Regression Model of Leading Edge Inflatable Kite Profile Aerodynamics

Kasper Masure

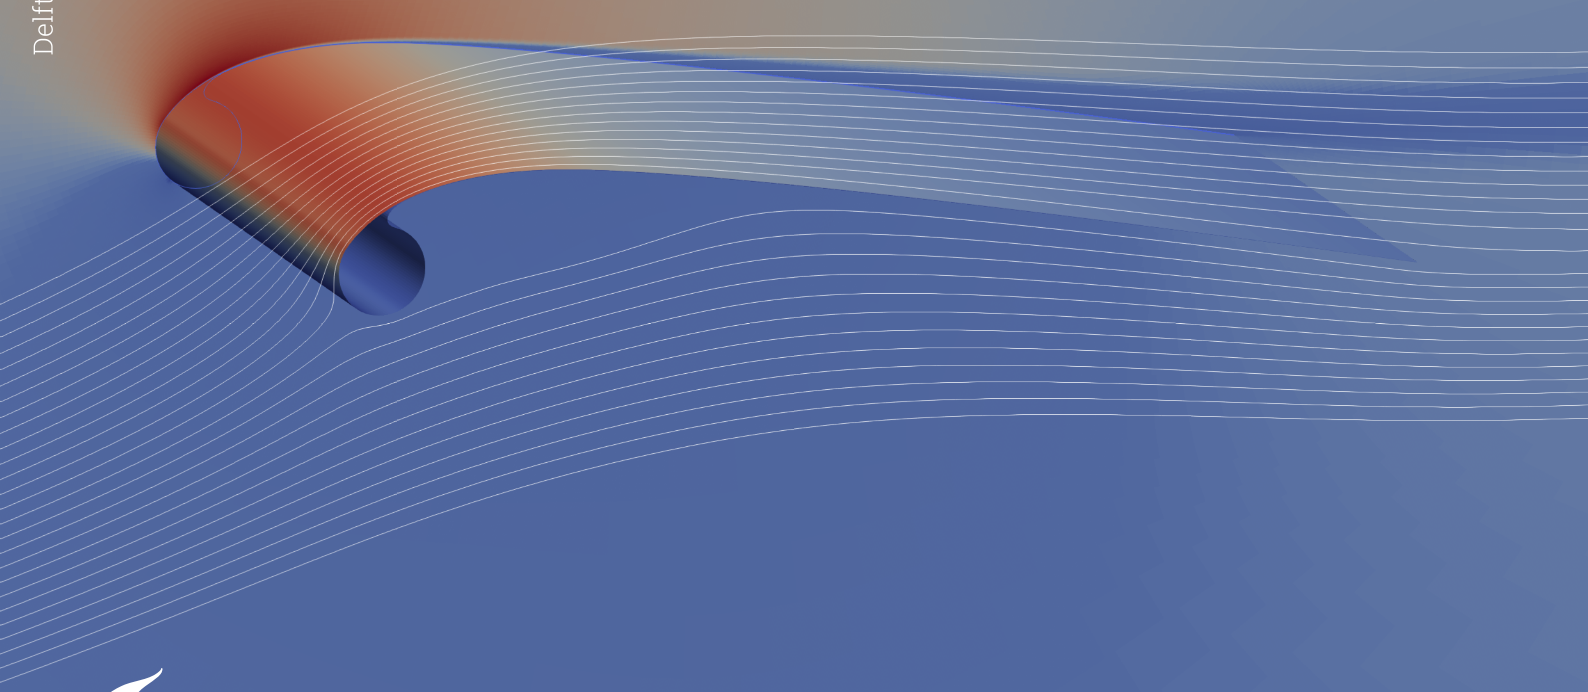

Delft University of Technology

TU Delft

Cover image: Contour plot of the flow field with streamlines around a two-dimensional leading edge kite profile at 10° angle of attack.

Regression Model of Leading Edge Inflatable Kite Profile Aerodynamics

by

Kasper Masure

to obtain the degree of Master of Science in Aerospace Engineering
at the Delft University of Technology,
to be defended publicly on Wednesday July 2, 2025 at 10:00.

Student number:	4861086	
Project duration:	August 13, 2024 - July 2, 2025	
Thesis committee:	Dr. F. Oliviero	TU Delft, chair
	Dr. F. de Prenter	TU Delft, examiner
	Dr.-ing. R. Schmehl,	TU Delft, supervisor
	MSc J.A.W. Poland,	TU Delft, supervisor

An electronic version of this thesis is available at `http://repository.tudelft.nl/`

TU Delft

Preface

I couldn't have imagined a better way to end my master's than by writing a thesis on my out-of-hand hobby: kitesurfing. If I weren't writing my thesis, chances were high that I was out on the water often with Jelle, my PhD supervisor.

Big thanks to Jelle for the supervision and for being the awesome person you are. Even though we only met at two very specific locations over the past ten months, either in Roland's office or on the beach, it feels like we spent a lot of time together, and I'm sure that won't end here. From talking kitesurfing during meetings to flexing on the beach about how many HPC CPUs I've got running, it's been a great trip.

Roland Schmehl, we first got in touch during DSE, and unsurprisingly, it was already about kites. Then came the great course you gave, where I got a deeper understanding of kite dynamics. Finally, I had the chance to bring it all together by doing my master's thesis under your supervision. I'm really happy with how I could integrate my practical knowledge from kitesurfing and design at Vantage Kites with your years of expertise in airborne wind energy. I honestly couldn't have wished for a better team together with Jelle and you. Thank you for all the knowledge, guidance, and inspiration you've shared throughout this journey.

I also want to thank Vantage Kites and, more specifically, Robin van de Putte for teaching me the practical side of kite design. Your insights helped me understand what to integrate into the numerical model for my thesis. On top of that, you provided a prototype kite which we could glue full of tufts, making it possible to visualise a specific flow phenomenon of my research.

The past ten months were mostly spent watching a kite fly or a kite model on my pc, with the occasional beers in between. It's not hard to imagine that most people got bored after the 1000th kite story. But one person who never did was Ivan Tamborero. I'm truly grateful for the trips to Leucate and Tarifa, the hours we spent goofing around on the water, and without a doubt, the endless discussions about kite profiles. One thing's for sure, I'll design you your own WR kite, just hope it brings you back to earth. #SDSB

Last but not least, I want to thank my friends at Mother Vaarse (Arnaud Mathieu, Oscar Carpentier, Niels Rubrecht, and Aurélie Soors), the scouts crew back home in Belgium, and my family led by my sheep Toine, for always being there for me, and for sharing a good beer together, Toine included!.

Kasper Masure
Delft, June 2025

Abstract

This study aims to develop a regression model of the aerodynamic coefficients for leading edge inflatable (LEI) kite profiles. Aerodynamic data is obtained by 2D computational fluid dynamics (CFD) simulations for different Reynolds numbers, angle of atack and profile configuration, leading to the following: lift, drag and moment coefficients C_l, C_d, C_m, as well as the surface distribution of pressure and skin friction coefficients (C_p and C_f). The machine learning regression model is trained on this multidimensional dataset to generate accurate 2D aerodynamic predictions, which serve as essential input for the vortex step method (VSM), a fast aerodynamic solver used in kite fluid-structure interaction (FSI) simulations.

The profile geometry parameterisation forms the backbone of the automated CFD toolchain. It provides an advanced and robust design framework for LEI kite profiles. The geometry is defined by its main components: a circular leading edge (LE) tube and a canopy, which is subdivided into two splines. The front spline connects the LE tube seam (i.e., the LE tube–canopy stitching connection) to the maximum camber point, while the rear spline extends from this point to the trailing edge (TE). Both splines are modelled as cubic Bézier curves with four control points, where the first and last points define the connection boundaries, enabling smooth and flexible surface shaping. The seam angle on the LE tube is dynamically calculated to ensure a smooth transition for any given configuration. The positions of the control points are governed by the following non-dimensionalised profile parameters, defined relative to the chord: LE diameter t, maximum camber chordwise position η, camber height κ, reflex angle δ, camber tension λ, and LE curvature ϕ. For meshing purposes, a finite thickness is assigned to the canopy to separate the upper and lower flow regions. Additionally, a LE fillet is added to the underside of the canopy to facilitate mesh smoothing at the sharp corner connection with the LE tube.

Aerodynamic data is collected from steady Reynolds-averaged Navier–Stokes (RANS) simulations, employing the $k-\omega$ shear stress transport (SST) turbulence model. The simulations are performed with the open-source CFD software OpenFOAM, using structured meshes generated in Pointwise. An extensive mesh sensitivity analysis was conducted, focusing on the effects of canopy thickness t_canopy, the LE fillet, and the resolution of the fully structured grid in both normal and tangential directions. Transition modelling was omitted based on the assumption that the boundary layer undergoes forced transition at the LE tube seams. This includes the LE tube-canopy connection on the upper side and the LE closing seam on the lower side, where numerical results indicated that the flow transitions due to seam-induced roughness. Since the region upstream of the seams is small, its impact is considered negligible, justifying the simplification.

Due to the large number of required simulations, computational resources from the high-performance computing (HPC) cluster of the Faculty of Aerospace Engineering at TU Delft were utilised. To define the parameter configurations for data collection, a trade-off was made between parameter resolution and computational cost. Parameters were sampled across the following ranges for three Reynolds numbers (Re $= 1\times10^6$, 5×10^6, and 2×10^7): α from -22° to -10° (13 values), t from 0.03 to 0.12 (5 values), η from 0.08 to 0.4 (8 values), κ from 0.04 to 0.16 (7 values), δ from -8° to 0° (4 values), and λ from 0.1 to 0.4 (4 values). The LE curvature ϕ was held constant at 0.65.

The flow fields were analysed for the effects on the newly introduced parameters in the updated profile geometry model: δ, λ, and ϕ. Downward deflections of the profile TE due to negative δ resulted in reduced lift (C_l) and increased drag (C_d) performance, while enhancing longitudinal pitching moment stability. In contrast, variations in λ showed the opposite effect; increased camber tension resulted in higher C_l and C_d values but diminished longitudinal stability. The parameter ϕ, having minimal geometric influence, caused negligible changes in aerodynamic performance, only slightly altering the pressure distribution. Consequently, ϕ was fixed in the regression model. Among all tested algorithms, the extra trees (ET) model achieved the highest predictive accuracy, with R^2 scores of 0.987 for Re $= 1\times10^6$, 0.988 for 5×10^6, and 0.989 for 2×10^7.

Contents

Nomenclature

Abbreviations

Abbreviation	Definition
2D	Two-dimensional
3D	Three-dimensional
AoA	Angle of attack
AWE	Airborne wind energy
CFD	Computational fluid dynamics
CST	Class shape transformation
CoP	Center of pressure
DNS	Direct numerical simulation
ET	Extra trees
FEM	Finite element method
FSI	Fluid-structure interaction
FVM	Finite volume method
HAWT	Horizontal-axis wind turbine
HPC	High performance computing
KNN	K-nearest neigbors
LE	Leading edge
LEI	Leading edge inflatable
LES	Large-eddy simulation
LLM	Lifting line method
LR	Linear regression
MLP	Multi-layer perceptron
MSE	Mean square error
NACA	National advisory committee for aeronautics
PCHIP	Piecewise cubic Hermite interpolating polynomial
RANS	Reynolds-averaged Navier-Stokes
RDP	Ramer-Douglas-Peucker
RF	Random forest
RMSE	Root mean square error
SIMPLE	Semi-implicit method for pressure linked equations
SST	Shear stress transport
SVR	Support vector regression
TE	Trailing edge
VSM	Vortex step method

Symbols

Symbol	Definition	Unit
c	Chord length	m
C_{d}	2D drag coefficient	-
C_{f}	Local skin friction coefficient	-
C_{l}	2D lift coefficient	-
C_{m}	2D moment coefficient	-
C_{p}	Pressure coefficient	-

Symbol	Definition	Unit
D	Aerodynamic drag	$\mathrm{kg\,m\,s^{-2}}$
f_{bl}	Velocity correction factor for boundary layer viscous effects	-
h	Roughness-element height in the boundary layer	-
I	Turbulence intensity	-
k	Turbulent kinetic energy	$\mathrm{m^2\,s^{-2}}$
k_i	Turbulent kinetic energy at the inlet	$\mathrm{m^2\,s^{-2}}$
L	Aerodynamic lift	$\mathrm{kg\,m\,s^{-2}}$
M	Aerodynamic moment	$\mathrm{kg\,m^2\,s^{-2}}$
n_{canopy}	Canopy node count	-
n_{LE}	LE node count	-
n_{TE}	TE node count	-
n_{total}	Profile node count	-
p	Pressure	$\mathrm{kg\,m^{-1}\,s^{-2}}$
p_∞	Free-stream pressure	$\mathrm{kg\,m^{-1}\,s^{-2}}$
p_{k}	Kinemtic pressure	$\mathrm{m^2/s^2}$
p_{s}	Static pressure	$\mathrm{kg\,m^{-1}\,s^{-2}}$
Re	Reynolds number	-
$\tilde{\text{Re}}_{\theta_t}$	Transition onset momentum thickness Reynolds number	-
Re_{h}	Roughness-height-based Reynolds number	-
U_{h}	Flow velocity at the roughness-element height	-
t	Non-dimensional tube diameter (relative to c)	-
t_{canopy}	Non-dimensional canopy thickness (relative to c)	-
t_{b}	Bézier curve interpolation parameter	
U_∞	Free-stream velocity magnitude	$\mathrm{m\,s^{-1}}$
$\mathbf{U}_{\text{i}}$	Velocity vector at inlet	$\mathrm{m\,s^{-1}}$
U_x	x-component of the velocity	$\mathrm{m\,s^{-1}}$
U_y	y-component of the velocity	$\mathrm{m\,s^{-1}}$
U_τ	Friction velocity	$\mathrm{m\,s^{-1}}$
x_{cp}	Chordwise position of the center of pressure	-
x_{ref}	Reference chordwise location for center of pressure	-
y^*	Non-dimensional wall distance, turbulent kinetic energy based	-
y^+	Non-dimensional wall distance, wall shear stress based	-
α	Angle-of-attack	$^\circ$
γ	Intermittency parameter	-
δ	Reflex angle	$^\circ$
δ_{ij}	Kronecker delta	-
ϵ	Turbulent dissipation rate	$\mathrm{m^2\,s^{-3}}$
η	Max. camber chordwise position (relative to c)	-
κ	Camber height (relative to c)	-
λ	Camber tension	-
μ	Dynamic viscosity	$\mathrm{kg\,m^{-1}\,s^{-1}}$
μ_∞	Free-stream dynamic viscosity	-
ν	Kinematic viscosity	$\mathrm{m^2\,s^{-1}}$
ν_∞	Free-stream kinematic viscosity	$\mathrm{m^2\,s^{-1}}$
ν_t	Eddy viscosity	$\mathrm{m^2\,s^{-1}}$
ρ	Density	$\mathrm{kg\,m^{-3}}$
ρ_∞	Free-stream density	$\mathrm{kg\,m^{-3}}$
σ	Standard deviation	-
τ_{w}	Wall shear stress	$\mathrm{kg\,m^{-1}\,s^{-2}}$
ϕ	LE curvature	-
ω	Specific turbulent dissipation rate	$\mathrm{s^{-1}}$
ω_{i}	Specific turbulent dissipation rate at the inlet	$\mathrm{s^{-1}}$

1 Introduction

With the accelerating impacts of climate change, the need for sustainable energy solutions is more urgent than ever. As global energy demand rises, transitioning from fossil fuels to renewable sources is critical. Wind energy is a leading contributor, primarily through the horizontal-axis wind turbines (HAWT). However, these are limited in height by structural and economic constraints. Airborne Wind Energy (AWE) systems present a promising alternative by accessing stronger, more consistent winds at higher altitudes. At around 500 m, wind speeds can be 10% to 100% higher than those at conventional turbine heights, depending on location [4]. This altitude advantage offers significant potential for more efficient and cost-effective energy generation.

AWE systems are typically classified as ground-generation (ground-gen) or onboard-generation (fly-gen), depending on where the electricity is produced. Each category includes various sub-systems [41, 14]. In ground-gen systems, for example, energy can be generated through crosswind flight or by maintaining a static position. During crosswind cycles, the kite generates energy by pulling a tether, which rotates a drum that is connected to a generator, hence delivering power. This work focuses on soft-wing kites, specifically leading edge inflatable (LEI) kites, which are also widely used in kitesurfing and can be applied in other areas such as boat towing. LEI kites consist of an inflated leading edge (LE) tube and spanwise-distributed inflatable struts that provide rigidity and shape the aerodynamic profile [47, 48].

Although the inflated structures provide stiffness to the kite, it remains a highly flexible system, with a thin canopy forming a lightweight, deformable platform. This flexibility is crucial for maintaining flight in low wind conditions and enables the deformations required for steering. As the wing deforms during flight and actuation, the airfoil shape changes locally, altering the aerodynamic loading. To accurately capture this aero-structural interaction, it's essential to include the effects of these shape changes in a coupled fluid-structure interaction (FSI) framework.

A high-fidelity approach to optimise soft-wing kites involves coupling computational fluid dynamics (CFD) with structural finite element method (FEM) simulations [24]. However, the high computational cost of CFD-FEM coupling necessitates the development of faster, low-fidelity models for efficient design exploration. At the TU Delft, the vortex step method (VSM) [11] and the particle system model [38] were developed for this purpose. Combined, they provide a rapid aero-structural simulation framework at significantly reduced computational cost.

As the airfoil shape changes dynamically during flight, the aerodynamic loading must be known for a wide range of possible configurations. To enable this, the VSM relies on viscous 2D airfoil polars to describe the local aerodynamic loading on the flexible profiles. These polars are currently obtained from look-up tables or polynomial regression models, first introduced by Breukels [9], using only camber magnitude and leading edge tube diameter as profile parameters. Kappel [27] and Cayon [11] recommended including the chordwise position of maximum camber, which was later implemented by Watchorn [50] in the current polynomial regression model. However, this model was developed using a limited dataset and low-order polynomial algorithms, resulting in low prediction accuracy and high variance for unseen profiles. This highlights the need for improvement and motivates the following research questions:

1. *How can the parametric model of the LEI kite profile be extended to cover the wide range of shape variations that occur during flight, while improving smoothness, robustness, and flexibility through additional shape parameters?*
2. *How can CFD simulations of LEI kite profiles be made more efficient to enable an automated process for meshing and determining the 2D flow field?*
3. *To what extent can machine learning enhance regression models for predicting the aerodynamic performance of LEI kite profiles?*

Machine learning (ML) has become widely adopted across various sectors, driven by the growing availability of large datasets and the increasing demand for performance optimisation. Recent advancements in computer science have led to the development of efficient ML algorithms that support scientists and engineers in data-driven analysis. A relevant example to this thesis, Teimourian et al. [45] applied ML regression techniques to aerodynamic data of airfoils, comparing different algorithms. Providing the foundation for the ML-based regression model developed in this thesis.

The thesis is structured as follows. Chapter 2 reviews previous work on parametric profile modelling, CFD analysis, and regression-based aerodynamic prediction, with a focus on key limitations. Chapter 3 presents a workflow diagram outlining the research method. Chapter 4 presents the extended parametric model, introducing additional shape parameters and ensuring robustness across configurations. Chapter 5 covers the CFD setup, including mesh sensitivity analysis and the rationale for omitting transition modelling. In Chapter 6, the aerodynamic influence of the new parameters is evaluated using Reynolds-averaged Navier-Stokes (RANS) simulations. Chapter 7 explores the development of an ML regression model for aerodynamic prediction, covering the full process from data preparation to algorithm evaluation. Chapter 8 demonstrates the regression model implementation for 2D LEI kite profile design optimisation. Finally, Chapter 9 concludes with key findings and outlines directions for future research.

2
Literature Study

This chapter reviews the current state of aerodynamic modelling for LEI kites. It begins with an overview of simplified FSI methods in Section 2.1, followed by a summary of past 2D RANS studies in Section 2.2 and a review of parametric profile geometry models and limitations in Section 2.3. It continues with a discussion of CFD practices covering meshing, turbulence, and transition modelling in Section 2.4, then reviews aerodynamic regression approaches in Section 2.5. The chapter concludes in Section 2.6 by identifying key gaps and formulating the research questions guiding this thesis.

2.1. Fluid Structural Interaction (FSI)

LEI kites, composed of a lightweight canopy supported by inflated struts and a LE tube, maintain structural stiffness while allowing the flexibility required for flight control and operations. However, their design is complex due to the strongly coupled FSI. Traditional high-fidelity FSI simulations, combining FEM with CFD solvers and robust mesh handling, are computationally expensive and unsuitable for design optimisation where many configurations must be evaluated [12].

To address this challenge, a potential flow method, the vortex step method (VSM), offers a solution by providing accuracy comparable to CFD simulations while significantly reducing computational cost [12]. In parallel, the particle system model is a simplified FEM approach that captures the kite's structural response [38]. When coupled in a loop, these models enable rapid prediction of kite aero-structural behaviour at a fraction of the computational cost of traditional FEM-CFD simulations.

The VSM extends the classic lifting line method (LLM) by computing force magnitude at the three-quarter chord and direction at the one-quarter chord, enhancing accuracy for unconventional geometries with low aspect ratios and high anhedral angles [12]. It models the wing as a spanwise arrangement of horseshoe vortices and uses local velocity and 2D viscous polar data to iteratively solve the aerodynamic loads.

These loads drive the particle system model, where the kite's structure is coarsely discretised into spanwise rectangular segments, defined by the strut locations. Particle points are placed at the LE and trailing edge (TE) of each strut and are connected by bridle lines modelled as spring-damper systems. The model includes additional springs at the LE and TE to provide structural stiffness, and diagonal springs within each segment to resist shear deformation. Load transfer is applied at the intersection points of the diagonal springs and is distributed approximately $75-25\%$ between the front and rear bridle points, respectively, according to the typical load distribution Poland et al. [38].

Following the implementation of the VSM and particle system model, integrating an ML-based aerodynamic regression model offers a promising path to enhance the accuracy of the FSI framework and increase flexibility in design optimisation, by replacing the current look-up tables or low-order polynomial regression models developed by Breukels [9] and revised by Watchorn [50].

With the availability of a predictive model for the C_p distribution, the particle system model could be improved in several ways. Adding particle points along the struts could enhance the representation of chordwise billowing in both the canopy and the struts. The current fixed $25 - 75\%$ load split between rear and front bridle points could be replaced by a dynamically computed distribution, informed by the predicted C_p or derived x_cp values from C_l and C_m. Additionally, increasing the spanwise particle resolution between struts guided by local C_p predictions could improve structural accuracy and reduce coupling errors with the more finely discretised VSM model.

On the aerodynamic side, incorporating predicted C_p values could refine the estimation of local velocities at the three-quarter and one-quarter chord points. With the local angle of attack (AoA) available per segment, a 2D polar regression model trained on a wide range of airfoils, AoA, and Reynolds numbers could provide more accurate aerodynamic loading predictions, significantly improving both global accuracy and design flexibility.

2.2. Overview 2D LEI Kite Profile Studies

This section presents an overview in Table 2.1, which summarises the two-dimensional RANS studies of LEI kite profiles and flexible membrane wings conducted to date. The overview was initially compiled by Demkowicz [18] and has since been extended to include more recent contributions. Collectively, these studies form the aerodynamic foundation upon which the present research is built. Among the listed works, only Breukels [9] and Watchorn [50] developed polynomial regression models based on their CFD results. Watchorn [50] further developed the model introduced by Breukels [9], improving its applicability by incorporating additional LEI kite profile geometry parameters. The RANS-based turbulence models, including the $k-\omega$ shear stress transport (SST) and the $k-\epsilon$ model, are introduced in Section 2.4.3.

Table 2.1: Overview of 2D RANS studies on flexible membranes and LEI kite profile.

Reference	Geometry	Steady / Unsteady	Solver	Turbulence Model	Transition Model	Mesh Type
Watchorn (2023) [50]	LEI kite	Steady	OpenFOAM simple-FOAM	RANS $k-\omega$ SST	-	Structured O-mesh
Folkersma et al. (2019) [23]	LEI kite	Steady	OpenFOAM simple-FOAM	RANS $k-\omega$ SST	γ-$\hat{\mathrm{Re}}_\theta$	Structured O-mesh
Steiner (2018) [44]	LEI kite	Unsteady	OpenFOAM preCICE	RANS $k-\omega$ SST	-	Structured O-mesh
Breukels (2011) [9]	LEI kite	Steady	Ansys Fluent	RANS $k-\omega$ SST	-	Hybrid mesh
Smith et al. (1996) [42]	Flexible membrane wing	Steady	-	RANS $k-\omega$ SST	-	Structured H-mesh
Smith et al. (1995) [43]	Flexible membrane wing	Steady	-	RANS $k-\epsilon$	-	Structured H-mesh

2.3. Geometry & Parametric Model

Developing a robust and flexible simulation setup depends on a solid understanding of existing modelling approaches and their limitations.

2.3.1. Geometry

LEI kite profiles differ significantly from conventional airfoils, mainly due to their distinct circular leading edge and attached canopy, which makes them easily recognisable. These profiles, also known as sailing airfoils, have been extensively studied in the context of sailing. Despite their widespread application across various fields, no standardised or universally accepted equations define the shape of LEI kite profiles.

In earlier studies, researchers often used the kite design software SurfPlan [1] to export profile coordinates, as done by Folkersma et al. [23], Deaves [16], and Sachdeva [40]. However, the limited number of profiles available in SurfPlan prompted Breukels [9] to develop a parametric model for LEI kite profiles. Similar to National advisory committee for aeronautics (NACA) airfoils [29], using shape parameters to generate coordinates via interpolation. Breukels' model provides more direct control over the profile using two parameters: the relative LE tube diameter t and camber height κ, both normalised by the chord length c, as shown in Figure 2.1. In contrast, Kappel [27] defined camber relative to the chord line, aligning with the NACA convention as adopted by Berens [5] and Bosch et al. [7], illustrated in Figure 2.2.

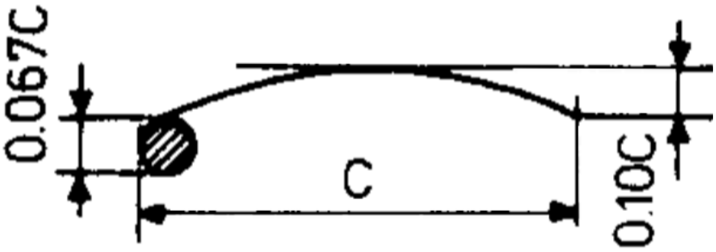

Figure 2.1: LEI kite profile parameterisation as defined by Boer [6].

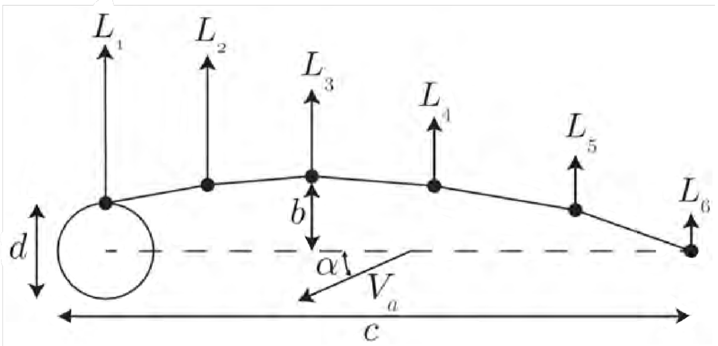

Figure 2.2: LEI kite profile parameterisation as defined by Kappel [27].

Additionally, Watchorn [50] introduced the chordwise position of the camber as a parameter to enhance design flexibility and expand the range of shape configurations. Including this parameter allows for better control, making the LEI airfoil more comparable to a NACA 4-digit airfoil with similar parameters [29]. Based on experimental studies of a sailing airfoil, Boer [6] observed that the camber height of a flexible membrane airfoil shifts toward the LE as the AoA increases. Following this insight, both Kappel [27] and Cayon [11] proposed incorporating the chordwise position of maximum camber as a key airfoil shape parameter in Breukels [9] model. As a result, the updated model from Watchorn [50] includes three parameters: camber height κ, LE tube diameter t, and the chordwise position of the maximum camber η, all normalised by the chord length, providing more precise control over the airfoil's shape.

Recently, Corentin [15] introduced a TE angle, also known as a reflex angle, into the parametric profile model, as shown in Figure 2.3. This feature is similar to the reflex commonly found in NACA 5-series airfoils [29]. For larger kites, the reflex angle helps prevent the kite from flying excessively high toward the zenith, a condition that can cause longitudinal pitching moment instabilities due to low apparent wind. By incorporating a reflex, the profile produces a positive moment and increases the angle of attack, reducing the risk of nose stalls.

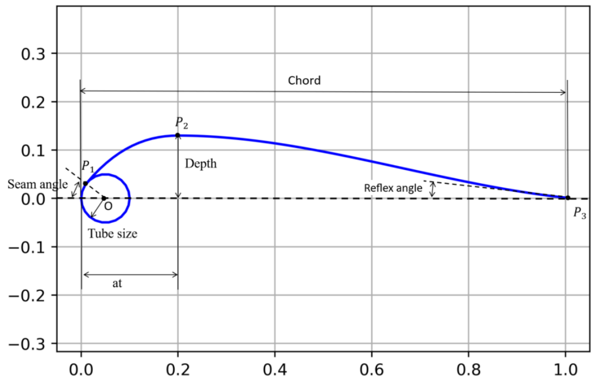

Figure 2.3: LEI kite profile parametric model with reflex angle [15].

The canopy is connected to the LE tube at a defined angle known as the seam angle, also visible in Figure 2.3. Various seam angles have been evaluated through trial and error for different maximum camber positions. In the model by Watchorn [50], the seam angle is fixed at $20°$, while Corentin [15] applies a value of $35°$. The canopy shape is constructed based on the reflex angle, LE seam angle, and maximum camber location, using either cubic polynomials or a combination of Bézier curves and cubic polynomials. To ensure a smooth geometric transition, tangency conditions are enforced at both the LE tube-canopy connection and the maximum camber point.

2.3.2. Limitations

The current parametric profile model already provides substantial geometric control. However, the reflex angle has not yet been incorporated into the regression model. Moreover, analysis of different LE tube diameters and camber locations reveals that the seam angle should vary accordingly. Fixing this parameter can result in issues such as reduced smoothness or invalid geometries. Incorporating a dynamic seam angle that adapts to LE tube diameter and camber would be beneficial for producing a broader range of smooth profile shapes.

At present, a large portion of the profile is constructed using cubic polynomials, which offer limited flexibility over the extended region defined by the parametric points. While additional points can be introduced to refine the shape, this makes it more difficult to maintain overall smoothness. Expanding control over LE canopy curvature, reflex angle, and canopy tension would broaden the range of achievable profile shapes, enabling more effective optimisation across various applications.

2.4. Computational Fluid Dynamics (CFD)

Due to the unconventional shape of the LEI kite profile and the large recirculation region on its lower surface, along with the need for accurate separation modelling near stall on the upper surface, XFOIL's potential flow assumptions are insufficient [19]. Therefore, CFD analysis is required.

2.4.1. Mesh Setup

To simulate the freestream interaction with the LEI kite profile, the surrounding space is discretised into a mesh. The quality of the mesh is crucial for accurately replicating actual conditions and relies on the level of grid refinement as well as specific quality parameters, such as skewness, orthogonality, aspect ratio, etc. Much of the meshing setup used here is based on Watchorn's [50] recent work, where he employed the commercial Pointwise [37] software controlled by Matlab, which inputs the profile and meshing characteristics.

