



Regime-Dependent Thrust-Driven Aerodynamic Optimization of a Two-Element Wingsail with a Bayesian–URANS Approach

Downloaded from: <https://research.chalmers.se>, 2026-07-03 03:02 UTC

Citation for the original published paper (version of record):


Yao, H. (2026). Regime-Dependent Thrust-Driven Aerodynamic Optimization of a Two-Element Wingsail with a Bayesian–URANS Approach. *Ocean Engineering*, 362.
<http://dx.doi.org/10.1016/j.oceaneng.2026.126337>

N.B. When citing this work, cite the original published paper.



Research paper

Regime-dependent thrust-driven aerodynamic optimization of a two-element wingsail with a Bayesian–URANS approach

Hua-Dong Yao 

Division of Marine Technology, Department of Mechanical Engineering, Chalmers University of Technology, Gothenburg, 41296, Sweden

ARTICLE INFO

Keywords:

Regime-dependent thrust
Aerodynamic optimization
Two-element wingsail
Bayesian optimization
URANS
Wind-assisted ship propulsion

ABSTRACT

Despite increasing research on rigid wingsail aerodynamics for wind-assisted ship propulsion, optimization of two-element configurations remains insufficiently explored, particularly under a wide range of apparent wind angles (AWA) and considering dynamic interactions between the elements. This study establishes a geometrically interpretable parameterization and develops an optimization framework using a regime-dependent thrust formulation that enables consistent optimization across AWA regimes. The parameterization comprises inter-element gap, main element thickness, flap thickness, flap-to-main chord ratio, flap hinge position, and flap deflection. Three-dimensional unsteady Reynolds-averaged Navier–Stokes equations, coupled with Bayesian optimization, are used. The thrust objective is defined through weighted averaging over AWA intervals, under fixed main-element angle of attack of 12° and Reynolds number of 5.7×10^6 . A full AWA range results in lift-dominated optimization. However, when emphasizing regimes to upwind or downwind (AWA below or above 90°), optimal configurations are different due to increased drag contribution to thrust. This infers the critical role of the flap deflection in drag production. Sensitivity analysis identifies the hinge position and gap as dominant parameters across regimes, while flap-related parameters are secondary and regime-dependent. This study demonstrates the importance of regime-dependence in wingsail optimization and provide physical insight into the transition between lift-dominated and drag-assisted propulsion mechanisms.

1. Introduction

Reducing greenhouse gas (GHG) emissions from maritime transport has stimulated growing interest in wind-assisted ship propulsion (WASP), which can complement conventional propulsion and support the decarbonization targets set by the International Maritime Organization (IMO) (Khan et al., 2021). Several WASP concepts have been proposed, including Flettner rotors, towing kites, soft sails, and rigid wingsails. Among these technologies, rigid wingsails have attracted attention because their airfoil-like geometry enables high lift generation while their structural rigidity reduces deformation, leading to relatively stable aerodynamic performance (Khan et al., 2021; Zhu et al., 2023; von Klemperer et al., 2023).

To enhance lift generation, two main aerodynamic strategies have been explored in the design of wingsail cross-sectional profiles. The first strategy introduces camber in a single-element profile, exemplified by so-called crescent-shaped or symmetrically cambered wingsails. This configuration has been investigated extensively through numerical and experimental studies (Zhu et al., 2023; van Reen et al., 2025), which have established its aerodynamic characteristics and propulsive performance. The second strategy adopts high-lift concepts widely used

in aircraft aerodynamics, such as flaps or slats, by employing multi-element configurations. For instance, a two-element wingsail consists of a main wing and a deflectable trailing flap separated with a gap (also termed slot) (Hillenbrand et al., 2024). The flap deflection introduces an additional degree of freedom that can increase the attainable lift coefficient and extend the operable range of apparent wind angles (AWAs). These characteristics make multi-element wingsails a promising concept for the WASP of commercial vessels. Additionally, Makram et al. (2023) investigated a specific configuration in which the flap merged with the main wing so that there is no gap between the two elements.

The aerodynamic advantage of the two-element configuration arises from the interaction between the main wing and the flap through the gap, a mechanism extensively studied in aeronautical high-lift systems. Classical analyses identified several physical effects through which the gap enhances lift, including circulation augmentation and boundary-layer re-energization (Smith, 1975). Subsequent studies showed that the aerodynamic performance of multi-element systems is highly sensitive to key gap geometry parameters, namely the width, overlap,

E-mail address: huadong.yao@chalmers.se.

<https://doi.org/10.1016/j.oceaneng.2026.126337>

Received 5 April 2026; Received in revised form 20 May 2026; Accepted 30 May 2026

Available online 11 June 2026

0029-8018/© 2026 The Author. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

and flap deflection angle δ (van Dam, 2002). Recent CFD investigations further indicated that the interactions among these parameters are strongly nonlinear, which highlights the importance of systematic exploration of the design space when analyzing two-element wingsails (Cerantola et al., 2025; Negoita et al., 2025).

Systematic aerodynamic investigations of two-element wingsails were initially motivated by their application in competitive sailing. The 34th America's Cup in 2013 introduced the AC72 class catamaran equipped with a two-element wingsail of approximately 40 m span, which stimulated a dedicated research programme at ISAE-Supéro in Toulouse. Using unsteady Reynolds-averaged Navier–Stokes (URANS) equations and large eddy simulation (LES), Chapin et al. (2015) analyzed an AC72-type configuration and identified laminar separation bubbles on both elements, together with the nonlinear joint effects of the flap deflection, gap width, and flap thickness governing stall onset. Subsequent wind-tunnel validation confirmed that URANS captures the main features of the gap-jet flow and aerodynamic loads (Fiumara et al., 2016). Further LES studies under gust excitation revealed transient leading-edge vortices and separation bubbles on the flap surface (Fiumara et al., 2017). Through direct comparison to LES, URANS simulations reproduced the primary unsteady loading behavior at significantly lower computational cost and therefore, are suitable for parametric wingsail studies (Fiumara et al., 2018).

The importance of unsteady flow physics becomes particularly evident near stall, where rapidly changing aerodynamic conditions place significant demands on automated trimming systems tasked with maintaining stable wingsail operation. Hillenbrand et al. (2024) carried out wind-tunnel experiments for a rectangular two-element wingsail representative of the Oceanbird WASP system, which consists of a NACA 0025 main element and NACA 0015 flap. They revealed a characteristic two-stage stall process. Stall first develops on the flap, followed by full-wing separation accompanied by strong hysteresis. This sequential stall mechanism reflects the progressive decoupling of the boundary layers on the main element and flap, and is intrinsically unsteady.

Several studies have examined two-element wingsails under operating conditions representative of merchant vessels. Using steady and unsteady RANS with the $k-\omega$ SST model, Li et al. (2020) performed a parametric study of camber, flap thickness, angle of attack, and hinge location, showing that positioning the hinge rotation axis at 85%–95% of the main-element chord can induce a sudden lift increase near stall by suppressing large-scale separation on the flap suction side, which indicates a strong nonlinear coupling between the camber and hinge position. The influence of atmospheric boundary-layer inflow was later investigated by Li et al. (2023), who compared uniform and logarithmic wind profiles. Their results showed that the vertical velocity gradient delays stall onset but leads to a more abrupt spanwise spread of separation once stall occurs. This effect implies that small changes in flap setting near the critical incidence can trigger deep stall. More recently, Li et al. (2025) explored bio-inspired leading-edge bulges inspired by humpback-whale tubercles using URANS and Improved Delayed Detached Eddy Simulation (IDDES) under gradient wind conditions. Their study demonstrated that an optimized configuration can reduce the post-stall lift drop, and that a passive flow-control strategy can widen the operational envelope of wingsails.

Despite the growing body of research reviewed above, the aerodynamic optimization of two-element wingsails remains an open topic. In particular, to the authors' knowledge, gapped two-element configurations have not been systematically investigated within a regime-dependent optimization framework for wind-assisted ship propulsion, for which the optimization objective differs fundamentally from conventional aircraft applications. Previous work has mainly examined single-element wingsails (Yao, 2025; van Reen et al., 2025). Although the integrated wing-flap configuration has been optimized by Makram et al. (2023), this configuration does not include a gap between the elements and therefore cannot represent the lift-enhancement mechanisms

associated with the gap. As a result, its applicability to two-element wingsails in general is limited.

The present study aims to address this gap by introducing a regime-dependent thrust formulation based on a weighted average of the thrust coefficient over prescribed apparent wind angle (AWA) intervals. This formulation enables a consistent evaluation of aerodynamic performance across upwind, beam-reach, and downwind conditions, while explicitly accounting for the coupled and regime-dependent contributions of lift and drag to propulsion. A Bayesian optimization framework coupled with three-dimensional URANS simulations is employed to explore a six-dimensional geometric design space, parameterized using a minimal set of geometrically interpretable variables derived from a NACA0012 reference geometry. This approach allows systematic investigation of two-element wingsail configurations while resolving the nonlinear aerodynamic interactions governing lift enhancement, drag production, and separation onset. The potential dependence of optimal designs on AWA ranges will be examined. The sensitivity of the geometric parameters will be systematically analyzed across different AWA regimes in relation to the regime-dependent role of drag in thrust generation.

