UNIVERSITÀ DEGLI STUDI DI PADOVA

Dipartimento di Ingegneria Industriale DII

Corso di Laurea Magistrale in Energy Engineering

Rotor Blade Optimization for Sub-200 kW Wind Turbines

Relatore:
Giorgio Pavesi

Co-Relatore:
Thomas Henrik Hansen

Studente con matricola:
Pietro Poletto
Student ID: 2129681

Anno Accademico 2025/2026

Contents

Abstract

This thesis presents a computational framework for the preliminary aerodynamic optimization of horizontal-axis wind turbine (HAWT) blades with power ratings below 200 kW. The optimization approach is divided into two stages: first, a Kriging surrogate model explores the design space of 4-digit NACA airfoils, parameterized through continuous variables, in order to identify an optimal initial profile; subsequently, a gradient-free optimization algorithm based on Class-Shape Transformation (CST) parametrization refines the airfoil geometry. The aerodynamic analysis couples XFOIL, used for the generation of two-dimensional airfoil polars, with the BEM solver CCBlade for rotor performance evaluation. A dual-airfoil strategy enables the sequential optimization of the root and tip sections, while the intermediate sections are obtained through interpolation of CST coefficients and aerodynamic polars. Chord and twist distributions are derived from Betz optimal relations, ensuring geometric and aerodynamic consistency among the candidate profiles. The framework, developed in Julia with an open-source perspective, is intended for applications in the sub-200 kW range and aims to provide a computationally efficient tool for the preliminary design of wind turbine rotors.

1 Introduction

When it comes to large wind turbines, manufacturers can afford to invest significant time and resources in blade design: high-fidelity CFD simulations, aeroelastic analyses, and complex optimization loops are standard practice. The situation is different for small and medium wind turbines below 1 MW. These machines, typically used for distributed generation or urban installations, are usually designed with much tighter budgets and shorter development cycles. As a consequence, their blades almost always use well-known airfoil families, chosen from published databases rather than optimized for the specific turbine. The profiles work, but they are rarely the best choice for the actual operating conditions of the machine.
This thesis tries to close this gap by building a fast, open-source tool for the preliminary optimization of rotor blades in the sub-200 kW range. The underlying concept is straightforward: instead of running expensive CFD campaigns, the framework uses a panel method for the aerodynamic analysis and a BEM solver, keeping the computational cost low enough to run a full optimization on a standard laptop. The optimization follows a two-stage approach. First, a Kriging surrogate model explores the NACA 4-digit design space to find a good starting airfoil. Then, a gradient-free optimizer refines the shape using a CST parametrization. The goal is to make this design process faster and more accessible without requiring heavy computational infrastructure.
The focus on sub-200 kW turbines comes from a practical limitation of the code: blade bending and aeroelastic deformation are not implemented. For longer blades these effects become important, but for the blade lengths typical of this power range, a rigid-blade BEM approach is accurate enough for preliminary design. Within this limit the framework is fully scalable, applicable from a 2 kW micro-turbine to a 150 kW machine by simply changing the input parameters.

The computational framework presented in this thesis was developed with the support of AI-assisted coding tools. Since my background is in energy engineering rather than computer science, and I had no prior experience with the Julia programming language, these tools were extremely useful for software implementation. Their role is analogous to that of a technical editor: the content and ideas originate from the author, while the tool helps translate them into correct and functional code. All engineering decisions were made independently, based on literature review, personal analysis, and discussions with the supervisors.

2 Wind turbine fundamentals

Before addressing the core topics of this thesis, it is necessary to establish a solid understanding of how wind turbines work and the physics governing their operation. This chapter breaks down the most relevant concepts needed to fully appreciate the methods and results presented in the following chapters.

2.1 Wind turbine components

Rather than following the conventional component by component description, the wind turbine architecture is here organised into two functional groups, *supporting structure* and *power conversion system*.

Figure 2.1: Main components of a horizontal-axis wind turbine (HAWT), including rotor, nacelle, tower and foundation.

2.1.1 Supporting structure

The supporting structure encompasses all components whose primary function is mechanical and structural: the foundation, the tower and the nacelle. Although these elements do not directly participate in the aerodynamic energy conversion, they define the environment in which the rotor operates and impose critical constraints on the overall turbine design.

Foundation. The foundation transfers all loads from the turbine to the ground or seabed. Onshore installations rely on reinforced concrete slab foundations or multi-pile arrangements, sized to keep the turbine stable under the worst load case assumed for the specific site [1]. Offshore turbines require more elaborate support structures: monopiles, large-diameter steel tubes driven into the seabed, are the dominant solution in shallow waters up to approximately 40 m depth. In deeper sites, jacket structures provide a stiffer and lighter alternative, while floating platforms, anchored by mooring lines, are the enabling technology for deep-water locations where fixed foundations become economically unfeasible.

Tower. The tower elevates the rotor where wind resources are stronger, more uniform and less affected by surface-level turbulence. It is typically a tubular steel structure, wider at the base to resist the large bending moments transmitted by wind and rotor loads [1]. Hub heights on modern onshore turbines range from 80 to 120 m, while offshore machines easily exceed 150 m. The tower must also be designed to avoid resonance: its natural frequency is tuned to fall outside the frequencies associated with rotor rotation and blade passing, which would otherwise cause fatigue-critical vibrations over the turbine lifetime. When all turbine dimensions are scaled proportionally and tip speed is maintained constant, the tower natural frequency varies inversely with rotor diameter, preserving the dynamic amplification factors.

Nacelle. The nacelle is the housing mounted at the top of the tower that encloses and protects the drivetrain and electrical generation equipment. It sits on a yaw bearing that allows it to rotate about the tower axis, keeping the rotor aligned with the wind direction. In big turbines this is achieved by a powered yaw drive consisting of an electric or hydraulic motor driving a pinion that engages with a gear ring on the tower, while in small machines the rotor is passively oriented by a tail vane. Internally, the main shaft connects the rotor hub to the drivetrain: in conventional configurations a multi-stage gearbox steps up the low rotor speed, typically between 5 and 20 rpm, to the higher speed required by the generator, while direct-drive designs eliminate the gearbox entirely and couple the shaft to a large-diameter low-speed generator. The nacelle also houses the mechanical brake, which provides an emergency stop capability independent of the pitch system, and the cooling and power electronics required to condition the electrical output before it is fed to the grid. The nacelle bedplate transfers the rotor loadings to the yaw bearing and provides mountings for the gearbox and generator [1]. Despite its structural simplicity, essentially a steel or fibreglass shell, the nacelle is one of the heaviest single components of the turbine, and its mass at hub height has a strong influence on the tower and foundation design.

2.1.2 Power conversion system

The conversion of wind kinetic energy into electricity involves two distinct stages: the aerodynamic extraction performed by the rotor, and the electromechanical conversion performed by the generator.

Generator. The generator converts the mechanical power delivered by the rotor shaft into electrical power. Historically, the dominant configuration relied on a multi-stage gearbox to step up the low rotor speed, typically between 5 and 20 rpm, to the higher speed required by a conventional induction generator. Fixed-speed wind turbines use induction rather than synchronous generators because the slip mechanism provides essential damping to the drive train against cyclic torque variations. However, the gearbox has long been identified as one of the most failure-prone components in the drivetrain, and its maintenance represents a significant fraction of the operational costs over the turbine lifetime, particularly in offshore installations where access is limited and expensive. For this reason, the industry has been progressively shifting towards direct-drive configurations, in which the rotor shaft is coupled directly to a large-diameter permanent magnet generator operating at low speed, eliminating the gearbox entirely. Since the power output of a rotating electrical machine scales as $P \propto D^2Ln$, where n is the rotational speed, D and L are respectively diameter and length of the generator. Reducing the rotational speed n requires a proportional increase in generator diameter, which is why direct-drive generators tend to have large diameters but limited axial length [1].

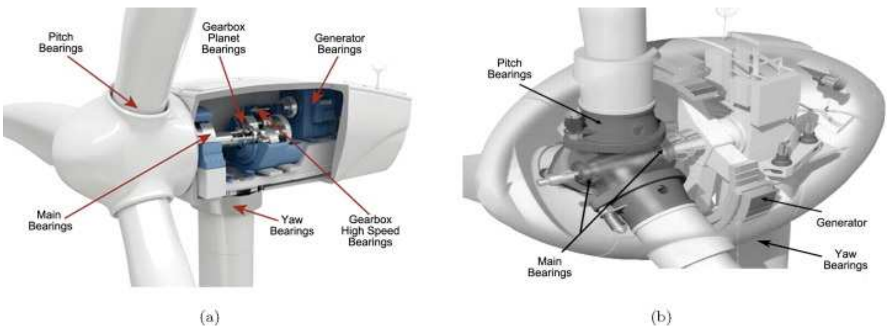

Figure 2.2: Comparison between geared and direct-drive drivetrain configurations.

This approach reduces mechanical complexity and improves reliability at the cost of a heavier and bulkier generator. Direct-drive technology is now the standard for large offshore turbines from manufacturers such as Siemens Gamesa and GE, while geared configurations remain common in the onshore segment where maintenance logistics are less constraining. While the choice of drivetrain architecture has a significant impact on nacelle layout and operational costs, it does not directly affect the aerodynamic design of the rotor and is therefore not analysed further in this work.

Rotor and blades. The rotor is the aerodynamic heart of the turbine and the component most directly relevant to this work. It consists of blades attached to a central hub, whose rotation sweeps a circular area, the rotor disc, through which the incoming wind passes. Modern utility-scale turbines adopt an upwind configuration, in which the rotor faces the wind in front of the tower; the alternative downwind arrangement allows passive yaw alignment but introduces periodic load fluctuations each time a blade passes through the tower wake, leading to higher vibration and lower energy yield.

The number of blades is a fundamental design choice that reflects a trade-off between aerodynamic performance, structural loads and cost. Fewer blades must rotate faster to compensate for the lower rotor solidity, so one and two-blade rotors operate at higher tip speed ratios ($\lambda \approx 15$ and $\lambda \approx 10$, respectively) compared to three-blade rotors ($\lambda = 7$–8). The three-blade configuration has become the industry standard for large machines, combining near-optimal aerodynamic efficiency with good performances flexibility and lower noise emission compared to two- and one-blade alternatives.

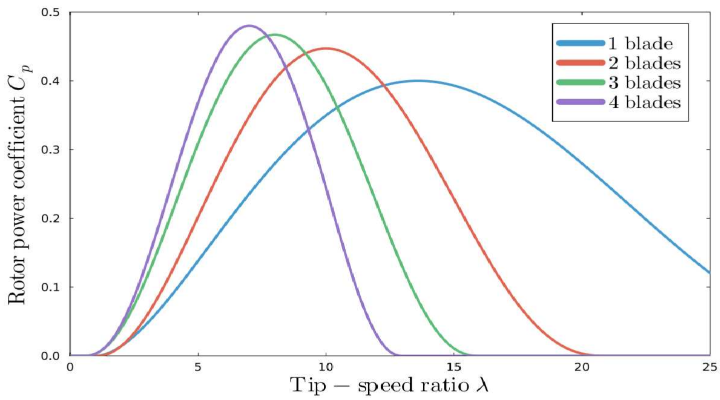

Figure 2.3: Rotor power coefficient C_p as a function of tip-speed ratio λ for different numbers of rotor blades. Higher blade counts achieve peak efficiency at lower tip-speed ratios but with a narrower operating range.

Each blade can be thought of as a slender wing that converts wind kinetic energy into shaft torque through lift. The incoming flow generates a lift force perpendicular to the local velocity vector; the tangential component of this force drives the rotation, while the axial component constitutes the thrust load on the rotor. The aerodynamic performance of each blade section is governed primarily by the lift-to-drag ratio C_l/C_d:a higher ratio generally translates into a higher rotor power coefficient C_P, making its maximisation one of the central objectives of airfoil optimisation. Unlike aircraft wings, wind turbine blades must operate over a wide range of angles of attack that extends well beyond stall, which can be broadly divided into three regimes: an attached-flow regime ($-15°$ to $+15°$), a high-lift stall-development regime ($+15°$ to $+30°$), and a fully stalled regime ($+30°$ to $+90°$). This wide operating range imposes specific requirements on airfoil design that differ significantly from those of conventional aircraft profiles. Early wind turbines used standard NACA asymmetric series; today the trend is towards purpose-designed profiles, such as the NREL S-series, optimised specifically for the aerodynamic and structural requirements of wind turbine blades.

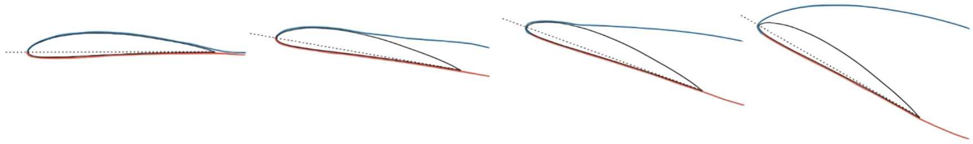

Figure 2.4: Evolution of flow separation with increasing angle of attack, from attached flow to fully stalled conditions.

Two geometric corrections are applied at the rotor to manage structural loads and operating conditions. A small cone angle of 2°–6° tilts the blades upwind relative to the rotor plane, aligning the resultant of centrifugal and aerodynamic forces with the blade longitudinal axis so that the dominant stress state is simple tension rather than bending. A tilt angle of approximately $6°$ inclines the entire rotor shaft upward, providing additional clearance between the blade tips and the tower without lengthening the main shaft.

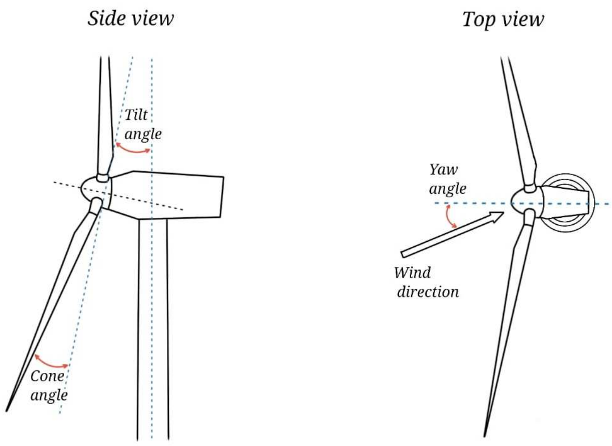

Figure 2.5: Definition of cone angle and shaft tilt angle in a horizontal-axis wind turbine rotor.

From a structural standpoint, the blades are the most mechanically demanding components of the turbine. A blade of 20–40 m length has a mass of 5–15 tonnes, and during rotation this weight generates a cyclic sequence of axial compression, bending and axial tension as the blade moves through its upper, horizontal and lower positions, respectively. This naturally periodic load pattern, combined with the irregular aerodynamic loads due to wind shear, tower shadow and atmospheric turbulence, makes fatigue the governing structural design criterion [1]. Modern blades are manufactured from glass-fibre reinforced epoxy composites, which offer a good combination of specific stiffness, fatigue resistance and cost. Carbon fibre composites provide superior specific strength and lower mass but remain significantly more expensive, and are used selectively in the most structurally critical regions.

2.2 One dimensional momentum theory

As in every form of energy conversion, extracting power from the wind inevitably involves losses. Unlike other conversion processes, however, wind energy is subject to a fundamental physical limitation that constrains the maximum achievable efficiency well below 100%, regardless of the machine design. It is possible to start with a simple 1-D model to study how a wind turbine extracts kinetic energy from the wind. In this case the turbine is modelled as an actuator disk acting as a drag device, under ideal flow conditions. The assumptions of this model are: the flow is steady, homogeneous, incompressible, and inviscid; the actuator disk is infinitely thin and permeable; and the static pressure far upstream and far downstream equals the undisturbed atmospheric pressure [2]. Before entering the demonstration, the term “a” must be defined, known as axial induction factor, a dimensionless number that represents the fractional decrease in wind speed at the disk. The velocities at the actuator disc and far downstream can be defined as function of the undisturbed wind speed and the induction factor:

$$U_d = U_\infty - U_\infty a = U_\infty(1-a) \tag{2.1}$$

$$U_w = U_\infty - 2U_\infty a = U_\infty(1-2a) \tag{2.2}$$

where U_d is the velocity at the disk, U_w is the wake velocity far downstream, and U_∞ is the undisturbed free-stream velocity. Since homogeneous, incompressible and steady-state flow conditions are considered, it is possible to apply the Bernoulli equation separately upstream and downstream of the disk (it cannot be applied across the disk itself, since work is extracted there):

$$p_\infty + \rho\frac{U_\infty^2}{2} = p_d^+ + \rho\frac{U_d^2}{2} \qquad p_w + \rho\frac{U_w^2}{2} = p_d^- + \rho\frac{U_d^2}{2} \tag{2.3}$$

Where the potential energy term ρgh is neglected, as is standard for horizontal-axis wind turbine analysis where elevation changes across the rotor are negligible.

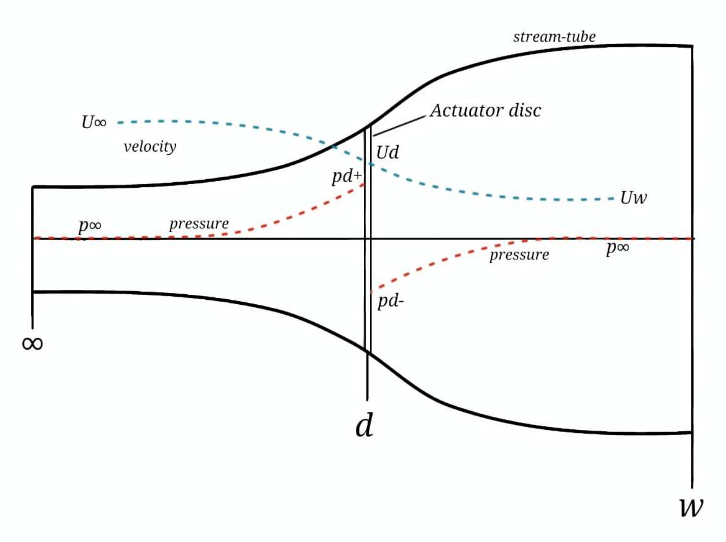

Figure 2.6: Actuator-disk representation of a wind turbine and corresponding streamtube expansion caused by the reduction in flow velocity.

The pressure difference as a function of the velocities can be expressed as:

$$(p_d^+ - p_d^-) = \frac{1}{2}\rho(U_\infty^2 - U_w^2) \tag{2.4}$$

Consequently the force applied in the streamwise direction is given by the pressure difference on the disk surface:

$$F = dp\, A_d = (p_d^+ - p_d^-)A_d \tag{2.5}$$

The downstream cross-section is larger than the disk while the upstream cross-section is smaller, because conservation of mass requires a constant mass flow rate through every section of the stream tube. The mass of air passing through a section per unit time is ρAU. Writing the mass flow equations for the three sections (upstream, disk, downstream) it follows that a decrease in velocity must be accompanied by an increase in cross-sectional area:

$$\dot{m} = \rho A_\infty U_\infty = \rho A_d U_d = \rho A_w U_w \tag{2.6}$$

The force equation can be written as:

$$F = \dot{m}(U_\infty - U_w) = 2\rho A_d U_\infty^2(1-a)a \tag{2.7}$$

Since the power can be evaluated as the force applied at the disc multiplied by the corresponding air velocity $P = F \cdot U_d$, substituting with equations (2.7) and (2.1):

$$P = 2\rho A_d U_\infty^3 (1-a)^2 a \tag{2.8}$$

To obtain the power coefficient based on the axial induction factor we need to divide the power by the available wind power,

$$P_{av} = \frac{1}{2}\rho U_\infty^3 A_d \tag{2.9}$$

$$C_p = 4a(1-a)^2 \tag{2.10}$$

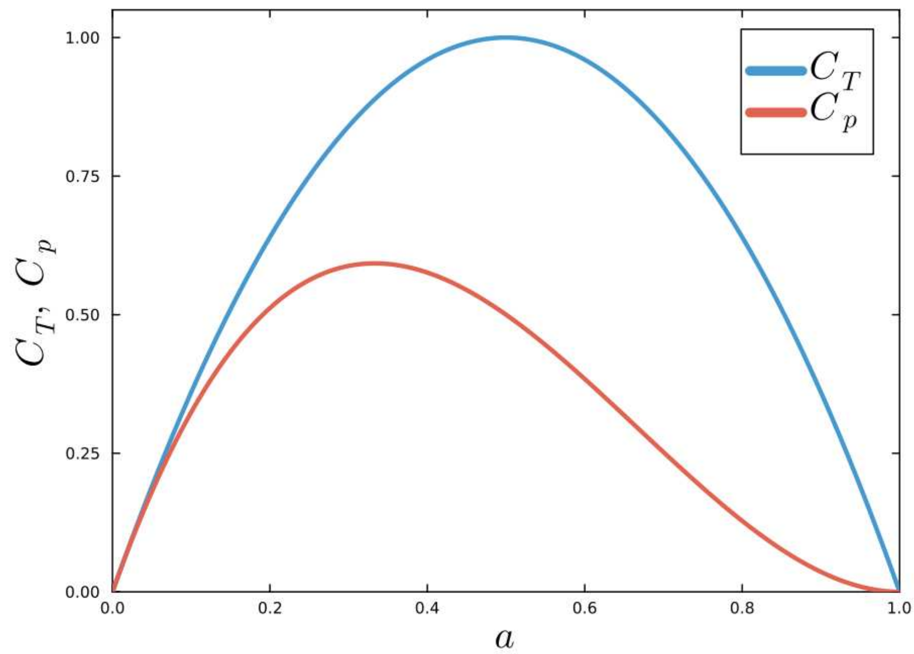

Figure 2.7: Power coefficient C_p and thrust coefficient C_t as functions of the axial induction factor a.

In the same way the thrust coefficient can be evaluated with the force on the actuator disk caused by the pressure drop divided by the dynamic pressure force on the disk area:

$$C_t = \frac{F}{\frac{1}{2}\rho U_\infty^2 A_d} = 4a(1-a) \tag{2.11}$$

The Betz limit is defined as the maximum value of C_p and it occurs when

$$\frac{dC_p}{da} = 4(1-a)^2 - 8a(1-a) = 0 \tag{2.12}$$

That leads to the solution of $a = \frac{1}{3}$, substituting in equation (2.10):

$$C_{p|max} = \frac{16}{27} = 0.593 \tag{2.13}$$

This limit is not a consequence of design choices, but of fundamental physical laws. The kinetic energy of the wind cannot be entirely captured because the air must retain enough velocity to flow away from the rotor; complete extraction would require zero wake velocity, which is physically impossible as it would imply an infinitely large wake cross-section. This specific condition occurs when the induction factor $a \geq \frac{1}{2}$, in this case the wake velocity $U_w = (1-2a)U_\infty$ becomes zero or negative, which is physically meaningless in the context of this model. It has been demonstrated experimentally that (2.11) is only valid for values of a smaller than 0.4 [2].

For higher values of a, the rotor enters the so-called *turbulent wake state*: the large velocity difference between the free stream and the wake destabilises the shear layer, generating turbulent eddies that transport momentum from the outer flow into the wake. In this regime, empirical corrections, such as Glauert's relation, are used instead of the momentum-theory expression.

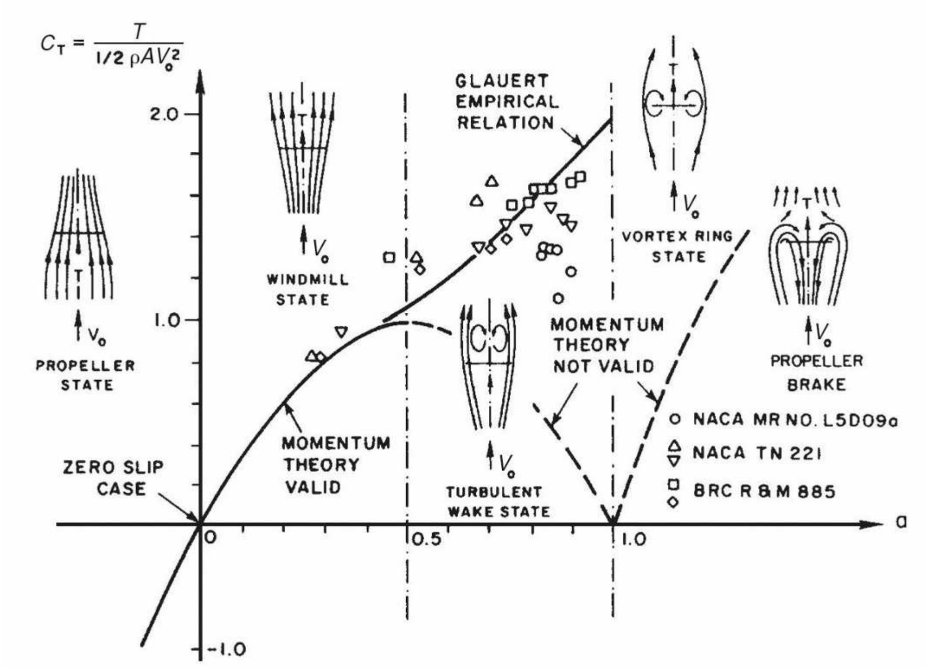

Figure 2.8: Thrust coefficient C_T as a function of the axial induction factor a. The momentum theory is valid for $a < 0.4$; beyond this value, the turbulent wake state is entered and empirical corrections are required. Source: [2].

The actuator disk model described above assumes that power is absorbed purely through a static pressure drop, without any rotational motion. The effect of rotation must however be taken into account in a real rotor. By Newton's third law, the torque exerted on the rotor blade produces an equal and opposite reaction torque on the air passing through the disk; resulting in a wake that will rotate in a direction opposite to that of the rotor. The imposition of a tangential component to air velocity leads to an increase in kinetic energy compensated by a drop in the static pressure downstream. As for the axial velocity, the change in tangential velocity is expressed in terms of the tangential induction factor a'. Specifically, a' is defined such that the tangential velocity imparted to the flow at the disk is $2\omega a' r$, where ω is the rotor angular velocity and r is the radial position.

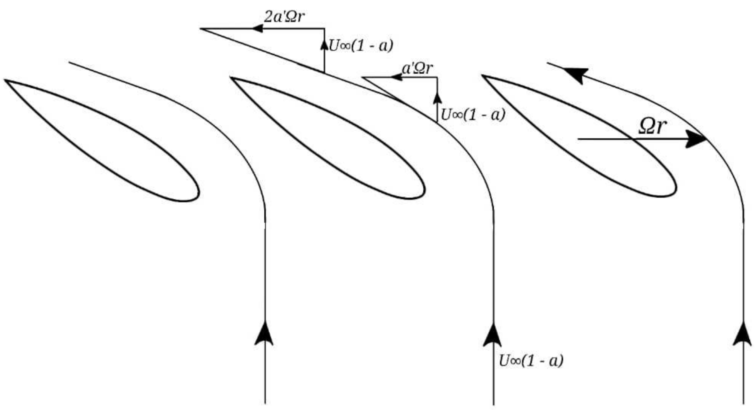

Figure 2.9: Velocity triangle at the rotor disk showing axial and tangential induction.

The torque imposed in the infinitesimal ring is equal to the rate of change of the angular momentum:

$$dM = \rho \, dA \, U_\infty (1-a) 2\omega a' r^2 \tag{2.14}$$

The extracted power is given by $dP = \omega \, dM$, and can also be expressed through the axial momentum balance:

$$dP = \omega \, dM = 2\rho \, dA \, U_\infty^3 a(1-a)^2 = 2\rho \, dA \, U_\infty (1-a) a' \omega^2 r^2 \tag{2.15}$$

The two rightmost expressions in (2.15) represent the same infinitesimal power, derived respectively from the axial momentum equation and from the angular momentum equation. Equating them yields the fundamental relationship between a and a'.

ωr is the local tangential velocity of the corresponding annulus, hence we can define the local speed ratio $x = \frac{\omega r}{U_\infty}$. At the tip when $r = R$ it is called tip speed ratio indicated by $\lambda = \frac{\omega R}{U_\infty}$.
From the previous equation (2.15):

$$U_\infty^2 a(1-a) = \omega^2 r^2 a'(1+a') \quad \rightarrow \quad a(1-a) = x^2 a'(1+a') \tag{2.16}$$

The area of the ring is $dA = 2\pi r\, dr$, thus from equation (2.14) the shaft power is evaluated as:

$$dP = \left[\frac{1}{2}\rho U_\infty^3 2\pi r\, dr\right] 4a'(1-a)x^2 \tag{2.17}$$

The first term between the square brackets represents the power flux through the annulus, the second part of the equation is the efficiency of the blade element,

$$\eta_r = 4a'(1-a)x^2 \tag{2.18}$$

By differentiating (2.18) and setting the result equal to zero:

$$(1-a)\frac{da'}{da} = a' \tag{2.19}$$

And differentiating (2.16) with respect to the axial induction factor:

$$x^2(1+2a')\frac{da'}{da} = 1-2a \tag{2.20}$$

By combining (2.19) and (2.20) the relation between the tangential and the axial induction factor, respectively a' and a, is obtained:

$$a' = \frac{1-3a}{4a-1} \tag{2.21}$$

Substituting equation (2.21) into equation (2.16), one finds that the value of a maximising the power coefficient is again $a = 1/3$, the same result obtained in the non-rotating analysis. At this optimum $a' = 0$, meaning that in the ideal limit no kinetic energy is lost to wake swirl. The total power is evaluated by integrating along the blade length from 0 to R:

$$P = 4\pi\rho\omega^2 U_\infty \int_0^R a'(1-a)\, r^3\, dr \tag{2.22}$$

And by using tip speed ratio $\lambda = \frac{\omega R}{U_\infty}$ and local speed ratio $x = \frac{\omega r}{U_\infty}$ it is possible to express the power coefficient in a non-dimensional form:

$$C_p = \frac{8}{\lambda^2}\int_0^\lambda a'(1-a)x^3\, dx \tag{2.23}$$

Substituting the optimal values of the induction factors confirms that the maximum C_p is $\frac{16}{27}$.

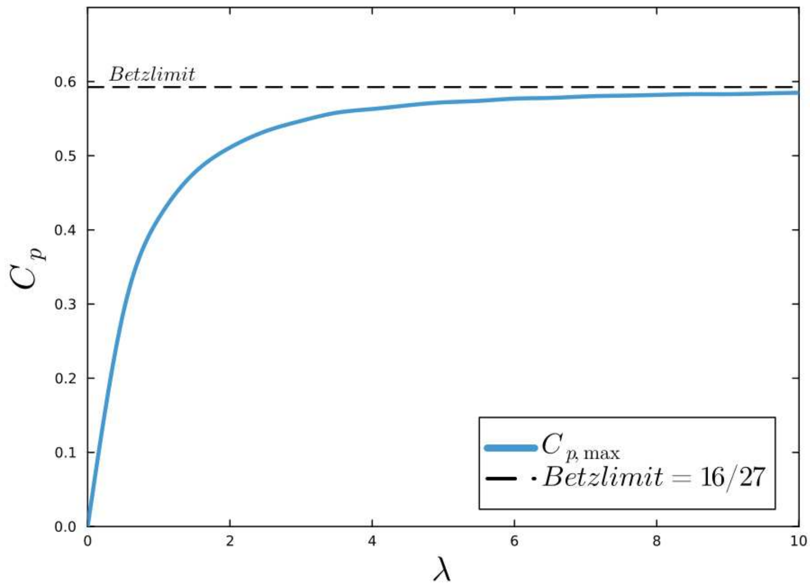

Figure 2.10: Maximum achievable power coefficient C_p as a function of the tip speed ratio λ, with and without wake rotation losses.

As the tip speed ratio λ increases, the tangential induction factor a' decreases, and consequently less kinetic energy is wasted in wake rotation. The theoretical optimum $C_p = 16/27$ is reached only in the limit $\lambda \to \infty$, i.e. when $a' \to 0$ everywhere along the blade. In practice, modern HAWTs operate at $\lambda \approx 7$–10, which is high enough that wake rotation losses are small and the achievable C_p is close to the Betz limit.

