Thesis for The Degree of Licentiate of Engineering

Aerodynamic Performance and Layout Optimisation of Symmetrically Cambered Wing Sails for Wind-Assisted Ship Propulsion

Stephan van Reen

Department of Mechanics and Maritime Sciences
Chalmers University of Technology
Gothenburg, Sweden, 2026

Aerodynamic Performance and Layout Optimisation of Symmetrically Cambered Wing Sails for Wind-Assisted Ship Propulsion

Stephan van Reen

Department of Mechanics and Maritime Sciences
Division of Marine Technology
Chalmers University of Technology SE-412 96 Göteborg,
Sweden
Phone: +46(0)31 772 1000

Printed by Chalmers Digitaltryck,
Gothenburg, Sweden 2026.

Aerodynamic Performance and Layout Optimisation of Symmetrically Cambered Wing Sails for Wind-Assisted Ship Propulsion

STEPHAN VAN REEN
Chalmers University of Technology
Department of Mechanics and Maritime Sciences
Division of Marine Technology

Abstract

A growing need for sustainable maritime transport has driven the development of rigid wing sails. Symmetrically cambered profiles offer high thrust potential across many apparent wind angles, but their aerodynamic design and installation layout can still be further optimised. This licentiate thesis addresses these challenges through two complementary studies: the development of an aerofoil optimisation framework and an exploratory analysis of multi-sail layouts.

The first study develops a surrogate-based Bayesian optimisation method to reduce reliance on costly CFD simulations. A Gaussian Process surrogate model uses prediction uncertainty to guide the search for maximum average thrust over apparent wind angles from 10° to 150°. A hybrid parametrisation enables flexible aerofoil geometry variation. The optimised configuration has a more uniform thickness distribution and achieves an 8% thrust increase over the D2R10 benchmark, validated by high-fidelity IDDES.

The second study investigates the aerodynamic interference in multi-sail layouts using a two-dimensional inviscid method. Two configurations are considered: a triple in-line and a quad-sail parallel layout. Under a fixed total spacing, optimisation yields only limited performance improvements. The in-line layout experiences thrust losses of up to 6% compared to isolated sails. The parallel layout exhibits larger reductions ranging from 10% to 28%. These findings indicate that aerodynamic interference significantly affects multi-sail performance.

This thesis develops a Bayesian optimisation framework for multi-fidelity wing sail design. It shows how using a fast lower-fidelity solver can deliver measurable 3D aerofoil performance gains. The new parametrisation enables broad shape variation, revealing how geometric characteristics affect thrust. The framework also identifies trends and limitations in different layout configurations. These results provide insights to guide future wing sail design strategies.

Keywords

Aerodynamics, installation layout, machine learning, multi-point optimisation, rigid wing sail, wind-assisted ship propulsion

Preface

This thesis presents the research conducted at the Division of Marine Technology at Chalmers University of Technology from November 2023 to the end of 2025, as part of the GEMINI project (GEneric Multidisciplinary optimisation for sail INstallation on wInd-assisted ships). The project focuses on optimising wing sail concepts for wind-assisted ship propulsion, with emphasis on aerodynamics and fluid-structure interaction across varying wind conditions. The research is in collaboration with RISE and funded by the Swedish Transport Administration.

I would like to express my deepest gratitude to my principal supervisor, Professor Hua-Dong Yao, for his exceptional guidance and support throughout this research. His mentorship has been invaluable, and I could not have done this without his expertise and encouragement.

I also wish to thank my co-supervisor, Dr Laura Marimon Giovannetti, and my examiner, Professor Jonas Ringsberg, for their valuable input. Dr Marimon Giovannetti's perspectives helped broaden my approach, while Professor Ringsberg's guidance on professional development and the academic aspects of the thesis ensured its rigour and quality.

A special thank you to my girlfriend, Smita, for her unwavering support and optimism. Your encouragement helped me stay focused and confident throughout the journey, and I look forward to a well-deserved break together. Thanks to my colleagues and friends at Chalmers for their companionship and support during this time.

Finally, I would like to thank my family and friends back in Norway, Switzerland, and, of course, the Netherlands for supporting me over all these years. Special thanks go to little Lucas, whose big smile brightens every day.

Contents

List of Publications

Appended publications

This thesis is based on the following publications:

[**Paper I**] **S. van Reen**, B. Serbülent, H.-D. Yao, *Machine learning-based multipoint optimisation for improving aerodynamics of symmetrically cambered wing sails in wind-assisted ship propulsion.*
Ocean Engineering 39 (December 2025), 133-144.
10.1016/j.oceaneng.2025.122829

[**Paper II**] **S. van Reen**, J. Lin, J. Niu, P. Sharpe, X. Li, H.-D. Yao, *Reducing Aerodynamic Interference Through Layout Optimization of Symmetrically Cambered Wingsails: A Comparative Study of In-Line and Parallel Configurations.*
Journal of Marine Science and Engineering 13(10) (October 2025), 1998.
10.3390/jmse13101998

Other publications

The following publication was published during my PhD study. However, it is not appended to this thesis, due to the contents overlapping those of the appended publications.

[**a**] **S. van Reen**, H. Zhu, J. Lin, J. Niu, P. Sharpe, H.-D. Yao, *Aerodynamic Optimization of In-Line and Parallel Layouts for Symmetric Cambered Wingsail Installation.*
Proceedings of the ASME 44th International Conference on Ocean, Offshore and Arctic Engineering (OMAE2025), Vancouver, Canada, June 22–27, 2025, Paper OMAE2025-155321. 10.1115/OMAE2025-155321

Nomenclature

Latin notations

C_d	Drag coefficient [-]
C_l	Lift coefficient [-]
C_t	Thrust coefficient [-]
C_p	Pressure coefficient [-]
$\overline{C}_t$	Average thrust coefficient [-]
$C_{t,\text{mult}}/C_{T,\text{sum}}$	Multiple sail thrust coefficient [-]
D	Drag force [N]
F_{res}	Resultant aerodynamic force [N]
I	Leading/trailing edge radius index [-]
L	Lift force [N]
M	Number of sails [-]
N	Number of wind angles [-]
N_{ribs}	Number of ribs [-]
Q	Number of control points [-]
T	Thrust force [N]
V_{AW}	Apparent wind vector [m/s]
V_∞	Air speed [m/s]
a	Coefficient [-]
b	Sail height [m]
c	Chord length [m]
h_c	Maximum camber [-]
h_t	Maximum thickness [-]
$k(\boldsymbol{\phi}, \boldsymbol{\phi}')$	Covariance function [-]
l	Length scale [-]
l_{ribs}	Rib spacing [m]
$l_{1/2}$	Sail unit distance parameter 1 [-]

k	Exploration–exploitation trade-off parameter [-]
t	Thickness [m]
x	Chord-wise coordinate [-]
y_c	Local camber height [-]
y_t	Local thickness [-]
$\mathcal{GP}_0$	Initial Gaussian Process surrogate model [-]

Greek notations

α	Angle of attack [°]
θ_{AW}	Apparent wind angle [°]
λ	Taper ratio [-]
μ_*	Mean [-]
ν	Smoothness parameter [-]
ρ	Air density [kg/m^3]
σ^2	Variance [-]
Φ	Input space [-]
ϕ/ϕ'	Input value [-]
ψ	Objective value [-]

Abbreviations

AOA	Angle of attack
AWA	Apparent wind angle
IDDES	Improved Delayed Detached Eddy Simulation
QP	Quad parallel layout
TL	Triple in-line layout
TWA	True wind angle

Chapter 1

Introduction

1.1 Background

In an effort to reduce greenhouse gas emissions, the modern shipping industry is exploring renewable technologies to replace or supplement fossil fuel-based propulsion systems (Zhu et al. 2023a). In this search, the old technology of wind propulsion is making its return. One of the most promising solutions is rigid wing sails supported by modern materials and control systems. They can be imagined as wings mounted vertically on a ship that produce thrust through aerodynamic lift and drag. This leads to a potentially substantial reduction in engine power demand and fuel consumption.

However, despite their conceptual similarity, wing sails do not operate like aircraft wings. The latter are designed under the assumption that the incoming flow is roughly parallel to the aircraft's movement, as the aircraft's forward speed dominates the local wind speed. This means that the generated lift force, which by definition is perpendicular to the flow direction, will provide an upward force on the aircraft. Contrarily, wing sails must operate in varying wind directions, and cannot guarantee that the lift vector aligns with the ship's movement direction. Instead, wing sails must be able to rotate to adapt to varying wind directions and harness lift and drag together to produce a thrust vector in the ship's movement direction. This puts different requirements on the aerodynamic shape of a wing sail and the arrangement of multiple wing sails on a ship.

To date, various industry groups have been developing prototypes and products with a range of cross-sectional aerofoil shapes. These include a two-element flapped aerofoil type by Oceanbird[1], a three-element type by Bar Technologies[2], and a symmetrically cambered crescent-shaped type by ScandiNAOS[3]. Much of the research in this field has focused on the aerodynamic performance of an individual wing sail and multiple sails working in tandem. However, there is relatively little research on how to optimise geometries

---

[1] Wallenius Marine AB 2025.

[2] BAR Technologies Ltd. 2025.

[3] ScandiNAOS AB 2025.

and layout arrangements to improve the aerodynamics. Furthermore, it is interesting to investigate hybrid optimisation involving both aerodynamics and structure.

1.2 Literature Study

The 2010 America's Cup marked a turning point in sailing technology with the introduction of rigid wing sails, which demonstrated superior aerodynamic performance and renewed interest in wing sail research (Blakeley et al. 2012). These sails typically featured a two-element configuration, with a larger wing and an attached flap. Early investigations into such sails were largely proprietary, limiting access to experimental data.

One of the first publicly available studies was by Blakeley et al. (2012), who conducted wind-tunnel tests on a two-dimensional, two-element wing sail model. They observed that the gap between the main element and the flap had a strong influence on the stall angle. Building on this work, Furukawa et al. (2015) found that, while the gap size and pivot location had relatively minor effects, the deflection of the flap was a dominant factor in determining lift and drag. They also reported that at high camber and angle of attack (AOA), low Reynolds numbers (below 6×10^5) could trigger early flow separation and reduce performance. Li et al. (2020) extended the analysis using both 2D and 3D CFD simulations. Their results indicated that at high flap deflections, moving the flap rotation axis aft increases the lift coefficient, particularly when the axis is positioned between 85% and 95% of the chord. More recently, Hillenbrand et al. (2024) investigated stall behaviour through wind tunnel experiments and showed that separation occurs in two stages: starting on the flap, then followed by the entire wing, with the latter exhibiting hysteresis effects.