The remainder of the paper is organized as follows. Section 2 presents the geometric parametrization of the two-element wingsail, followed by Section 3, which introduces the optimization framework, including the regime-dependent thrust formulation. Section 4 describes the computational methodology of URANS along with numerical verifications. Section 5 presents the optimization results and discussion, including the baseline lift-dominated case and the regime-dependent analysis across AWA intervals. Finally, Section 6 summarizes the main findings and outlines directions for future work.

2. Parametrization of two-element wingsails

2.1. Parameterization method

The two-element wingsail configuration consists of a main element (foil 1) and a trailing flap (foil 2), as illustrated in Fig. 1. Both elements are obtained by scaling a reference NACA 0012 airfoil, with a chord length c_0 and a maximum thickness Y_0 . Hereafter c_0 is taken as the reference length for defining all geometric parameters and for nondimensionalizing the flow quantities.

The thickness distribution of a NACA 4-digit airfoil is expressed as (Mason, 2018):

$$\frac{y_t(x)}{c} = 5 \left(\frac{Y}{c} \right) \left[a_0 \sqrt{\frac{x}{c}} + a_1 \left(\frac{x}{c} \right) + a_2 \left(\frac{x}{c} \right)^2 + a_3 \left(\frac{x}{c} \right)^3 \right], \quad 0 \leq \frac{x}{c} \leq 1 \quad (1)$$

where Y is the maximum thickness. The coefficients are specified as:

$$a_0 = 0.2969, \quad a_1 = -0.1260, \quad a_2 = -0.3516, \quad a_3 = 0.2843, \quad a_4 = -0.1036 \text{ or } -0.1015. \quad (2)$$

In the present work, $a_4 = -0.1036$ is chosen to construct a closed trailing edge (i.e., $y_t = 0$ at $x/c = 1$), whereas -0.1015 is disregarded because it results in a finite thickness at this position.

In the NACA 4-digit series, the thickness distribution function is fixed, since the polynomial coefficients are not adjustable. This means the shape of the thickness distribution is the same for all NACA 4-digit airfoils, only scaled by the maximum thickness parameter Y .

Accordingly, the geometric parameters of the two elements are defined as:

$$c_1 = k_{c1} c_0, \quad c_2 = k_{c2} c_0, \quad Y_1 = k_{y1} Y_0, \quad Y_2 = k_{y2} Y_0, \quad (3)$$

where c_1 and c_2 are the chord lengths of the main element and flap, and Y_1 and Y_2 are the maximum thicknesses. The scaling factors k_{c1} , k_{c2} , k_{y1} and k_{y2} are defined relative to the reference profile characterized by the reference lengths (c_0 , Y_0).




Fig. 1. Parameterization of the two-element wingsail configuration.

The angular orientation of the configuration is prescribed by two parameters. The geometric angle of attack of the main element, α , is defined as the angle between its chord line and the freestream direction. The flap deflection angle, β , is defined as the angle between the flap chord line and that of the main element, with the center of rotation (CoR, i.e., the hinge point) located at X_{CoR} . A positive β corresponds to a downward deflection of the flap, toward the pressure side of the main element. The hinge point location is measured from the leading edge of the main element as:

$$X_{\text{CoR}} = p_1 c_1, \quad (4)$$

where $p_1 \in [0, 1]$ denotes the fractional chordwise position.

The gap between the trailing edge of the main element and the leading edge of the flap is denoted by g_{12} and is defined when the chord lines of the two elements are aligned. Thus, the distance from the flap leading edge to its CoR is $c_1 - X_{\text{CoR}} + g_{12}$.

Since the flap chord length c_2 is a variable parameter, the reference length used in the calculation of the force coefficients is defined as $c_{\text{tot}} = c_1 + c_2$, ensuring consistent aerodynamic evaluation across configurations.

2.2. Operating conditions

The flow is characterized by a chord-based Reynolds number, $Re_c \approx 5.7 \times 10^6$, defined with respect to the chord length of the main element. The main-element chord is fixed as $c_1 = c_0$ (i.e., $k_{c1} = 1.0$) for the convenience of sampling and optimization. The Reynolds number corresponds to a full-scale configuration with $c_{\text{tot}} = 14$ m at a freestream velocity of $U_\infty \approx 6.4$ m/s, assuming the air properties such as the density $\rho = 1.184$ kg/m³ and the dynamic viscosity $\mu = 1.855 \times 10^{-5}$ Pa · s. The conditions are consistent with those reported in previous studies on crescent-shaped wingsails (Zhu et al., 2023).

3. Optimization framework

The aerodynamic shape optimization of the two-element wingsail is formulated as a constrained single-objective maximization problem over a six-dimensional design space. Since each candidate evaluation requires a full computational fluid dynamics (CFD) simulation using URANS at a cost of 90–160 core hours (see details in Section 4.2), Bayesian optimization (BO) is adopted following the efficient global optimization (EGO) paradigm (Jones et al., 1998). The present optimization problem involves a computationally expensive black-box objective function with moderate design dimensionality and no analytically available gradients. Under these conditions, BO provides an efficient surrogate-based framework for reducing the number of required CFD evaluations while systematically exploring the design space. The URANS method is presented in Section 4.

The overall procedure is summarized in Fig. 2 and organized into three phases: initialization (Phase 1), the sequential BO loop (Phase 2),

Table 1

Design variables and their prescribed bounds.

Symbol	Description	Lower	Upper	Unit
X_{CoR}	Centre of rotation along main element chord	0.50	0.95	–
g_{12}	Gap spacing (main element to flap)	0.005	0.08	–
k_{y1}	Main element thickness scale ratio	0.833	1.667	–
k_{y2}	Flap thickness scale ratio	0.667	2.000	–
k_{c2}	Flap chord scale ratio	0.40	0.80	–
β	Flap deflection angle	10	25	°

and post-processing (Phase 3). In Phase 2, a GP surrogate is first fitted and the acquisition function is minimized to propose a candidate design \mathbf{x}^* . Geometric feasibility is then evaluated prior to any CFD simulation. If the design is infeasible, a penalized observation is recorded, and the workflow returns to the proposal step without performing the CFD. For geometrically feasible designs, a URANS simulation is carried out. The resulting performance is subsequently checked against aerodynamic constraints. Designs that violate these constraints are excluded from the GP training, whereas feasible evaluations are used to update the surrogate. The iteration proceeds with checkpoints saved periodically, until convergence is reached. In Phase 3, the best design is returned. And convergence diagnostics and reports are generated, followed by saving the final checkpoint state.

3.1. Optimization objective based on AWA regime partitioning

The optimization objective is formulated as a weighted aggregation of time-averaged thrust performance over N prescribed AWA conditions or intervals,

$$\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) = \sum_{i=1}^N w_i \overline{C_T}^{(i)}(\mathbf{x}), \quad (5)$$

where $\mathbf{x} \in \mathbb{R}^6$ denotes the geometric design vector (see Table 1), \mathcal{X} is the box-constrained feasible design space, and w_i are prescribed weights satisfying $\sum_{i=1}^N w_i = 1$. This formulation enables a regime-dependent evaluation in which different AWA ranges contribute with varying importance.

The force decomposition based on the AWA and AOA is illustrated in Fig. 3. The local operating AOA is defined in the wingsail frame, whereas the AWA is defined in the ship frame. In ship operations, as the AWA varies, the rigid wingsail is assumed to rotate relative to the ship such that its AOA remains unchanged, thereby maintaining robust aerodynamics. The AWA is applied to projecting the combined lift and drag forces onto the ship-forward direction, to construct the thrust as the optimization objective.

The thrust coefficient at a given AWA is defined as the streamwise projection of the aerodynamic force coefficients,

$$\overline{C_T}(\theta) = \overline{C_L} \sin \theta - \overline{C_D} \cos \theta, \quad (6)$$

where θ is the AWA, and $\overline{C_L}$ and $\overline{C_D}$ are the time-averaged lift and drag coefficients. This definition indicates that the relative contributions of the lift and drag to the thrust vary with the AWA, and that the contribution of drag may change sign depending on the angular range considered.

For an AWA interval $[\theta_1, \theta_2]$, the aggregated thrust coefficient is defined as the normalized angular average,

$$\overline{C_T}^{(i)} = \frac{1}{\Delta\theta} \int_{\theta_1}^{\theta_2} \overline{C_T}(\theta) d\theta, \quad (7)$$

where $\Delta\theta = \theta_2 - \theta_1$ ensuring consistency with respect to the interval size.

