the Eppler 387

The Eppler 387 (e387) is an airfoil design specifically developed for low-Reynolds applications (Fig. 4.1). The flow over this airfoil perfectly exemplifies the low-Reynolds regime dynamics covered in Section 2.1; the flow's sensibility noted in experimental campaigns [5, 36, 37] provide a challenging benchmark test case for any kind of numerical simulation.

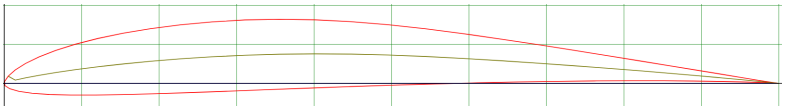

Figure 4.1: Eppler 387 airfoil geometry.

The experimental studies mentioned above were conducted at three different facilities, namely the LTPT at NASA Langley Research Center [5], the Low-Turbulence Tunnel at the Delft University of Technology [36] and the Model Wind Tunnel at Stuttgart [37]. The present work replicated the test conditions of the the three experimental studies with both XFOIL and CFD simulations, for an incidence range of $0° - 12°$ at a Reynolds number value of 60,000, in essentially incompressible flow conditions, and reproducing the turbulence intensity levels measured at the three facilities ($Tu \approx 0.0003$ at Stuttgart, $Tu \approx 0.0008$ at Delft and $Tu = 0.0018$ at NASA Langley). The numerical methodology is the same as the one explained for the polar data generation blocks in the modular pipeline.

For the CFD calculation, a mesh of 500x240 cells was used[1], based off of comparable simulations from Zhang et al. [38]. The mesh was generated by the custom mesher for the e387 airfoil, specifically for $Re = 60{,}000$ (3.8). The target turbulent length scale was set to $l_t = 0.04c$, to relfect the small scale of vortices in a low-turbulence wind tunnel.

XFOIL's discretization employed 200 panels distributed on the airfoil, with standard clustering at the leading edge and at the trailing edge.

[1]The dimensions indicate $n_{tangential} \times n_{radial}$ cells.

4.1.1 Comparison to experimental data

Aerodynamic coefficients

The comparison of the numerical polar data and experimental results is shown in Fig. 4.2. As anticipated by the theoretical discussion in Section 2.1, the low-Reynolds regime is highly sensitive to freestream conditions, surface roughness, and geometry; consequently, the three experimental datasets do not agree perfectly with one another, although the small differences in experimental conditions. Both the NASA Langley and Delft datasets exhibit an anomalous drop in C_L accompanied by a sharp rise in C_D in the interval $\alpha \approx 5°$–$7°$ (Fig. 4.2), indicative of complete laminar separation without reattachment. As discussed in Section 2.1, at $Re = 60\,000$ the boundary layer is insufficiently energized to sustain turbulent transition in the free shear layer, and the laminar separation bubble may burst entirely. The NASA Langley campaign additionally reports static hysteresis at the boundaries of this interval, where the flow alternates between reattachment and complete separation depending on the direction of the angle-of-attack sweep [5]; this bistable behavior is consistent with the separation mechanism discussed in the theory.

Considering the lift coefficient, both XFOIL and CFD provide generally satisfactory predictions throughout the attached-flow range, capturing the overall slope and magnitude of $C_L(\alpha)$ reasonably well (Figs. 4.2a–4.2e), although matching the experimental turbulence intensity doesn't seem to improve accuracy with respect to the corresponding experimental data. The drag coefficient proves considerably more challenging for both methods, with relative errors that are substantially larger than those observed for C_L, particularly at high angles of attack. This is consistent with expectations: at high incidence the adverse pressure gradient intensifies, boundary layer dynamics become more severe, and the laminar separation bubble grows, all of which place greater demands on the underlying turbulence and transition models.

Regarding the anomalous separation interval, the complete separation occurring at $\alpha \approx 5°$–$7°$ proves very difficult to capture for both solvers. Interestingly, the XFOIL calculation at $Tu = 0.0003$ (Figs. 4.2a, 4.2b) comes closest to reproducing the trend in both C_L and C_D in this range; however complete separation was observed experimentally at NASA Langley and Delft, which operated at $Tu =$

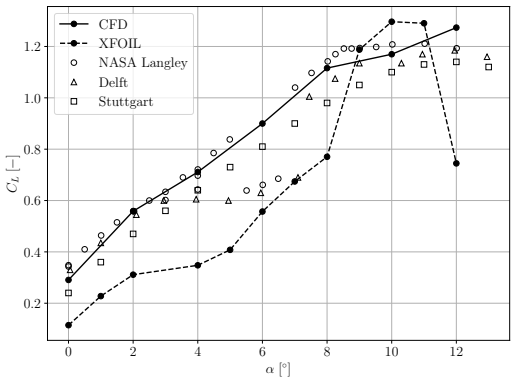

(a) $C_L(\alpha)$ at $Tu = 0.0003$.

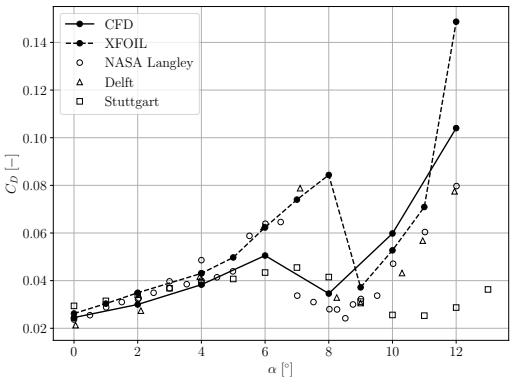

(b) $C_D(\alpha)$ at $Tu = 0.0003$.

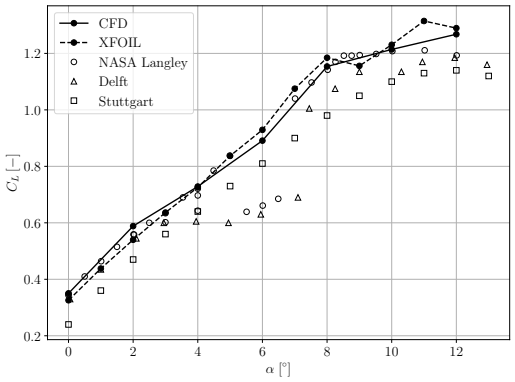

(c) $C_L(\alpha)$ at $Tu = 0.0008$.

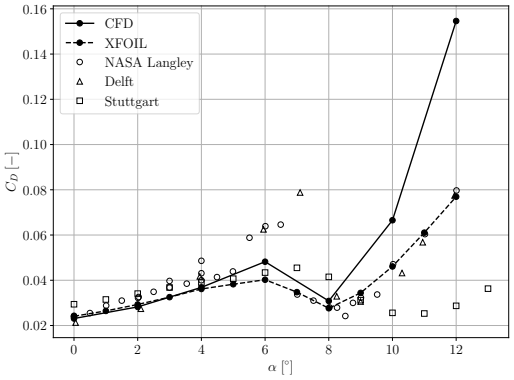

(d) $C_D(\alpha)$ at $Tu = 0.0008$.

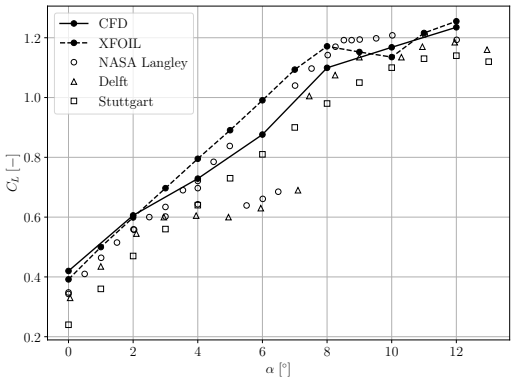

(e) $C_L(\alpha)$ at $Tu = 0.0018$.

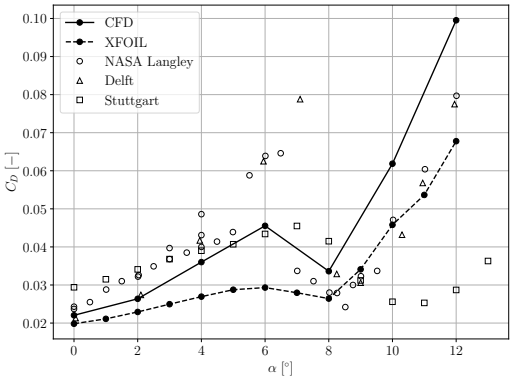

(f) $C_D(\alpha)$ at $Tu = 0.0018$.

Figure 4.2: Numerical aerodynamic coefficients against experiments.

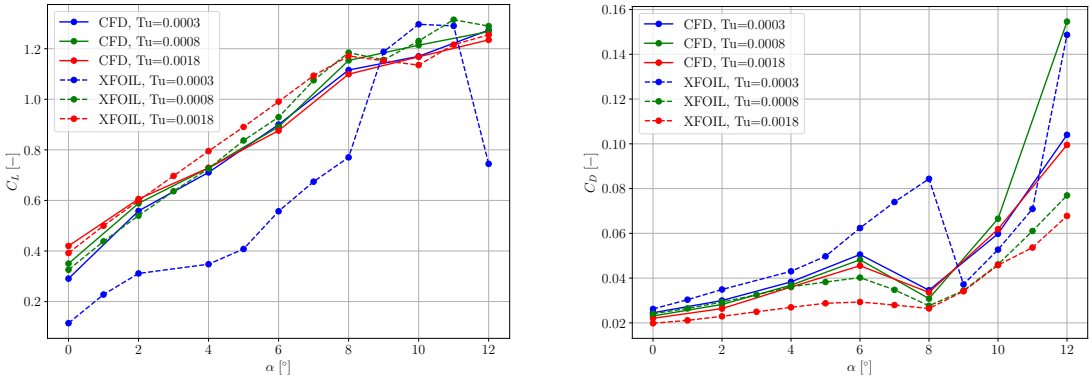

(a) $C_L(\alpha)$. **(b)** $C_D(\alpha)$.