Before generating the mesh, slight adjustments are made to the profile to enhance mesh quality by addressing the parameters mentioned. Mesh quality tends to deteriorate near sharp corners and edges, such as the area where the tube and canopy profiles intersect. A fillet is introduced to mitigate this, balancing the trade-off between mesh quality and representing the actual geometry. Deaves [16] compared aerodynamic performance with and without a fillet across various angles with RANS simulations, concluding that the lift and drag differences were within 5%, which falls within the error range when comparing the RANS simulation of a NACA0012 profile with experimental data. The fillet is visible in Figure 2.4(c). Leloup et al. [33] also conducted an analysis using XFOIL [19]. However, the results are of limited validity, as XFOIL cannot accurately capture the recirculation region on the pressure side of the airfoil.

The canopy thickness t_{canopy}, considered negligible, is included in the work of Watchorn [50], Deaves [16], Folkersma et al. [23], Demkowicz [18], to distinguish airflow between the upper and lower surfaces. A larger t_{canopy} improves mesh quality at the TE but is less representative of the actual membrane. Watchorn found 0.001 m to be a reasonable compromise but did not assess its impact on residual convergence or aerodynamic performance.

A finite canopy thickness introduces a gap in the mesh at the TE. To address this, Watchorn [50] implemented a semi-circular TE with a diameter equal to t_{canopy}. This rounded geometry is well-suited for generating an O-grid mesh, offering improved mesh quality near the TE. In contrast, closing the gap with a straight line would be more appropriate for a C-grid mesh topology. The semi-circular TE implementation is illustrated in Figure 2.4(d).

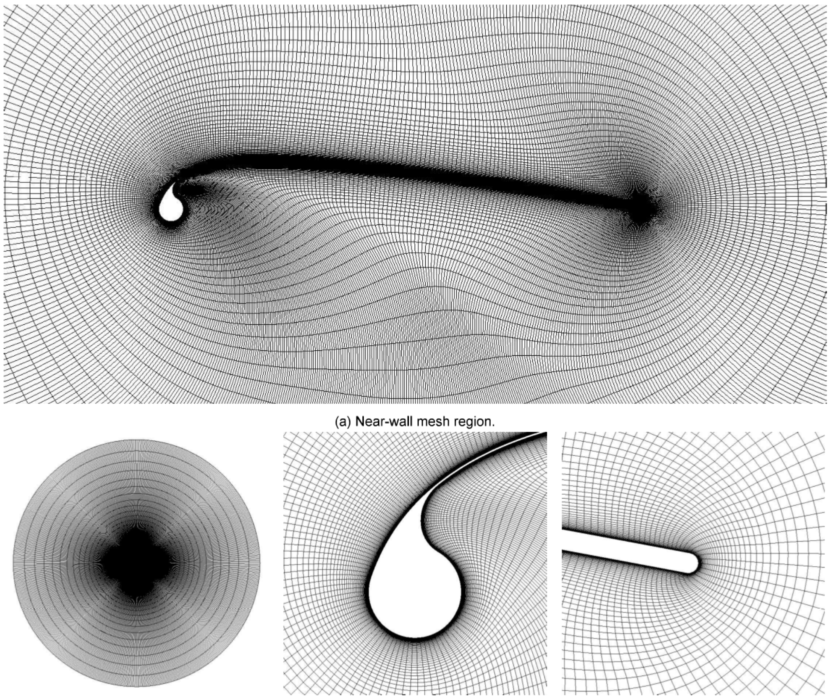

(a) Near-wall mesh region.

(b) Full mesh domain.

(c) Circular LE tube and smooth edge.

(d) Semi-circular TE.

Figure 2.4: Structured mesh with O-grid topology of a LEI kite profile with LE fillet and round TE.

Regarding the meshing structure, there are options to use a fully structured, unstructured, or hybrid mesh, which combines both types. Watchorn opted for a structured mesh, which offers better accuracy but requires more intricate work to achieve high precision. For the topology, Watchorn used an O-grid, as shown in Figure 2.4. Folkersma et al. [23], Steiner [44] also used this same topology in there analyses. A summary of previous studies and their corresponding mesh types is provided in Table 2.1.

Another commonly used topology for discretising the space around an airfoil is the C-grid. Both C-grid and O-grid topologies have their advantages and drawbacks. The main reason for choosing an O-grid over a C-grid is that the flow remains orthogonal to the mesh cells at both the inflow and the wake, ensuring better alignment. In contrast, a C-grid requires more effort to maintain mesh quality, particularly at larger angles of attack, where wake alignment becomes necessary to improve orthogonality.

However, the O-grid's circular nature results in equal grid lengths in all directions, which is less efficient, as the far-field at the TE requires a larger area to properly capture the wake region compared to the front, lower, and upper sections of the mesh. In summary, while a C-grid can offer higher accuracy for the same resolution, it involves the additional complexity of aligning the mesh with the wake.

Although Watchorn's mesh analysis was thorough, no sensitivity study was carried out using alternative meshing algorithms, topologies, or mesh types. Moreover, the influence of canopy thickness t_{canopy} on mesh behaviour was not explored. Incorporating such sensitivity analyses into the design process could improve simulation accuracy and lead to better residual convergence.

2.4.2. Direct Numerical Simulation (DNS) vs Large Eddy Simulation (LES)

Direct numerical simulation (DNS) offers the highest fidelity in turbulence modelling by solving the Navier–Stokes equations without approximations, resolving all spatial and temporal scales down to the Kolmogorov scale. However, its computational cost scales with Re^3 [22], making it impractical for high Reynolds number flows such as those around LEI airfoils, which can reach 10^7. This limitation becomes even more restrictive when regression models need to be developed over a wide range of geometries and operating conditions.

Nevertheless, DNS studies have been performed on LEI kite profiles at low Reynolds numbers (e.g., $\text{Re} = 1000 - 5000$), providing valuable insight into flow behaviour in simplified scenarios [46].

Large-eddy simulation (LES) offers a compromise between accuracy and efficiency by resolving only the larger, more energetic eddies while modelling the smaller scales through subgrid-scale models [22]. Although less expensive than DNS, LES still requires a fine spatial grid and time step resolution to capture unsteady wake dynamics, particularly at high Reynolds numbers.

2.4.3. Steady Reynolds-Averaged Navier-Stokes (RANS) Simulation

Reynolds-averaged Navier-Stokes (RANS) methods are a widely used approach for predicting average flow properties, such as lift and pressure distributions, at a fraction of the computational cost of DNS or LES [22]. By time-averaging the Navier–Stokes equations, all turbulent fluctuations are treated statistically, resulting in the loss of unsteady flow details and requiring closure models for the unresolved Reynolds stresses.

Despite reduced accuracy, RANS remains the standard in engineering applications due to its efficiency, and is commonly used in LEI kite simulations by Folkersma et al. [23], Watchorn [50], and others listed in Table 2.1.

Turbulence Modelling

Turbulence models are essential for closing the RANS equations. Among the most widely used are the two-equation models: $k - \epsilon$ and $k - \omega$.

The $k\epsilon$ model [31] solves transport equations for turbulent kinetic energy (k) and its dissipation rate (ϵ). It is robust and efficient for fully developed turbulence in free-stream flows but performs poorly near walls or in flows with strong pressure gradients.

The $k - \omega$ model, introduced by Wilcox [52], replaces ϵ with the specific dissipation rate (ω), improving accuracy in near-wall and adverse pressure gradient regions, though it is sensitive in free-stream conditions.

To address these limitations, the $k - \omega$ SST model [34] blends both approaches: it uses the $k - \omega$ model near walls and transitions to $k - \epsilon$ in the free stream. It also accounts for shear stress transport, making it highly effective for separated and complex boundary layer flows.

The SST model is the preferred turbulence model in LEI kite analyses and has been applied by Folkersma et al. [23], Watchorn [50], Steiner [44], and Breukels [9], as summarised in Table 2.1.

Transition modelling

Accurately modelling the boundary layer transition from laminar to turbulent flow is challenging due to the various factors that can trigger this transition. Typically, three primary modes of laminar-turbulent transition are considered, each driven by different influences [23].

1. Natural transition: Perturbations from the exterior flow, such as turbulence, propagate into the boundary layer and force the laminar flow to transition into turbulent flow.

2. Bypass transition: The freestream disturbances are high (turbulence intensity) and propagate directly into the boundary layer, or the rough surface triggers the transition.
3. Separation-induced transition: A laminar boundary layer cannot withstand much adverse pressure gradient, so it separates rather easily. Separated flow transitions quickly, and the turbulent flow may reattach from the enhanced momentum.

Folkersma et al. [23] conducted a study investigating the impact of incorporating a transition model (natural transition) versus a fully turbulent flow on an LEI kite profile across a wide range of Reynolds numbers. Representative of the traction and retraction phases in airborne wind energy (AWE) systems, covering Re values between 1×10^5 and 5×10^7. This study implemented a transition model, such as the $\gamma - \tilde{R}e_{\theta t}$ model [30], to predict where flow over an airfoil transitions from laminar to turbulent. Both approaches exhibit quite different results with distinct flow characteristics given in Figure 2.6 and Figure 2.7. The simulation results were compared against experimental data at a low Reynolds number of Re $= 10^5$.

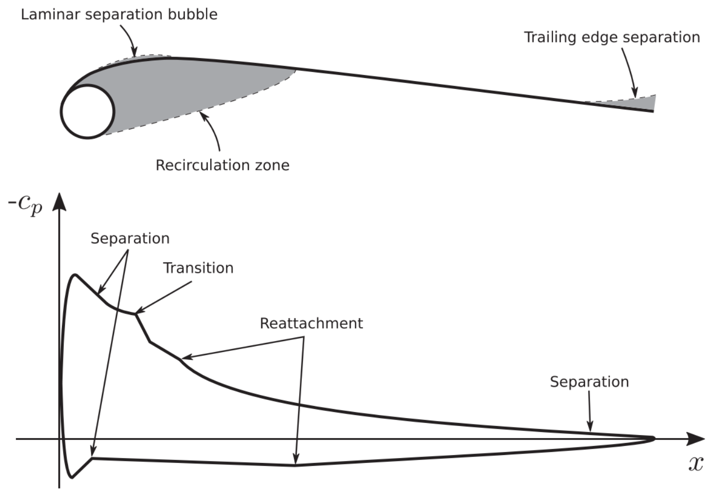

Figure 2.5: Flow topology around a LEI kite airfoil, adapted from Folkersma et al. [23]

The results demonstrated that the simulations with a transition model more closely matched the experimental data than those without. However, it should be noted that the wind tunnel tests were conducted using a LEI kite profile with a smooth surface at low turbulence intensity (I), approximately 2%. Under these conditions, it is logical that simulation with a transition model lies closer to the experimental data. The transition model can capture the laminar regions, turbulent regions, and the laminar separation bubble, which are prominent at low Reynolds numbers. Unlike the fully turbulent simulation, a different pressure distribution is achieved without a laminar separation bubble and laminar flow.

A similar phenomenon appears at the pressure side for the large recirculation region depicted in Figure 2.5. With the transition model, the laminar boundary layer separates earlier than the turbulent boundary layer in the fully turbulent model. Creating a larger recirculation region with the transition model for low Reynolds numbers around Re $= 10^5$.

Folkermsa's report comprehensively analyses the observed differences with and without the transition model. Still, the key distinctions are highlighted here to underscore the importance of selecting the most accurate model that reflects real-world conditions.

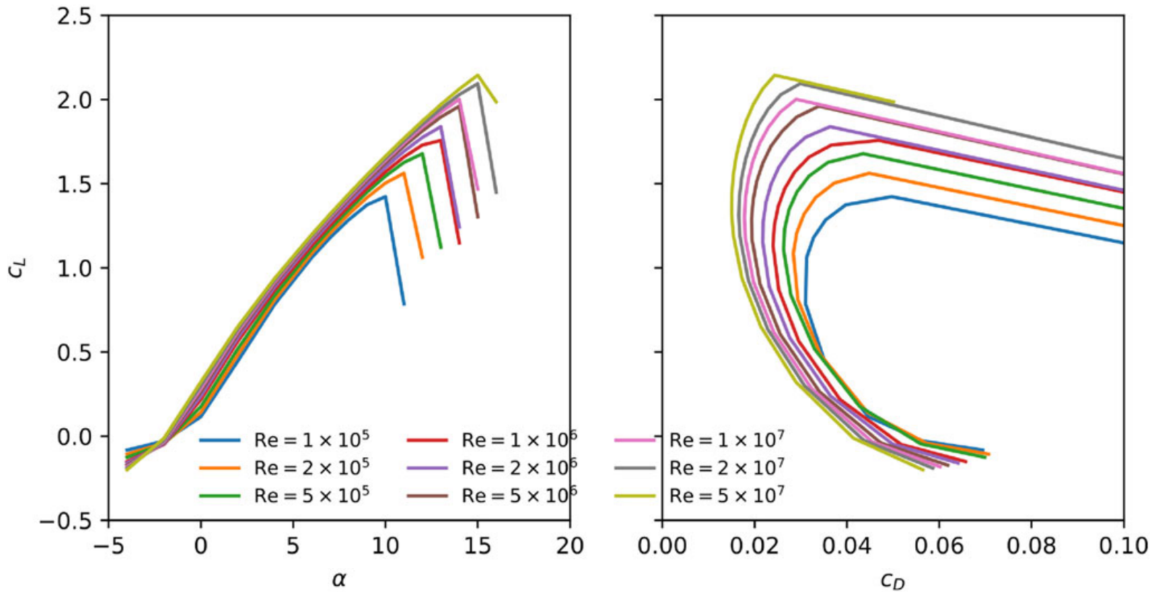

Figure 2.6: Lift polar (left) and drag polar (right) for a LEI kite profile at various Reynolds numbers, computed without transition model [23].

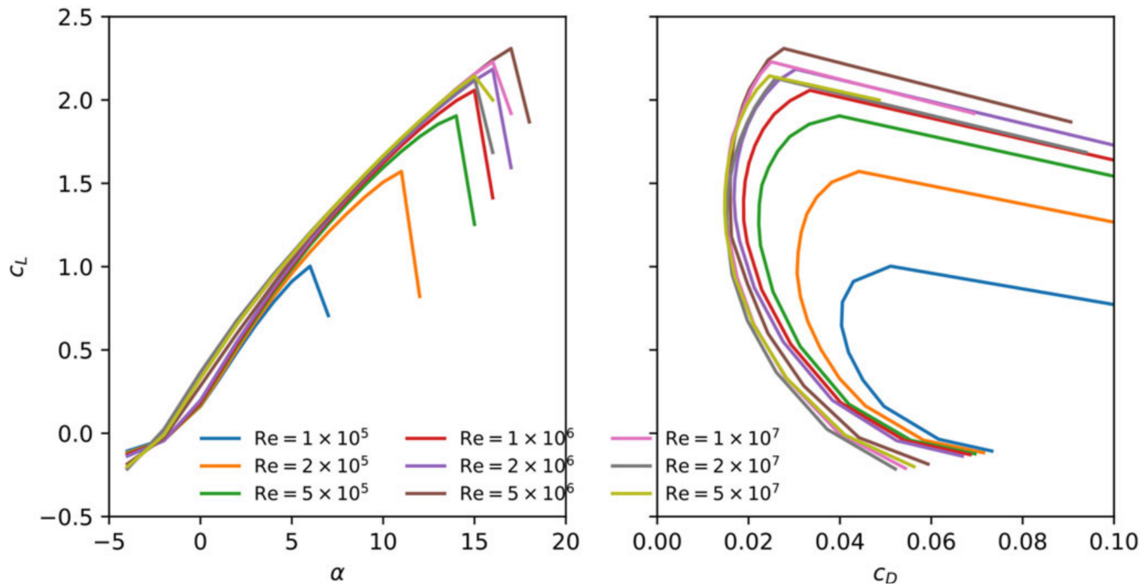

Figure 2.7: Lift polar (left) and drag polar (right) for a LEI kite profile at various Reynolds numbers, computed with transition model [23].

For Reynolds numbers below 5×10^6, the aerodynamic performance predicted using the transition model in Figure 2.7 shows lower lift and higher drag. This is due to the formation of a long laminar separation bubble on the suction side of the airfoil, as opposed to the fully turbulent flow seen in Figure 2.6. For higher Reynolds numbers, the distinction in maximum lift and drag coefficients is due to the presence of a laminar region when using a transition model different to fully turbulent flow. Beyond 2×10^7, the performance is similar in both cases as the flow naturally transitions at the front of the airfoil, inducing fully turbulent flow. Highlighting the importance of considering transition effects at Reynolds numbers below 2×10^7.

Figure 2.8: Stitching seams with canopy overlay near the upper surface of the LE tube.

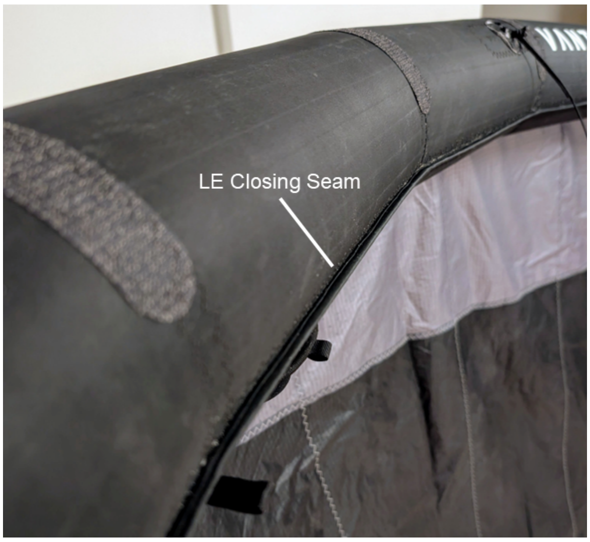

Figure 2.9: LE tube closing seam located on the underside of the LE tube.

Unlike the smooth surface assumed in the validation model by Folkersma et al. [23], the kite's canopy is made of textile material with stitched seams, as shown on both the upper and lower sides of the LE tube in Figures 2.8 and 2.9. Due to the presence of seams near the LE, Watchorn [50] omitted a detailed transition model, assuming that the surface irregularities would cause a bypass transition, immediately triggering the boundary layer to transition from laminar to turbulent flow.

2.5. Regression Model

First, Breukels [9] introduced the use of polynomial curve fitting to model aerodynamic performance as a function of AoA, κ, and the LE tube diameter. This approach was later refined by Watchorn [50], who maintained the polynomial structure while improving the model's applicability.

In Watchorn's revised model, the aerodynamic coefficients (C_l, C_d, C_m) are predicted using polynomial regressions: C_l is fitted with a cubic polynomial in α, while C_d and C_m are approximated using quadratic terms. The model was trained on 256 CFD data points derived from 64 unique LEI kite profiles, each evaluated at four angles of attack ($\alpha = 0°$, $5°$, $10°$, $15°$) at a constant Reynolds number of 5×10^6. As shown in Table 2.2, the R^2 and root mean square error (RMSE) values do not always show consistent trends. This is expected, as RMSE is sensitive to the scale of the output variable: for instance, C_l typically larger in magnitude naturally yields a higher RMSE despite a strong R^2. Additionally, the lack of variance or confidence intervals makes it difficult to assess the statistical reliability of the results.

Future improvements could include expanding the dataset with more diverse input combinations and transitioning from polynomial regression to ML algorithms, which are better suited for capturing complex nonlinear relationships.

Table 2.2: Polynomial regression model performance for the aerodynamic coefficients, developed by Watchorn [50].

Coefficient	R^2	RMSE
C_l	0.9582	0.1181
C_d	0.8971	0.0184
C_m	0.9791	0.0088

2.6. Research Question

Based on the latest parametric model by Watchorn [50] and the limitations identified in the literature review (Section 2.3.2), the current model lacks support for incorporating a reflex angle at the TE and offers limited control over the rest of the profile both of which are essential for capturing the wide range of shape variations caused by the flexible membrane. Additionally, there is insufficient control over the seam angle and LE fillet geometry, affecting smoothness and robustness. These shortcomings lead to the following research questions for further development of the profile geometry model:

1.1 How can the parametric model of the LEI kite profile be extended to cover the wide range of shape variations that occur during flight?

1.2 What is the aerodynamic effect of incorporating a reflex angle into the LEI kite profile?

1.3 What is the aerodynamic effect of the LE fillet size?

The literature highlights uncertainty around the transition of the upper and lower seams present on the LE tube. This unresolved aspect motivates the following CFD research questions:

2.1 Do the LE tube-canopy intersection seam and the closing seam located on the upper and lower surfaces of the LE tube, respectively, trigger the boundary layer to transition from laminar to turbulent flow?

2.2 Is the boundary layer transition behaviour associated with the flow conditions, and to what extent?

Finally, answering the research questions above provides the foundation of aerodynamic data used for the regression model, enabling improvements to the revised polynomial regression model by Watchorn [50]. The following research questions specifically address how the regression model itself can be further enhanced:

3.1 To what extent can machine learning enhance regression models for predicting the aerodynamic performance of LEI kite profiles?

3.2 What are the optimal parameter bounds and resolution requirements needed to develop a reliable and applicable aerodynamic prediction model?

3 Research Method

The research methodology is summarised in the flowchart in Figure 3.1, which outlines three main building blocks. These guide the investigation into improving the parametric profile model, enhancing CFD accuracy, and enabling automated data generation for more reliable and accurate regression models.

The first part of the workflow focuses on extending the profile geometry model by introducing additional parameters and spline interpolation methods to capture a wider range of profile shapes. It also includes an investigation into potential dependencies for implementing a robust, dynamic LE fillet and seam angle. These steps are examined in detail in Chapter 4.

In parallel, mesh sensitivity studies and CFD setup verification are carried out. These tasks, covered in Chapters 4 and 5, address the implementation of transition modelling and the generation of automated meshes capable of reliably capturing recirculation and separation regions.

The next phase involves analysing the aerodynamic influence of the newly introduced parameters across the defined parameter space, as detailed in Chapter 6. The data generated will be used to explore the development of an ML regression model. This involves investigating various preprocessing techniques, algorithm choices, and postprocessing methods, as well as evaluating model performance using appropriate scoring metrics, all of which are discussed in Chapter 7.

Finally, the regression model implementation for 2D LEI kite profile design optimisation is investigated in Chapter 8.

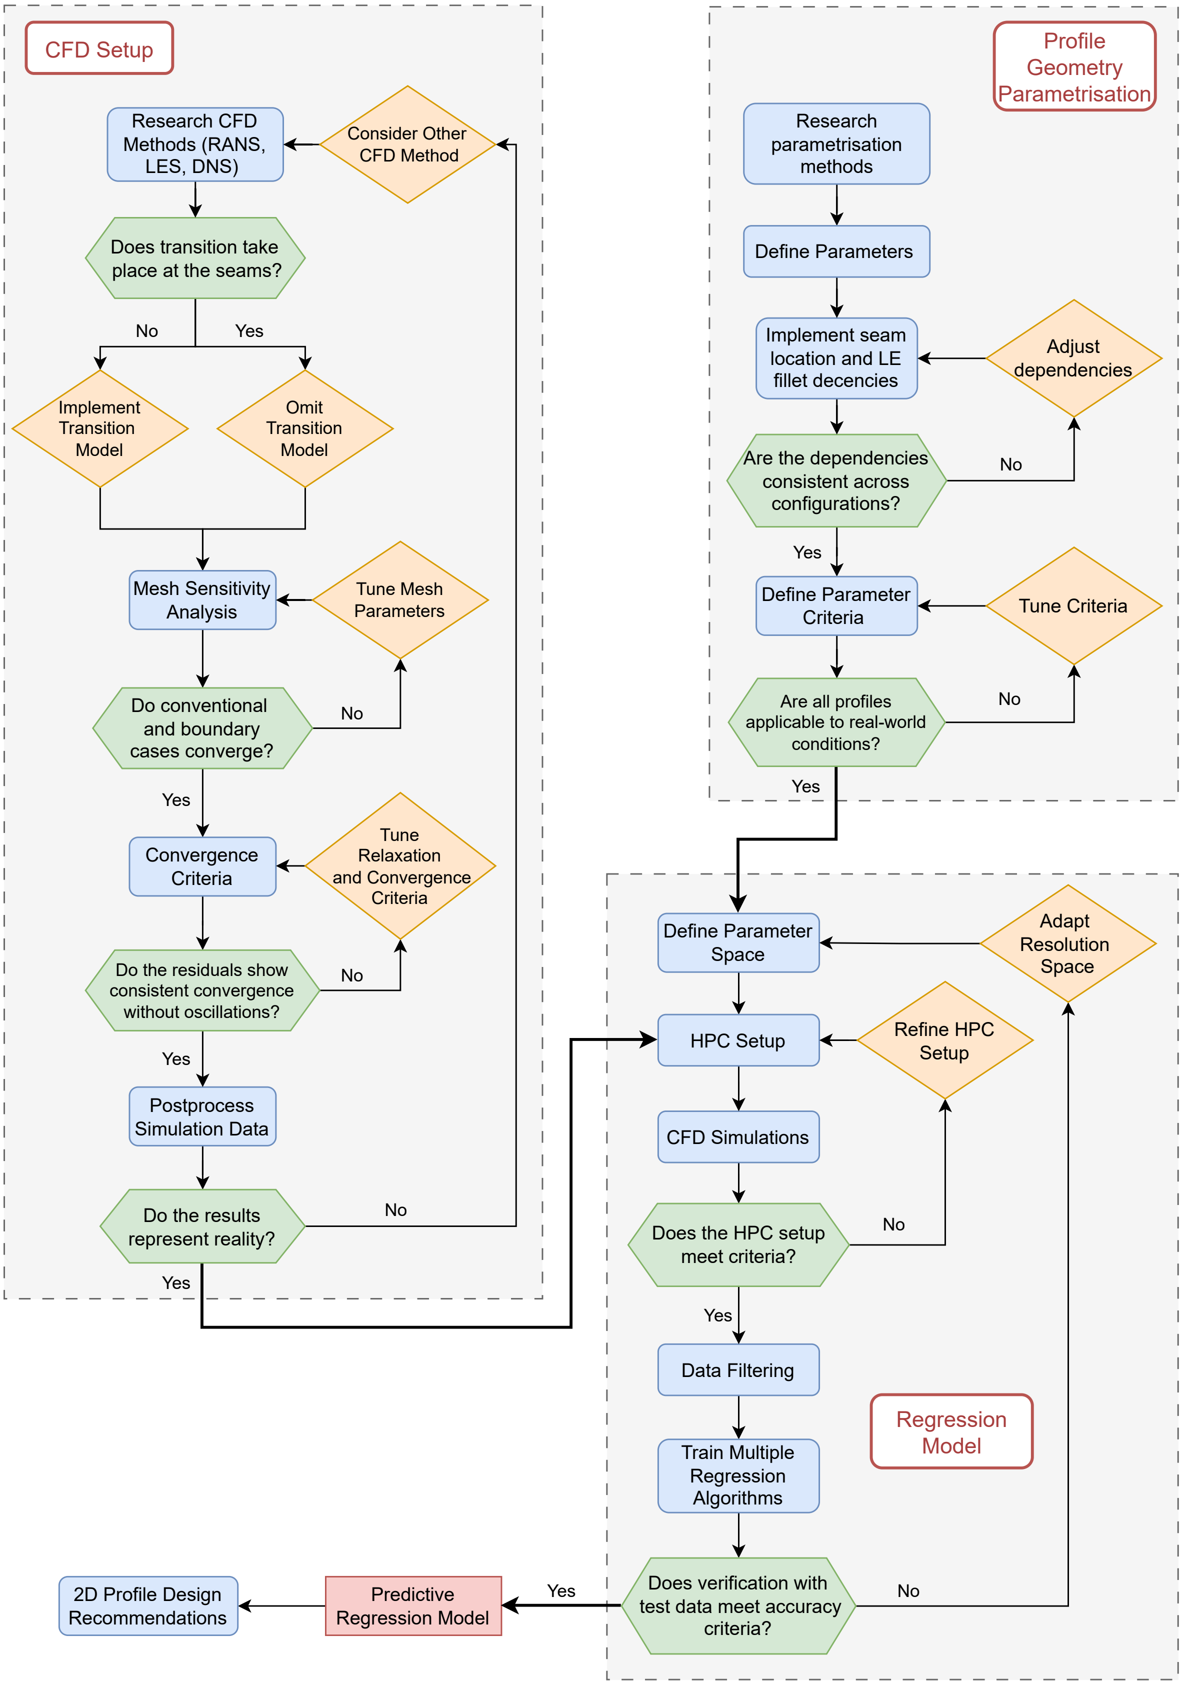

Figure 3.1: Workflow diagram of the methodology.

4

Profile Geometry Parameterisation and Mesh Generation

Generating a high-quality mesh for CFD simulations is a demanding task that requires careful attention to detail, particularly when working with the unconventional geometries of LEI kite profiles. Since a large number of profiles need to be analysed, a parametric model is first developed in Section 4.1 to enable automated geometry and mesh generation. The corresponding mesh topology and structure are then addressed in Section 4.2, with a comparison of the commercial software Pointwise and the open-source Python-based application Gmsh. To ensure accurate and reliable results across the design space, mesh parameters and smoothing strategies are fine-tuned in Section 4.3, where a comprehensive mesh sensitivity study is performed.

4.1. Profile Geometry Parameterisation

This section establishes the foundation for a broadly applicable regression model by developing a parametric profile geometry model with multiple design parameters, with a focus on surface mesh reliability and robust automation.

4.1.1. Geometric Components, Constraints & Parameters

The literature study provides a comprehensive overview of the evolution of the kite profile model in Section 2.3, beginning with Breukels [9], who laid the foundation by defining the circular LE tube diameter t and the camber height κ. Based on earlier experimental studies of flexible sailing airfoils, Boer [6] observed that the maximum camber position shifts toward the leading edge as the angle of attack increases. Building on this insight, Kappel [27] and Cayon [11] proposed incorporating the chordwise position of maximum camber η as an additional design parameter, an idea that Watchorn [50] later implemented to enhance design flexibility. More recently, Corentin [15] introduced a reflex angle δ at the TE, which increases the longitudinal pitching moment C_m by locally deflecting the canopy at the TE downward, thereby helping prevent kites from climbing too high toward the zenith and entering an unrecoverable front stall.

In addition to the current model, this work introduces a dynamic seam angle function to improve smoothness at the LE tube-canopy intersection, along with two cubic Bézier curves, which increase geometric flexibility through additional control point parameters and ensure smooth curvature along the canopy profile.

Constraints

Before introducing these new implementations, it is essential to establish the constraints that the parametric model must adhere to. The airfoil chord is set to a unit length of one, with all parameters non-dimensionalised relative to the chord length. TE is positioned at the standard airfoil coordinate $(1, 0)$, while the LE is located at $(0, 0)$. A point constraint is applied to the centre of the circular tube, which has a non-dimensional radius of $t/2$, placing it at $(t/2, 0)$ to ensure that the leftmost point, the LE, remains fixed at $(0, 0)$.

The profile is constructed by first defining the circular LE tube, followed by the front spline extending from the LE tube-canopy seam to the maximum camber location, and finally the rear spline connecting the maximum camber point to the TE. Tangent continuity constraints are enforced at both junctions to ensure smooth transitions between the segments.