Assuming that $\overline{C_L}$ and $\overline{C_D}$ are approximately constant within each interval, Eq. (7) reduces to

$$\overline{C_T}^{(i)} = w_L \overline{C_L} - w_D \overline{C_D}, \quad (8)$$

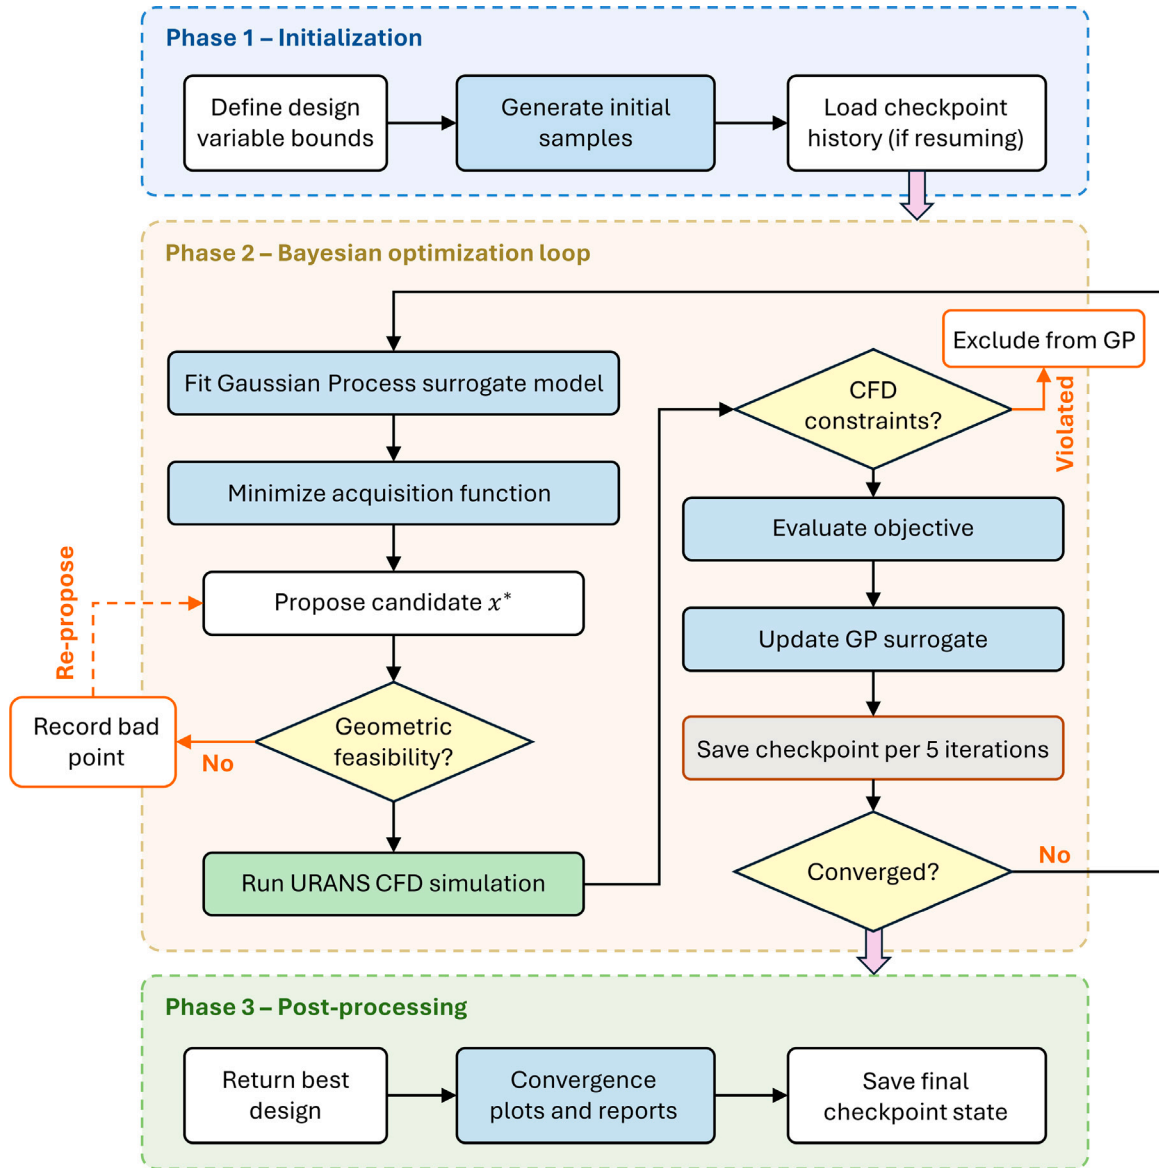


Fig. 2. The flowchart of the Bayesian optimization procedure.




Fig. 3. Schematic of the AWA, AOA, wind directions (gray arrows), and aerodynamic force decomposition.

where

$$w_L = \frac{\cos \theta_1 - \cos \theta_2}{\Delta \theta}, \quad w_D = \frac{\sin \theta_2 - \sin \theta_1}{\Delta \theta}. \quad (9)$$

The coefficients w_L and w_D depend only on the AWA bounds and characterize the regime-dependent roles of the lift and drag. In particular, w_D changes sign across AWA regimes, such that the drag acts as a penalty for $\theta < 90^\circ$ and contributes positively to the thrust for $\theta > 90^\circ$. This behavior is expected to give rise to qualitatively different optima across regimes, as demonstrated in Section 5.2.

3.2. Physical and geometric constraints

The main-element angle of attack is fixed at $\alpha = 12^\circ$. At this incidence, the flow remains attached over the primary lifting surface, providing a margin relative to stall onset. This choice reflects a compromise between aerodynamic performance and operational robustness, to ensure that moderate perturbations in inflow conditions, such as wind gusts or small variations in effective angles of attack, do not trigger large-scale separation.

To exclude non-physical CFD solutions, the following aerodynamic constraints are imposed:

$$\bar{C}_L \geq 0.50, \quad \bar{C}_D \leq 0.25, \quad C_L^{\text{rms}} \leq 0.1, \quad C_D^{\text{rms}} \leq 0.01, \quad \frac{A_{\text{sep}}}{A_{\text{foil2}}} \leq 0.50 \quad (10)$$

where C_L^{rms} and C_D^{rms} are the root-mean-square values for the fluctuations of the lift and drag coefficients, and $A_{\text{sep}}/A_{\text{foil2}}$ is the normalized projected separation area on the flap.

The reasons for only considering the flap separation are two folds. Flows passing the main element, which is shaped in symmetric NACA 4-digit profiles, remains largely attached at the fixed AOA of 12° . Moreover, since the flap angle is deflectable in the optimization, it is more risky to be subjected to stall, which is expected to control for robust aerodynamics. The aerodynamic constraints are not arbitrary values, but engineering filtering criteria informed by preliminary parameter-space exploration. During the initial exploration, configurations violating these constraints were frequently associated with strongly separated flow ($A_{\text{sep}}/A_{\text{foil2}}$), large force fluctuations (root-mean-square values), or poor propulsive performance (C_L and C_D). The constraints are therefore introduced to exclude non-physical or impractical operating conditions from the optimization space.

The constraints for the design variables in the geometric parameterization are listed in Table 1. The lower and upper bounds are chosen to span practical two-element wingsail proportions while excluding configurations that are geometrically infeasible. Specifically, the hinge position $X_{\text{CoR}} \in [0.50, 0.95]$ excludes configurations that would cause mechanical interference between the flap structure and the mast. The gap distance is bounded within a positive range $g_{12} \in [0.005, 0.08]$, and therefore the overlapping of the two elements is excluded by construction within the admissible design space. The thickness ratios $k_{y1} \in [0.833, 1.667]$ and $k_{y2} \in [0.667, 2.000]$ correspond approximately to NACA-0010 to NACA-0020 thickness ranges for both elements. The flap chord ratio $k_{c2} \in [0.40, 0.80]$ avoids excessively small and large flaps in reference to the main element, while the flap deflection angle $\beta \in [10^\circ, 25^\circ]$ covers the typical operational envelope below the onset of flap stall. Moreover, to prevent disproportionate thickness ratios between the two elements, a geometric feasibility constraint is imposed such that

$$0.2 \leq Y_2/Y_1 \leq 1.0. \quad (11)$$

3.3. Constraint handling and convergence

The constraints are enforced in two sequential stages. Prior to each simulation, the geometric constraint $0.2 \leq Y_2/Y_1 \leq 1.0$ is evaluated. Designs that violate this condition are rejected, and a penalized

observation is defined as

$$p(\mathbf{x}) = \frac{0.05}{1 + \gamma d(\mathbf{x})^2}, \quad (12)$$

where $d(\mathbf{x})$ denotes the normalized distance to the nearest feasible boundary in the Y_2/Y_0 -dimension. The adaptive scaling factor is given by $\gamma = \max(f^* - f_{\min}, 1)/f^*$, where f^* and f_{\min} represent the current best and worst observed objective values with $\gamma = 1$ at initialization. This formulation introduces a repulsive effect that discourages sampling in infeasible regions. A maximum of 50 rejection-resampling attempts is allowed per iteration, after which the run is aborted.