Figure 4.3: Direct comparison of aerodynamic coefficients from CFD and XFOIL.

0.0018 and $Tu = 0.0008$ respectively, conditions where the higher freestream turbulence intensity should promote earlier transition and thus favor reattachment rather than full separation. Furthermore, this apparent match comes at the cost of significant errors throughout the rest of the incidence range in the lift coefficient, again proving that accurate modeling of complex low-Reynolds phenomena is very difficult to achieve.

In the post-stall regime, XFOIL surprisingly yields better drag predictions than CFD, which may reflect the known tendency of two-dimensional URANS simulations to over-predict large-scale separation structures, as previously discussed in this chapter. On the other hand, at the highest turbulence intensity ($Tu = 0.0018$), CFD produces drag predictions that agree more closely with the experimental data than XFOIL (Figs. 4.2f), suggesting that the URANS transition model captures the enhanced mixing effect of elevated freestream turbulence more faithfully in attached and mildly separated conditions.

Comparing the two numerical methods more directly (Fig. 4.3), XFOIL exhibits greater sensitivity to the prescribed turbulence intensity than the URANS simulations: the spread between the $Tu = 0.0003$ and $Tu = 0.0018$ coefficients' curves is noticeably wider in XFOIL than in CFD across both C_L and C_D. This might be a consequence of the viscous–inviscid interaction framework underlying XFOIL, in which the e^N transition criterion with Mack's correlation (Eq. 2.10) directly dictates the transition location in response to N_{crit}.

Pressure and friction coefficient distributions

The pressure coefficient distributions across the incidence range at $Tu = 0.0018$, together with the friction coefficient distributions from the numerical solvers, are reported in Fig. 4.5. The experimental pressure data from NASA Langley are included for comparison; friction coefficient data are not available experimentally.

The whole incidence range examined is governed by laminar separation dynamics. The clearest indicator is the characteristic plateau in the C_p distribution on the suction side: as discussed in Section 2.1.2, a laminar separation bubble suppresses the pressure recovery of the outer flow, imposing an approximately constant pressure over the bubble extent. This is visible at all angles of attack in the numerical results, and is broadly reproduced by both CFD and XFOIL. The friction coefficient distributions confirm this picture, as regions of negative C_f on the suction side indicate reverse flow, the hallmark of a recirculation bubble.

With incidence at 2° and 4°, both solvers predict laminar separation near the trailing edge with no reattachment; the bubble extends to the trailing edge, consistent with the long-bubble regime expected at $Re = 60\,000$, as discussed in the theory. The C_p plateau is visible in the aft portion of the suction side, and the experimental data are well reproduced by both methods. As the angle of attack increases from 2° to 4°, the mean laminar separation point moves upstream, as visible from the friction coefficient distributions of both CFD and XFOIL. This is the classical behaviour described in Section 2.1.3: the increasing adverse pressure gradient on the suction side destabilises the boundary layer earlier, causing separation to advance toward the leading edge.

At $\alpha = 6°$, the NASA Langley experiment records a dramatic degradation in both C_L and C_D, as already noted in the polar analysis. The C_p distribution confirms this: the experimental data show a near-complete loss of suction peak, indicative of massive laminar separation without reattachment over most of the chord. The CFD calculation at $Tu = 0.0018$ completely fails to reproduce this behaviour, predicting instead a moderate bubble and substantially higher suction levels, consistent with what was observed in the aerodynamic coefficient comparison. Interestingly, the XFOIL calculation at $Tu = 0.0003$ (Fig. 4.7) captures this behaviour remarkably well: the predicted C_p distribution is in close agreement with the experimental data, and the friction coefficient reveals a very early laminar separation at approximately

$x/c \approx 0.25$ with no subsequent reattachment. The absence of pressure recovery and the large separated region directly explain the lift deficit and drag surge seen in the polar data. As noted in the discussion on theory, the physical cause is the inability of the laminar shear layer to transition and reattach at this Reynolds number when separation occurs too far upstream, leading to complete bursting of the laminar separation bubble.

At $\alpha = 8°$, the numerical calculations regain accuracy. CFD in particular performs well, predicting two distinct laminar separation events on the suction side followed by reattachment, as evidenced by two negative C_f regions with subsequent recovery. It should be noted, however, that the CFD result at this incidence represents a time-averaged field of an intrinsically unsteady flow: as will be discussed in the next section, at $\alpha = 8°$ the solver predicts significant vortex shedding and dynamic loading. The averaged C_p therefore reflects the mean effect of large vortical structures being periodically shed, rather than a steady separation bubble. XFOIL, whose viscous formulation is inherently steady, cannot fully capture this averaged distribution.

Furthermore, we note that XFOIL's predictions for the C_f distribution are more moderate than the CFD's across the whole range, consistently with its known inability to model large-scale recirculation and shed vortices.

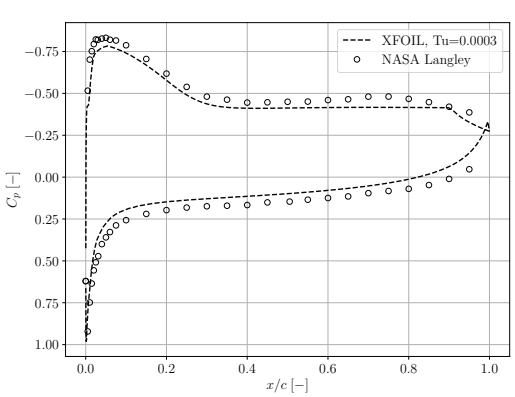

(a) Pressure coefficient distribution.

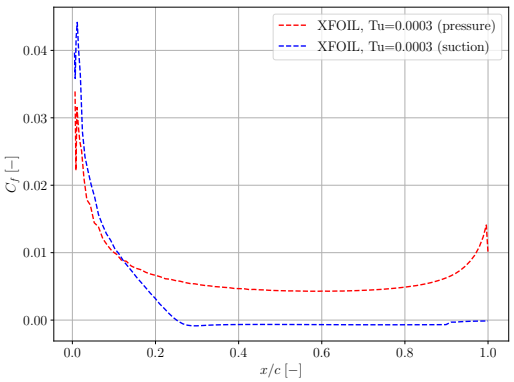

(b) Friction coefficient distribution (XFOIL only).

Figure 4.4: Pressure and friction coefficients comparison at $\alpha = 6°$, NASA Langley and XFOIL ($Tu = 0.0003$).

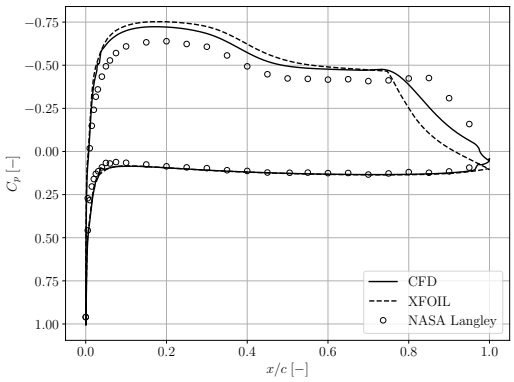

(a) $\alpha = 2°$.

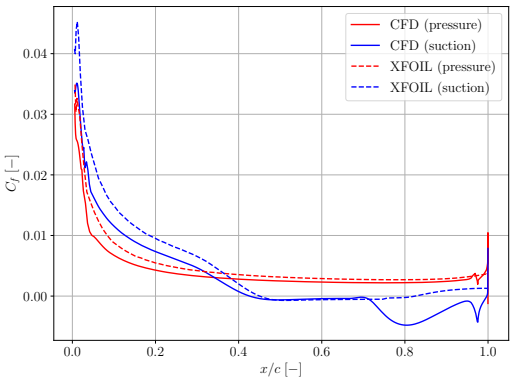

(b) $\alpha = 2°$.

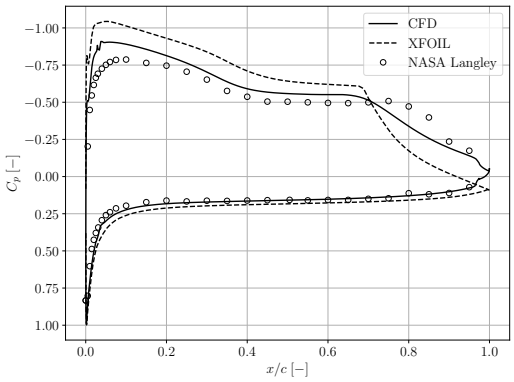

(c) $\alpha = 4°$.

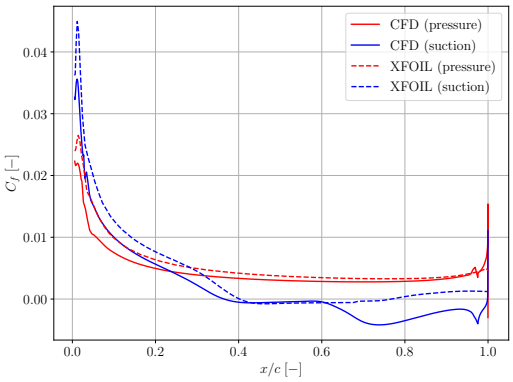

(d) $\alpha = 4°$.