1.2.1 Symmetrically Cambered Wing Sail

As an alternative to conventional aerofoils, the symmetrically cambered wing sail concept is receiving increasing attention. One of the earlier studies on this concept was by Ouchi et al. (2016), who explored a ship with nine symmetrically cambered wings whose height could be telescopically adjusted. This concept has since developed into a four-wing sail configuration known as the "Wind Challenger", which is estimated to achieve a 30% cut in energy consumption on a trip between Yokohama and Seattle (Ouchi et al. 2016). A physical prototype of the wind challenger sail has been in operation since 2022, having been fitted to the forecastle of a bulk carrier called the Shofu Maru (Ouchi et al. 2023). According to the manufacturer[4], the prototype saves an average of 5% to 8% in fuel consumption per voyage.

Furthermore, Ma et al. (2019) introduced a coupled wing sail design that integrates an arc-shaped wing with a NACA 0018 aerofoil to enhance aerodynamic performance. The geometry was described using 14 parameters based on a modified PARSEC method. Particle Swarm Optimisation (PSO) was

---

[4] Mitsui O.S.K. Lines, Ltd. 2025.

employed to maximise the power factor, defined as $C_L^{1.5}/C_D$, averaged over AOAs ranging from -15° to 40°. The optimal configuration achieved a 30% improvement in average power factor over the baseline arc-shaped sail.

The Swedish company ScandiNAOS developed a new type of symmetrically cambered wing sail, often referred to as the "crescent-shaped" wing sail. Zhu et al. (2023a) conducted a parametric CFD study on this new design. This study employed the unsteady Reynolds-Averaged Navier–Stokes (RANS) equations to examine different geometric configurations, resulting in a shape called the "D2R10" that generated the highest lift. The configuration was then further analysed with Improved Delayed Detached Eddy Simulation (IDDES) (Zhu et al. 2024) in order to confirm the improved performance. The D2R10 would serve as the benchmark configuration to be improved with the optimisation method developed by van Reen et al. (2025), which is included in this thesis as Paper I.

A separate study by Guzelbulut et al. (2024) also carried out a parametric shape optimisation of the D2R10 profile. Three design variables were proposed to be varied across three levels to either enhance thrust or reduce propeller power. Rather than evaluating all 27 possible combinations generated from the profile parametrisation, they selected nine representative samples for CFD analysis and examined them at a range of wind directions and AOAs. For each configuration, the most favourable angles were averaged to characterise the aerodynamic behaviour. These results were then used to construct a second-order polynomial surrogate model. The thrust-focused design yielded a 12.3% improvement, whereas the power-focused variant achieved a 22% reduction in required propeller power relative to the original D2R10.

1.2.2 Multiple Wing Sails

Real-world ships will not always have just one wing sail mounted on them, but will likely operate with multiple sails. These large structures affect the airflow around them, and thus, to get a good understanding of wing sail performance in the field, it is crucial to understand how they interfere with each other. Research in this area dates back more than a decade, but it is still relatively sparse. It had seen an earlier introduction of optimisation models. Jo et al. (2013) harnessed a genetic algorithm with a Kriging surrogate model to minimise interference between three morphing flapped wing sails in an in-line layout. They varied the flap deflection angle, flap length and AOA, and found that optimising these parameters could yield a 10% thrust increase.

Using steady RANS simulations, Hussain et al. (2021) evaluated the thrust performance of a bulk carrier fitted with five rigid wing sails. The sails were designed with a NACA 4412 profile, optimised for wind coming from a single board side. The simulations predicted a reduction in required brake power of approximately 30% under zero-drift conditions and 20% at larger drift angles compared to a ship without wing sails. Later, Giovannetti et al. (2022) experimentally examined how three rigid wing sails interact under turbulent flow to optimise thrust for wind-powered ships. Their results showed that adjusting individual sheeting angles and wing placement significantly improves efficiency.

Makram et al. (2023) explored the design and aerodynamic optimisation of flapped wing sails for a Suezmax-class tanker. The first step involved a comparative analysis of 13 symmetrical aerofoil profiles using the XFOIL tool (Drela 1989). A single profile was selected for its favourable aerodynamic characteristics, namely, strong lift, low drag, small pitching moment, and moderate stall characteristics from medium to high Reynolds numbers. This selected profile was subsequently employed in the development of nine three-dimensional wing sail arrangements, each differing in aspect ratio, taper ratio, and the number of sail pairs. Notably, the research underscored the importance of accounting for aerodynamic interference between sails, as omitting these effects was shown to significantly over-predict the thrust and lateral force.

In the same year, wind tunnel tests for a triple in-line wing sail layout were performed by Zhu et al. (2023b). The sails had the D2R10 profile and were tested under varying wind speeds, apparent wind angles (AWA), and AOAs. The results showed that the front sail produced the greatest thrust, the middle sail exhibited the strongest structural vibrations, and the aft sail produced the smallest thrust. Moreover, Malmek (2023) developed a lifting line model for multiple wing sails, which was then further improved by Malmek et al. (2024) by adding a potential flow-based interaction method, as well as an optional boundary layer correction. This allows for better 3D aerodynamic force prediction in upwind conditions without sacrificing computational efficiency.

That same year Yasuda et al. (2024) investigated a two-sail configuration of symmetrically cambered wing sails with CFD, looking to understand how the two sails affect each other at large AWAs of 150° and 180° (tailwind). They sought to optimise each sail's AOA to achieve maximum thrust at the two wind angles, and concluded that they require opposite AOA combinations: fore small and aft large at 150° and fore large and aft small at 180°.

1.3 Goals and Objectives

The overarching aim of this research is to develop a general and robust optimisation framework for the design of individual wing sails and multi-sail arrangements, and to ultimately incorporate structural responses into the framework. This licentiate thesis represents the first stage of the work. It establishes the overall framework, demonstrates its feasibility and provides initial insights through systematic studies of symmetrically cambered aerofoils and multi-sail layouts. The preliminary steps based on aerodynamic analysis are intended to identify which modelling assumptions, performance metrics, and design variables are most influential, and to clarify the main trade-offs that need to be addressed in later structural response extensions of the framework. The current stage also defines suitable parametrisations, numerical settings, and optimisation strategies that remain computationally tractable in subsequent extensions to coupled fluid-structure solvers. The specific objectives are:

1. Develop a general, modular optimisation framework that integrates surrogate models and numerical solvers, which can be implemented with

different algorithms or fidelity levels. The goal is to reduce reliance on computationally expensive simulations while maintaining accuracy and robustness, and to validate the resulting optimal designs using higher-fidelity CFD simulations.

2. Construct an efficient parametrisation method for symmetrically cambered aerofoil shapes and optimise these geometries to achieve the largest thrust under realistic constraints. This includes assessing the sensitivity of performance to key geometric parameters to identify compact but expressive design spaces.

3. Perform an exploratory optimisation of multi-sail layouts to assess aerodynamic interference effects across a representative range of spacings. Identify trends and layout principles for efficient deck arrangements, which could support industrial design and preliminary layout decisions for wing sail-equipped vessels.

1.4 Limitations and Assumptions

The present study is carried out under a number of assumptions and limitations that should be kept in mind when interpreting the results. These choices are partly dictated by the need to keep the optimisation and validation computationally tractable, and partly by practical constraints related to available computational resources. They also indicate directions for future extensions of the optimisation framework.

- The present optimisation is based on two-dimensional reduced-order aerodynamic modelling, with validation via three-dimensional CFD simulations. This reduces the computational cost but means that fully three-dimensional effects, such as spanwise flow and tip vortices, are only indirectly accounted for through these validation cases rather than being embedded in the optimisation loop.

- No experimental validation has been performed. Modern wing sails are very large structures; for example, the reference wing sail in the validation is 72 m tall with a 14 m chord, corresponding to a Reynolds number of order 10^7. There are currently no full-scale facilities capable of testing such a large structure, and achieving a comparable Reynolds number at model scale in laboratory tests would require conditions that significantly alter the Mach number or introduce other scaling distortions. As a result, the assessment of accuracy relies entirely on numerical simulations with thorough method validation, such as a mesh sensitivity study.

- The aerodynamic analyses are performed over a limited number of AOAs and inflow conditions. Therefore, the behaviour outside the sampled range, especially near stall or in highly unsteady regimes, is not fully resolved.

- Side forces induced by the wing sail are not included in the optimisation, since no representative ship configuration is specified from which to derive the influence of side forces on overall propulsion performance. The present results should therefore be interpreted primarily in terms of thrust generation rather than full vessel-level performance.

- The atmospheric boundary layer is not explicitly modelled in the three-dimensional validation simulations. A uniform inflow profile is used, consistent with the two-dimensional optimisation, which cannot represent height-dependent velocity profiles. This choice simplifies the comparison between two- and three-dimensional results but limits the direct applicability of the simulations to realistic atmospheric boundary-layer conditions.

- The aerofoil design optimisation is only done in two dimensions. While three-dimensional effects such as spanwise vortices influence performance, high-fidelity CFD validation confirmed that the 2D-optimised aerofoil maintains superior thrust compared to the benchmark design in three dimensions.

- The multi-sail optimisation focuses on layout design, aiming to evaluate the aerodynamic interference among wing sails for different spacings. Operational variables, such as the individual sheeting angles for each wing sail, are not considered because their necessity depends on the level of aerodynamic interference.

- Given computational constraints and the exploratory nature of this stage, the structural model of the wing sail will be deliberately simplified for the next-step development of the optimisation framework. The model will not capture detailed internal layouts, local stiffness variations, or joint conditions. Consequently, deformation-induced aerodynamic changes, stress concentrations, and structural failure modes such as buckling or fatigue are excluded from the optimisation loop.

- Only steady or quasi-steady aerodynamic forces are taken into account for the structural analysis of the simplified wing sail models. Unsteady phenomena such as gust loading, dynamic stall, motion-induced effects, and fully coupled FSI would be either neglected or modelled in a highly simplified manner in the optimisation. Consequently, the optimisation framework cannot address dynamic structural responses.

Chapter 2

Methodology

The modular optimisation framework developed in this thesis integrates a range of numerical methods and tools to perform aerodynamic optimisation. To demonstrate its modularity and interchangeability, different geometric parametrisations, surrogate modelling algorithms, and flow solvers have been implemented and tested within the same framework, which are presented in the published papers. This chapter elaborates on the methods that are used to achieve this goal. For brevity, only the methods down-selected for the final framework are presented.

It starts with definitions of the relevant wind conditions and force components acting on a wing sail in section 2.1. This is followed by the parametrisation of the wing sail profile geometry and the multi-wing sail layout used for optimisation in section 2.2. Next, the two aerodynamic analysis methods employed in Papers I and II are introduced in section 2.4. The optimisation model is described in section 2.5. Lastly, the CFD simulation setup for validating the optimised wing sail profiles is presented in section 2.6.

2.1 Sailing Forces

Two wind angle definitions are commonly used in sailing: the true wind angle (TWA) and the apparent wind angle (AWA). The TWA is the wind angle measured by a fixed external observer, for example, on shore, whereas the AWA is the wind angle relative to the moving ship itself. When the ship is stationary, the AWA and TWA are identical. Once the ship moves, the AWA results from the vector sum of the true wind and the ship velocity. For aerodynamic analysis, only the apparent wind needs to be considered, as it directly determines the wind loads acting on the sails (Kimball 2009). In this study, the AWA is defined as positive in the counter-clockwise direction when viewed from the top of the ship, as illustrated in Figure 2.1.