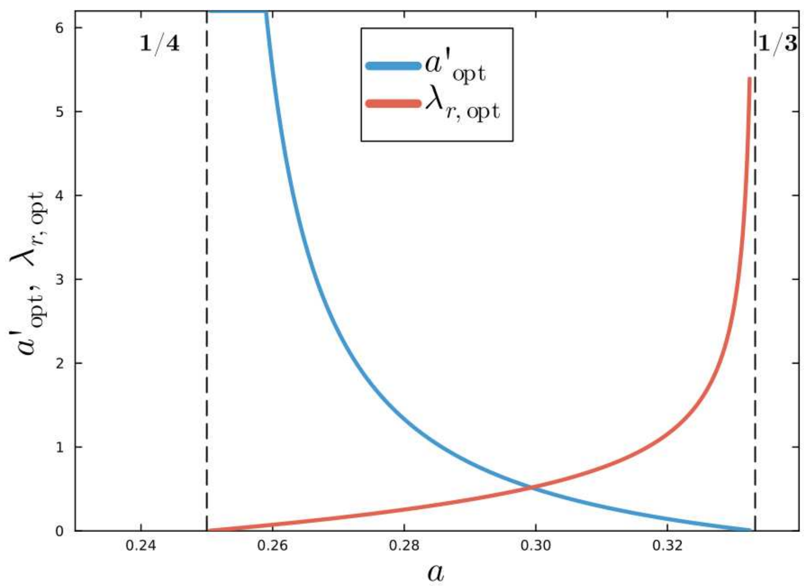

Figure 2.11: Optimal tangential induction factor a'_{opt} and local speed ratio $\lambda_{r,\text{opt}}$ as functions of the axial induction factor a. As $a \to 1/3$, both a' and λ_r diverge.

2.3 Airfoil aerodynamics

A truly two-dimensional aerodynamic flow can only exist around an airfoil of infinite span; it is therefore an idealisation. Nevertheless, wind turbine blades behave like long, slender wings moving through the air, where the streamwise velocity component is much larger than any spanwise component. Under these conditions, the flow over each cross-section is largely independent of what occurs at neighbouring sections, making the two-dimensional approximation a well-justified and widely adopted simplification. In the context of Blade Element Momentum (BEM) theory, this assumption is applied independently to each radial blade section, allowing the local aerodynamic forces to be determined from two-dimensional airfoil data alone.

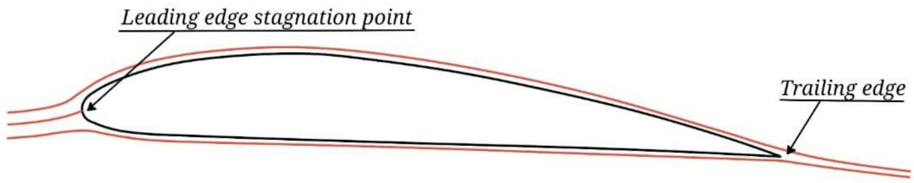

Figure 2.12: Two dimensional representation of the flow around an airfoil, showing the stagnation point and stagnation streamline.

When an airfoil is immersed in a moving fluid, the oncoming flow encounters the leading edge and divides at the stagnation point into two streams, one travelling along the upper surface and one along the lower surface, before rejoining at the trailing edge. Due to the asymmetric geometry of a typical airfoil, the two streams travel different paths lengths.

To quantify this effect, Bernoulli's equation can be applied along a streamline connecting the undisturbed free stream to a point on the upper surface (subscript u), and separately to a point on the lower surface (subscript l). Assuming steady, incompressible, and inviscid flow, one can write:

$$p_\infty + \frac{1}{2}\rho V_\infty^2 = p_u + \frac{1}{2}\rho v_u^2 = p_l + \frac{1}{2}\rho v_l^2 \tag{2.24}$$

Considering constant density ρ and neglecting the potential energy contribution, the asymmetric shape of the airfoil and its angle of attack cause the flow over the upper surface to accelerate relative to the lower surface. From (2.24), a higher velocity corresponds to a lower static pressure, resulting in a net pressure difference that generates an upward force. It should be noted that the common explanation based on "different path lengths" is an oversimplification. The actual mechanism involves the airfoil's curvature deflecting the flow downward (downwash), which by Newton's third law produces a reaction force (lift). The Bernoulli equation provides a consistent pressure-field description of this phenomenon, but the root cause is flow deflection, not path-length difference [3].

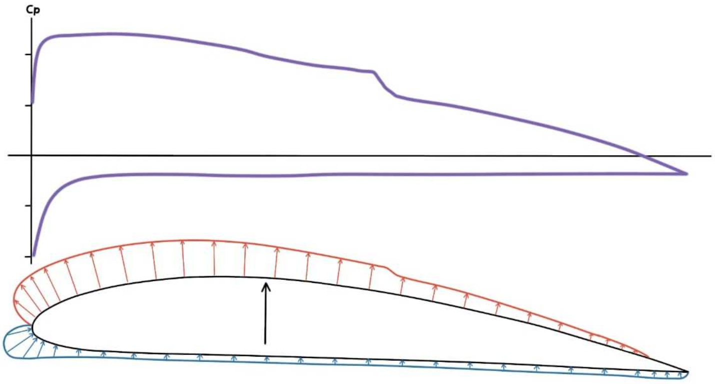

Figure 2.13: Qualitative pressure distribution around an airfoil responsible for lift generation.

The pressure distribution acting on the airfoil surface can be integrated to yield a single resultant force. This force is conveniently decomposed into two components: one perpendicular to the free-stream direction, known as lift, and one parallel to it, known as drag. In the case of an aircraft wing, it is the lift force that balances the weight of the vehicle, while the drag force opposes the motion and must be overcome by the engine thrust. In the context of a wind turbine, the lift force on each blade section is the primary driver of rotor torque, while drag represents a parasitic loss.

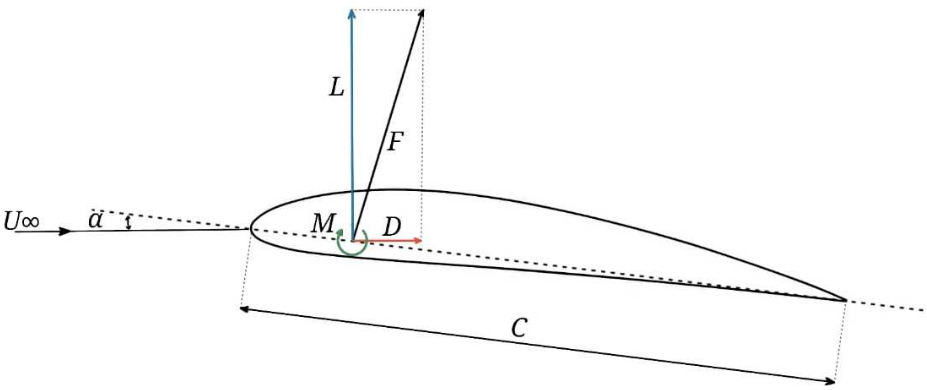

Figure 2.14: Lift and drag forces acting on an airfoil at an angle of attack α.

Lift and drag coefficients can be evaluated in relation of chord length c and air velocity V_∞:

$$C_l = \frac{l}{\frac{1}{2}\rho V_\infty^2 c} \tag{2.25}$$

$$C_d = \frac{d}{\frac{1}{2}\rho V_\infty^2 c} \tag{2.26}$$

where l and d are the lift and drag forces per unit span, respectively. To know where along the chord line those forces are applied, a moment coefficient C_m can be evaluated, if the coefficient is positive the airfoil tends to rotate nose-up (clockwise in the convention of Figure 2.16).

$$C_m = \frac{m}{\frac{1}{2}\rho V_\infty^2 c^2} \tag{2.27}$$

This coefficient becomes useful whenever we want to evaluate the pitching moment acting on the blade axis, as mentioned in section 5.3.2 about blade structural constraints.

The coefficients C_l, C_d and C_m depend on the angle of attack α between the chord line and the free stream, on the Reynolds number Re and on the Mach number Ma. In wind turbine applications where the speed is well below the speed of sound, lift and drag only depend on α and Re.

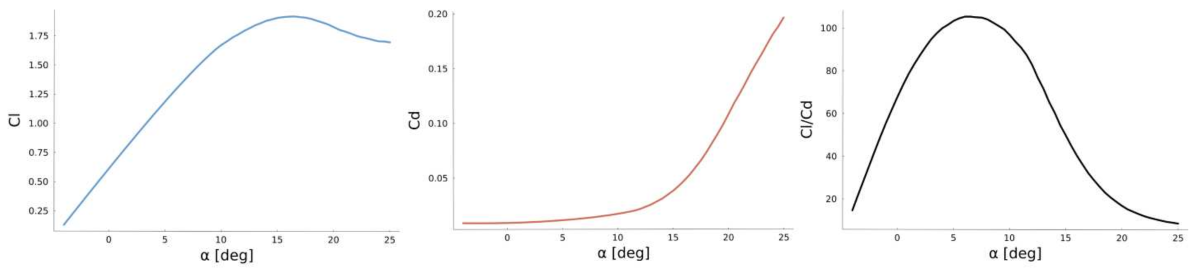

Figure 2.15: Representative airfoil polar showing lift, drag and lift-to-drag ratio as functions of angle of attack.

The Reynolds number $Re = \rho V c/\mu$ characterises the ratio of inertial to viscous forces and governs the nature of the boundary layer. Higher Reynolds numbers promote earlier transition from laminar to turbulent flow, which affects both the maximum lift and the stall behaviour of the airfoil. At high angles of attack, the adverse pressure gradient on the upper surface intensifies until the boundary layer can no longer remain attached. This phenomenon, known as *stall*, causes a sharp drop in lift and a rapid increase in drag. The stall angle and its characteristics (abrupt vs. gradual) depend on the airfoil geometry and the Reynolds number. Stall behaviour is of central importance in wind turbine design: it limits the maximum operating C_l and, in stall-regulated machines, is the primary mechanism for limiting power output above the rated wind speed [2], [3]. In the optimization process particular attention has been paid not only on the airfoil performances but also in the stall behaviour trying to promote airfoil with smooth stall rather than ”super performing” profiles but with very unstable lift trend.

2.4 Blade Element Momentum Theory

The aerodynamic analysis of the rotor relies on the Blade Element Momentum (BEM) method, a classical approach that couples one-dimensional momentum theory with local blade element aerodynamics. Although more accurate than pure one-dimensional momentum theory, BEM still relies on a simplified model in which three-dimensional effects are neglected; corrections can however be applied to improve the fidelity of the rotor analysis. Despite these limitations, the method remains computationally inexpensive and well suited for preliminary design optimisation loops where thousands of evaluations are required [4].
The rotor plane is divided into N independent annular elements of radial thickness $\mathrm{d}r$, each treated in isolation with no radial interaction between adjacent strips. By applying conservation of linear and angular momentum to a single annulus, the incremental thrust and torque are:

$$\mathrm{d}T = 4\pi r \rho U_\infty^2\, a(1-a)\,\mathrm{d}r, \tag{2.28}$$

$$\mathrm{d}M = 4\pi r^3 \rho U_\infty \omega\, a'(1-a)\,\mathrm{d}r, \tag{2.29}$$

where a is the axial induction factor and a' the tangential induction factor defined previously, quantifying how much the flow decelerates axially and accelerates tangentially due to the rotor.

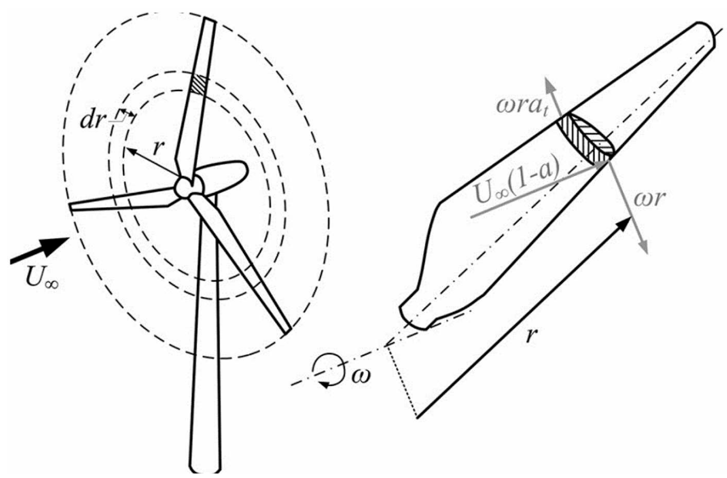

Figure 2.16: Visual representation of the annular sections along the bladespan. Source [4].

Each annular strip intersects the blade at a radial station r. The relative velocity seen by the section has an axial component $U_\infty(1-a)$ and a tangential component $\omega r(1+a')$, giving the inflow angle:

$$\tan\phi = \frac{U_\infty(1-a)}{\omega r(1+a')}. \tag{2.30}$$

The local angle of attack is $\alpha = \phi - \theta$, where θ is the local pitch angle. Given the sectional lift and drag coefficients $C_l(\alpha)$ and $C_d(\alpha)$ from airfoil polars, the normal and tangential force coefficients are:

$$C_n = C_l \cos\phi + C_d \sin\phi, \tag{2.31}$$

$$C_t = C_l \sin\phi - C_d \cos\phi, \tag{2.32}$$

and the incremental thrust and torque from blade element theory:

$$\mathrm{d}T = \tfrac{1}{2}\rho U_{\mathrm{rel}}^2\, B\, c\, C_n\, \mathrm{d}r, \tag{2.33}$$

$$\mathrm{d}M = \tfrac{1}{2}\rho U_{\mathrm{rel}}^2\, B\, c\, C_t\, r\, \mathrm{d}r, \tag{2.34}$$

where B is the number of blades, c the local chord, and $U_{\mathrm{rel}} = U_\infty(1-a)/\sin\phi$.
Equating the momentum and blade element expressions and introducing the local solidity $\sigma = cB/(2\pi r)$, the induction factors become:

$$a = \frac{1}{\dfrac{4\sin^2\phi}{\sigma C_n} + 1}, \tag{2.35}$$

$$a' = \frac{1}{\dfrac{4\sin\phi\cos\phi}{\sigma C_t} - 1}. \tag{2.36}$$

Since a and a' depend on the aerodynamic coefficients, which in turn depend on α (and therefore on the induction factors themselves), the system is solved iteratively: starting from $a = a' = 0$, the algorithm computes ϕ, looks up C_l and C_d, updates C_n and C_t, recalculates a and a', and repeats until convergence.
Two corrections are applied to improve the accuracy of the basic model.

Prandtl tip loss correction. The momentum equations above assume that the rotor loading is distributed uniformly over the full azimuth, which is equivalent to a rotor with an infinite number of blades. In practice, the flow can pass around the blade tips without being deflected, creating a region of reduced loading near the tip and root. This effect becomes stronger as the number of blades decreases or as the local inflow angle increases. Prandtl's correction models this by introducing a factor $F \in [0, 1]$ that multiplies the momentum-side thrust and torque [2]. The factor is computed as:

$$F = \frac{2}{\pi}\cos^{-1}\!\left(e^{-f}\right), \qquad f = \frac{B}{2}\,\frac{R-r}{r\sin\phi}, \tag{2.37}$$

where R is the rotor radius. At the tip ($r \to R$) the exponent f vanishes and $F \to 0$, effectively unloading the blade tip. Away from the tip F approaches unity and the correction becomes negligible. With this factor included, the corrected induction factors become:

$$a = \frac{1}{\dfrac{4F\sin^2\phi}{\sigma C_n} + 1}, \tag{2.38}$$

$$a' = \frac{1}{\dfrac{4F\sin\phi\cos\phi}{\sigma C_t} - 1}. \tag{2.39}$$

Glauert correction for high axial induction. When the axial induction factor exceeds approximately $a \approx 0.4$, the rotor enters the turbulent wake state and the linear momentum relation between thrust coefficient and induction factor is no longer valid [5]. The Glauert empirical correction, using the formulation by Spera (1994) [6] with a critical value $a_c \approx 0.2$, replaces the thrust coefficient with:

$$C_T = \begin{cases} 4a(1-a)F & a \le a_c, \\ 4\big(a_c^2 + (1-2a_c)\,a\big)F & a > a_c. \end{cases} \tag{2.40}$$

This ensures a smooth and physically realistic transition into the high-loading regime without the divergence that the standard momentum equation would produce.

Once converged at every station, the total thrust and power follow from spanwise integration:

$$T = B\int_{R_{\text{hub}}}^{R_{\text{tip}}} \mathrm{d}T, \tag{2.41}$$

$$P = \omega\, B\int_{R_{\text{hub}}}^{R_{\text{tip}}} \mathrm{d}M, \tag{2.42}$$

with the power and thrust coefficients defined as:

$$C_P = \frac{P}{\frac{1}{2}\rho V_0^3 \pi R^2}, \qquad C_T = \frac{T}{\frac{1}{2}\rho V_0^2 \pi R^2}. \tag{2.43}$$

2.4.1 Iterative solution procedure

The system of equations presented above cannot be solved in closed form, since the induction factors a and a' depend on the aerodynamic coefficients C_l and C_d, which in turn depend on the angle of attack α and therefore on the induction factors themselves. The standard approach is an iterative fixed-point procedure applied independently at each radial station [5]. The algorithm proceeds as follows:

1. Initialisation. Set $a = 0$ and $a' = 0$ at every radial station.
2. Inflow angle. Compute the flow angle ϕ from Equation (2.30) using the current values of a and a'.
3. Angle of attack. Evaluate $\alpha = \phi - \theta$, where θ is the local pitch angle (sum of the blade twist and the collective pitch setting).
4. Aerodynamic lookup. Recover the sectional lift and drag coefficients $C_l(\alpha)$ and $C_d(\alpha)$ from the airfoil polar table, using linear interpolation between the tabulated angles of attack.
5. Force coefficients. Compute the normal and tangential force coefficients C_n and C_t from Equations (2.31)–(2.32).
6. Induction update. Update a and a' from the corrected expressions (2.38)–(2.39), including the Prandtl tip-loss factor. If $a > a_c$, apply the Glauert correction to obtain the thrust coefficient and derive a accordingly.
7. Convergence check. If the change in both a and a' between successive iterations is smaller than a prescribed tolerance (typically 10^{-6}), the solution at the current station is considered converged. Otherwise, return to step 2 and repeat.

The convergence rate of this fixed-point scheme depends on the operating conditions and on the blade geometry. In the attached-flow regime, where the aerodynamic coefficients vary smoothly with the angle of attack, convergence is typically achieved within 10–20 iterations. Near stall, or at very high induction factors, the iteration may require relaxation or a larger number of steps to avoid oscillatory behaviour. In the present framework, convergence issues are largely avoided by the use of CCBlade, which reformulates the BEM system in terms of the single variable ϕ and solves it using a robust root-finding algorithm (Section 4.5), guaranteeing convergence under all operating conditions.

It is worth noting that the independence of the annular elements is a fundamental assumption of the BEM method. Each radial station is solved in isolation, with no communication of pressure or velocity information between adjacent strips. This assumption is justified for slender blades operating in attached flow, where the radial velocity component is small compared with the axial and tangential components. However, it breaks down in the presence of significant radial flow, such as occurs near the blade root due to the centrifugal pumping effect, or in deeply stalled conditions where the separated flow has a strong three-dimensional character.

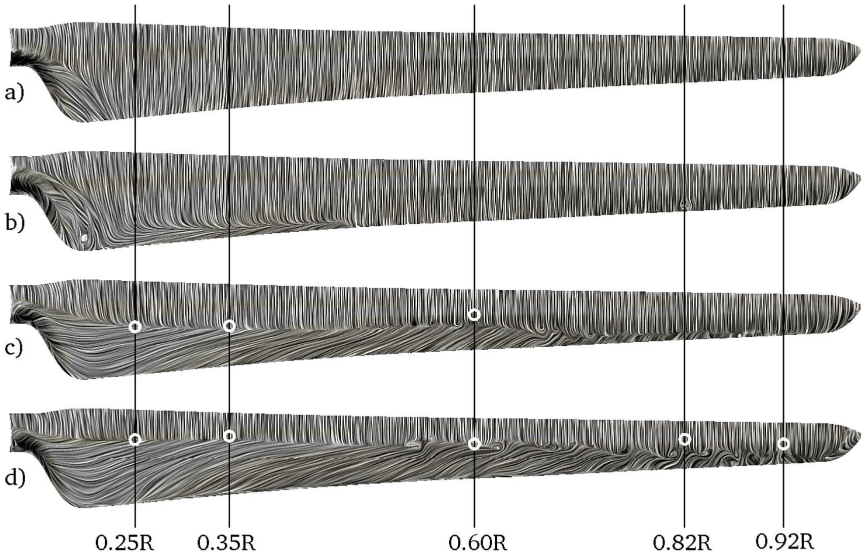

Figure 2.17: Radial flows and stall on the suction side at increasing wind speeds, from attached flow (a) to fully stalled conditions (d). From [7].

These limitations are inherent to the BEM formulation and are not addressed by the corrections described above; they represent the fundamental accuracy boundary of the method and motivate the restriction of the present framework to preliminary design rather than detailed performance prediction.

The oil-flow visualisation of Figure 2.17 shows the computed wall shear stress on the suction side of the mexico [8] rotor blade for four wind speeds. At low wind speeds the flow is nearly attached and two-dimensional, with streamlines oriented predominantly in the chordwise direction. As the wind speed increases and the blade enters stall, large regions of the suction side become dominated by strong radial flows directed outboard. This phenomenon, known as radial pumping, is caused by the centrifugal force acting on the separated flow and is directly linked to the rotational effects, stall delay and lift increase, that occur in the inboard blade region.

2.5 Wind resource analysis

The design of a wind turbine begins with a characterisation of the wind resource at the installation site. Unlike other energy sources, wind is intrinsically stochastic: its speed varies continuously over timescales ranging from seconds (turbulent gusts) to months (seasonal patterns). A statistical description of the wind regime is therefore essential for the design process.

2.5.1 Wind speed distribution

The probability distribution of wind speed at a given site is most commonly described by the Weibull distribution, a two-parameter continuous distribution whose probability density function is

$$f(v) = \frac{k}{c}\left(\frac{v}{c}\right)^{k-1} \exp\left[-\left(\frac{v}{c}\right)^{k}\right], \tag{2.44}$$

where $k > 0$ is the dimensionless *shape parameter* and $c > 0$ is the *scale parameter*.

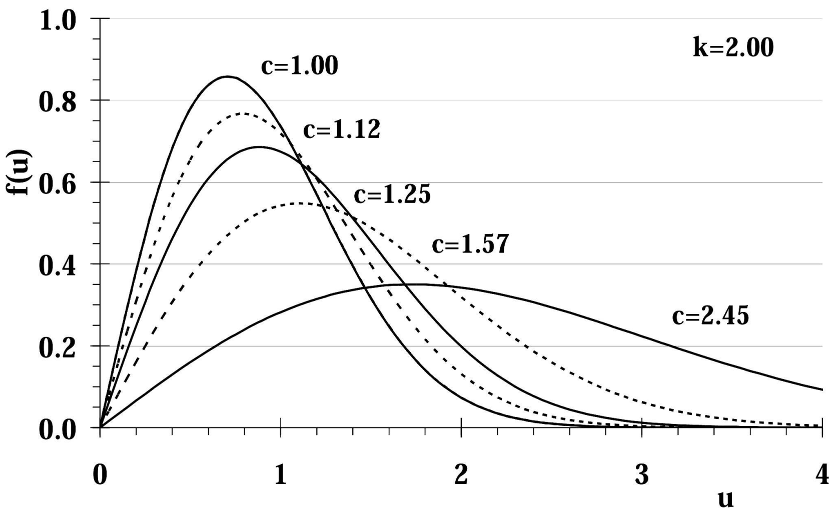

Figure 2.18: Effect of the scale parameter c on the Weibull probability density function at constant shape parameter k.

The shape parameter controls the spread of the distribution: $k \approx 1$ corresponds to a broad, exponential-like distribution (highly variable wind), while $k \approx 3$–4 produces a narrow, nearly symmetric bell curve (steady wind). Most continental sites fall in the range $k \approx 1.5$–2.5. The scale parameter is closely related to the mean wind speed and determines the horizontal position of the distribution along the velocity axis.

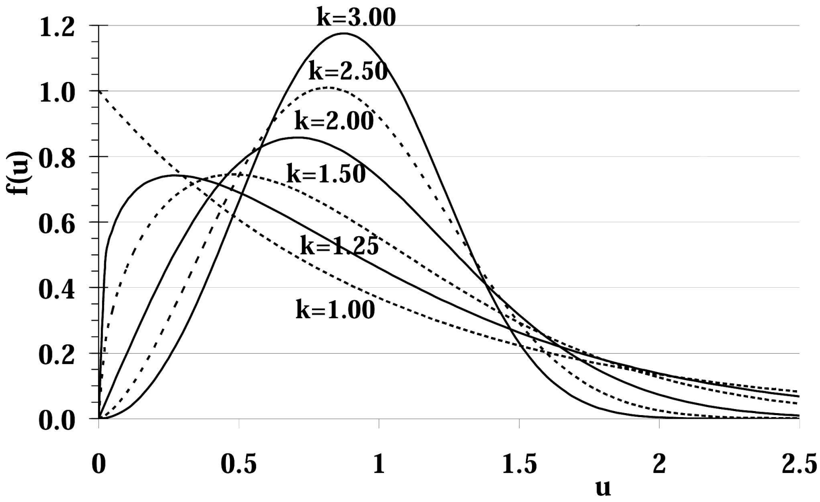

Figure 2.19: Effect of the shape parameter k on the Weibull probability density function at constant scale parameter c.

Two characteristic speeds derived from the Weibull parameters are frequently used in turbine design. The arithmetic mean wind speed is

$$\bar{v} = c\,\Gamma\left(1 + \frac{1}{k}\right), \tag{2.45}$$

where Γ denotes the Euler gamma function, and represents the long-term average velocity at the site. The most probable wind speed (the mode of the distribution) is

$$v_d = c\left(\frac{k-1}{k}\right)^{1/k}, \tag{2.46}$$

which corresponds to the wind speed that occurs most frequently. For the typical range of shape parameters encountered in practice, v_d is 10–30 % lower than the mean $\bar{v}$.

The Weibull parameters k and c are estimated from measured anemometric data, typically consisting of 10-minute averaged wind speed records spanning at least one full year. The most common estimation methods are maximum likelihood estimation (MLE) and the method of moments; the former is generally preferred for its statistical efficiency and is used in this work.

Once the wind distribution is characterised, several design speeds must be defined. The *rated wind speed* v_r is the speed at which the turbine reaches its nominal power output; above this speed the control system limits the power to protect the generator and drivetrain. A widely used engineering rule of thumb places the rated speed at

$$v_r \approx (1.3\text{–}1.5)\,\bar{v}, \tag{2.47}$$

balancing the desire for high energy capture (which favours a lower v_r) against the cost of over-sizing the drivetrain and electrical system (which favours a higher v_r). The *cut-in speed* v_{in} is the minimum wind speed at which the turbine starts producing useful power (typically 3–4 $\mathrm{m\,s^{-1}}$), and the *cut-out speed* v_{out} is the upper limit beyond which the turbine is shut down for structural safety (typically 25 $\mathrm{m\,s^{-1}}$ for small and medium turbines).

The choice of the aerodynamic design speed v^*, i.e. the wind speed at which the blade chord and twist are optimised, involves a trade-off between rotor size and energy yield. Designing at the most probable speed v_d maximises the hours per year at peak efficiency but requires a large rotor to reach the target rated power; designing at v_r minimises the rotor size but sacrifices efficiency at the most frequently occurring speeds. A compromise value between v_d and v_r is typically adopted in practice.

2.5.2 Rotor sizing

For a given design wind speed v^* and target power P_{target}, the required rotor radius follows directly from the definition of the available wind power:

$$P = \frac{1}{2}\,\rho\,v^{*3}\,\pi R^2\,C_p\,\eta, \tag{2.48}$$

where η accounts for drivetrain and electrical losses (typically $\eta \approx 0.85$–0.95). Solving for R:

$$R = \sqrt{\frac{2\,P_{\text{target}}}{\rho\,v^{*3}\,\pi\,C_p\,\eta}}. \tag{2.49}$$

This expression shows the strong sensitivity of rotor size to the design wind speed: since $R \propto v^{*-3/2}$, halving the design speed increases the required radius by a factor of $2\sqrt{2} \approx 2.8$.

For small wind turbines a common rule of thumb sets the hub height equal to the rotor diameter, $z_{\text{hub}} \approx 2R$ [1], ensuring adequate ground clearance while keeping the tower cost proportional to the rotor size.

2.5.3 Wind shear and vertical profile

The variation of wind speed with height above the ground is a well-documented atmospheric phenomenon driven by the friction between the moving air mass and the surface. Close to the ground the wind speed is reduced by surface roughness elements such as vegetation, buildings, and terrain features; with increasing altitude the retarding effects become smaller and the wind speed approaches the geostrophic value.

For engineering purposes the vertical profile is commonly described by either a logarithmic law or a power law. The logarithmic wind profile, derived from boundary-layer theory under the assumption of neutral atmospheric stability, expresses the mean wind speed at height z as

$$v(z) = v_{\text{ref}} \frac{\ln(z/z_0)}{\ln(z_{\text{ref}}/z_0)}, \tag{2.50}$$

where z_{ref} is the reference measurement height and z_0 is the aerodynamic roughness length, a parameter that characterises the surface type.

Type of terrain	Roughness length z_0 (m)
Cities, forests	0.7
Suburbs, wooded countryside	0.3
Villages, countryside with trees and hedges	0.1
Open farmland, few trees and buildings	0.03
Flat grassy plains	0.01
Flat desert, rough sea	0.001

Table 2.1: Typical Surface Roughness Lengths [1].

Typical values range from $z_0 \approx 0.001$ m for calm open water to $z_0 \approx 0.7$m for urban areas with dense buildings.

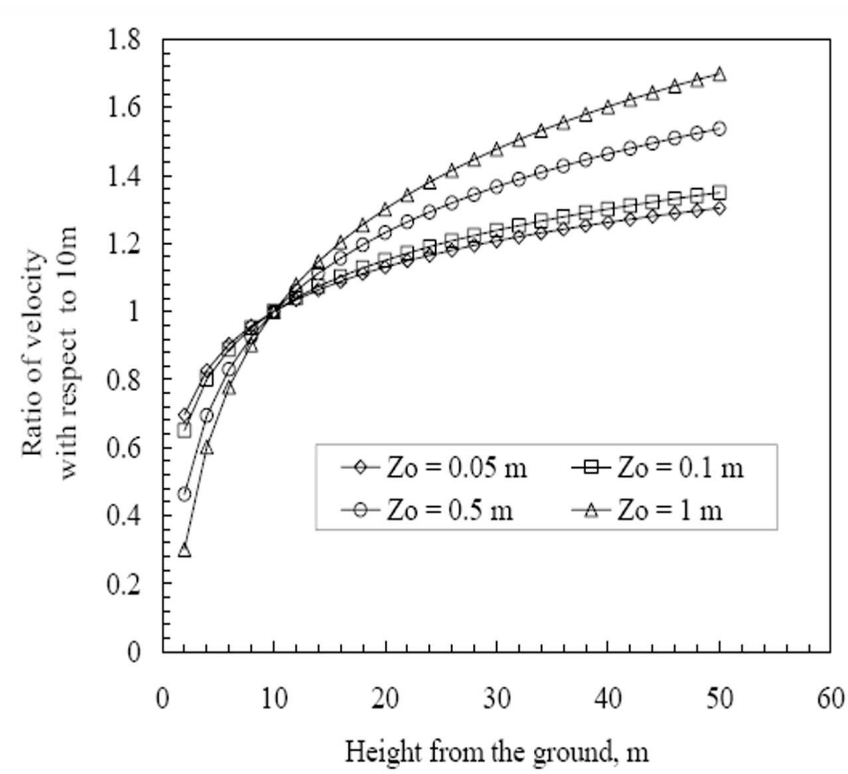

Figure 2.20: Ratio of wind velocity with respect to 10 m height as a function of height from the ground, for different surface roughness lengths z_0.