Following each successful simulation, designs that violate any aerodynamic constraint in Eq. (10) are excluded from the GP training set. The optimization is terminated when a maximum budget of $N = 160$ function evaluations is reached, or when the maximum relative improvement over a sliding window of $W = 5$ iterations satisfies

$$\max_{1 \leq j \leq W-1} \left| \frac{f_{i-j+1} - f_{i-j}}{f_{i-j}} \right| < \epsilon, \quad \epsilon = 10^{-5}. \quad (13)$$

Robustness in long-duration HPC runs is maintained by retrying failed simulations up to three times using an exponential back-off strategy, by checkpointing the optimizer state every five iterations, and by persisting infeasible-design records across restarts.

3.4. Surrogate model and acquisition

BO constructs a Gaussian process (GP) surrogate (Rasmussen and Williams, 2006):

$$\hat{f}(\mathbf{x}) \sim \mathcal{GP}(\mu_0(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')), \quad (14)$$

where $\hat{f}(\mathbf{x})$ denotes the surrogate prediction of the objective, $\mu_0(\mathbf{x})$ is the prior mean function, and $k(\mathbf{x}, \mathbf{x}')$ is the covariance (kernel) function between design points \mathbf{x} and \mathbf{x}' . The kernel is composed of a Matérn term of $\nu_M = 5/2$ together with a Gaussian noise term to account for numerical variability in URANS simulations:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 k_{M5/2}(\mathbf{x}, \mathbf{x}') + \sigma_n^2 \delta(\mathbf{x}, \mathbf{x}'), \quad (15)$$

where σ_f^2 is the signal variance (constant kernel), $k_{M5/2}$ is the Matérn-5/2 covariance with isotropic length scale ℓ , σ_n^2 is the noise variance (white kernel), and δ is the Kronecker delta. The hyperparameter bounds are set as follows: $\sigma_f^2 \in [10^{-3}, 10^3]$, $\ell \in [0.1, 1.5]$, and $\sigma_n^2 \in [10^{-6}, 10^{-2}]$. The upper bound on ℓ is deliberately tightened from the library default to prevent the surrogate from converging to an over-smoothed plateau, which was observed in preliminary runs with the default setting. Hyperparameters are estimated by maximizing the log-marginal likelihood with five random restarts, and response normalization is applied for numerical stability.

The next evaluation point is selected by minimizing the lower confidence bound (LCB) (Srinivas et al., 2010):

$$\alpha_{\text{LCB}}(\mathbf{x}) = \mu(\mathbf{x}) - \kappa \sigma(\mathbf{x}), \quad (16)$$

where $\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$ are the GP posterior mean and standard deviation, respectively, and $\kappa = 3.0$ controls the trade-off between exploration and exploitation. The acquisition function is optimized by random sampling over \mathcal{X} , following the implementation in `scikit-optimize` (Head et al., 2020).

The surrogate is trained on the complete set of CFD evaluations accumulated during the Bayesian optimization campaign (approximately 155 cases, including both feasible and infeasible evaluations). No separate design-of-experiments phase is used, and the dataset is the direct product of the sequential optimization loop described in Section 3.3.

Prior to the first GP-guided iteration, an initial set of $N_0 = 10$ evaluations is generated to seed the optimizer. The sampling distributions are chosen to concentrate candidates away from the parameter bounds, where solver failures are more likely. Specifically, the hinge position X_{CoR} is drawn from a truncated normal distribution centred at

the midpoint of its range, with standard deviation equal to one-third of the range. The flap chord ratio k_{c2} and the flap deflection angle β are drawn from a Beta(2,2) distribution mapped to their respective bounds. The remaining shape and thickness parameters (k_{y1} , k_{y2}) are drawn from an 80/20 mixture: 80% of samples are taken from the interior of each parameter range (excluding a 10% boundary margin at each bound), and 20% from the full range. This strategy reduces early infeasible evaluations caused by geometric constraint violations near the parameter boundaries.

During the optimization, additional random candidates are periodically injected to enhance exploration, with higher frequency in the early iterations and reduced frequency thereafter. Additional injections are triggered after repeated infeasible proposals. These candidates are sampled from a 50/50 mixture of a Gaussian distribution centred on the current best design, with per-variable standard deviation $\sigma_{s,k} = (x_k^{\text{ub}} - x_k^{\text{lb}})/4$, and a uniform distribution over \mathcal{X} , where x_k^{ub} and x_k^{lb} are the upper and lower bounds of the k th design variable. Duplicate candidates are removed within numerical tolerance.

The five AWA-regime cases defined in Section 5.2 are evaluated as independent single-run optimizations: each case loads the same accumulated CFD dataset, trains a dedicated GP surrogate, and applies the acquisition function with the objective weights w_L and w_D corresponding to the respective AWA interval. No transfer of surrogate state occurs between the cases.

4. CFD method of URANS

4.1. Turbulence modeling

The flow around the two-element wingsail is considered to be incompressible. The turbulence modeling adopts URANS. Turbulence closure is provided by the shear-stress transport (SST) k - ω model proposed by Menter (1994). The model combines the near-wall accuracy of the k - ω formulation with the free-stream robustness of a k - ϵ -type behavior through a blending function, enabling reliable predictions across both boundary-layer and outer-flow regions.

Several model extensions are activated to improve predictive capability for the separated-flow conditions considered in the present study. A curvature correction following the approach of Spalart and Shur (1997), as adapted for the SST model by Smirnov and Menter (2009), is employed to account for the influence of streamline curvature on turbulence production. In addition, the Durbin realizability limiter (Durbin, 1996) is applied to bound the turbulent time scale and enforce realizability constraints, thereby preventing the spurious build-up of turbulent kinetic energy in stagnation regions. Low-Reynolds-number damping is also included to extend the validity of the model into the viscous sublayer.

An all- y^+ wall treatment is adopted, which blends viscous-sublayer and logarithmic-layer formulations and allows the turbulence model to adapt automatically to the local near-wall mesh resolution.

The suitability of URANS has been supported by a series of investigations of symmetrically cambered wingsails (Zhu et al., 2023, 2024), in which both two- and three-dimensional simulations were validated against experimental data at model scale. These studies further showed that URANS reproduces similarly accurate predictions of time-averaged lift and drag forces compared with hybrid approaches such as IDDES, although the latter can resolve large-scale vortex shedding and deep-stall separation structures beyond the capabilities of URANS alone. The scaling effects from model to full scale for this wingsail configuration were subsequently quantified by Zhu et al. (2025). Collectively, these results indicate that URANS retains sensitivity to the dominant aerodynamic mechanisms governing wingsail performance.

4.2. Numerical schemes and settings

The governing equations are solved using a segregated pressure-velocity coupling based on the semi-implicit method for pressure-linked equations (SIMPLE). Second-order upwind schemes are applied to the convection terms in the momentum and turbulence-transport equations. Gradients are reconstructed using the Venkatakrishnan limiter to suppress spurious oscillations while retaining second-order accuracy, and cross-diffusion gradients are included.

Time advancement uses a second-order implicit scheme with a time step controlling the Courant–Friedrichs–Lewy (CFL) number of the freestream flow to $\text{CFL}_\infty = 0.9 \times 10^{-2}$ and locally $\text{CFL}_{\text{max}} < 1.0$. Up to seven inner iterations are performed per time step. Under-relaxation factors are set to 0.8 for velocity, 0.2 for pressure. All linear systems are solved using an algebraic multigrid (AMG) method.

The statistics are collected after the flow reaches a statistically stationary state. Convergence is assessed based on the time-averaged lift and drag coefficients, which are evaluated over 11 convective flow-through periods. Here one convective period is defined as the time required for the freestream to traverse a distance equal to the wingsail chord length. The solution is considered converged when the variation of these time-averaged coefficients remains within 2.0% over the most recent 4.5 periods. Once this criterion is satisfied, the flow is regarded as converged, and the corresponding statistics are extracted. This approach is further exemplified in Section 4.3.1.

A velocity inlet is prescribed upstream to impose a uniform freestream velocity corresponding to $Re_c \approx 5.7 \times 10^6$. The top, bottom, and downstream boundaries are specified as pressure outlets, the spanwise boundaries as periodic, and the foil surfaces as no-slip walls.

The computational domain is defined as a rectangular region enclosing the wingsail. The upstream inlet boundary is located at a distance of 8 chord lengths from the wingsail, while the downstream outlet is placed 17 away. The top and bottom boundaries are each positioned at a distance of 8 from the wingsail, ensuring minimal blockage and confinement effects. The spanwise extent is determined based on a dedicated convergence study, as detailed in Section 4.3.2. Based on this assessment, a spanwise length of 0.2 is selected, which is sufficient to capture the relevant three-dimensional flow structures while maintaining computational efficiency.

Within the optimization loop, candidate geometries are regenerated and re-meshed automatically using a parametric meshing template, without manual intervention. The meshing procedure includes consistent surface refinement, prism-layer resolution near the walls with $y^+ \leq 1$, and local volumetric refinement in regions of strong flow interaction, including the gap region, leading/trailing edges, and wake region. Automatic mesh-quality checks on skewness, orthogonality, and aspect ratio are also performed before each simulation is launched.