Bézier Curve Splines

The two canopy splines are defined using cubic Bézier curves, as given in Equation 4.1, with parameter t_b controlling the point interpolation along each spline. The four control points, $[P_0, P_1, P_2, P_3]$, define the shape of the curve: the first and last points set the start and end positions of the spline, while the two intermediate points control the tangency and curvature at the intersections and the TE reflex angle.

With two Bézier splines used to define the profile, four additional parameters are introduced through the placement of the control points. Depending on t_b, the Bézier curve equation generates a dense set of points along the profile in the form of $\begin{bmatrix} x & y \end{bmatrix}^\mathsf{T}$ vectors.

$$\begin{bmatrix} x \\ y \end{bmatrix} = (1 - t_\mathsf{b})^3 P_0 + 3(1 - t_\mathsf{b})^2 t_\mathsf{b} P_1 + 3(1 - t_\mathsf{b}) t_\mathsf{b}^2 P_2 + t_\mathsf{b}^3 P_3 \,. \tag{4.1}$$

During the decision-making process for spline interpolation, class shape transformation (CST) variables were also considered, as they produce exceptionally smooth profiles and yield excellent results for optimisation. However, they did not meet the criteria for this study. While an existing profile can be easily represented by a set number of CST variables, directly deriving these variables from the profile parameters to accommodate reflex angle modifications or shifts in the maximum camber position is not feasible. This does not resolve the issue, as an initial spline still needs to be generated precisely for the profile geometry.

Dynamic seam angle

Watchorn [50] primarily studied highly cambered profiles, assuming a constant seam angle of $20°$. However, for flatter profiles, this low seam angle often leads to the front spline intersecting the circular LE tube, resulting in invalid geometry. To overcome this, a higher seam angle is required. Therefore, a dynamic seam angle is introduced, enabling automated adjustment across a variety of profile shapes necessary for the regression model.

The dynamic seam angle function first generates a cubic interpolation spline that is tangent to both the LE tube and the maximum camber location. The function iterates through a range of seam angles, starting from $0°$ at $(0, 0)$ up to $90°$ at $(t/2,\, t/2)$. If the seam angle is too low, the spline may create a bump higher than the maximum camber point, which is not feasible. The function continuously monitors the highest point of the spline and identifies the optimal seam angle when it becomes equal to or lower than the maximum camber location.

Parameters

With the additional control points introduced by the two Bézier splines, the profile shape can be defined with greater flexibility. The standard parameters already implemented in Watchorn's model include the LE tube diameter t, the Camber height κ, and its chordwise position η, all non-dimensionalised relative to the chord length to allow simulations based solely on Re. Additionally, the reflex angle at the TE, denoted as δ and introduced by Corentin [15], will be incorporated into the model.

The Bézier formulation also enables the introduction of two additional parameters: λ, the camber tension, and ϕ, the LE curvature. All parameters are listed below and visualised in Figure 4.1.

- t: LE tube diameter
- η: Chordwise position of the maximum chamber location.
- κ: Camber height
- δ: Reflex angle
- λ: Camber tension
- ϕ: LE curvature

The reflex angle δ is defined as the angle between the line connecting the second rear spline control point to the TE, and the line connecting the third control point to the TE. A downward deflection of the third control point corresponds to a negative reflex angle. The distance from the third control point to the TE is 20% of the distance between the second control point and the TE. This placement reduces the number of free parameters and enhances the robustness of the model.

The camber tension λ improves the smoothness of the transition from the maximum camber point to the TE by locally tensioning the spline. It is defined as the non-dimensional distance from the maximum camber point to the second control point, relative to the unit chord length.

The LE curvature parameter ϕ defines the position of the intermediate control points of the front canopy spline. These points are constrained to move along two tangent lines originating from the LE tube and the maximum camber location. The non-dimensional parameter ϕ determines the relative distance of each control point along its tangent line, with larger values shifting the control points closer to their intersection. As a result, larger ϕ values produce more curvature near the centre of the spline, while smaller values concentrate curvature toward the endpoints.

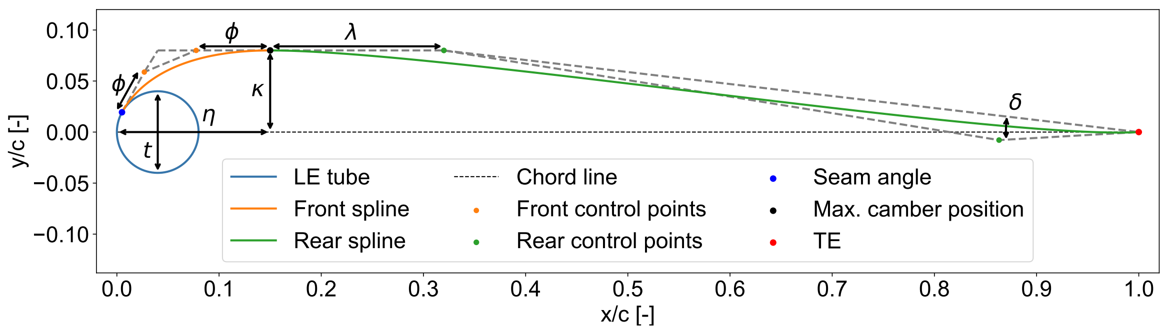

Figure 4.1: LEI kite profile parametrisation using Bézier curves and control points.

A special profile configuration is the flat, zero-camber profile commonly found near the wing tips of a kite. This profile depends solely on the LE tube diameter and the reflex angle, with other parameters fixed for simplicity and robustness. When the maximum camber location falls below the LE tube radius, this flat profile is generated. The seam angle is then positioned at the most suitable tube angle, slightly aft, of $(t/2,\, t/2)$. An example of this configuration, using $t = 0.08$ and $\delta = -10°$, is shown in Figure 4.2. With η fixed at $2t$ and λ at 0.2.

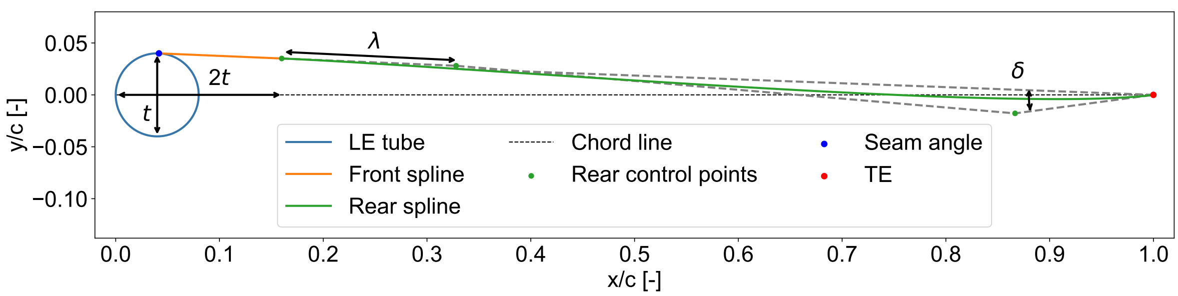

Figure 4.2: Parametrisation of a flat wingtip LEI kite profile using Bézier curves and defined control points.

4.2. Mesh Structure & Profile Geometry Adaptation

With the profile parameters defined, the geometry requires adaptation for CFD analysis. A smooth surface is essential to support mesh generation and accurately capture the boundary layer. This section outlines the mesh type, topology, and automation-friendly meshing platform designed to balance accuracy and computational efficiency.

4.2.1. Geometry Adaptation for Mesh Smoothing

With the parametric profile established, as shown in Figure 4.1, it cannot be directly used for CFD simulations due to sharp corners that cause mesh distortions and the infinitely thin canopy, which allows adjacent mesh cells on the upper and lower surfaces to share flow information. Although the canopy has negligible physical thickness, a finite thickness must be introduced to avoid this issue. This is defined as the non-dimensional canopy thickness, t_{canopy}, relative to the chord length c. Increasing this gap improves mesh quality at the TE but deviates from physical accuracy; its impact is assessed in the sensitivity study in Section 4.3.

To construct the lower canopy surface, two Bézier curves are used: existing control points from the upper surface are shifted downward by t_{canopy} at the maximum camber location and at the second control point of the rear spline. The lower TE point is placed using an orthogonal offset of t_{canopy}, as is the third control point of the rear spline near the TE. For the front spline, all spline interpolation points are uniformly shifted downward by t_{canopy}. The lower splines are presented with an exaggerated canopy thickness of $t_{\text{canopy}} = 0.006$ in Figure 4.3.

The connection between the lower front spline and the tube creates a sharp corner, causing mesh irregularities in this region. This issue is mitigated by incorporating an LE fillet, which facilitates smooth transitions. The LE fillet is constructed using a cubic Bézier curve, ensuring tangent continuity at its connections with both the circular tube and the lower front spline. Similar to the canopy thickness, increasing the fillet region enhances mesh quality but introduces deviations from physical reality. A detailed sensitivity analysis of this effect is provided in Section 4.3. Lastly, the TE is rounded to a semi-circular TE fillet for mesh smoothing in an O-grid topology mesh. The resulting adapted profile is illustrated in Figure 4.3.

Improving spline Smoothness, point density along the Bézier curve splines, t_{b} is allocated as follows: 80 points for the LE tube and the upper front spline, 100 points for both the upper and lower rear splines, 30 points for the rounded TE, and 50 points for the LE fillet.

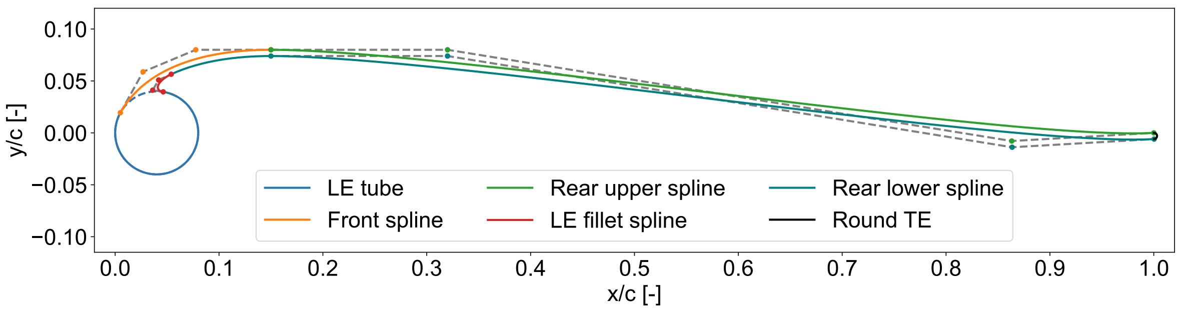

Figure 4.3: Profile modification to achieve smooth and mesh-compatible geometry for CFD with exaggerated $t_{\text{canopy}} = 0.006$.

4.2.2. Pointwise versus Gmsh Mesh Generator Platforms

The search for an automated, high-quality mesh generator focused on the commercial software Pointwise [37] and the open-source application Gmsh [25], both of which provide Python APIs for mesh parameterisation. The main advantage of Gmsh is its free access, extensive documentation, and native integration with OpenFOAM. Pointwise, while less accessible and more complex to automate via Python, offers the ability to generate fully structured meshes using a hyperbolic solver. This results in superior mesh quality, especially for high-curvature geometries like LEI kite profiles. In contrast, Gmsh lacks support for structured hyperbolic meshing. Although structured mesh generation was attempted in Gmsh via automation scripts, the resulting mesh quality did not match that achievable with Pointwise.

As part of the sensitivity analysis, a comparison is conducted between a fully structured grid generated in Pointwise and a hybrid grid created in Gmsh. In the hybrid mesh, structured cell layers are applied near the profile surface to accurately resolve the boundary layer, while the rest of the domain is filled with an unstructured triangular mesh bounded by a circular outer edge in the O-grid topology. The hybrid mesh is illustrated in Figure 4.4, and the structured mesh is shown in Figure 4.5.

The profile used for the mesh comparison in this section is shown in detail in Figure 4.13, corresponding to fillet number 3. It is defined for Re $= 5 \times 10^6$, with the following specifications: $t = 0.08$, $\eta = 0.15$, $\kappa = 0.08$, $\delta = 0°$, $\lambda = 0.2$, $\phi = 0.65$, and $t_{\text{canopy}} = 0.0005$.

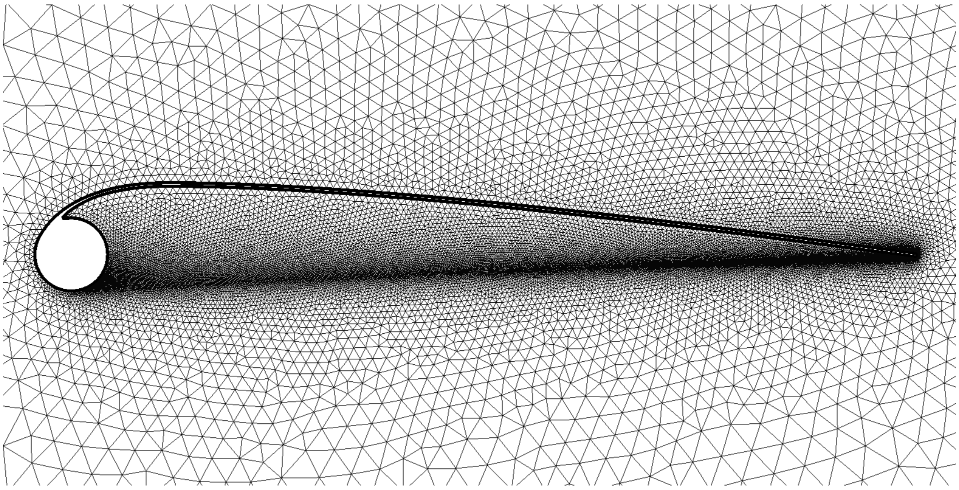

(a) Near airfoil mesh region with structured grid along the surface and unstructured towards the outer boundary.

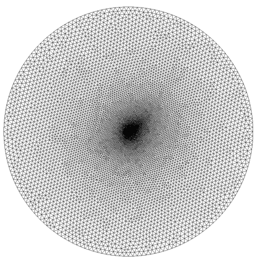

(b) Full mesh domain.

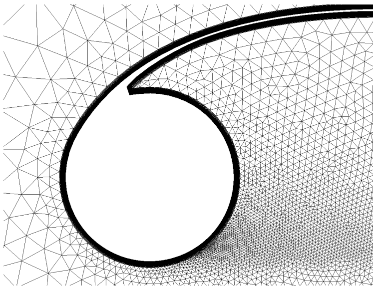

(c) Mesh distribution along the LE section.

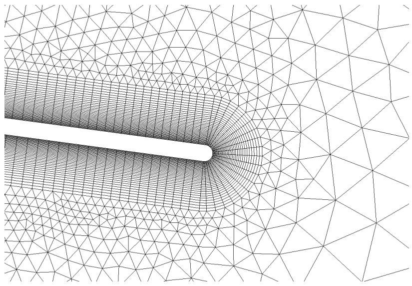

(d) Mesh distribution along the TE.

Figure 4.4: Hybrid mesh with O-grid topology generated in Gmsh [25], with a domain radius of $150c$, and 575 nodes along the profile surface.

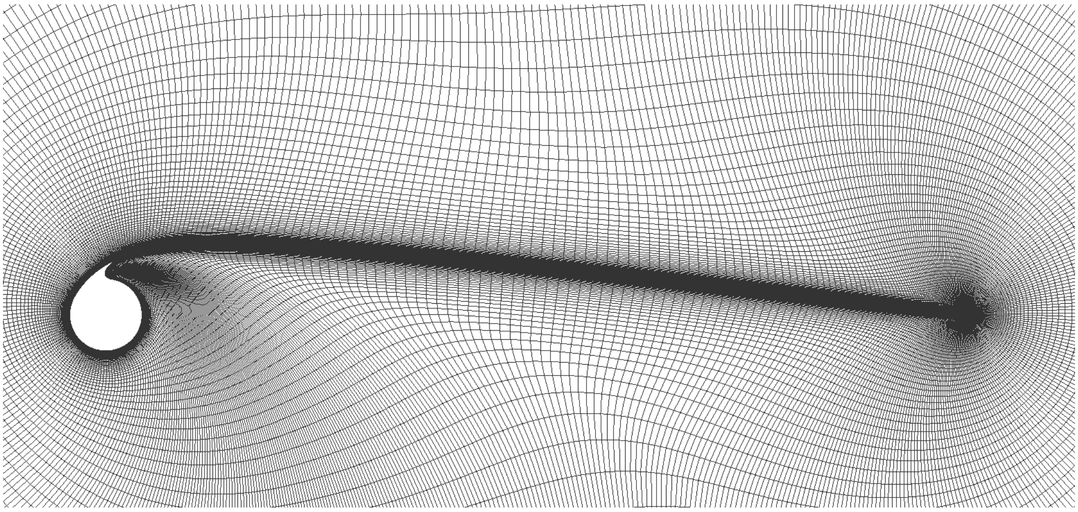

Figure 4.5: Fully structured mesh with O-grid topology generated in Pointwise [37] (575×201) for a LEI kite profile with shape parameters: $t = 0.08$, $\eta = 0.15$, $\kappa = 0.08$, $\delta = 0°$, $\lambda = 0.2$, and $\phi = 0.65$.

The first issue encountered with the hybrid mesh in Gmsh is its inability to adapt the end-node spacing for the hyperbolic distribution along the upper and lower surfaces, which connects the uniform node distributions at the TE and LE. In contrast, Pointwise supports this adaptation, allowing for smooth transitions between regions. Additionally, the structured grid in the hybrid mesh is not dynamically optimised for cell quality, resulting in high aspect ratio cells when the mesh is extruded in the normal direction from the surface. This is illustrated in Figure 4.4(d). Gmsh also performs poorly at high Reynolds numbers of 5×10^6, where the initial cell height at the surface must be very small to achieve $y^+ < 1$. This limitation is not observed when meshing with Pointwise.

The lift, drag, and moment polars, shown in Figures 4.6, 4.7, and 4.8, are examined using the CFD setup described in Chapter 5. The mesh generated in Pointwise demonstrates a smoother trend when a linear spline is interpolated through the CFD data points, in contrast to the results obtained with the Gmsh mesh. Within the operational flight regime, the aerodynamic coefficients are generally in close agreement, with minor discrepancies observed in C_d; specifically, the Gmsh mesh predicts slightly higher drag than the Pointwise mesh.

The number of iterations required for the initial residuals to fall below 8×10^7 is shown in Figure 4.9. The Pointwise mesh consistently requires fewer iterations, with a general trend of decreasing iteration count corresponding to the reduction of the recirculation bubble on the lower surface as the AoA increases up to the stall point. In comparison, the Gmsh mesh generally requires more iterations, and no clear dependency on AoA is noticeable. Additionally, for the Gmsh mesh, the simulation diverged for $\alpha = -2°$, resulting in no data being obtained for this angle. These results further support the reliability of the fully structured Pointwise mesh, which also offers advantages for automation, requiring fewer iterations across a wide range of simulations.

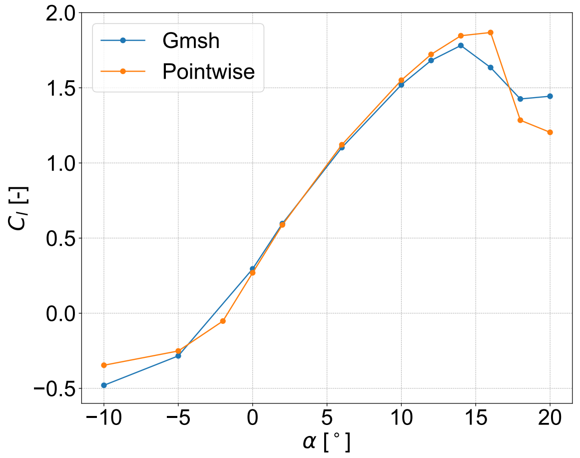

Figure 4.6: Lift polar comparison between hybrid mesh (Gmsh) and structured mesh (Pointwise).

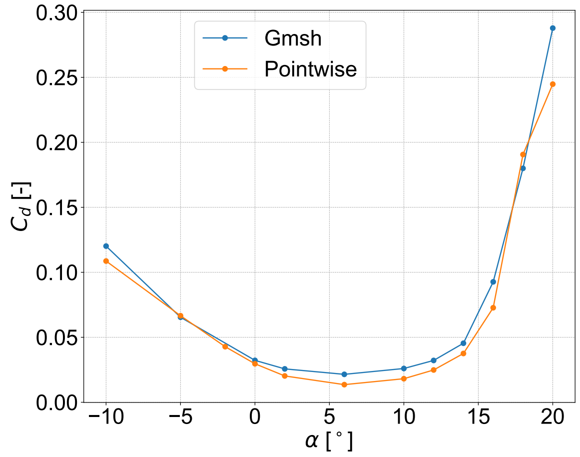

Figure 4.7: Drag polar comparison between hybrid mesh (Gmsh) and structured mesh (Pointwise).

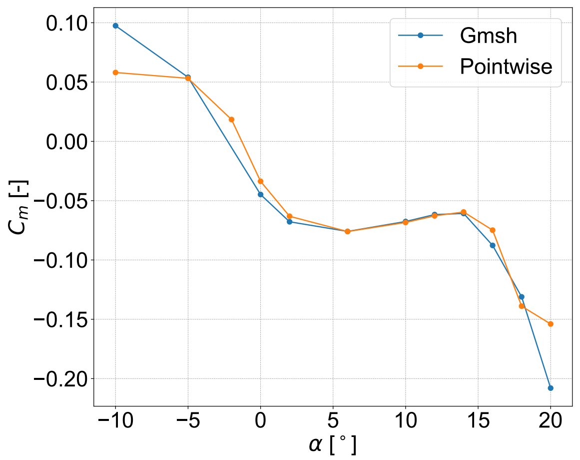

Figure 4.8: Moment polar comparison between hybrid mesh (Gmsh) and structured mesh (Pointwise).

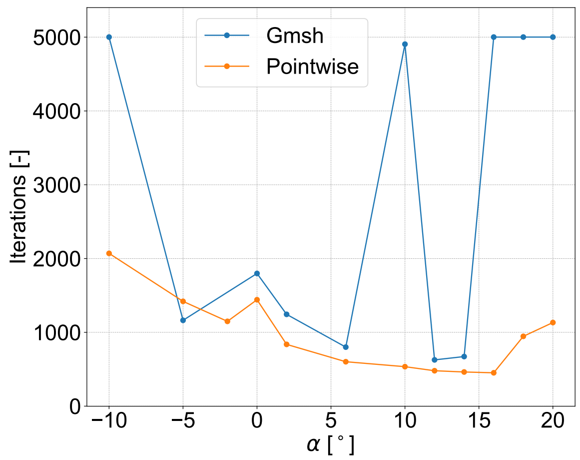

Figure 4.9: Iteration comparison between hybrid mesh (Gmsh) and structured mesh (Pointwise).

Finally, the residual plots for $\alpha = 10°$ are presented in Figure 4.10 for the Gmsh mesh and Figure 4.11 for the Pointwise mesh. Once again, the structured Pointwise mesh shows greater stability, with already lower residual values at the start of the simulation and a steeper convergence slope.

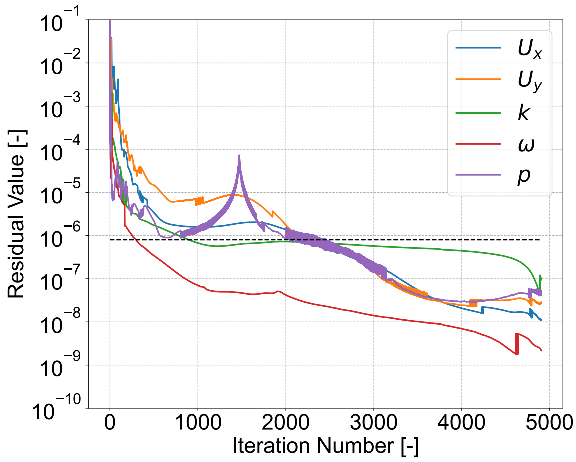

Figure 4.10: Gmsh mesh simulation residuals at $\alpha = 10°$.

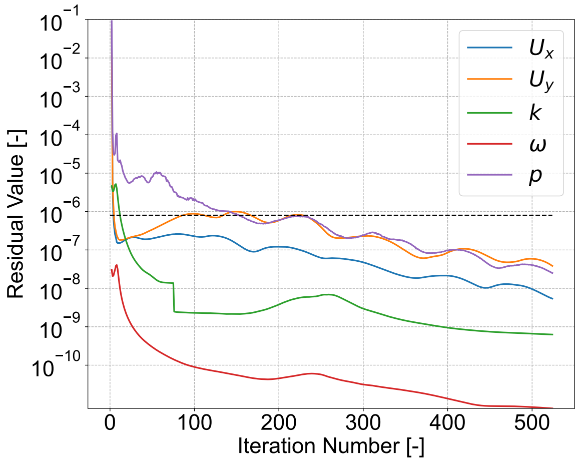

Figure 4.11: Pointwise mesh simulation residuals at $\alpha = 10°$.

4.2.3. Selected Mesh Type and Topology

The structured mesh generated in Pointwise has been shown to outperform the hybrid mesh from Gmsh in terms of accuracy and numerical robustness, as discussed in Section 4.2.2. This aligns with general findings that structured meshes better resolve boundary layers and reduce numerical diffusion [22]. Regarding mesh topology, the large number of simulations and the need for automation across varied geometries and angles of attack motivated the use of an O-grid topology. This choice ensures mesh independence from AoA, as the grid lines naturally follow the flow direction. While a C-grid topology may offer improved wake resolution, it requires more manual tuning to align the mesh with the wake and is therefore less suitable for automated workflows, making the O-grid a more practical and reliable choice for this study.

4.2.4. Profile Surface Node Distribution

Conventional airfoils, such as NACA profiles, typically use cosine spacing to interpolate nodes along the surface [29]. This method distributes grid points more densely near the LE and TE, where curvature and pressure gradients are highest, thereby improving accuracy

$$x_i = \frac{1 - \cos\left(\frac{i\pi}{N}\right)}{2} c. \tag{4.2}$$

The unconventional and highly asymmetric shape of a LEI kite profile presents challenges for mesh discretisation along the chord line. The lower surface spans a larger area with slower-moving flow in the recurvature region, while the upper surface requires sufficient resolution up to the point of maximum camber to capture the accelerated flow accurately. To address these issues, Watchorn [50] implemented a custom node distribution: uniform spacing near the LE, a hyperbolic tangent distribution along the canopy, and another uniform spacing around the semi-circular TE. The tanh spacing is designed to start and end with the same spacing as the adjacent uniform regions near the LE and TE, thereby avoiding abrupt transitions in grid density.

The transition between the uniform and hyperbolic tangent distributions occurs near the end of the LE fillet on both surfaces. An example of this node distribution, using a total of 575 points, is shown in Figure 4.12. The number of nodes near the LE is defined as 28% of the total amount, $n_{\text{LE}} = 0.28 n_{\text{total}}$, providing good results across a range of LE tube diameters. For the TE region, a fixed number of $n_{\text{TE}} = 15$ nodes is used to improve mesh density and prevent failures in the hyperbolic extrusion process, especially in cases where small initial layer thicknesses would otherwise violate the Jacobian criteria. The remaining nodes are distributed along the canopy region, resulting in

$$n_{\text{canopy}} = \left\lfloor \frac{1}{2} \left(n_{\text{total}} - n_{\text{TE}} - n_{\text{LE}} \right) \right\rfloor \tag{4.3}$$

nodes on each side, ensuring a balanced distribution. Here, the floor operator $\lfloor \cdot \rfloor$ denotes rounding down to the nearest integer.

(a) Tanh node distribution along the upper and lower canopy surface.

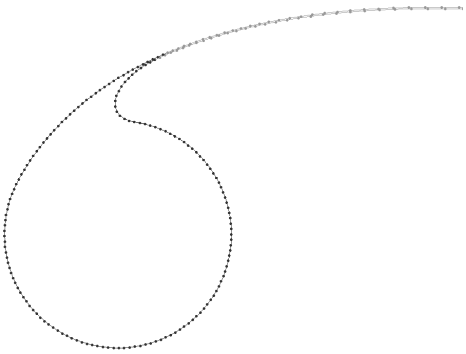

(b) Uniform node distribution along the LE section.

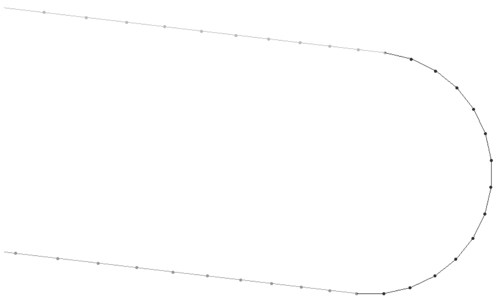

(c) Uniform node distribution along the TE section.

Figure 4.12: Grid point distribution along the surface of an LEI kite profile with 575 nodes. Shape parameters: $t = 0.08$, $\eta = 0.15$, $\kappa = 0.08$, $\delta = 0°$, $\lambda = 0.2$, and $\phi = 0.65$.

4.2.5. Mesh Layer Spacing

Pointwise employs a hyperbolic extrusion algorithm to generate mesh layers in the wall-normal direction, using a progression-based method. This algorithm requires three key parameters to control the extrusion process. The cell growth rate is set to 1.1, consistent with both Pointwise's default and the setting used in Watchorn's study [50]. The number of layers is discussed in Section 4.3, while the third parameter is the initial step size y, which directly affects the non-dimensional wall distance y^+.

The y^+ value is a dimensionless quantity used to assess whether the near-wall mesh resolution is sufficient to capture boundary layer effects in turbulence modelling. It accounts for the no-slip condition at the wall, where the velocity is zero. The friction velocity is defined as $U_\tau = \sqrt{\tau_\mathrm{w}/\rho_\infty}$, with τ_w representing wall shear stress. The resulting expression for y^+ is

$$y^+ = \frac{\rho_\infty y U_\tau}{\mu_\infty} = \frac{\rho_\infty \sqrt{\tau_\mathrm{w}/\rho_\infty}}{\mu_\infty} \,. \tag{4.4}$$

The general rule for achieving a well-resolved boundary layer in turbulence modelling is to ensure that the surface grid cells satisfy the condition $y^+ < 1$ [22]. The wall shear stress τ_w, which is needed for calculating y^+ in Equation 4.4, is defined as

$$\tau_\mathrm{w} = \frac{1}{2} C_\mathrm{f} \rho_\infty U_\infty^2 \,, \tag{4.5}$$

where C_f is the skin friction coefficient, ρ_∞ is the freestream density, and U_∞ is the freestream velocity.

The skin friction coefficient is estimated using an empirical relation from turbulent flat-plate boundary layer theory, commonly known as Prandtl's one-seventh-power law [2]:

$$C_\mathsf{f} = \frac{0.027}{{Re_\mathsf{x}}^{1/7}} \,. \tag{4.6}$$

The Reynolds number based on the chord length is defined in terms of either the dynamic or kinematic viscosity as

$$\mathsf{Re} = \frac{\rho_\infty U_\infty c}{\mu_\infty} = \frac{U_\infty c}{\nu_\infty} \,, \tag{4.7}$$

where μ_∞ and ν_∞ represent the freestream dynamic and kinematic viscosity, respectively.

Assuming $U_\infty = 1\,\mathrm{m\,s^{-1}}$ and $\rho_\infty = 1\,\mathrm{kg\,m^{-3}}$, the Reynolds number determines μ_∞. Rather than relying on fixed freestream conditions, this approach focuses on non-dimensional aerodynamic coefficients, which are more relevant for the analysis. By assembling the previously defined relations and enforcing a y^+ value below 1, the required initial step size can be computed. This ensures that the y^+ value at each surface node remains below 1 during the simulation. As such, this relation is of critical importance for enabling automated CFD workflows across varying Reynolds numbers.