The power-law approximation, widely used in wind energy engineering for its algebraic simplicity, relates the wind speed at two heights through an empirical exponent α_s:

$$v(z) = v_{\text{ref}} \left(\frac{z}{z_{\text{ref}}} \right)^{\alpha_s}. \tag{2.51}$$

A value of $\alpha_s = 1/7 \approx 0.143$ is often cited as representative of neutral stability over flat, open terrain; measured values can change depending on the site.

The wind shear has two main consequences for rotor design. First, the wind speed at hub height differs from the speed measured at the standard anemometer height (typically 10 m), and a height correction must be applied to the measured data. Second, the wind speed varies across the rotor disk, creating an asymmetric loading on the blade as it rotates: the upper blade tip sees a higher velocity than the lower tip, producing a once-per-revolution cyclic load that is relevant for fatigue analysis. In a BEM framework the power-law formulation can be used to adjust the inflow velocity at each radial station according to its instantaneous vertical position relative to the hub.

2.5.4 Annual energy production

The annual energy production of a wind turbine is obtained by integrating the power curve $P(v)$ weighted by the wind speed probability density:

$$E_{\text{annual}} = 8760 \int_{v_{\text{in}}}^{v_{\text{out}}} P(v)\, f(v)\, \mathrm{d}v, \tag{2.52}$$

where $f(v)$ is the Weibull density function and the factor 8760 converts from hours per year. The capacity factor is defined as the ratio of actual energy produced to the energy that would be produced if the turbine operated continuously at rated power:

$$\text{CF} = \frac{E_{\text{annual}}}{P_{\text{rated}} \times 8760}. \tag{2.53}$$

Typical capacity factors for well-sited small wind turbines range from 0.20 to 0.40, depending on the wind resource and the turbine characteristics [1].

3 Numerical optimization fundamentals

The optimization algorithms adopted in this work are established methods available through open-source libraries. The objective of this thesis is not the development of new optimization algorithms, but their application to a wind turbine blade optimization problem within a coherent engineering framework. For this reason, the focus is placed on understanding the principles, limitations, and practical configuration of the adopted methods rather than on their mathematical derivation or implementation details.

In an ideal scenario, numerical optimization (in this case applied to an airfoil) refers to a process in which, starting from a random initial configuration and a defined set of parameters, commonly called Degrees of Freedom (DoF), the combination of these parameters that maximizes a given objective is sought. From a mathematical perspective, the optimization problem can be formulated as the search for the best configuration within the space of all possible geometries. For every specific application and its associated objectives, there exists, in principle, an optimal airfoil that outperforms all the others. If the design space is imagined as a landscape containing all possible geometries, this optimal configuration would correspond to the highest peak. However, since optimization problems are often formulated as minimization problems for convenience, it can equivalently be described as the deepest valley in that landscape. The goal of the optimization algorithm is therefore to reach this deepest valley as efficiently as possible, while avoiding getting trapped in shallower local valleys that may prevent the exploration of better solutions located beyond them.
In some cases, however, rather than searching directly for the optimum, it can be advantageous to first build a representation of the entire solution space. Following the definition given in [9], a surrogate consists of a model capable of reaching the optimum much faster than the original function. Instead of using a proper optimization loop, it can sometimes be advantageous to implement a surrogate model. In particular, the main advantages are:

- When the problem is computationally expensive, in particular if the number of evaluations needed to build the surrogate model is less than that needed for direct optimization
- When noisy data must be handled
- If understanding of the solution space is also required

Since a surrogate model needs to build an initial solution space, a significant limitation is given by the number of inputs: the more the DoF, the more samples are needed to build the starting space.

3.1 Kriging surrogate model

A very commonly used method for the approximation of nonlinear problems is the Kriging model. Instead of "driving" towards the optimum like a classic gradient-based method, Kriging builds the entire map first, then looks for the optimum. Rather than trying to describe what the function is, it tries to describe how the function behaves. Most Kriging models are variations of the following formulation:

$$F(x) = \mu(x) + Z(x) \qquad Z(x) \sim \mathcal{N}(0, \sigma^2) \tag{3.1}$$

Where, if x represents the location, F is the study variable, μ is the mean value at the location and Z is the residual (difference between F and μ at x) that follows a normal distribution with mean zero and variance σ^2. In most cases the mean μ is considered constant, because $Z(x)$, if built properly, can be a powerful tool itself to follow the function. This version is called Ordinary Kriging:

$$F(x) = \mu + Z(x) \tag{3.2}$$

Before discussing how the Kriging model works, the kernel functions must be introduced: a family of functions that correlate the residual $Z(x)$ with its corresponding position x. If two points $x^{(i)}$ and $x^{(j)}$ are close to each other, $f(x^{(i)})$ and $f(x^{(j)})$ are also expected to have similar values.

$$K(x^{(i)}, x^{(j)}) = \text{corr}(Z(x^{(i)}), Z(x^{(j)})) \tag{3.3}$$

The most commonly used kernel is:

$$K(x^{(i)}, x^{(j)}) = \exp\left(-\sum_{l=1}^{n_d} \theta_l |x_l^{(i)} - x_l^{(j)}|^{p_l}\right) \tag{3.4}$$

- n_d is the number of dimensions (inputs)
- $\theta_l > 0$ tells how rapidly the correlation changes along dimension l
- $0 \leq p_l \leq 2$ is an indicator of the smoothness

From the equation it is noticeable that when $x^{(i)} = x^{(j)}$ the $|x^{(i)} - x^{(j)}| \longrightarrow 0$ and so $K = 1$. When the term θ_l increases, the function becomes more sensitive in that direction, leading to a

large decrease in the correlation for a small variation in that variable. As smoothness operator, when $p_l = 2$ gradual variations in the function are expected, and when p_l is close to 0 the function is likely to have irregularities and abrupt variations.
Kriging involves two main steps. First, the statistical model parameters $(\mu,\ \sigma,\ \theta_1, \ldots, \theta_n,\ p_1, \ldots, p_{n_d})$ are estimated using the data. The second step uses these parameters and the statistical model to make predictions.
The estimation of the statistical model parameters is based on the maximum likelihood method. Without going into the mathematical details of the derivation, the optimization process leads to the following analytical solutions for μ and σ^2:

$$\mu^* = \frac{e^\top K^{-1} f}{e^\top K^{-1} e} \tag{3.5}$$

$$\sigma^{*2} = \frac{(f - e\mu^*)^\top K^{-1} (f - e\mu^*)}{n_s} \tag{3.6}$$

where K is the correlation matrix constructed using the kernel function, f is the vector of observed values, and e is a vector of ones [9]. The parameters θ_l and p_l of the kernel function do not have an analytical solution and must be determined numerically by maximizing the concentrated likelihood function:

$$\ell(\theta, p) = -\frac{n_s}{2} \ln(\sigma^{*2}) - \frac{1}{2} \ln |K| \tag{3.7}$$

Once all the parameters have been estimated, the Kriging model is fully defined and can be used to make predictions. The mean value of the Kriging prediction is:

$$f_p = \mu^* + k^\top K^{-1} (f - e\mu^*) \tag{3.8}$$

While the mean square error is:

$$\sigma_p^2 = \sigma^{*2} \left[1 - k^\top K^{-1} k + \frac{(1 - k^\top K^{-1} e)^2}{e^\top K^{-1} e} \right] \tag{3.9}$$

For production-level surrogate-based optimization, the SMT Python library provides Kriging with analytical gradients, enabling efficient gradient-based optimization of the surrogate model [10] [11].

3.2 Sampling

A sampling method allows the selection of the points needed in the building of the surrogate model. This procedure can be crucial in the utilisation of the surrogate model; two main characteristics are required:

- Achieve the most uniform coverage of the sampling space
- Minimise the total number of points required

Instead of random sampling algorithms, Latin Hypercube Sampling (LHS) has been chosen for being more efficient and effective. Other methods based on low-discrepancy sequences also present good performance, but they are more useful in iterative optimization processes such as the Monte Carlo method. In Latin Hypercube Sampling, to ensure a constant density distribution all samples must be chosen beforehand. Each dimension is divided into equally spaced sections, with the constraint that exactly one point can exist in each section of each dimension. Moving to a bi-dimensional space (Latin Square) for visualization simplicity, given a square divided into rows and columns, only one point can exist in each row or column. Extending this to higher dimensions yields the so-called Latin Hypercube.

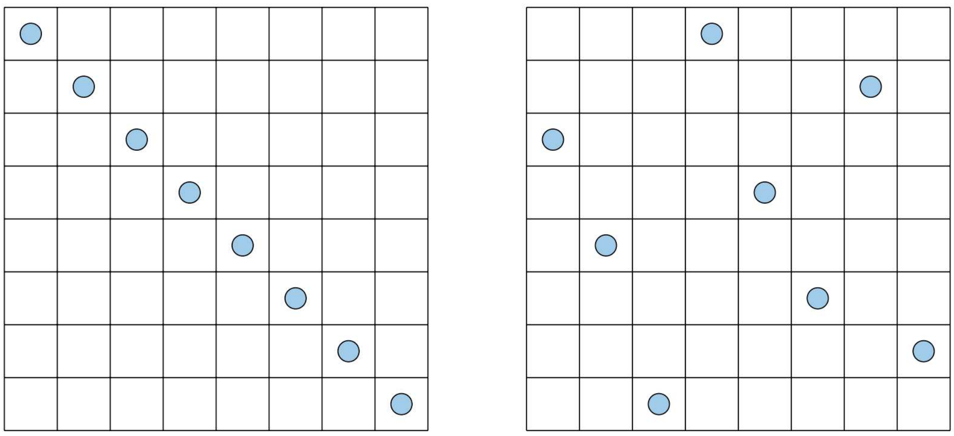

Figure 3.1: Two possible solutions of the LHS; the right one is more effective [9].

As shown in Figure 6.3, there are many ways to satisfy the LHS constraints, some more effective than others. In some ways the generation of the space is an optimization process itself, in which the distance between the points is maximised. Even if treating the LHS as an optimization problem leads to the most spread configuration, different approaches are used, less effective but more practical and simpler than the optimization problem itself [9]. Those methods are based on the probability function used to distribute the points along the intervals; the most common one exploits the uniform distribution in which every point has the same probability of being chosen.

This is particularly useful when an unknown region must be explored, spreading the points all over the search area. As an example, consider the study of different configurations of maximum camber (A) and camber position (B) of a NACA 4-digit airfoil, assuming 8 sampling points and chosen the intervals $A \in [2, 8]$ and $B \in [3, 7]$.
Each interval is divided into 8 equal sub-intervals and a random value is drawn from each one, as shown in Table 3.1.

Sample	A interval	B interval	(A_i, B_i)
1	[2.0, 2.8]	[4.5, 5.0]	(3.0, 4.8)
2	[2.8, 3.5]	[6.0, 6.5]	(3.5, 6.2)
3	[3.5, 4.3]	[3.0, 3.5]	(4.0, 3.2)
4	[4.3, 5.0]	[5.5, 6.0]	(5.0, 5.8)
5	[5.0, 5.8]	[4.0, 4.5]	(5.5, 4.3)
6	[5.8, 6.5]	[6.5, 7.0]	(6.5, 7.0)
7	[6.5, 7.3]	[3.5, 4.0]	(7.0, 3.7)
8	[7.3, 8.0]	[5.0, 5.5]	(8.0, 5.2)

Table 3.1: Latin Hypercube Sampling example with 8 points.

One random value is drawn from each sub-interval for both A and B. The values of one variable (e.g. B) are then randomly permuted, and each A_i is paired with the re-ordered B_i to form the initial sampling duplets.

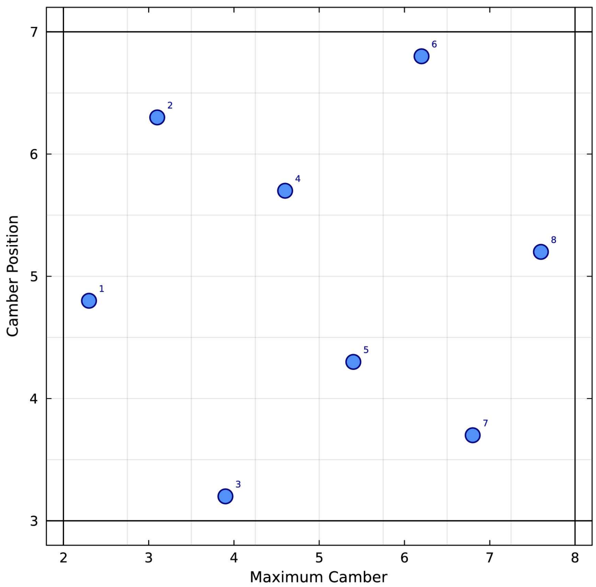

Figure 3.2: Example distribution.

3.3 Gradient-free optimization

When optimizing an airfoil based on CST coefficients, the number of degrees of freedom can increase very quickly. For high dimensional problems, the Kriging surrogate model encounters significant limitations: the number of operations required to minimize the prediction error increases exponentially with the number of variables, making the approach computationally impractical beyond a moderate number of parameters. A further constraint arises from the choice of flow solver: using XFOIL as the sole source of aerodynamic data prevents the use of gradient-based optimization algorithms, since the panel method does not provide gradient information and the function must therefore be treated as a black box. Although gradients can in principle be approximated via finite differences, this approach introduces inaccuracies and requires $N+1$ function evaluations per step, where N is the number of design variables. As reported in [9], the number of function evaluations required to construct the gradients scales with an exponent of approximately 2.52 with respect to the number of variables, compared to 0.37 for gradient-based methods with analytical gradients. For these reasons, a purely gradient-free approach is adopted: rather than relying on any derivative information, these algorithms guide the search using only function values. They are generally easier to implement, but require more evaluations as the dimension of the problem increases. Some gradient-free methods feature global search strategies that increase the likelihood of finding the global optimum; others perform only a local search and must therefore be combined with global initialization strategies to mitigate convergence to local minima.

Among local search methods, two main families can be identified based on how they use function evaluations. Direct search methods, such as the Nelder–Mead algorithm, Generalized Pattern Search (GPS), and Mesh Adaptive Direct Search (MADS); explore the design space by evaluating the objective function directly at a set of trial points, without constructing any internal model. The Nelder–Mead algorithm operates heuristically by manipulating a simplex of $N+1$ points through geometric operations (reflection, expansion, contraction, and shrinking); GPS and MADS instead rely on mathematical polling strategies along positive spanning directions. Model-based methods take a different approach: at each iteration, an analytic approximation of the objective function is built; referred to as a surrogate model, which is then used to guide the search. This model should be smooth, easy to evaluate, and accurate in the neighbourhood of the current point [9]. Trust-region algorithms and implicit filtering belong to this family; the trust-region approach can be considered gradient-free when the surrogate model is constructed using only function evaluations, without requiring gradient information. COBYLA belongs to this model-based family, sharing the characteristics of the trust-region row in Table 3.2: local search, mathematical rationale, surrogate-based function evaluation, and deterministic behaviour.

	Search		Algorithm		Function evaluation		Stochasticity	
	Local	Global	Mathematical	Heuristic	Direct	Surrogate	Deterministic	Stochastic
Nelder–Mead	•			•	•		•	
GPS		•	•		•		•	
MADS		•	•		•			•
Trust region	•		•			•	•	
Implicit filtering	•		•			•	•	
DIRECT		•	•		•		•	
MCS		•	•		•		•	
EGO		•	•			•	•	
Hit and run		•		•	•			•
Evolutionary		•		•	•			•

Table 3.2: Classification of gradient-free optimization methods, adapted from [9]. COBYLA [12] shares the characteristics of the Trust region row.

COBYLA (Constrained Optimization BY Linear Approximations) was originally developed by Powell [12] and operates by constructing linear polynomial approximations to both the objective and the constraint functions, using interpolation at the vertices of a simplex. A simplex in n dimensions is the convex hull of $n+1$ points in the design space, each representing a complete set of design variables [13].
At each iteration, the function values $F(\underline{x}_i)$, $i = 0, 1, \ldots, n$, evaluated at the simplex vertices are used to fit a unique linear polynomial $L(\underline{x})$. A trust region of radius $\Delta > 0$ is maintained around the best point found so far, $\underline{x}_0$. The next candidate point $\underline{\hat{x}}$ is obtained by minimizing $L(\underline{\hat{x}})$ subject to $\|\underline{\hat{x}} - \underline{x}_0\| \leq \Delta$, where the norm is Euclidean. In the unconstrained case this reduces to a steepest-descent step of length Δ along the gradient of the linear model:

$$\underline{\hat{x}} = \underline{x}_0 - \frac{\Delta}{\|\nabla L\|} \nabla L \tag{3.10}$$

After evaluating $F(\underline{\hat{x}})$, one vertex of the simplex is replaced by $\underline{\hat{x}}$; the choice of which vertex to drop is governed by the need to prevent the simplex from degenerating, i.e. from collapsing to a lower-dimensional object [13]. In the constrained version, linear approximations are built for each constraint function as well, and the trust region subproblem is solved subject to all linearised constraints simultaneously; constraints are therefore handled directly within the subproblem, without any penalty reformulation. The trust region radius Δ is regulated by a lower bound ρ, which is never increased; reductions occur only when the current value appears to be preventing further progress, and convergence is declared when ρ falls below a prescribed tolerance.

Despite its efficiency, COBYLA presents a fundamental limitation: being a local search method, it converges to the nearest local optimum from the given starting point, with no guarantee of finding the global optimum. Furthermore, as acknowledged by Powell himself, formal convergence guarantees within the standard conventions of optimization theory cannot be provided for COBYLA. The algorithm is also sensitive to the quality of the initial simplex: a poorly chosen starting point may lead to convergence in a shallow local minimum, far from the true optimum. These limitations motivate the initialization strategy adopted in this work. The starting geometry is not chosen arbitrarily: the best NACA airfoil identified through the Kriging surrogate map (Section 5.1) provides the baseline shape from which the CST parametrization is initialized, ensuring that the optimization begins in a region of the design space already known to perform well. From this baseline, a pool of candidate solutions is generated using Latin Hypercube Sampling over the CST coefficient bounds; the candidates are ranked by a preliminary evaluation, and COBYLA is applied independently to each of the top-ranked points. This multi-start approach significantly reduces the risk of premature convergence to a suboptimal local minimum, at the cost of a proportionally higher number of function evaluations.

4 Computational framework

The aim of this chapter is to explain how the calculations for the optimization loops have been carried out and how the code is structured. To this end, a waterfall approach is adopted, following the logical path that the code itself follows in the calculations, analysing each step by step.The framework has been developed following the GitHub format and is intended to be open source on the platform. Everything is ready to run by simply downloading and extracting a `.zip` folder, with no additional installation steps required. In this thesis the code "syntax" is not discussed in depth since the aim is not about how the code looks but about what the code does.
The program is available at `https://github.com/pietropoletto/BladeOptim.jl`;

the README file contains all the instructions needed to operate and modify the code parameters. The codebase is organised into sub-modules, each with a clearly defined and independent role, so that any component can be modified or replaced without affecting the rest of the framework. The `examples/` folder contains all the executables, while under the `src/` (source) folder all the sub-function modules are nested, each responsible for performing specific tasks. For instance, if one wants to modify the way the blade geometry is generated, it is sufficient to change the `BladeGeom/` module only, without affecting the rest of the code.

```
BladeOptim.jl/
├─ examples/
├─ Results/
├─ src/
│  ├─ BladeGeom/
│  ├─ CCBlade/
│  ├─ CST/
│  ├─ XFOIL/
│  ├─ OptimCST/
│  ├─ OptimNACA/
│  ├─ Plotting/
│  ├─ BladeOptim().jl
│  ├─ Config.jl
│  ├─ input_turbine.jl
│  └─ input_optimizer.jl
├─ tools/
└─ README.md
```

The following subsections mirror step by step the Baseline executable, which represents the main computational path from the airfoil geometry to the turbine power coefficient.

4.1 Airfoil parametrization

An airfoil shape is geometrically a continuous curve, and describing it exactly would require an infinite number of parameters. In an optimization context, this is not feasible: the design space must be finite and compact enough for the optimizer to explore efficiently. Airfoil parametrization addresses this by representing the shape through a small set of parameters, each controlling a specific geometric feature. A good parametrization should be expressive enough to capture a wide variety of physically relevant shapes, while keeping the number of degrees of freedom low. In this work two parametrization strategies are employed at different stages of the optimization pipeline: the NACA 4-digit family for the preliminary surrogate-based search, and the Class-Shape Transformation (CST) method for the gradient-free refinement.

4.1.1 NACA airfoils

The NACA 4-digit series is one of the oldest and most widely used families of standardized airfoil geometries, developed by the National Advisory Committee for Aeronautics in the 1930s. Its appeal lies in the direct geometric interpretation of its parameters and in the large body of experimental and numerical data available in the literature, which makes it a reliable and well-understood starting point for aerodynamic design.
In the 4-digit nomenclature the first number (A) indicates the maximum camber as a percentage of the chord, the second number (B) specifies the maximum camber position in tenths of the chord starting from the leading edge, and the last two numbers (C) indicate the maximum thickness as percentage of the chord [14]. For example, the NACA 6409 has a maximum thickness of 9% and 6% camber located at 40% of the chord length.

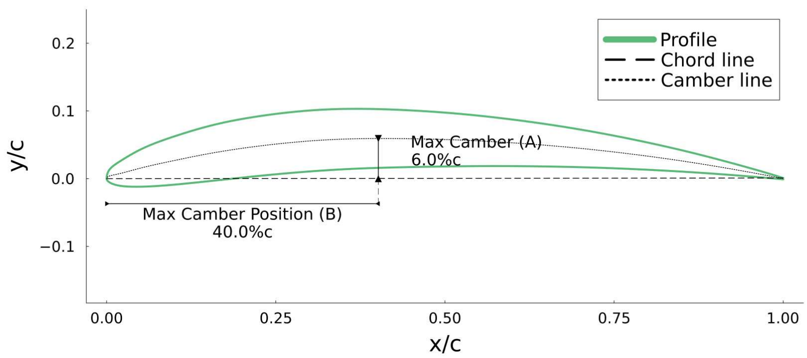

Figure 4.1: NACA 6409.

Considering constant thickness, the optimization problem relies only on two DoF: A and B. This is one of the reasons why this kind of airfoil has been chosen for the preliminary selection;

this optimization also serves as a visual representation of the optimization process, since the 3D nature of the problem makes it well suited for visualization. Another reason is that NACA 4-digit profiles are well-documented geometries, generally robust and stable, which provides a solid starting point.

4.1.2 NACA parametrization

To enhance the results of the analysis and to convey the capabilities of a Kriging model for such analysis, the limitation of unit intervals between possible values of camber (A) and camber position (B) is avoided. Instead of loading the NACA profile directly into XFOIL, a function is used that, given the combination of A and B as floats, generates the coordinate set in a `.dat` file. This file is then loaded into the flow solver in order to simulate a *float-parameters NACA profile*. According to [15], [14] it is possible to parametrize a 4-digit NACA profile in relation to the three main characteristics: A, B, C.
The first step consists in generating n nodes along the chord length (from 0 to 1); the spacing follows a cosine distribution to ensure a denser fitting near the edges.

$$x_i = \frac{1-\cos(\theta_i)}{2}, \qquad i = [1,\ldots,n], \qquad \theta_i \in [0,\pi] \tag{4.1}$$

The symmetric thickness distribution around the median line is given by the polynomial:

$$y_t(z) = 5t(0.2969\sqrt{z} - 0.1260z - 0.3516z^2 + 0.2843z^3 - 0.1015z^4) \tag{4.2}$$

The camber line is obtained by combining two parabolic functions, divided at the point of maximum camber p:

$$y_c(z) = \begin{cases} \dfrac{m}{p^2}(2pz - z^2) & z < p \\ \dfrac{m}{(1-p)^2}\left[(1-2p) + 2pz - z^2\right] & z \geq p \end{cases} \tag{4.3}$$

The thickness is then distributed in the normal direction from the camber line, both for the lower and upper surface:

$$\begin{aligned} x_u &= z - y_t \sin\theta, \quad y_u = y_c + y_t \cos\theta \\ x_l &= z + y_t \sin\theta, \quad y_l = y_c - y_t \cos\theta \end{aligned} \tag{4.4}$$

The points are placed starting from the trailing edge, following the whole profile around the leading edge and back to the starting point. This is due to the constraints imposed by `xfoil.jl` on the `.dat` file format.

At this point the LHS can generate points with much more freedom without being limited to integer values; this allows the Kriging model to explore the empty spaces between adjacent integer parameters, enhancing the quality of the map.

4.1.3 CST parametrization

The Class-Shape Transformation (CST) method, proposed by Kulfan and Bussoletti [16], provides an elegant way to represent airfoil geometries through a small set of parameters. The core idea is to decompose the airfoil surface into two components: a class function $C(x)$, which enforces the fundamental geometric constraints of the shape (such as a round leading edge and a sharp trailing edge), and a shape function $S(x)$, which modulates the specific geometry through a set of free coefficients. These coefficients represent the degrees of freedom (DoF) of the optimization problem. The more the DoF, the better the ability to shape the geometry, but the complexity of the solution space increases exponentially.
The general formulation is:

$$y(x) = C(x) \cdot S(x) + x \cdot \Delta z \tag{4.5}$$

Where Δz defines the thickness of the trailing edge, ideally 0, practically set to be feasible from a manufacturing point of view. The class function is defined as:

$$C(x) = x^{N_1}(1-x)^{N_2} \tag{4.6}$$

The parameters N_1 and N_2 define the type of geometry; for airfoils $N_1 = 0.5$ and $N_2 = 1$, since $\sqrt{x}$ leads to a round leading edge while $(1-x)$ gives a sharp trailing edge.
The shape function can be written as:

$$S(x) = \sum_{j=0}^{n} w_j \binom{n}{j} x^j (1-x)^{n-j} \tag{4.7}$$

The w_j's are the weight coefficients; upper and lower surfaces have two different sets, respectively w_u and w_l. The total number of coefficients is given by $2(n+1)$, so for $n = 2$ there are 6 DoF.
The x coordinates are distributed following a cosine distribution, with higher density at the edges since these areas are subject to large aerodynamic gradients, improving both the geometric fidelity and the flow solver convergence:

$$x_i = \frac{1}{2}\left(1 + \cos\left(\frac{2\pi i}{N}\right)\right), \quad i = 0, \ldots, N \tag{4.8}$$

Varying N changes the number of points. Once the points are generated, they are scaled to match the target maximum thickness-to-chord ratio t/c. The maximum thickness is defined as:

$$t = \max_i \left(y_u(x_i) - y_l(x_i)\right) \tag{4.9}$$

And the scaling of the y coordinates:

$$y_{\text{scaled}} = y \cdot \frac{(t/c)_{\text{target}}}{t} \tag{4.10}$$

The set of inputs results in 5 parameters: the upper and lower surface CST coefficients w_u and w_l, the trailing-edge thickness Δz, the target thickness-to-chord ratio Δt, and the number of discretization points N.

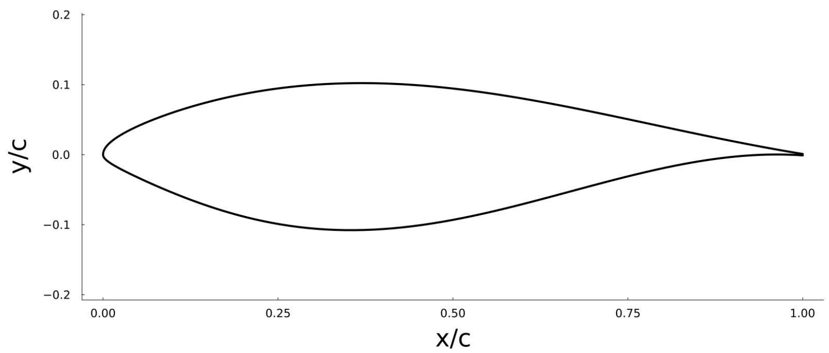

Figure 4.2: S809 airfoil parametrized by means of 4 coefficients both for upper and lower surface.

4.2 Flow solver

The aerodynamic polars were computed using `Xfoil.jl`, a Julia wrapper that exposes the panel method solver natively within the Julia environment, eliminating the need to call the official XFOIL executable as an external process and simplifying data exchange between modules.

Reynolds Number. CCBlade relies only on the BEM theory to evaluate the performance of the turbine; it neglects every interaction between different sections and radial movements. As input, BEM requires only the polar of the airfoil for different angles of attack. If a single polar file obtained with a single Reynolds number is adopted, all the blade sections refer to the same Reynolds number. Recalling the Reynolds formulation for a wind turbine:

$$Re = \frac{\rho \cdot W \cdot c}{\mu} \tag{4.11}$$

Where: ρ is the air density, c is the chord length, μ is the dynamic viscosity of the air and W is the relative velocity hitting the blade, given by:

$$W = \sqrt{V_\infty^2 + (\Omega r)^2} \tag{4.12}$$

Substituting (4.12) in (4.11):

$$Re = \frac{\rho \cdot \sqrt{V_\infty^2 + (\Omega r)^2} \cdot c(r)}{\mu} \tag{4.13}$$

It is easy to see, considering constant density, viscosity, wind speed and rotational speed, that the Reynolds number increases with the radius and decreases with the chord length. The contributions of radius and chord length do not have the same magnitude; in fact Re increases more quickly with the radius than it decreases with the chord. For very small turbines with a blade span around 1 meter, the combination of the Betz chord distribution and the limited span allows the Reynolds variation in the radial direction to be neglected. When the blade length increases, this approximation is no longer acceptable. The approach adopted here is the same used for the polars, airfoil shape and twist distribution, in which two sections of the blade are analyzed (at 40% and 90% of the bladespan) and all the middle sections' values are interpolated in between.

Reynolds number iteration. From eq. (4.11) the dependency of Re on the chord length is evident. The chord length is evaluated by means of the optimal Betz chord distribution eq. (**??**), which requires the polars c_l and c_d, leading to the following iterative loop:

$$Re \quad \longrightarrow \quad Polars(c_l, c_d) \quad \longrightarrow \quad Chord\, length \quad \longrightarrow \quad Re \tag{4.14}$$

The induction factors and the flow angle Φ can be evaluated as shown in the next section 4.3; at this point the relative velocity (4.12) becomes:

$$W = \sqrt{\left[(V_\infty(1-a))^2 + (\omega r(1+a'))^2\right]} \tag{4.15}$$

The iterative loop starts with initial values $Re = 10^6$, $c_l = 0.9$ and $c_d = 0.01$. The first association $(Re, \Phi) \leftrightarrow (c_l, c_d)$ is then established, enabling the evaluation of c_n and the chord length, which in turn yields the next Reynolds number.

The tip and root Re values obtained are then used throughout the whole optimization process (the chord length changes slightly with different profiles, but the difference is negligible).

Airfoil transition. The transition criterion used in XFOIL is based on the e^N method, where the parameter N_{crit} represents the logarithm of the amplification factor of Tollmien–Schlichting instabilities at the point of laminar-to-turbulent transition. A higher value of N_{crit} corresponds to a lower freestream turbulence intensity, meaning transition occurs further downstream and the boundary layer remains laminar over a larger portion of the surface [17]. The choice of N_{crit} should reflect the operational environment of the turbine rather than its size directly; however, for small-to-medium wind turbines in the range up to 200 kW, the two are often correlated. Smaller rotors operate at lower chord Reynolds numbers, typically between 10^5 and 10^6, and are more commonly deployed in sites with higher atmospheric turbulence intensity, such as urban or complex terrain environments,compared to large utility-scale machines installed at open, low-turbulence sites. An approximate relationship between freestream turbulence intensity Tu and N_{crit} is given by [18]:

$$N_{crit} = -8.43 - 2.4 \ln\left(\frac{Tu}{100}\right) \tag{4.16}$$

For a turbulence intensity of Tu = 0.1% (low-turbulence wind tunnel or calm open-air conditions) this yields $N_{crit} \approx 12$, while Tu = 1% gives $N_{crit} \approx 6$.