All simulations are performed using the CFD solver Simcenter STAR-CCM+, version 20.04.008 (Siemens Digital Industries Software, 2025). Mesh generation, flow solution, and post-processing are carried out within the same integrated environment. A computational node consisting of two Intel Xeon Gold 6130 CPUs with a total of 32 cores is used for each simulation case. The typical computational cost per case is approximately 3–5 h of wall-clock time, depending on the degree of flow unsteadiness associated with phenomena such as flow separation. This corresponds to approximately 96–160 core-hours per case.

4.3. CFD validation

The tested configuration is defined by the parameter set $\alpha = 15^\circ$, $\beta = 10^\circ$, $X_{\text{CoR}} = 0.9$, $g_{12} = 0.1$, $k_{y1} = 2$, $k_{y2} = 1.2$, and $k_{c2} = 0.5$. The main-element angle of attack is set to $\alpha = 15^\circ$, which exceeds the nominal value of 12° . This choice imposes a more stringent flow condition, to ensure that the numerical setup remains stable and accurate in the presence of stronger nonlinear effects. Adequate performance under

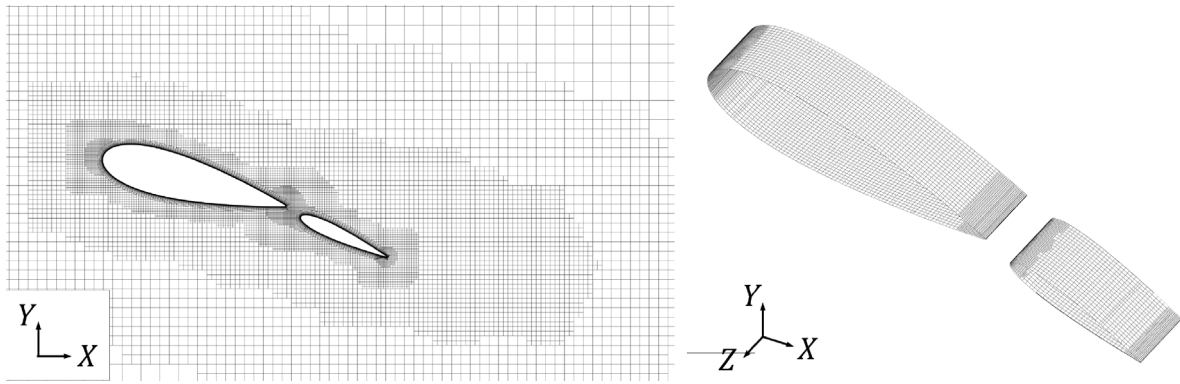


Fig. 4. The baseline mesh used in the mesh sensitivity study: (left) in the mid-span xy -plane and (right) on the configuration surface in an isometric view.




Fig. 5. Time histories of the lift coefficient (left) and drag coefficient (right) computed with the baseline mesh. The shaded zones outline the time-averaging window of 4.5 periods. The dashed lines indicate the resultant time-averaged values.

Table 2
Results of the mesh sensitivity study.

Case	$r_{i,i+1}$	$N_{\text{cell}} (\times 10^6)$	\overline{C}_L	$p_{\text{obs},L}$	$\text{GCI}_{21,L}$ (%)	\overline{C}_D	$p_{\text{obs},D}$	$\text{GCI}_{21,D}$ (%)
1	1.5	0.31	1.953			0.121		
2	1.3	0.34	1.939			0.123		
3	1.1	0.46	1.946			0.120		
4 (Baseline)	1.0	0.52	1.953			0.120		
5	0.9	0.64	1.962			0.120		
6	0.7	1.05	1.961	1.30	0.10	0.119	1.05	0.73
7	0.5	2.20	1.963	1.75	0.08	0.117	1.98	0.56

this condition implies sufficient accuracy at lower angle of attack such as $\alpha = 12^\circ$.

4.3.1. Mesh sensitivity study

The spanwise length of the computational domain is set to 0.25, which is slightly larger than the converged value of 0.2 obtained from the spanwise domain-size study in following Section 4.3.2. This choice provides a margin relative to the minimum required extent and reduces potential confinement effects associated with the finite spanwise domain.

Trimmed unstructured meshes are generated. The baseline mesh is shown in Fig. 4. On the wingsail surfaces, the target cell size is controlled to 0.01, and the minimum surface size to 0.002. Meanwhile, curvature refinement is enforced with 36 points per circle. The first-cell height is selected to achieve $y^+ \leq 1$ over most of the surface. Wake refinement zones are prescribed downstream with a streamwise extent of 1.5 and an isotropic target size of 0.03. This ensures adequate resolution of the wakes induced by the two elements and their interaction

region. The resulting mesh comprises approximately 0.53×10^6 cells, 1.57×10^6 interior faces, and 0.58×10^6 vertices.

Six additional meshes are generated by globally refining the baseline mesh with cell-size refinement ratios of $r_{i,i+1} = 1.5, 1.3, 1.1, 0.9, 0.7,$ and 0.5 . A grid-convergence study is conducted to assess the numerical uncertainty. The observed order of accuracy p_{obs} and the Grid Convergence Index (GCI) are evaluated following the procedure of Celik et al. (2008), which is based on the original formulation of Roache (1994). The GCI provides a conservative estimate of the discretization error via the Richardson extrapolation. This estimate accounts for the grid refinement ratio and convergence behavior.

The time histories of the lift coefficient, C_L , and the drag coefficient, C_D , computed using the baseline mesh are shown in Fig. 5. The simulation is performed for more than 60 periods, defined based on the freestream velocity and reference chord length. After the convergence criterion is satisfied, the statistical quantities are computed over the last 4.5 periods.

The results for all meshes, including p_{obs} and GCI_{21} , are summarized in Table 2, where the total number of cells is denoted by N_{cell} . The grid

Table 3
Results of the spanwise domain length study.

L_z	0.2	0.4	0.6	0.8	1.0
$\overline{C_L}$	1.953	1.953	1.953	1.953	1.953
$\overline{C_D}$	0.120	0.120	0.120	0.120	0.120
$\Delta\overline{C_L}/\overline{C_L} _{ref}$ (%)	0.07	0.07	0.07	0.07	0.07
$\Delta\overline{C_D}/\overline{C_D} _{ref}$ (%)	0.00	0.00	0.00	0.00	0.00
L_z	1.2	1.4	1.6	1.8	2.0 (Ref.)
$\overline{C_L}$	1.952	1.951	1.951	1.952	1.952
$\overline{C_D}$	0.120	0.120	0.120	0.120	0.120
$\Delta\overline{C_L}/\overline{C_L} _{ref}$ (%)	0.00	-0.03	-0.03	0.00	-
$\Delta\overline{C_D}/\overline{C_D} _{ref}$ (%)	0.00	0.00	0.00	0.00	-

Table 4
Results of the time-step sensitivity study.

CFL_∞ ($\times 10^{-2}$)	1.80	1.58	1.35	1.13	0.90 (Ref.)	0.68	0.45
$\Delta\overline{C_L}/\overline{C_L} _{ref}$ ($\times 10^{-5}$)	-3.8	-2.9	-1.4	-0.5	$\overline{C_L} _{ref} = 1.953$	1.5	2.7
$\Delta\overline{C_D}/\overline{C_D} _{ref}$ ($\times 10^{-5}$)	3.5	3.1	3.6	1.6	$\overline{C_D} _{ref} = 0.120$	-0.7	-0.2

convergence behavior is evaluated using the finest mesh (Case 7). For the lift coefficient, p_{obs} increases from approximately 1.30 to 1.75 from Case 6 to Case 7, accompanied by a decrease in GCI_{21} from 0.10% to 0.08%, which indicates weak sensitivity to further mesh refinement and near grid independence. For the drag coefficient, p_{obs} increases from approximately 1.05 to 1.98 over the same cases, approaching the nominal second-order accuracy, while the corresponding GCI_{21} decreases from 0.73% to 0.56%, indicating reduced discretization uncertainty and convergence toward the asymptotic range. The coarser meshes (Cases 1–3) do not show monotonic convergence, whereas consistent behavior is observed from Case 4 onward. Overall, the discretization uncertainty remains below 0.1% for lift and 1% for drag. The baseline mesh setup (Case 4) is therefore considered sufficiently accurate for the subsequent optimization, which provides a suitable balance between computational cost and numerical accuracy.

4.3.2. Convergence of spanwise domain length

Based on the baseline mesh setup introduced in Section 4.3.1, ten spanwise domain lengths are examined, ranging from $L_z = 0.2$ to 2.0 in increments of 0.2. The corresponding time-averaged lift and drag coefficients, $\overline{C_L}$ and $\overline{C_D}$, are reported in Table 3. The results show that $\overline{C_L}$ varies only slightly across the tested range, with a maximum deviation within $\pm 0.07\%$ relative to the reference case with $L_z = 2.0$. In contrast, $\overline{C_D}$ remains invariant for all cases. These results indicate that the solution is weakly sensitive to the spanwise domain length over the entire range considered. Therefore, the smallest spanwise length of 0.2 is adopted in the subsequent optimization, as it provides sufficient accuracy while minimizing computational cost.