4.2.6. Smoothing Parameters

An LEI kite profile differs significantly from a conventional airfoil, primarily due to its large LE tube and thin canopy, which form a sharp corner on the lower surface near the junction.

To prevent errors in Pointwise's hyperbolic extrusion method, smoothing parameters are carefully adjusted to suit the profile. Proper tuning of these parameters not only helps to mitigate numerical instabilities but also ensures a smooth and well-structured mesh. The software provides four types of smoothing techniques for this purpose: explicit smoothing, implicit smoothing, Kinsey-Barth smoothing and volume smoothing.

Both the explicit and implicit smoothing parameters operate in the transverse direction, with the implicit coefficient always set to twice the value of the explicit coefficient. In contrast, the Kinsey-Barth smoothing parameter functions in the wall-normal direction and is typically applied when dealing with severe concavity.

Compared to the smoothing parameters used by Watchorn [50], only the explicit parameter is increased to 8, while the implicit parameter is adjusted accordingly. This modification results in slightly improved smoothness in localised high-curvature regions. As shown in Table 4.1, these settings are consistently applied across all configurations to ensure uniform mesh quality and stability.

Table 4.1: Hyperbolic extrusion smoothing parameters.

Smoothing Parameter	Value	Notes
Explicit	8.0	
Implicit	16.0	Double the explicit coefficient
Kinsey Barth	5.0	High concavity $\rightarrow \geq 3.0$
Volume	0.5	Default value

4.3. Mesh Sensitivity Analysis

This section presents a sensitivity analysis of the fully structured mesh designed for the LEI kite profile at Re $= 5 \times 10^5$. It evaluates mesh resolution convergence (normal and tangential), LE fillet effects, and canopy thickness. Their impact on aerodynamic performance and residual convergence is assessed for both a conventional and a highly cambered airfoil, as shown in Figures 4.13 and 4.14.

To support mesh automation across different profile geometries, these two contrasting profiles were selected to test mesh performance, ranging from conventional shapes to extreme cases like the highly cambered airfoil. The aim is to identify optimal mesh settings and assess the need for dynamic adaptation based on profile parameters.

This approach aims to minimise numerical errors and ensure an accurate representation of real-world aerodynamic behaviour. For all cases, percentage differences in force and moment coefficients are analysed, along with residual convergence plots for U_x, U_y, k, p, and ω.

Standard convergence guidelines suggest residuals should reach approximately 10^{-6} for momentum, pressure, and turbulence equations, and 10^{-4} for force coefficients [22]. To meet these criteria consistently across configurations, an initial residual threshold of 8×10^{-7} is used.

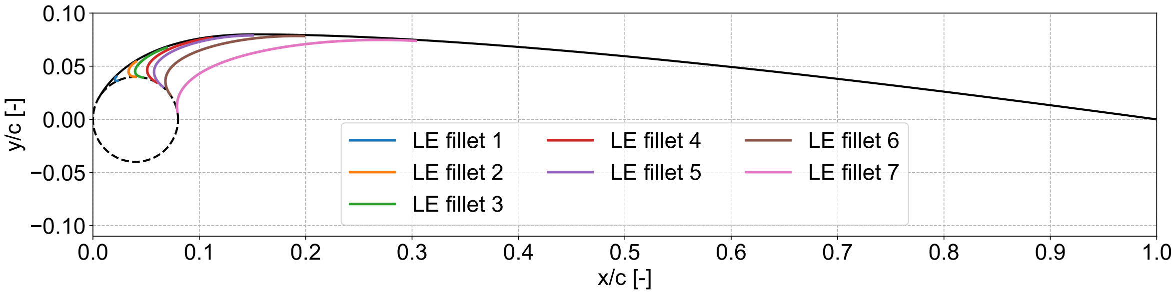

Figure 4.13: Conventional LEI kite profile with varying LE fillet configuration and fixed shape parameters: $t = 0.08$, $\eta = 0.15$, $\kappa = 0.08$, $\delta = 0°$, $\lambda = 0.2$, and $\phi = 0.65$.

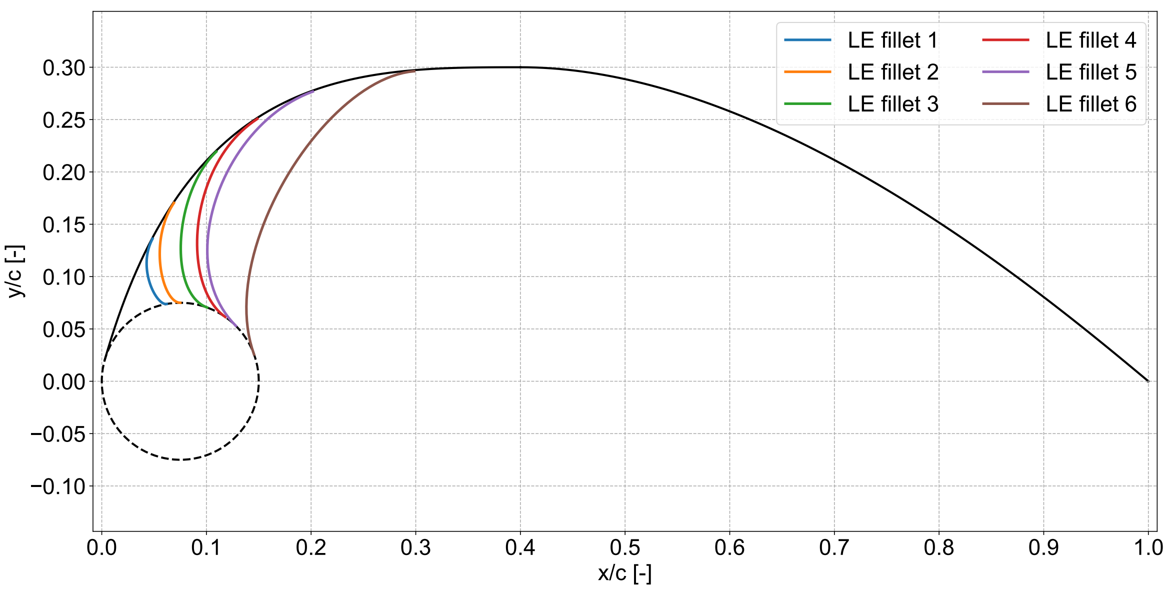

Figure 4.14: Highly cambered LEI kite profile with varying LE fillet configuration and fixed shape parameters: $t = 0.15$, $\eta = 0.4$, $\kappa = 0.3$, $\delta = 0°$, $\lambda = 0.4$, and $\phi = 0.65$.

4.3.1. Mesh Domain Convergence

The aerodynamic performance and residual convergence of the mesh domain are analysed by varying the tangential and normal mesh resolutions for two specific configurations: the conventional profile with fillet 3 and the highly cambered profile with fillet 2, as shown in Figures 4.24 and 4.14. The influence of tangential resolution is shown in Figures 4.15 and 4.16 for both profiles. Except at $\alpha = 0°$ for the highly cambered case, a tangential resolution of 475 nodes yields good convergence across both configurations. However, to ensure consistency and remain aligned with Watchorn's study, a resolution of 575 nodes is selected. As the optimal resolution is less clearly defined for the highly cambered profile, the 575-node distribution is applied uniformly across all cases.

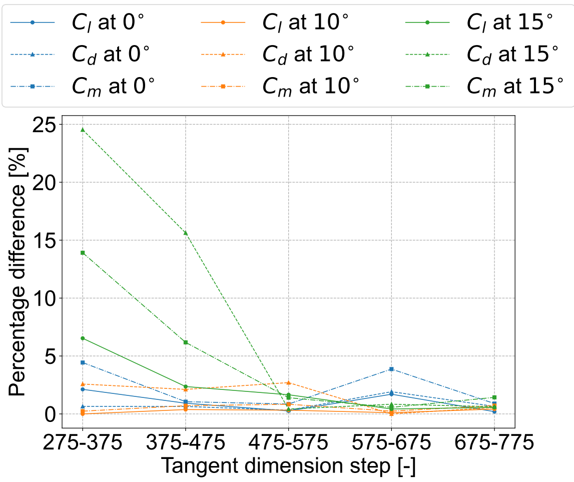

Figure 4.15: Absolute percentage differences in (C_l, C_d, and C_m) for various tangential node resolutions for a conventional LEI kite profile. For a fixed normal distribution of 201 layers and $t_\text{canopy} = 0.0001$.

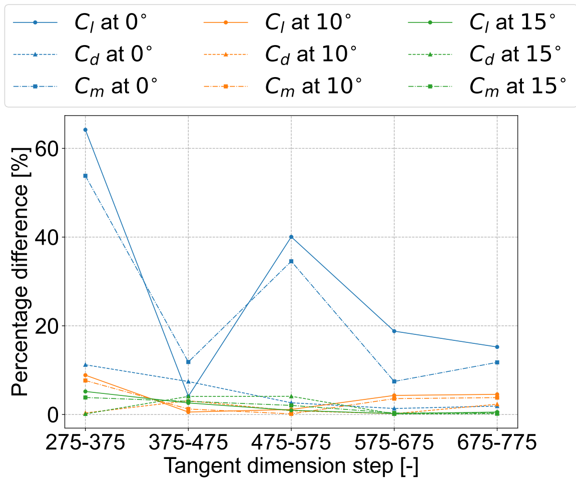

Figure 4.16: Absolute percentage differences in (C_l, C_d, and C_m) for various tangential node resolutions for a highly cambered LEI kite profile. For a fixed normal distribution of 201 layers and $t_\text{canopy} = 0.0001$.

In contrast, establishing convergence for aerodynamic performance with respect to wall-normal cell layers is more challenging, as shown in Figures 4.17 and 4.18. Nevertheless, a mesh with 201 layers appears sufficient to capture the relevant flow features, showing minimal differences compared to adjacent configurations. This choice also aligns with Watchorn's sensitivity analysis, which used the same number of layers. The pronounced peak at steps 191-201 for the highly cambered profile in Figure 4.18 occurs at $\alpha = 0°$, likely due to the large recirculation region extending far toward the TE, which can cause the results to deviate using steady RANS simulations.

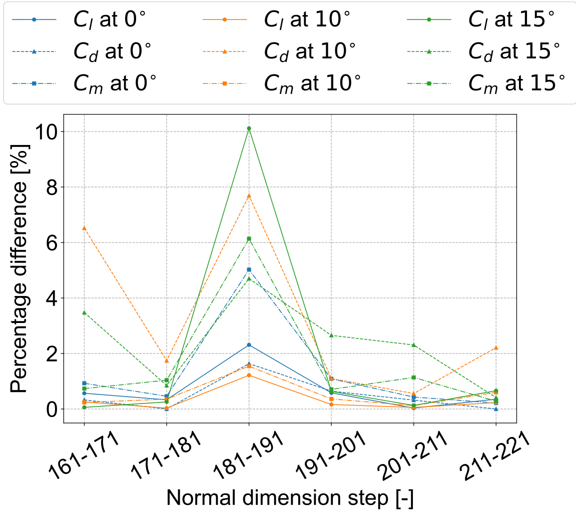

Figure 4.17: Absolute percentage differences in (C_l, C_d, and C_m) for various wall-normal cell layers for a conventional LEI kite profile. For fixed tangential resolution of 575 nodes and $t_\text{canopy} = 0.0001$.

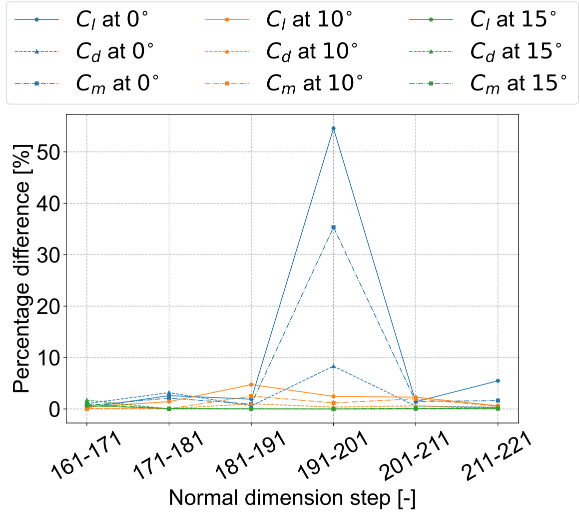

Figure 4.18: Absolute percentage differences in (C_l, C_d, and C_m) for various wall-normal cell layers for a highly cambered LEI kite profile. For fixed tangential resolution of 575 nodes and $t_\text{canopy} = 0.0001$.

Regarding the number of simulation cycles for the tangential and wall-normal sensitivity analyses in Figures 4.19 and 4.20, no clear correlation is observed between tangential resolution and convergence rate. However, for the wall-normal resolution, a trend of earlier convergence emerges with increasing layer count. While this improves CFD automation, it also raises computational cost due to a higher cell count. To balance accuracy and efficiency, the wall-normal layer count is set to 201.

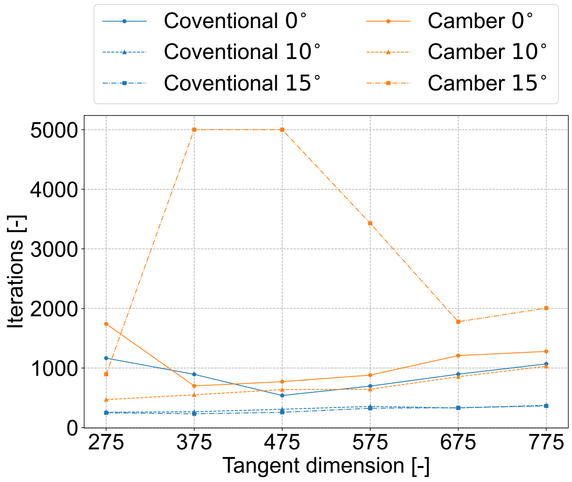

Figure 4.19: Iteration count to reach 8×10^{-7} residual for varying tangential node distributions on conventional and cambered LEI kite profiles. For a fixed normal dimension of 201 nodes and $t_\text{canopy} = 0.0001$.

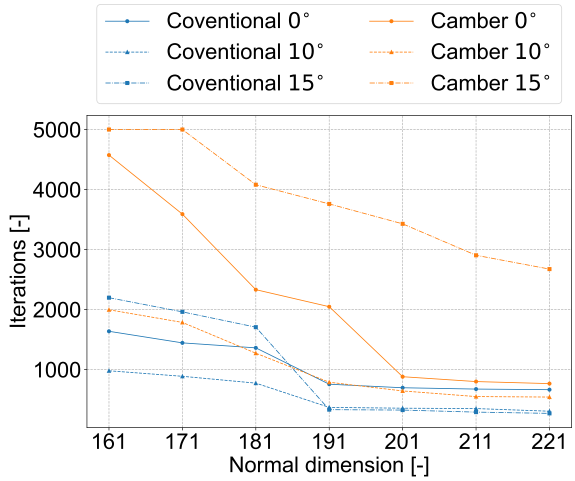

Figure 4.20: Iteration count to reach 8×10^{-7} residual for varying wall-normal cell layers on conventional and cambered LEI kite profiles. For fixed tangential resolution of 575 nodes and $t_\text{canopy} = 0.0001$.

4.3.2. Leading Edge Fillet Analysis

In reality, no fillet is present on the lower surface near the LE. However, in CFD simulations, introducing a fillet is necessary to prevent numerical errors caused by grid compression at the sharp intersection of the front canopy spline and the LE tube.

This analysis evaluates the impact of varying fillet sizes on aerodynamic performance and flow residual convergence, while also examining potential correlations between them. The fillet configurations for both conventional and highly cambered profiles are shown in Figure 4.13 and 4.14.

To quantify the fillet geometry, Delaunay triangulation is used to generate non-overlapping triangular mesh elements from a set of contour points, ensuring that no point lies within the circumcircle of any triangle [32]. The total LE fillet area is obtained by summing the areas of these triangles, as illustrated in Figure 4.21. The calculated areas for the different fillet configurations are listed in Table 4.2.

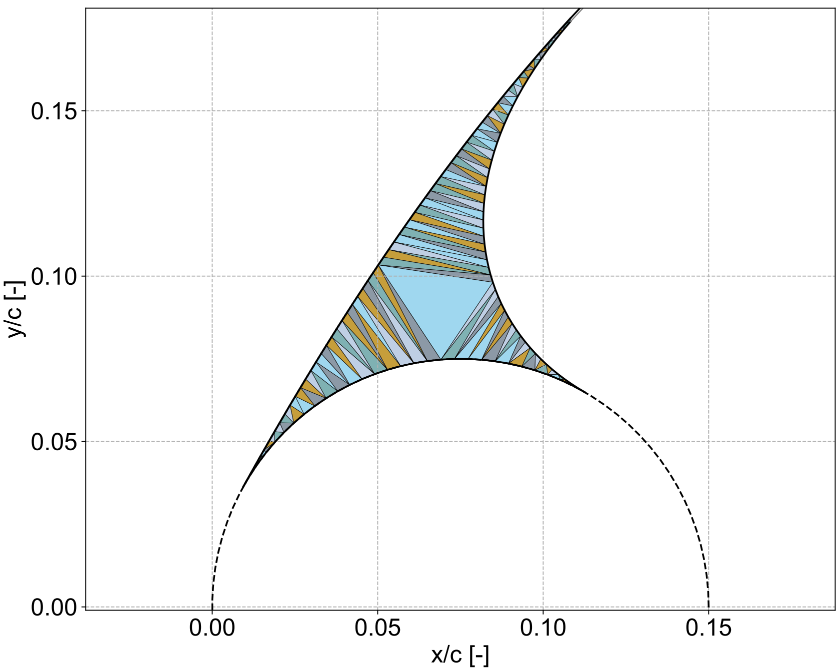

Figure 4.21: Delaunay triangulation method applied for area calculation on the LE fillet.

Table 4.2: LE fillet enclosed areas for the conventional and highly cambered profile.

Conventional Profile	Area $[\mathrm{cm}^2]$	Highly Cambered Profile	Area $[\mathrm{cm}^2]$
LE fillet 1	0.13	LE fillet 1	4.22
LE fillet 2	0.98	LE fillet 2	8.58
LE fillet 3	2.26	LE fillet 3	26.12
LE fillet 4	5.04	LE fillet 4	36.60
LE fillet 5	9.45	LE fillet 5	47.90
LE fillet 6	19.31	LE fillet 6	137.12
LE fillet 7	45.52		

To gain a better understanding of the mesh structure at the LE fillet, a close-up view of the conventional profile with LE fillet number 1 and 2 is presented in Figures 4.22 and 4.23, with a tangential and normal mesh resolution of 575×201. Both meshes exhibit decent orthogonality despite the presence of a sharp LE corner, showcasing the capabilities of the hyperbolic mesh generator in Pointwise.

In this analysis, it is assumed that the smallest LE fillet would provide the best representation of real-world conditions while also considering the numerical error introduced with squeezed mesh quality. Therefore, the second smallest mesh for the conventional profile Figure 4.23 is used as the baseline for aerodynamic comparisons in terms of C_l, C_d, and C_m.

For the highly cambered profile, the smallest fillet demonstrated a smooth local mesh, ensuring reliable flow representation. The aerodynamic performance comparisons for both the conventional and highly cambered profiles are shown in Figures 4.24 and 4.25, where the LE fillet number is represented on the x-axis as the reference for comparison with LE fillet 2 for the conventional profile and LE fillet 1 for the highly cambered one.

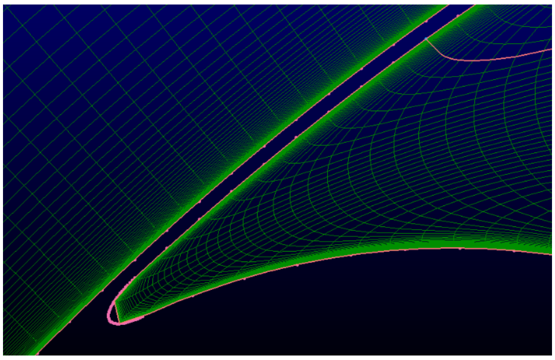

Figure 4.22: Conventional profile with LE fillet configuration 1, with mesh resolution 575×201.

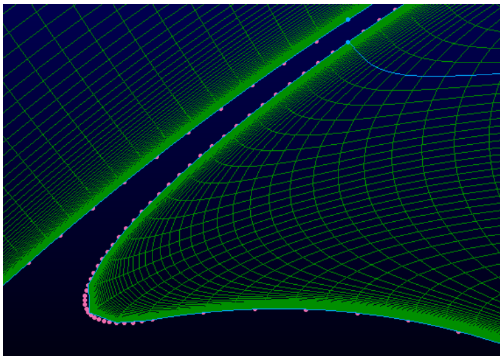

Figure 4.23: Conventional profile with LE fillet configuration 2, with mesh resolution 575×201.

A clear diverging trend in aerodynamic performance is noticeable at $\alpha = 0°$. However, no clear correlation is observed for other AoAs. For the highly cambered profile (Figure 4.25), a difference in aerodynamic performance is observed only for the largest fillet, affecting all coefficients across the AoA range.

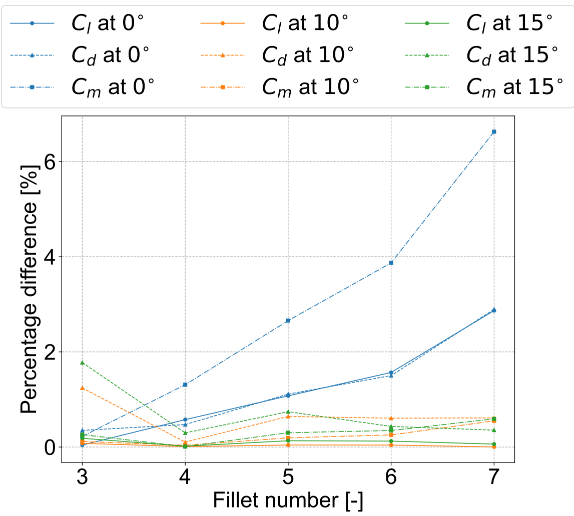

Figure 4.24: Relative percentage differences in (C_l, C_d, and C_m) for the conventional LEI kite profile across various LE fillets compared to LE fillet number 2.

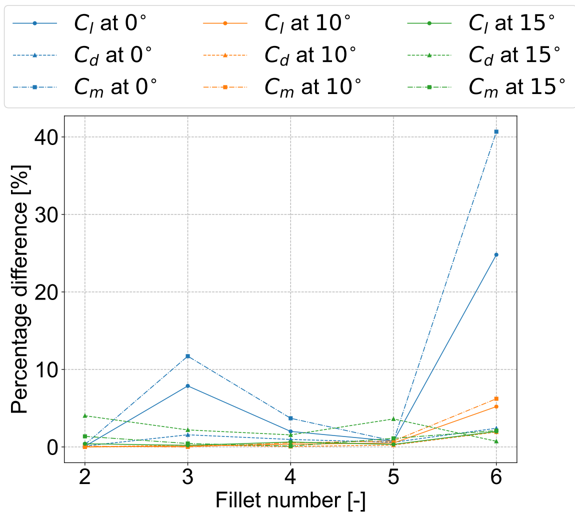

Figure 4.25: Relative percentage differences in (C_l, C_d, and C_m) for the highly cambered LEI kite profile across various LE fillets compared to LE fillet number 1.

Concerning the convergence iteration count for the highly cambered profile in Figure 4.27, convergence at $\alpha = 15°$ is achieved only for fillet configurations 3 and 6. While no clear trend is observed across all fillets, fillet 6 appears to offer greater stability. However, it also more closely resembles a double-skin LEI kite, which may not reflect the target design profile.

For the conventional LEI kite profile, the required number of iterations is shown in Figure 4.26. A general trend of faster convergence is visible with increasing fillet size, except at $\alpha = 15°$. The convergence trend obtained for the $\alpha = 0°$ and $10°$ is likely due to the larger recirculation region present at these angles, in which the LE fillet can help streamline the flow.

Ultimately, considering the already smooth surface mesh obtained with LE fillet 2 in Figure 4.23 and the minimal impact of LE fillet size on aerodynamic performance, fillet configuration 3, offering slightly improved smoothness, is selected for the conventional LEI kite profile. While even larger fillets showed improved convergence, they deviated more from the realistic geometry. Therefore, fillet 3 is chosen as a balanced solution and will be applied consistently across all other profile configurations.

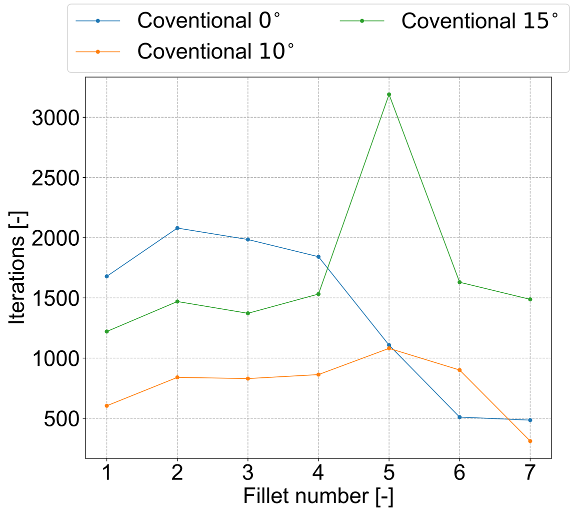

Figure 4.26: Iteration count to reach an initial residual threshold of 8×10^{-7} for the conventional LEI kite profile with varying LE fillets.

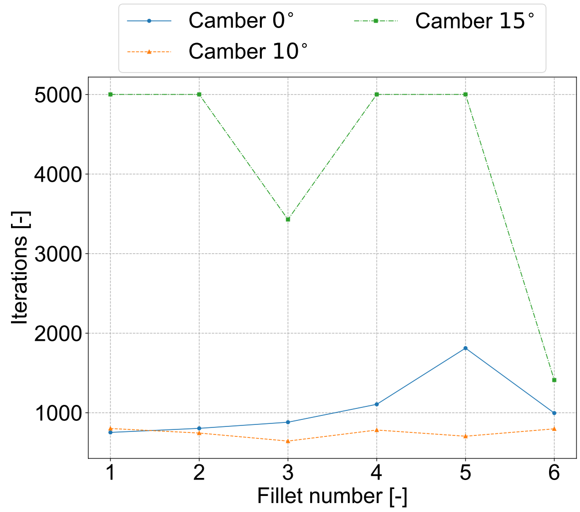

Figure 4.27: Iteration count to reach an initial residual threshold of 8×10^{-7} for the highly cambered LEI kite profile with varying LE fillets.

4.3.3. Canopy Thickness Analysis

The canopy sensitivity analysis focuses only on the conventional profile, as the large camber could introduce additional interference. While the canopy is small relative to the chord, its thickness is not negligible. This analysis compares the smallest canopy thickness against larger values, with the relative aerodynamic performance differences shown in Figure 4.28. A clear divergence in performance appears as thickness increases, highlighting the importance of selecting a representative canopy thickness. All meshes are assumed to be of smooth quality due to the use of hyperbolic mesh generation. Although experimental validation is ideal.

In real-world kites, canopy fabric can be as thin as $0.1\ \mathrm{mm}$, but reinforcement near the TE often increases this to around $1\ \mathrm{mm}$ to prevent damage from flutter. When scaled to typical centre chord lengths, $1.8\ \mathrm{m}$ for kitesurfing kites and $2.7\ \mathrm{m}$ for KitePower systems [12], this corresponds to non-dimensional canopy thicknesses of $t_{\text{canopy}} = 0.00055$ and 0.00037, respectively. In comparison, Watchorn used a larger value of $t_{\text{canopy}} = 0.001$ [50]. The thinner, more realistic values derived here are therefore included in this sensitivity analysis.

The iteration count demonstrated in Figure 4.29 remains roughly constant at $\alpha = 10°$ and $15°$. At $\alpha = 0°$, however, a trend emerges: thicker canopies tend to generate more turbulence at the TE, requiring more iterations for convergence.

This leads to the selection of a more representative canopy thickness of $t_{\text{canopy}} = 0.0005$, balancing slightly improved convergence behaviour with reduced deviation from realistic aerodynamic performance.

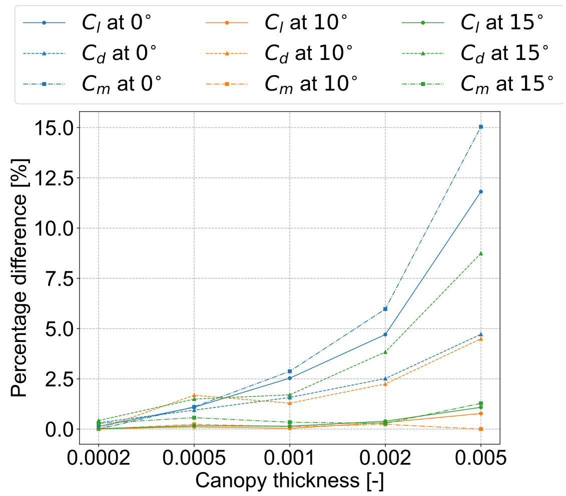

Figure 4.28: Relative percentage differences in (C_l, C_d, and C_m) for the conventional LEI kite profile across various canopy thicknesses compared to ($t_\text{canopy} = 0.0001$). Mesh resolution 575×201.

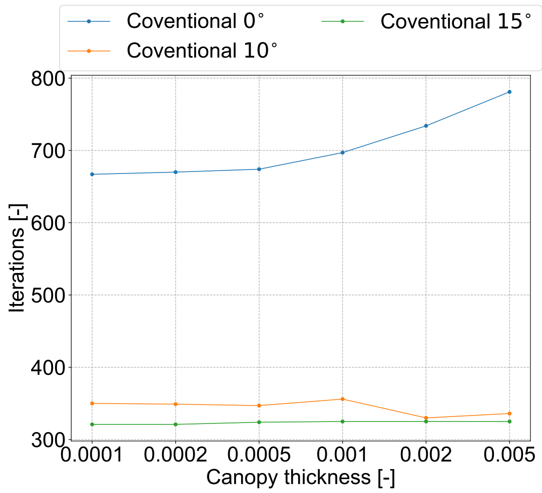

Figure 4.29: Iteration count to reach an initial residual threshold of 8×10^{-7} for the conventional LEI kite profile with varying canopy thickness.

5

Computational Fluid Dynamics & Simulation Set-Up

This chapter presents the computational methodology used to simulate the aerodynamic behaviour of LEI kite profiles. It includes the justification for using the RANS equations and the selected turbulence modelling approach. The omission of transition modelling is supported by a numerical analysis of the LE tube-canopy connection and the lower surface closing seam. These elements, along with the open-source CFD solver OpenFOAM used to solve the RANS equations, are discussed in Section 5.1. The simulation setup in Section 5.2 includes the use of the high-performance computing (HPC) cluster and outlines the applied boundary conditions. The chapter concludes with a discussion of the convergence criteria and the postprocessing methods used to extract aerodynamic performance coefficients.

5.1. Computational Fluid Dynamics

Due to the unconventional shape of LEI kite profiles, flow phenomena such as lower surface recirculation and upper surface separation at higher angles of attack arise, both of which are inherently viscous and rotational. As these effects cannot be captured by conventional potential flow tools like Xfoil [19], CFD simulations are necessary for accurate analysis.