4.3 Blade geometry

In this section, blade geometry refers to the spanwise distributions of chord and twist angle along the blade radius. Rather than adopting fixed distributions or including the geometry directly among the optimization parameters, a hybrid alternative is pursued. Chord and twist are not fixed and vary depending on the airfoil and its polar. Once the optimal angle of attack (depending on the objective function taken into consideration) of an airfoil is found, the geometry is calculated as a consequence of those parameters, following the Betz rules for the optimal chord and twist distribution [2]. From equations (2.20), (2.21) a polynomial equation linking the axial induction factor a and the local speed ratio x can be obtained:

$$16a^3 - 24a^2 + a(9 - 3x^2) - 1 + x^2 = 0 \qquad x = \frac{\omega r}{V_\infty} \tag{4.17}$$

After solving the third-degree polynomial for a, the value of a' can be evaluated using eq. (2.21). Once both axial and tangential induction factors are known, the flow angle ϕ follows from eq. (2.30), rewritten in terms of the local speed ratio $x = \omega r / V_\infty$:

$$\tan\phi = \frac{(1-a)\,V_0}{(1+a')\,\omega r} = \frac{(1-a)}{(1+a')\,x} \tag{4.18}$$

The optimal blade twist Θ_{opt} is by definition:

$$\Theta_{opt} = \phi - \alpha_{opt} \tag{4.19}$$

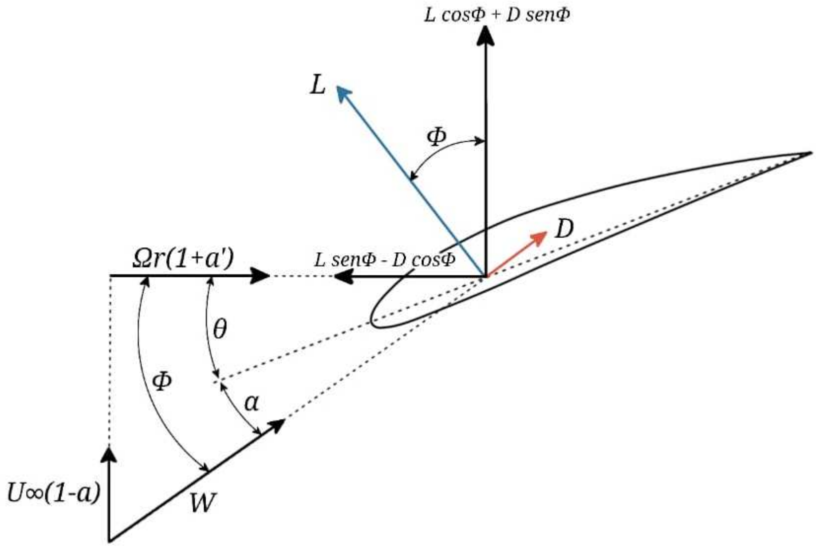

Figure 4.3: Velocity triangle and aerodynamic forces acting on a blade element. The flow angle Φ is the sum of the local twist angle Θ and the angle of attack α.

And the optimal chord distribution is given by:

$$\frac{c(x)}{R} = \frac{8\pi F a x \sin^2 \varphi}{(1-a)B\lambda C_n} \tag{4.20}$$

Where the normal force coefficient C_n is a combination of lift and drag contributions:

$$C_n = C_{l,opt} \cos(\phi) + C_{d,opt} \sin(\phi) \tag{4.21}$$

The term F accounts for Prandtl's tip loss, as defined in eq. (2.37).
For a micro-wind turbine ($< 2\,$kW), a single airfoil with a single thickness can be adopted for simplicity. For larger turbines where a thickness distribution is required, the situation is different. When the size of the turbine reaches the order of tens of kW, the use of a single thickness becomes impractical and counterproductive, and a thickness distribution along the blade must be introduced. The use of different thicknesses introduces differences in the airfoil shape, whether the same profile is scaled in thickness or the geometry is changed entirely. Ideally, the results of every section defined by BEM theory should rely on the polar of the corresponding airfoil. This method (considering a minimum of 10 sections) requires more than $10\times$ the number of XFOIL simulations; unsustainable in an optimization process where hundreds of evaluations

are needed.
The solution is to use an interpolation approach for both polars and airfoil shapes. Considering CST-parametrized airfoils, given root and tip geometries, all the n sections in between can be represented by the n-weighted mean values of the CST parameters. An example is given below:

```
AIRFOIL_ROOT = (                  AIRFOIL_TIP = (                  AIRFOIL_MID = (
wu = [0.17, 0.37, 0.23],          wu = [0.19, 0.30, 0.21],         wu = [0.19, 0.34, 0.23],
wl = [-0.12, -0.29, -0.44],       wl = [-0.07, -0.12, -0.14],      wl = [-0.10, -0.21, -0.29],
dz = 0.004,                       dz = 0.004,                      dz = 0.004,
dt = 0.23,                        dt = 0.12,                       dt = 0.175,
N  = 220)                         N  = 220)                        N= 220 )
```

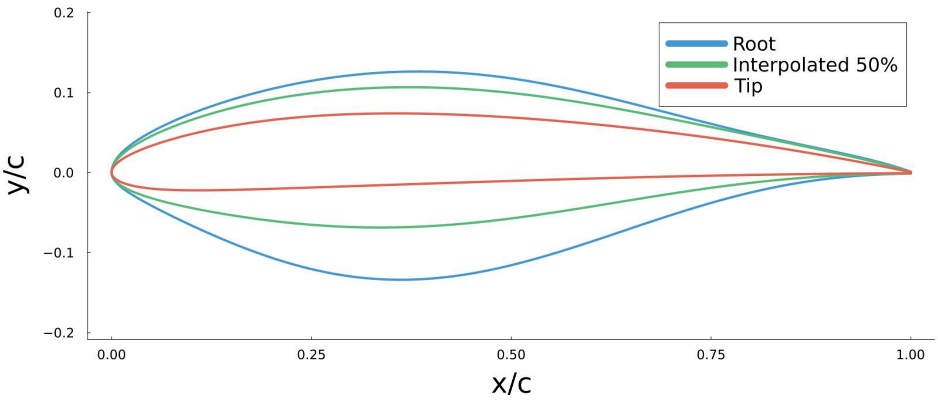

Figure 4.4: Interpolated airfoil between root and tip.

When defining the thickness distribution of the blade, some basic concepts must first be clarified. The first part of the blade (roughly the inner 40 percent) has mainly structural duties, so the thickness must ensure a geometry capable of withstanding flapwise and edgewise loads, which are maximum at the root. In this area aerodynamics plays a secondary role: lower lever arm and lower velocities result in a low torque contribution. In the external region, moving towards the tip, the blade carries less weight and load, and aerodynamics becomes increasingly relevant; here the thickness can decrease in order to achieve higher values of C_l/C_d and thus increase the turbine's power coefficient. Two polars are evaluated at 40 percent (root geometry) and 90 percent (tip geometry), dividing the blade into three main regions.

Root section. The geometry evaluated at 40% is used, projected all the way to the hub. Normally the full profile is scaled with the chord length, but when the chord becomes quite small close to the hub, a linear transition to a cylinder attached to the hub is applied.

This leads to an easy manufacturing joint between hub and blade root, and at the same time avoids thin trailing-edge sections, increasing the strength of the first sections.

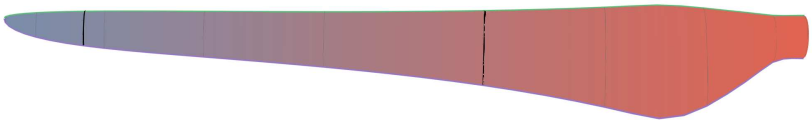

Figure 4.5: Example of the transition to a cylinder.

Tip section. From 0.9 to 1.0 of the blade radius, the airfoil evaluated at 90% is used. This profile, called tip airfoil, is optimised to be as efficient as possible without considering structural aspects. The only constraints concern the feasibility of the shape itself, a minimum thickness is imposed to avoid geometrically unrealistic sections. The airfoil is not evaluated at 100% because at the very tip the aerodynamic regime is overwhelmed by tip vortices and losses; it is therefore a deliberate choice to set the reference profile at 90% of the blade, still very close to the tip but safe from tip loss phenomena.

Mid section. The sections in this region are obtained by interpolating between the root and tip geometries, gradually linking the lower and upper parts of the blade. A similar approach is used for the polars: once the root and tip airfoils are evaluated, all mid-section results are obtained by blending between the two. The first step is interpolating C_l and C_d of both sections to a common α grid, to handle non-converged points in the alpha span. For the sections before the root station, the root airfoil's polars are used; the same applies for sections beyond the tip station. The polars of the sections in between are obtained by weighted average between root and tip.
Some considerations about twist and chord distribution are in order. The outer part of the blade is responsible for most of the power production, while the part close to the hub has mainly structural purpose. For the twist, a distinction must be made between a pitch-controlled turbine and a stall-controlled turbine. In a pitch-regulated machine the blade always faces the wind in an optimal way; the twist is therefore set to be optimal all along the blade. Recalling (4.19), in the case of a dual geometry the optimal angle of attack may differ between tip and root; the two optimal twists are evaluated as:

$$\Theta_{opt,root} = \phi - \alpha_{opt,root} \qquad \Theta_{opt,tip} = \phi - \alpha_{opt,tip} \tag{4.22}$$

$\Theta_{opt,root}$ propagates to the root, $\Theta_{opt,tip}$ is used at the tip, and in the middle the two twist values are interpolated, following the same logic applied to polars and airfoil geometry.
In a stall-controlled turbine the blade is fixed to the hub and the power decrease after the nominal

wind speed is caused by stall processes occurring as the wind speed increases. Rather than avoiding stall as in pitch control, the aim is to exploit and control it. Since the tip is responsible for the majority of the power production, that part should work as efficiently as possible. The goal is to achieve smooth stall propagation from the root to the tip, also to avoid violent transient loads on the blade following rapid wind fluctuations. To this end, the twist at the root can be increased by a fixed offset in order to bring it closer to stall, then propagating this offset through the blade and reducing it to zero at the tip:

$$\Theta_{opt,root} = \phi - (\alpha_{opt,root} - \delta) \qquad \Theta_{opt,tip} = \phi - \alpha_{opt,tip} \tag{4.23}$$

And δ goes from its maximum value at the root to zero at the tip.
The inner part of the blade, being thicker, produces less lift. From (4.21) it follows that the higher the lift, the lower the chord length. Using a chord distribution referred to the tip, the root of the blade presents a smaller chord than its optimal. This is actually an advantage because, given the lower aerodynamic contribution at the root, it is common practice to reduce the chord length close to the hub; this allows less material to be used and simplifies the manufacturing process.

4.4 Viterna extrapolation

A BEM solver such as CCBlade (next section) requires the 2D airfoil performance data. These data can be taken from tables or, as in the present case, evaluated by means of a tool such as `xfoil.jl`; in either case only a limited range of angles of attack is typically covered. During normal wind turbine operation, especially at startup and shutdown or under extreme conditions, angles of attack well beyond the flow solver range can be encountered. To handle this, BEM uses the Viterna extrapolation method, which extends the available airfoil polars to the full 360° range by blending the measured data with flat-plate aerodynamic theory at high angles of attack. This ensures that the BEM solver always has valid aerodynamic coefficients, avoiding lookup failures or unphysical discontinuities, while still capturing the essential behaviour of a deeply stalled airfoil. The only input required by the Viterna method is an initial dataset; the airfoil's polars serve this purpose. Viterna's formulation is given by [19]:

$$C_L = A_1 \sin 2\alpha + A_2 \frac{\cos^2\alpha}{\sin\alpha} \quad , \quad C_D = B_1 \sin^2\alpha + B_2 \cos\alpha \tag{4.24}$$

Where:

$$C_{D_{\max}} \simeq 1.11 + 0.018\,AR \tag{4.25}$$

$$A_1 = \frac{C_{D_{\max}}}{2} \tag{4.26}$$

$$B_1 = C_{D_{\max}} \tag{4.27}$$

$$A_2 = (C_{L_{\text{stall}}} - C_{D_{\max}} \sin \alpha_{\text{stall}} \cos \alpha_{\text{stall}}) \frac{\sin \alpha_{\text{stall}}}{\cos^2 \alpha_{\text{stall}}} \tag{4.28}$$

$$B_2 = \frac{C_{D_{\text{stall}}} - C_{D_{\max}} \sin^2 \alpha_{\text{stall}}}{\cos \alpha_{\text{stall}}} \tag{4.29}$$

The symbol AR represents the aspect ratio of the blade section, which in a BEM context accounts for finite blade length effects on the flat-plate analogy. In CCBlade, the recommended value is AR=10, as variations of this parameter have a negligible influence on the final results. For angles of attack beyond $\alpha > 90°$ and below $\alpha < \alpha_{\text{min}}$, the Viterna equations are applied symmetrically by reflecting the computed coefficients. It is important to note that the Viterna method does not attempt to resolve the pressure or skin friction distribution over the airfoil; rather, it relies on simplified analytical expressions calibrated to match the known stall-point data. Despite this approximation, the extrapolated polars provide estimates that are sufficiently accurate for BEM-based analyses, making the method well suited for preliminary design and performance evaluation of wind turbine blades within the CCBlade framework.

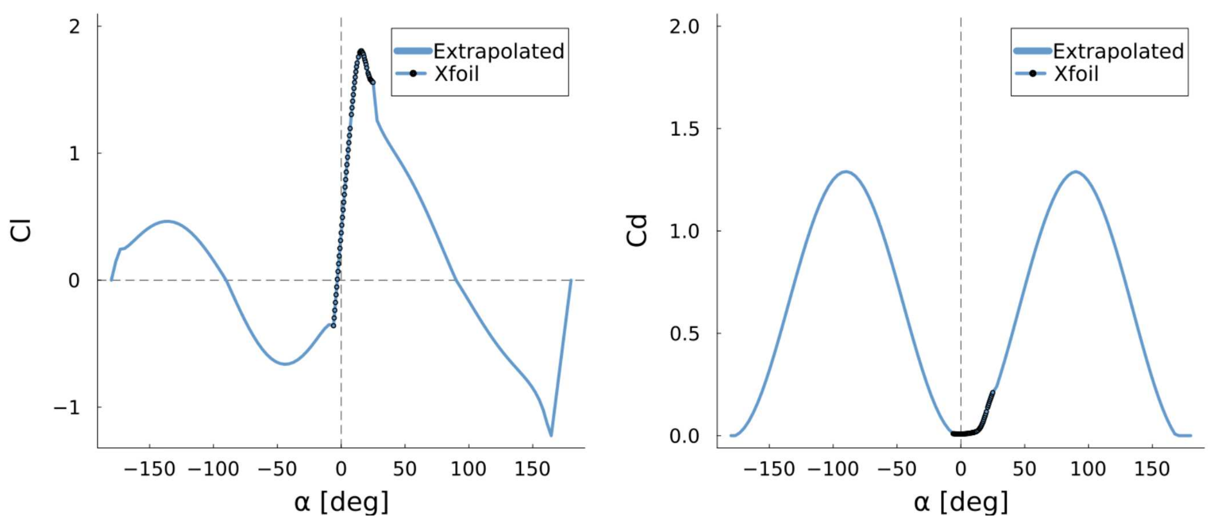

Figure 4.6: Example distribution.

4.5 BEM method

CCBlade is an open-source BEM solver developed by Andrew Ning in 2013 [20]. CC stands for Continuity and Convergence. It is available as a Julia package and its implementation is relatively straightforward. Only three types of inputs are required:

Input Category	**Description**
Geometrical	Radial sections r, chord distribution $c(r)$, twist distribution $\theta(r)$, rotor radius R_{tip}, hub radius R_{hub}, number of blades B
Polars	Airfoil aerodynamic data (α, C_l, C_d) from XFOIL analysis, extrapolated to $\pm 180°$ via Viterna
Operative	Freestream wind speed V_∞, rotational speed Ω, air density ρ, pitch angle θ_p, hub height, wind shear

Table 4.1: Input data required by CCBlade.

CCBlade differs from traditional BEM solvers in the way the 2D system is solved. Rather than using induction factors (a, a'), the entire system is solved by means of a single variable: the flow angle ϕ between the relative velocity and the rotor plane. Instead of an iterative process between a and a', the solver only needs to find the value of ϕ that closes the system. Ning demonstrated that a solution in ϕ can always be found, which guarantees continuity. This ensures convergence regardless of how the configuration changes, making CCBlade particularly well suited for optimization loops.

#	**Step**	**Description**	`CCBlade call`
1	**Rotor**	Build the rotor object from global geometry (R_{hub}, R_{tip}, B, precone)	`Rotor()`
2	**Sections**	Define each blade section with its local chord, twist, and airfoil polar	`Section.()`
3	**Op. Point**	Construct operating conditions for each section at given V_∞ and Ω	`windturbine_op.()`
4	**Solve**	Solve BEM equations; integrate for thrust T, torque Q, C_P, C_T	`solve.(), nondim()`

Table 4.2: CCBlade workflow: from inputs to performance coefficients.

5 Rotor blade optimization

Some parts of this section of the thesis concerning the numerical optimization of the airfoil were developed in collaboration with another master's student working on a similar topic applied to VAWTs. Since the final objective functions differ according to the nature of the machines studied, the collaboration was limited to the aerodynamic analysis of the airfoil and to the implementation of the numerical optimization loop. As mentioned in Chapter 2, the selection of a good starting point is an important aspect when setting up a numerical optimization problem. Optimizing an airfoil over six, seven, or more Degrees of Freedom (DoF) leads to a design space that grows exponentially with the number of parameters, making the search computationally expensive. To mitigate this, the optimization is structured in two sequential routines. The first identifies the NACA 4-digit airfoil closest to the optimal configuration, reducing the preliminary search to only two parameters. The second routine then performs the full airfoil optimization using a CST parametrization, starting from the geometry obtained in the first step.

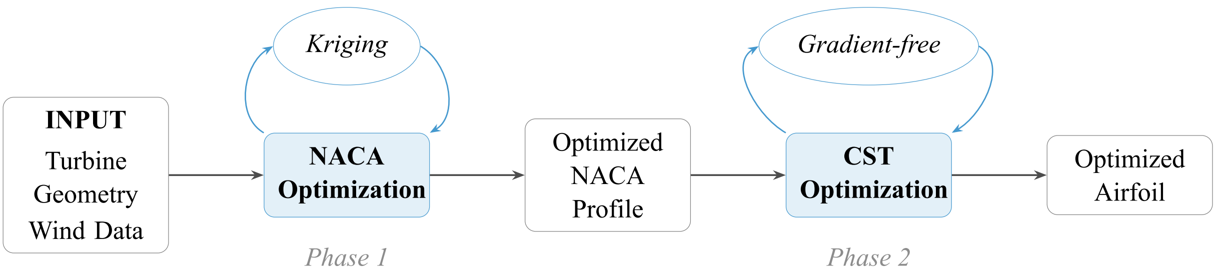

Figure 5.1: Overall optimization pipeline.

5.1 Kriging applied to NACA

The entire design space for the chosen objective function can be evaluated across all NACA airfoils within the defined boundaries. The only required inputs are the camber and camber position limits, the Reynolds number and XFOIL settings, and the airfoil thickness. Given only two DoF and the narrow parameter ranges for camber $(2-9)$ and camber position $(2-8)$, the initial number of samples can be kept relatively low. The lower bounds exclude camber values of 0 and 1, which produce nearly symmetric profiles with negligible lift, and a camber position of 1, which places the maximum camber at the leading edge resulting in an aerodynamically degen-

erate shape. The upper bounds reflect the validity range of the NACA 4-digit parametrization: camber values above 9% and positions beyond 80% of the chord fall outside the region in which the two-parabola camber-line formulation is reliable. It has been observed that a number of initial samples around 16 is more than enough for a good result. The number of initial samples is not critical since after the sampling, the Kriging loop starts minimizing the error; with few starting points more iterations are required to reach convergence, while increasing the number of initial samples beyond a certain threshold leads to wasted evaluations, as the map is already well defined with fewer samples.
To generate the DOE (Design of Experiments), the Latin Hypercube sampling julia package is used, requiring only the number of initial samples and the boundaries of each dimension as inputs. After evaluating the objective function at each of the starting points, the surrogate begins refining the solution space by identifying the point with the greatest variance, solving for the corresponding airfoil, and updating the Kriging model. The process stops when the prescribed minimum error is reached.
As a representative example, the best performing 12% thickness airfoil in terms of aerodynamic efficiency (C_l/C_d) is sought within a tolerance of 0.5% of the mean value of the objective function. All the following results are obtained at a Reynolds number of 4×10^5. After the first 16 samples shown in Figure 5.3a, the error goal has been reached in 55 iterations, with a computation time on the order of 1–2 minutes.

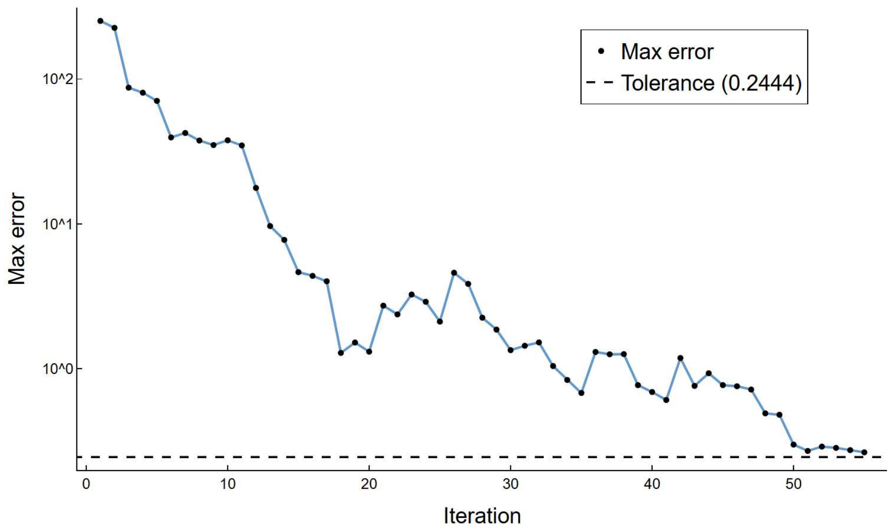

Figure 5.2: Kriging error minimization convergence: maximum squared error at each iteration. The dashed line indicates the tolerance threshold.

As visible in Figure 5.2, the error occasionally increases after an iteration. The Kriging surrogate uses a smoothness parameter $p = 2$ corresponding to a Gaussian correlation function, which enforces a twice-differentiable, infinitely smooth response surface. While this produces well-behaved interpolation in regular regions, it can cause the surrogate to over-smooth local features; when a new sample is added in a high-error region, the model must reconcile the new point with its global smoothness constraint, which can temporarily redistribute the variance and cause the maximum error to increase in a neighboring region. This behavior is analogous to pressing down on a tarp with trapped air underneath: flattening one spot forces the bulge to migrate elsewhere rather than disappear.

The resulting maps of all 12% thickness NACA airfoils within the boundaries are shown below. A comparison between the map before and after the error minimization illustrates the refinement achieved by the surrogate.

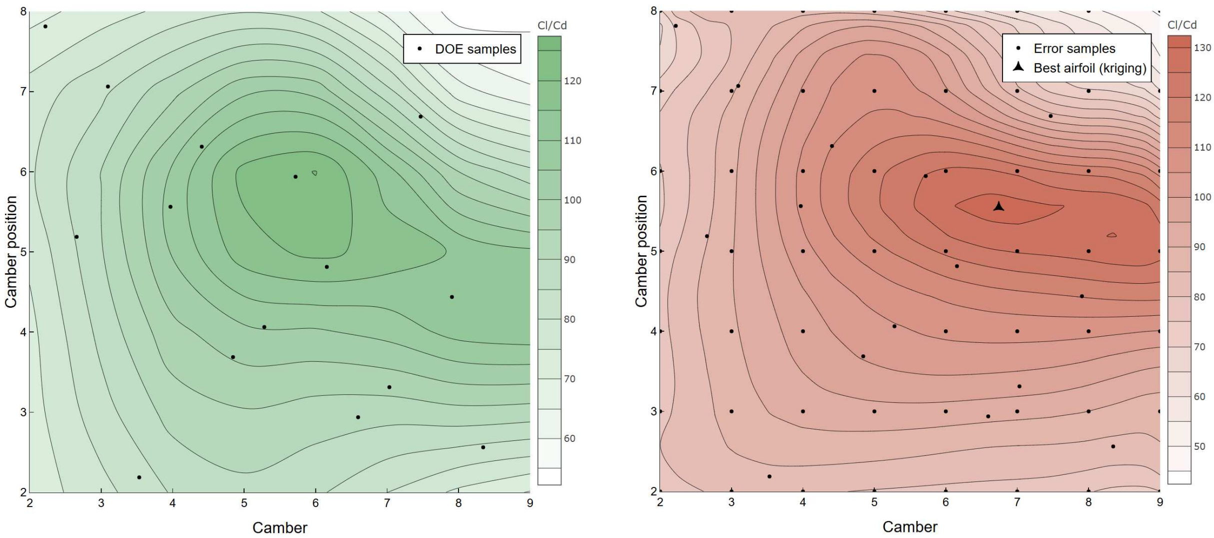

(a) Kriging from 16 initial DOE samples

(b) Kriging after error minimization

Figure 5.3: C_l/C_d Kriging surrogate maps for 12% thickness NACA airfoils, $Re = 4 \times 10^5$.

The Kriging surrogate identifies the optimal airfoil at camber $A = 6.74\%$ and camber position $B = 5.55\%$, a configuration that does not correspond to any integer NACA combination. This result highlights the advantage of treating A and B as continuous parameters: the true optimum lies in the interior of the design space, between the integer grid points, and would have been missed by a purely discrete search. The surrogate map confirms that the region around $A \approx 6$–9 and $B \approx 5$–6 consistently yields the highest C_l/C_d values, with the peak predicted by the model reaching $C_l/C_d = 131.9$, compared to the best configuration found during the DOE: NACA (5.72-5.94-12) at $C_l/C_d = 126.5$.

The Kriging surrogate was validated by performing a full enumeration of the design space: XFOIL was run for every integer combination of camber and camber position within the specified bounds, yielding a complete 7×8 solution map of 56 evaluations. This brute-force map serves as the ground truth against which the surrogate predictions are compared, allowing a direct assessment of the Kriging accuracy across the entire domain.

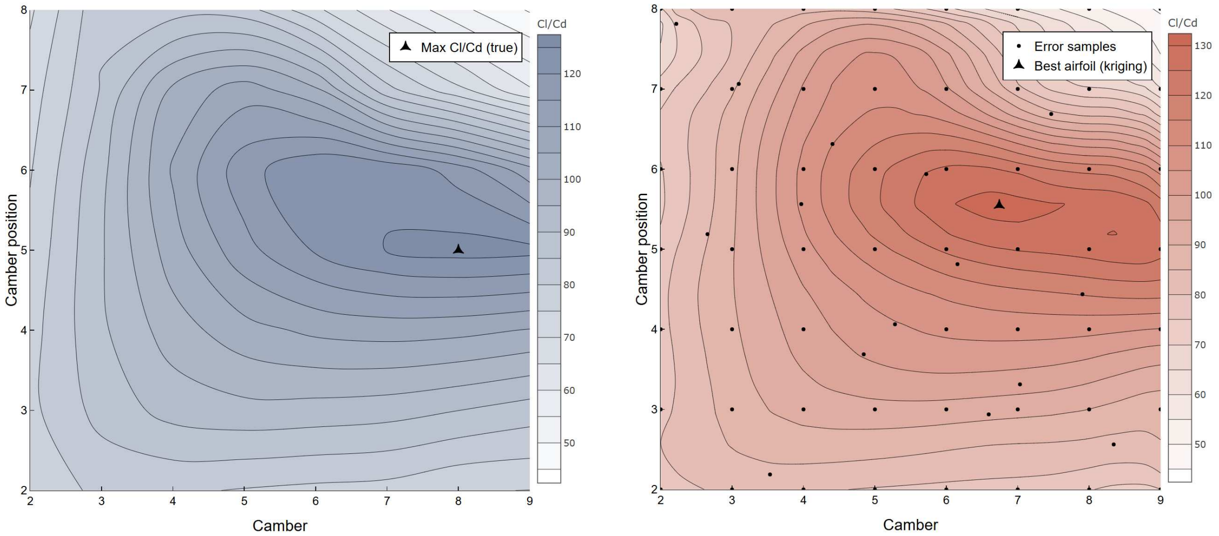

(a) True C_l/C_d map (full enumeration)

(b) Kriging after error minimization

Figure 5.4: Validation: true NACA map vs. Kriging surrogate, 12% thickness, $Re = 4 \times 10^5$.

The comparison between the true NACA map and the Kriging surrogate reveals overall excellent agreement across the entire domain. A notable discrepancy is observed in the high-performance region, specifically for camber position $B \in [5, 6]$ and maximum camber $A \in [5, 8]$, where the surrogate predicts slightly higher C_l/C_d values than the integer-sampled ground truth (Figure 5.5). This is consistent with the identified optimum at $A = 6.74\%$, $B = 5.55\%$ falling precisely within this region, between integer grid points, confirming that the Kriging model captures and exploits the smooth variation of the objective function in areas not directly covered by the discrete NACA configurations.

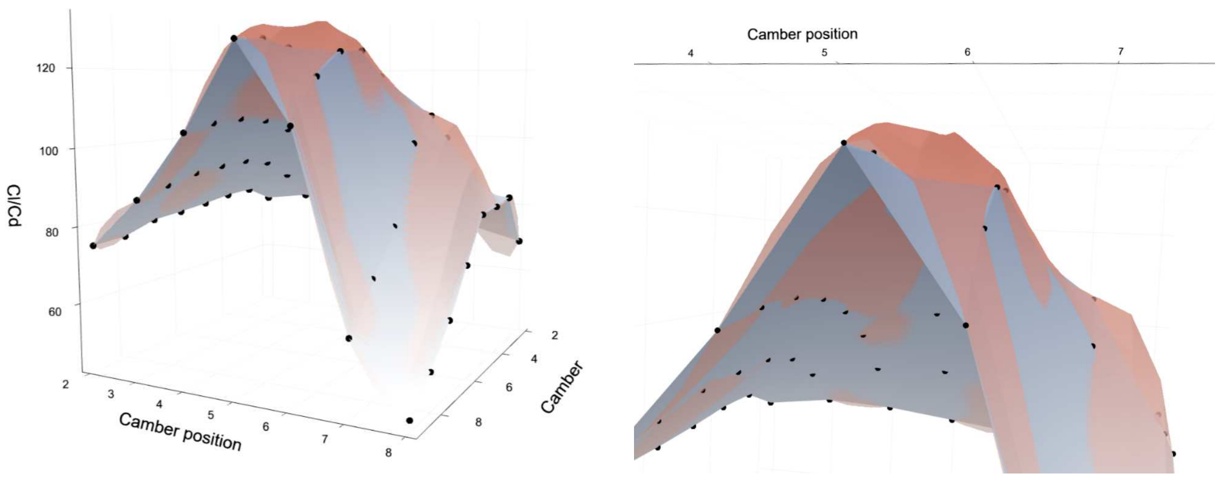

(a) View 1

(b) View 2

Figure 5.5: 3D surface comparison: true NACA map vs. Kriging surrogate.

The procedure described in this section is fully general: by changing the thickness parameter C and the objective function in the configuration file, the same Kriging-based routine can be applied to map the entire design space for any NACA 4-digit family. This allows the generation of a library of optimal starting profiles, one for each combination of thickness and objective, to be used as initial geometries for the subsequent CST-parametrized optimization.

5.2 Optimization algorithm

This section describes how the COBYLA optimizer is integrated into the framework and applied to the CST-parametrized airfoil optimization. The implementation follows a multi-start strategy in order to mitigate the limitation of COBYLA as a local optimizer, while keeping the total computational budget within a range that is tractable on a standard laptop.