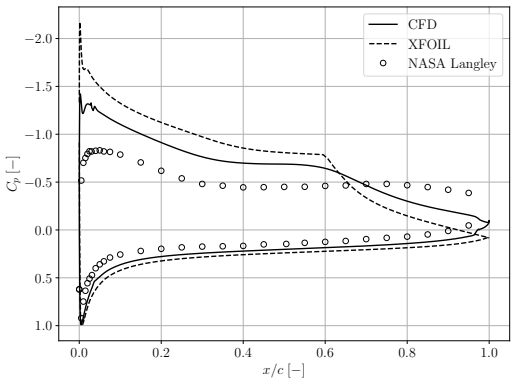

(e) $\alpha = 6°$.

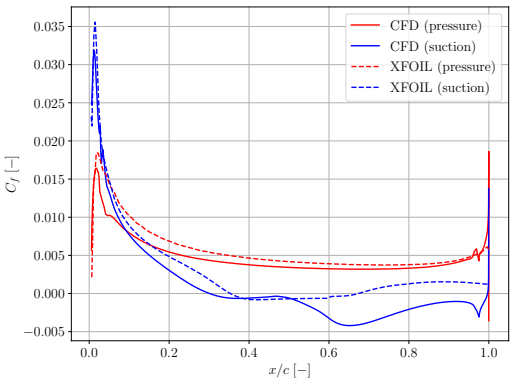

(f) $\alpha = 6°$.

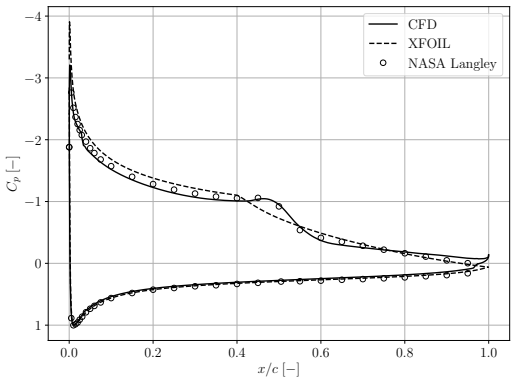

(g) $\alpha = 8°$.

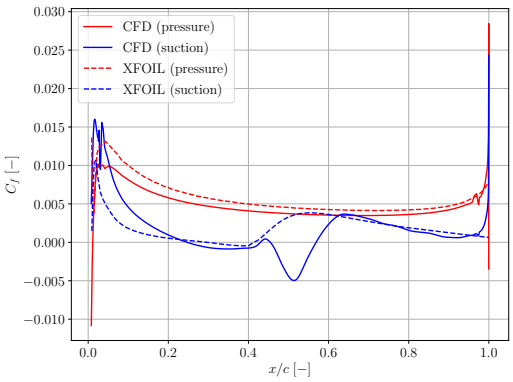

(h) $\alpha = 8°$.

Figure 4.5: Pressure and friction coefficients comparison, at $Tu = 0.0018$.

Unsteady loading

The drag coefficient histories from CFD at $Tu = 0.0018$ are reported in Fig. 4.6 for $\alpha = 4°$ and $\alpha = 8°$, together with instantaneous velocity magnitude fields with superimposed streamlines in Figs. 4.7 and 4.8.

At $\alpha = 4°$, the drag coefficient exhibits a clearly periodic trend, upon which minor oscillations of smaller amplitude are occasionally superimposed (Fig. 4.6a). The velocity field snapshots in Fig. 4.7 illustrate the physical mechanism underlying this behaviour. The first snapshot (Fig. 4.7a), captured at a drag minimum, shows a compact laminar separation bubble on the suction side, with the flow reattaching ahead of the trailing edge; the reduced bubble extent limits the dead-air region and its associated pressure drag penalty. The second snapshot (Fig. 4.7b), captured at a drag maximum coinciding with the minor oscillations, reveals instead a significantly enlarged separation bubble, with the flow failing to reattach to the body and the wake exhibiting lateral oscillations. The dominant low-frequency periodicity is consistent with the bubble flapping mechanism described in Section **??**: the reattachment point oscillates back and forth, alternately shrinking and enlarging the bubble at a characteristic frequency. The superimposed minor oscillations are attributed to small-scale vortical structures shed in the wake during the phases of enlarged separation.

At $\alpha = 8°$, the unsteady character of the flow is considerably more pronounced, and the drag signal in Fig. 4.6b reflects a well-established vortex shedding regime. As discussed in Section 2.1, Kelvin–Helmholtz instabilities in the laminar shear layer following separation are amplified and roll up into discrete vortices, which are periodically shed into the wake. The velocity snapshot in Fig. 4.8 captures one such shedding event, with coherent vortical structures clearly visible downstream of the separation point. The enhanced turbulent mixing operated by these vortices promotes momentum transport toward the wall, favouring mean reattachment despite the severity of the separation; this is consistent with the negative C_f regions followed by recovery observed in the time-averaged friction coefficient distribution (Fig. 4.5h). The time-averaged C_p and C_f distributions discussed in the previous section therefore reflect the statistical effect of this periodic shedding rather than a steady flow field, which explains both the good mean agreement between CFD and experiment and the inability of the steady XFOIL formulation to fully reproduce

the averaged distributions at this incidence.

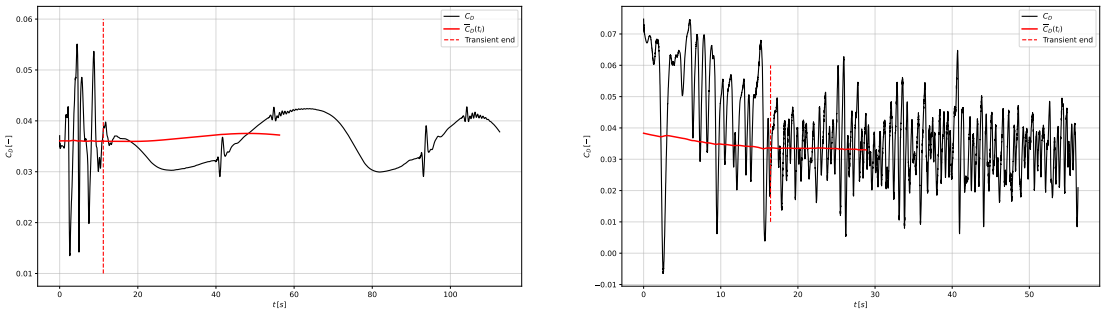

(a) $\alpha = 4°$. **(b)** $\alpha = 8°$.

Figure 4.6: Drag coefficients in time from CFD, at $Tu = 0.0018$.

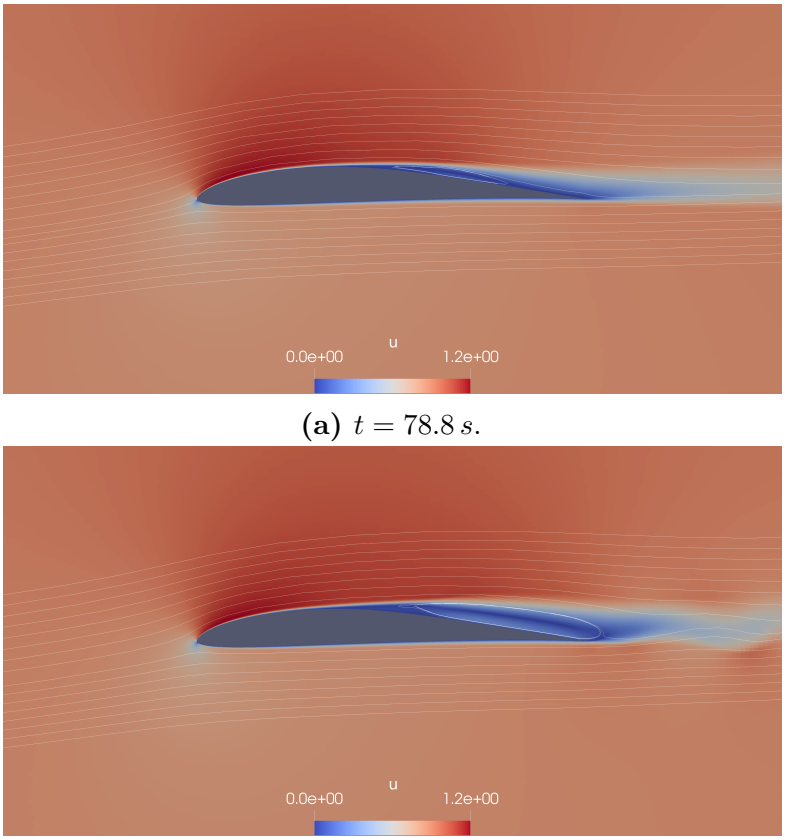

(a) $t = 78.8\,s$.

(b) $t = 56.3\,s$.

Figure 4.7: Velocity magnitude snapshots from CFD, $Tu = 0.0018$, $\alpha = 4°$.

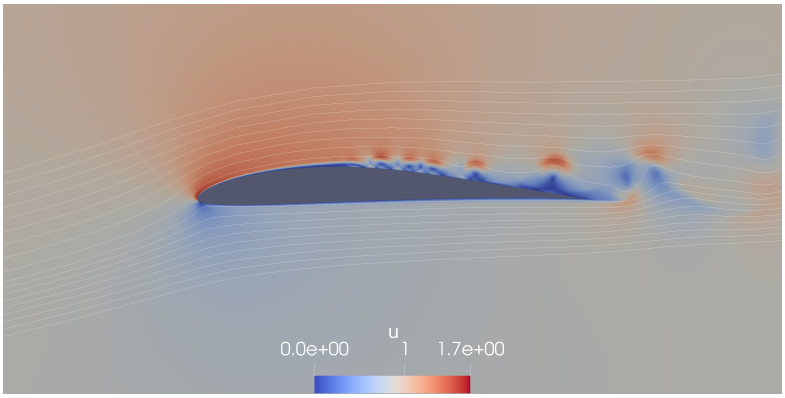

Figure 4.8: Velocity magnitude snapshot, $t = 45.0\,s$.