4.3.3. Convergence of time step

A time-step sensitivity study is conducted using the baseline mesh in Section 4.3.1. There are seven cases ranging from $CFL_\infty = 0.45 \times 10^{-2}$ to 1.80×10^{-2} , defined by the time-step size and freestream velocity. The relative differences in $\overline{C_L}$ and $\overline{C_D}$ compared with the reference case of $CFL_\infty = 0.90 \times 10^{-2}$ are summarized in Table 4. The time-averaged coefficients exhibit negligible sensitivity to the selected time step sizes. The relative differences in $\overline{C_L}$ are within $\pm 3.8 \times 10^{-5}$, and those in $\overline{C_D}$ within $\pm 3.6 \times 10^{-5}$. This indicates that temporal discretization errors are insignificant over the tested range. Considering the numerical accuracy and computational efficiency, the reference time step corresponding to $CFL_\infty = 0.90 \times 10^{-2}$ is adopted in all subsequent simulations.

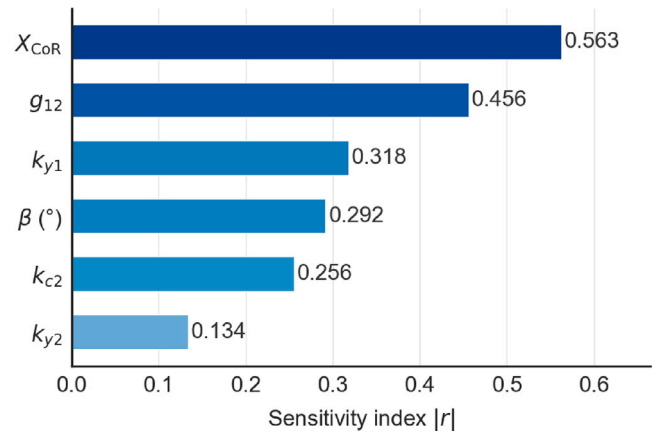


Fig. 6. Parameter sensitivity based on the absolute Pearson correlation coefficient.

5. Results and discussion

This section presents the optimization results with respect to the AWA-dependent objective formulation introduced in Section 3.1. The analysis is structured to highlight how different objective definitions influence the resulting aerodynamic performance and optimal configurations. A baseline case based on a single lift-dominated objective is first examined, followed by regime-partitioned formulations that account for varying contributions of lift and drag across AWA ranges.

5.1. Optimization under a single lift-dominated objective

The baseline optimization scenario is defined over the AWA interval of 30° – 150° . Under this condition, the integrated drag weighting becomes zero ($w_D = 0$), as seen in Eq. (8). Accordingly, the thrust objective reduces to the maximization of the time-averaged lift coefficient, $\overline{C_L}$, with the lift weighting of $w_L = 0.827$.

The relative influence of the design variables on the objective $\overline{C_L}$ is quantified using the absolute Pearson correlation coefficient $|r|$, as shown in Fig. 6. The center of rotation, X_{CoR} , exhibits the strongest correlation of $|r| = 0.563$, followed by the inter-element gap g_{12} of 0.456. A second tier of variables includes k_{y1} ($|r| = 0.318$), β ($|r| = 0.292$), and k_{c2} ($|r| = 0.256$), which show moderate influence on lift. In contrast, k_{y2} ($|r| = 0.134$) has a comparatively weak effect. This ranking indicates that lift variation is primarily governed by the relative positioning between the main element and flap, controlled by X_{CoR} and g_{12} . The remaining parameters act as secondary modifiers, adjusting camber distribution and local loading without significantly altering the overall lift level.

The structure of the design space is illustrated by parallel coordinates in Fig. 7. The top-performing designs cluster toward high values of X_{CoR} and low values of g_{12} . This behavior is consistent with the sensitivity analysis. This concentration indicates a well-defined primary design direction associated with increased lift. In contrast, the remaining parameters exhibit a broader spread within the high-performance subset. In particular, k_{c2} , k_{y1} , and β span a wide range of values among the top designs, indicating that multiple combinations of these variables can achieve comparable lift levels. The parameter k_{y2} shows a tendency toward lower values, although with greater dispersion than X_{CoR} and g_{12} . The best-performing design (the black line) follows the same overall trend, with high X_{CoR} and low g_{12} , while the remaining parameters lie within the envelope of the top-performing set. To summarize the distribution of the top-performing subset robustly, the shaded band in the figure represents the interquartile range (IQR) of the top 20 designs on each axis, defined as the interval between the 25th and 75th percentiles. It therefore encloses the central 50% of the values

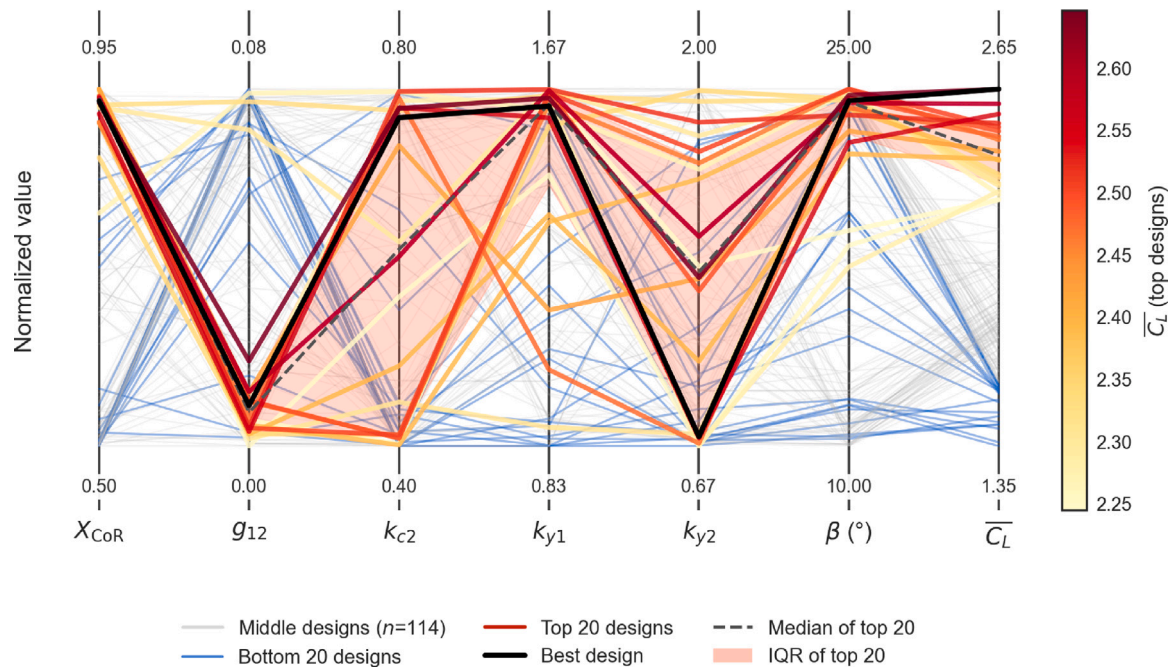


Fig. 7. Parallel coordinates plot showing the distribution of design variables for different performance levels.