5.2.1 Overall structure

The optimization of each airfoil (tip or root) proceeds in two stages. First, a set of candidate starting points is generated using Latin Hypercube Sampling over the CST coefficient bounds. Each candidate undergoes the geometric feasibility checks described in Section 5.4 before being evaluated; candidates that fail any check are discarded and replaced. The baseline airfoil from the Kriging phase is always included as the first point in the DOE, ensuring that the search starts from a region already known to perform well. All surviving candidates are evaluated through the full pipeline, CST coordinate generation, XFOIL analysis, Betz chord and twist calculation, polar blending, Viterna extrapolation, and CCBlade BEM solve, and ranked by the penalized objective value.

In the second stage, the top-ranked N_{TOP} candidates from the DOE are used as independent starting points for N_{TOP} parallel COBYLA runs. Each run is allowed a maximum of N_{eval} function evaluations and is subject to the convergence tolerances described below. The best result across all restarts is selected as the final optimized airfoil for the current phase. This multi-start approach significantly reduces the risk of converging to a shallow local minimum: even if one COBYLA instance gets trapped, another starting from a different region of the design space may find a better solution.

The total computational budget per phase is therefore

$$N_{\text{budget}} \;=\; N_{\text{DOE}} \;+\; N_{\text{TOP}} \times N_{\text{eval}} \tag{5.1}$$

function evaluations. The dominant cost per evaluation is the XFOIL analysis, which typically requires 2–4 seconds depending on the airfoil geometry and the convergence behaviour of the viscous solver; the CCBlade BEM solve adds approximately 0.01 seconds, which is negligible by comparison. With typical settings of $N_{\text{DOE}} = 40$, $N_{\text{TOP}} = 4$, and $N_{\text{eval}} = 80$, a single phase requires approximately 360 evaluations and completes in about 22 minutes. Since the optimization involves two sequential phases (tip, then root), the total wall time for a full blade optimization is approximately 45 minutes on a standard laptop with a single CPU core, consistent with the design goal stated in the introduction of running on accessible hardware.

5.2.2 COBYLA configuration

The optimizer is implemented through the NLopt library, which provides a COBYLA implementation accessible from Julia, `LN_COBYLA` (Local, No-gradient, Constrained Optimization BY Linear Approximations). The algorithm constructs a linear surrogate of the objective function at each iteration using the $n+1$ vertices of a simplex in the n-dimensional CST coefficient space, and solves a trust-region subproblem to propose the next candidate point. The box constraints from Equations (5.7)–(5.8) are enforced as lower and upper bounds on each design variable. The leading-edge and trailing-edge inequality constraints defined in Section 5.4 are passed to NLopt as explicit nonlinear constraints; COBYLA linearises them at each iteration and handles them directly within the trust-region subproblem, without penalty reformulation.
The optimizer maximizes the penalized objective value, which is either the peak power coefficient C_P multiplied by the stall penalty factor $(1-p)$ (see Section 5.3.1), or the penalized C_l/C_d ratio depending on the configuration. Since NLopt's COBYLA implementation supports direct maximization through the `max_objective` interface, no sign inversion is required.

Initial step size. The trust-region radius ρ is initialized as a fraction of the mean bound width:

$$\rho_0 = f_\rho \cdot \overline{(ub - lb)} \tag{5.2}$$

where f_ρ is a user-defined fraction, typically set to 0.25. This parameter controls how aggressively the first iterations explore the neighbourhood of the starting point: a value of 0.05 produces a conservative search that converges quickly near the DOE point, while 0.30 allows a broader initial exploration at the cost of more evaluations before convergence. The value of 0.25 was found to provide a good balance between exploration and convergence speed, prioritizing a larger initial seeking area to avoid local a optimum.

Convergence criteria. Two stopping conditions are enforced simultaneously. The relative function tolerance $f_{\text{tol}} = 10^{-4}$ triggers termination when the objective value changes by less than 0.01% between consecutive iterations, indicating that the optimizer has effectively converged.

The relative variable tolerance $x_{\text{tol}} = 10^{-3}$ stops the optimization when all design variables change by less than 0.1% of their current values, detecting stagnation in the design space. In addition, a hard limit of N_{eval} function evaluations is imposed to cap the computational cost regardless of convergence.

Parameter	Value	Description
Algorithm	`LN_COBYLA`	Local, gradient-free, constrained
Design variables	7	3 upper + 4 lower CST coefficients
f_ρ	0.1–0.25	Initial trust-region radius
f_{tol}	10^{-4}	Relative function tolerance
x_{tol}	10^{-3}	Relative variable tolerance
N_{eval}	80	Maximum evaluations per restart
N_{DOE}	40–60	LHS screening candidates per phase
N_{TOP}	3–6	Multi-start restarts (top DOE points)
Constraint tol.	10^{-8}	Nonlinear inequality constraint tolerance

Table 5.1: COBYLA optimizer parameters used in the next sections for the optimization.

A sanity check is applied to every returned value: results outside the physically meaningful range ($-0.5 \leq C_P \leq 0.58$ for the C_P objective, or $-50 \leq C_l/C_d \leq 200$ for the aerodynamic efficiency objective) are replaced by the penalty value. This prevents numerical artefacts from corrupting the optimizer's internal model.

5.2.3 Sequential dual-airfoil optimization

The framework optimizes two airfoils: the tip profile (evaluated at $\eta = 0.90$) and the root profile (evaluated at $\eta = 0.40$). Since both airfoils influence the same objective function through the blended polar, optimizing them simultaneously would double the design space to 14 variables, making the problem significantly harder for a local optimizer. Instead, the optimization follows a sequential strategy organized in two phases:

Tip optimization. The root airfoil is fixed at its baseline geometry and its polar is cached. The optimizer varies only the 7 CST coefficients of the tip airfoil, while the root contribution to the blended sections remains constant. Since the tip sections operate at higher relative velocities and contribute the majority of the rotor torque, this phase typically produces the largest improvement in C_P.

Root optimization. The tip airfoil is now fixed at the optimized geometry found in Phase 1, and its polar is cached. The optimizer varies the root CST coefficients, seeking a root profile that complements the optimized tip. Different geometric constraints apply: the root airfoil is allowed a higher maximum thickness (30% vs. 20%) and maximum camber (7% vs. 8%) to accommodate the structural requirements of the inner blade region.

This sequential approach reduces each optimization to a 7-variable problem, which lies well within the tractable range for COBYLA. The polar caching ensures that each function evaluation requires only one XFOIL call (for the variable airfoil) rather than two, halving the per-evaluation cost. The sequential structure also reflects the physical hierarchy of the blade design: the tip airfoil dominates the aerodynamic performance, while the root airfoil plays a secondary role and can be adapted to the tip geometry once it is finalized.

5.3 Objective function

The aim of this thesis is not the airfoil's performance optimization itself but the optimization of the wind turbine's rotor blade. The primary goal is the maximization of the coefficient of performance of the machine. The sole maximization of C_p, however, can lead to unstable or impossible turbine configurations; it is therefore necessary to define the boundaries within which the optimization operates and the characteristics the final result must satisfy. For mini urban wind turbines (up to some kW) without any control system, the most important characteristic is a C_p as constant as possible for different wind speeds, and consequently a smooth stall transition. For large turbines equipped with pitching control and able to maintain optimal conditions under varying operating conditions, the most important parameter is the maximization of the power coefficient.
As discussed, C_p alone is not sufficient. The optimization is formulated as a single-objective problem. The primary objective is the maximization of C_p; secondary criteria are incorporated as penalty terms that scale the objective value, rather than as independent objectives in a Pareto sense.

5.3.1 Stall penalty

Even in the presence of a pitch control system, violent transients during wind fluctuations should be avoided. For this reason, the stall behaviour of each airfoil is monitored and penalized. The C_l envelope is strictly related to the airfoil geometry. To enhance the stability of the blade, C_l must follow an ideal path with constant slope until a smooth and gradual decrease (stall). In the optimization process, reshaping an airfoil surface can generate local stall points, which manifest as local decreases in lift followed by a recovery. These points can cause fluctuations and erratic behaviour during normal operation of the wind turbine. The control of the lift curve shape has been implemented as a penalty on the C_p.

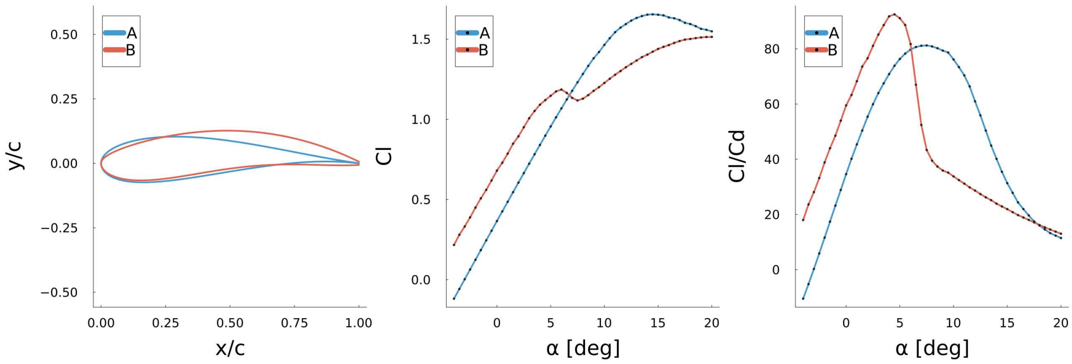

Figure 5.6: Comparison of two airfoils: although profile B achieves a higher C_l/C_d, its stall behaviour is abrupt and unstable around the stall dip ($\sim 5°$), Re = 400000.

For every generated geometry, the gradient along the lift curve is calculated. A constant gradient up to a point of decrease near the global maximum lift indicates a so-called "good stall behaviour". The algorithm detects discrepancies in the gradient and assigns a penalty to C_p whenever local valleys are identified. This penalty ranges from 0 to 1 and depends on the severity of the local valley, rather than using a binary approach that discards airfoils violating the constraint. In this way, a high-performance airfoil with poor stall behaviour is penalized relative to a lower-performance airfoil with a smooth stall. The weight and extent of this selection are analysed in the following.

Gradient calculations. From the polars, the array index corresponding to the maximum C_l (the global maximum) is identified; all verifications are performed on points between index 1 and the index of $C_{l,\text{max}}$, i.e. $i \in \{1, \ldots, i_{C_{l,\text{max}}}\}$. In this region C_l should behave as a monotonically increasing function, as shown by the blue curve in Figure 5.6. The local gradient is evaluated for every couple of adjacent points:

$$\frac{\mathrm{d}C_l}{\mathrm{d}\alpha} = \frac{C_l[i] - C_l[i-1]}{\alpha[i] - \alpha[i-1]} \tag{5.3}$$

Whenever the local gradient is negative, the square of its value is accumulated in a sum:

$$S = \sum_{\frac{dC_l}{d\alpha}<0} \left(\frac{dC_l}{d\alpha}\right)^2 \tag{5.4}$$

This allows quantification of the magnitude of local valleys. The sum S is then mapped to the interval $[0, 1]$ through an exponential function:

$$p = 1 - \exp^{(-k\,S)} \tag{5.5}$$

where k is the sensitivity parameter. For large values of S the penalty tends to 1, while it remains zero if $S = 0$.

k value	Behaviour	Sensitivity
$k = 0$	No penalty	Debugging
$k = 20$	Soft penalty	Only deep dips
$k = 50$	Balanced	General use
$k = 100$	Hard penalty	Also small dips

Table 5.2: Effect of the sensitivity parameter k on the penalty function.

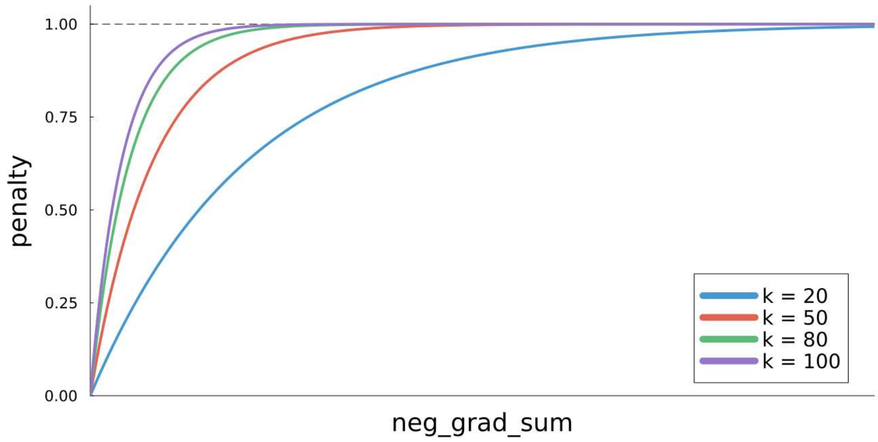

Figure 5.7: Penalty function $p = 1 - \exp(-k\, S)$ for different values of k.

Penalty application. The penalty for each airfoil lies in the interval $[0, 1]$, where 1 corresponds to complete penalization.

$$C_{p,\,\text{penalized}} = C_p\,(1 - p) \tag{5.6}$$

This gradually scales the C_p value for profiles with a non-optimal C_l trend.

5.3.2 Pitching-Moment limitation

Since no structural analysis is performed in the current framework, a simplified constraint was introduced to prevent the optimizer from selecting excessively cambered airfoil shapes that would result in large torsional loads on the blade. A hard cap was imposed on the pitching-moment coefficient C_m at the optimal angle of attack: any candidate whose C_m falls below a user-defined threshold is rejected by assigning a large negative objective value, effectively removing it from the design space. Separate thresholds are defined for the tip and root airfoils, with the root cap set slightly more permissive to accommodate the thicker, more cambered sections typical of inboard stations.

In the absence of a coupled aero-structural model, the specific values of these thresholds are chosen empirically rather than derived from a structural load analysis. Nevertheless, the implementation is fully parametric: adjusting the caps in the input file is sufficient to adapt the constraint to a different rotor or to incorporate results from a future structural study, with no changes required in the optimization logic.

5.4 Geometrical constraints

In the NACA 4-digit optimization described in Section 5.1, the design space consisted of only two free variables, each bounded within a compact interval. Moving to a CST-parametrized airfoil substantially increases the dimensionality of the problem: as the number of coefficients grows, the space of admissible shapes expands exponentially, making any exhaustive exploration computationally intractable.

To illustrate this, Figure 5.8 shows two CST fits of the same reference airfoil using 10 coefficients and 2 coefficients per surface, respectively. A higher coefficient count yields a more faithful reconstruction of complex shapes, but at the cost of a much larger search space.

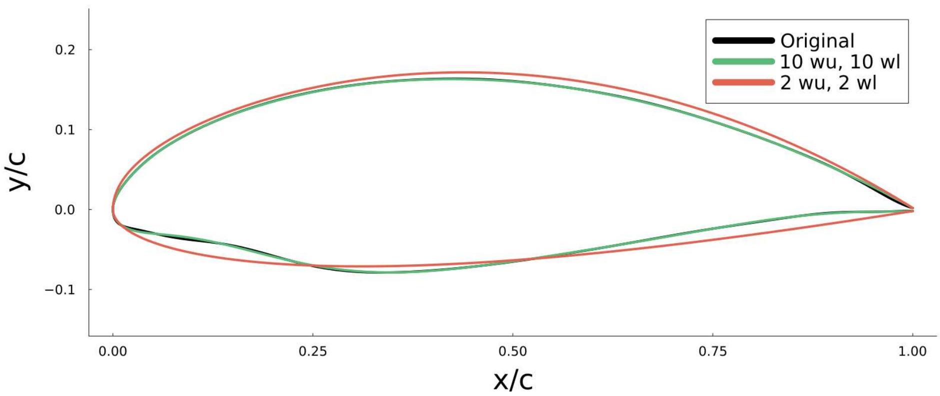

Figure 5.8: CST interpolation of a reference airfoil using 10 coefficients and 2 coefficients per surface.

Choosing the right number of coefficients is therefore a trade-off between expressiveness and tractability. From the analysis of a representative set of wind turbine airfoils, it was observed that the upper surface typically follows a smooth, nearly semi-parabolic profile, while the lower surface often exhibits more complex curvature changes with one or more inflection points. For this reason, the optimization assigns 3 coefficients to the upper surface ($\mathbf{w}_u$) and 4 to the lower surface ($\mathbf{w}_l$), yielding 7 degrees of freedom in total. This asymmetric parametrization allows

greater flexibility on the pressure side, where the aerodynamic behaviour is more sensitive to shape variations, without unnecessarily inflating the design space.
The selection was guided by systematic interpolation tests on representative wind turbine profiles. On the upper surface, 2 coefficients produced an overly simplified shape unable to capture the smooth curvature distribution near the leading edge, while 3 coefficients proved sufficient to reproduce the nearly semi-parabolic profile with good fidelity. On the lower surface, 3 coefficients could not reproduce the concave region typically present near the leading edge of cambered airfoils, resulting in a flattened pressure side that deviates significantly from the target geometry. Using 5 or more lower-surface coefficients improved the fit only marginally while increasing the design space from $80^7 \approx 2 \times 10^{13}$ to $80^9 \approx 10^{17}$ candidate configurations. The selected count of 4 lower-surface coefficients provides enough flexibility to capture one to two inflection points, which is consistent with the curvature patterns observed in the reference profiles.
Figure 5.9 shows an example of the resulting CST fit.

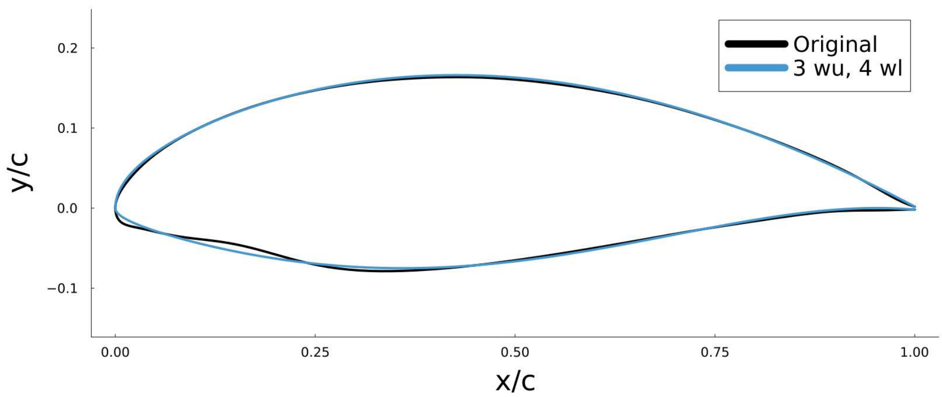

Figure 5.9: CST parametrization adopted in the optimization: 3 upper-surface coefficients (w_u) and 4 lower-surface coefficients (w_l).

With a perturbation range of $\Delta = \pm 0.4$ around the reference airfoil and a discretization step of 0.01, the resulting design space contains on the order of $80^7 \approx 2 \times 10^{13}$ candidate configurations. The introduction of geometric constraints is therefore essential to restrict the search space to physically meaningful and aerodynamically viable airfoil shapes. These constraints are organized into three layers, applied at different stages of the optimization pipeline.

Box constraints. The first layer consists of bound constraints on the individual CST coefficients. Each coefficient is allowed to vary within a symmetric window of amplitude Δ around its reference value:

$$w_{u,i} \in \left[\max(w_{u,i}^0 - \Delta,\, 0.05),\ \min(w_{u,i}^0 + \Delta,\, 0.8)\right], \tag{5.7}$$

$$w_{l,j} \in \left[\max(w_{l,j}^0 - \Delta,\, -0.8),\ \min(w_{l,j}^0 + \Delta,\, 0.8)\right]. \tag{5.8}$$

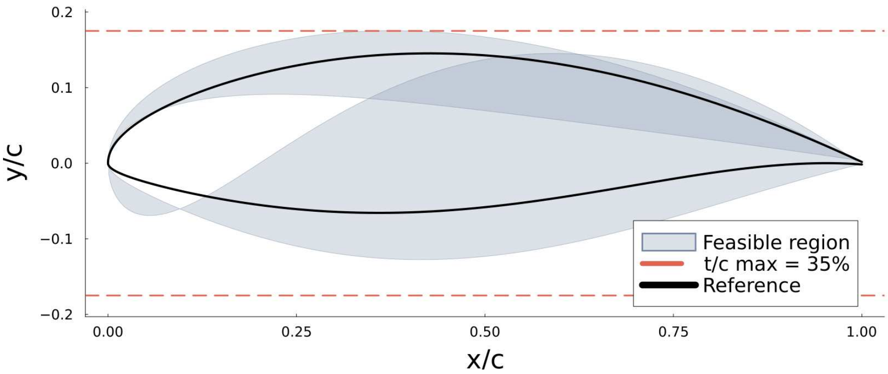

Figure 5.10: Feasible design space around the reference airfoil for $\Delta = \pm 0.3$. Dashed lines indicate the maximum thickness constraint.

where $w_{u,i}^0$ and $w_{l,j}^0$ are the reference coefficient values of the starting airfoil. In addition to the window bounds, absolute limits are enforced on the first coefficient of each surface: $w_{u,1} \geq 0.10$ and $w_{l,1} \leq -0.10$. The first CST coefficient governs the leading-edge curvature, and these limits ensure that the profile retains a well-defined, physically meaningful nose on both sides.

Nonlinear leading- and trailing-edge constraints. Two additional inequality constraints are passed directly to the NLopt solver. The first enforces that the leading-edge radius of the upper surface is at least as large as that of the lower surface:

$$w_{u,1} \ \geq\ |w_{l,1}|. \tag{5.9}$$

This prevents the airfoil from becoming concave near the nose, which would produce geometries that are both aerodynamically poor and unrealistic. The second constraint addresses the trailing edge:

$$w_{u,2} \ \geq\ |w_{l,n_{wl}}|, \tag{5.10}$$

where n_{wl} is the index of the last lower-surface coefficient. This condition is a practical heuristic to avoid excessively sharp or inverted trailing-edge geometries; it is not geometrically symmetric in a strict sense but has proved effective in keeping the optimizer within the region of feasible shapes.

Geometric feasibility filters (DOE phase). During the Latin Hypercube Sampling phase, an additional set of hard geometric checks is applied to each candidate airfoil before it is admitted to the pool. Any candidate that violates one or more of the following conditions is discarded:

(a) Maximum thickness $t_{\text{max}} \leq 35\%$ of the chord. Prevents physically unreasonable shapes with excessively blunt cross-sections.

(b) Maximum camber $f_{\text{max}} \leq 15\%$ of the chord. Avoids highly cambered profiles that would be structurally and aerodynamically impractical.

(c) Minimum aft thickness at $x/c = 0.90$. Ensures a sufficient cross-section near the trailing edge, compatible with manufacturing tolerances.

(d) No surface crossover. The upper surface must lie strictly above the lower surface at every chord station, guaranteeing a physically closed airfoil.

These filters are applied only at the DOE stage and not enforced as explicit constraints inside the COBYLA loop. Evaluating them at every function call would roughly triple the computational cost per iteration, since each check requires a full geometric reconstruction of the airfoil. The box bounds and the nonlinear NLopt constraints are sufficient to keep the optimizer away from extreme violations during the local search, while the hard filters ensure that all starting points fed to COBYLA are geometrically valid.

6 Results

This chapter presents the results obtained with the optimisation framework described previously. The analysis is divided into two parts. First, the BEM-based solver is validated against the experimental data from the NTNU Blind Test 1 campaign, comparing the predicted power and thrust coefficients with wind-tunnel measurements over a wide range of tip-speed ratios. This step verifies that the developed tool, produces reliable results, in particular, near the design operating conditions since, relying on panel methods and BEM theory it is expected to fail far from the design point, in the deep stall regimes. Second, the full optimisation framework is applied to a representative turbine configuration to assess the improvements that can be achieved through airfoil shape optimisation.

6.1 Validation NTNU 1

To check the reliability of the baseline framework, the test environment of the NTNU blindtest 1 is replicated and the results compared. The geometry and the experiment conditions are taken from the report [21].

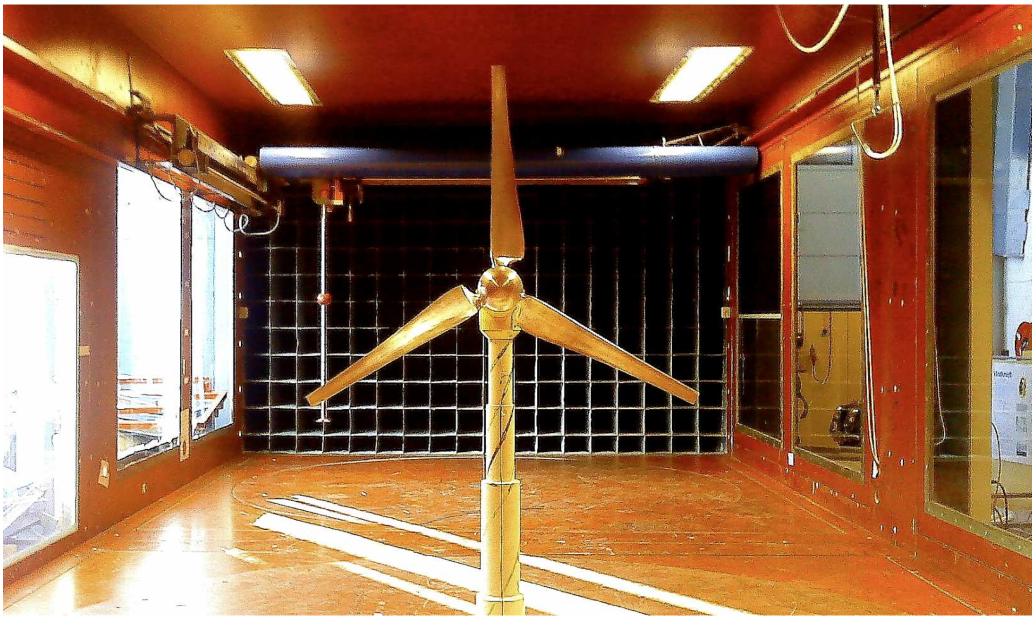

Figure 6.1: Picture of the real turbine in the wind tunnel.

Inputs

Tunnel and Turbine. The test is about a small rotor with 3 blades, the nacelle has a diameter $d = 90mm$ while the rotor diameter $D = 0.894m$. The tunnel has a width section of 2.17m and a length of 11.15m, the flexible roof has been adjusted to zero pressure gradient at a bulk tunnel speed of approximately U = 14m/s. The turbine is located at mid height and a distance from the inlet of 3.66m The height of the tunnel slightly vary along the length (from 1.801 to 1.851), the value in correspondence of the turbine location (1.801m) has been chosen.

Airfoil. The airfoil used is the NREL s826, to avoid any discrepancy given from CST approximation, the airfoil coordinates reported on the documentation have been used.

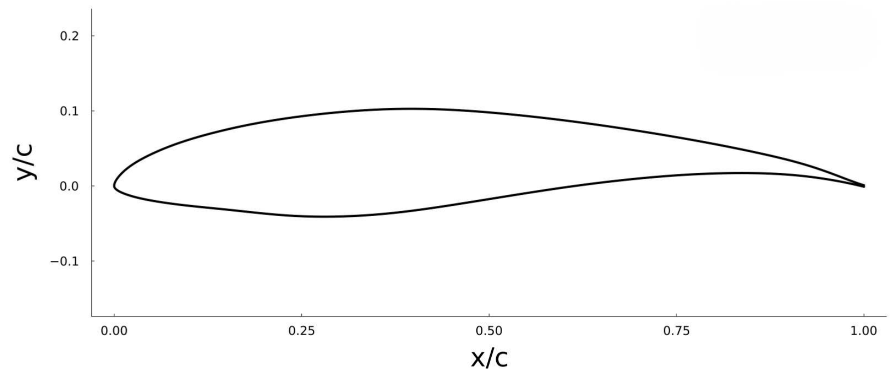

Figure 6.2: Representation of NREL s826 airfoil used in the test.

Blade Geometry. In the report are also provided chord and twist data for each section and they have been hardcoded to avoid any difference as done for the airfoil. In the calculation the blade taping to the hub has been neglected, in the real model the last section close to the hub are circular, while in our simulation the airfoil shape is keep untill the hub.

Test Conditions. As specified in [21], the reference wind speed is $U_{\text{ref}} = 10.0$ m/s and the air density $\rho = 1.2$ kg/m^3. The design tip speed ratio is

$$\lambda = \frac{\Omega R}{U_{\text{ref}}} = 6, \tag{6.1}$$

where Ω is the rotor angular velocity and R the tip radius. This gives a tip chord Reynolds number of $Re_c = \lambda\, U_{\text{ref}}\, c_{\text{tip}}/\nu \approx 103\,600$, where c_{tip} is the tip chord length and ν is the kinematic viscosity of air.

Blockage Correction. During tunnel simulations, the results do not represent the real operating condition in a free air stream, the walls constrain the flowing air and so the streamtube cannot expand as it would do in normal conditions. It is working in a way that reminds of a shrouded rotor. This leads to an increase in the performances because of an higher inflow velocity seen by the rotor. This phenomenon is called ”Blockage Effect”. Mikkelsen [22] gives a method to apply the blockage correction to a rotor disc inside a constant cross-section tunnel. Considering uniform velocity V_0 and continuity and incompressibility:

$$\begin{aligned} u_1 S_1 &= uS, \\ u_2(C - S_1) &= V_o C - uS. \end{aligned} \tag{6.2}$$

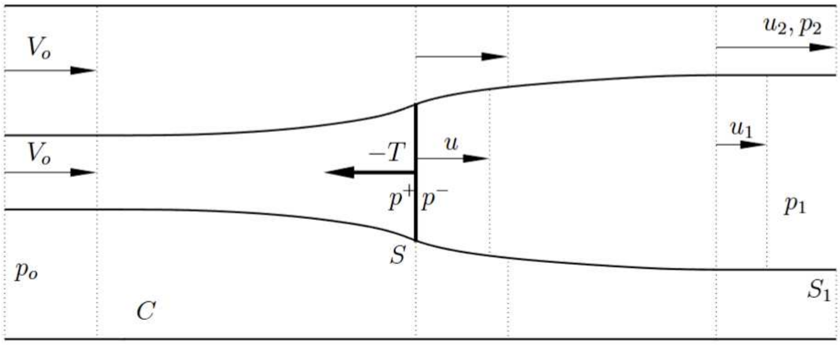

Figure 6.3: Rotor disc in a tunnel. V_0, u represent the velocities; C, S, S1 are the areas and p is the pressure.

where u is the velocity at the rotor disc, u_1, u_2 are, respectively, the velocities in the far wake inside and outside the streamtube. The disc area is S, the total tunnel cross-sectional area is C and S_1 is the streamtube area in the far wake, defined by the limiting streamline passing through the disc edge. Outside the streamtube the total pressure head is conserved:

$$p_o + \frac{1}{2}\rho V_o^2 = p_2 + \frac{1}{2}\rho u_2^2. \tag{6.3}$$

Inside the streamtube, the total head decreases by the pressure drop across the disc:

$$p_o + \frac{1}{2}\rho V_o^2 + p'_+ = p'_- + p_1 + \frac{1}{2}\rho u_1^2, \tag{6.4}$$

with $p_1 = p_2$. The quantities (p'_+, p'_-) denote the pressure jump across the disc, which, as seen in Section 2.2 is proportional to the thrust: $T/S = p'_+ - p'_-$.

Applying Bernoulli and the thrust definition:

$$T = \frac{1}{2}\rho S\left(u_1^2 - u_2^2\right), \tag{6.5}$$

$$\Delta p = p_1 - p_o = \frac{1}{2}\rho\left(V_o^2 - u_2^2\right). \tag{6.6}$$

Applying the momentum equation over the full tunnel cross-section:

$$T - (p_1 - p_o)C = \rho u_1 S_1(u_1 - V_o) - \rho u_2(C - S_1)(V_o - u_2). \tag{6.7}$$

This gives a system of five equations in five unknowns: u, u_1, u_2, S_1, p_1. After non-dimensionalization with:

$$\sigma = \frac{S_1}{S} \tag{6.8}$$

$$\alpha = \frac{S}{C} \tag{6.9}$$

The system reduces to a single expression for the disc velocity:

$$u = \frac{\sigma(\alpha\sigma^2 - 1)}{[\sigma\alpha(3\sigma - 2) - 2\sigma + 1]} \tag{6.10}$$

and so the equivalnet free stream air velocity V'/V_0 is given by:

$$\frac{V'}{V_0} = u - \frac{1}{4}\left(\frac{C_T}{u}\right) \tag{6.11}$$

Since Equation (6.11) requires knowledge of the slipstream ratio σ, Glauert (1934) derived a simpler approximate expression depending only on C_T and the tunnel geometry:

$$\frac{V'}{V_o} = 1 - \frac{1}{4}\left(\frac{R_{\text{rot}}}{R_{\text{tun}}}\right)^2 \frac{C_T}{\sqrt{1 + C_T}}. \tag{6.12}$$

Mikkelsen [22] validates this expression against both the exact momentum solution and Navier-Stokes actuator disc computations. For a constant-loaded rotor at $C_T = 0.8$ and $R_{\text{tun}}/R_{\text{rot}} = 2.5$, Equation (6.12) agrees with the exact solution to within 2–4 percentage points, which is acceptable for engineering purposes. The expression becomes singular at $C_T = 1$, however for the NTNU BT1 operating range ($C_T \approx 0.6$–1.0) the error remains within acceptable bounds.