4.1.2 Tools assessment

Overall, both XFOIL and CFD demonstrate a reasonable capacity to predict airfoil aerodynamic coefficients in the low-Reynolds regime, despite falling short of accurately capturing its full phenomenological complexity. The anomalous separation interval at $\alpha \approx 5°$–$7°$ is a clear illustration of the extreme sensitivity of this regime: neither solver reliably reproduces the complete laminar separation observed experimentally, and the accidental agreement of XFOIL at mismatched turbulence conditions underlines how difficult it is to predict these phenomena in a physically consistent way. Away from this interval, the lift coefficient is predicted satisfactorily by both methods across the attached-flow range, while drag proves more challenging, as errors grow with incidence as boundary layer dynamics become more severe, though predictions remain acceptable at moderate angles of attack where the flow is largely attached. These results provide sufficient confidence to proceed with the application of both tools to the T1045-II propeller sections: XFOIL and CFD can be used to generate polar data for BEMT, but the limitations identified here, particularly around drag accuracy and the breakdown of both methods in strongly separated conditions, must be kept in mind when interpreting the propeller-level calculations that follow.

4.2 Numerical polar data for T1045-II

The T1045-II is a commercially available small propeller representative of those mounted on small multi-rotor platforms. Its geometry was acquired via structured-light scanning at the VKI, and the airfoil sections were extracted using the procedure described in Section 3.2. As discussed there, the bulk of thrust and torque in a propeller is generated in the mid-to-outer span region, approximately $r_{nd} \approx 0.5$–0.8; a single representative section was therefore selected for the polar map calculations, located at the non-dimensional radial coordinate

$$r_{nd} = \frac{r - R_{hub}}{R - R_{hub}} = 0.65 \,, \tag{4.1}$$

where r is the local radius, R_{hub} is the hub radius, and R is the propeller tip radius. This section, hereafter referred to as a0.65, is shown in Fig. 4.9 and is well representative of the local blade geometry in the thrust-generating region.

Two aerodynamic coefficient datasets were constructed for a0.65 following the numerical methodology described in Chapter 3: one using XFOIL and one using CFD. The XFOIL dataset covers a Reynolds number range of $Re = 10\,000$–$100\,000$ and an incidence range of $\alpha = -5°$–$15°$, while the CFD dataset is restricted to $Re = 50\,000$–$80\,000$ and $\alpha = 0°$–$10°$, reflecting the higher computational cost of the simulations. Both datasets adopt a freestream turbulence intensity of $Tu = 0.001$, representative of a low-turbulence atmospheric environment. XFOIL employed a discretization of 200 panels on the airfoil surface, with the same setup described in Section 3.4. The CFD mesh (Fig. 4.10) was generated with the same mesher and settings adopted for the Eppler 387 validation, with the wall-normal first cell height sized for $Re = 70\,000$ to maintain $y^+ \approx 1$ across the full Reynolds number range considered.

4.2.1 XFOIL and CFD comparison

The aerodynamic coefficient maps produced by XFOIL and CFD are shown in Figs. 4.11 and 4.12 respectively, while the absolute difference $\Delta C = C^{\mathrm{XFOIL}} - C^{\mathrm{CFD}}$ is reported in Fig. 4.13 for both C_L and C_D. It should be noted that the XFOIL dataset covers a substantially wider range of Reynolds numbers and angles of attack

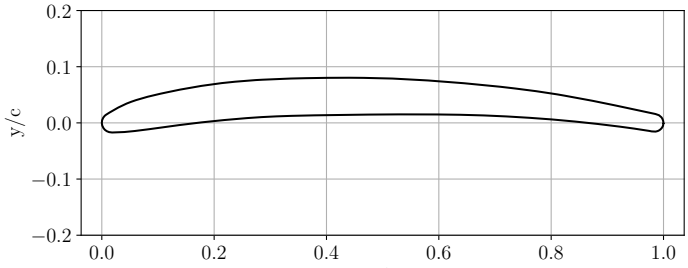

Figure 4.9: Geometry of the a0.65.

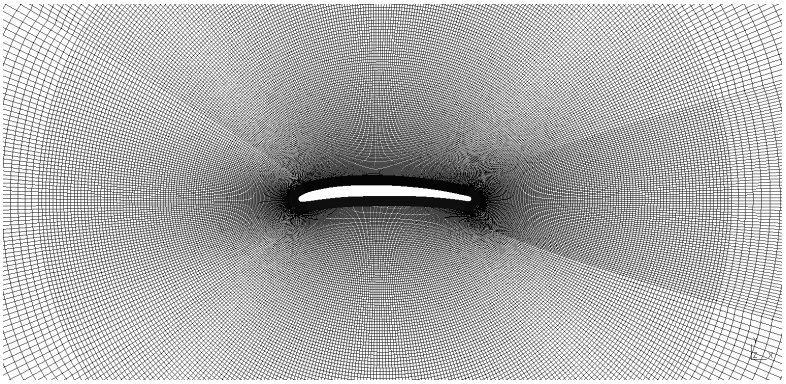

Figure 4.10: Mesh of the a0.65.

than the CFD dataset, reflecting the enormous difference in computational cost between the two methods.

Both solvers show a consistent degradation of aerodynamic performance as the Reynolds number decreases, with the trend being most pronounced in XFOIL below $Re \approx 40\,000$, where the low-Reynolds effects discussed in Section 2.1.3 become increasingly severe. The same tendency is present in the CFD results, although it is less immediately apparent due to the narrower Reynolds number range computed.

Regarding the drag coefficient difference (Fig. 4.13b), the two methods show good agreement for the range $4° - 6°$, where the flow remains largely attached and the boundary layer dynamics are relatively benign. However, at very low incidence discrepancies grow, reaching about 20%. This also happens beyond $\alpha \approx 6°$–$8°$, where the onset of unsteady dynamics place greater demands on both solvers. These discrepancies were to be expected from the Eppler 387 validation case.

The lift coefficient difference (Fig. 4.13a) also presents some discrepancies. It has to be noted that the agreement between XFOIL and CFD actually improves at higher angles of attack, where one might expect the largest discrepancies. More remarkably, at low incidence the two methods diverge significantly: $\Delta C_L \approx 0.5$ at $\alpha = 0°$ across the entire Reynolds number range, a substantial error given that $C_{L,\max} \approx 1.3$ in the domain considered.

The pressure and friction coefficient distributions at $Re = 60\,000$ are reported in Fig. 4.14 for $\alpha = 0°$, $2°$, and $4°$, which seems to be the range to further investigate, given the good agreement that was registered for the Eppler 387 in this same range. Both the pressure and friction coefficients exhibit meaningful discrepancies at very low incidence, particularly along the pressure side of the airfoil, with the magnitude of the error decreasing as the angle of attack increases. At $\alpha = 4°$ the two methods agree reasonably well, while at $\alpha = 0°$ the differences are most pronounced.

At low incidence, both the pressure and friction distributions indicate the presence of a laminar separation bubble on the pressure side of the airfoil in the vicinity of the leading edge, with subsequent reattachment. This structure is captured by both solvers, but in substantially different ways. XFOIL predicts a larger bubble, which prevents proper pressure recovery on the pressure side and maintains a higher local pressure there; simultaneously, the suction peak on the suction side is greatly reduced relative to the CFD prediction. The combined effect of these two discrepancies — excessive pressure-side pressure and a weaker suction peak — directly explains XFOIL's large underprediction of C_L at low incidence with respect to CFD. As the angle of attack increases, the pressure-side bubble retreats and eventually disappears, and XFOIL and CFD recover agreement on both the pressure distribution and the integrated coefficients, consistent with the ΔC_L map in Fig. 4.13a.

The velocity magnitude snapshots in Fig. 4.15 confirm the presence of the laminar separation bubble on the pressure side at low incidence. The snapshots also reveal a feature not observed on the Eppler 387: the a0.65 exhibits vortex shedding from the trailing edge even at $\alpha = 0°$, producing unsteady loading also at low incidence. This behaviour is attributable to the blunt trailing edge geometry of the section, which acts as a bluff body and triggers a Kármán-like shedding process independently of the angle of attack. At higher incidence, the separation

point migrates progressively toward the leading edge, consistent with the behaviour described in Section 2.1.3. At $\alpha = 8°$, separation occurs just downstream of the leading edge and large vortical structures begin to form in the shear layer; the associated entrainment and mixing favour mean reattachment and partially recover aerodynamic performance, as also observed on the Eppler 387 at comparable conditions.

The airfoil characterization results for the a0.65 section reveal a more challenging picture than the Eppler 387 validation suggested. The unforeseen discrepancy at low incidence, driven by the pressure-side laminar separation bubble, represents an additional source of uncertainty that was not anticipated from the validation case alone, and highlights how the specific geometry of a propeller section can introduce difficulties beyond those observed on canonical airfoils. Between the two methods, CFD is likely to provide more reliable predictions: its higher-fidelity transition model resolves the bubble dynamics more accurately, and its unsteady formulation is better suited to the vortex shedding observed even at low angles of attack on this geometry. XFOIL remains useful for its computational efficiency and broader operative range, but its polar data should be interpreted with caution in the low-incidence regime where the pressure-side bubble dominates. Despite these limitations, both datasets will be employed in the BEMT calculations of the following chapter, where the propeller performance predictions are compared against experimental measurements. This comparison will also serve to assess how sensitive the BEMT predictions are to the choice of polar source, and therefore whether the differences identified here between XFOIL and CFD actually propagate into meaningful differences at the propeller level.