The thrust vector is defined as the projection of the resultant aerodynamic force F_{res} onto the ship's forward sailing direction, representing the portion of the aerodynamic load that contributes directly to propulsion. This definition is shown in Figure 2.2. Here, the AOA is denoted by α, the apparent wind vector and angle by V_{AW} and θ_{AW}, respectively, the lift force by L (perpendicular to V_{AW}), the drag force by D (aligned with V_{AW}), and the resulting thrust by T. For a given apparent wind condition, the balance between L and D, together with their orientation relative to the ship's motion, determines the magnitude and sign of T, and thus the effective propulsive performance of the wing sail.

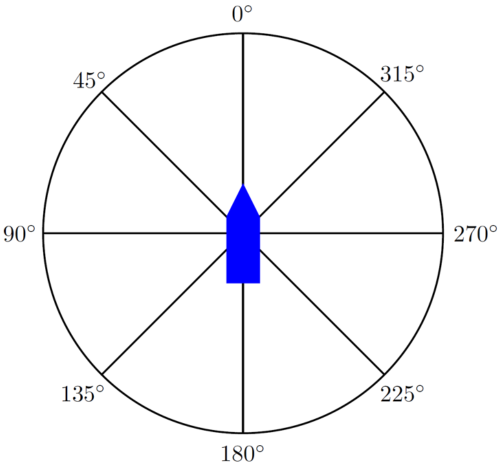

Figure 2.1: Definition of the apparent wind angle from a top-down view, where the ship is in blue.

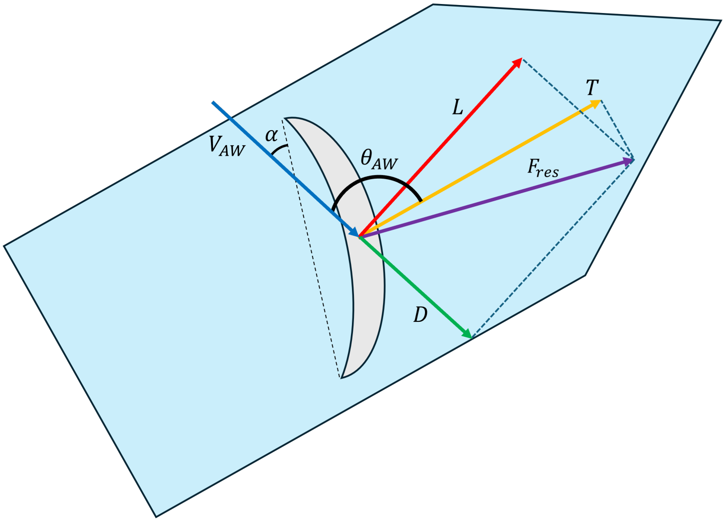

Figure 2.2: Diagram of sail force decomposition, outlined with the AOA (α), the lift (L), the drag (D), the resultant aerodynamic force (F_{res}) and the thrust (T) with respect to the apparent wind vector (V_{AW}) at the AWA (θ_{AW}).

From the definition illustrated in Figure 2.2, the thrust can be derived as:

$$T = L(\alpha) \cdot \sin(\theta_{AW}) - D(\alpha) \cdot \cos(\theta_{AW}) \ , \tag{2.1}$$

where L and D are dependent on α. When non-dimensionalised, the thrust, lift, and drag coefficients are defined as:

$$C_t = \frac{T}{\frac{1}{2}\rho V_\infty^2 c}, \qquad C_l = \frac{L}{\frac{1}{2}\rho V_\infty^2 c}, \qquad C_d = \frac{D}{\frac{1}{2}\rho V_\infty^2 c}, \tag{2.2}$$

where ρ is the air density, V_∞ is the speed of the incoming flow and c is the chord of the aerofoil.

According to Equation 2.1, the thrust coefficient varies directly with the AWA and indirectly with the AOA through C_l and C_d. A wing sail can fully rotate in the horizontal plane, that is, be freely adjustable to any AOA. Thus, for an arbitrary AWA, the thrust is calculated for the AOA that maximises thrust. This thrust variance means that one cannot simply assign a single thrust value to a certain geometry. Instead, the optimisation objective is not strictly the thrust but the average thrust over a range of AWAs. This yields:

$$\overline{C}_t = \frac{1}{N}\sum_{i=1}^{N} C_{t,i} \ , \tag{2.3}$$

where i represents the ith AWA of interest, and N is the total number of angles involved in the evaluation. This thesis considers AWAs from 10° to 150°. Beyond 150°, drag becomes a significant thrust component (Zhu et al. 2023a), allowing deep-stall configurations. Forward motion makes these angles less common, so they are less critical for analysis. The 10° lower bound reflects the negligible thrust in near-headwinds, when wing sails are typically retracted.

In the case that multiple wing sails are arranged on the deck, it makes sense to define the total thrust produced by all wing sails together at an arbitrary θ_{AW} as $C_{t,\mathrm{mult}}$. The definition is presented in Equation 2.4, where j is the jth sail unit, and M is the total number of sails. Note that $C_{t,\mathrm{mult}}$ is written as $C_{T,\mathrm{sum}}$ in Paper II.

$$C_{t,\mathrm{mult}} = \sum_{j=1}^{M} C_{t,j} = \sum_{j=1}^{M} [C_{l,j}\sin(\theta_{AW}) - C_{d,j}\cos(\theta_{AW})]. \tag{2.4}$$

2.2 Wing Sail Profile Parametrisation

To optimise a wing sail aerofoil, its shape needs to be parametrised. In Paper I, a new hybrid parametrisation method for symmetrically cambered aerofoils proposed by Yao (2025) is used to generate various geometric samples during the optimisation process. The parametrisation is formulated in reference to NACA-series foils and cambers.

The camber line is based on the NACA 4-digit definition (Jacobs et al. 1933). Since the crescent aerofoil is symmetric, the front section from the leading edge to mid-chord is used, and then mirrored to form a complete camber line. The resulting definition for the local camber height y_c is:

$$\frac{y_c}{c} = \frac{4h_c}{c}\left[\frac{x}{c} - \left(\frac{x}{c}\right)^2\right], \quad 0 \leq \frac{x}{c} \leq 1 \tag{2.5}$$

where h_c is the maximum camber, c the chord length and x is the chord-wise coordinate.

From this reference, the actual camber line in the new parametrisation is constructed using a Bézier curve defined by control points. These points shape the curvature of the front half of the camber line, while the rear half is generated by mirroring the front. Any number Q of control points can be used, but control points 1 and Q are always the same. Control Point 1 (CP 1) is fixed at the leading edge of the foil with nondimensionalised coordinates $(x/c, y/c) = (0, 0)$, and Control Point Q (CP Q) is fixed at $(x/c, y/c) = (0.5, h_c/c)$. The remaining $Q - 2$ control points are left unfixed to allow for flexibility in adjusting the curvature. The initial Bézier curve is obtained by fitting it to the reference camber line defined in Equation 2.5.

The thickness distribution of the new symmetrically cambered aerofoil parametrisation is derived from a modified NACA 4-digit series (Mason 2018). The front half of the thickness distribution is defined as:

$$\frac{y_t}{c} = 5\left(\frac{h_t}{c}\right)\left[a_0\sqrt{\frac{x}{c}} + a_1\left(\frac{x}{c}\right) + a_2\left(\frac{x}{c}\right)^2 + a_3\left(\frac{x}{c}\right)^3\right], \quad 0 \leq \frac{x}{c} \leq 0.5 \tag{2.6}$$

where h_t is the maximum thickness. The coefficient a_0 is given as:

$$a_0 \approx 0.296904 \cdot \chi_{LE}\,, \tag{2.7}$$

and

$$\chi_{LE} = \begin{cases} I/6\,, & \text{for } I \leq 8 \\ 10.3933\,, & \text{for } I = 9 \end{cases} \tag{2.8}$$

where I is the index used to specify the leading/trailing edge radius.

The other coefficients are:

$$a_1 \approx 0.477 - 2.650a_0 \tag{2.9}$$

$$a_2 \approx -0.708 + 3.536a_0 \tag{2.10}$$

$$a_3 \approx 0.308 - 2.121a_0 \tag{2.11}$$

According to Equation 2.6–2.11, the thickness distribution requires h_t and I as inputs. The front half is mirrored about the mid-chord to form a symmetric profile, following the same approach used for the camber line. The resulting camber and thickness distributions are then combined to produce a symmetrically cambered aerofoil.

In Paper I, the sample aerofoils are generated with $Q = 4$ control points. A geometry example and its parameters are presented in Figure 2.3. The variable parameters are I, h_t, and the x and y-coordinates for CP 2 and CP 3 of the Bézier curve.

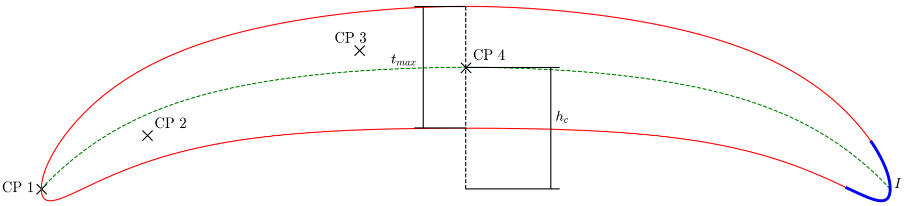

Figure 2.3: A parametrised crescent aerofoil with a camber that is symmetric about the mid-chord. Control points such as CP 1–CP 4 are used to define a Bézier curve for half of the camber line.

The parameter ranges for the proposed parametrisation are constrained using the D2R10 as a reference. The minimum thickness is set to 14.29% of the chord, matching the D2R10 to allow mast integration, while the upper bound is 20% above this value. The maximum camber height remains fixed at 14.25% of the chord, as preliminary studies show its strong influence on aerodynamic forces, which could obscure the effects of other parameters. The coordinates of the four control points (x_2, y_2, x_3, y_3) may vary by ±5% from those fitted to the NACA 4-digit camber line. Finally, I can vary from 4 to 8.

2.3 Multi-Sail Layout Parametrisation

To study the interaction between multiple wing sails in different configurations, two installation layouts are developed. These layouts are investigated in Paper II and are presented in Figure 2.4.