5.1.1. RANS k-ω SST Simulation

Based on the discussion presented in Section 2.4, Steady RANS simulations provide an effective balance between computational cost and accuracy compared to LES and DNS, making it particularly suitable for automating large datasets of profiles intended for regression modelling. The assumption made by [23] further supports employing steady-state RANS simulations, as the flow experiences a low Mach number and can thus be considered incompressible.

The steady RANS equations are derived by applying Reynolds averaging to the Navier–Stokes equations, which removes turbulent fluctuations. This process introduces additional unknowns, leaving the system unclosed. To achieve closure, empirical approximations in the form of turbulence models are introduced, enabling the equations to be solved [22].

Among the various turbulence models discussed in more detail in the literature Section 2.4 and previously implemented in RANS simulations by Folkersma et al. [23], Demkowicz [18], and Viré et al. [47], the $k-\omega$ SST model by Menter [34] is selected as the most suitable option. This model is widely adopted in engineering applications due to its ability to accurately capture flow behaviour both near walls and in the far field. For a more in depth explanation of the RANS $k-\omega$ SST equations as implemented in the solver, the reader is referred to the work of Watchorn [50].

5.1.2. Transition Modelling

The influence of a transition model versus fully turbulent flow on a LEI kite profile has been investigated by Folkersma et al. [23] using RANS simulations with results compared to experimental data as summarised in Section 2.4.3. The study demonstrated that incorporating a transition model produced results more consistent with experiments, particularly for smooth profiles tested at low turbulence intensity ($I = 0.02$).

In contrast, Watchorn [50] noted the presence of a seam at the LE, connecting the canopy to the LE tube (see Figure 2.8), which acts as a disturbance triggering bypass transition. This observation has led to the assumption that the flow transitions to a turbulence boundary layer immediately at the seam, with the seam located within the first 2% of the chord length. Justifying the omission of a transition model. However, no direct evidence currently confirms that the seam initiates the bypass transition outlined in Section 2.4.3.

When deciding whether to implement a transition model or assume fully turbulent flow, key parameters such as the free-stream Reynolds number (Re), turbulence intensity, and seam geometry must be considered. Recent field measurements from the Kitepower V9 system, conducted in 2023 and 2024, report turbulence intensities ranging from 5% up to 30%, with a median around 8% [13]. These values indicate a significantly more turbulent environment compared to the 2% turbulence level assumed in the study by Folkersma et al. [23].

Regarding surface imperfection-induced disturbances, bypass transition can be triggered by the LEI seam. The seam on the upper side of the LE consists of two main components: a forward-facing step and a zig-zag stitching pattern, both resulting from the overlapping canopy fabric seen in Figure 2.8. A study involving LES simulations of a forward-facing step with a step-height Reynolds number ($\mathrm{Re_h}$) of 720 revealed the formation of turbulent structures, indicating its potential to trigger transition [53].

The step-height Reynolds number, or more generally, the roughness-height-based Reynolds number, is defined by the following relation [8],

$$\mathrm{Re_h} = \frac{u_\mathrm{h} h}{\nu}. \tag{5.1}$$

where h is the obstacle height, and u_h is the local flow velocity at that height. u_h can be estimated using the pressure coefficient

$$u_h = f_\mathrm{bl}\left(U_\infty\sqrt{1 - C_\mathrm{p}}\right). \tag{5.2}$$

Here, f_bl is a boundary layer correction factor accounting for viscous effects, as the obstacle is located within the boundary layer. The term in brackets is derived from Bernoulli's equation.

Additionally, the zig-zag stitching pattern observed on the LEI seam, as shown in Figure 2.8, is commonly used in aerospace applications to induce a fixed transition and is therefore important to account for. While the same methodology and equations apply, the zig-zag pattern acts as a three-dimensional roughness element and is generally more effective at triggering transition than a simple forward-facing step. Consequently, it typically requires a smaller height to initiate transition over the profile [8].

Determining the critical roughness-height-based Reynolds number ($\mathrm{Re_{h,crit}}$) required to trigger transition is complex, as various studies report different threshold values. For the current configuration featuring a forward-facing step combined with a zig-zag stitching pattern, no experimental data exist that also account for the influence of elevated turbulence intensity. However, literature provides some guidance. Braslow et al. [8] reported typical critical values ranging between 300 and 600. For three-dimensional roughness elements such as zig-zag or wavy patterns, Balakumar [3] adopted a value of 300, while other studies by Rooij et al. [39] and Elsinga et al. [20] found that values as low as 200 may be sufficient to trigger transition.

Considering the elevated turbulence intensity and presence of a two-layer roughness, an estimate of $\mathrm{Re_{h,crit}} = 200$ is assumed. The obstacle height, representing the stitching and step at the LE seam, is measured at $0.6\ \mathrm{mm}$. The Reynolds number formulation given in Equation 5.1 is applied across a freestream Reynolds number range from 10^5 to 10^7, with pressure coefficient values ranging from $C_\mathrm{p} = -2$ to $C_\mathrm{p} = -4.5$, corresponding to low and high angles of attack. Chord lengths of $1.8\ \mathrm{m}$ and

$5\ \mathrm{m}$ are used to represent a typical kitesurfing kite and an airborne wind energy kite, respectively, as illustrated in Figure 5.1.

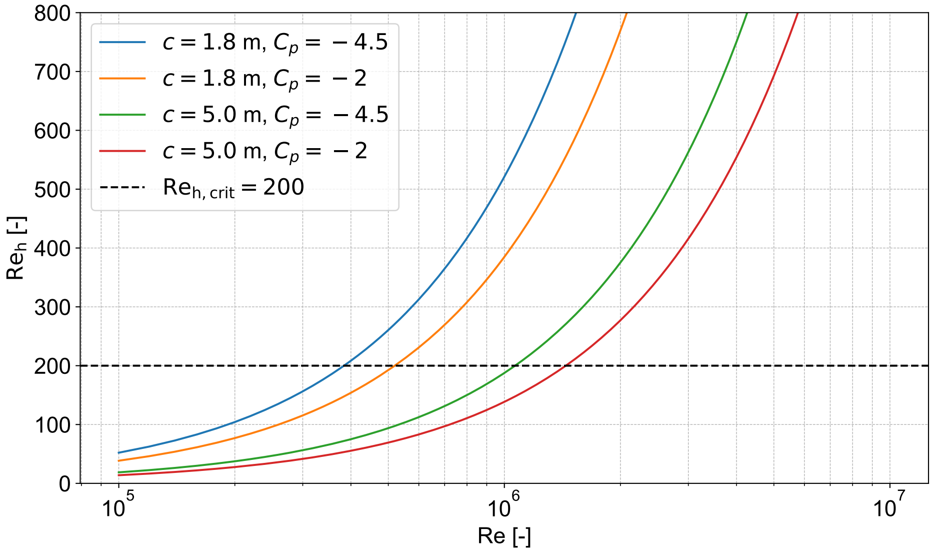

Figure 5.1: Roughness-height-based Reynolds number versus flow Reynolds number for various c and C_p values, with a seam height of $0.6\ \mathrm{mm}$ and $f_\mathrm{bl} = 0.8$.

In the most transition-prone case, with a chord length of $5\ \mathrm{m}$, $C_\mathrm{p} = -2$, and a boundary layer factor $f_\mathrm{bl} = 0.8$, transition would occur beyond $\mathrm{Re} = 1.5 \times 10^6$, representative of low wind speeds during the reel-in phase. However, during power generation in crosswind flight, this leads to significantly higher Reynolds numbers. Since the LE seam is located within the first 2% of the chord, the influence of any upstream laminar region is minimal.

Regarding the closing seam on the underside of the LE tube, shown in Figure 2.9, its height, approximately $6\ \mathrm{mm}$, is around 10 times greater than that of the LE tube-canopy intersection. Based on Equations 5.1 and 5.2, this configuration yields transition as early as $\mathrm{Re} = 1.5 \times 10^5$ under most transition-prone conditions. Since this seam is also located within the first 2% of the chord, both geometric features support the omission of laminar flow modelling. Consequently, assuming a fully turbulent boundary layer from the LE onward is both reasonable and justified.

5.1.3. CFD Solver OpenFoam

OpenFOAM is an open-source CFD software widely utilised in Universities and industry for aerodynamic simulations. It is notable for its flexibility and innovative object-oriented structure built on the C++ programming language [51].

The software utilises the finite volume method (FVM) to discretise and solve the governing partial differential equations of fluid flow, accurately capturing the physics within each control volume. Flow variables are computed at the centroids of these volumes and interpolated through quadrature formulas, inherently maintaining flux conservation across neighbouring control volumes.

Although OpenFOAM can simulate compressible flows, the low free-stream Mach number associated with flow over a LEI kite allows the assumption of incompressibility [23]. This simplification leads to a reduction in the complexity of the RANS equations. By further assuming steady-state conditions, the simplified equations become suitable for simulation using the SimpleFOAM solver library, as demonstrated by Folkersma et al. [23].

SimpleFOAM is based on the semi-implicit method for pressure linked equations (SIMPLE) algorithm, which couples pressure and velocity fields [10]. For a detailed overview of its structure and implementation, readers are directed to Watchorn's master's thesis [50]. The solver is designed exclusively for 3D domains, so for 2D simulations, the profile domain is extruded outward by one chord length.

OpenFOAM integrates smoothly with the meshing tool Pointwise, and simulation results can be easily analysed and visualised using built-in utilities or external programs like ParaView [28].

5.2. Simulation Set-Up

This section outlines the simulation toolchain used to run multiple simulations in parallel on the HPC system at the faculty of Aerospace Engineering. The boundary conditions used in the simulations are discussed in Section 5.2.2. The numerical schemes employed in the OpenFOAM setup were adopted from Watchorn's work, to which the reader is referred for further details. The convergence criteria are outlined in Section 5.2.3, and the aerodynamic data postprocessing is presented in Section 5.2.4.

5.2.1. Simulation Platform

Developing an accurate regression model for multiple input parameters defining the LEI kite profile, using aerodynamic performance coefficients as output, requires substantial computational resources due to the intensive CFD simulations needed for various profile configurations. Running these calculations locally would be impractical, making high-performance computing resources essential. Therefore, the HPC-12 cluster at Delft University of Technology, Faculty of Aerospace Engineering, is employed.

The cluster connects nodes managed by a central master, with access via SSH keys, and jobs are submitted through PBS files that specify the simulation tasks and resource requirements. Operating in a Linux environment with OpenFOAM pre-installed, the cluster efficiently distributes simulations across nodes, automatically assigning new tasks as cores become available.

Mesh generation is performed on a local device and uploaded to the cluster, where simulations are executed and monitored. Once completed, the aerodynamic data is transferred back to the local device for further analysis and regression model development. For verification, local runs using parallel processing are suitable for a small set of AoAs, but for multiple profile configurations, it is advisable to utilise the supercomputer.

5.2.2. Boundary Conditions & Initial Values

This section outlines the boundary types and values employed in the OpenFOAM simulation setup, grounded in the methodology introduced by [23] and verified through subsequent studies by [18] and [50]. The discussion covers the various regions of the simulation domain, including the far-field, profile, and side boundary conditions, which are summarised with their respective types and values in Table 5.1.

Far field

The far-field boundary is modelled as a cylindrical region extending 201 cell layers outward from the profile, with a unit depth of length 1. The inlet velocity scalar U_∞, set to a unit magnitude of 1, is projected in the direction determined by the AoA to define the vector inlet velocity:

$$\mathbf{U}_\mathrm{i} = U_\infty \cdot \begin{bmatrix} \cos\alpha \\ \sin\alpha \\ 0 \end{bmatrix}. \tag{5.3}$$

Since the analysis is two-dimensional, there is no flow variation in the normal z-direction.

The turbulent kinetic energy at the inlet k_i is defined by the free-stream velocity U_∞ and the turbulence intensity I, as

$$k_i = \frac{3}{2}(U_\infty I)^2. \tag{5.4}$$

The specific turbulent dissipation rate at the inlet ω_i is related to the initial turbulent kinetic energy k_i and the free-stream eddy viscosity ratio ν_t/ν_∞, given by

$$\omega_\mathrm{i} = \frac{k_i}{\nu_\infty}\left(\frac{\nu_t}{\nu_\infty}\right)^{-1}. \tag{5.5}$$

These expressions follow the formulation adopted by Watchorn [50], which has been adopted in this work. Given the parameter ν_t/ν_∞, Folkersma et al. [23] reported negligible sensitivity to changes in the eddy viscosity ratio across a range of 0.1 to 10. Similarly, Demkowicz [18] observed minimal influence when varying this ratio between 1 and 50, and for the turbulence intensities ranging from 0.5% to 20%. Based on these findings, both studies selected constant values of $\nu_t/\nu_\infty = 10$ and $I = 0.02$, a choice also adopted by Watchorn and consistently applied throughout this work.

Under the assumption of incompressible flow for LEI kites [23], the fluid density ρ is treated as constant. The pressure field is expressed in terms of the kinematic pressure, $p_k = p_s/\rho$, where p_s is the static pressure. This formulation, adopted in `simpleFoam`, preserves the correct physical pressure gradient while simplifying its numerical treatment.

For the boundary conditions, the `inletOutlet` type is applied to the velocity field **U**, k, and ω. This condition uses a `fixedValue` at inflow and a `zeroGradient` at outflow [10]. For the pressure field, a `fixedValue` condition is imposed at the outlet ($p_k = 0$), while the inflow uses a `zeroGradient` condition, implemented via the `outletInlet` boundary type.

Profile
The LEI kite profile is defined as a stationary, impermeable solid boundary around which the flow moves. Due to fluid viscosity and impermeability, the no-slip boundary condition is applied, meaning both tangential and normal velocity components at the wall are zero. Under the assumption of incompressible flow, the pressure gradient normal to the wall is considered negligible and is set to zero.

Turbulence fluctuations diminish near impermeable surfaces, making the flow essentially laminar in close proximity to the wall. The boundary layer adjacent to the wall is known as the viscous sub-layer, where k and eddy viscosity (ν_t) are effectively zero. However, the specific turbulent dissipation rate ω is not zero at the wall due to the damping effects of the solid boundary on turbulence fluctuations. [22]

Side
The SimpleFOAM solver is limited to simulating three-dimensional domains. To accommodate two-dimensional airfoil simulations, the profile is extruded by one chord length in the z-direction, forming a thin 3D domain [10]. The two side boundaries created through this extrusion are defined as empty, specifying a two-dimensional simulation setup in OpenFOAM. This boundary condition ensures uniform flow conditions along the domain thickness, preventing flow variation in the normal z-direction.

Table 5.1 summarises the boundary conditions applied at the far-field, profile, and side surfaces, including their respective types and values.

Table 5.1: Boundary conditions for the flow variables [50].

Variable	Far-field Type	Far-field Value	Profile Type	Profile Value	Side Type
U, m/s	inletOutlet	inletValue = $\mathbf{U}_i$ Initial value = $\mathbf{U}_i$	fixedValue	(0, 0, 0)	empty
$p_k\,(p_s/\rho)$, m^2/s^2	outletInlet	outletValue = 0 Initial value = 0	zeroGradient	–	empty
k, m^2/s^2	inletOutlet	inletValue = k_i Initial value = k_i	fixedValue	Value = 0	empty
$\omega, 1/$s	inletOutlet	inletValue = ω_i Initial value = ω_i	omegaWallFunction	Initial value = ω_i	empty
ν_t, m^2/s	calculated	Initial value = 0	nutkWallFunction	Initial value = 0	empty

5.2.3. Convergence Monitoring

A general rule of thumb for convergence criteria is that the final residuals for momentum, pressure, and turbulence equations represented by U_x, U_y, p, k, and ω should fall below 10^{-6}. Meanwhile, the force residuals, corresponding to C_l, C_d, and C_m, should typically reach the order of 10^{-4} [22].

In OpenFOAM, residuals are monitored using the initial residual, which is slightly higher than the final value. In cases of rapid convergence, force coefficients may not reach the 10^{-4} threshold. To ensure consistent behaviour, an initial residual threshold of 8×10^{-7} was selected. Each simulation was limited to 5,000 iterations. Exceeding this indicates non-convergence, typically at negative or post-stall AoA, where steady RANS accuracy declines due to large flow separation.

The final flow and force residuals for a conventional LEI kite profile at $\alpha = 6°$ are shown in Figures 5.2 and 5.3, respectively. Compared to the setup used by Watchorn, the relaxation factor for the velocity component U was reduced from 0.9 to 0.8, which helped smooth out residual noise and led to a more stable convergence. Notably, the final flow residuals fall below the order of 10^{-7}, while the initial flow residual reached 8×10^{-7}. The force residuals also dropped just below the convergence threshold of 10^{-4}.

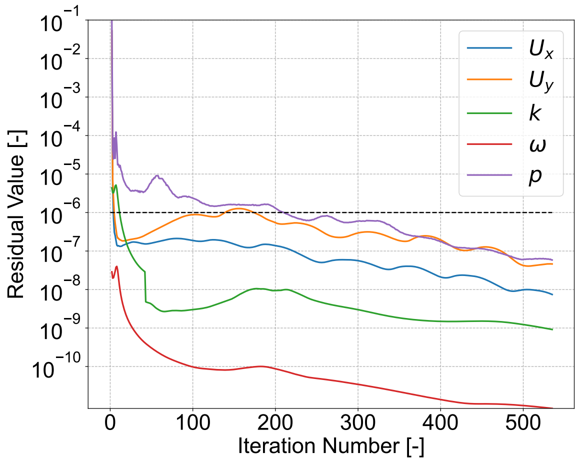

Figure 5.2: Final flow residuals at $\alpha = 6°$ for the following profile parameters: $t = 0.09$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0$, $\lambda = 0.2$, and $\phi = 0.65$.

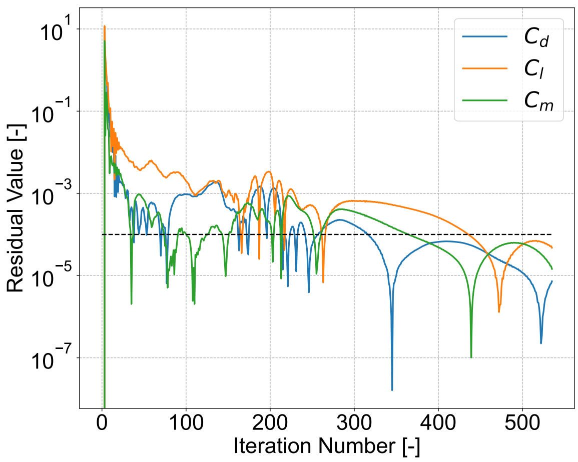

Figure 5.3: Force residuals at $\alpha = 6°$ for the following profile parameters: $t = 0.09$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0$, $\lambda = 0.2$, and $\phi = 0.65$.

5.2.4. Postprocessing

For each simulation, data containing the profile configuration, Reynolds number, AoA, force, and moment coefficients are recorded. For verification, the maximum iteration number, corresponding residuals, and the maximum Y^+ value are also included. A VTK file can be exported to visualise all relevant flow variables, for example, the velocity field with streamlines at $\alpha = 10°$, as shown in Figure 5.4. The recirculation region is clearly visible, comprising a primary recirculation zone and a smaller secondary region near the LE tube-canopy intersection.

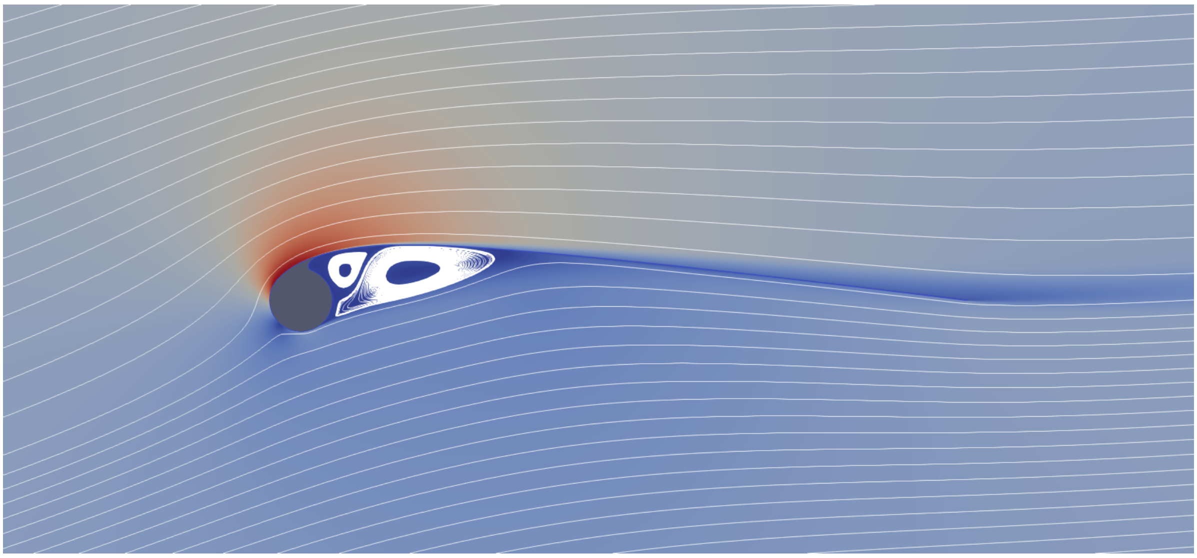

Figure 5.4: Contour plot of the velocity field of a LEI kite profile including the streamlines around the profile at $\alpha = 10°$ for the following profile parameters: $t = 0.09$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0$, $\lambda = 0.2$, and $\phi = 0.65$.

Pressure coefficient (C_p) and skin friction coefficient (C_f) distributions are compiled at each AoA, calculated using Equations 5.6 and 5.7,

$$C_\mathrm{p} = \frac{p - p_\infty}{\frac{1}{2}\,\rho_\infty\,U_\infty^2}, \tag{5.6}$$

$$C_\mathrm{f} = \frac{\tau_\mathrm{w}}{\frac{1}{2}\,\rho_\infty\,U_\infty^2}. \tag{5.7}$$

Where τ_w is wall shear stress, p_∞ the free-stream pressure, and p the local pressure. Since regression modelling involves many simulations, reducing the dataset size without significant information loss is essential. The Ramer-Douglas-Peucker (RDP) algorithm, implemented via the RDP Python library [26], was used to simplify both C_p and C_f curves while preserving their shape. As shown in Figures 5.5 and 5.6, the data was reduced to about 10% of the original points using epsilon values of 0.004 for C_p and 0.0001 for C_f. No significant visual differences were observed.

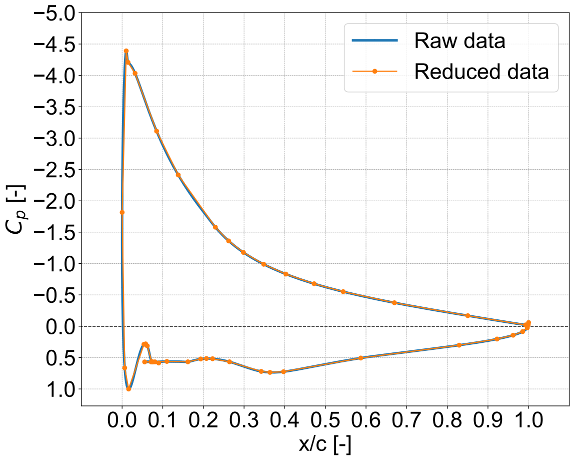

Figure 5.5: Pressure coefficient distribution at $\alpha = 10°$ and $\text{Re} = 5 \times 10^6$ for the following profile parameters: $t = 0.09$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0$, $\lambda = 0.2$, and $\phi = 0.65$.

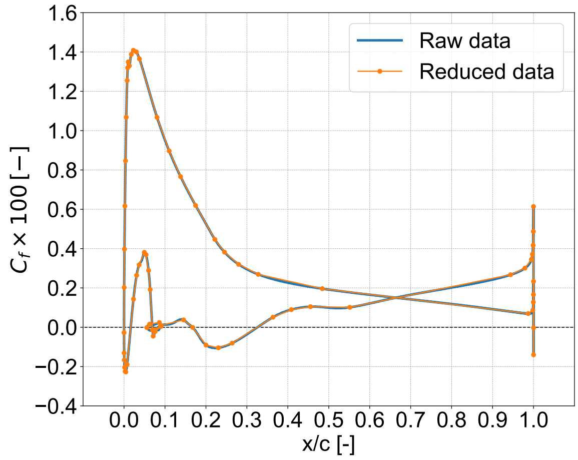

Figure 5.6: Skin friction coefficient distribution at $\alpha = 10°$ and $\text{Re} = 5 \times 10^6$ for the following profile parameters: $t = 0.09$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0$, $\lambda = 0.2$, and $\phi = 0.65$.

The chordwise position of the centre of pressure (x_{cp}) can be determined by integrating the pressure distribution over the airfoil surface. The centre of pressure (CoP) represents the mean location of aerodynamic pressure and is a key metric for assessing airfoil stability. Alternatively, x_{cp} can be calculated using force and moment coefficients with α, as defined in Equation 5.8 [2]. A common simplified form, neglecting drag and α, is $x_{\text{cp}} = x_{\text{ref}} - \frac{C_{\text{m}}}{C_{\text{l}}}$. Figure 5.7 compares all three methods. The results show minimal differences between the direct, simplified, and C_{p} integration-based method. These small differences may arise from numerical integration errors or from reducing the resolution of the C_{p} data.

Since the data for the direct method is available for all simulations, it is used throughout this work. However, care is needed near angles of attack where C_{l} approaches zero. For instance, at $\alpha = -2°$, a sudden shift in x_{cp} occurs. This sensitivity is particularly evident in the simplified expression, $x_{\text{cp}} = x_{\text{ref}} - \frac{C_{\text{m}}}{C_{\text{l}}}$, where small values of C_{l} can cause large fluctuations.

$$x_{\text{cp}} = x_{\text{ref}} - \frac{C_{\text{m}}\, c}{\sqrt{C_{\text{l}}^2 + C_{\text{d}}^2}\, \sin\!\left(\tan^{-1}\frac{C_{\text{l}}}{C_{\text{d}}} + \alpha\right)} \tag{5.8}$$

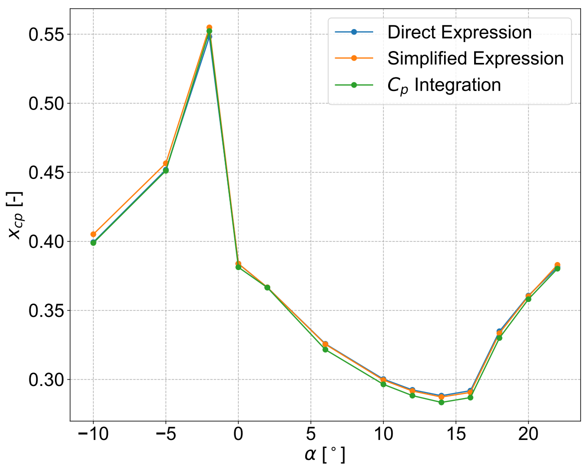

Figure 5.7: CoP versus α at $\text{Re} = 5 \times 10^6$ for the following profile parameters: $t = 0.09$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0$, $\lambda = 0.2$, and $\phi = 0.65$.

6

LEI Kite profile Aerodynamic Analysis

Optimisation of the LEI kite profiles begins with understanding how individual profile parameters influence the flow field and, consequently, aerodynamic performance. All simulations presented in this chapter are conducted at Re $= 5 \times 10^6$. The parameters with the most significant impact, being the non-dimensional LE tube diameter t, chordwise position of maximum camber η, and camber height κ, have been thoroughly analysed by Watchorn [50] and are therefore not discussed in detail here.

Similarly to Watchorn's approach, the C_l, C_d, and C_m coefficients, as well as the C_p and C_f distributions, are analysed here for the updated parametric model, which includes the reflex angle δ in Section 6.1, camber tension λ in Section 6.2, and LE curvature ϕ in Section 6.3. Additionally, the CoP is included as a key stability performance parameter.

6.1. Reflex Angle δ

The reflex angle δ influences the TE profile and is introduced to modify the pitching moment, thereby affecting the CoP. It serves as a passive design feature to increase the nose-up pitching moment, helping to ensure that the kite does not fly too close to the zenith but instead glides naturally to a slightly lower position in the wind window. This behaviour helps to prevent an unrecoverable front stall, which could pose safety risks for users or bystanders. For these reasons, the effect of δ is worth investigating.

The reflex angle is analysed for four small deflection variations, as visualised in Figure 6.1, which result in noticeable changes in both the lift and moment coefficients. The definition of δ is detailed in Section 4.1, where a deflection of $\delta = 0°$ corresponds to a straight canopy, and negative values indicate a downward deflection.

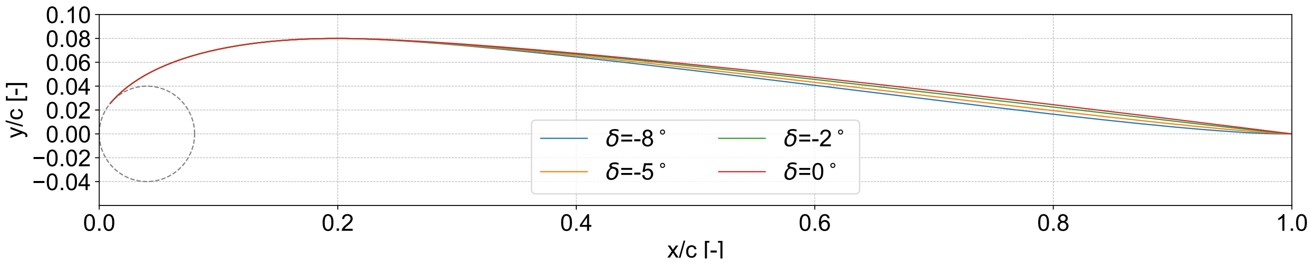

Figure 6.1: LEI kite profile for various δ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\lambda = 0.2$, and $\phi = 0.65$.

Considering the C_d in Figure 6.3, only negligible differences are observed across the simulated reflex angles. However, C_l in Figure 6.2 shows a consistent variation over the entire AoA range, indicating that reflex can negatively impact performance. For this reason, only small deflection angles were considered.

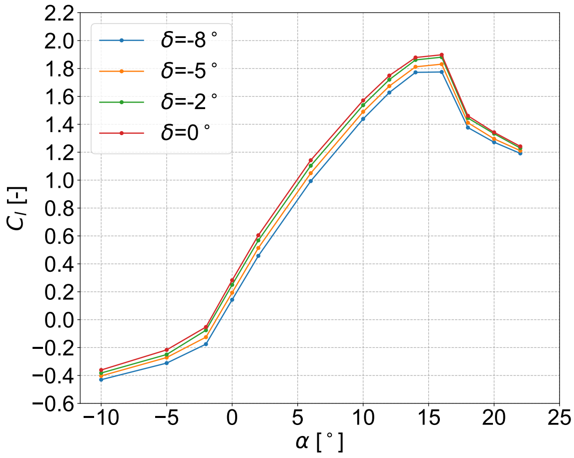

Figure 6.2: C_l versus α for various δ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\lambda = 0.2$, and $\phi = 0.65$.