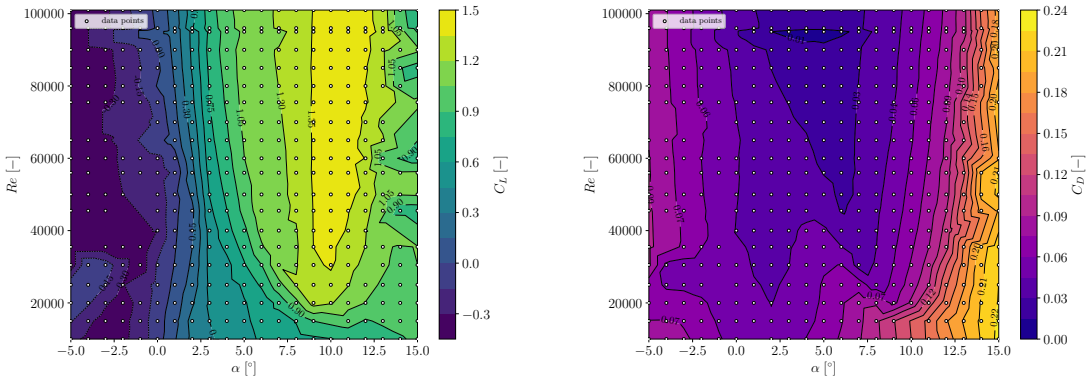

(a) $C_L(\alpha, Re)$. **(b)** $C_D(\alpha, Re)$.

Figure 4.11: Airfoil coefficients maps for the a0.65, by XFOIL.

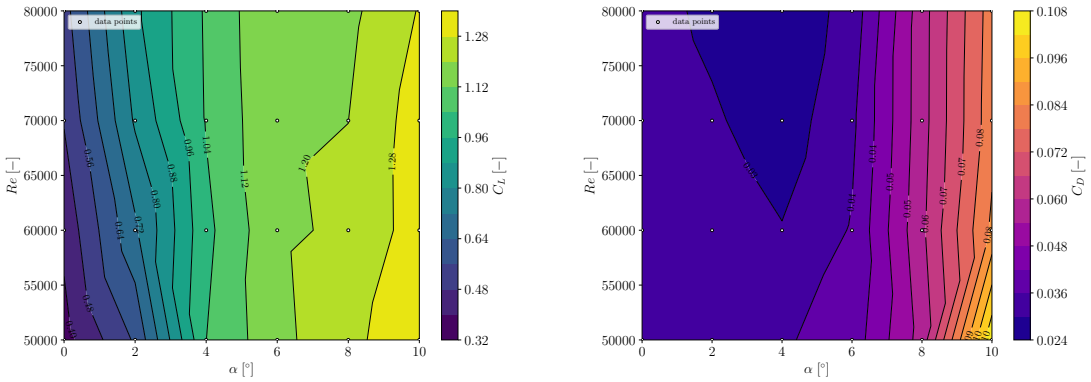

(a) $C_L(\alpha, Re)$. **(b)** $C_D(\alpha, Re)$.

Figure 4.12: Airfoil coefficients maps for the a0.65, by CFD.

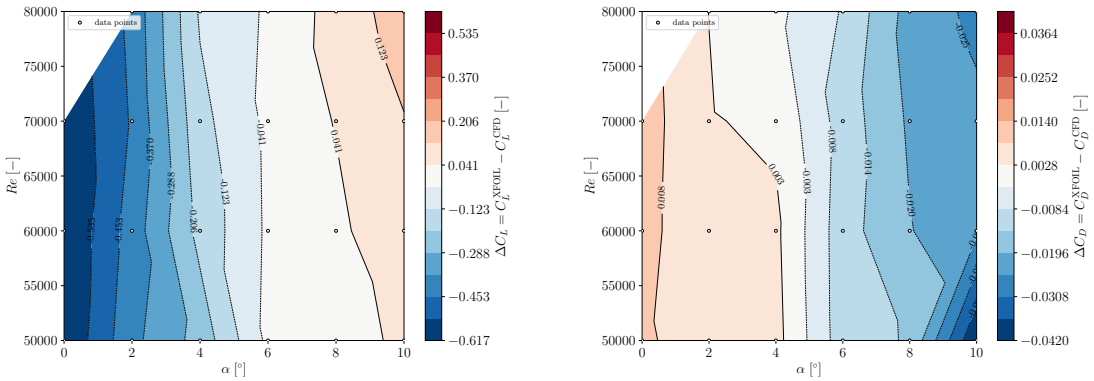

(a) $\Delta C_L(\alpha, Re)$. **(b)** $\Delta C_D(\alpha, Re)$.

Figure 4.13: Absolute difference in airfoil coefficients between CFD and XFOIL.

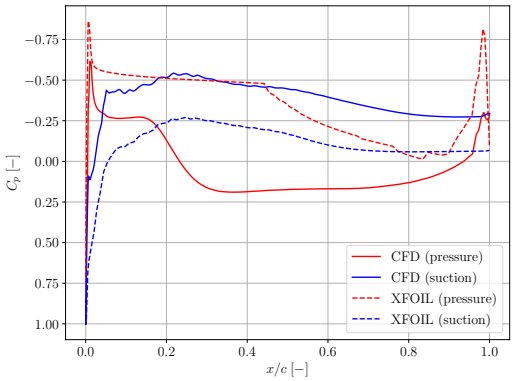

(a) $\alpha = 0°$.

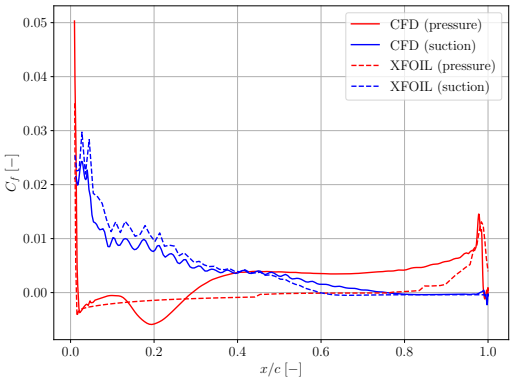

(b) $\alpha = 0°$.

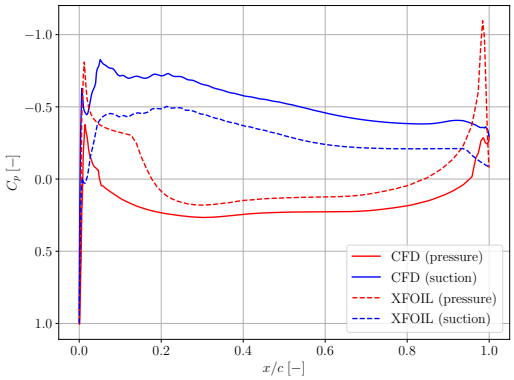

(c) $\alpha = 2°$.

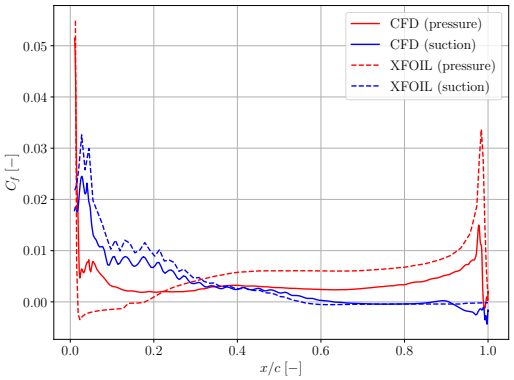

(d) $\alpha = 2°$.

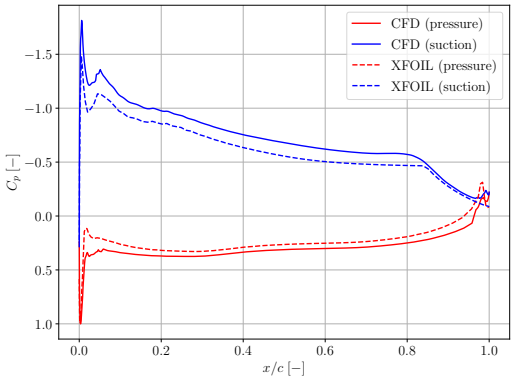

(e) $\alpha = 4°$.

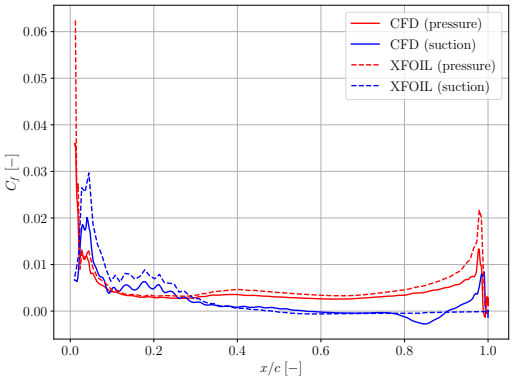

(f) $\alpha = 4°$.

Figure 4.14: Pressure and friction coefficients comparison at $Re = 60{,}000$, CFD against XFOIL.