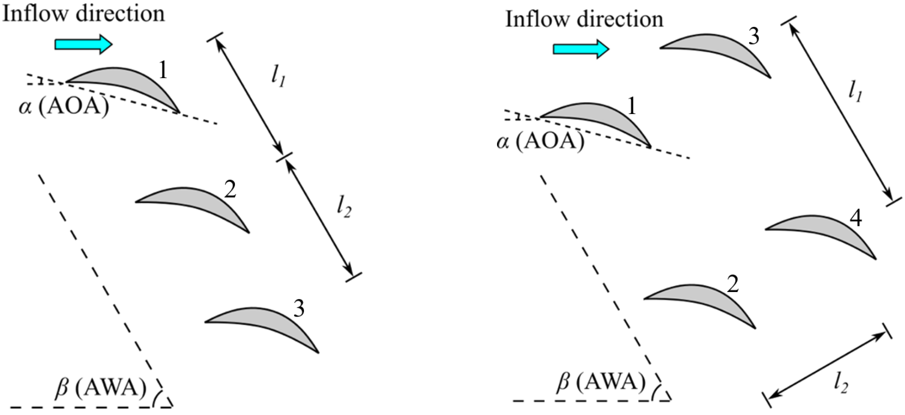

(a) Triple layout (b) Quad layout

Figure 2.4: Schematics of conceptual wing sail installation layouts: (a) the triple units in-line (TL), and (b) the quad units in parallel (QP).

The first layout assembles three sails in a triple in-line layout (TL), and the second layout has two lines of two sail units in parallel, referred to as quad in parallel (QP). The distances between the sail units, l_1 and l_2, are defined in different ways for the two layouts. In the first layout, l_1 is the distance between the first and second sail, and l_2 is the distance between the second and third sail. In the second layout, l_1 and l_2 are the distances between the longitudinal and lateral rows, respectively.

In Paper II, both layouts are optimised under the constraint that the sum of l_1 and l_2 remains constant at 2.4 chord lengths. This constraint reflects the structural, maintenance, and stability requirements, resulting in compact, retrofit-friendly layouts. By enforcing these limits, both configurations operate under the same practical conditions, enabling a fair comparison between TL and QP setups.

Figure 2.4 also illustrates the definitions of the AWAs and AOAs. Three AWAs are analysed: 60°, 90°, and 120°. These angles are selected because they represent the area of highest thrust generation according to Paper I. Accordingly, only the higher AOA of 15° is considered, as it is high enough to produce potential interference effects. This choice avoids coming too close to the expected stall angle of the D2R10 at 20° (Zhu et al. 2023a), which the current aerofoil resembles.

2.4 Aerodynamic Analysis

Paper I and Paper II use two different modules from the open-source aerodynamic analysis tool AeroSandbox (Sharpe 2024) to calculate lift and drag. Paper I uses Neuralfoil (Sharpe et al. 2025), a physics-informed neural network-based fast aerofoil analysis tool. Paper II uses the inviscid module for two-dimensional aerodynamics.

2.4.1 NeuralFoil Analysis

NeuralFoil (Sharpe 2023) is a machine learning model trained on data generated by XFOIL, designed for two-dimensional aerodynamic analysis under both viscous and inviscid conditions, as well as steady, incompressible, and compressible flows. Its predictions match the accuracy of XFOIL while offering significantly faster computation times (Sharpe et al. 2025). Furthermore, XFOIL can occasionally fail to converge when handling complex simulations (Sharpe et al. 2025), making it difficult to analyse unconventional symmetrically cambered wing sails. In contrast, NeuralFoil produces reliable results across a wide range of geometries. NeuralFoil also integrates analytical relations developed by Truong (2020) into its architecture, extending its capabilities beyond XFOIL-based predictions. These formulations enable the estimation of C_l and C_d across a full 360° AOA range and include models that capture moderate to deep stall behaviour.

2.4.2 Inviscid Analysis

The inviscid solver in AeroSandbox applies a two-dimensional vortex-panel formulation to model multiple independent lifting surfaces. Each panel is associated with a degree of freedom, represented by its vortex strength, and a corresponding equation enforcing the no-penetration condition. The Kutta condition introduces an additional constraint, which is implemented via a distributed source term. This results in a dense linear system that is solved to reconstruct the flow field. While this inviscid method is computationally efficient, it neglects viscous phenomena, such as the boundary layer, separation, and wake interactions, thereby limiting the accuracy of drag predictions. Nevertheless, it provides reliable lift estimates and computes much faster than Reynolds Averaged Navier-Stokes (RANS)-based CFD approaches.

2.5 Optimisation Model

2.5.1 Optimisation Algorithm

The optimisation model developed for this research uses the Bayesian optimisation method (Snoek et al. 2012; Hebbal et al. 2021), where an objective function is minimised by approximating it with a probabilistic surrogate model. A surrogate model is a mathematical approximation of a complex system, such as one represented by CFD simulations or experimental data, that enables the estimation or prediction of outputs at new points not included in the original dataset. In this case, the surrogate model is a Gaussian Process (GP) (Snoek et al. 2012).

A GP is defined as a collection of random variables, where any finite subset of these variables has a joint Gaussian distribution (Rasmussen 2004; Hebbal et al. 2021). It is specified by a prior mean function $m(\boldsymbol{\phi})$ and a prior covariance function $k(\boldsymbol{\phi}, \boldsymbol{\phi}')$, also known as a kernel, where $\boldsymbol{\phi}$ and $\boldsymbol{\phi}'$ are inputs in the input space Φ:

$$f(x) = GP(m(\boldsymbol{\phi}), k(\boldsymbol{\phi}, \boldsymbol{\phi}')) \tag{2.12}$$

The mean function describes the prior belief about the trend of the function that the surrogate model seeks to approximate. If nothing is known, a constant mean function may be used. The covariance function captures the prior belief about the correlation between individual samples (Shan et al. 2025; Hebbal et al. 2021).

This study uses the Matérn kernel (Rasmussen et al. 2006), which is particularly suitable for modelling functions that are not overly smooth. The Matérn kernel with smoothness parameter $\nu = 5/2$ and a length scale $l = 1$ is defined as (Rasmussen et al. 2006):

$$k(\boldsymbol{\phi}, \boldsymbol{\phi}') = (1 + \sqrt{5}\,|\boldsymbol{\phi} - \boldsymbol{\phi}'| + \frac{\sqrt{5}\,(\boldsymbol{\phi} - \boldsymbol{\phi}')^2}{3}) \exp(-\sqrt{5}\,|\boldsymbol{\phi} - \boldsymbol{\phi}'|) \tag{2.13}$$

The algorithm that creates a GP is called the Gaussian Process Regression (Rogers et al. 2023). When the GP observes sample data $S_o(\boldsymbol{\phi}_o, \boldsymbol{\psi}_o)$, with observed inputs $\boldsymbol{\phi}_o$ and observed objective values $\boldsymbol{\psi}_o$, it updates its prior beliefs to form a posterior distribution. This posterior is defined by a predictive mean $\mu_*(\boldsymbol{\phi}_*)$ and variance $\sigma_*^2(\boldsymbol{\phi}_*)$ (Rasmussen 2004). These functions describe a Gaussian distribution over the objective value at any new input $\boldsymbol{\phi}_* \in \Phi$, capturing both prediction and uncertainty (Snoek et al. 2012; Hebbal et al. 2021). The functions are defined as:

$$\mu_*(\boldsymbol{\phi}_*) = \boldsymbol{k_*}^T(\boldsymbol{K} + \sigma_n^2 I)^{-1}\boldsymbol{\psi}_o \tag{2.14}$$

$$\sigma_*^2(\boldsymbol{\phi}_*) = k(\boldsymbol{\phi}_*, \boldsymbol{\phi}_*) - \boldsymbol{k_*}^T(\boldsymbol{K} + \sigma_n^2 I)^{-1}\boldsymbol{k_*} \tag{2.15}$$

To guide the optimisation, an acquisition function is derived from the GP posterior. This function determines where to sample next by balancing exploration (uncertainty) and exploitation (high predicted value) (Wang et al. 2023). The acquisition function used in this study is the Upper Confidence Bound (UCB), defined as:

$$\text{UCB}(\boldsymbol{\phi}_*) = (\mu_*(\boldsymbol{\phi}_*) + k\sigma_*(\boldsymbol{\phi}_*)), \tag{2.16}$$

The parameter k controls the trade-off between exploration and exploitation (Snoek et al. 2012). A higher value of k encourages exploration by favouring regions with high uncertainty, while a lower value promotes exploitation of regions with high predicted objective values. In this study, k is set to 2.576, corresponding to the 99% confidence interval of a standard normal distribution.

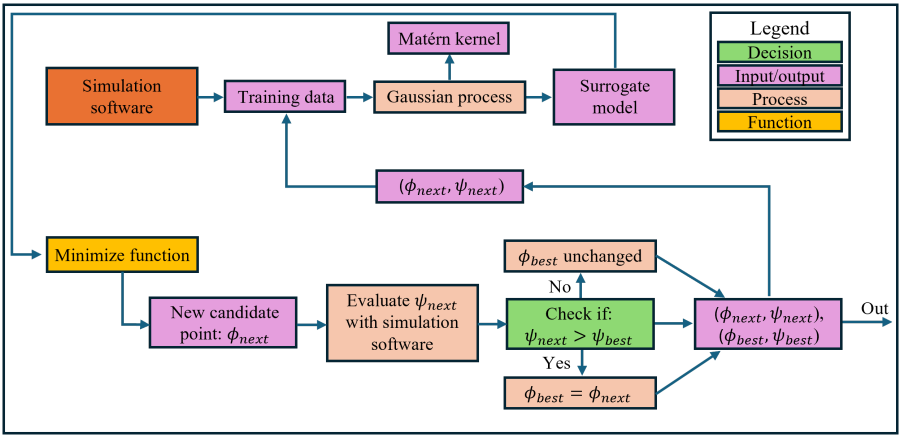

Figure 2.5: The flowchart of the optimisation process.

2.5.2 Optimisation Process

The optimisation process is illustrated in Figure 2.5. The process begins by generating the initial dataset S_o using simulation software. The best sample in S_o, defined as the one with the highest $\boldsymbol{\psi}_o$, in this case $\overline{C}_t$, is selected as $\boldsymbol{\phi}_{best}$. This initial dataset is then used to train the Gaussian Process surrogate model $\mathcal{GP}_0$. Both $\mathcal{GP}_0$ and $\boldsymbol{\phi}_{best}$ are passed to the *minimize* function from Scipy.

To maximise UCB($\boldsymbol{\phi}$), the function is set to minimise $a(\boldsymbol{\phi}) = -\mathrm{UCB}(\boldsymbol{\phi}_*)$, using the SLSQP algorithm (Virtanen et al. 2020; Kraft 1988). The *minimize* function returns the minimised sample $\boldsymbol{\phi}_{next}$. The simulation software is then used to evaluate $\boldsymbol{\psi}_{next}$. If $\boldsymbol{\psi}_{next}$ is greater than $\boldsymbol{\psi}_{best}$, then $\boldsymbol{\phi}_{next}$ becomes the new $\boldsymbol{\phi}_{best}$. Otherwise, $\boldsymbol{\phi}_{best}$ remains unchanged. In either case, the pair $(\boldsymbol{\phi}_{next}, \boldsymbol{\psi}_{next})$ is added to the dataset to refine the surrogate model. In the next iteration, $\boldsymbol{\phi}_{best}$ is input to the *minimize* function. This process is repeated for a specified number of iterations. In summary, the corresponding algorithm is shown in Algorithm 1 as follows.