The corrected performance coefficients are recovered by scaling with the ratio $v_r = V'/V_o$:

$$C_{P,\text{tun}} = C_{P,\text{free}}\, v_r^3, \tag{6.13}$$

$$C_{T,\text{tun}} = C_{T,\text{free}}\, v_r^2. \tag{6.14}$$

Results

The BEM model described in the previous sections was evaluated over a range of tip-speed ratios from $\lambda = 1$ to $\lambda = 12$, both in free-air conditions and with the Glauert blockage correction applied. Figure 6.4 reports the computed aerodynamic coefficients at different values of λ. The corresponding power and thrust coefficient curves are compared against the experimental data from the NTNU Blind Test 1 campaign [21]. In free-air conditions the model predicts a peak power coefficient of $C_{P,\,\text{max}} \approx 0.43$ near $\lambda \approx 5$; when the blockage correction is included, the peak rises to $C_{P,\,\text{max}} \approx 0.46$ at a similar tip-speed ratio, in closer agreement with the experimental value of $C_{P,\,\text{max}} \approx 0.45$ around $\lambda = 6$. The thrust coefficient follows a monotonically increasing trend with λ in the experimental datasets while for the numerical results the thrust coefficient experiences a decrease in the increment leading to a descending trend after $\lambda \approx 10$.

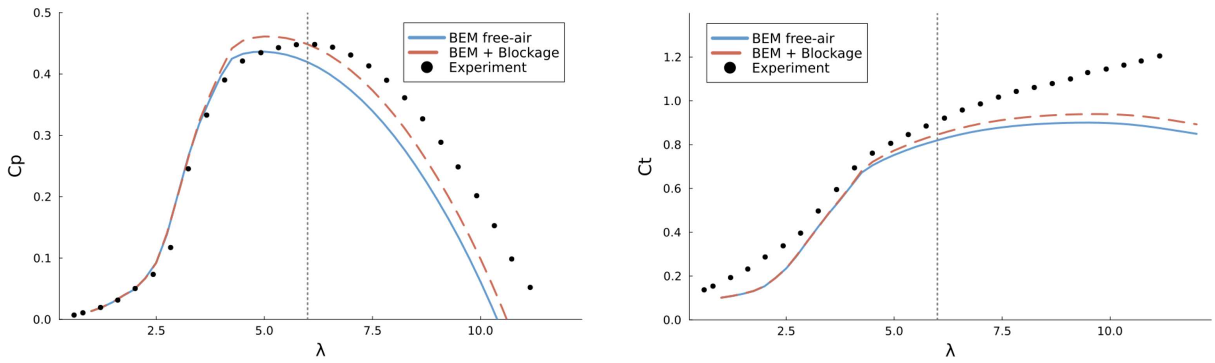

Figure 6.4: Power and thrust coefficients versus tip-speed ratio: BEM free-air, BEM with Glauert blockage correction, and experimental data from the NTNU Blind Test 1 campaign.

The comparison between the BEM predictions and the experimental data reveals a clear pattern. Overall, the match is satisfactory in the range $3 < \lambda < 7$, which covers the design operating region of the turbine. At the design point $\lambda = 6$, the blockage-corrected C_P matches the experimental value almost exactly. At higher tip-speed ratios, however, the model progressively underestimates both C_P and C_T, with the discrepancy growing as λ increases.

Reynolds number sensitivity. During the experimental data validation the dual section with interpolation in between carried out in the main code hasn't been implemented, this mainly due to model accuracy, since airfoil geometry, chord distribution and twist distribution are fixed, following the real model geometry. For this reason the aerodynamic polars used in the BEM calculation are evaluated at a fixed Reynolds number corresponding to the design tip-speed ratio. This approach implicitly assumes that the Reynolds number distribution along the blade does not change with the operating conditions. In practice, however, the local Reynolds number varies substantially along the span. Table 6.1 reports the estimated chord-based Reynolds number at the innermost and outermost airfoil sections for integer values of λ. The relative velocity at each section is computed using the iterative procedure described previously, accounting for the axial and tangential induction factors:

$$W(r) = \sqrt{\left[V_\infty(1-a)\right]^2 + \left[\omega\, r\,(1+a')\right]^2}, \qquad \mathrm{Re}(r) = \frac{\rho\, W(r)\, c(r)}{\mu}. \tag{6.15}$$

The table shows that the tip Reynolds number increases significantly between $\lambda = 1$ and $\lambda = 12$, reaching approximately 2.1×10^5 at the highest tip-speed ratio. Since increasing Reynolds number leads to higher maximum lift coefficients and lower drag coefficients, using polars evaluated at a single operating condition systematically penalises the predicted performance at high λ, where the outer sections, which experience the highest Reynolds numbers, dominate the integrated loads.

λ	W_{root} [m/s]	$\mathrm{Re}_{\text{root}}$	W_{tip} [m/s]	Re_{tip}
1	8.9	29 600	13.6	23 600
2	9.3	30 900	21.9	38 100
3	9.8	32 500	31.2	54 200
4	10.4	34 500	40.7	70 800
5	11.1	36 800	50.4	87 600
6	11.8	39 300	60.1	104 600
7	12.7	42 000	69.9	121 600
8	13.5	45 000	79.8	138 700
9	14.5	48 100	89.6	155 800
10	15.5	51 300	99.4	172 900
11	16.5	54 600	109.3	190 100
12	17.5	58 000	119.2	207 200

Table 6.1: Chord-based Reynolds number at the root ($r = 55$ mm, $c = 49.5$ mm) and tip ($r = 442.5$ mm, $c = 25.9$ mm) airfoil sections for different tip-speed ratios.

Should be noted that at the design tip speed ratio's Reynolds number, the evaluated power coefficient perfectly match the experimental datas, underlying that the missmatch involves the off design operating point. A fully Reynolds-aware approach, in which CCBlade interpolates between polars at multiple Reynolds numbers at every BEM iteration, would allow the solver to select the appropriate polar for the local flow conditions at each section and each tip-speed ratio, and would alleviate this limitation.
This limitation affects the validation predictions more than the optimisation results. The optimiser seeks the airfoil shape that maximises C_P at the design point ($\lambda \approx 6$), and the optimal geometry is largely insensitive to small variations in the polar Reynolds number: a profile with high C_l/C_d at Re $= 100\,000$ retains its aerodynamic advantage at neighbouring Reynolds numbers. The absolute C_P value may carry a small bias, but the ranking between candidate shapes, and therefore the optimisation outcome, remains reliable.

Simplified blockage correction. The wind-tunnel blockage correction implemented in the present work follows the classical Glauert model [22], which treats the rotor as an actuator disc in a circular tunnel of equivalent cross-sectional area. The NTNU wind tunnel, however, has a rectangular test section of 2.71×1.80 m, resulting in a geometric blockage ratio of approximately 13%. At such a non-negligible blockage level, several simplifying assumptions of the Glauert model become questionable: the velocity distribution across the rotor plane is not uniform, and the wall confinement effects in a rectangular section differ from those in an equivalent circular duct. Furthermore, at high tip-speed ratios the thrust coefficient becomes large, and the simple one-dimensional Glauert formula is no longer accurate enough to capture the real blockage effect.

Three-dimensional effects. The BEM method relies on two-dimensional airfoil polars and treats each annular element as independent, neglecting rotational augmentation effects (Himmelskamp effect) that increase the effective lift on the inboard sections of a rotating blade [23]. To assess the impact of this simplification, two stall-delay correction models were tested: the Raj model (applied as a pre-processing step to the polars) and the Du-Selig/Eggers correction built into CCBlade (applied on-the-fly during the BEM iterations). In both cases the agreement with the experimental data worsened, with the corrected curves underpredicting C_P across the full λ range. This is attributed to the fact that the empirical constants of these models were calibrated on full-scale turbines (NREL Phase VI, $D \approx 10$ m), where the chord-to-radius ratios are significantly larger and the rotational effects more pronounced. On the present laboratory-scale rotor ($D = 0.894$ m), the 3D corrections are too aggressive and were therefore not retained in the final validation.

An interesting detail emerging from the results is that the underestimation in C_T is proportionally larger than that in C_P at high tip-speed ratios. This can be explained by noting that C_T depends almost entirely on the lift force projected along the rotor axis, so any error in the polar translates directly into an error in thrust. The power coefficient, by contrast, depends on the difference between the driving (lift) and resisting (drag) components of the aerodynamic force. Because lift and drag errors act in opposite directions on the torque, they partially cancel out, making C_P less sensitive to polar inaccuracies than C_T.
The good match in the range $3 < \lambda < 7$ confirms that the BEM model, in conjunction with the real S826 coordinates and the Glauert blockage correction, captures the essential physics of the rotor near its design point. The progressive divergence at higher tip-speed ratios is the cumulative result of well-known limitations as seen above.

6.2 Scenario optimisation

Wind turbines in the sub-200 kW class find their most compelling application in the autonomous electrification of remote communities. With a total annual supply that falls squarely within the reach of small wind technology. It is precisely this context that motivates the present scenario: the preliminary aerodynamic design of a wind turbine sized for rural village energy supply.
The optimisation framework is applied to the design of a small wind turbine for rural village energy supply. A typical off-grid settlement of 20–30 households consuming 3 000–5 000 kWh/yr each requires 60–150 MWh/yr; with a capacity factor of 0.30–0.40, the required rated power is approximately 20–55 kW. A minimum aerodynamic target of 25 kW at the design point with a blade length of approximately 6 meters is adopted. The study focuses exclusively on aerodynamic optimisation; structural analysis, aeroelastic simulation, and IEC 61400-2 certification are outside its scope.

6.2.1 Site data and design choices

Given the availability of a wind measurement dataset from a site in Turkey, the choice of location follows naturally. The anemometric dataset consists of $n = 50\,520$ valid 10-minute averaged wind speed records from the site in Turkey, assumed to be measured at the standard WMO height of $z_{\text{ref}} = 10\,\text{m}$. Table 6.2 summarises the Weibull fit and the derived design speeds (see Section 2.5 for the underlying definitions).

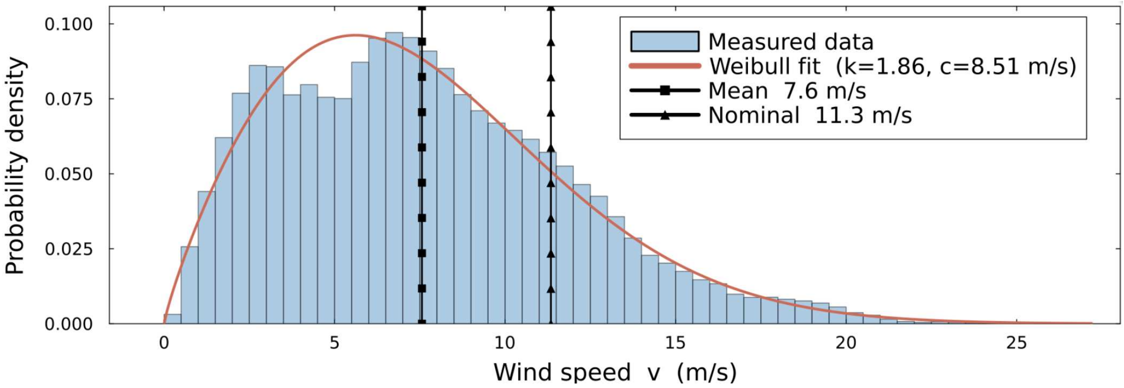

Figure 6.5: Weibull distribution fit to the measured wind speed data. Vertical lines indicate the arithmetic mean and the rated wind speed.

Quantity	**Value**	**Unit**
Shape parameter k	1.857	–
Scale parameter c	8.515	m/s
Mean wind speed $\bar{v}$	7.56	m/s
Standard deviation σ	4.23	m/s
Maximum recorded speed	25.2	m/s
Most probable speed v_d	5.61	m/s
Rated speed $v_r = 1.5\bar{v}$	11.34	m/s

Table 6.2: Weibull fit results and derived design speeds for the Turkish site.

Design wind speed. Designing at the most probable speed $v_d = 5.61$ m/s would require a rotor radius of approximately 14 m, far exceeding the size target. Designing strictly at v_r would concentrate efficiency at the upper end of the operating range. As a compromise, an aerodynamic design point of $v^* = 10$ m/s is adopted, keeping the rotor radius consistent with the 25 kW target while limiting the efficiency penalty at intermediate wind speeds. Table 6.3 quantifies the rotor sizing trade-off for the three candidate speeds, assuming $C_p = 0.40$ and drivetrain efficiency $\eta = 0.90$ (neglecting the wind shear factor).

Design point	**Wind speed**	**Required R**
Most probable v_d	5.61 m/s	∼14 m
Compromise v^*	10.0 m/s	∼6 m
Rated $v_r = 1.5\bar{v}$	11.34 m/s	∼5.3 m

Table 6.3: Rotor radius required for 25 kW at different design wind speeds ($C_p = 0.40$, $\eta = 0.90$).

Rotor sizing and hub height. A blade radius of $R = 6\,\text{m}$ satisfies the power requirement at $v^* = 10\,\text{m/s}$. Following the rule of thumb $z_{\text{hub}} \approx 2R$ [1], the hub height is set to $z_{\text{hub}} = 12\,\text{m}$. The logarithmic height correction from 10 m to 12 m (with roughness length $z_0 = 0.05\,\text{m}$ for open rural terrain) yields a factor of approximately 1.03, which is considered negligible; the measured speeds are taken as representative of hub-height conditions.

Wind shear. The power-law profile with $\alpha_s = 0.2$ is used in CCBlade to account for the variation in wind speed across the rotor disk. With $D/z_{\text{hub}} = 1.0$, the lower blade tip passes at 6 m above ground, where the wind speed is approximately 13 % lower than at hub height; at the upper tip (18 m) it is about 7 % higher, giving a total variation of roughly 20 % across the disk.

Turbulence. Freestream turbulence is characterised indirectly through the XFOIL critical amplification factor $N_{\text{crit}} = 6$, corresponding to a turbulence intensity of approximately 1 %, consistent with a moderately turbulent rural environment.

Control strategy. A variable-speed fixed-pitch (VS-FP) configuration is adopted. Below v^* the rotor speed tracks λ_{opt} for maximum C_p; above v^* the speed is held constant and the blade sections progressively stall, passively limiting the power. This avoids the cost and complexity of a pitch actuator, which is advantageous in the sub-200 kW class.

6.2.2 Geometry setup

The aerodynamic design of the rotor blade is carried out in two sequential stages. First, a Kriging surrogate model explores the NACA 4-digit design space to identify a well-performing starting airfoil for both the root and tip sections. Then, a gradient-free COBYLA optimiser refines the blade shape using a CST parametrisation. All stages share the common set of turbine parameters listed in Table 6.4.

Tip-speed ratio. A design tip-speed ratio $\lambda_{\text{opt}} = 6$ is selected, consistent with the lower end of the range typical for three-blade rotors ($\lambda = 6$–8). The corresponding tip speed at the design point is $\Omega R = \lambda_{\text{opt}}\, v^* = 60\,\text{m/s}$, and the design rotational speed is $\Omega = 10\,\text{rad/s}$ ($\approx 95\,\text{rpm}$). Choosing $\lambda_{\text{opt}} = 6$ rather than a higher value reduces tip-noise emission and centrifugal loads, which is advantageous for a small rural installation.

Thickness distribution. The tip section thickness ratio is set to $(t/c)_{\text{tip}} = 0.13$, a value that offers a good compromise between low aerodynamic drag and sufficient structural stiffness. The root section, which must carry the full bending moment of the blade, is assigned a thickness ratio $(t/c)_{\text{root}} = 1.8 \times (t/c)_{\text{tip}} \approx 0.234$. The spanwise interpolation follows a power law with exponent $n_{t/c} = 1.5$, producing a rapid thickness increase near the hub where structural demands are highest.

Parameter	**Symbol**	**Value**	**Unit**
Operating conditions			
Design wind speed	v^*	10.0	m/s
Air density	ρ	1.225	kg/m^3
Rotor geometry			
Number of blades	B	3	–
Blade radius	R	6.0	m
Hub radius	R_{hub}	0.35	m
Hub height	z_{hub}	12.0	m
Precone / tilt / yaw	–	0	deg
Aerodynamic design			
Design tip-speed ratio	λ_{opt}	6	–
Twist law	–	Betz ($\Delta\alpha = 0$)	–
Root chord (hub)	c_{hub}	0.26	m
Root taper extent	η_{taper}	0.07	–
Airfoil thickness			
Tip section t/c	$(t/c)_{\text{tip}}$	0.13	–
Root section t/c	$(t/c)_{\text{root}}$	0.234	–
Spanwise exponent	$n_{t/c}$	1.5	–
Flow solver			
XFOIL N_{crit}	–	6	–
BEM radial sections	–	40	–

Table 6.4: Turbine design parameters used in the blade optimisation.

Root taper. A linear taper is applied over the innermost $\eta_{\text{taper}} = 0.07$ of the span, blending the Betz chord smoothly down to a fixed cylindrical hub profile of $c_{\text{hub}} = 0.26\,\text{m}$. This region contributes negligibly to the aerodynamic torque but is critical under structural aspects.

Reynolds number interpolation. The aerodynamic polars are not evaluated at a single global Reynolds number. Instead, two section-specific Reynolds numbers are computed iteratively at the root ($\eta = 0.40$) and tip ($\eta = 0.90$) stations, accounting for the local chord and relative velocity through the procedure described in Section 4.2. The resulting values are used for the respective XFOIL evaluations, and the Reynolds number at intermediate blade sections is obtained by linear interpolation, without requiring a separate XFOIL run for each of the 40 radial stations.

6.2.3 Airfoil selection via Kriging surrogate

The starting airfoil geometry for both the root and tip sections is determined through a Kriging surrogate model applied to the NACA 4-digit design space. The two free parameters are the maximum camber and its chordwise position; the thickness is fixed at the values prescribed in Table 6.4 ($t/c = 0.234$ at the root, $t/c = 0.13$ at the tip). Because this stage analyses a single airfoil in isolation rather than a complete rotor, the objective function is the maximum lift-to-drag ratio C_l/C_d, optimising directly for the rotor power coefficient C_p would require a full BEM evaluation and would implicitly assume a single repeated airfoil along the entire span, which is not representative of the dual-airfoil design adopted here.

The surrogate is built from an initial design of experiments (DOE) of 20 samples made from a space-filling Latin Hypercube design over the parameter ranges camber $\in [2\%, 9\%]$ and camber position $\in [2, 8]$. The model is then refined adaptively: at each iteration the point of maximum surrogate prediction error is evaluated and added to the dataset, until the maximum error falls below a convergence threshold of 0.5% of the mean C_l/C_d of the current dataset.

Table 6.5 summarises the outcome for both sections.

Section	**Iterations**	**Tolerance**	$(C_l/C_d)_{\max}$	t/c	Re	**Selected profile**
Root	43	0.105	79.2	0.234	6.85×10^5	NACA(8.65)(5.73)(23.4)
Tip	55	0.514	203.2	0.130	1.40×10^6	NACA(8.54)(6.33)(13.0)

Table 6.5: Kriging airfoil selection results for the root and tip sections.

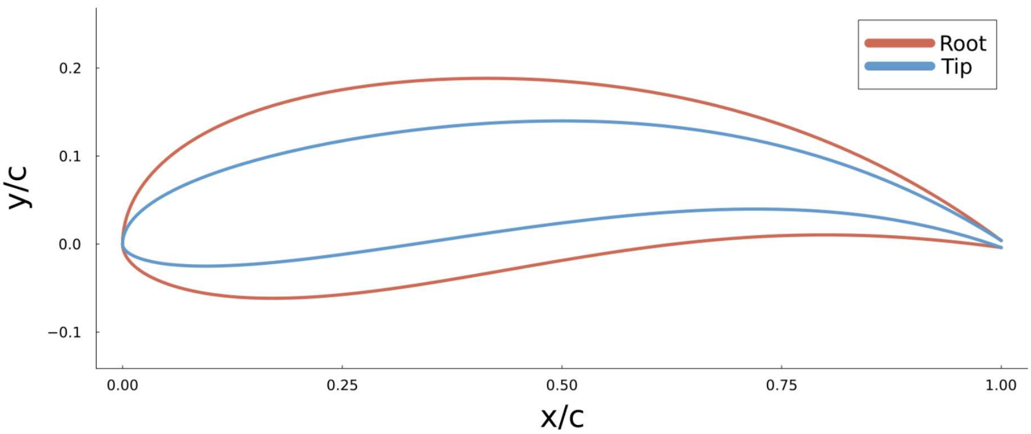

Figure 6.6: Optimal NACA profiles selected by the Kriging surrogate for the root section and the tip section.

The tip section achieves a significantly higher C_l/C_d than the root (203.2 vs. 79.2), which is expected for two reasons. First, the tip operates at a Reynolds number roughly twice as large, so viscous losses are proportionally smaller. Second, the root section is constrained to a much thicker profile to satisfy structural requirements, and thicker airfoils inherently produce higher drag. Both selected profiles converge to similar camber geometries (camber $\approx$ 8.5–8.6%, position $\approx$ 5.7–6.3%), suggesting that this region of the NACA 4-digit space is aerodynamically favourable at both Reynolds numbers. The high cambered profile are delicate in terms of structural reliability as they lead to high pitching moment, it will be the CST optimizer's task to limit this phenomenon.

Before proceeding to the CST shape optimisation, the two selected NACA profiles are assembled into the full dual-airfoil blade and evaluated with the BEM solver to establish a reference performance. The root polar is computed at $\alpha = 0.0°$ and the tip polar at $\alpha = 0.5°$, corresponding to the operating points that maximise C_l/C_d at the respective Reynolds numbers. Table 6.6 reports the resulting rotor performance; this constitutes the baseline against which the CST-optimised blade is subsequently compared.

Quantity	**Root**	**Tip**
Operating α	$0.0°$	$0.5°$
C_l/C_d	62.3	191.5
C_m	-0.246	-0.333
Rotor		
$C_{p,\text{max}}$	0.484	
λ_{opt}	6.68	
Power at v^*	38.3 kW	

Table 6.6: Kriging airfoil selection results for the root and tip sections.

6.2.4 CST shape optimisation

Starting from the obtained blade, the shape is refined by a gradient-free COBYLA optimiser acting on the CST coefficients of each section. The two sections are treated sequentially. Each phase runs a Latin Hypercube DOE of 60 samples, retains the top 4 starting points, and launches up to 4 COBYLA restarts of at most 100 evaluations each this limit is hardcoded to avoid loops in the process, the algorithms exit the optimization once achieved the convergence tolerance the majorityof the times. The objective function is $C_{p,\text{max}}$ evaluated by the full BEM solver; a hard penalty on the pitching moment ($C_m tip > -0.19 \quad and \quad c_m root > -0.22$) and a local stall penalty is applied to discourage geometries with excessive moment or separated flow. Convergence tolerances are $f_{\text{tol}} = x_{\text{tol}} = 10^{-3}$ Section 5.2.2. Table 6.7 summarises the optimiser performances.

Phase	DOE	Restarts	Total evals	Best $C_{p,\max}$	Stop criterion
Tip	60	4	460	0.4905	`FTOL_REACHED`
Root	60	4	448	0.4913	`FTOL_REACHED`
Total wall time			0 h 46 min 10 s		

Table 6.7: COBYLA optimiser performances for the two sequential phases.

Power coefficient improvement. The sequential optimisation raises $C_{p,\max}$ from 0.484 (NACA baseline) to 0.491, a gain of $+1.51\,\%$ corresponding to an additional 1.7 kW at the design wind speed.

Quantity	NACA baseline	CST optimised	Δ
$C_{p,\max}$	0.484	0.491	$+1.51\,\%$
C_T	0.792	0.802	$+1.3\,\%$
λ_{opt}	6.68	6.21	-0.47
Power at v^*	38.3 kW	40.0 kW	$+1.7$ kW

Table 6.8: Rotor performance comparison: NACA baseline vs. CST-optimised blade.

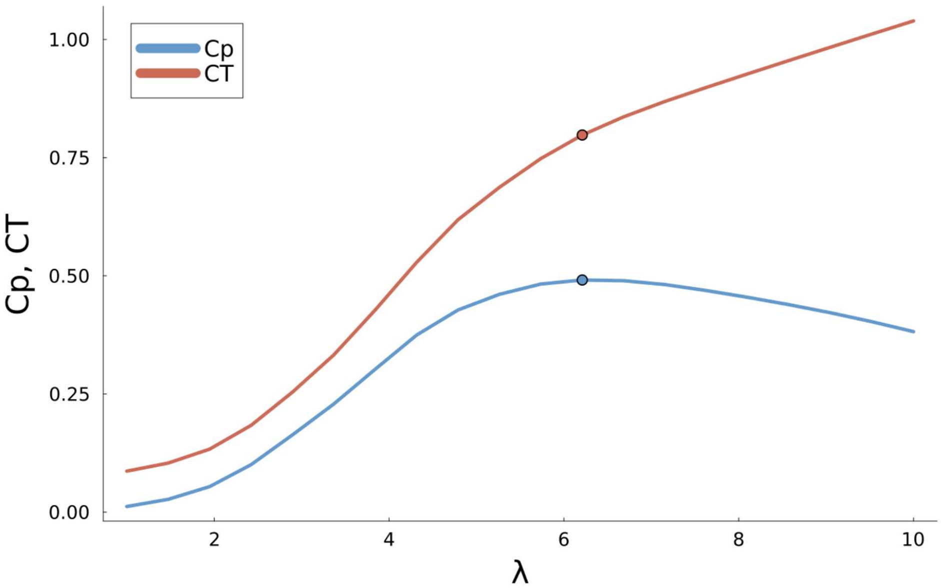

Figure 6.7: Power coefficient C_p and thrust coefficient C_T as functions of the tip-speed ratio λ for the CST-optimised blade.

Optimised section geometry. Figure 6.8 compares the NACA baseline and CST-optimised profiles for both the root and tip sections. The CST parametrisation (3 upper and 4 lower coefficients per section) allows departures from the NACA 4-digit family that are inaccessible to the two-parameter camber model: in particular, the optimiser is free to modify the camber line independently of the thickness envelope and to redistribute curvature chordwise in ways that NACA families cannot represent.

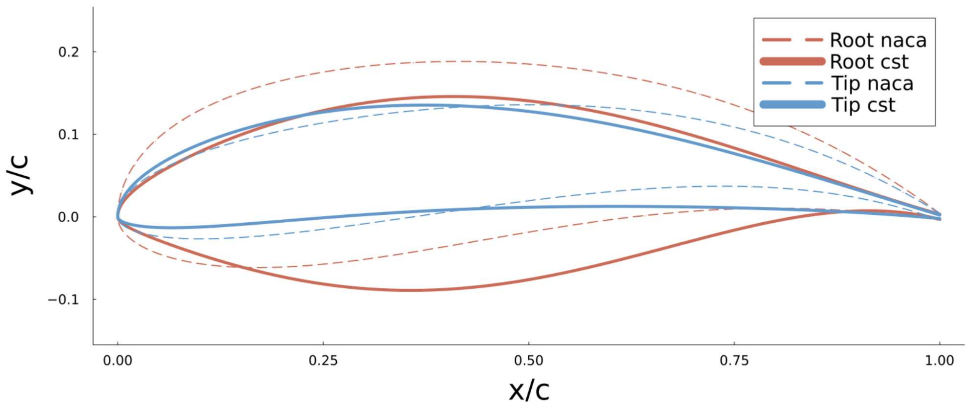

Figure 6.8: Root and tip airfoil geometry: NACA baseline (dashed) vs. CST-optimised profile (solid).

Chord and twist distribution. Figures 6.9 and 6.10 show the optimised blade planform. The chord reaches its maximum of approximately 0.85 m at $r \approx 0.79$ m from the axis and then tapers to zero at the tip. The innermost segment is governed by the structural root-taper constraint. The twist distribution decreases from $41.1°$ at the hub attachment to $1.8°$ at the blade tip, following the aerodynamic washout required to align each section with the relative wind at $\lambda = 6$.
A secondary but practically relevant effect of the CST optimisation is related to the chord distribution. Because the CST-optimised root achieves a C_l/C_d much closer to that of the tip (109.7 vs. 166.6, compared with 62.3 vs. 191.5 for the NACA baseline), the BEM-computed chord varies more gradually along the span: the root no longer requires an oversized chord to compensate for a disproportionately low aerodynamic efficiency. A more uniform chord taper brings several benefits. From a structural standpoint, smoother transitions in cross-sectional area reduce stresses and simplify the spar design. From a manufacturing perspective, a more regular planform leads to gentler mould curvatures and facilitates composite layup. Aerodynamically, a gradual chord variation attenuates the spanwise circulation gradients that drive trailing vorticity inboard, thereby reducing the induced-drag penalty associated with abrupt planform changes.

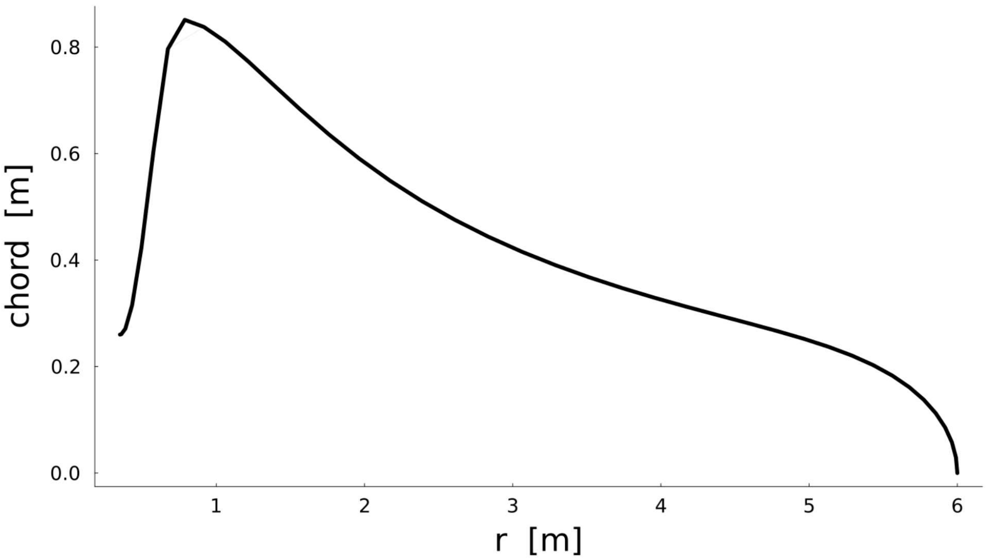

Figure 6.9: Chord distribution along the blade span following the optimal Betz chord distribution law.