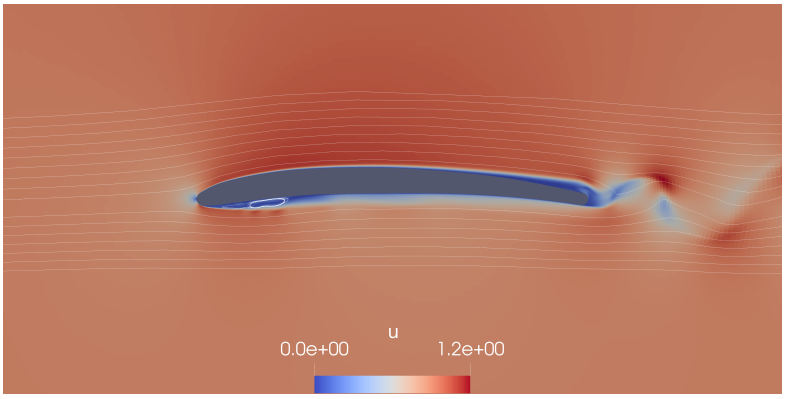

(a) $\alpha = 0°$.

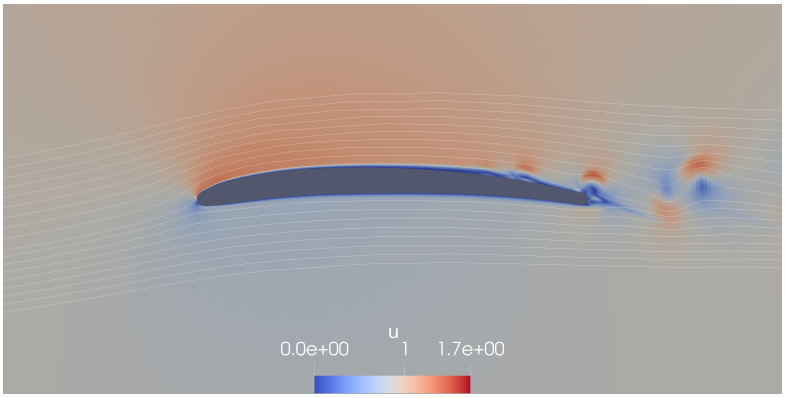

(b) $\alpha = 6°$.

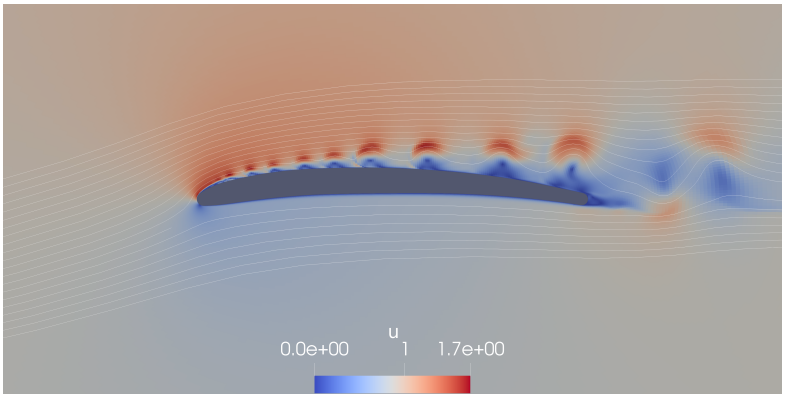

(c) $\alpha = 8°$.

Figure 4.15: Velocity magnitude snapshots.

Chapter 5

Propeller performance results

The complete BEMT methodology, employing aerodynamic polar data computed with either XFOIL or CFD as described in the previous chapter, is compared to experimental measurements on the T1045-II propeller, which are considered sufficiently accurate to serve as a reliable benchmark. The experiments were conducted at the von Kármán Institute for Fluid Dynamics (VKI) in hover conditions, in a controlled indoor environment, varying the rotational speed across a range of RPM values. Hover is the most representative operating condition for small quadrotor platforms and was therefore selected as the benchmark case. Fluid property data were not available for the experimental campaign; standard atmospheric conditions were assumed ($\rho = 1.225\,\mathrm{kg/m^3}$, $\mu = 1.81 \times 10^{-5}\,\mathrm{Pa\,s}$, at $15\,^{\circ}\mathrm{C}$ and $1\,\mathrm{atm}$). The BEMT used the a0.65 airfoil as the single representative geoemtry of the blades' sections. A freestream turbulence intensity of $Tu = 0.001$ was adopted as the baseline, reproducing a low-turbulence atmospheric environment consistent with the polar data presented in the previous chapter. An additional BEMT calculation was performed using XFOIL polar data computed at $Tu = 0.01$, in order to assess the sensitivity to turbulence intensity and to better represent the experimental setting: a propeller operating in a confined room eventually ingests its own wake, so the flow over the blade sections is expected to have higher turbulence intensity than in free-field conditions.

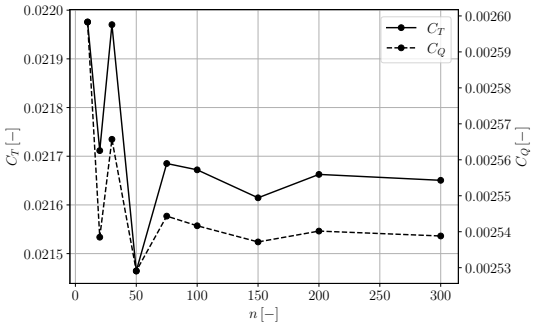

Figure 5.1: BEMT convergence analysis.

A radial discretization convergence analysis on the BEMT with polar maps from XFOIL was carried out to adequately represent the propeller. The thrust and torque coefficients C_T and C_Q were computed as a function of the number of equispaced radial annuli n used to discretize the blade, with results shown in Fig. 5.1. Both coefficients stabilize after approximately $n = 75$ annuli; $n = 100$ was selected and applied to all calculations presented in this chapter.

The thrust and torque predictions of the BEMT are compared against experimental data in Fig. 5.2. Regarding thrust (Fig. 5.2a), both BEMT with XFOIL and BEMT with CFD show very good agreement with the experimental measurements across the full RPM range, demonstrating that the aerodynamic polar data, regardless of the method used to compute it, is sufficient to accurately capture the thrust-generating behaviour of the propeller. The torque prediction (Fig. 5.2b) presents a substantially different picture: the experimental torque is considerably higher than all BEMT predictions, with a consistent discrepancy of approximately 30% relative to the CFD-based calculations across the whole range. This suggests that the BEMT method coupled with numerical airfoil data, while capturing the dominant lift-driven thrust contribution accurately, underestimates the total drag on the blades, which is more sensitive to the limitations of both the aerodynamic model and the two-dimensional strip-theory assumption. Finally, the sensitivity of the BEMT results to the turbulence intensity of the polar data is found to be small: the XFOIL calculations at $Tu = 0.001$ and $Tu = 0.01$ yield very similar predictions, and their difference is in fact smaller than the discrepancy between the XFOIL and CFD results at the same $Tu = 0.001$. The choice of polar data source

has thus a larger influence on the outcome than the assumed turbulence level.

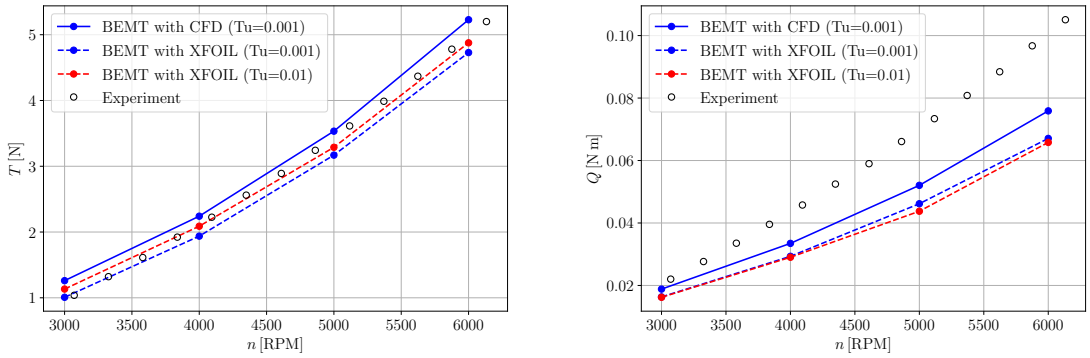

(a) $T(n)$. **(b)** $Q(n)$.

Figure 5.2: Thrust and torque comparison, BEMT against experiment.

Figures 5.3–5.5 report the radial distributions of the main local BEMT variables at 6000 RPM, namely the incremental thrust and torque coefficients $\mathrm{d}C_T/\mathrm{d}(r/R)$ and $\mathrm{d}C_Q/\mathrm{d}(r/R)$ as defined in Eq. **??**, **??**, the airfoil axial and tangential force coefficients c_z and c_θ, and the local angle of attack α and Reynolds number Re which together define the operative point on the two-dimensional aerodynamic coefficient maps.

The incremental thrust and torque distributions (Fig. 5.3) exhibit the characteristic propeller loading shape: both quantities are negligible near the hub and rise steeply in the outer portion of the blade, with the bulk of thrust and torque production concentrated in the range $r_{nd} \approx 0.5$–0.9, consistent with the expectations. The agreement between the CFD-based and XFOIL-based BEMT calculations is generally good across the span, with some discrepancies appearing beyond $r_{nd} \approx 0.7$.

The same pattern is reflected in the c_z and c_θ distributions (Fig. 5.4), which measure how the polar data gets translated into local propeller sections. Although the two polar maps had big discrepancies, the local coefficients instead show close agreement between the two methods over most of the blade span, with the same divergence in the outer region.

The angle of attack distribution (Fig. 5.5a) finally reveals the reason of this good agreement despite bad agreement in the polar maps: the blade operates at

approximately $\alpha \approx 5°$ over the range $0.4 \lesssim r_{nd} \lesssim 0.7$, where the two methods agree closely, while beyond $r_{nd} \approx 0.7$, the angle of attack decreases toward lower values. It is precisely in this outer region that the two calculations diverge: the outer blade stations, operating at reduced incidence, fall squarely in the regime where the two polar datasets disagree most, explaining the discrepancy in the local load distributions at $r_{nd} > 0.7$ and its propagation into the global torque prediction. The Reynolds number distribution (Fig. 5.5b) is virtually identical between the two calculations, probably due to how the solver iterates on this variable.