Algorithm 1: Bayesian optimisation using minimize

```
Input: Initial dataset $S_o(\boldsymbol{\phi}_o, \boldsymbol{\psi}_o)$, number of iterations $T$, constraints $\mathcal{C}$
Output: Best found point $\boldsymbol{\phi}_{best}$
Train Gaussian Process surrogate model $\mathcal{GP}_0$ on $S_o$
Set $\boldsymbol{\phi}_{best} \leftarrow \arg\max_{\boldsymbol{\phi}_o \in S_o} \boldsymbol{\psi}_o$
for $t \leftarrow 1$ to $T$ do
    Set $\boldsymbol{\phi}_0 \leftarrow \boldsymbol{\phi}_{best}$
    Define acquisition function $a(\boldsymbol{\phi}_*) = -\mathrm{UCB}(\boldsymbol{\phi}_*)$
    $\boldsymbol{\phi}_{next} \leftarrow \mathsf{minimize}(a(\boldsymbol{\phi}), \boldsymbol{\phi}_0, \mathcal{GP}_0, \mathcal{C})$
    $\boldsymbol{\psi}_{next} \leftarrow \mathsf{Simulation\ software}(\boldsymbol{\phi}_{next})$
    if $\boldsymbol{\psi}_{next} > \boldsymbol{\psi}(\boldsymbol{\phi}_{best})$ then
        $\boldsymbol{\phi}_{best} \leftarrow \boldsymbol{\phi}_{next}$
    end
    Add $(\boldsymbol{\phi}_{next}, \boldsymbol{\psi}_{next})$ to dataset
    Retrain Gaussian Process on updated dataset
end
return $\boldsymbol{\phi}_{best}$
```

2.6 CFD Validation Setup

The optimisation results from Paper I were validated with three-dimensional CFD simulations using IDDES. This is a hybrid method that combines Large Eddy Simulation (LES) in separated or free-shear regions and the RANS equations near wall boundaries. In the present work, the $\mathrm{k}-\omega$ SST turbulence model is employed within the IDDES. The simulation setup closely follows the methodology described by Zhu et al. (2024) and Zhu et al. (2023b), where the current setup has been validated. The Reynolds number is 1×10^7. The commercial software STAR-CCM+ is used for the simulations.

The three-dimensional wing sail model is generated by scaling the optimised aerofoil profile to a chord length of 14 m and extruding it vertically to a height of 72 m along the z-axis. The computational domain, illustrated in Figure 2.6, is a rectangular box measuring 600 m in length, 300 m in width, and 360 m in height. The wing sail is placed 150 m downstream of the inlet, which is the upstream face of the domain.

A velocity inlet boundary condition is applied to the inlet. A symmetry boundary condition is imposed on the bottom face to avoid modelling the free surface. All remaining outer faces are assigned pressure outlet boundary conditions, allowing the flow to convect out of the domain with minimal reflection. A no-slip wall condition is applied to the surface of the wing sail. The bottom end of the wing sail is fixed to the domain bottom to represent a mast or support structure anchored to the deck.

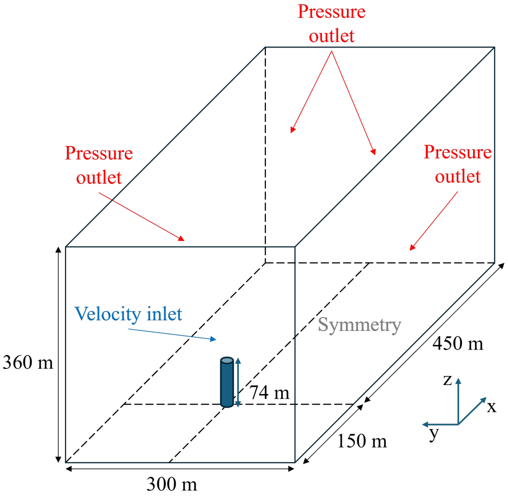

Figure 2.6: Schematic of the computational domain, including its dimensions and boundary conditions, for the CFD simulations. Note that the cylinder is representative of the wing sail.

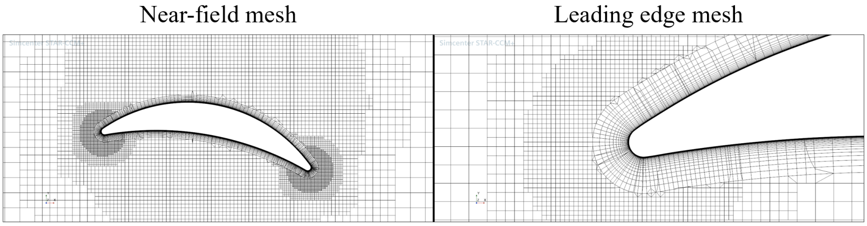

Figure 2.7: Cut-plane views of the mesh, with the D2R10 (Zhu et al. 2023a) as example. Includes the near-field view and the view around the leading edge at $\alpha = 10°$.

An unstructured mesh consisting of approximately 21 million cells is employed for validation simulations. To accurately resolve the boundary layer, 65 prism layers are applied along the geometry walls, with a total thickness of 0.5 m. The mesh is refined near critical regions, such as the leading and trailing edges, with cell sizes around 0.08 m. Figure 2.7 presents an overview of the near-field mesh, along with a close-up view of the mesh refinement at the leading edge, for the D2R10. The mesh independence was previously confirmed by Zhu et al. (2023a).

Chapter 3

Summary of Papers

This section summarises the main results and findings of Papers I and II to show how they contribute to the overall optimisation framework developed in this thesis. Together, the two papers form the aerodynamic stage of the framework, focusing on the design and optimisation of single-sail profiles and multi-sail layouts.

Paper I develops the Bayesian optimisation framework for individual symmetrically cambered wing sail profiles. It combines a dedicated geometry parametrisation with a surrogate aerodynamic model to maximise average thrust over a relevant range of apparent wind angles, and it demonstrates that the modular setup can be linked to higher-fidelity CFD for validation. In this way, Paper I establishes the basic optimisation workflow for aerofoil profile design and identifies the most influential geometric parameters for thrust production.

Paper II builds on the application of the optimisation framework at the layout level for multiple wing sails installed within a limited deck area. Using a reduced-order aerodynamic model and a genetic algorithm, it examines how different spacings and arrangements affect thrust through aerodynamic interference. The results provide preliminary layout guidelines and show how the same modular framework can be used to address both shape and layout design.

3.1 Summary of Paper I

The Bayesian optimisation framework presented in section 2.5.1 is developed in this paper. It introduces the wing sail profile parametrisation and uses NeuralFoil to develop a profile shape that maximises average thrust across a wide range of AWAs. The initial dataset contains 100 samples and wing sails are evaluated at AOAs between 10° and 15°. This avoids reaching the stall angle of the D2R10 at 20°, which is impractical due to gust risks and varying stall behaviour. The AWAs are limited to 10° to 150°, since these angles are more likely to be encountered. The Reynolds number of the flow is 1×10^7, and the flow is assumed incompressible.

Figure 3.1 shows the current and best candidate at each iteration (red dashed and green solid lines, respectively). Iteration zero corresponds to the best initial sample (Initial Aerofoil) with $\overline{C}_t = 1.58$. The algorithm explores the design space for 17 iterations before converging, with small variations in objective value, indicating a focus on exploitation over exploration.

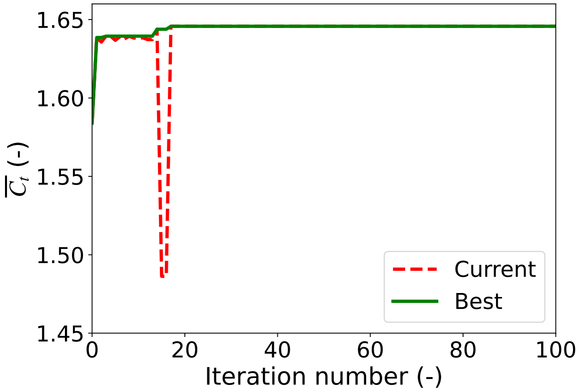

Figure 3.1: Predicted thrust results for each candidate aerofoil over 100 iterations (dashed), along with the best candidate identified so far at each iteration (solid).

The final optimised aerofoil (BN4) is compared with the Initial Aerofoil and the D2R10 in Figure 3.2. Compared to the D2R10, the optimised foils are thicker away from the chord centre, due to their blunter edges and their larger values of y_2, which is the y-coordinate of CP 2. The differences between the BN4 and the Initial Aerofoil are smaller, as the BN4 has slightly larger y_2 and reduced h_t. The predicted $\overline{C}_t$ is 1.65, a 4% and 26% improvement over the Initial Aerofoil and the D2R10, respectively.

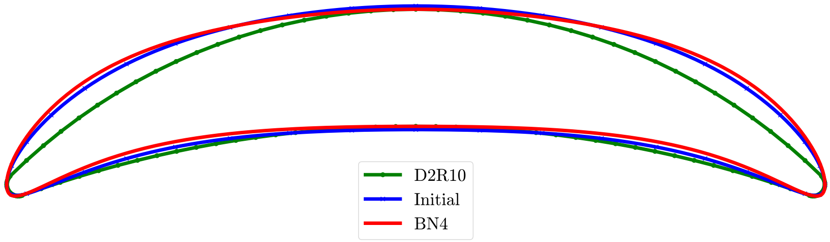

Figure 3.2: The three foil profiles, including the D2R10 (Zhu et al. 2023a), the Initial Aerofoil and the optimised BN4.

Sensitivity analysis further revealed that the average thrust indeed increases notably with y_2, which reflects a more uniform thickness distribution away from the chord centre. This strong influence of y_2 explains why the BN4, with its higher CP 2 position, achieves a significant thrust advantage over the D2R10. The maximum thickness shows a slight negative correlation with thrust, while the leading/trailing edge radius index exhibits a slight positive correlation.

The C_l–α and C_d–α trends for all three aerofoils are shown in Figure 3.3. Across the entire AOA range, the BN4 delivers the highest C_l, with the Initial Aerofoil close behind and the D2R10 consistently lower. The drag behaviour contrasts with this: the D2R10 exhibits the largest C_d, while the Initial Aerofoil and BN4 alternate in performance. Up to $\alpha = 13°$, the BN4 produces slightly more drag than the Initial Aerofoil, after which the trend flips. This pattern may stem from limitations in NeuralFoil's drag-prediction accuracy.