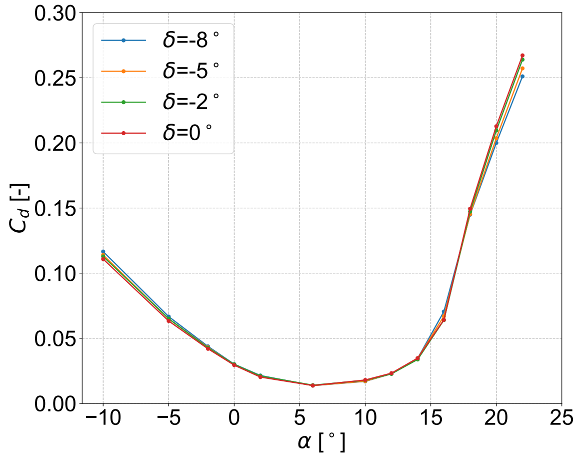

Figure 6.3: C_d versus α for various δ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\lambda = 0.2$, and $\phi = 0.65$.

Examining C_m and the chordwise position of the centre of pressure x_cp in Figures 6.4 and 6.5, the motivation for introducing the reflex becomes clearer. As the downward deflection angle (negative δ) increases, C_m becomes more positive across the full AoA range, indicating a nose-up pitching moment. Consequently, x_cp shifts forward along the chord, which contributes to improved stability characteristics.

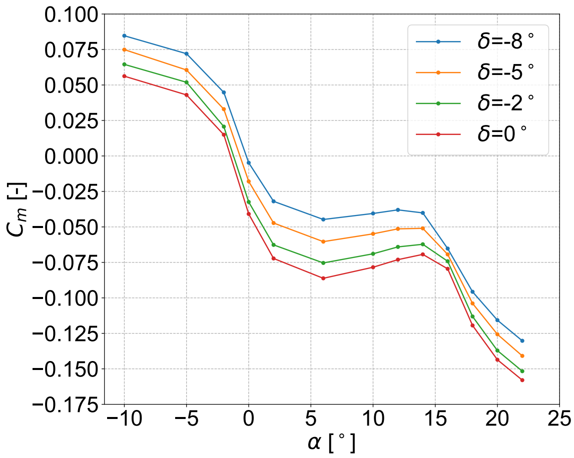

Figure 6.4: C_m versus α for various δ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\lambda = 0.2$, and $\phi = 0.65$.

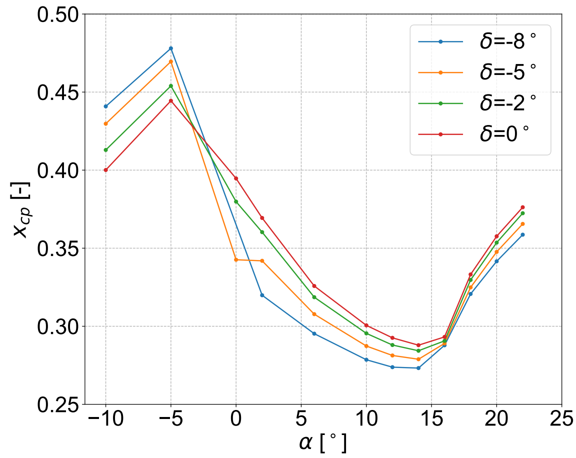

Figure 6.5: x_cp versus α for various δ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\lambda = 0.2$, and $\phi = 0.65$.

An interesting observation arises in the pressure distribution shown in Figure 6.6. While increasing the reflex angle δ results in more positive C_m, one might expect this to correlate with a stronger suction peak near the LE. However, the opposite occurs. As the profile becomes more concave at the upper surface of the TE due to increased reflex, suction on the aft section decreases, and pressure on the lower surface drops due to the locally convex shape, both contributing to an increase in C_m. Surprisingly, suction near the LE also weakens, which would typically reduce C_m. Despite this opposing effect, the overall C_m still increases. This behaviour is explained by global pressure coupling in the flowfield, the pressure field adapts over the entire profile to satisfy continuity and momentum conservation [2].

The C_f distribution is closely related to the C_p plot, as visualised in Figure 6.7. Lower pressure corresponds to locally higher flow velocity, which results in increased skin friction. Additionally, the recirculation region on the lower surface is evident from the negative C_f values, with the flow reattaching around $x/c = 0.4$.

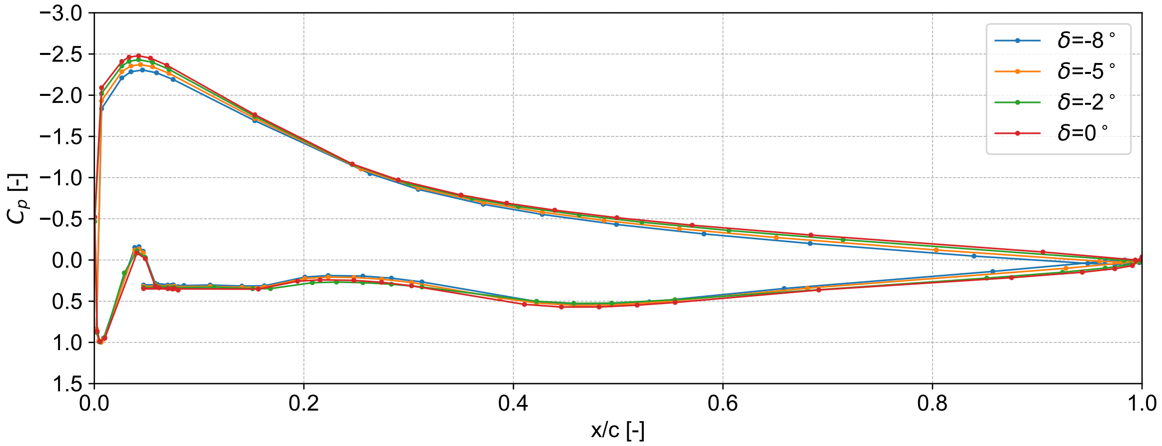

Figure 6.6: C_p distribution at $\alpha = 10°$ for various δ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\lambda = 0.2$, and $\phi = 0.65$.

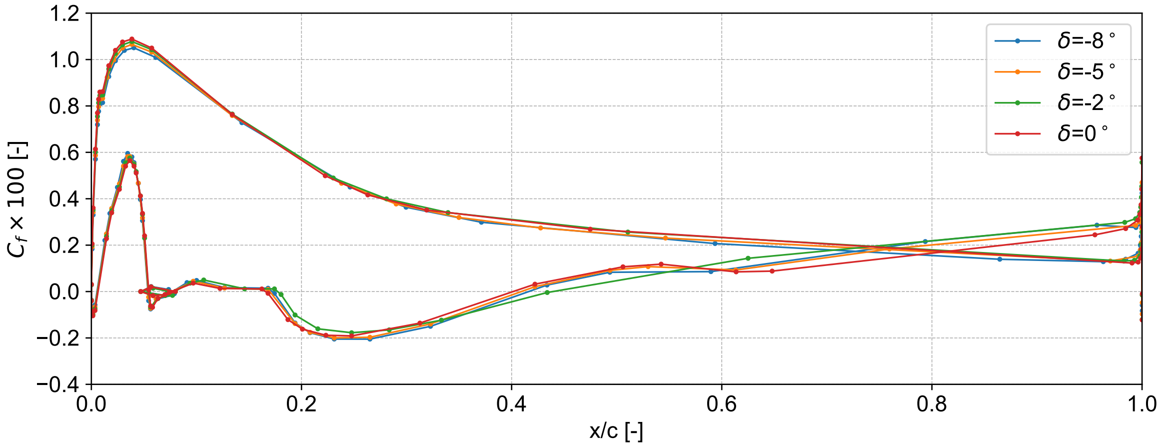

Figure 6.7: C_f distribution at $\alpha = 10°$ for various δ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\lambda = 0.2$, and $\phi = 0.65$.

6.2. Camber Tension λ

The camber tension parameter λ was introduced based on the observation of an abrupt change in curvature from the maximum camber location toward the TE in profiles designed using Surfplan. To address this, a Bézier curve was implemented, using control points to define the rear spline of the profile from the maximum camber location towards the TE. The control point, located near the maximum camber position, defines the so-called camber tension λ. Larger values of λ increase the curvature tension, resulting in smoother rear profiles, as visualised in Figure 6.8. The role of this tension control point is discussed in more detail in Section 4.1.

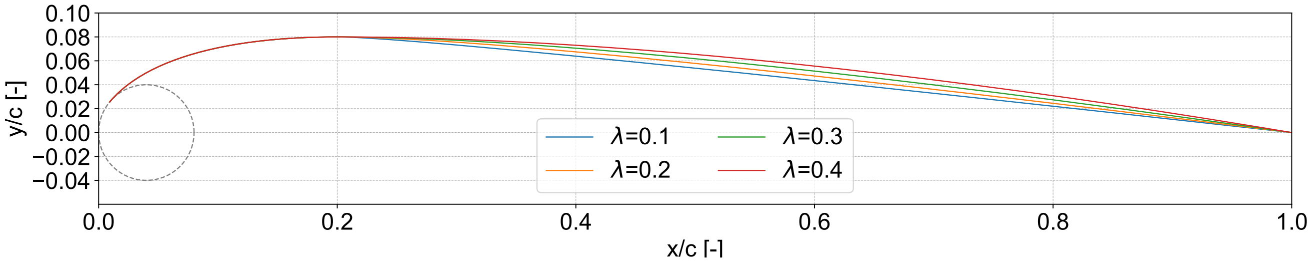

Figure 6.8: LEI kite profile for various λ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, and $\phi = 0.65$.

The lift and drag coefficients, shown in Figures 6.9 and 6.10, reveal a slight increase in C_l across the full range of AoA as the camber tension increases. The C_d remains mostly unchanged within the normal operating range, but exhibits beneficial behaviour at higher angles of attack, particularly at $16°$ and $18°$ AoA.

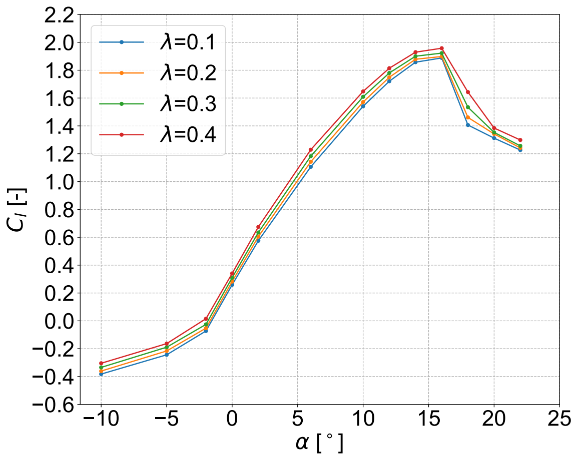

Figure 6.9: C_l versus α for various λ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, and $\phi = 0.65$.

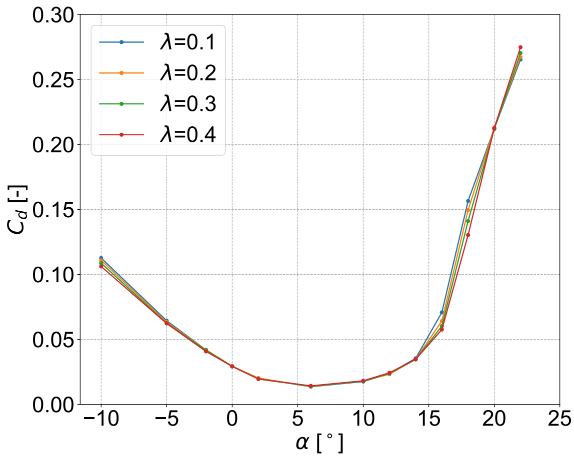

Figure 6.10: C_d versus α for various λ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, and $\phi = 0.65$.

The C_m and CoP, shown in Figures 6.11 and 6.12, exhibit trends similar to those discussed in Section 6.1. For stability reasons, it is desirable to maintain more positive values of C_m. Although the lift increases with higher camber tension λ, the C_m becomes more negative across all angles of attack, indicating reduced stability. As a consequence, the CoP shifts rearward with increasing λ.

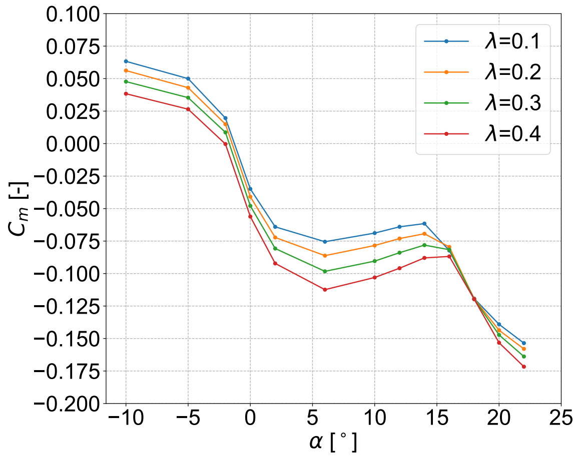

Figure 6.11: C_{m} versus α for various λ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, and $\phi = 0.65$.

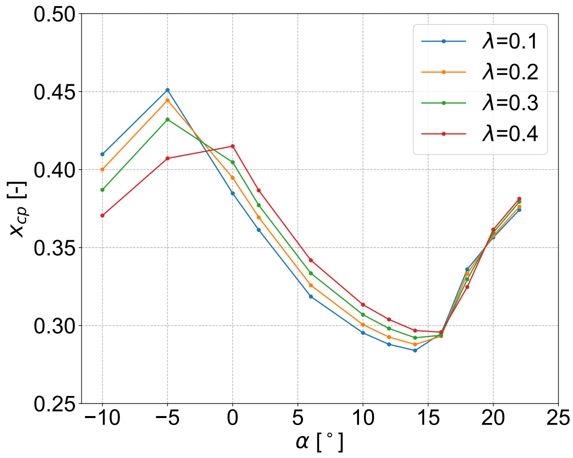

Figure 6.12: x_{cp} versus α for various λ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, and $\phi = 0.65$.

The gain in C_{l} can be explained by examining the C_{p} distribution in Figure 6.13. At higher values of λ, the maximum camber position transitions more smoothly toward the TE, forming a continuous convex spline over a longer portion of the profile. This contrasts with the case of $\lambda = 0.1$, which results in a straighter rear profile. The extended convex curvature increases suction on the upper surface and enhances pressure on the lower surface, thereby contributing to higher lift. Additionally, a slight increase in suction is observed upstream near the LE for larger λ by global pressure coupling in the flowfield, further supporting the lift gain.

With respect to the C_{m}, the additional lift generated by the rear part of the profile at higher λ dominates over the increased suction near the LE, leading to a more negative C_{m}, indicating reduced pitch stability as λ increases.

For the strongest curvature at $\lambda = 0.1$, a distinct adverse pressure gradient step is noticeable in both the C_{p} and C_{f} plots around the maximum camber location at $x/c = 0.2$. As shown in Figure 6.14, slightly higher C_{f} values are observed for larger λ, indicating a modest increase in overall skin friction. The recirculation region on the lower surface remains visible, with negative C_{f} values followed by reattachment occurring around $x/c = 0.4$.

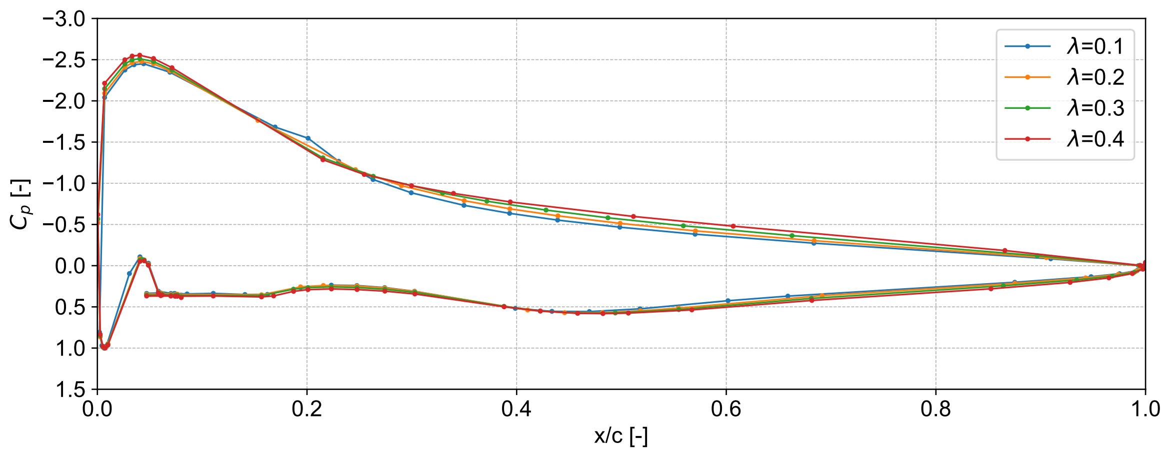

Figure 6.13: C_{p} distribution at $\alpha = 10°$ for various λ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, and $\phi = 0.65$.

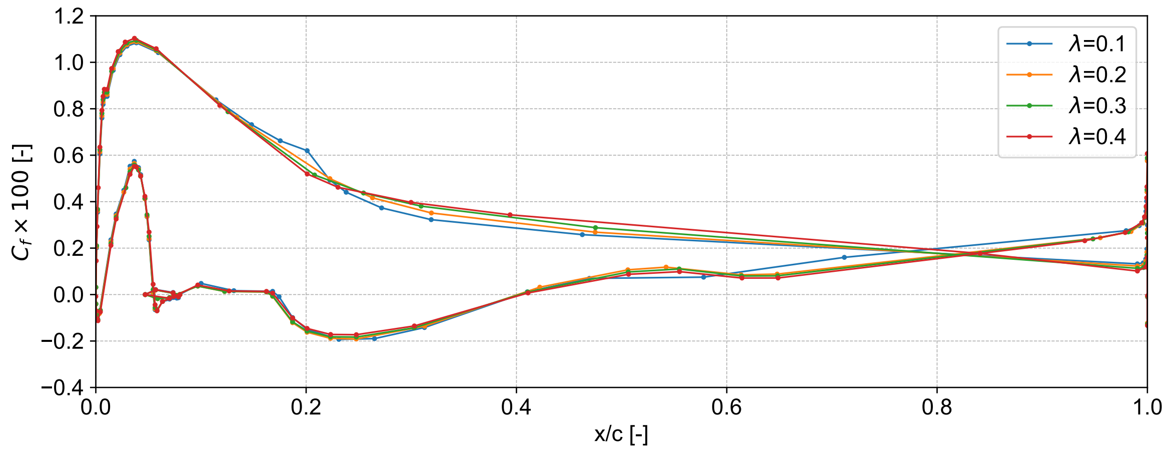

Figure 6.14: C_f distribution at $\alpha = 10°$ for various λ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, and $\phi = 0.65$.

6.3. Leading Edge Curvature ϕ

The LE curvature, parametrised by ϕ, is explained and visualised using control points in Section 4.1. In short, larger values of ϕ move the control points away from the maximum camber position and LE seam, introducing curvature closer to the centre of the front spline. In contrast, smaller values concentrate the curvature near the endpoints of the front spline, specifically at the seam and maximum camber locations. This behaviour is illustrated in Figure 6.15 for various values of ϕ.

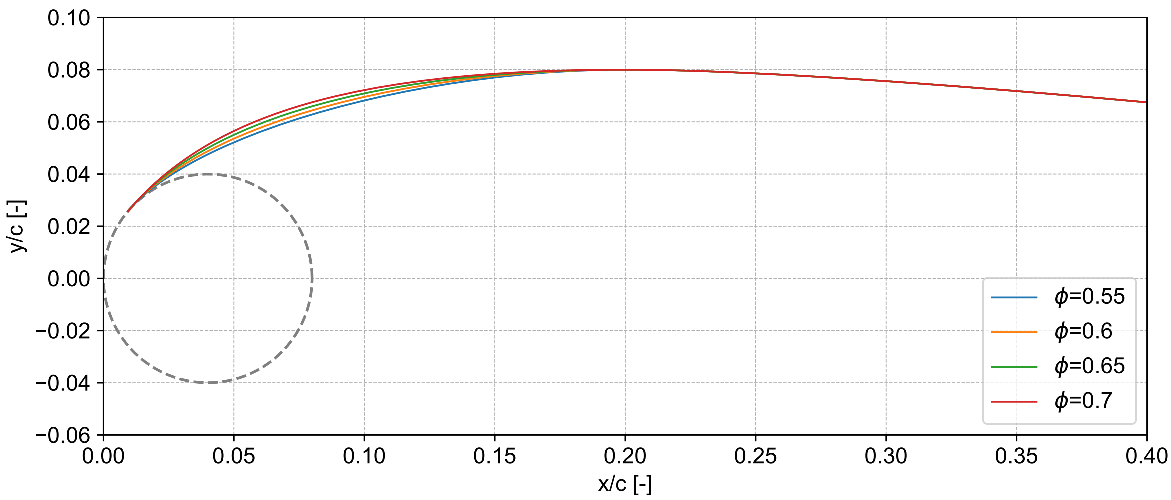

Figure 6.15: LEI kite profile for various ϕ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, and $\lambda = 0.2$.

Analysing the C_l and C_d curves reveals negligible differences throughout the AoA range, except for minor variations beyond stall at $\alpha = 22°$. These differences at high AoA should be interpreted with caution and ideally verified using transient simulations, as steady RANS approaches lack reliability at high AoA.

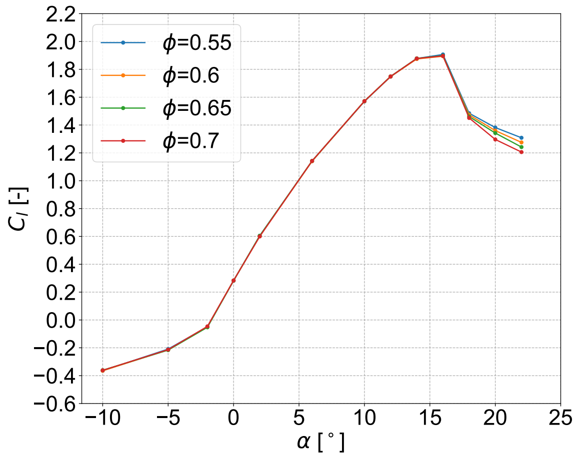

Figure 6.16: C_l versus α for various ϕ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, and $\lambda = 0.2$.

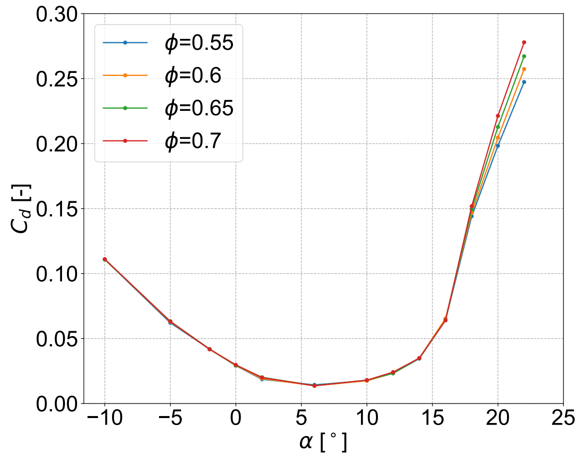

Figure 6.17: C_d versus α for various ϕ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, and $\lambda = 0.2$.

Focusing on C_m and x_cp within the operating range of $0-16°$ AoA, the moment coefficient C_m becomes slightly more positive for larger ϕ values, resulting in a small forward shift of the CoP. This behaviour can be explained by examining the C_p distribution in Figure 6.20. Larger ϕ values produce sharper curvature and more pronounced suction peaks, while smaller values tend to create a broader pressure plateau with two smaller peaks caused by the dual convex regions associated with reduced curvature in the centre of the front spline, as seen for $\phi = 0.55$. Additionally, the C_f distribution for the upper surface, shown in Figure 6.21, closely reflects the trends in the C_p distribution.

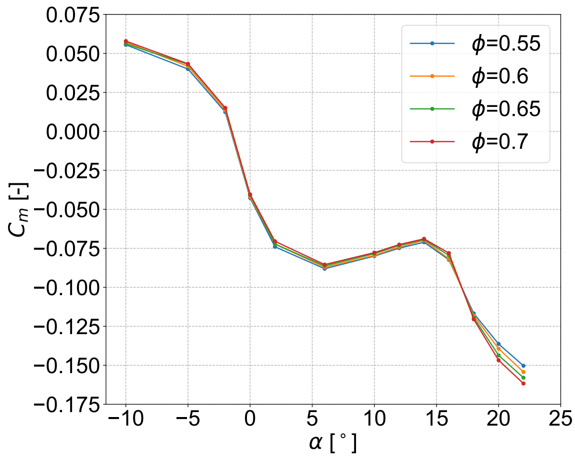

Figure 6.18: C_m versus α for various ϕ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, and $\lambda = 0.2$.

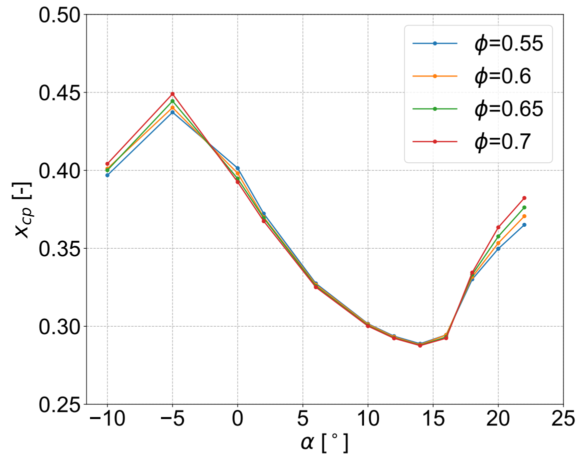

Figure 6.19: x_cp versus α for various ϕ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, and $\lambda = 0.2$.

Although the overall lift coefficient C_l remains nearly constant, this redistribution of pressure causes a slight shift in the CoP without significantly affecting the total suction force. Given that this performance variation is relatively small. Overall, it can be concluded that the LE curvature ϕ has only a minor effect on the aerodynamic performance and can be omitted as a variable, instead being fixed at $\phi = 0.65$.

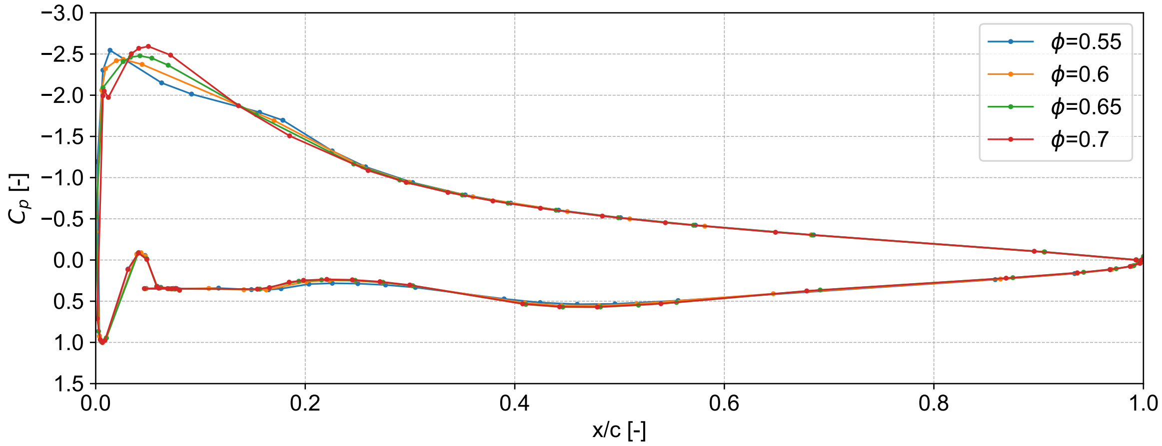

Figure 6.20: C_p distribution at $\alpha = 10°$ for various ϕ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, and $\lambda = 0.2$.

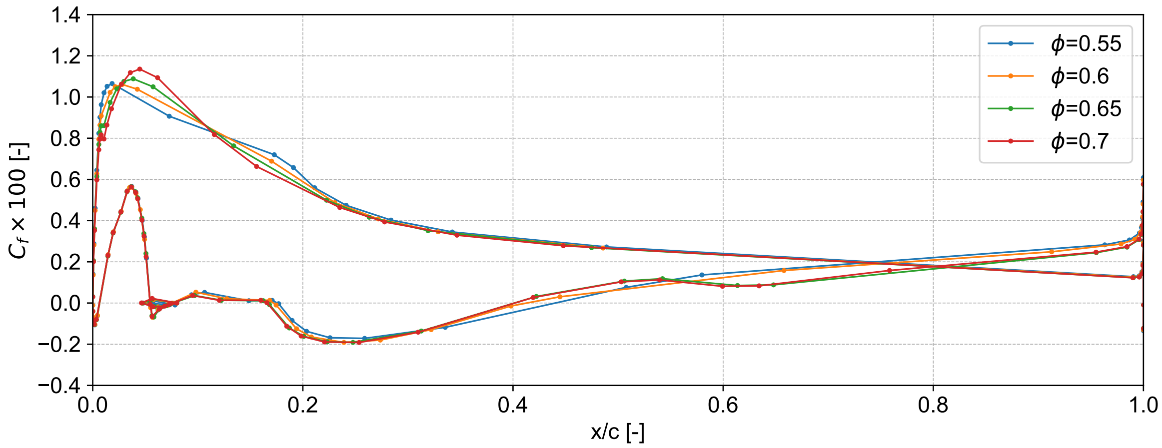

Figure 6.21: C_f distribution at $\alpha = 10°$ for various ϕ, with fixed profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, and $\lambda = 0.2$.

7

Machine Learning for Aerodynamic Regression

This chapter introduces the application of machine learning (ML) algorithms and hyperparameter tuning to develop regression models for predicting aerodynamic performance coefficients (C_l, C_d, C_m), as well as surface distributions of C_p and C_f, for LEI kite profile configurations at various AoA and Reynolds numbers. In Section 7.1, the six-dimensional space comprising geometric profile parameters and AoA is defined across a representative range of Reynolds numbers. The data generation and filtering process, described in Section 7.2, ensures high-quality CFD data by removing non-converged simulations and statistical outliers. In Section 7.3, multiple ML models are trained and tuned using cross-validation. The model's performance is then further improved through spline interpolation. Finally, Section 7.4 evaluates the most promising models on a randomly selected test set, with the tuned extra trees model achieving the highest accuracy, and is visually demonstrated through selected profile comparisons.

7.1. Parameter Definition

With the Reynolds number varying significantly during the crosswind flight and stationary retraction phases of LEI kites, it is worthwhile to investigate the performance differences across this range. Therefore, the lift, drag, and moment polars, as well as the drag-lift polar, are presented in Figures 7.1, 7.2, 7.3, and 7.4 respectively. These results correspond to a conventional profile with following parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.09$, $\delta = 0$, $\lambda = 0.2$ and $\phi = 0.65$, across a Reynolds number range from 6×10^5 to 6×10^7.

A convergence in aerodynamic performance is observed at higher Reynolds numbers, while a greater spread appears at lower values. Even lower Reynolds numbers may occur at the wingtips during static flight due to shorter chord lengths and reduced apparent wind speeds. However, these scenarios are considered less critical for optimisation and are therefore omitted. Instead, three representative models are considered: one at the lower bound, corresponding to smaller kites or static flight (1×10^6), one near the typical in-flight operational conditions for kites (5×10^6), and one representing larger kites or higher wind conditions, and future designs with extended chord lengths (2×10^7) [13].

As expected, the same aerodynamic performance trends were observed, verifying the findings of Folkersma without transition modelling [23]. The maximum lift coefficient $C_{l,\mathrm{max}}$ increases with Reynolds number, while the C_d decreases. The new analysis further reveals that C_m decreases, corresponding to an increasing nose-down pitching moment as the Reynolds number rises. At higher Re, the profile achieves a higher $C_{l,\mathrm{max}}$ and maintains a stable C_m over a longer range before stall. Beyond stall, a sharp drop in C_l is accompanied by a plunge into more negative C_m values.