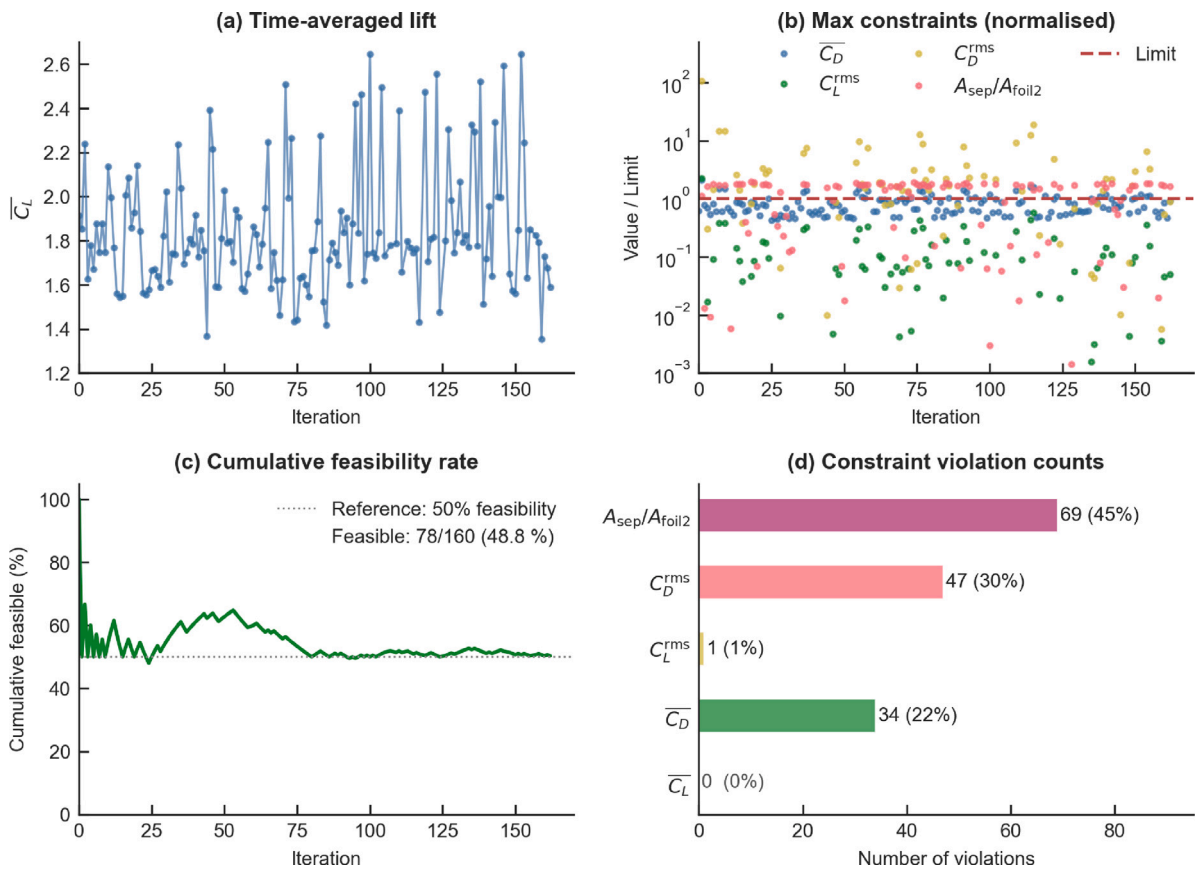


Fig. 8. Constraint evolution and feasibility statistics during the optimization process.

while remaining relatively insensitive to extreme designs. The median trend and IQR further support the observations above, showing a tight concentration for X_{CoR} and g_{12} , where the band collapses into a narrow ribbon, and comparatively larger variability for the other variables, where the band broadens substantially. This suggests that X_{CoR} and g_{12}

play dominant roles in the primary lift-generation mechanism, whereas the remaining parameters mainly provide secondary tuning within a broader feasible design region.

The evolution of objective and constraint values during the optimization is shown in Fig. 8. While the lift constraint remains inactive




Fig. 9. Response surface analysis about the influences of selected parameter pairs on $\overline{C_L}$.

throughout the process, the feasible design space is primarily shaped by flow-related constraints associated with drag and separation. The time-averaged lift consistently exceeds the minimum requirement (Fig. 8(a)), indicating that lift generation is not a limiting factor under the present objective. Fig. 8(b) presents the normalized constraint values, obtained by dividing each instantaneous constraint metric by its prescribed limit. A normalized value below unity therefore denotes a feasible condition, whereas values exceeding unity indicate constraint violation. The logarithmic scale highlights the wide variation in violation magnitude across the explored designs. It suggests that the violations are predominantly associated with the separation area and drag-related quantities, which frequently approach or exceed their prescribed limits. Fig. 8(c) shows the cumulative feasibility rate during the optimization process, defined as the percentage of sampled designs satisfying all imposed constraints up to a given iteration. The feasibility rate gradually stabilizes at approximately 48.8%, indicating that nearly half of the explored designs are rejected due to violations of the flow-related constraints. This behavior is further reflected in the violation statistics (Fig. 8(d)), where the separation area ratio and RMS drag coefficient account for the majority of constraint violations, with additional contributions from the mean drag coefficient. In sum, these results indicate that the optimization operates in a regime where lift is readily achieved, and performance is instead limited by the onset of separation and increased drag. The constraints therefore act to bound the admissible high-lift configurations, restricting the design space to regions where lift enhancement can be maintained without incurring excessive flow degradation.

Based on the sensitivity ranking, X_{CoR} is selected as the reference variable to examine its interaction with representative parameters. The corresponding response surfaces are shown in Fig. 9, revealing a consistent design trend governed by the joint effect of X_{CoR} and the secondary variables. In the (X_{CoR}, g_{12}) plane (Fig. 9(a)), $\overline{C_L}$ increases monotonically toward high X_{CoR} and low g_{12} , forming a well-defined gradient that aligns with both the sensitivity ranking and the clustering observed in the parallel coordinates plot. The high-performance region is concentrated near this boundary, indicating a strong and coupled influence of these two variables. In contrast, the dependence on k_{y1} (Fig. 9(b)) is comparatively weak, with a relatively flat response over a wide range of values. This suggests that k_{y1} primarily acts as a secondary tuning parameter once the dominant configuration set by X_{CoR} is established. The flap deflection angle β (Fig. 9(c)) exhibits a moderate influence, with higher lift generally associated with larger β , although without a sharply defined optimum. The irregular boundaries visible in the response surfaces, particularly in Fig. 9(a) and (c), correspond to infeasible regions of the design space. These truncated regions indicate that the achievable high-lift configurations are constrained not

Table 5

Geometric parameters of selected designs such as best (rank 1) and two worst within the top 20 (rank 19 and rank 20).

Design	X_{CoR}	g_{12}	k_{c2}	k_{y1}	k_{y2}	β (°)
Rank 1	0.935	0.0135	0.768	1.627	0.702	24.5
Rank 19	0.930	0.0774	0.776	1.652	1.954	24.5
Rank 20	0.942	0.0104	0.402	1.659	1.338	19.1

only by aerodynamic trends but also by geometric and flow-related limitations, which restrict the accessible portion of the design space near the apparent optimum.

Table 5 lists the geometric parameters of three representative designs within the Pareto-ranked set. Moreover, Table 6 highlights the role of flow attachment in determining aerodynamic performance through URANS and IDDES. Since experimental data at the present full-scale Reynolds number, $Re_c \approx 5.7 \times 10^6$, are not available in the public literature for the exact two-element geometry considered here, we do not claim direct experimental validation of the optimized cases. For example, existing two-element wingsail experiments (Fiumara et al., 2016, 2018; Hillenbrand et al., 2024) were conducted at substantially lower Reynolds numbers. Therefore, IDDES computations are performed at the same Reynolds number, geometry, and AOA as the URANS simulations. This provides a more direct assessment of the model-form uncertainty associated with URANS-SST under the present operating conditions. The mesh resolution of the IDDES computations is 1.5 times finer than that of URANS, and the numerical setup follows the previous studies (Zhu et al., 2023; van Reen et al., 2025).

As seen in Table 6, the best design (rank 1) achieves the highest $\overline{C_L}$ while maintaining fully attached flow, as indicated by $A_{sep}/A_{foil2} \approx 0$, and correspondingly exhibits negligible unsteady content with vanishingly small C_L^{rms} and C_D^{rms} . In contrast, rank 19 is characterized by a substantial separation region covering approximately 45% of the second element. This is associated with a reduction in $\overline{C_L}$ and the emergence of weak unsteadiness, indicating partial flow detachment. Despite having a similar lift level to rank 19, rank 20 maintains attached flow with $A_{sep}/A_{foil2} \approx 0$ and negligible fluctuations, indicating that comparable performance can be achieved without separation. These comparisons infer that high-lift performance is consistently associated with attached-flow configurations, while partial separation introduces unsteadiness without providing a compensating lift benefit. Flow attachment therefore acts as a key constraint in the selection of optimal designs within the high-performance set.

The time-averaged flow fields for three representative designs are compared in Fig. 10 to elucidate the mechanisms governing their aerodynamic performance. The best design (rank 1) achieves the highest




Fig. 10. Time- and span-averaged pressure coefficient \overline{C}_p (left) and normalized velocity magnitude \overline{U}/U_∞ with streamlines (right) for three representative designs (Rank 1, Rank 19, Rank 20).

Table 6
Aerodynamic coefficients and separation area of selected designs using different simulation methods.

Design	Method	\overline{C}_L	C_L^{rms}	\overline{C}_D	C_D^{rms}	A_{sep}/A_{foil2}
Rank 1	URANS	2.646	≈ 0	0.2355	≈ 0	≈ 0
	IDDES	2.633	≈ 0	0.2410	≈ 0	≈ 0
Rank 19	URANS	2.268	3.95×10^{-3}	0.1920	2.22×10^{-3}	0.453
	IDDES	2.189	1.13×10^{-2}	0.2095	2.86×10^{-3}	0.460
Rank 20	URANS	2.265	≈ 0	0.1403	≈ 0	≈ 0
	IDDES	2.191	≈ 0	0.1466	≈ 0	≈ 0

lift ($\overline{C}_L = 2.646$) through a strong suction peak on the leading element, followed by a smooth pressure recovery over the second element, resulting in fully attached flow on both surfaces ($A_{sep}/A_{foil2} \approx 0$). In contrast, rank 19 exhibits a fundamentally different flow structure, with a large separation region covering approximately 45% of the suction side of the second element. This is reflected by low-velocity regions and recirculating streamlines, together with a weakened pressure gradient. The associated loss of suction leads to a reduction of about 14% in \overline{C}_L and is accompanied by the onset of weak unsteadiness. Rank 20 again maintains fully attached flow, as indicated by continuous streamlines and the absence of recirculation. However, the suction levels over

both elements are consistently weaker than in rank 1. This behavior is consistent with its less aggressive geometry, characterized by lower flap deflection ($\beta = 19.1^\circ$) and reduced camber ($k_{c2} = 0.402$), which limit circulation without inducing separation. As a result, a lower lift coefficient is obtained, comparable to rank 19 but arising from a different mechanism. The comparison distinguishes two limiting factors for lift performance: excessive loading leading to separation and insufficient loading leading to weak circulation. The optimal configuration balances these effects by maintaining attached flow while sustaining strong suction over both elements.