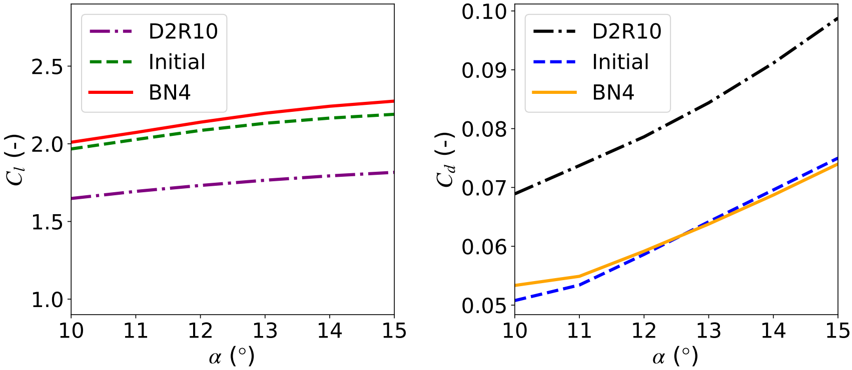

Figure 3.3: (left) The lift coefficient and (right) the drag coefficient of the D2R10, the Initial Aerofoil and the optimised BN4, predicted using NeuralFoil.

The thrust coefficient is presented in Figure 3.4 as a function of θ_{AW} and shows that the BN4 consistently outperforms the other aerofoils across the entire wind angle range. The Initial Aerofoil generates slightly lower thrust, while the D2R10 shows significantly poorer performance. These findings suggest that, within the θ_{AW} range of 10° to 150°, higher lift is associated with increased thrust. The greatest differences in thrust occur in the mid-range of θ_{AW}, roughly from 50° to 110°, corresponding to crosswind conditions where wing sails generate significant thrust through lift. In this region, the BN4 and the Initial Aerofoil outperform the D2R10 due to higher lift production. Outside this range, lift plays a smaller role relative to drag, and since the aerofoils' drag values are similar, the difference in thrust becomes smaller.

The three-dimensional IDDES simulations for the BN4 and D2R10 confirmed the improvement in thrust production by the former, but with a more modest 8% increase. The BN4 results showed reduced flow separation and a larger low-pressure region across the suction side, with a more uniform pressure distribution. Together, these effects contribute to improved thrust performance.

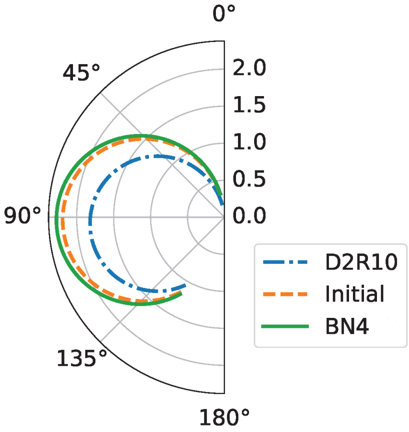

Figure 3.4: The distribution of the thrust coefficient, C_t, as a function of the AWA, for the three different foil profiles.

3.2 Summary of Paper II

The study covers a preliminary analysis of the aerodynamic interaction between multiple symmetrically cambered wing sails. It investigates the thrust of two different layouts and compares them with individual sails. The optimal values for l_1 and l_2 of each layout are determined using a generic genetic algorithm. The aerodynamic analysis is performed with a 2D inviscid model that provides quantitative predictions of the lift production but does not provide accurate drag predictions. The results are therefore primarily applicable to wind angles where the lift dominates the sail thrust, and as an initial screening of different layouts. The analysis uses a reference aerofoil with the same camber height and maximum thickness as the D2R10. The aerodynamic performance analysis is conducted at specific AWAs that should be most relevant to practical operations. A matrix of these angles of high interest is defined in Table 3.1.

Table 3.1: The matrix of AOAs and AWAs.

Layout type	**AOA**	**AWA**
TL	15°	60°, 90°, 120°
QP	15°	60°, 90°, 120°

The results for the TL layout are listed in Table 3.2, in the form of the ratio between $C_{t,\mathrm{mult}}$ and the sum of C_t for M individual sails. The optimised layout has l_1 is 1.151, and l_2 is 1.249. The optimisation leads to a significant improvement of the thrust ratio at $\theta_{AW} = 60°$, but no changes at 90° and 120°. The optimal layout shows similar values for each AWA, indicating improved robustness. The TL layout shows a 4–6% penalty in thrust generation compared to three individual sails, indicating aerodynamic interference between the sails. The penalty can be explained by looking at Figure 3.5, which illustrates the pressure coefficient C_p distribution of all three wing sails and the reference sail. Sail 1 experiences a marked increase in suction, especially over $x = 0.3$–0.7, which boosts its thrust. Conversely, Sails 2 and 3 display progressively weaker suction, reflecting reduced thrust generation downstream. Overall, the TL arrangement benefits the first sail but increasingly penalises the subsequent sails, resulting in a total thrust below that of three isolated sails.

Table 3.2: The ratio of the thrust coefficient $C_{t,\text{mult}}/(C_t M)$ of the TL layout at different AWAs.

	l_1	l_2	60°	90°	120°
Equal spacing	1.200	1.200	0.790	0.949	0.933
Optimal	1.151	1.249	0.961	0.949	0.933

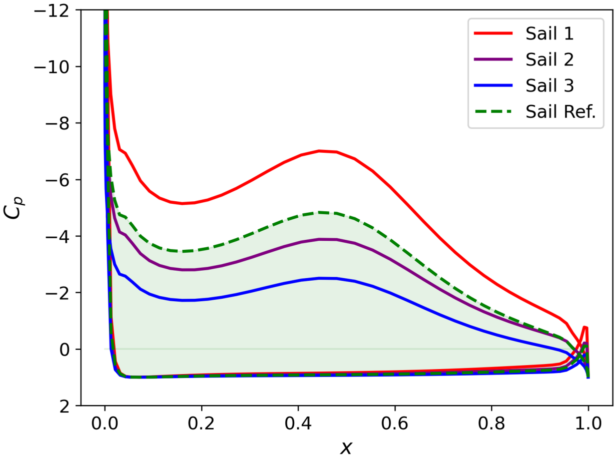

Figure 3.5: The pressure coefficients, C_p, along the chord for the sails in the optimised TL layout (Sails 1–3) compared to the reference individual sail, at $\alpha = 15°$ and $\theta_{AW} = 90°$.

The QP layout optimisation yields l_1 is 1.132 and l_2 is 1.268. The improvement from optimising the distances is minimal: thrust at 60° and 90° remains unchanged, with only a 2% increase at 120°. Compared with the sum of four individual sails, the QP layout produces substantially less thrust, particularly at 60° (around 28% lower), and reductions at 90° and 120° are also notable, approximately 17% and 10% respectively. Figure 3.6 presents the C_p distributions for the optimised QP layout. The suction side of Sail 1 shows a clear increase in negative pressure compared with the individual sail. Sail 3 experiences increased negative pressure on the rear half of the suction side ($x = 0.5$–1.0), but reduced negative pressure near the leading edge and reduced positive pressure on the pressure side, due to the wake from Sail 1, resulting in lower thrust than the reference sail. Sail 2 shows a decrease in suction-side pressure, similar to the TL layout. Sail 4 generates markedly less pressure overall, particularly in the front half of the suction side, due to wake interference from Sail 2.

Table 3.3: The ratio of the thrust coefficient $C_{t,\text{mult}}/(C_t M)$ of the QP layout at different AWAs.

	l_1	l_2	60°	90°	120°
Equal spacing	1.200	1.200	0.722	0.833	0.889
Optimal	1.132	1.268	0.722	0.833	0.910

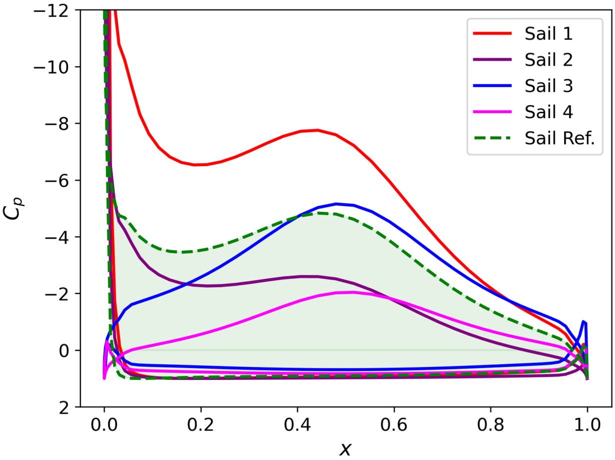

Figure 3.6: The pressure coefficients, C_p, along the chord for the sails in the optimised QP layout (Sails 1–4) compared to the reference individual sail, at $\alpha = 15°$ and $\theta_{AW} = 90°$.

To gauge the influence of viscous drag on the inviscid results, a generic viscous correction $C_d = C_{d0} + k\,C_l^2$ (typical bounds $C_{d0} \in [0.01, 0.02]$, $k \in [0.01, 0.02]$) is added to the thrust prediction, using Equation 2.1, with the reported C_l. The correction reduces absolute thrust but does not change layout rankings, showing that the inviscid $C_{T,\text{proxy}}$ is sufficient for comparative screening.

The symmetrically cambered profile generates lift mainly in the mid-chord region, which is largely unaffected by interference, supporting its use in multi-sail arrangements. The TL layout consistently produces less interference than the QP layout, which suffers a particularly large penalty for Sail 4. Finally, spacing optimisation showed limited effect, only meaningfully benefiting the TL layout at $\theta_{AW} = 60°$.

Chapter 4

Conclusions

Wing sails offer significant potential for reducing fuel consumption in the shipping industry. The symmetrically cambered aerofoil concept has demonstrated the ability to generate high thrust over a wide range of apparent wind angles. However, its geometry is still far from fully optimised, and there remains room to enhance thrust production and improve performance reliability. At the same time, ships have limited deck space, so multiple wing sails must often be installed in close proximity. A preliminary understanding of the aerodynamic interactions between multiple symmetrically cambered aerofoils is therefore needed to quantify interference effects and to inform the design of efficient multi-sail configurations that can be used in realistic ship layouts.

This licentiate thesis represents the first stage of addressing these challenges, focusing on establishing the optimisation framework and on conducting preliminary studies of symmetrically cambered aerofoils and multi-sail layouts. The goals can be summarised as:

1. Develop a modular optimisation framework that combines surrogate models with numerical solvers to reduce computational cost while maintaining accuracy, validated against high-fidelity CFD.

2. Create an efficient parametrisation for symmetrically cambered aerofoils and optimise these shapes for maximum thrust under realistic geometric and operational constraints.

3. Explore multi-sail layout optimisation to quantify aerodynamic interference across different spacings and identify design principles for efficient arrangements.

A surrogate-based Bayesian optimisation framework was developed for wing sail design. By employing a Gaussian Process as the probabilistic surrogate model, the method efficiently approximated the complex thrust objective function while quantifying prediction uncertainty. This approach enabled the optimiser to balance exploration and exploitation within the design space, significantly reducing the number of computationally expensive CFD simulations required. NeuralFoil served as a rapid aerodynamic analysis tool in the framework, providing accurate lift predictions across a wide range of apparent wind angles and supporting the iterative optimisation process.