5.0.1 Final remarks

The results of this chapter present an overall encouraging picture for the BEMT methodology, while also highlighting its limitations in the low-Reynolds propeller regime. Thrust is predicted accurately by both polar sources across the full RPM range, and the close agreement between BEMT with CFD and BEMT with XFOIL is itself a positive result, suggesting that the methodology is internally consistent and that the choice of polar source does not dramatically alter the outcome in the operative range considered. Together, these results indicate that the proposed methodology has potential as a lightweight tool for propeller performance prediction, although more studies across different operative conditions and propeller geometries would be needed to build confidence in its generality. Torque prediction is less satisfactory, with a systematic underprediction of approximately 30% relative to the experimental measurements. This discrepancy should however be interpreted with caution, as the experimental torque data was not corrected for mechanical losses such as bearing friction and motor dissipation; these contributions could account for a non-negligible fraction of the measured value, and their removal might reduce the observed offset significantly. Among the aerodynamic sources of error, the two-dimensional strip-theory assumption inherent to BEMT is a plausible candidate: three-dimensional rotational effects on the blade boundary layer modify the effective drag in ways that no two-dimensional polar map can capture, regardless of how accurately it is computed. Fully resolving the torque discrepancy would likely require a three-dimensional rotor CFD simulation alongside experimentally corrected torque measurements, and is left for future work. Finally, it is worth noting that hover represents a favorable condition for BEMT: the axisymmetric,

steady inflow eliminates many of the unsteady and three-dimensional complications present in forward flight or maneuvering conditions, where the model's assumptions would be more severely tested.

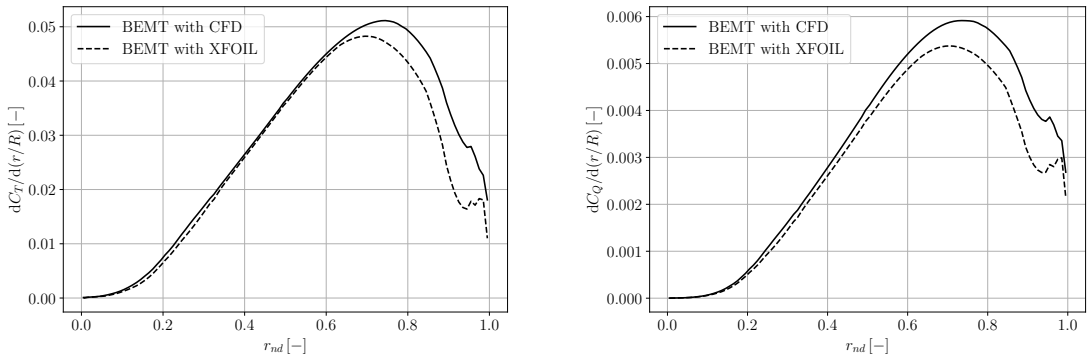

(a) Local thrust coefficient. **(b)** Local torque coefficient.

Figure 5.3: Local coefficients comparison, CFD against XFOIL.

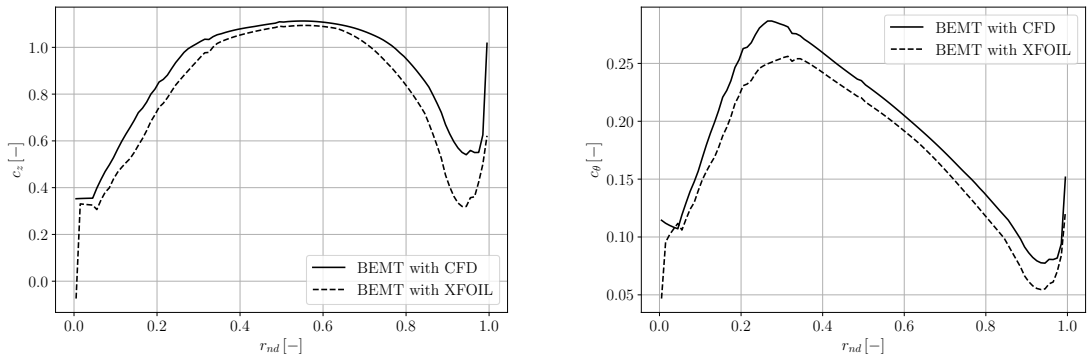

(a) Airfoil axial coefficient. **(b)** Airfoil tangential coefficient.

Figure 5.4: Airfoil coefficients comparison, CFD against XFOIL.

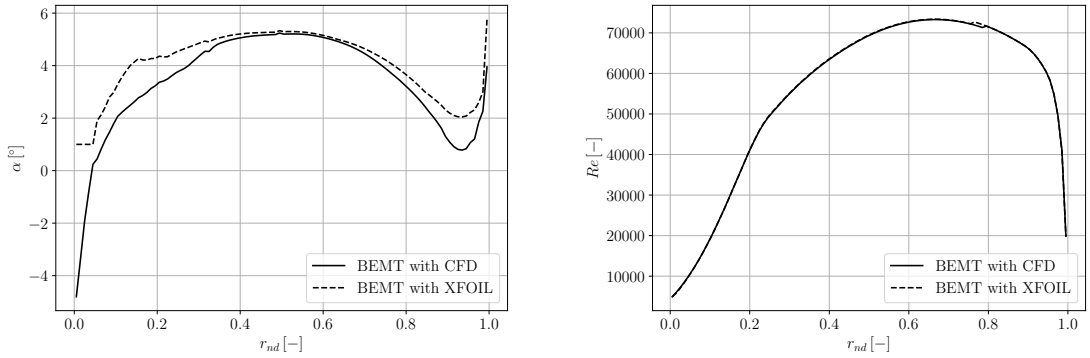

(a) Angle of attack.

(b) Reynolds number.

Figure 5.5: Airfoil operative parameters comparison, CFD against XFOIL.

Chapter 6

Conclusions

This work investigated the limitations of low-Reynolds airfoil characterization and their influence on propeller performance prediction within a BEMT-based framework, applied to the T1045-II small drone propeller. A modular numerical pipeline was developed and employed to generate aerodynamic polar maps using both XFOIL and CFD, which were then fed into a custom BEMT solver and compared against experimental measurements in hover conditions.

Airfoil characterization

On the Eppler 387 test case, both tools demonstrated satisfactory overall agreement with experimental data in the attached-flow regime, correctly capturing the lift curve slope and reproducing the qualitative features of the pressure distributions across the incidence range. Drag proved consistently more challenging for both methods, with errors growing at higher angles of attack as boundary layer dynamics intensified. The anomalous complete separation interval observed experimentally around $\alpha \approx 5°$–$7°$ could not be reliably reproduced by either solver, underscoring the extreme sensitivity of low-Reynolds flows to conditions that neither tool fully controls.

On the T1045-II representative section (a0.65), the comparison between XFOIL and CFD revealed a more significant disagreement than observed on the documented test case. A systematic offset in lift coefficient of approximately $\Delta C_L \approx 0.5$ was

found at low incidence across the full Reynolds number range considered, and while less pronounced, a meaningful discrepancy in drag was present at low incidence as well. The pressure and friction coefficient distributions suggest that both effects are primarily attributable to differing resolution of a laminar separation bubble on the pressure side near the trailing edge: XFOIL predicts a larger bubble that suppresses pressure recovery and reduces the suction peak, while CFD resolves a more compact structure, leading to higher integrated lift and somewhat different drag. As incidence increases and the pressure-side bubble retreats, the two methods converge on both coefficients, consistent with this interpretation. The finding highlights that on undocumented, geometrically complex sections, the reliability of XFOIL's predictions cannot be taken for granted even in nominally benign conditions, and that CFD likely provides the more physically faithful characterization.

Propeller performance

At the propeller level, the BEMT results present an encouraging and nuanced picture. Thrust predictions showed very good agreement with experimental measurements across the full RPM range, regardless of whether XFOIL or CFD polar data were used. This robustness suggests that the dominant lift-driven thrust contribution is captured faithfully by both polar sources, and that the BEMT modeling framework itself is sound. The close agreement between the XFOIL-based and CFD-based BEMT calculations further supports this conclusion: despite the meaningful differences in the polar maps, the propeller-level thrust prediction is not strongly sensitive to the choice of aerodynamic tool.

Torque prediction told a different story, with all BEMT calculations underestimating the experimental values by approximately 30%. Mechanical losses in the experimental setup, although corrected for, likely account for part of this discrepancy. The polar data inaccuracies identified at the airfoil level — both in lift and drag at low incidence — may contribute a further share. However, since the blade operates predominantly at moderate angles of attack where the two polar datasets are in closer agreement, polar inaccuracy alone is unlikely to fully explain the gap. Three-dimensional rotational effects on the blade boundary layer, which modify the

effective sectional drag beyond what two-dimensional strip theory can capture, and limitations of the simplified momentum wake model in representing the angular momentum of the swirling hover wake, are considered the most plausible remaining contributors. The torque underprediction therefore likely reflects a combination of all these sources, and its full resolution would require dedicated investigation beyond the scope of this work.

Methodological contribution and broader outlook

Beyond the specific results, the modular numerical pipeline developed in this work proved to be a flexible and effective tool for the analysis. The clean separation between geometry extraction, polar generation, and BEMT integration allowed different aerodynamic solvers to be compared systematically without altering the core propeller modeling logic, and the high degree of automation significantly reduced the manual effort required for what is inherently a large-scale parametric study. The framework is well positioned to serve as the basis for a more comprehensive BEMT implementation. Natural extensions include the incorporation of additional corrective models — such as wake contraction, swirl correction, and dynamic inflow models — which would improve prediction fidelity across a broader range of operating conditions, as well as the use of multiple blade sections distributed along the span to better capture the geometric variation of real propeller blades.