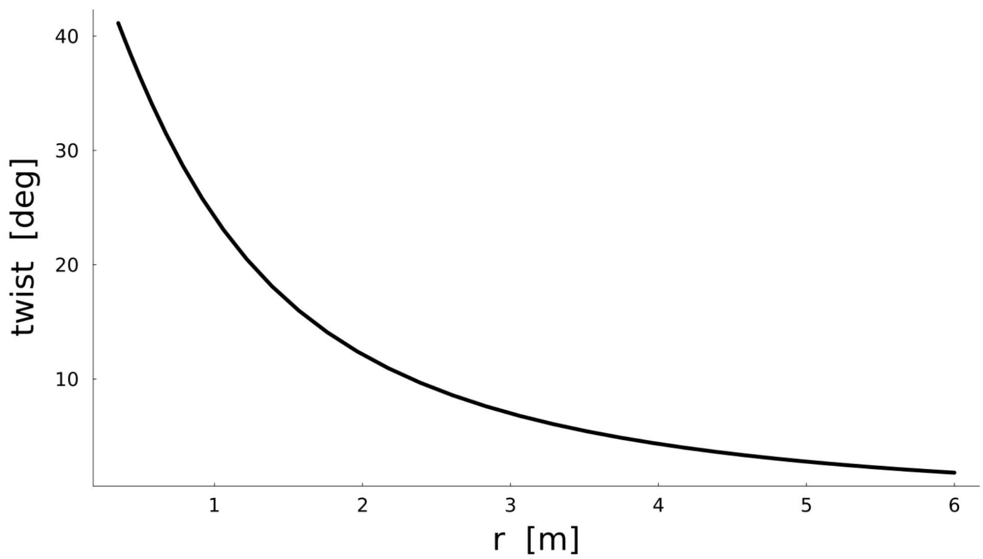

Figure 6.10: Twist distribution along the blade span following the optimal angle of attack along the blade span.

Lift polar comparison. Table 6.9 and Figure 6.11 compare the section aerodynamics before and after optimisation. The two sections respond differently to the CST reshaping.
At the *root*, the CST profile produces a notably smooth $C_l(\alpha)$ curve with a monotonic rise up to a clearly defined stall at $\alpha \approx 10°$, followed by a lift drop. This is a marked improvement over the NACA baseline, whose geometry yields a dashed curve that continues to climb without a well-defined maximum. The earlier but cleaner stall break of the CST root gives the BEM

solver a more physically consistent polar and places the operating point ($\alpha = 6.0°$) well within the linear regime, with a comfortable stall margin of roughly $4°$, necessary for a stall-regulated turbine.

At the *tip*, the CST polar shows a regular monotonic lift increase, compared with his NACA counterpart that clearly present a early separation point around $\alpha \approx 5°$ that causes local lift decrease, a phenomenon is preferable to avoid. The CST profile is characterised by a very gentle stall propagation as is possible to see from the gradual lift decrease after $\alpha \approx 10°$. However at $\alpha \approx 12°$ it is possible to notice a small depression in the lift curve, this may have been caused bya a laminar separation bubble formation. It would be very interesting to delve deeper into this phenomenon and see the performance of the airfoil using a zigzag trip strip.

	Root		**Tip**	
	NACA	CST	NACA	CST
α_{opt}	$0.0°$	$6.0°$	$0.5°$	$4.5°$
C_l	–	1.361	–	1.359
C_d	–	0.0124	–	0.0082
C_l/C_d	62.3	109.7	191.5	166.6
C_m	-0.246	-0.178	-0.333	-0.172

Table 6.9: Section aerodynamics: NACA baseline vs. CST-optimised blade.

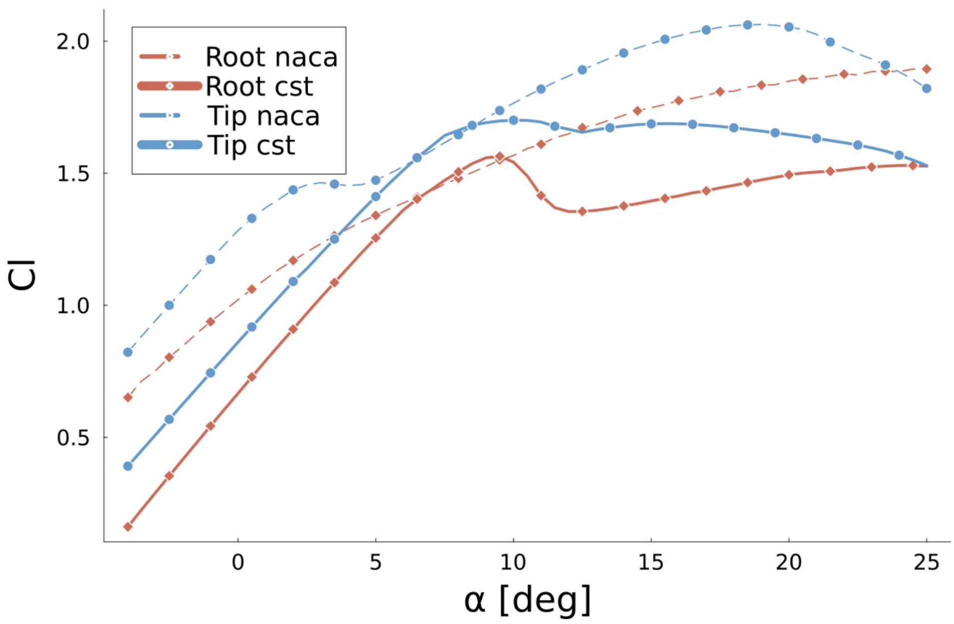

Figure 6.11: Lift coefficient $C_l(\alpha)$: NACA baseline (dashed) vs. CST-optimised profile (solid).

Pitching moment comparison. Figure 6.12 shows the $C_m(\alpha)$ polars for both sections. The pitching moment constraint ($C_m tip > -0.19 \quad and \quad c_m root > -0.22$) was active during the optimisation and has a pronounced effect: the CST profiles reach values of $C_m = -0.178$ (root) and $C_m = -0.172$ (tip) at their respective operating angles of attack, compared with -0.246 and -0.333 for the NACA baseline. The reduction in magnitude, roughly 28 % at the root and 48 % at the tip, is significant from a structural perspective: a lower (less negative) pitching moment reduces the torsional loading on the blade. The tip shows the largest improvement, consistent with the observation that the NACA tip was already operating largely above the constraint and the optimiser found a geometry that satisfies the same lift objective with lower torsional loading.

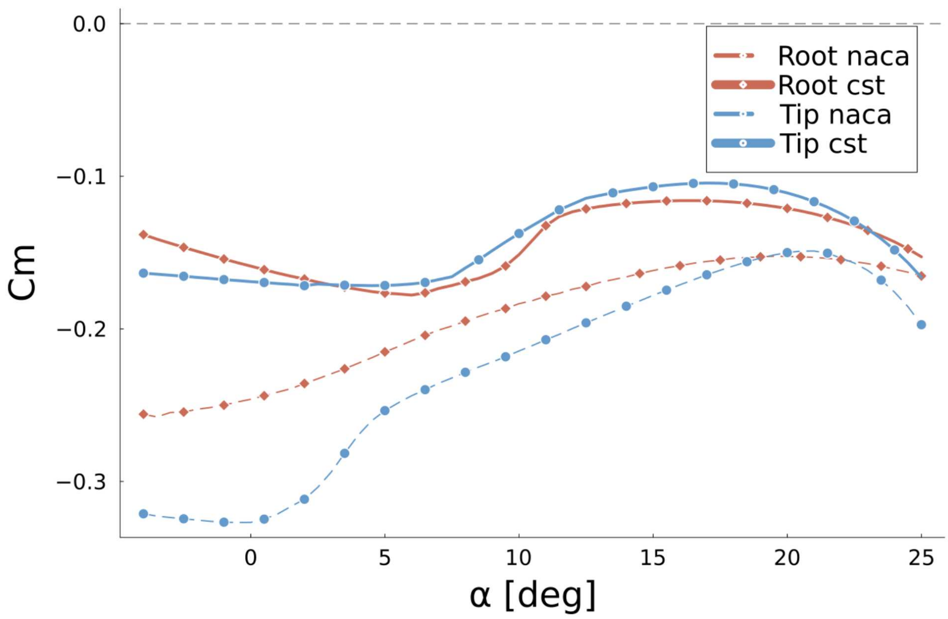

Figure 6.12: Pitching moment coefficient $C_m(\alpha)$: NACA baseline (dashed) vs. CST-optimised profile (solid).

Power curve and off-design behaviour. Figure 6.13 shows the power output as a function of wind speed for the VS-FP control strategy. Below $v^* = 10$ m/s, the rotational speed tracks $\lambda_{\text{opt}} = 6$, keeping C_p at its peak value of 0.491 and producing a cubic power curve. Above v^*, Ω is held constant at $\Omega_N = 10.0$ rad/s (95.5 rpm); the effective tip-speed ratio drops and the blade sections progressively enter aerodynamic stall, reducing C_p (Figure 6.14).

Despite the C_p reduction above v^*, the aerodynamic power continues to increase because the V^3 term in $P = \frac{1}{2}\rho V^3 A\, C_p$ grows faster than C_p decreases. The BEM predicts approximately 106 kW at 25 m/s, roughly three times the design-point power. This behaviour is characteristic of VS-FP turbines and represents a known limitation of passive stall regulation: unlike pitch-controlled machines, which can cap the power at rated by feathering the blades, a fixed-pitch rotor relies entirely on aerodynamic stall to limit loads. For turbines in this power class, the standard industrial solution is to combine the variable-speed generator with a mechanical protection system, typically passive furling, tip brakes, or an electrical dump load, that limits the rotor speed and power above a defined threshold.

The post-stall aerodynamic predictions carry significant uncertainty. The airfoil polars beyond the XFOIL convergence range ($\alpha > 15$–$20°$) are supplied by the Viterna extrapolation, which was designed to prevent BEM solver failures rather than to provide accurate deep-stall aerodynamics. In reality, three-dimensional rotational effects, massive flow separation, and unsteady phenomena in deep stall would further reduce the rotor torque compared to the two-dimensional Viterna estimates. The power curve above v^* should therefore be interpreted as an upper bound; a more accurate prediction would require experimental post-stall data or higher-fidelity CFD simulations.

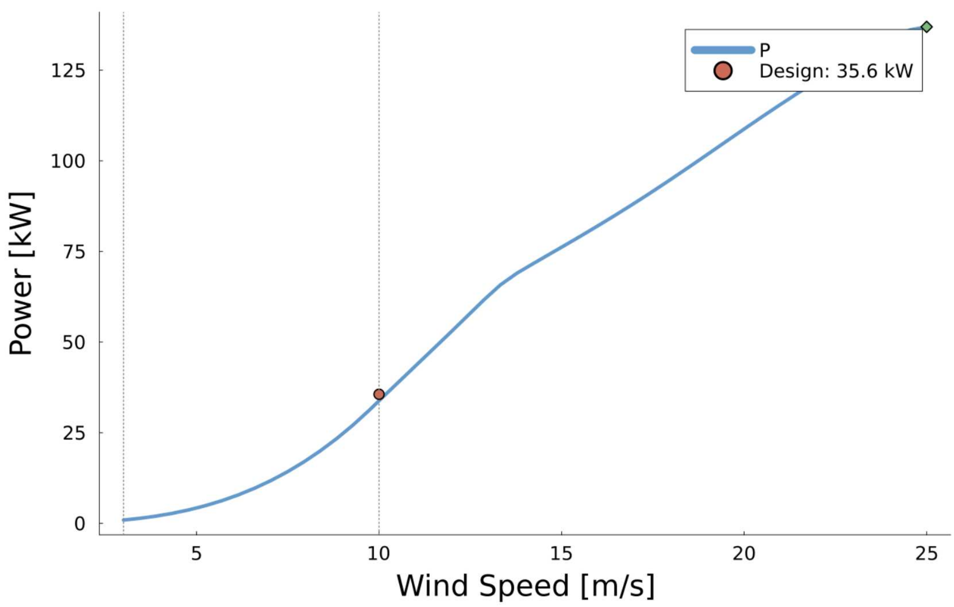

Figure 6.13: Power curve $P(V)$ for the CST-optimised blade under VS-FP control. Below $v^* = 10$ m/s the rotor tracks λ_{opt}; above v^* the rotational speed is fixed and the blades passively stall.

The power curve exhibits an unusual and unrealistic trend, with power increasing almost linearly up to approximately 130 kW. The main explanation lies in the low-fidelity tools employed: panel methods (XFOIL) with polar extrapolation throughout viterna and BEM theory are both unable to accurately predict deep-stall conditions during off-design operation of the machine.

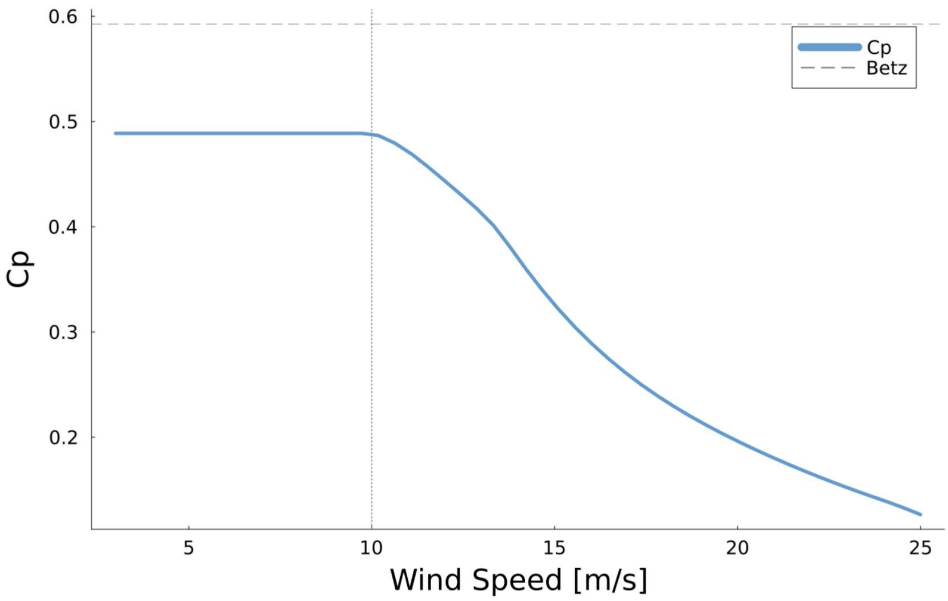

Figure 6.14: Power coefficient C_p as a function of wind speed. C_p is constant at 0.491 below v^* (variable-speed region) and decreases monotonically above it as the blade sections enter passive stall.

Annual energy production. The power curve can be combined with the Weibull wind distribution ($k = 1.857$, $c = 8.515$ m/s) to estimate the annual energy production (AEP):

$$E_{\text{annual}} = 8760 \int_{v_{\text{in}}}^{v_{\text{out}}} P(v)\, f(v)\, \mathrm{d}v, \tag{6.16}$$

where $f(v)$ is the Weibull probability density function, $v_{\text{in}} = 3$ m/s is the cut-in speed and $v_{\text{out}} = 25$ m/s the cut-out speed. Evaluating Equation (6.16) numerically with the BEM-computed power curve yields $E_{\text{annual}} \approx 143$ MWh/yr, corresponding to a capacity factor of

$$\text{CF} = \frac{E_{\text{annual}}}{P_{\text{design}} \times 8760} = \frac{143\,000}{40.0 \times 8760} \approx 0.41. \tag{6.17}$$

This value is consistent with the high mean wind speed of the site ($\bar{v} = 7.56$ m/s) and the favourable shape factor of the Weibull distribution. For context, a small rural village of 20–30 households with an average consumption of 3 000–5 000 kWh/yr per household would require 60–150 MWh/yr, placing the estimated AEP well within the target range defined at the beginning of this section.

Methodological limitations. The present study focuses on aerodynamic optimisation and does not include structural analysis and aeroelastic simulation. The power curve above the design wind speed relies on Viterna-extrapolated polars whose accuracy in deep stall is limited. The Reynolds number is evaluated at the design operating point and held constant during the optimisation; while the ranking between candidate airfoils is largely insensitive to moderate Reynolds variations, the absolute C_p values carry a small bias at off-design tip-speed ratios. A panel method itself introduces uncertainty near stall, where boundary-layer separation is not fully captured; this limitation affects the thick root section ($t/c = 0.234$) more than the thinner tip, and may contribute to the apparent improvement in polar regularity observed for the CST root, part of the smoother $C_l(\alpha)$ curve may reflect better XFOIL convergence rather than a genuine aerodynamic improvement. These limitations do not invalidate the optimisation results at the design point but should be considered when interpreting the off-design performance.

6.3 Parametric sensitivity analysis

The following section presents a series of comparisons between the results obtained by first varying the number of CST coefficients and subsequently modifying the design tip speed ratio. These comparisons aim to illustrate and justify the choices made in the previous sections, providing a clearer understanding of how each parameter influences the final outcome.

6.3.1 Effect of the number of CST coefficients

The choice of how many CST coefficients to assign to the upper and lower surfaces is one of the most consequential decisions in the optimisation setup, as it directly controls both the expressiveness of the parameterisation and the dimensionality of the design space. To investigate this trade-off quantitatively, the full dual-airfoil optimisation was repeated with three different coefficient configurations, keeping all other parameters identical (same turbine geometry, same Reynolds numbers, same COBYLA settings with 60 DOE samples and 4 restarts of 100 evaluations each):

- **7 DoF** (3 upper + 4 lower coefficients per surface) — the baseline configuration adopted throughout this work;
- **10 DoF** (4 upper + 6 lower) — an intermediate expansion intended to capture more complex curvature on the pressure side;
- **13 DoF** (5 upper + 8 lower) — a high-dimensional configuration that maximises geometric flexibility.

	7 DoF	10 DoF	13 DoF
	(3wu + 4wl)	(4wu + 6wl)	(5wu + 8wl)
Rotor performance			
$C_{p,\max}$	0.4913	0.4916	0.4931
ΔC_p vs baseline	$+1.51\,\%$	$+1.57\,\%$	$+1.88\,\%$
C_T	0.802	0.800	0.823
λ_{opt}	6.21	6.21	6.68
Power [kW]	40.0	40.0	39.7
Root aerodynamics			
α_{opt}	$6.0°$	$5.0°$	$5.5°$
C_l/C_d	109.7	114.2	107.3
C_m	-0.178	-0.181	-0.120
Tip aerodynamics			
α_{opt}	$4.5°$	$4.0°$	$5.5°$
C_l/C_d	166.6	163.7	186.7
C_m	-0.172	-0.184	-0.181
Computational cost			
Wall time	0 h 46 min	1 h 06 min	1 h 36 min

Table 6.10: Optimisation results for different numbers of CST coefficients.

The three configurations achieve remarkably similar peak power coefficients: $C_{p,\max} = 0.4913$, 0.4916, and 0.4931 for 7, 10, and 13 DoF respectively. The absolute spread is less than $0.4\,\%$, which is well within the uncertainty margin of the XFOIL-based aerodynamic analysis and suggests that the design space in the vicinity of the Kriging-selected starting point is relatively flat with respect to C_p. However, the paths taken by the optimiser and the resulting airfoil shapes differ substantially, revealing important practical trade-offs that go beyond the headline C_p figure.

Airfoil geometry. Figures 6.15 and 6.16 compare the optimised root and tip profiles for the three configurations.

At the *root* (Figure 6.15), all three configurations produce similar upper surfaces in general, the 13 Dof configuration present a more pronounced upper surface in corrspondence of the leading edge. The main differences are concentrated on the lower (pressure) side. The 7-DoF profile exhibits a smooth, gently concave lower surface that transitions cleanly from the leading edge to the trailing edge. The 10-DoF profile shows a more pronounced concavity in the aft region, a feature enabled by the additional lower-surface coefficient. The 13-DoF profile, develops a even more curved lower-surface geometry: the pressure side is flatter near the leading edge and more cambered in the mid-chord region. This additional geometric freedom allows the 13-DoF parameterisation to redistribute the camber along the chord, achieving a higher C_l/C_d at the root (107.3) with a simultaneously lower pitching moment ($C_m = -0.120$), which is the most favourable value among the three configurations.

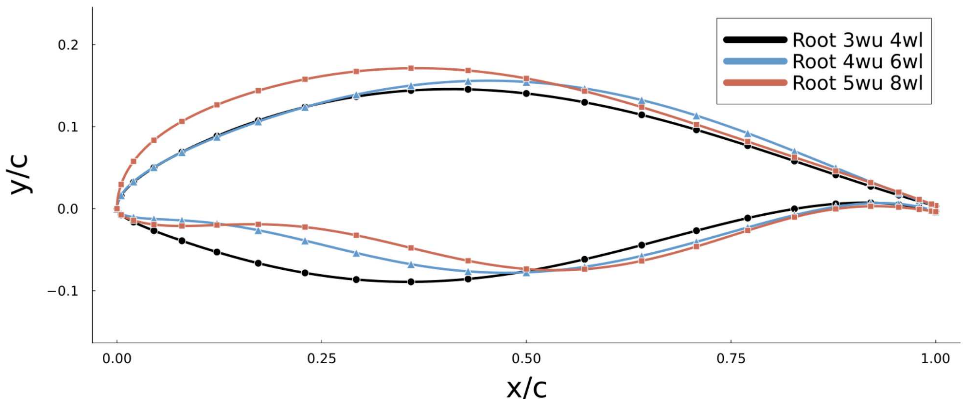

Figure 6.15: Optimised root airfoil geometry for 7, 10, and 13 CST degrees of freedom.

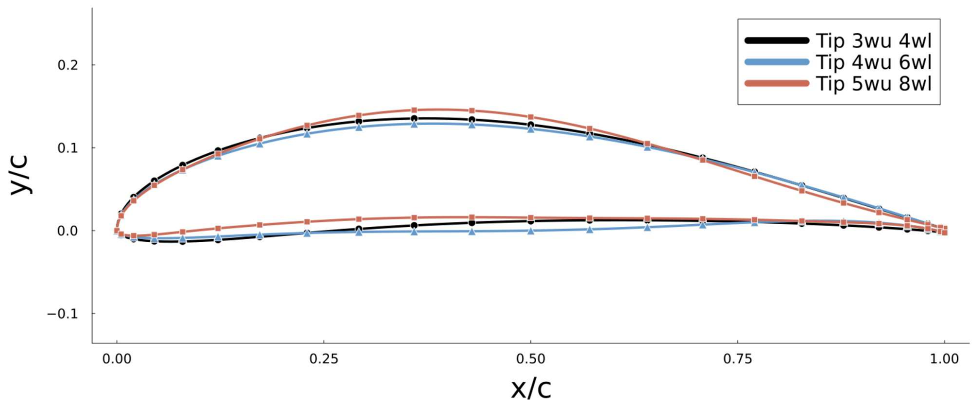

Figure 6.16: Optimised tip airfoil geometry for 7, 10, and 13 CST degrees of freedom.

At the *tip* (Figure 6.16), the differences between configurations are less pronounced than at the root. The 7-DoF and 10-DoF profiles are quite similar in both thickness and camber, whereas the 13-DoF profile departs more noticeably from the other two, exhibiting a reduced camber. The 13-DoF tip, despite its geometric departure, recovers a high $C_l/C_d = 186.7$ by exploiting its additional degrees of freedom to shape a more refined pressure distribution.

Lift polar comparison. Figures 6.17 and 6.18 show the lift coefficient as a function of the angle of attack for the root and tip sections, respectively.

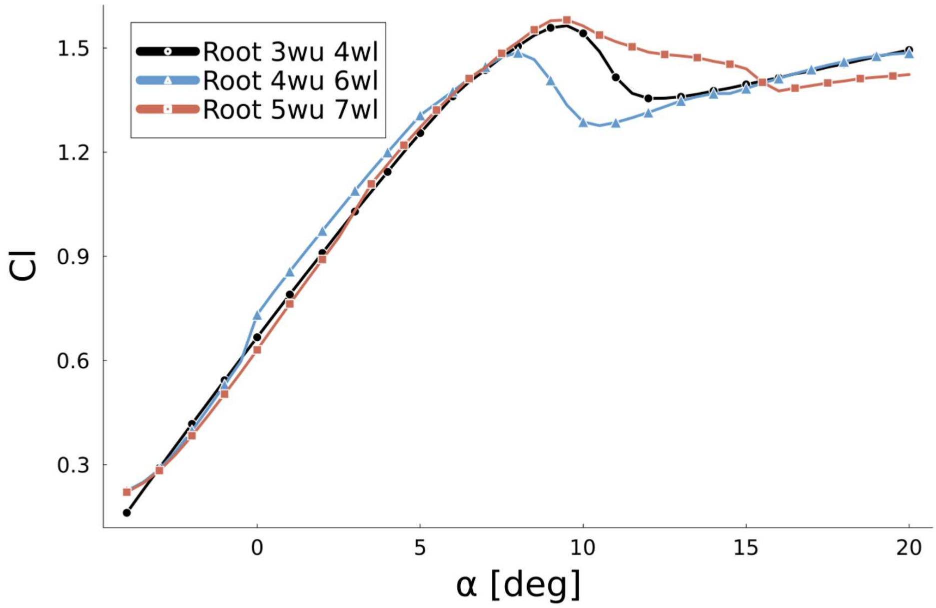

Figure 6.17: Lift coefficient $C_l(\alpha)$ of the optimised root airfoil for 7, 10, and 13 CST degrees of freedom.

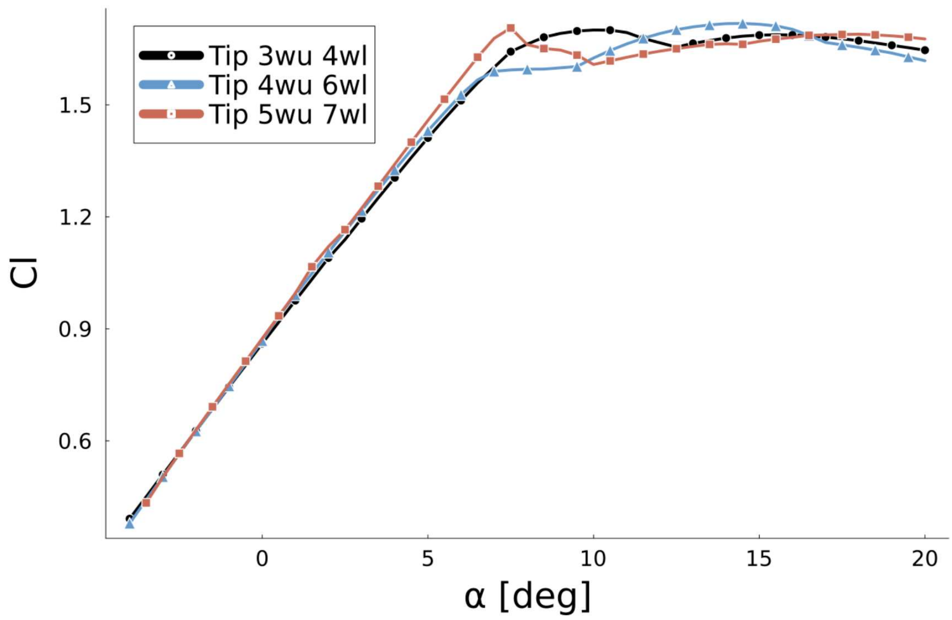

Figure 6.18: Lift coefficient $C_l(\alpha)$ of the optimised tip airfoil for 7, 10, and 13 CST degrees of freedom.

The root lift polars (Figure 6.17) reveal differences among the three configurations. The 7-DoF profile (black) exhibits the smoothest pre-stall behaviour, with a clean monotonic rise up to the stall angle at approximately $10°$, followed by a gradual lift decrease. The 10-DoF profile (blue) reaches a lower $C_{l,\,\mathrm{max}}$ with an earlier and more abrupt stall break near $\alpha \approx 7.5°$; it aslo present

some irregularities at $\alpha \approx 0°$. The 13-DoF profile (blue) displays the least regular behaviour, with small changes in the lift slope angle.
At the tip (Figure 6.18), the differences are less striking for the lower angles of attack, all the configuration present monotonic and linear behaviour, wit a little exception for the 13-DoF configuration that presents a little dip around $\alpha \approx 2°$. The 10-DoF (blue) profile shows a flat plateau between $\alpha \approx 5°$ and $\alpha \approx 10°$, indicating a possible early flow separation. This local valley is precisely the type of feature that the stall penalty function is designed to penalise, and its presence in the 10-DoF result indicates that the higher-dimensional design space contains regions where geometrically complex shapes can produce non-monotonic lift behaviour that the optimiser struggles to avoid entirely.
The 13-DoF tip closely follow the 7-DoF's lift curve but shows a distinctive kink near $\alpha \approx 7°$ followed by a secondary lift increase. As for the 7-DoF, this suggest the formation of a laminar separation bubble, anyway the more abrupt lift decrease in this case can be caused by the formation of the bubble in correspondence of the stall propagation line or a non reliable XFOIL convergence. These irregular polars are a direct consequence of the increased geometric freedom: with more coefficients, the optimiser can create subtle surface undulations that produce favourable pressure gradients at the design angle of attack but introduce flow instabilities at neighbouring incidences.

Power and thrust curves. The $C_p(\lambda)$ curves of the three configurations are practically indistinguishable for $\lambda > 6$, i.e. above the design tip-speed ratio, confirming that in the high-λ regime the rotor performance is governed primarily by the planform geometry and is insensitive to the CST parameterisation level. Below the design point, however, differences emerge: the 7-DoF and 10-DoF configurations produce the highest C_p at every sub-design λ, followed by the 13-DoF curve. The $C_T(\lambda)$ curves follow the same pattern, with virtually identical thrust coefficients above $\lambda \approx 6$ and.
A plausible explanation for this behaviour lies in the relationship between geometric simplicity and off-design robustness. At the design tip-speed ratio the local angles of attack along the blade are close to the values at which all three airfoils have been optimised, so the modest geometric differences between configurations have little effect on the integrated rotor loads. At low λ, however, the blade sections operate at substantially higher angles of attack, well away from the design incidence. The smoother, less complex profiles produced by the 7-DoF parameterisation, with their gentle curvature transitions and absence of localised surface undulations, are inherently more tolerant of these off-design conditions, maintaining attached flow and effective lift generation over a wider incidence range.

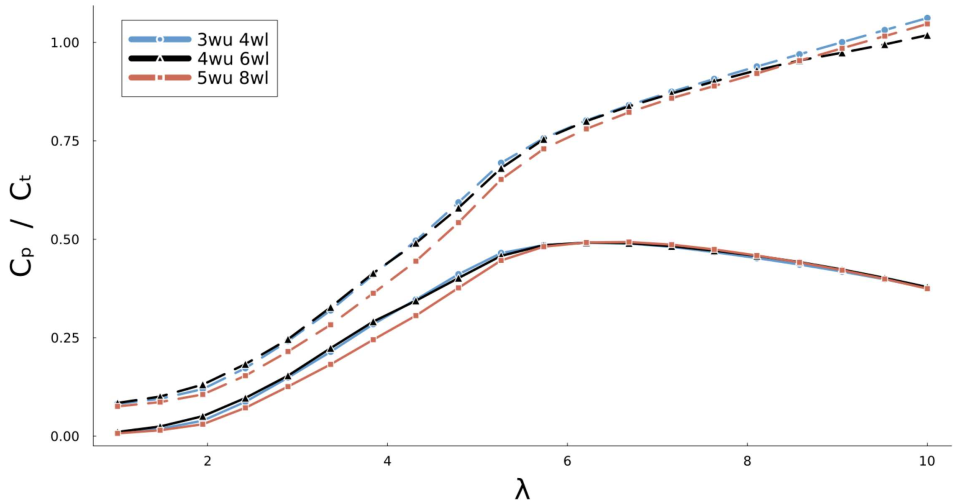

Figure 6.19: Power coefficient C_p (solid) and thrust coefficient C_T (dashed) versus tip-speed ratio λ for 7, 10, and 13 CST degrees of freedom.

The higher-order parameterisations, by contrast, exploit their additional degrees of freedom to fine-tune the pressure distribution at the design angle of attack, but the resulting subtle geometric features (sharper curvature variations, localised concavities) can trigger premature separation or laminar-separation-bubble instabilities when the airfoil operates at the elevated incidences typical of low tip-speed ratios. This explains why the 13-DoF configuration, despite achieving the highest $C_{p,\max}$ at the design point, underperforms the simpler parameterisations across the low-λ range.