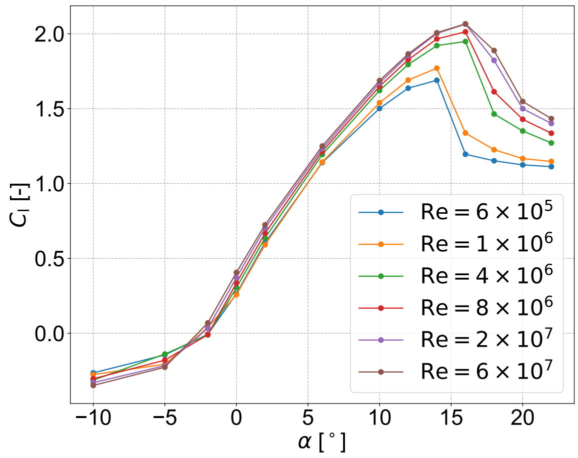

Figure 7.1: C_l versus AoA for varying Reynolds numbers, given the following profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.09$, $\delta = 0°$, $\lambda = 0.2$, and $\phi = 0.65$.

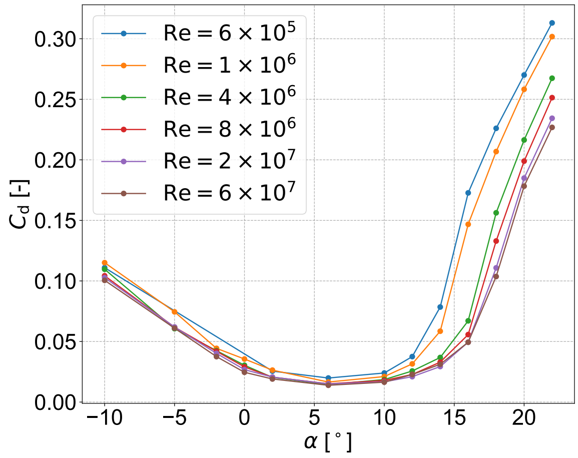

Figure 7.2: C_d versus AoA for varying Reynolds numbers, given the following profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.09$, $\delta = 0°$, $\lambda = 0.2$, and $\phi = 0.65$.

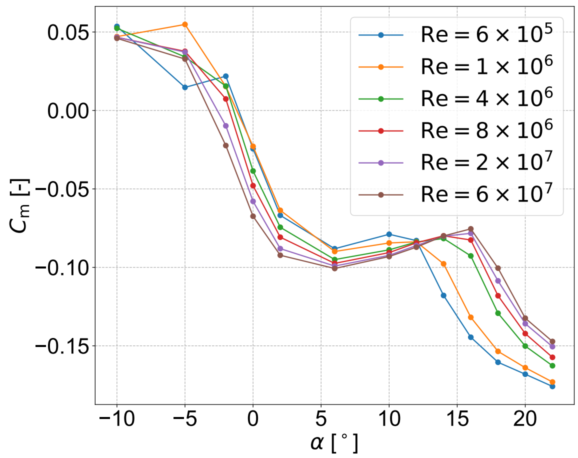

Figure 7.3: C_m versus AoA for varying Reynolds numbers, for the following profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.09$, $\delta = 0°$, $\lambda = 0.2$, and $\phi = 0.65$.

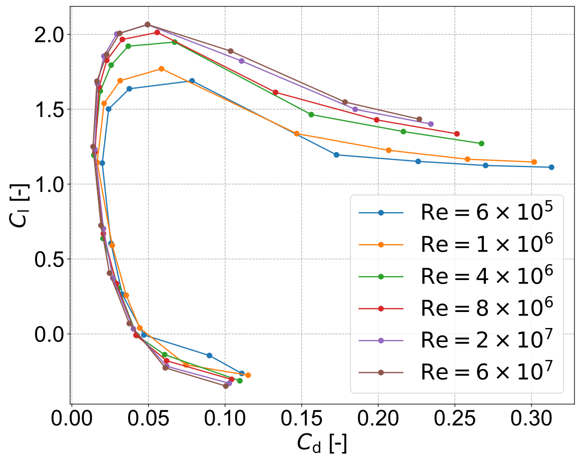

Figure 7.4: C_d versus C_l for varying Reynolds numbers, given the following profile parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.09$, $\delta = 0°$, $\lambda = 0.2$, and $\phi = 0.65$.

With the flow conditions defined, the profile parameters and AoA need to be selected to sufficiently cover the design space and achieve an accurate model. Through discussions with kite designers across applications such as boat towing, airborne wind energy, and kitesurfing, a practical range for each profile parameter was defined, carefully balancing computational cost with data resolution across the design space. To ensure a uniform data distribution, a complete grid covering all possible parameter combinations was generated, creating a dedicated model for each Re. Additionally, based on the aerodynamic analysis in Section 6.3, the parameter ϕ was found to have minimal impact on aerodynamic performance. Therefore, considering computational efficiency, this parameter was fixed at a value of 0.65. Allowing greater refinement of other, more influential parameters. Regarding the AoA range, the distribution was chosen to optimise accuracy in the lift polar while maintaining computational efficiency. Thus, fewer data points were used in the linear lift region (around $\alpha = 2°$ to $\alpha = 10°$). Negative AoA were also sampled more coarsely due to less importance. Around the stall region, a finer AoA step size of $2°$ was chosen to accurately capture the stall behaviour. Ultimately, the finalised parameter ranges are summarised in Table 7.1. With ϕ fixed and a separate model constructed for each Re, this results in a six-dimensional profile geometry space.

Table 7.1: Design parameter ranges of the dataset.

Parameter	Variables	Count
t [-]	0.03, 0.06, 0.08, 0.1, 0.12	5
η [-]	0.08, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35, 0.4	8
κ [-]	0.04, 0.06, 0.08, 0.1, 0.12, 0.14, 0.16	7
δ [°]	-8, -5, -2, 0	4
λ [-]	0.1, 0.2, 0.3, 0.4	4
ϕ [-]	0.65	1
α [°]	-10, -5, -2, 0, 2, 6, 10, 12, 14, 16, 18, 20, 22	13
Re [-]	1×10^6, 5×10^6, 2×10^7	3

A preprocessing step was performed to remove profile configurations outside the applicable design space. Profiles were filtered based on the following criteria:

- Profiles where the camber tension control point λ is positioned too far aft, i.e., when $\eta + \lambda \geq 0.8$.
- Profiles for which the camber height κ is excessively large compared to thickness t, specifically when $\kappa > 2.5, t$.
- Additionally, profiles with κ less than $\frac{t}{2}$ were considered flat (no camber). These profiles are simplified to the LE tube diameter t and the reflection angle δ, as other parameters have minimal influence. A flat profile representation can be seen in Figure 4.2.

Profiles with larger non-dimensional tube diameters and low camber are retained, as they represent wingtip configurations necessary for the 3D aerodynamic analysis using the VSM model. The total number of filtered profile configurations is summarised in Table 7.2. A total of 3,464 configurations were evaluated across 13 angles of attack, leading to 45,032 simulations per Reynolds number (Re).

Table 7.2: Number of profile configurations per t in the dataset.

Thickness t [-]	Count
0.03	256
0.06	768
0.08	896
0.10	772
0.12	772
Total	3,464

7.2. Data Generation and Preprocessing

In chapter 5, the steady RANS $k-\omega$ SST simulations are described in detail, including the use of the OpenFOAM solver and the HPC-12 supercomputer platform to compute the aerodynamic coefficients for all profile configurations. This also includes the boundary condition setup and postprocessing steps to extract surface distributions of C_p and C_f. The resulting simulation data, along with the integrated force coefficients C_d, C_l, and C_m, is loaded for further analysis.

Although a thorough sensitivity analysis was conducted using a fully structured hyperbolic mesh and solver tuning, including relaxation factor adjustments and a low initial residual convergence threshold of 8×10^{-7}, a limited number of outliers and non-converged cases remained. To further improve the quality of the model input, a data filtering step was applied to ensure a cleaner training dataset.

All non-converged simulations, defined as cases reaching a maximum of 5000 iterations, were removed. This was especially relevant at high AoA values beyond stall, where flow separation and large recirculation regions occur. Such conditions cannot be accurately resolved by a steady-state RANS solver and would instead require a transient simulation approach. The effect of removing these non-converged cases is illustrated in the lift, drag, and moment polars shown in Figures 7.5, 7.6, and 7.7.

To identify and exclude outliers in the force coefficient output data, a z-score filtering Equation 7.1 approach was used [17]. Specifically, for each AoA, any data point with an absolute z-score greater than approximately 3 standard deviations $\pm 3\sigma$ from the mean (μ) was classified as an outlier and removed. These outliers are visualised as orange markers near the distribution boundaries per AoA in Figures 7.5, 7.6, and 7.7, and their removal is quantified in Table 7.3. The percentage of excluded points remains low, confirming that aside from the non-converged simulations, no significant outliers were present. This mild filtering approach improves the overall model accuracy and standard deviation.

$$z = \frac{x - \mu}{\sigma} \tag{7.1}$$

As described in Section 7.1, the profile geometry dataset spans a large six-dimensional grid defined by all possible combinations of the selected parameter variables [1]. Building a predictive model within this space also requires a test dataset to evaluate the model's accuracy and standard deviation. To ensure reliable verification of the model across the grid, the test set was constructed from randomly sampled profile configurations that satisfy the three configuration criteria outlined in Section 7.1, while using AoA data points that differ from those in the training set. From this point forward, the data used to train the model is referred to as the training data, and the verification set as the test data.

Table 7.3 presents the percentage of non-converged and outlier-removed samples for both datasets. After filtering, the final dataset consists of 42,183 training samples and 5,016 test samples, corresponding to a test-train split of 10.6%.

Table 7.3: Data filtering steps and resulting sizes for Re $= 5 \times 10^6$ dataset.

Step	Train Data	Test Data
Simulation data	45,032	5,187
Not converged (removed)	2,452 (5.45%)	118 (2.27%)
Outliers (removed)	397 (0.93%)	53 (1.05%)
Final dataset size	42,183	5,016

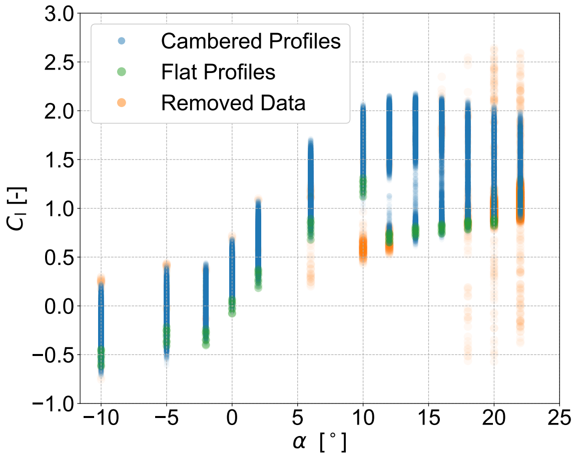

Figure 7.5: C_l versus α for the Re $= 5 \times 10^6$ dataset, illustrating filtering of cambered, flat, and removed profiles.

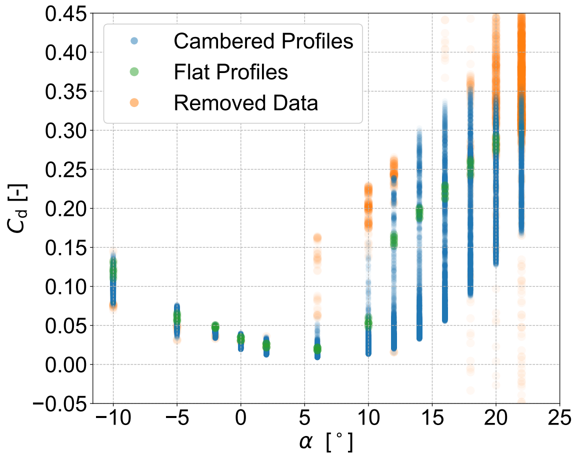

Figure 7.6: C_d versus α for the Re $= 5 \times 10^6$ dataset, illustrating filtering of cambered, flat, and removed profiles.

[1]The source code for the CFD simulations, figure generation, and prediction models is available on GitHub: `https://github.com/awegroup/Pointwise-Openfoam-toolchain`. The CFD simulation dataset used to train the ML model is archived on Zenodo: `https://doi.org/10.5281/zenodo.15593140`.

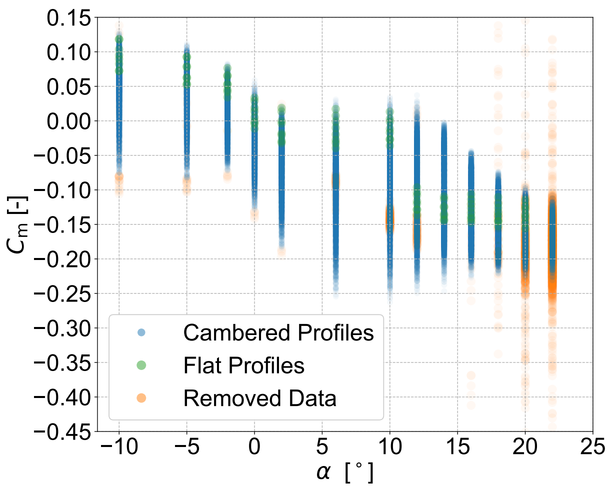

Figure 7.7: C_m versus α for the Re $= 5 \times 10^6$ dataset, illustrating filtering of cambered, flat, and removed profiles.

To visualise the flat (no camber) profiles described in Section 7.1, these cases are highlighted in green in the figures above. Only 104 samples, corresponding to approximately 0.23% of the training data, are classified as flat profiles. Analysing their aerodynamic behaviour reveals that flat profiles, as expected, generate less lift and tend to stall earlier, around $\alpha = 10°$. Consequently, they exhibit relatively high C_d beyond stall, and show a higher C_m, indicating a stronger nose-up tendency in the pre-stall flight regime.

7.3. Machine Learning Model Setup and Postprocessing

The machine learning model setup for aerodynamic force prediction is organised as a modular pipeline, illustrated in Figure 7.8. The process begins with the standardisation of the filtered training datasets, using a Standard Scaler to ensure consistent scaling of all parameters.

Initially, several regression algorithms are evaluated using their default settings. The selected algorithms include k-nearest neighbors (KNN), random forest (RF), extra trees (ET), support vector regression (SVR), and multi-layer perceptron (MLP), all implemented via the scikit-learn library [36]. These models were chosen based on their strong default regression performance, as reported in [45] and facilitated by the `lazypredict` library [35]. Linear regression (LR) is also included for baseline comparison.

Subsequently, for each algorithm, a dedicated hyperparameter search space was defined and systematically explored using grid search cross-validation (GridSearchCV). This process enables effective fine-tuning of each model's configuration based on a selected scoring metric. The coefficient of determination (R^2) was chosen over root mean squared error (RMSE), as it provides a more informative measure of model performance for this dataset. Unlike RMSE, R^2 takes into account the variance of the true target values, as reflected in the denominator of Equation 7.2 compared to the definition of RMSE provided in Equation 7.3. Both expressions are adapted from A Modern Introduction to Probability and Statistics [17].

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_\mathrm{i} - \hat{y}_\mathrm{i})^2}{\sum_{i=1}^{n}(y_\mathrm{i} - \overline{y})^2} \tag{7.2}$$

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_\mathrm{i} - \hat{y}_\mathrm{i})^2} \tag{7.3}$$

In the above equations, y_i denotes the true value of each test data point, $\hat{y}_\mathrm{i}$ represents the corresponding model prediction. $\overline{y}$ is the mean value of the true data, and n is the total data size in the test set.

It is important to note that in the context of grid search cross-validation, the term validation does not refer to comparing model predictions against experimental data. Instead, in ML language, validation refers to the internal process of fine-tuning model hyperparameters using subsets of the training data, prior to evaluating the final model on an independent test set.

During grid search cross-validation, the training data is split into multiple subsets, known as folds. This technique, referred to as k-fold cross-validation, typically involves 5 to 10 folds, depending on the size of the dataset. In this study, the dataset was sufficiently large to support a more robust evaluation, and a configuration of $k = 10$ was selected. This means that in each iteration, 90% of the data is used for training and 10% for validation. While increasing the number of folds can enhance R^2 scoring, it may also introduce greater variance in the results. However, given the size of the dataset, this trade-off was carefully evaluated, and the chosen setting was considered appropriate.

In Table 7.4, the machine learning algorithms and their corresponding hyperparameter settings that achieved the highest mean R^2 values are summarised, with extra trees and Random Forest emerging as the top-performing models.

Table 7.4: Overview of ML models, hyperparameters, and R^2 values.

ML Model	Hyperparameter	Value	R^2
K-nearest neighbors (KNN)	`n_neighbors`	8	0.945
Linear regression (LR)	–	–	0.617

Random forest (RF)	`n_estimators`
	`max_features`
	`min_samples_split`
		130
		sqrt
		2
			0.993

Extra trees (ET)	`n_estimators`
	`max_depth`
	`max_features`
	`min_samples_split`
		130
		None
		log2
		2
			0.991

Support vector regression (SVR)	`kernel`
	`C`
	`epsilon`
	`gamma`
		rbf
		1
		0.1
		scale
			0.737

Multi-layer perceptron (MLP)	`hidden_layer_sizes` `activation`	100 tanh	0.941

An additional postprocessing step was introduced to improve prediction accuracy at AoA values that fall between the discrete data points used during model training, visualised in the flowchart 7.8. Most of the ML models employed, except for SVR, lack smooth continuity in AoA, leading to less reliable predictions in between training points, as observed in the random test data used for model verification. Given that AoA strongly influences aerodynamic performance, as supported by the parameter importance weights shown in Table 7.5, a more refined approach is required. To address this, a piecewise cubic Hermite interpolating polynomial (PCHIP) spline is applied, using the `scipy.interpolate` module [49]. This enables smooth interpolation and more accurate aerodynamic performance predictions for any given profile configuration and AoA.

For the C_p and C_f distribution predictions, each distribution consists of approximately 80 data points, too many for conventional ML approaches. Instead, the KNeighbors algorithm is employed to identify, for a given configuration and AoA, the closest simulated C_p and C_f distributions from the filtered training data. Since certain geometric parameters have a greater influence on the prediction than others, a weighted KNeighbors model is utilised, with the weights determined via permutation importance. Permutation importance estimates feature importance by randomly shuffling each parameter across the dataset and measuring the resulting drop in model performance, as indicated by the reduction in R^2 score. A larger drop signifies a more influential feature. In this study, permutation importance was calculated using the scikit-learn library [36]. The resulting parameter weights, presented in Table 7.5, reflect the relative importance of each parameter in predicting the C_p and C_f distributions.

Table 7.5: Weights used for each parameter in the KNeighbors algorithm.

α	t	η	κ	δ	λ
1.758	0.0470	0.2441	0.250	0.0170	0.0278

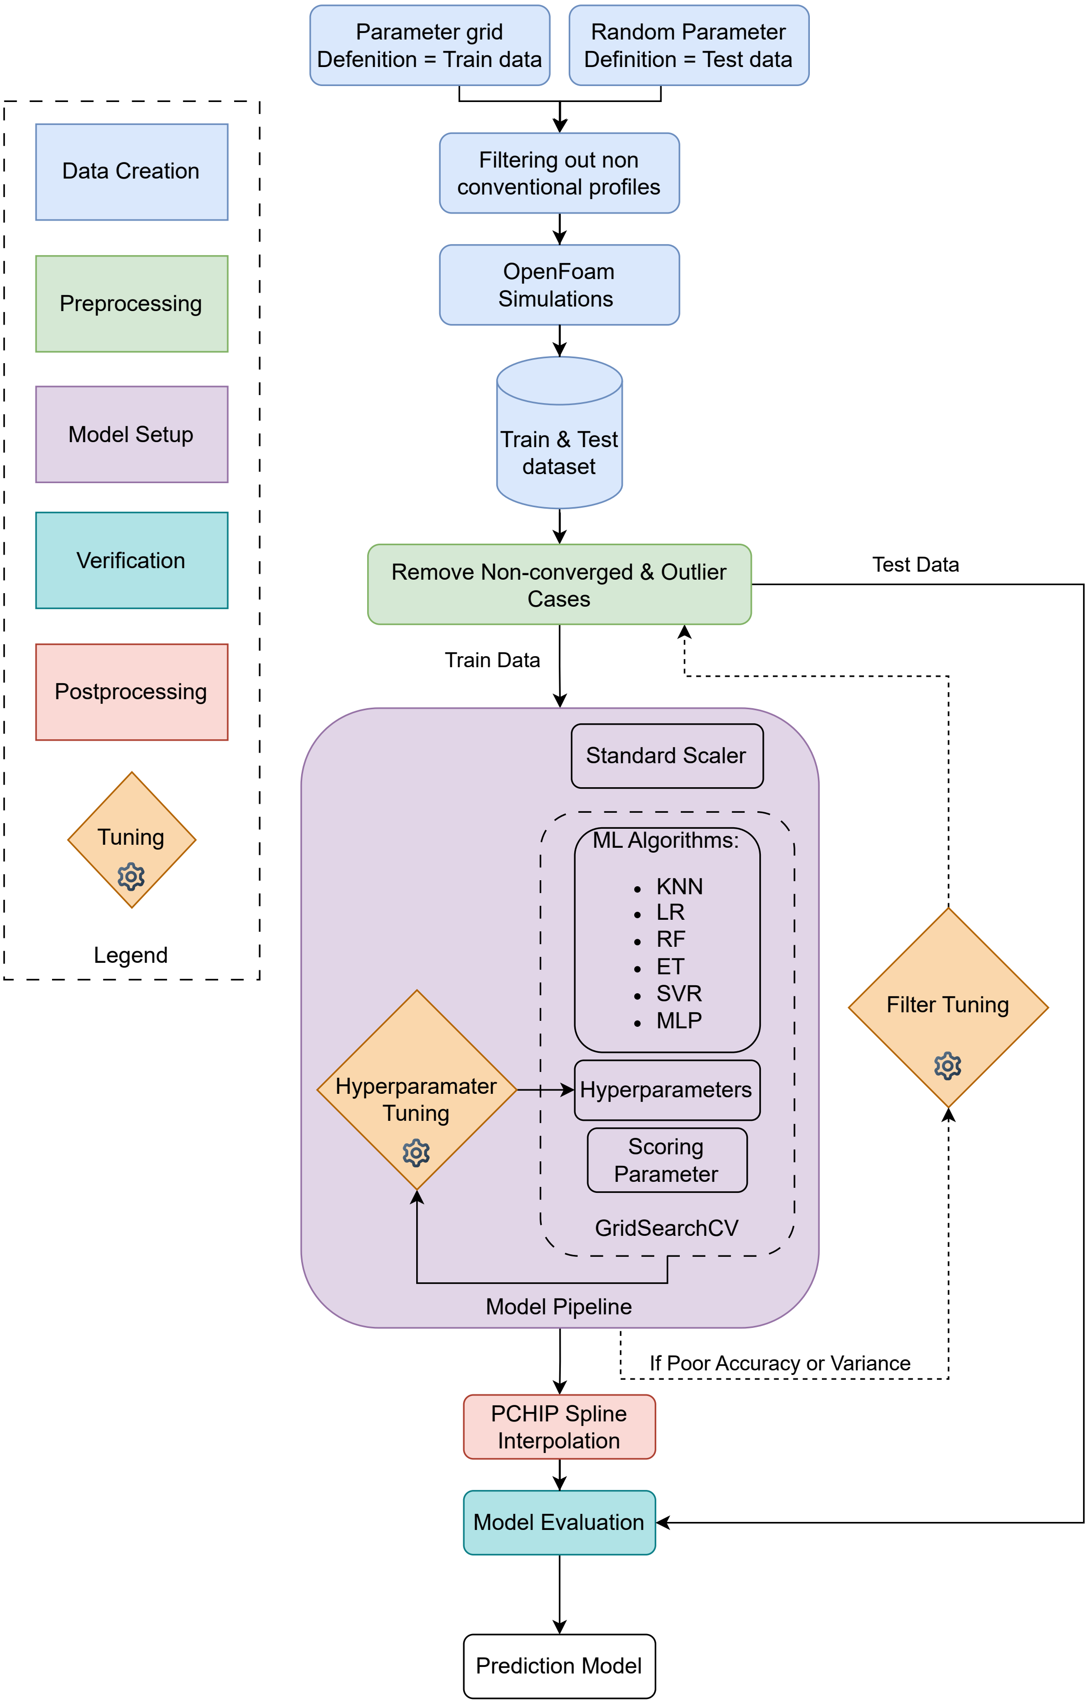

Figure 7.8: Machine learning model setup flowchart, from data preprocessing to model evaluation.

7.4. Model Evaluation

The final dataset that remains completely unseen by the ML models is the test dataset, as described in Section 7.2. After filtering the data, a test-train split of 10.6% remains, which, given the large overall dataset size of approximately 47,000 data points, was deemed sufficient to ensure statistical reliability without significantly increasing variance. To verify whether the high R^2 values observed for Random Forest and extra trees on the grid dataset indicate true model reliability, these models will be further evaluated using the randomly selected test set. For evaluation, both R^2 and the root mean square error (RMSE) as a percentage are reported, to accurately capture differences in the magnitude of aerodynamic forces. Additionally, the percentage of the standard deviation (σ), as defined in Equations 7.4 and 7.5 and following the approach from A Modern Introduction to Probability and Statistics [17], is also presented to provide a comprehensive assessment of model performance.

$$\text{RMSE } [\%] = \frac{\sqrt{\text{MSE}}}{\langle |y_{\text{test}}| \rangle} \times 100 \tag{7.4}$$

$$\sigma \ [\%] = \frac{\sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(r_{\text{i}} - \bar{r})^2}}{\langle |y_{\text{test}}| \rangle} \times 100 \tag{7.5}$$

In the above equations, n is the number of test samples. The mean square error (MSE) is the average squared difference between the true (y_{i}) and predicted ($\hat{y}_{\text{i}}$) values. $\langle |y_{\text{test}}| \rangle$ is the mean absolute value of the test targets, used to normalise both RMSE and σ so they are expressed as percentages. In Equation 7.5, r_{i} are the residuals and $\bar{r}$ is their mean.

The performance of each machine learning model on the test dataset evaluated using R^2, RMSE, and σ is summarised in Table 7.6. Among all models, extra trees stands out by achieving the highest R^2 and the lowest RMSE and variance, demonstrating both high accuracy and consistent predictions. While Random Forest also performs well, it is clearly outperformed by extra trees across all metrics. KNN and MLP show moderate performance with higher prediction errors, whereas LR and SVR yield the least accurate results. Based on these findings, the extra trees model is selected as the preferred prediction method.

Extra trees (extremely randomised trees) build an ensemble of decision trees using two key sources of randomness: it selects split points completely at random rather than optimising them, and it uses the entire training dataset instead of bootstrap samples [21]. These differences make the algorithm both faster and less prone to overfitting. Its high level of randomisation enhances generalisation, particularly in high-dimensional spaces, as demonstrated in Table 7.6.

Table 7.6: ML model verification using R^2, percentage RMSE, and σ, for the Re $= 5 \times 10^6$ model.

ML Algorithm	R^2	RMSE [%]	σ [%]	Training Time [s]
KNN	0.9576	18.54	18.28	0.05
LR	0.6169	55.50	55.08	0.04
RF	0.9442	21.15	24.08	10.6
ET	0.9857	10.87	9.83	9.1
SVR	0.7470	43.10	18.23	46.3
MLP	0.9366	22.42	17.02	7.4

To visualise the performance of the extra trees model[2], three representative profiles were selected from the test dataset: one with a small LE tube diameter ($t = 0.032$), one with a medium diameter ($t = 0.086$), and one with a large diameter ($t = 0.115$). The third profile is notable as it represents a flat configuration, providing insight into the model's performance on edge cases. The prediction metrics for these profiles are summarised in Table 7.7.

Profile 2 demonstrates excellent prediction accuracy, with high R^2 values and low RMSE percentages across all aerodynamic coefficients. These results are further illustrated for C_l and C_d in Figure 7.9, and for C_m in Figure 7.10. In these plots, the PCHIP spline forms a smooth connection between the ML predictions at the AoA values present in the training data, allowing aerodynamic coefficients to be extrapolated for comparison with the true test data.

Generally, a lower R^2 is associated with a higher RMSE, although RMSE is more sensitive to outliers. For example, in the moment polar for Profile 1, the predictions are mostly accurate, but a single outlier near $\alpha = -3°$ causes an increase in RMSE. Similarly, Profile 3 shows an outlier around $\alpha = 11°$. In the drag polar for Profile 1, a larger outlier at $\alpha = 18°$ leads to an RMSE for C_d as high as 25.3%.

Table 7.7: Prediction performance metrics (R^2, RMSE) for Profiles 1-3, listed from top to bottom.

t	η	κ	δ	ϕ	λ	$R^2_{C_\mathrm{l}}$	$R^2_{C_\mathrm{d}}$	$R^2_{C_\mathrm{m}}$	RMSE$_{C_\mathrm{l}}$ [%]	RMSE$_{C_\mathrm{d}}$ [%]	RMSE$_{C_\mathrm{m}}$ [%]
0.032	0.308	0.047	-3.5	0.65	0.134	0.980	0.959	0.931	11.4	25.3	25.8
0.086	0.192	0.08	-5.6	0.65	0.325	0.998	0.996	0.999	3.3	6.7	2.3
0.115	0	0	-7.7	0	0	0.950	0.996	0.955	19.6	6.0	31.0

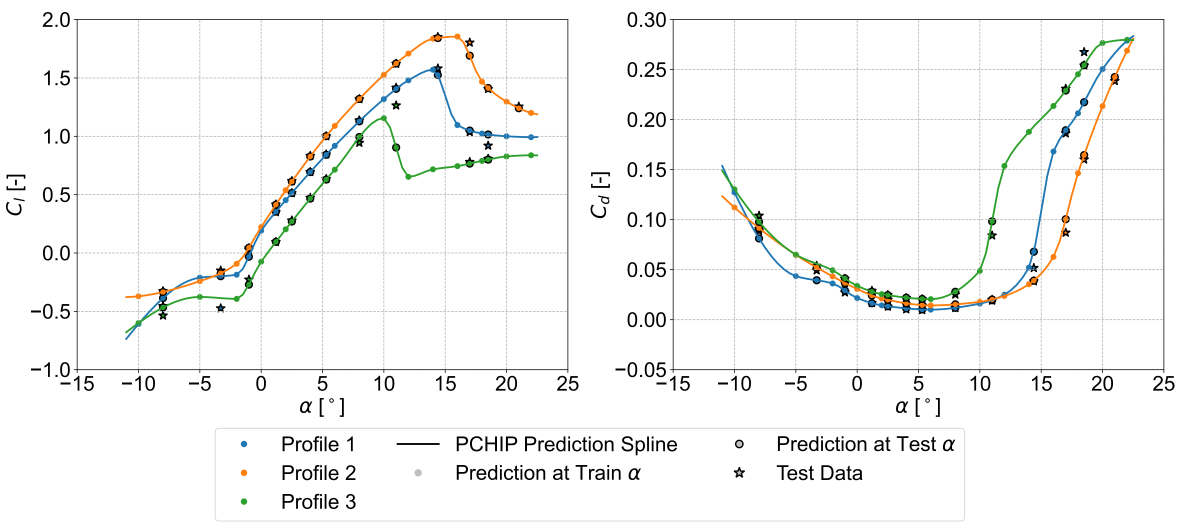

Figure 7.9: Verification of the ML model using lift and drag polars (left and right, respectively) for the $\mathrm{Re} = 5 \times 10^6$ dataset.