The integration of this optimisation framework with a novel hybrid parametrisation method allowed for systematic variation of key geometric features under realistic constraints. A sensitivity analysis revealed that the y-coordinate of the camber line control point CP 2 strongly influences thrust by promoting a more uniform thickness distribution. This explains the superior performance of the optimised BN4 aerofoil compared to the D2R10. In the three-dimensional validation, the BN4 achieved an 8% improvement in predicted average thrust, demonstrating the effectiveness of combining surrogate models with Bayesian optimisation for wing sail design. Future work will extend this methodology to multi-objective optimisation and incorporate high-fidelity CFD data to enhance accuracy and robustness.

The thrust production efficiency of two multi-sail layouts was compared by optimising sail spacing for three different AWAs. The first layout consisted of three sails arranged in a line, while the second layout had four sails in a rectangular pattern. The aerodynamic analyses were carried out with a two-dimensional inviscid method. Under a fixed total spacing, the optimisation led to only limited gains. The rectangular layout showed no significant improvement overall, and the in-line layout only achieved a meaningful increase in efficiency at a single wind angle.

Compared to the sum of three individual sails, the single line layout produced up to 6% less thrust for the given AWAs. The rectangular layout performed consistently worse, with a thrust loss of about 10–28% relative to four individual sails. A semi-empirical drag correction confirmed this qualitative trend. This indicates that the conclusions are robust to the inclusion of viscous effects.

In summation, the novel points that this thesis has achieved can be summarised as follows:

- Introduced a new hybrid geometry parametrisation method for symmetrically cambered wing sails. Sensitivity analysis addressed key parameters affecting thrust production.
- Developed a Bayesian optimisation framework with a Gaussian process surrogate model, which integrated NeuralFoil for rapid 2D aerofoil analysis to reduce CFD reliance. The viability of this framework was demonstrated.
- Optimised multi-sail layouts for crosswind AWAs using a 2D inviscid method. It was found that the single-line layout had less aerodynamic interference than the rectangular layout, given the constraint of fixed spacing.

Chapter 5

Future Work

5.1 Fluid Structure Interactions

The ultimate goal is to integrate Fluid–Structure Interaction (FSI) into the optimisation process to address its influence on wing sail design. This is because aerodynamic loads deform a wing sail by bending and twisting the geometry and changing the effective angle of attack, and the structural deformations, in turn, alter the flow, redistributing the loads. Unlike conventional rigid-body assumptions in flow simulations, FSI requires coupling flow and structural deformation in simulations that resolve displacement, strain and internal structural forces. This coupling can be achieved in a weak form, where the information exchange between fluid and structure is performed once per time step with minimal iteration, or in a strong form, where fluid and structural solutions are iterated to mutual convergence at each time level to improve accuracy and stability (Venkatesh et al. 2024). The solving and coupling procedure is typically implemented either in a partitioned framework, where separate fluid and solid solvers are linked through interface conditions, or in a monolithic framework, where a single solver advances the coupled fluid–structure system in one step. Each approach involves different trade-offs in terms of numerical robustness, implementation complexity, and computational cost. A key challenge is to identify FSI strategies that are sufficiently accurate for design optimisation while remaining affordable for large sets of candidate geometries and operating conditions.

To make a coupled aerodynamic and structural optimisation, FSI effects must be expressed as objectives and design parameters. The most intuitive objective would be structural weight minimisation. FSI is inherently three-dimensional, requiring a three-dimensional wing sail model. This model can be split into an external aerodynamic shape and an internal structural model. The structural model ought to be simplified to a large extent in order to reduce the complexity of the FSI problem and to save computational resources.

The external shape for the 3D wing sail has four key parameters: the height b, the root chord length c_{root}, the taper ratio λ, and the cross-sectional aerofoil profile. The last parameter encompasses the aerofoil design, which is developed separately with the optimisation method from Paper I. The reason to keep the profile optimisation separate from the other parameters is that the effects of spanwise coherent vortex structures and wing tip vortices are much smaller than the streamwise flow passing over the aerofoil profile.

The complexity and computational costs of optimisation are, in general, increased by the number of optimised parameters. Given that FSI involves both flow and structural aspects, it is intuitive to take into account only the most critical design parameters and to consider simple models, while preserving physical fidelity. The structural model consists of three major components: the sail skin, the ribs and the mast. The sail skin and ribs use shell models. The mast has a circular pipe cross-section and is represented by a beam model. The pipe diameter is equal to the maximum thickness of the tip aerofoil. Therefore, the key variable parameters for the structural design include the number of ribs N_{ribs}, their spacing l_{ribs}, the shell thickness of the sail skin t_{skin} and ribs t_{ribs}, and the wall thickness of the mast t_{mast}. A schematic drawing of all the parameters is shown in Figure 5.1. It is worth noting that N_{ribs} and l_{ribs} will be fixed during the optimisation process, while further investigation will be conducted after a feasible design is down-selected from the optimisation.

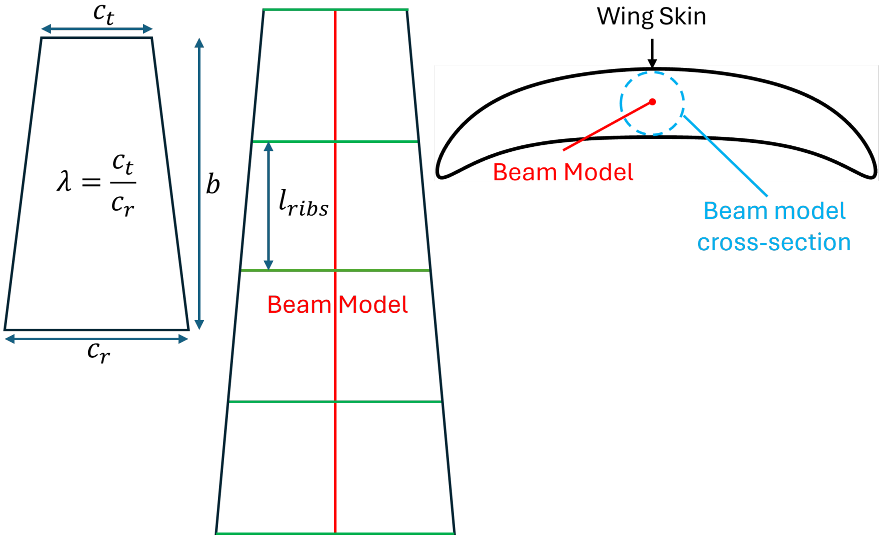

Figure 5.1: Schematic of the external wing sail shape and its simplified internal structural model, outlined with the key design parameters.

In the context outside the current scope, the structural dynamics properties could be explored in more detail, including the role of bending and torsional stiffness, damping coefficients, and other parameters governing the modal response. Variations in stiffness distribution along the span and chord may shift natural frequencies and mode shapes towards or away from critical excitation ranges. This directly affects the risk of resonance and flutter in the coupled FSI. Likewise, material and structural damping are able to control the decay of oscillations and the amplitude of limit-cycle responses under unsteady loading. A systematic parametric study of these properties, together with reduced-order modelling, would help identify combinations that lead to robust, well-damped structural behaviours without excessive weight penalties. On this basis, structural failure modes such as fatigue, buckling, and ultimate strength exceedance should be assessed using appropriate stress-cycle (S-N) curves, under both steady design loads and realistic gust or manoeuvre load histories. Integrating these failure metrics into the optimisation framework would allow the design space to be constrained not only by instantaneous stress limits but also by long-term durability and safety margins.

5.2 Future Research

Beyond FSI analysis, several directions for future work are identified. A first interesting study is to incorporate structural weight as an additional design objective in the newly developed optimisation framework. The problem is reformulated with multi-objectives to find the performance balance between structural weight and time-averaged thrust. Suitable normalisation measures of the objective functions need to be defined, so that their relative importance can be compared in a physically meaningful way. On this basis, different scaling approaches or Pareto-based formulations can be tested to quantify acceptable trade-offs and to map out design families that are structurally efficient without degrading propulsion capability or compromising dynamic response in FSI.

A related task is to understand how the structural-weight objective interacts with constraints such as maximum stress, strain, or tip deflection. Introducing weight as an explicit objective may shift the location of the optimal designs towards lighter, more flexible solutions, which, in turn, may alter the aeroelastic behaviour. Systematic sensitivity studies are therefore needed to assess how changes in material properties and thickness distributions influence both the structural weight and the thrust response, and to identify regions in the design space where small mass reductions cause disproportionate losses in thrust. Such analyses would provide a clearer basis for deciding how strongly structural weight should be prioritised in different application scenarios.

The surrogate model can also be refined by training it on higher-fidelity CFD data, for example, viscous RANS or hybrid RANS–LES simulations, covering the full (0°–360°) wind-angle range. Extending the database in this way would improve the reliability of performance predictions under off-design inflow conditions, including strongly separated flows and highly three-dimensional wake structures. It would allow the model to capture directional asymmetries that are invisible in a limited-angle training set. At the same time, generating such a database is substantially more expensive than the current low-fidelity simulations, so there is a need for sparse sampling and possibly adaptive refinement driven by surrogate error indicators. The goal is to identify a sampling strategy that delivers the most informative high-fidelity data given limited computational resources.

Finally, the thrust production trends reported in Paper II should be reexamined using high-fidelity viscous simulations. Such simulations would resolve boundary-layer development, separation and reattachment, vortex shedding, and three-dimensional effects that are either absent or only approximately represented in the original study. Comparing thrust, power, and unsteady load histories across fidelities would clarify to what extent the main conclusions of Paper II remain valid once realistic viscous and turbulent effects are included. Any systematic discrepancies could then be analysed to refine the underlying physical interpretation and to calibrate correction factors. This step would provide a more reliable foundation for subsequent optimisation and FSI studies and help bridge the gap between idealised numerical analyses and practical design conditions.

Bibliography

BAR Technologies Ltd. (2025). *WindWings*. Accessed: 2025-12-10. URL: `https://www.bartechnologies.uk/commercial-ships/dnv-validates-energy-savings-of-windwings/`.

Blakeley, AW, RGJ Flay and PJ Richards (2012). "Design and optimisation of multi-element wing sails for multihull yachts". *18th Australasian Fluid Mechanics Conference, Launcestion, Australia*, pp. 3–7.

Drela, M. (1989). "XFOIL: an analysis and design system for low Reynolds number airfoils". *Conference on Low Reynolds Number Airfoil Aerodynamics, University of Notre Dame.* DOI: `10.1007/978-3-642-84010-4_1`.

Furukawa, Hiroyuki, Alexander W. Blakeley, Richard G. J. Flay and Peter J. Richards (2015). "Performance of wing sail with multi element by two-dimensional wind tunnel investigations". *Journal of Fluid Science and Technology* 10.2, pp. 15–00256. DOI: `10.1299/jfst.2015jfst0019`.