Computational cost and convergence. The wall time increases substantially with the number of degrees of freedom: from 46 minutes for 7 DoF to 66 minutes for 10 DoF and 96 minutes for 13 DoF. This increase is partly due to the larger simplex that COBYLA must maintain in a higher-dimensional space (requiring more function evaluations per iteration) and partly due to XFOIL convergence difficulties on the more geometrically complex profiles generated during the search.

The convergence behaviour is also informative. With 7 DoF, all COBYLA restarts terminate with the `FTOL_REACHED` criterion, indicating that the optimiser has found a local minimum within the allocated evaluation budget. With 10 DoF, two of the four tip restarts and two root restarts hit the maximum evaluation limit (`MAXEVAL_REACHED`) without satisfying the function tolerance, suggesting that the design space is not fully explored. With 13 DoF, three of the four tip restarts reach the evaluation limit, and the spread between the best and worst restart results increases, indicating that the optimiser is struggling to converge in the larger search space. This is consistent with the known scaling behaviour of gradient-free methods, whose computational cost grows approximately as $N^{2.5}$ with the number of design variables.

Assessment. The comparison across the three parameterisation levels leads to a clear conclusion: the 7-DoF configuration (3 upper + 4 lower coefficients) provides the best overall trade-off for this application. While it achieves a marginally lower $C_{p,\,\mathrm{max}}$ than the 13-DoF configuration (0.4913 vs. 0.4931, a difference of $0.4\,\%$), it produces substantially smoother lift polars, more regular airfoil geometries, and reliable optimiser convergence at roughly half the computational cost. The 10-DoF configuration does not offer a clear advantage over the 7-DoF baseline: its C_p is essentially identical, but the tip C_l/C_d is lower (163.7 vs. 166.6) and the root stall behaviour is less smooth.

The 13-DoF configuration achieves the highest headline C_p and the best root pitching moment, but these gains come at the cost of irregular lift polars, geometrically complex profiles that may be difficult to manufacture, and an optimiser that does not reliably converge within the allocated budget. In a practical design context where robustness, manufacturability, and predictable aerodynamic behaviour are at least as important as peak efficiency, the 7-DoF parameterisation is the preferred choice. This conclusion is consistent with the general principle that, for gradient-free optimisation with a limited evaluation budget, a compact parameterisation that covers the essential geometric features is preferable to a high-dimensional one that offers marginal accuracy improvements at a disproportionate computational cost.

6.3.2 Effect of the design tip-speed ratio

A parametric study was carried out by varying the design tip-speed ratio across four values $\lambda_d = 5,\ 6,\ 7,\ 8$, while keeping all other turbine parameters fixed (rotor radius $R = 6\,\mathrm{m}$, three blades, $V_\infty = 10\,\mathrm{m\,s^{-1}}$, hub radius $r_h = 0{,}35\,\mathrm{m}$, 40 spanwise sections). Each configuration was assigned to a separate optimisation run using the 7-DoF CST parametrisation ($3\mathrm{wu} + 4\mathrm{wl}$), and the NACA starting profiles were re-evaluated at the Reynolds numbers specific to each design TSR before launching the DOE phase.

TSR	Re_{root}	Re_{tip}
5	721 998	1 679 949
6	684 583	1 404 628
7	649 613	1 206 438
8	616 985	1 057 063

Table 6.11: Reynolds numbers at root and tip sections for each design TSR.

Optimised airfoil shapes. Figure 6.20 compares the root airfoils obtained for the four design TSRs. The optimised root airfoil profiles are broadly similar in shape across the four design

TSRs, with the exception of TSR 7, which exhibits a noticeably wider trailing-edge region compared to the others. TSR 8 departs slightly from the common geometry by raising the surface near the leading edge. The optimal angles of attack follow a decreasing trend with increasing TSR, consistent with the higher local velocity seen at greater tip-speed ratios. The lift-to-drag ratio C_l/C_d peaks at TSR 6 for the root section, a consequence of the relatively higher angle of attack at that operating condition.

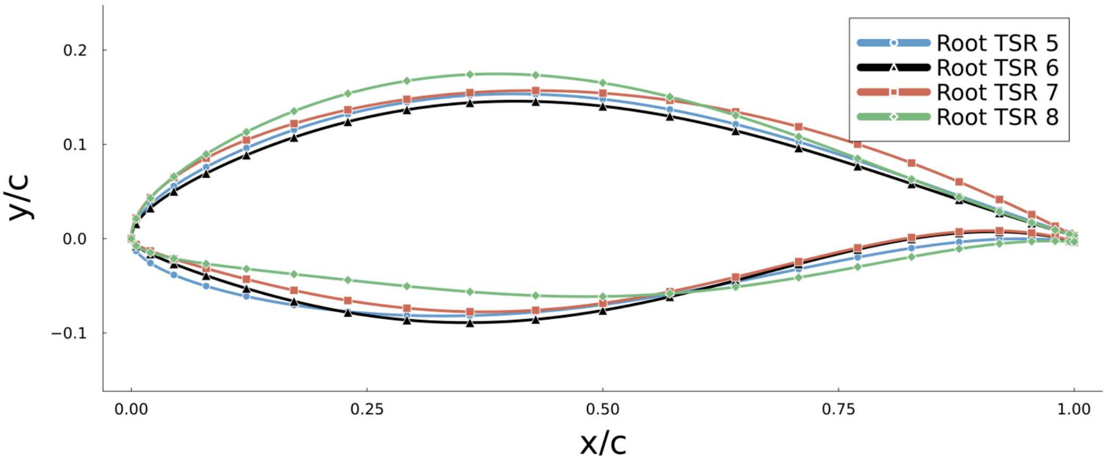

Figure 6.20: Comparison of optimised root airfoil shapes for $\lambda_d = 5, 6, 7, 8$.

At the tip (Figure 6.21), the four profiles cluster around a common geometry once TSR 6 is set aside; the TSR 6 tip displays a slightly more pronounced camber compared to the remaining three. A clear tendency toward flatter profiles is observed as TSR increases, which is directly consistent with the progressively lower optimal angles of attack recorded at higher tip-speed ratios, since a less cambered section generates sufficient lift at reduced incidence thanks to the higher local Reynolds number and dynamic pressure.

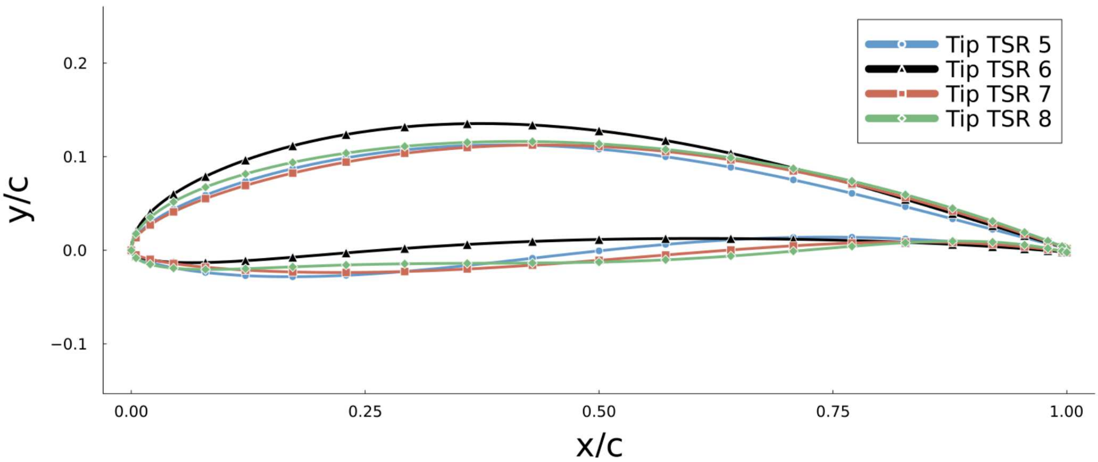

Figure 6.21: Comparison of optimised tip airfoil shapes for $\lambda_d = 5, 6, 7, 8$.

Performance comparison. Table 6.12 summarises the key results extracted from the four optimisation reports.

	TSR 5	**TSR 6**	**TSR 7**	**TSR 8**
Rotor performance				
$C_{p,\,\mathrm{max}}$	0.4844	0.4913	0.4940	0.4930
C_T @ λ_{opt}	0.807	0.802	0.805	0.818
λ_{opt}	5.74	6.21	7.16	8.11
Power [kW]	38.8	40.0	40.3	40.6
Root aerodynamics				
α_{opt}	5.0°	6.0°	3.0°	4.5°
C_l/C_d	100.7	109.7	93.9	90.4
C_m	−0.156	−0.178	−0.190	−0.118
Tip aerodynamics				
α_{opt}	4.5°	4.5°	3.5°	2.0°
C_l/C_d	166.6	166.6	161.9	143.5
C_m	−0.190	−0.172	−0.190	−0.189
Computational cost				
Wall time	0 h 45 min	0 h 46 min	0 h 40 min	1 h 04 min

Table 6.12: Optimisation results for different design tip-speed ratios (7-DoF CST parametrisation, 3wu + 4wl).

Power and turst coefficients curves. Figure 6.22 presents the power and thrust coefficient curves across the full λ range (1–10) for all four optimised designs. Solid lines represent C_p, dashed lines represent C_T; colours follow the convention established earlier (blue – TSR 5, black – TSR 6, red – TSR 7, green – TSR 8). At their respective design points all four rotors achieve remarkably similar peak power coefficients (slightly increasing with the design tip spee ratio), confirming that the BEM optimisation converges to near-identical aerodynamic efficiency regardless of the chosen tip-speed ratio, within the moderate λ range considered here. The thrust coefficient at the design point, on the other hand, decreases monotonically with design TSR, a direct consequence of the reduced rotor solidity associated with higher-TSR blade geometries. This implies that the choice of design TSR involves a trade-off not only in the shape of the $C_p(\lambda)$ curve but also in the structural loading imposed on the tower and foundation across the full operating envelope.

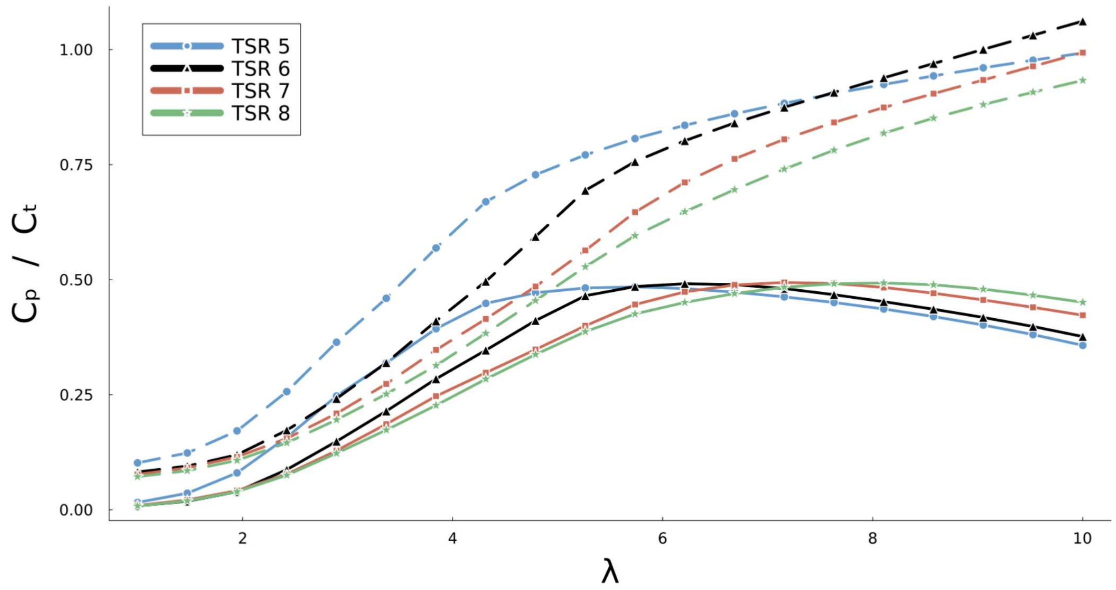

Figure 6.22: Power coefficient (solid) and thrust coefficient (dashed) versus tip-speed ratio for the four optimised rotors.

Several observations stand out:

1. **The C_p peak shifts rightward with design TSR.** As expected. Each rotor is geometrically optimised for its design point, so the $C_p(\lambda)$ curve peaks near the design λ. The differences in peak C_p across the four configurations are however negligible, with all four designs achieving $C_{p,\max} \approx 0.49$.

2. **Lower design TSRs yield flatter C_p curves.** Rotors optimised for TSR 5 and TSR 6 maintain competitive power extraction over a wider range of below-design tip-speed ratios, making them more suitable for sites with variable or gusty wind conditions. Conversely, the TSR 7 and TSR 8 rotors produce flatter curves on the high-λ side of their peak, offering better off-design performance at above-design wind speeds.

3. **The thrust coefficient C_T decreases with increasing design TSR.** This is a direct consequence of rotor solidity: BEM-optimal chord distributions scale inversely with the design tip-speed ratio, so higher-TSR rotors have narrower blades and loads a smaller fraction of the rotor disc, resulting in lower axial thrust at any given operating point.

Chord distributions. The chord distributions (Figure 6.23) reveal one of the most visually striking effects of the design TSR. As λ_d decreases, the optimal chord increases substantially across the entire span. The TSR 5 blade reaches a peak chord of approximately $1{,}04$ m at $r \approx 0{,}9$ m, while the TSR 8 blade peaks at only $0{,}78$ m at the same radial station. At the outer span ($r = 4$ m), the chord is $0{,}47$ m for TSR 5 versus $0{,}24$ m for TSR 8, nearly a factor of two.

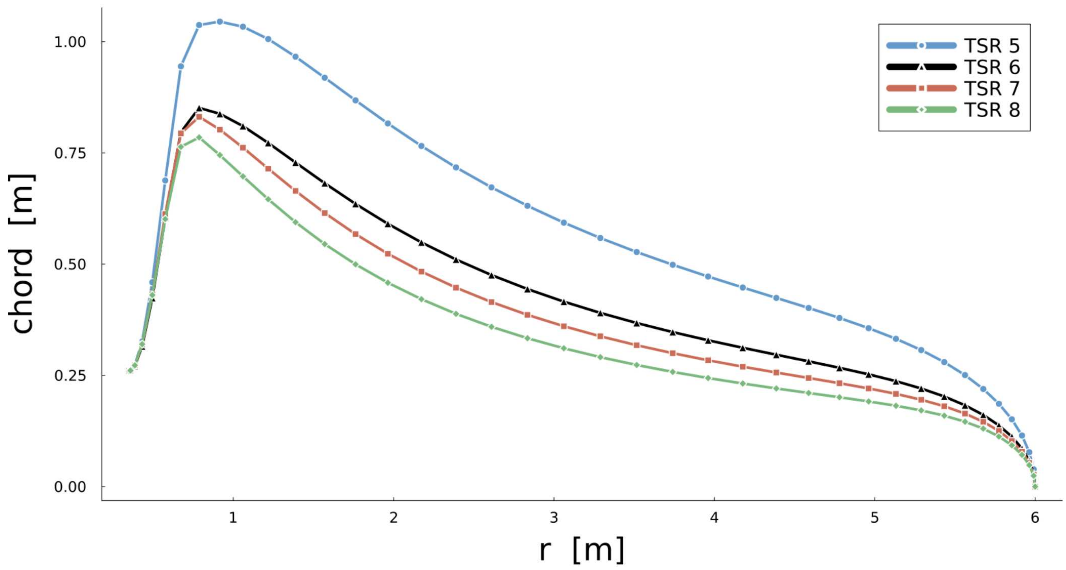

Figure 6.23: Chord distribution along the blade span for the four optimised rotors.

This behaviour follows directly from BEM theory. Recalling eq. (4.20):

$$\frac{c(x)}{R} = \frac{8\pi F a x \sin^2 \varphi}{(1-a)B\lambda C_n} \tag{6.18}$$

where B is the number of blades. For a given C_l and r, lower λ_d produces wider chords.

Acoustic implications. From an acoustic standpoint, lower-TSR designs are favourable for noise *despite* the wider chords they require. The dominant noise source in modern wind turbines is broadband turbulent-boundary-layer trailing-edge (TBL-TE) noise [24], whose level scales with the fifth power of the local flow speed [24], [25]. Since the tip speed is $V_{\text{tip}} = \lambda V_\infty$, a TSR 5 rotor operating at $V_\infty = 10\,\mathrm{m\,s^{-1}}$ has a tip speed of $50\,\mathrm{m\,s^{-1}}$ compared with $80\,\mathrm{m\,s^{-1}}$ for TSR 8. Applying the fifth-power scaling gives a theoretical noise reduction of

$$\Delta\text{SPL} \approx 50\log_{10}\frac{V_{\text{tip},8}}{V_{\text{tip},5}} = 50\log_{10}\frac{80}{50} \approx 10\,\text{dB}. \tag{6.19}$$

The wider chord of the lower-TSR blade does increase the noise-producing surface area and raises low-frequency TBL-TE contributions through a thicker boundary layer displacement thickness δ^*, which scales with chord and Reynolds number. However, as shown by Oerlemans et al. [24], blade noise is produced primarily at the outer portion of the blade where local flow speeds are highest, so the fifth-power velocity term dominates. The chord-related increase is of the order of $10\log_{10}(c_1/c_2) \approx 2$–$3\,\text{dB}$ and is far outweighed by the velocity reduction of Eq. (6.19). In summary, lower design TSR leads to wider chords but substantially slower rotation and, therefore, significantly less aerodynamic noise.

Shift between design and optimal λ. An interesting observation from the optimisation reports is that the operating TSR at which C_p is maximised consistently exceeds the design value:

λ_d	λ_{opt}	$\Delta\lambda/\lambda_d$
5	5.74	$+14.8\,\%$
6	6.21	$+3.5\,\%$
7	7.16	$+2.3\,\%$
8	8.11	$+1.4\,\%$

This may not be a numerical artefact but a consequence of the difference between the idealised assumptions used in the inverse BEM design and the more realistic forward BEM analysis. Three effects contribute to the upward shift:

1. **Tip and root loss corrections.** The Prandtl correction reduces the effective annular loading near the blade extremities, lowering the axial induction factor a below its design value. A lower a means that the rotor extracts less momentum at the design rotational speed, and the C_p peak shifts to a slightly higher λ where the remaining blade sections can compensate.

2. **Profile drag.** Viscous drag subtracts from the tangential force component (torque) but adds to the axial component (thrust). Since drag reduces torque preferentially, the rotor needs to spin slightly faster to recover the lost torque, again pushing λ_{opt} upward.

3. **Airfoil optimisation.** The CST-optimised profiles may achieve their best C_l/C_d at a slightly lower angle of attack than assumed during the initial inverse design. If the blade operates at lower α than expected it is effectively under-loaded at λ_d, and the true optimum shifts to a higher λ where the geometric angle of attack increases back toward the aerodynamic sweet spot.

The shift magnitude ($\Delta\lambda \approx 0.1$–$0.7$) is fully consistent with values reported in the BEM literature [26]. The effect is most pronounced for the lowest design TSR ($+14.8\,\%$ for $\lambda_d = 5$) because the absolute tip losses are larger and the drag penalty is proportionally higher on wider, slower blades. As λ_d increases the blade becomes more slender, the Prandtl correction weakens, and the shift diminishes to just $+1.4\,\%$ for $\lambda_d = 8$.

7 Conclusions

This thesis set out to build a fast, open-source tool for the preliminary aerodynamic optimization of wind turbine blades in the sub-200 kW range, closing the gap left in this class, who usually rely on off-the-shelf airfoils rather than a design optimized for the actual operating conditions. The goal was to reach this without CFD-level computational cost, so that the whole optimization could run on a normal laptop. As shown throughout the thesis, this objective was achieved.

7.1 Current state of the code

The validation against the NTNU Blind Test 1 campaign showed that the framework reproduces the essential aerodynamics of the rotor around its design point, with predicted and experimental power coefficients matching closely in the range $3 < \lambda < 7$. The dual-airfoil optimization strategy, combined with the preliminary starting point design built on the Kriging baseline, worked well at avoiding premature convergence to bad local optima, while keeping the number of function evaluations low. The resulting blade designs also behaved consistently across the parametric sensitivity analysis: as the design tip-speed ratio changes, the working conditions of the blade, like the chord distribution, adapt in a way that matches standard BEM theory.
At the same time, the framework has a set of limitations that were already discussed through the thesis, and that define its current scope of application:

- **No structural or aeroelastic model.** Blade bending and aeroelastic deformation are not included, and this is the main reason behind the sub-200 kW limit. The threshold itself is not a mathematically precise boundary, but more a qualitative estimation: without any structural or aeroelastic study, the code cannot be trusted to give reliable results for larger, more flexible blades.

- **Reynolds number held constant.** During optimization the Reynolds number is fixed at the design point instead of being interpolated locally, which introduces a small bias at off-design tip-speed ratios.

- **Viterna extrapolation in deep stall.** Off-design results in deep stall are not reliable. XFOIL does not converge at high angles of attack, so there are no real polar data in that region to begin with. Viterna extrapolation is used to fill this gap and keep CCBlade from crashing, but it only gives a rough, qualitative estimate of the polars, not an accurate one. This means that the power curve above the design wind speed is built on extrapolated, not computed, data.

- **Local optimizer.** COBYLA gives no formal guarantee of finding the global optimum; the multi-start strategy only partially reduces this risk.

None of these limitations invalidate the results obtained at the design point, but they do mark the boundaries within which the current version of the code can be trusted without further validation.

7.2 Future developments

The limitations above point to a clear direction for future work: moving from the current, fairly monolithic implementation to a modular architecture built around a small, well-tested core. Instead of expanding the main BEM/optimization loop to cover every possible case, the idea is to keep the core simple and reliable, and attach interchangeable packages depending on the size and type of turbine being analyzed.
The clearest example is blade geometry generation itself: the blade geometry (hub attachment blending, thicknes, etc.) that work for a 500 W micro-turbine are very different from those of a 500 kW machine, to the point that one single geometry routine is unlikely to stay optimal across the whole size range. More dedicated geometry packages, chosen based on the turbine's power class, would let the core focus only on the optimization logic, while the size-specific geometric assumptions live in separate, swappable modules.
The same idea can be applied to most of the limitations:

- **3D-correction module**, calibrated separately for small- and large-scale rotors, instead of reusing corrections tuned for utility-scale turbines;
- **Reynolds-aware polar module**, interpolating between multiple Reynolds numbers during the BEM iterations instead of freezing it at the design point;
- **Structural/aeroelastic module**, to be turned on for blade lengths where rigid-blade BEM is no longer good enough, extending the framework beyond the sub-200 kW range;
- **Blade blending / tip-treatment**, packages, needed for larger rotors where the root section has to be thick for structural reasons while the tip needs to stay thin and aerodynamic. In the current framework the thickness distribution along the span is fixed a priori, based on the turbine size and general geometry choices. A more advanced module could instead compute the thickness distribution directly from the actual blade loads, so that root and

tip thickness (and the blending between them) come from a structural requirement rather than a predefined assumption. This would also naturally include tip-specific treatments (e.g. swept or tapered tips) to reduce tip losses and noise, both of which are less critical on the small blades considered in this thesis.

- **LBM-based aerodynamic module**, coupling the framework with a Lattice Boltzmann solver for selected high-fidelity runs. Unlike XFOIL, an LBM solver would not fail to converge at high angles of attack, and could resolve 3D rotational effects directly instead of relying on empirical stall-delay corrections. This would not replace the fast BEM/XFOIL loop used for optimization, but could be used to validate or refine the off-design and deep-stall portion of the power curve, where Viterna extrapolation is currently the only option.

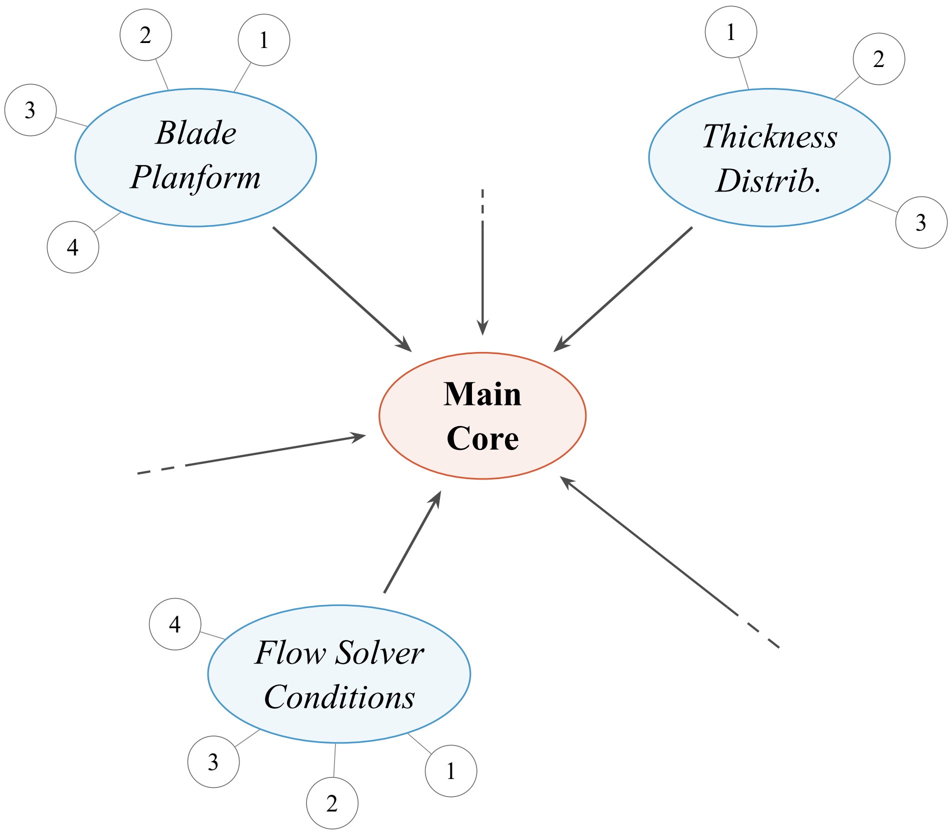

Figure 7.1: Example of a possible layout: the main core draws from several sub-modules, each of which admits variants depending on the case under analysis.

This modular structure would keep the framework's original goal: a fast, open-source, laptop-friendly tool; while making it possible to widen its range of application without putting at risk the reliability of the validated core.

Bibliography

[1] T. Burton, D. Sharpe, N. Jenkins, and E. Bossanyi, *Wind Energy Handbook*. John Wiley & Sons, 2001.

[2] M. O. L. Hansen, *Aerodynamics of Wind Turbines*, 3rd ed. Routledge, 2015.

[3] J. D. Anderson, *Fundamentals of Aerodynamics*, 6th ed. McGraw-Hill, 2017.

[4] S. Raghavendra, T. R. Ravikumar, G. V. Gnaendra Reddy, K. N. Manjunatha, and S. V. Madhusudhana, "Design of wind blades for the development of low-power wind turbines using Betz and Schmitz methods," *Advances in Materials and Processing Technologies*, 2020, Accepted 5 October 2020. DOI: 10.1080/2374068X.2020.1833605

[5] E. Branlard, *Wind Turbine Aerodynamics and Vorticity-Based Methods: Fundamentals and Recent Applications*. Springer, 2017. DOI: 10.1007/978-3-319-55164-7

[6] D. A. Spera, *Wind Turbine Technology: Fundamental Concepts of Wind Turbine Engineering*. New York: ASME Press, 1994, p. 638.

[7] I. Herráez, B. Stoevesandt, and J. Peinke, "Insight into rotational effects on a wind turbine blade using Navier–Stokes computations," *Energies*, vol. 7, no. 10, pp. 6798–6822, 2014. DOI: 10.3390/en7106798

[8] J. Schepers and H. Snel, "Model experiments in controlled conditions," Energy Research Centre of the Netherlands (ECN), Tech. Rep. ECN-E–07-042, 2007, EU FP5 Project MEXICO, Contract No. ENK5-CT-2000-00309.

[9] J. R. R. A. Martins and A. Ning, *Engineering Design Optimization*. Cambridge University Press, 2021, ISBN: 9781108833417.

[10] M. A. Bouhlel, J. T. Hwang, N. Bartoli, R. Lafage, J. Morlier, and J. R. R. A. Martins, "A Python surrogate modeling framework with derivatives," *Advances in Engineering Software*, vol. 135, p. 102 662, 2019, ISSN: 0965-9978. DOI: 10.1016/j.advengsoft.2019.03.005

[11] P. Saves et al., "SMT 2.0: A surrogate modeling toolbox with a focus on hierarchical and mixed variables Gaussian processes," *Advances in Engineering Software*, vol. 188, p. 103 571, 2024. DOI: 10.1016/j.advengsoft.2023.103571

[12] M. J. D. Powell, "A direct search optimization method that models the objective and constraint functions by linear interpolation," in *Advances in Optimization and Numerical Analysis*, S. Gomez and J.-P. Hennart, Eds., Dordrecht: Kluwer Academic Publishers, 1994, pp. 51–67.

[13] M. J. D. Powell, "A view of algorithms for optimization without derivatives," Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Tech. Rep. DAMTP 2007/NA03, 2007, Presented as a William Benter Distinguished Lecture at the City University of Hong Kong.

[14] C. L. Ladson, C. W. Brooks, A. S. Hill, and D. W. Sproles, *Computer program to obtain ordinates for NACA airfoils*, 1996. [Online]. Available: `https://ntrs.nasa.gov/citations/19970008124`

[15] Stanford University. "The NACA airfoil series." Course material, AA200, Accessed: Feb. 17, 2025. [Online]. Available: `https://web.stanford.edu/~cantwell/AA200_Course_Material/The%20NACA%20airfoil%20series.pdf`

[16] B. M. Kulfan and J. E. Bussoletti, "Fundamental parametric geometry representations for aircraft component shapes," in *11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, AIAA 2006-6948, 2006. DOI: `10.2514/6.2006-6948`

[17] M. Drela and H. Youngren, *XFOIL 6.9 user primer*, MIT Department of Aeronautics and Astronautics, Nov. 2001. [Online]. Available: `https://web.mit.edu/drela/Public/web/xfoil/xfoil_doc.txt`

[18] L. M. Mack, "Transition and laminar instability," Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, Tech. Rep. JPL Publication 77-15, 1977, NASA-CP-153203, N77-24058.

[19] L. A. Viterna and D. C. Janetzke, "Theoretical and experimental power from large horizontal-axis wind turbines," NASA, Cleveland, Tech. Rep. DOE/NASA/20320-41, 1982.

[20] S. A. Ning, "A simple solution method for the blade element momentum equations with guaranteed convergence," *Wind Energy*, vol. 17, no. 9, pp. 1327–1345, 2014. DOI: `10.1002/we.1636`

[21] P.-Å. Krogstad, P. E. Eriksen, and J. A. Melheim, ""blind test" workshop calculations for a modell wind turbine," Department of Energy and Process Engineering, NTNU, Tech. Rep., 2011.

[22] R. F. Mikkelsen, "Actuator disc methods applied to wind turbines," PhD thesis, Technical University of Denmark, Department of Mechanical Engineering, Lyngby, Denmark, 2004, ISBN: 87-7475-296-0.

[23] S.-P. Breton, F. N. Coton, and G. Moe, "A study on rotational effects and different stall delay models using a prescribed wake vortex scheme and NREL phase VI experiment data," *Wind Energy*, vol. 11, no. 5, pp. 459–482, 2008. DOI: `10.1002/we.269`

[24] S. Oerlemans, P. Sijtsma, and B. Méndez López, "Location and quantification of noise sources on a wind turbine," *Journal of Sound and Vibration*, vol. 299, no. 4–5, pp. 869–883, 2007. DOI: `10.1016/j.jsv.2006.07.032`

[25] T. F. Brooks, D. S. Pope, and M. A. Marcolini, "Airfoil self-noise and prediction," NASA, Hampton, Virginia, Tech. Rep. NASA-RP-1218, 1989.

[26] J. F. Manwell, J. G. McGowan, and A. L. Rogers, *Wind Energy Explained: Theory, Design and Application*, 2nd ed. Chichester: John Wiley & Sons, 2009.