[2]The source code for the CFD simulations, figure generation, and prediction models is available on GitHub: `https://github.com/awegroup/Pointwise-Openfoam-toolchain`. The CFD simulation dataset used to train the ML model is archived on Zenodo: `https://doi.org/10.5281/zenodo.15593140`.

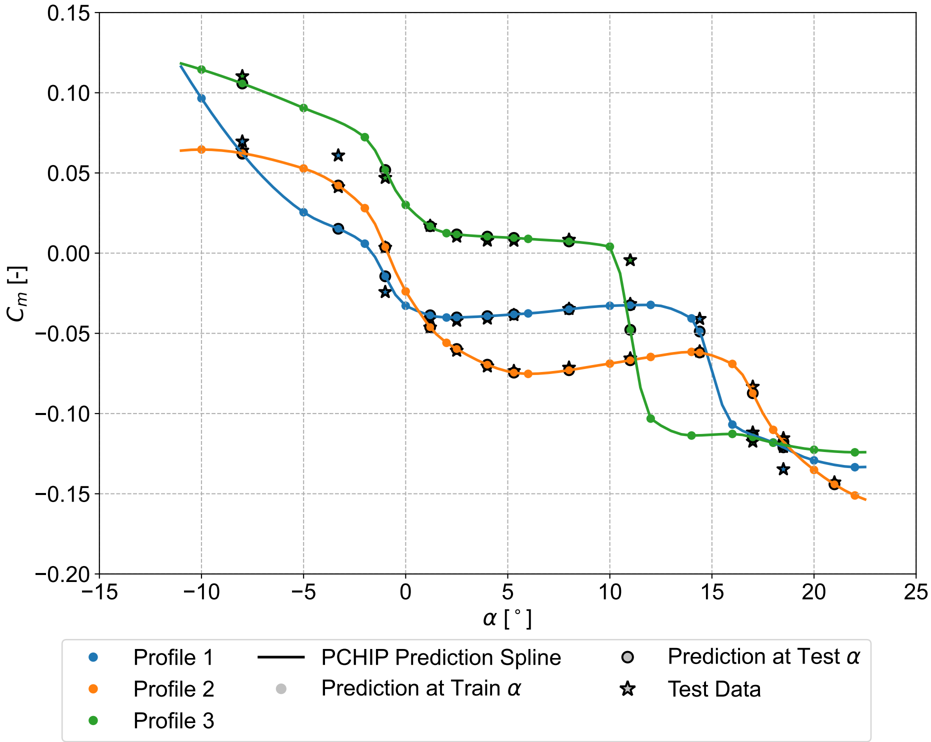

Figure 7.10: Verification of the ML model using moment polar for the Re $= 5 \times 10^6$ dataset.

To assess the C_p and C_f distributions, as illustrated in Figure 7.11, three random profiles from the test dataset are analysed, each with a distinct AoA to ensure clear differentiation. The AoA was selected as the varying parameter, as it has the most significant influence on aerodynamic performance, as shown in Table 7.5. The chosen AoA values were $\alpha = 5.3°$, $11°$, and $14.4°$. The selected profiles and their corresponding k-nearest neighbor are listed in Table 7.8.

The predicted C_p distributions show strong agreement with the reference profiles at each AoA, demonstrating that the k-nearest neighbors model successfully identifies the closest possible matches within the dataset, as also confirmed in Table 7.8. Moreover, the weighted parameter search proves particularly beneficial when certain reference profiles are excluded or unavailable in the reference dataset.

With respect to the C_f distribution, the model yields accurate predictions, especially on the upper surfaces of the second and third profiles, where excellent agreement with the reference data is observed. Minor discrepancies are noted on the lower surfaces, primarily due to the effects of vorticity in the recirculation region, which introduces localised deviations.

Table 7.8: Profile configurations and α values for test data and KNeighbors predictions from training data.

Profile No.	Test Profile						KNeighbors Prediction					
	α	t	η	κ	δ	λ	α	t	η	κ	δ	λ
1	5.3	0.03	0.315	0.051	-7.6	0.109	6	0.03	0.3	0.06	-8	0.1
2	11.0	0.08	0.311	0.062	-5.6	0.302	10	0.08	0.3	0.06	-5	0.3
3	14.4	0.085	0.31	0.12	-4.2	0.23	14	0.08	0.3	0.12	-5	0.2

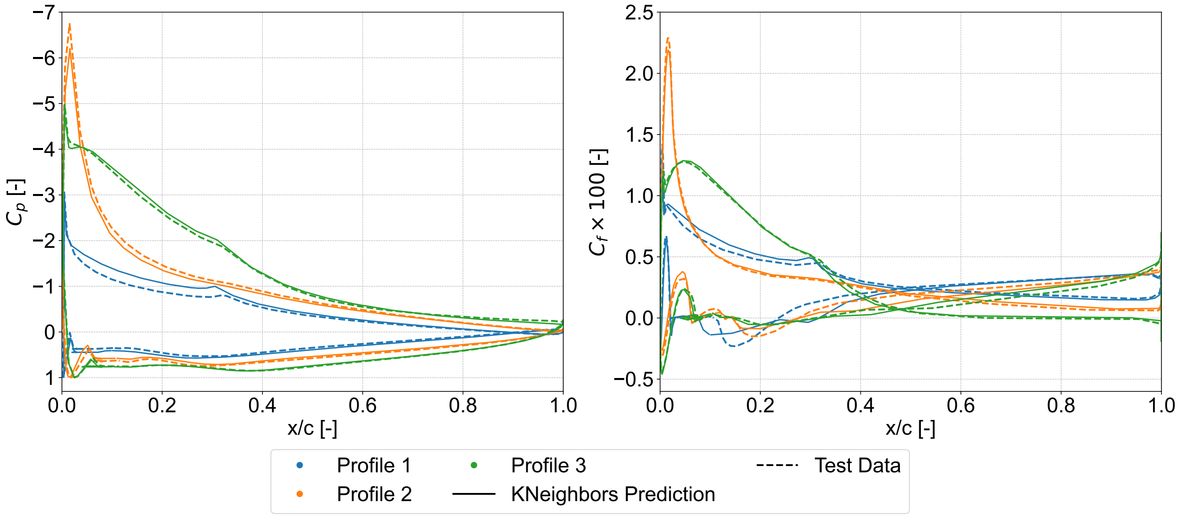

Figure 7.11: Verification of C_p and C_f distributions for the test data and KNeighbors predictions at Re $= 5 \times 10^6$.

The overall regression model results presented in Table 7.9 demonstrate consistently high predictive performance across all Reynolds numbers, as indicated by the total R^2 values above 0.98. The R^2 performance metric measures how well the predicted values match the actual data. With, R^2 values close to 1 across all aerodynamic coefficients (C_l, C_d, and C_m) confirm the model's strong accuracy and applicability.

The total standard deviation (σ) values, expressed in percentage, show a decreasing trend with increasing Reynolds number, indicating improved prediction consistency at higher Re. Lower variance in the residuals at 2×10^7 reflects closer alignment of predictions with actual values, enhancing confidence in the model's output for high Re scenarios.

Table 7.9: Regression model performance metrics.

Re	R^2	σ [%]	$R^2_{C_\mathrm{l}}$	$R^2_{C_\mathrm{d}}$	$R^2_{C_\mathrm{m}}$
1×10^6	0.987	10.520	0.988	0.989	0.985
5×10^6	0.988	9.430	0.993	0.984	0.987
2×10^7	0.989	7.677	0.995	0.983	0.990

8

LEI Kite Profile Design Recommendations

Given the diverse applications of LEI kites, ranging from boat towing to airborne wind energy generation and kitesurfing, profile design must be tailored to the specific operational context. A key consideration in all cases is the variation in Reynolds number across different flight phases, as aerodynamic performance shifts across the range. This sensitivity is illustrated in Figures 7.1, 7.2, 7.3, and 7.4.

Among the aerodynamic coefficients, C_m and the associated CoP position are of particular importance. In large-scale systems such as airborne wind energy platforms and boat-towed kites, a strong negative C_m, corresponding to a large nose-down pitching moment, poses a serious risk of front stall at the edge of the wind window. This failure mode must be avoided at all costs. It can be mitigated by shifting the CoP forward through adjustments to the profile parameters. More positive pitching moments also reduce the force required on the rear bridle lines due to a longer moment arm, thereby improving both mechanical efficiency and longitudinal stability. These forces are visualised in Figure 8.1, where the red arrow indicates the rear bridle force, the green arrow the front bridle force, and the blue arrow the lift force, located at the CoP. The CoP shifts with AoA, as illustrated by the dashed blue line.

It is important to note that a forward-shifted CoP position often comes at the cost of reduced aerodynamic efficiency, particularly in terms of C_l and C_d. Optimal aerodynamic performance typically requires larger values of η and λ without reflex deflection, which tend to shift the CoP rearward as discussed in Chapter 6 and demonstrated by Watchorn in the context of η variation [50].

An example of the wide CoP variation with α, ranging from $x/c = 0.28$ to 0.45, is shown in Figure 8.1 for a representative LEI kite profile with the following parameters: $t = 0.08$, $\eta = 0.2$, $\kappa = 0.08$, $\delta = 0°$, $\lambda = 0.2$, and $\phi = 0.65$. The corresponding CoP location across the AoA range is plotted in Figure 6.19, illustrating the relatively rearward position of x_cp in LEI kite profiles. At $\alpha = 10°$, for instance, the CoP is located at $x/c = 0.3$, which is significantly aft compared to conventional NACA airfoils, where the CoP typically aligns near the quarter chord [2].

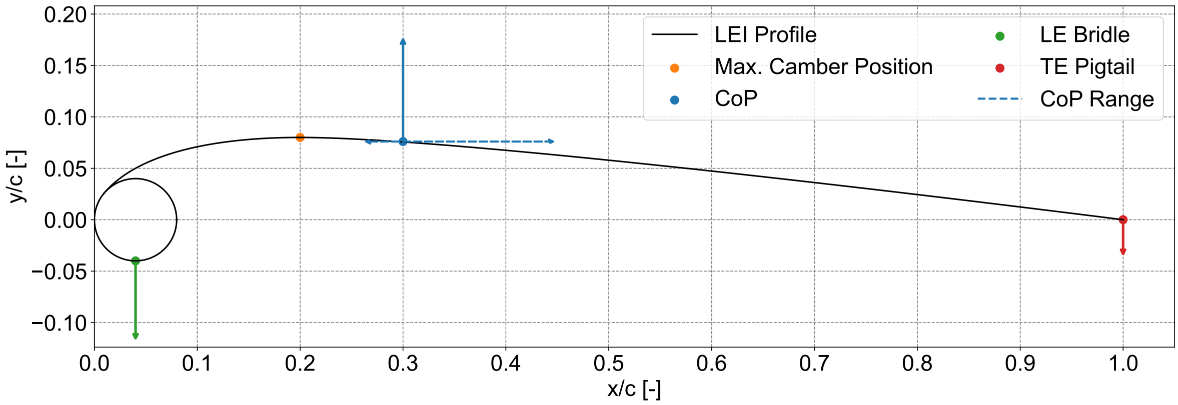

Figure 8.1: 2D free-body diagram of a LEI kite showing rear (red), front (green), and lift (blue, at CoP) forces.

The recommended procedure for implementing the regression model[1] is as follows. First, define a grid of profile parameters for which the prediction model should compute the aerodynamic coefficients. Next, specify acceptable bounds for C_m within a given range of α. Depending on the performance objectives, such as maximising lift, optimising the lift-to-drag ratio, or a weighted combination, a suitable profile can be identified. Alternatively, performance weights can be applied across a range of AoA values to prioritise specific flight regimes.

After this selection, the results can be postprocessed by examining the C_p distribution predicted via k-nearest neighbours from the reference dataset. Once an approximate optimum is identified, a finer parameter grid analysis can be conducted using steady RANS simulations to achieve higher accuracy and perform final fine tuning. Finally, the optimal profile contour coordinates can be directly exported from the parametric model.

[1]The source code for the CFD simulations, figure generation, and prediction models is available on GitHub: `https://github.com/awegroup/Pointwise-Openfoam-toolchain`. The CFD simulation dataset used to train the ML model is archived on Zenodo: `https://doi.org/10.5281/zenodo.15593140`.

9 Conclusion & Recommendations

The parametric model improvements began with the implementation of a dynamic seam angle locator, which exhibited a smooth and robust dependency on other geometric parameters. A similar approach was applied to the lower surface fillet at the LE seam, whose size was thoroughly assessed through sensitivity analysis across various configurations. Although smooth meshes could be achieved with small LE fillets, experimental validation is recommended to assess the aerodynamic influence of the fillet with a real-life configuration without a fillet.

Further profile smoothness was achieved by replacing the single cubic Bézier curve used by Watchorn [50] with two connected cubic Bézier curve segments. This increased the number of control points and enabled greater flexibility and control through the introduction of additional shape parameters: reflex angle δ, camber tension λ, and LE curvature ϕ.

A mesh sensitivity study revealed slight variations in aerodynamic performance and convergence iterations due to canopy thickness. Given its influence on the results and to more accurately represent the actual canopy thickness, the non-dimensional thickness was reduced relative to Watchorn's, reaching a value of $0.0005c$. Mesh resolution in the normal and tangential directions was found to be sufficiently converged. Therefore, the same grid density as used by Watchorn [50] (201×575) was maintained to balance accuracy and computational cost.

The decision to omit a transition model due to the presence of overlapping fabric and stitching seams was verified numerically, confirming that bypass transition is triggered by the roughness elements. The overlap acts as a forward-facing step, while the stitching behaves like a three-dimensional roughness element. Given the elevated turbulence intensity $(\approx 8\%)$, the estimated critical roughness Reynolds number ($\mathrm{Re}_{\mathrm{h,\,crit}} = 200$), and trip height ($h = 0.6\ \mathrm{mm}$), transition is predicted for flows above approximately Re $> 1 \times 10^6$. However, since this approach is based on assumptions, experimental validation is recommended, specifically, wind tunnel testing across a range of Reynolds numbers and seam configurations. Techniques such as fluorescent oil film and infrared thermography could be considered for boundary layer visualisation.

Numerical results also suggest a bypass transition at the closing seam on the lower side of the LE tube, which behaves similarly to a trip wire. The larger seam height on the lower surface, measuring $6\ \mathrm{mm}$, induces transition at even lower Reynolds numbers than the upper seam. Experimental validation is again recommended across different seam placements, angles of attack, and Reynolds numbers. Finally, it remains uncertain whether a fixed transition location should be enforced in simulations or if the short upstream laminar region can be neglected under the assumption that a fully turbulent boundary layer from the LE suffices for aerodynamic predictions.

The latest LEI kite profile parametric model by Watchorn [50] provided a solid foundation but lacked key geometric parameters, limiting aerodynamic optimisation. One critical addition is the reflex angle δ, which increases C_m and shifts the CoP forward, an important feature for improving static longitudinal stability and reducing rear bridle forces. Further refinement was achieved by enhancing control over the

Bézier curves that define the front and rear canopy splines. Leading to a camber tension parameter, λ, adjusting the Bézier control point at the location of maximum camber near the TE. This enables improvements in lift coefficient C_l, with a trade-off in slightly lower C_m values, while maintaining a stable drag coefficient C_d. Additional LE curvature control using the front spline control points allowed for a smoother geometric transition at the LE seam. Although this modification had a negligible effect on aerodynamic performance.

Building upon the polynomial regression model by Watchorn [50], the present approach introduces significant improvements in both model performance and reliability. The dataset was expanded from 256 to 126,500 CFD simulations[1], covering a broader range of Reynolds numbers, Re $= 1\times10^6$, 5×10^6, and 2×10^7. This led to an overall regression accuracy improvement, with the mean R^2 increasing from 0.945 in Watchorn's results to 0.988 using the extra trees ML model. Although RMSE is sensitive to the scale of the output variable, it remains a useful metric for comparison. In this case, the RMSE values reported by Watchorn were approximately 50% higher for both C_l and C_d compared to those obtained with the extra trees model, indicating a noticeable improvement in accuracy, while the RMSE for C_m was identical.

Model confidence was verified using randomised test configurations comprising 10.6% of the total dataset. The resulting relative error σ was 10.5% for Re $= 1\times10^6$, 9.4% for 5×10^6, and as low as 7.68% for 2×10^7, indicating greater reliability at higher Reynolds numbers likely due to delayed stall. In contrast, Watchorn's results lack reported variance and confidence intervals, making it difficult to assess the statistical reliability of the findings. Furthermore, a k-nearest neighbours (kNN) model was employed for C_p and C_f predictions, achieving good agreement with the randomised test data.

Despite these advancements, several areas for further improvement remain. Outlier detection and initial profile filtering could be enhanced, for instance, by excluding cases with low η where $\kappa > 2.5\times t$. Additionally, the AoA distribution, though effective for C_l, showed limitations for C_d and C_m. A more uniformly and densely distributed AoA range could help reduce the variance and improve regression accuracy across all coefficients.

Future Perspective

The developed aerodynamic regression model offers a valuable opportunity to enhance existing FSI frameworks. For structural modelling, incorporating more particle points along the struts and replacing the current fixed load split with a dynamically derived distribution from the C_cp k-nearest neighbours model or from x_cp obtained with C_l and C_m could significantly improve accuracy. Increasing spanwise resolution guided by predicted C_p values would also better align with the VSM model and reduce coupling errors.

On the aerodynamic side, replacing traditional polar lookup tables or less accurate polynomial regression models with an ML regression model trained across a wide range of angles of attack, Reynolds numbers, and profile parameters would enable more accurate force predictions per segment. This approach would align aerodynamic force estimation more closely with 3D CFD results and improve adaptability to various kite configurations.

Overall, these improvements point to a robust path forward for applying the regression model in more advanced, data-driven FSI simulations.

[1]The source code for the CFD simulations, figure generation, and prediction models is available on GitHub: `https://github.com/awegroup/Pointwise-Openfoam-toolchain`. The CFD simulation dataset used to train the ML model is archived on Zenodo: `https://doi.org/10.5281/zenodo.15593140`.

References

[1] D. Aberdeen. *SurfPlan: Kite and Wing Design Software*. Professional CAD software for designing inflatable and ram-air kites and wings, featuring one-click pattern generation, bridling tools, and structural control. 2024. URL: http://www.surfplan.com.au/sp/ (visited on 06/16/2025).

[2] J.D. Anderson. *Fundamentals of Aerodynamics*. 6th ed. New York: McGraw-Hill Education, 2017. ISBN: 978-1-259-25134-4.

[3] P. Balakumar. “Laminar to turbulence transition in boundary layers due to tripping devices”. In: *AIAA Scitech 2021 Forum*. 2021. DOI: https://doi.org/10.2514/6.2021-1948.

[4] P. Bechtle et al. “Airborne wind energy resource analysis”. In: *Renewable Energy* 141 (Oct. 2019). research about the gain at heigher altitudes, pp. 1103–1116. ISSN: 18790682. DOI: 10.1016/j.renene.2019.03.118.

[5] J. Berens. “Dynamic Nonlinear Aeroelastic Behaviour of Flexible Wings in an Airflow”. retrieved from John- reworked breukels model. MA thesis. Tu Delft, Mar. 2015.

[6] R.G. den Boer. *Low speed aerodynamic characteristics of a two-dimensional sail wing with adjustable slack of the sail*. Tech. rep. Technische Hogeschool Delft, Oct. 1980. URL: https://repository.tudelft.nl/record/uuid:18ae2cc6-434e-49c8-9296-d3fa450850a5.

[7] A. Bosch et al. “Dynamic Nonlinear Aeroelastic Model of a Kite for Power Generation”. In: *Journal of Guidance, Control, and Dynamics* 37.5 (2014), pp. 1426–1436. ISSN: 1533-3884. DOI: 10.2514/1.G000545.

[8] A.L. Braslow et al. *Simplified method for determining critical height of distributed roughness particles for boundary-layer transition*. Tech. rep. TN 4363. NACA, 1958. URL: https://ntrs.nasa.gov/citations/19930085292.

[9] J. Breukels. “An engineering methodology for kite design.” Breukel PHD work, Johns’s predecessor. PhD thesis. Tu Delft, Jan. 2011. ISBN: 9789088912306.

[10] L.S. Caretto et al. “Two Calculation Procedures for Steady, Three-Dimensional Flows With Recirculation”. In: *Numerical Methods in Fluid Mechanics*. Ed. by J. N. Reddy. Vol. 19. Lecture Notes in Physics. Springer, 1972, pp. 60–68.

[11] O. Cayon. “Fast aeroelastic model of a leading-edge inflatable kite”. About VSM model retrieved from John- incorporating max camber location. MA thesis. Tu Delft, DTU, July 2022. URL: https://repository.tudelft.nl/.

[12] O. Cayon et al. “Fast Aero-Structural Model of a Leading-Edge Inflatable Kite”. In: *Energies* 16 (7 Apr. 2023). Latest VSM method paper. ISSN: 19961073. DOI: 10.3390/en16073061.

[13] O. Cayon et al. “Kite as a Sensor: Wind and State Estimation in Tethered Flying Systems”. In: *Wind Energy Science Discussions* (2025). in review. DOI: 10.5194/wes-2024-182. preprint: wes-2024-182. URL: https://wes.copernicus.org/preprints/wes-2024-182/.

[14] A. Cherubini et al. “Airborne Wind Energy Systems: A review of the technologies”. In: *Renewable and Sustainable Energy Reviews* 51 (July 2015). different AWE systems, pp. 1461–1476. ISSN: 13640321. DOI: 10.1016/j.rser.2015.07.053. URL: http://linkinghub.elsevier.com/retrieve/pii/S1364032115007005.

[15] T. Corentin. *Kite Design, Methods and Digital Simulations End-of-study internship report Fluid Dynamics-Energy Transport and Environment*. Tech. rep. Corentin ref about 2D kite profiles from BTS comparison Xfoil/gridpro/GMSH. Institut Supérieur de l’Aéronautique et de l’Espace (Accessible on request: Beyond-The-Sea.com), Dec. 2023.

[16] M. Deaves. “An Investigation of the Non-Linear 3D Flow Effects Relevant for Leading Edge Inflatable Kites”. consulted from John as source for the canopy thickness t- sensativity analysis on the LE filling- also uses the RANS SST model- does sensitivity mesh analysis on naca 0012 not on LEI. MA thesis. Tu Delft, Sept. 2015.

[17] F.M. Dekking et al. *A Modern Introduction to Probability and Statistics: Understanding Why and How*. 1st ed. London: Springer, 2005. ISBN: 9781852338961. DOI: 10.1007/1-84628-168-7.

[18] P. Demkowicz. “Numerical analysis of the flow past a leading edge inflatable kite wing using a correlation-based transition model”. 3D and 2d flow- hybrid mesh- good explanation of mesh quality, modeling and pre simulation. MA thesis. Tu Delft, Dec. 2019. URL: http://repository.tudelft.nl/..

[19] M. Drela. “XFOIL: An Analysis and Design System for Low Reynolds Number Airfoils”. In: *Conference on Low Reynolds Number Airfoil Aerodynamics*. University of Notre Dame, June 1989.

[20] G.E. Elsinga et al. “Tomographic-PIV measurement of the flow around a zigzag boundary layer trip”. In: *Experiments in Fluids* 52.4 (2012), pp. 865–876. DOI: https://doi.org/10.1007/s00348-011-1153-8.

[21] P. Geurts D. Ernst et al. “Extremely Randomized Trees”. In: *Machine Learning* 63.1 (Apr. 2006), pp. 3–42. DOI: 10.1007/s10994-006-6226-1. URL: https://doi.org/10.1007/s10994-006-6226-1.

[22] J.H. Ferziger et al. *Computational Methods for Fluid Dynamics*. 4th ed. Cham, Switzerland: Springer Nature Switzerland AG, 2020. ISBN: 978-3-319-99691-2. DOI: 10.1007/978-3-319-99693-6. URL: https://doi.org/10.1007/978-3-319-99693-6.

[23] M. Folkersma et al. “Boundary layer transition modeling on leading edge inflatable kite airfoils”. In: *Wiley* 7 (Feb. 2019), pp. 908–921. DOI: 10.1002/we.2329. URL: https://doi.org/10.1002/we.2329.

[24] M. Folkersma et al. “Steady-state aeroelasticity of a ram-air wing for airborne wind energy applications”. In: *Journal of Physics: Conference Series* 1618 (3 Sept. 2020). ramair 3d cfd and fsi. ISSN: 17426596. DOI: 10.1088/1742-6596/1618/3/032018.

[25] C. Geuzaine et al. “Gmsh: A three-dimensional finite element mesh generator with built-in pre- and post-processing facilities”. In: *International Journal for Numerical Methods in Engineering* 79.11 (2009), pp. 1309–1331. DOI: 10.1002/nme.2579.

[26] F. Hirschmann. *rdp: Ramer–Douglas–Peucker algorithm implementation*. Version 0.8. [Computer software]. 2016. URL: https://pypi.org/project/rdp.

[27] R.H.M Van Kappel. “Aerodynamic analysis tool for dynamic leading edge inflated kite models: a nonlinear vortex lattice method”. Retreived from John- reworked the model of Breukel. MA thesis. Tu Delft, Oct. 2012.

[28] Kitware Inc. *ParaView*. https://www.paraview.org. Version 5.13. 2024.

[29] C.L. Ladson et al. *Development of a Computer Program to Obtain Ordinates for NACA 4-Digit, 4-Digit Modified, 5-Digit, and 16-Series Airfoils*. Technical Memorandum NASA TM X-3284. NASA Technical Memorandum L-10375. Hampton, VA: NASA Langley Research Center, 1975.

[30] R.B. Langtry et al. “Correlation-Based Transition Modeling for Unstructured Parallelized Computational Fluid Dynamics Codes”. In: *AIAA Journal* 47.12 (2009), pp. 2894–2906. ISSN: 0001-1452. DOI: 10.2514/1.42362.

[31] B.E. Launder et al. “The numerical computation of turbulent flows”. In: *Computer Methods in Applied Mechanics and Engineering* 3 (2 1974), pp. 269–289.

[32] D.T. Lee et al. “Two Algorithms for Constructing a Delaunay Triangulation”. In: *International Journal of Computer and Information Sciences* 9.3 (1980), pp. 219–242. ISSN: 0091-7036. DOI: 10.1007/BF00977785.

[33] R. Leloup et al. “Modelling approach and numerical tool developments for kite performance assessment and mechanical design; application to vessels auxiliary propulsion”. PHD report referred by Corentin, about sensitivity analysis of the pressure side edge fillet. PhD thesis. University of Western Brittany, Oct. 2014.

[34] F.R. Menter. “Two-equation eddy-viscosity turbulence models for engineering applications”. In: *AIAA Journal* 32.8 (1994), pp. 1598–1605. DOI: 10.2514/3.12149.

[35] S.R. Pandala. *lazypredict: Lazy Machine Learning*. Version 0.2.12. GitHub repository. 2021. URL: https://github.com/shankarpandala/lazypredict.

[36] F. Pedregosa et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.

[37] Pointwise Inc. *Pointwise: Computational Fluid Dynamics (CFD) Meshing Software*. https://www.pointwise.com. Accessed: 2025-06-02. 2023.

[38] J.A.W. Poland et al. “Modelling Aero-Structural Deformation of Flexible Membrane Kites”. In: *Energies* 16 (14 July 2023). Particle spring model latest version. ISSN: 19961073. DOI: 10.3390/en16145264.

[39] R.P.J.O.M. van Rooij et al. “Roughness sensitivity considerations for thick rotor blade airfoils”. In: *Journal of Solar Energy Engineering* 125.4 (2003), pp. 468–478. DOI: https://doi.org/10.2514/6.2003-350.

[40] S. Sachdeva. “Impact of Turning-Induced Shape Deformations on Aerodynamic Performance of Leading Edge Inflatable Kites”. MA thesis. TU Delft, 2017.

[41] R. Schmehl, ed. *Airborne Wind Energy. Advances in Technology Development and Research*. 1st ed. Green Energy and Technology. Singapore: Springer Singapore, 2018, pp. XXVII, 752. ISBN: 978-981-10-1947-0. DOI: 10.1007/978-981-10-1947-0. URL: https://doi.org/10.1007/978-981-10-1947-0.

[42] R. Smith et al. “Computation of aerodynamic coefficients for a flexible membrane airfoil in turbulent flow: A comparison with classical theory”. In: *Physics of Fluids* 8.12 (1996), pp. 3346–3353.

[43] R. Smith et al. “Computation of aerodynamic coefficients for a flexible two-dimensional membrane wing in turbulent flow”. In: *ASME-PUBLICATIONS-FED* 234 (1995), pp. 3–10.

[44] J. Steiner. “Towards the CSM-CFD modelling of membrane wings at high Reynolds numbers”. retreived from Demkowicz tabel of previous work, 2D structured mesh, but also lot of FSI. MA thesis. Tu Delft, Jan. 2018.

[45] A. Teimourian et al. “Airfoil aerodynamic performance prediction using machine learning and surrogate modeling”. In: *Heliyon* 10 (8 Apr. 2024), p. 10. ISSN: 24058440. DOI: 10.1016/j.heliyon.2024.e29377.

[46] A. Viré et al. “Direct Numerical Simulations of Flow Past a Leading-Edge Inflatable Wing”. In: *7th International Airborne Wind Energy Conference (AWEC 2017)*. Conference abstract. Freiburg, Germany, 2017, p. 143. URL: https://www.awec2017.com.

[47] A. Viré et al. “Effect of Chordwise Struts and Misaligned Flow on the Aerodynamic Performance of a Leading-Edge Inflatable Wing”. In: *Energies* 15 (4 Feb. 2022). 3D hybrid meshReff from John, consulted for the decision of transition model integration. ISSN: 19961073. DOI: 10.3390/en15041450.

[48] A. Viré et al. “Reynolds-averaged Navier-Stokes simulations of the flow past a leading edge inflatable wing for airborne wind energy applications”. In: *Journal of Physics: Conference Series* 1618 (3 Sept. 2020). for introduction text, Folkersma 3D model. ISSN: 17426596. DOI: 10.1088/1742-6596/1618/3/032007.

[49] P. Virtanen et al. *SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python*. 2020. DOI: 10.1038/s41592-019-0686-2. URL: https://scipy.org/.

[50] J.P. Watchorn. “Aerodynamic Load Modelling for Leading Edge Inflatable Kites”. Delft University of Technology, Aug. 2023. URL: https://repository.tudelft.nl/record/uuid:42f611a2-ef79-4540-a43c-0ea827700388.

[51] H.G. Weller et al. “A Tensorial Approach to Computational Continuum Mechanics Using Object-Oriented Techniques”. In: *Computers in Physics* 12.6 (Nov. 1998), pp. 620–631.

[52] D.C. Wilcox. “Reassessment of the scale-determining equation for advanced turbulence models”. In: *AIAA Journal* 26.11 (1988), pp. 1299–1310.

[53] H. Zhu et al. “Forward-Facing Steps Induced Transition in a Subsonic Boundary Layer”. In: *EU-CASS* (2015). study about forward facing steps. DOI: 10.13009/EUCASS2017-682.