Giovannetti, Laura Marimon, U. Dhome, K. Malmek, A. Persson and C. Wielgosz (June 2022). "Multi-Wing Sails Interaction Effects". *SNAME 24th Chesapeake Sailing Yacht Symposium.* SNAME Chesapeake Sailing Yacht Symposium, D011S002R002. DOI: `10.5957/CSYS-2022-006`.

Guzelbulut, Cem, Timoteo Badalotti and Katsuyuki Suzuki (2024). "Optimization techniques for the design of crescent-shaped hard sails for wind-assisted ship propulsion". *Ocean Engineering* 312, p. 119142. ISSN: 0029-8018. DOI: `10.1016/j.oceaneng.2024.119142`.

Hebbal, Ali, Loïc Brevault, Mathieu Balesdent, El-Ghazali Talbi and Nouredine Melab (2021). "Bayesian optimization using deep Gaussian processes with applications to aerospace system design". *Optimization and Engineering* 22.1, pp. 321–361. ISSN: 1573-2924. DOI: `10.1007/s11081-020-09517-8`.

Hillenbrand, Antonia, Laura Marimon Giovannetti, Ulysse Dhomé and Jakob Kuttenkeuler (Oct. 2024). "Wind Tunnel Tests of a Two-Element Wingsail with Focus on Near-Stall Aerodynamics". *Journal of Sailing Technology* 9.01, pp. 110–127. ISSN: 2475-370X. DOI: `10.5957/jst/2024.9.1.110`.

Hussain, Md Daluar and Osman Md Amin (2021). "A Comprehensive Analysis of the Stability and Powering Performances of a Hard Sail–Assisted Bulk Carrier". *Journal of Marine Science and Application* 20.3, pp. 426–445. ISSN: 1993-5048. DOI: `10.1007/s11804-021-00219-w`.

Jacobs, Eastman N., Kenneth E. Ward and Robert M. Pinkerton (1933). *The characteristics of 78 related airfoil sections from tests in the variable-*

density wind tunnel. Research report 460. National Advisory Committee for Aeronautics.

Jo, Yeongmin, Hakjin Lee, Seongim Choi, Jongoh Kwon and Sungmok Ahn (June 2013). "Aerodynamic Design Optimization of Wing-sails", pp. 2013–2524. DOI: `10.2514/6.2013-2524`.

Kimball, John (2009). *Physics of Sailing*. CRC Press. ISBN: 978-1-4200-5964-4.

Kraft, Dieter (1988). *A software package for sequential quadratic programming*. Köln. URL: `https://www.tib.eu/de/suchen/id/TIBKAT%3A016896521`.

Li, Chen, Hongming Wang and Peiting Sun (2020). "Numerical Investigation of a Two-Element Wingsail for Ship Auxiliary Propulsion". *Journal of Marine Science and Engineering* 8.5. ISSN: 2077-1312. DOI: `10.3390/jmse8050333`.

Ma, Yong, Huaxiong Bi, Mengqi Hu, Yuanzhou Zheng and Langxiong Gan (2019). "Hard sail optimization and energy efficiency enhancement for sail-assisted vessel". *Ocean Engineering* 173, pp. 687–699. ISSN: 0029-8018. DOI: `10.1016/j.oceaneng.2019.01.026`.

Makram, Theodoros Evangelos, Pericles Panagiotou and Dimitrios Mattheou (2023). "Wingsail layout design and shape optimization using a CFD-aided Taguchi approach: The Aegean Marathon case study". *Ocean Engineering* 276, p. 114055. ISSN: 0029-8018. DOI: `10.1016/j.oceaneng.2023.114055`.

Malmek, Karolina (2023). "Rapid Aerodynamic Method for Interacting Sails". PhD thesis. Gothenburg, Sweden: Chalmers University of Technology.

Malmek, Karolina, Lars Larsson, Sofia Werner, Jonas W. Ringsberg, Rickard Bensow and Christian Finnsgård (2024). "Rapid aerodynamic method for predicting the performance of interacting wing sails". *Ocean Engineering* 293, p. 116596. ISSN: 0029-8018. DOI: `10.1016/j.oceaneng.2023.116596`.

Mason, William H. (2018). *Configuration Aerodynamics*. Blackburg VA: Virginia Tech.

Mitsui O.S.K. Lines, Ltd. (2025). *Wind Challenger - Energy Saving Technologies*. Accessed: 2025-11-05. URL: `https://www.mol-service.com/en/services/energy-saving-technologies/wind-challenger`.

Ouchi, K., K. Shima and K. Kimura (2023). ""WIND HUNTER"- the zero emission cargo ship powered by wind and hydrogen energy". *Conference Proceedings The Japan Society of Naval Architects and Ocean Engineers 37*.

Ouchi, Kazuyuki, Kiyoshi Uzawa, Akihiro Kanai and Masahiko Katori (2016). "Wind Challenger: The Next Generation Hybrid Sailing Vessel". *Smatech Symposium*. Glasgow, UK.

Rasmussen, Carl Edward (2004). "Gaussian Processes in Machine Learning". *Advanced Lectures on Machine Learning: ML Summer Schools 2003, Canberra, Australia, February 2 - 14, 2003, Tübingen, Germany, August 4 - 16, 2003, Revised Lectures*. Ed. by Olivier Bousquet, Ulrike von Luxburg and Gunnar Rätsch. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 63–71. ISBN: 978-3-540-28650-9. DOI: `10.1007/978-3-540-28650-9_4`.

Rasmussen, Carl Edward and Christopher K. I. Williams (2006). *Gaussian Processes for Machine Learning*. The MIT Press.

Rogers, T. J., J. Mclean, E. J. Cross and K. Worden (2023). "Gaussian Processes". *Machine Learning in Modeling and Simulation: Methods and Applications*. Ed. by Timon Rabczuk and Klaus-Jürgen. Bathe. Cham:

Springer International Publishing, pp. 121–147. ISBN: 978-3-031-36644-4. DOI: `10.1007/978-3-031-36644-4_3`.

ScandiNAOS AB (2025). *WindStruc.* Accessed: 2025-12-10. URL: `https://www.scandinaos.com/projects.html#heading_windstruc`.

Shan, Yunan, Shengqiang Zhao, Fangyu Peng, Rong Yan, Xiaowei Tang, Juntong Su and Hao Sun (2025). "Sparse Representation of Robotic Machining Deformation Based on Key Points Determination". *Intelligent Robotics and Applications.* Ed. by Xuguang Lan, Xuesong Mei, Caigui Jiang, Fei Zhao and Zhiqiang Tian. Singapore: Springer Nature Singapore, pp. 3–18. ISBN: 978-981-96-0774-7. DOI: `10.1007/978-981-96-0774-7_1`.

Sharpe, Peter (2023). *NeuralFoil: An airfoil aerodynamics analysis tool using physics-informed machine learning.* `https://github.com/peterdsharpe/NeuralFoil`.

Sharpe, Peter and R. John Hansman (2025). "NeuralFoil: An Airfoil Aerodynamics Analysis Tool Using Physics-Informed Machine Learning". *arXiv.* DOI: `10.48550/arXiv.2503.16323`.

Sharpe, Peter D. (2024). "Accelerating Practical Engineering Design Optimization with Computational Graph Transformations". PhD thesis. Massachusetts Institute of Technology.

Snoek, Jasper, Hugo Larochelle and Ryan P Adams (2012). "Practical Bayesian Optimization of Machine Learning Algorithms". *Advances in Neural Information Processing Systems.* Ed. by F. Pereira, C.J. Burges, L. Bottou and K.Q. Weinberger. Vol. 25. Curran Associates, Inc.

Truong, V. K. (May 2020). "An analytical model for airfoil aerodynamic characteristics over the entire 360° angle of attack range". *Journal of Renewable and Sustainable Energy* 12.3, p. 033303. ISSN: 1941-7012. DOI: `10.1063/1.5126055`.

van Reen, Stephan, Berken Serbülent and Hua-Dong Yao (2025). "Machine learning-based multipoint optimisation for improving aerodynamics of symmetrically cambered wing sails in wind-assisted ship propulsion". *Ocean Engineering* 342, p. 122829. ISSN: 0029-8018. DOI: `https://doi.org/10.1016/j.oceaneng.2025.122829`. URL: `https://www.sciencedirect.com/science/article/pii/S0029801825025120`.

Venkatesh, A., J. Niu, X. Xue, Z.-W. Chen and H.-D. Yao (2024). "Coexistence of passive vortex-induced vibrations and active pitch oscillation triggered by a square cylinder attached with a deformable splitter plate". *Physics of Fluids* 36.4, p. 043612. DOI: `10.1063/5.0200566`.

Virtanen, Pauli, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt and SciPy 1.0 Contributors (2020). "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. *Nature Methods* 17, pp. 261–272. DOI: `10.1038/s41592-019-0686-2`.

Wallenius Marine AB (2025). *The Oceanbird.* Accessed: 2025-11-23. URL: `https://www.theoceanbird.com/`.

Wang, Xilu, Yaochu Jin, Sebastian Schmitt and Markus Olhofer (July 2023). “Recent Advances in Bayesian Optimization”. *ACM Computing Surveys* 55.13s. ISSN: 0360-0300. DOI: `10.1145/3582078`.

Yao, H.-D. (2025). *Hybrid Parametrization of Symmetrically Cambered (Crescent-Shaped) Airfoil Profiles for Rigid Wingsail Design in Wind-Assisted Ship Propulsion.* Tech. rep. TR-2025-1. Gothenburg, Sweden: Chalmers University of Technology. DOI: `10.63959/m2.techreport/2025.1`.

Yasuda, Akane, Tomoki Taniguchi and Toru Katayama (2024). “Numerical Investigation of Aerodynamic Interactions between Rigid Sails Attached to Ship”. *Journal of Marine Science and Engineering* 12.8, p. 1425. ISSN: 2077-1312. DOI: `10.3390/jmse12081425`.

Zhu, Heng, Valery Chernoray, Hua-Dong Yao, Jonas W. Ringsberg and Bengt Ramne (2023b). “Experimental study on structure responses of triple wing sails to turbulence flows at multiple apparent wind angles”. *Advances in the Analysis and Design of Marine Structures.* CRC Press, pp. 781–787.

Zhu, Heng, Hua-Dong Yao and Jonas W. Ringsberg (2024). “Unsteady RANS and IDDES studies on a telescopic crescent-shaped wingsail”. *Ships and Offshore Structures* 19.1, pp. 134–147. DOI: `10.1080/17445302.2023.2256601`.

Zhu, Heng, Hua-Dong Yao, Fabian Thies, Jonas Ringsberg and Bengt Ramne (July 2023a). “Propulsive performance of a rigid wingsail with crescent-shaped profiles”. *Ocean Engineering* 285.2, p. 115349. DOI: `10.1016/j.oceaneng.2023.115349`.