A promising direction for improving polar map quality is the combination of XFOIL and CFD data into a single unified dataset, exploiting the complementary strengths of the two approaches: XFOIL's ability to cover wide ranges of Reynolds number and angle of attack at negligible cost, and CFD's higher physical fidelity in the regions where the two methods diverge most. Such a strategy could draw on data fusion or multi-fidelity modeling techniques to produce polar maps that are both accurate and computationally affordable. Looking further ahead, higher-fidelity approaches beyond two-dimensional URANS — such as wall-resolved large eddy simulation or direct numerical simulation on representative sections — could provide ground truth data to anchor and validate lower-fidelity predictions, progressively building a more reliable characterization of low-Reynolds propeller aerodynamics.

The present work represents a necessary step toward reliable low-cost propeller

modeling for small drone platforms. The steady-state, uniform inflow assumptions underlying the BEMT naturally limit its scope: capturing the unsteady aerodynamic effects introduced by wind gusts, rapid maneuvers, or turbulent inflow — all of which are relevant to the applications pursued at VKI — will require substantial further development of both the propeller model and the airfoil characterization methodology. Nevertheless, the results demonstrate that even in its current form, a BEMT-based framework built on numerically computed polar data can provide a useful and computationally efficient baseline, and a solid foundation on which more capable models can be built.

Bibliography

[1] S. A. H. Moshan; N. Q. H. Othman; Y. Li; M. H. Alsharif; M. A. Khan. «Unmanned aerial vehicles (UAVs): practical aspects, applications, open challenges, security issues, and future trends». In: *Intelligent Service Robotics* (2023) (cit. on pp. 2, 3).

[2] Fact.MR. *Quadcopter Market*. Tech. rep. 2022 (cit. on p. 2).

[3] R. Merkert; J. Bushell. «Managing the drone revolution: A systematic literature review into the current use of airborne drones and future strategic directions for their effective control». In: *Journal of Air transport* (2020) (cit. on p. 2).

[4] J. Winslow; H. Otsuka; B. Govindarajan; I. Chopra. «Basic Understanding of Airfoil Characteristics at Low Reynolds Number ($10^4 - 10^5$)». In: *Journal of Aircraft* (2018) (cit. on pp. 3, 6, 13–15).

[5] R. J. McGhee; B. S. Walker; B. F. Millard. *Experimental Results for the Eppler 387 Airfoil at Low Reynolds Numbers in the Langley Low-Turbulence Pressure Tunnel*. Tech. rep. NASA: Scientific and Technical Information Division, 1988 (cit. on pp. 3, 14, 49, 50).

[6] O. Bergmann; F. Götten; C. Braun; F. Janser. «Comparison and evaluation of blade element methods against RANS simulations and test data». In: *CEAS Aeronautical Journal* (2022) (cit. on p. 3).

[7] Fact sheet of ERC project: "REinforcement TWInning SysTems". At: `https://cordis.europa.eu/project/id/101165479` (cit. on p. 4).

[8] L. Schena; P. Marques; R. Poletti; S. Ahizi; J. Van den Berghe; M. A. Mendez. «Reinforcement Twinning: from digital twins to model-based reinforcement learning». In: *Journal of Computational Sciences* (2024) (cit. on p. 4).

[9] Ludwig Prandtl. «Über Flüssigkeitsbewegung bei sehr kleiner Reibung». In: *Verhandlungen des III. Internationalen Mathematiker-Kongresses, Heidelberg*. English translation: "On Fluid Motion with Very Small Friction". Leipzig, Germany: Teubner, 1904, pp. 484–491 (cit. on p. 6).

[10] H. Schlichting; K. Gersten. *Boundary-Layer Theory*. 9th. Springer-Verlag, 2017 (cit. on pp. 7–10).

[11] W. Tollmien. «Über die Entstehung der Turbulenz». In: *Nachrichten der Gesellschaft der Wissenschaften zu Göttingen* (1929) (cit. on p. 9).

[12] J. Pretsch. «Die Stabilität einer ebenen Laminarströmung bei Druckgefälle und Druckanstieg». In: *Jahrbuch der deutschen Luftfahrtforschung* 1 (1942) (cit. on p. 10).

[13] A. M. O. Smith and N. Gamberoni. «Transition, Pressure Gradient, and Stability Theory». In: *Douglas Aircraft Report ES-26388* (1956) (cit. on p. 10).

[14] J. L. van Ingen. «A Suggested Semi-Empirical Method for the Calculation of the Boundary Layer Transition Region». In: *Delft University of Technology Report VTH-74* (1956) (cit. on p. 10).

[15] L. M. Mack. «Transition and Laminar Instability». In: *NASA Contractor Report* 153203 (1977) (cit. on p. 10).

[16] T. von Kármán. «Über laminare und turbulente Reibung». In: *Zeitschrift für Angewandte Mathematik und Mechanik* 1 (1921) (cit. on p. 11).

[17] K. Pohlhausen. «Zur näherungsweisen Integration der Differentialgleichung der laminaren Reibungsschicht». In: *Zeitschrift für Angewandte Mathematik und Mechanik* 1 (1921) (cit. on p. 11).

[18] Image from: `https:/en.wikipedia.org/wiki/File:Natureof_laminar_separation_bubble_on_a_turbine_blade.gif#filelink` (cit. on p. 12).

[19] M. Jahanmiri. *Laminar Separation Bubble: Its Structure, Dynamics and Control*. Tech. rep. Chalmers University of Technology, 2011 (cit. on p. 13).

[20] W. J. F. Koning; E. A. Romander; H. V. Cummings; B. N. Perez Perez. «On Improved Understanding of Airfoil Performance Evaluation Methods at Low Reynolds Numbers». In: *Journal of Aircraft, Vol. 60, No. 3* (2023) (cit. on pp. 13, 14, 41).

[21] Y. Zhang; Z. Zhou; K. Wang; X. Li. «Aerodynamic Characteristics of Different Airfoils under Varied Turbulence Intensities at Low Reynolds Numbers». In: *Applied Sciences* (2020) (cit. on p. 14).

[22] M. Drela. «XFOIL: An Analysis and Design System for Low Reynolds Number Airfoils». In: Lecture Notes in Engineering (1989). Ed. by Thomas J. Mueller (cit. on p. 16).

[23] M. Drela. «Viscous-Inviscid Analysis of Transonic and Low Reynolds Number Airfoils». In: *AIAA Journal* (1987) (cit. on p. 16).

[24] XFOIL's documentation, at `https://v0xnihili.github.io/xfoil-docs` (cit. on p. 18).

[25] M. S. Selig; C. A. Lyon; P. Giguére; C. P. Ninham; J. J. Guglielmo. *Summary of Low-Speed Airfoil Data*. Tech. rep. University of fllinois at Urbana-Champaign, 1996 (cit. on p. 18).

[26] F. R. Menter. «Two-equation eddy-viscosity turbulence models for engineering applications». In: *AIAA journal* (1994) (cit. on pp. 19, 20).

[27] F. R. Menter; B. R. Langtry; SR Likki; Y. Suzen; P. Huang; S. Völker. «A correlation-based transition model using local variables: part I—model formulation». In: *Journal of turbomachinery* (2006) (cit. on p. 21).

[28] R. B. Langtry; F. R. Menter. «Correlation-based transition modeling for unstructured parallelized computational fluid dynamics codes». In: *AIAA journal* (2009) (cit. on pp. 21, 42).

[29] Hermann Glauert. *The Elements of Aerofoil and Airscrew Theory*. Cambridge, UK: Cambridge University Press, 1926 (cit. on pp. 23, 28).

[30] R. Tognaccini. *Aerodinamica dell'Ala Rotante*. Università degli Studi di Napoli Federico II, 2023 (cit. on pp. 23, 26).

[31] E. Branlard. *Wind Turbine Aerodynamics and Vorticity-Based Methods*. Springer, 2017 (cit. on pp. 25–27).

[32] A. Ning. «A Simple Solution Method for the Blade Element Momentum Equations with Guaranteed Convergence». In: *Wind Energy* (2014) (cit. on p. 28).

[33] OpenFOAM documentation, at `https://doc.cfd.direct/openfoam/user-guide-v13/index`. (cit. on p. 40).

[34] Documentation on OpenFOAM's kLowReWallFunction, at `https://www.openfoam.com/documentation/guides/v2112/doc/guide-bcs-wall-turbulence-kLowReWallFunction.html`. (cit. on p. 41).

[35] Documentation on OpenFOAM's omegaWallFunction, at `https://www.openfoam.com/documentation/guides/v1912/doc/guide-bcs-wall-turbulence-omegaWallFunction.html`. (cit. on p. 41).

[36] D. F. Volkers. *Preliminary Results of Windtunnel Measurements on Some Airfoil Sections at Reynolds Numbers Between* 0.6×10^5 *and* 5.0×10^5. Tech. rep. Technische Hogeschool Delft, Vliegtuigbouwkunde, 1977 (cit. on p. 49).

[37] D. Althaus. *Profilpolaren für den Modellflug: Windkanalmessungen an Profilen im kritischen Reynoldszahlbereich*. 1980 (cit. on p. 49).

[38] Y. Zhang; Z. Zhou; K. Wang; X. Li. «Aerodynamic Characteristics of Different Airfoils under Varied Turbulence Intensities at Low Reynolds Numbers». In: *Applied Sciences* (2020) (cit. on p. 49).