To further compare the aerodynamic characteristics of the representative optimized geometries, Fig. 11 presents the time- and span-averaged pressure coefficient distributions on the upper and lower surfaces. Rank 1 exhibits the strongest suction peak near the leading edge of the main element, followed by the most extended low-pressure regions over both main element and flap. This leads to the largest lift generation. In comparison, rank 19 shows weaker suction levels together with more rapid pressure recovery, particularly on the flap upper surface due to the flow separation observed in Fig. 10. Given that rank 20 employs the smallest flap chord, the suction level over the flap is weaker on both surfaces. Additional differences are observed near the flap leading edge, where rank 1 maintains a sharper suction peak associated with stronger slot-driven flow acceleration. The other designs




Fig. 11. Time- and span-averaged pressure coefficient distributions on the surfaces for the best design (Rank 1) and the two worst designs within the top 20 (Ranks 19 and 20). The solid and dashed lines denote the lower and upper surfaces. The chordwise coordinates are normalized by the reference chord length, c_0 .

exhibit weaker suction levels in this region, suggesting reduced flow energization through the gap. Therefore, the improved aerodynamic performance is dependent on maintaining strong suction and smooth pressure recovery across both elements while avoiding large-scale flow separation.

5.2. Adaptation of optimization objectives to variable AWA regimes

To resolve the regime-dependent nature of thrust generation, five AWA cases are defined to enable a consistent interpretation of optimization results across operating conditions. Since the time-averaged thrust coefficient $\overline{C_T}(\theta)$ varies systematically with AWA (Eq. (6)), evaluating performance over a single broad interval can obscure the distinct roles of lift and drag in propulsion. Partitioning the AWA domain into targeted sub-ranges preserves this dependence and allows the underlying mechanisms to be examined separately. For each case, the angular-averaged thrust coefficient in Eq. (7) is expressed in the weighted form in Eq. (8), where the coefficients w_L and w_D quantify the relative contributions of lift and drag. In particular, the sign and magnitude of w_D indicate whether drag acts as a penalty or contributes positively to thrust, providing a direct link between AWA regime and propulsive mechanisms. The normalized weights for each AWA case are summarized in Table 7. Each AWA regime was optimized independently using the same BO evaluation budget and optimization settings, with a maximum of 160 evaluations. This formulation establishes a consistent basis for comparing optimal designs across regimes, and enables the identification of regime-dependent trends in aerodynamic performance and configurations.

Case 1 (AWA in 30° – 150°) spans the full AWA domain and serves as a drag-neutral reference with $w_D = 0$, resulting in a purely lift-driven objective. Cases 2 (30° – 90°) and 3 (30° – 60°) represent upwind conditions, where $w_D > 0$ and the drag acts as a penalty. This penalty increases as the AWA range shifts toward smaller angles, with Case 3 corresponding to equal weighting of the lift and drag, i.e., $|w_L| = |w_D|$. Case 4 (90° – 150°) represents the transition to downwind conditions, where $w_D < 0$ and drag contributes positively to thrust. This contribution becomes more pronounced in Case 5 (120° – 150°), where the lift and drag again exhibit equal weighting, but with the drag contributing constructively rather than penalizing. Taken together, the five cases define a continuous transition from lift-dominated to drag-assisted propulsion, to provide a consistent framework for interpreting regime-dependent optimization outcomes.

Cases 1–3 yield the same optimal geometry, as presented in Table 8. This confirms that when drag penalizes or does not contribute to thrust

($w_D \geq 0$), the design is governed by lift maximization. The optimal configuration consistently attains $\overline{C_L} = 2.646$ with moderate drag of $\overline{C_D} = 0.236$. Cases 4 and 5 converge to a distinct configuration in which k_{y2} increases substantially from 0.702 to 1.296, while other design parameters remain nearly unchanged. This modification yields a slight increase in $\overline{C_D}$ (0.238 vs. 0.236), detrimental in upwind conditions but beneficial under downwind weighting, resulting in a higher $\overline{C_T}$. The near invariance of $\overline{C_L}$ across all cases indicates that the performance difference is governed by drag rather than lift, revealing a regime-dependent bifurcation controlled by the sign of w_D . The sensitivity analysis in Fig. 12 provides further insight into this regime dependence.

Across all AWA cases, X_{CoR} and g_{12} consistently emerge as the dominant design parameters ($|r| > 0.4$), indicating that the primary aerodynamic response is governed by the relative positioning between the main element and flap. The parameter k_{y1} ranks third in importance and remains comparatively stable across regimes, acting as a secondary adjustment. The most pronounced regime-dependent variation is observed in β , whose sensitivity increases from 0.13–0.29 in the upwind cases (Cases 1–3) to 0.38–0.43 in the downwind cases (Cases 4 and 5). This trend indicates that sail twist becomes more influential when drag contributes positively to thrust, consistent with the shift from lift-dominated to drag-assisted propulsion. In contrast, k_{y2} exhibits consistently low global sensitivity ($|r| \leq 0.17$), yet plays a key role in distinguishing the optimal configurations between regimes. As shown in Table 8, k_{y2} increases markedly from the upwind cases (Cases 1–3) to the downwind cases (Cases 4 and 5), while the remaining parameters remain nearly unchanged. This behavior reflects the distinction between global sensitivity, which characterizes influence across the design space, and local optimality, where targeted variation in k_{y2} produces a significant change in $\overline{C_D}$ under downwind weighting. The results discussed above indicate that the regime-dependent transition in optimal designs is achieved through a localized adjustment of k_{y2} superimposed on an otherwise consistent baseline configuration.

Fig. 13 displays the time- and span-averaged pressure coefficient distributions for the two optimal designs, which are obtained from the lift-dominated AWA ranges (cases 1–3 in Table 8) and from the drag-assisted AWA ranges (cases 4 and 5). The designs show very similar distributions over the entire main element surfaces and the lower flap surface. The most noticeable differences appear on the flap upper surface, where the lift-dominated design has slightly stronger suction and a sharper suction peak near the flap leading edge. This behavior infers a more aggressive aerodynamic loading on the flap. In contrast, the drag-assisted design exhibits slightly weaker suction together with smoother pressure recovery downstream of the flap

Table 7
Normalized integration weights for different AWA ranges, computed from Eq. (8).

Case	AWA range	$\Delta\theta$ (rad)	w_L	w_D	Physical interpretation
1	30°–150°	2.094	0.827	0.000	Pure lift contribution (drag-neutral)
2	30°–90°	1.047	0.827	0.477	Lift-dominated; drag as a penalty
3	30°–60°	0.524	0.699	0.699	Large drag penalty (equal weighting, upwind)
4	90°–150°	1.047	0.827	−0.477	Drag-assisted for positive contribution to thrust
5	120°–150°	0.524	0.699	−0.699	Large drag-assisted (equal weighting, downwind)

Table 8
Optimized geometric parameters and aerodynamic performance of the best design for each thrust objective case.

Case	AWA range	X_{CoR}	g_{12}	k_{c2}	k_{y1}	k_{y2}	β (°)	\overline{C}_L	\overline{C}_D	\overline{C}_T
1	30°–150°	0.935	0.0135	0.768	1.627	0.702	24.51	2.646	0.236	2.646
2	30°–90°	0.935	0.0135	0.768	1.627	0.702	24.51	2.646	0.236	2.076
3	30°–60°	0.935	0.0135	0.768	1.627	0.702	24.51	2.646	0.236	1.685
4	90°–150°	0.933	0.0229	0.778	1.646	1.296	24.74	2.645	0.238	2.301
5	120°–150°	0.933	0.0229	0.778	1.646	1.296	24.74	2.645	0.238	2.015




Fig. 12. Sensitivity index ($|r|$, the Pearson correlation) of the six design variables across five thrust objective cases defined by different AWA ranges.

leading edge, suggesting a modest reduction in aerodynamic loading. This observation is consistent with the optimized geometric parameters summarized in Table 8, where the different objectives defined by the AWA ranges lead to similar multi-element loading characteristics with moderate adjustments in the flap loading.

The results in this section indicate a consistent hierarchy of design variables across both lift- and thrust-based objectives. The hinge position X_{CoR} and the inter-element gap g_{12} remain the dominant parameters in all AWA regimes, governing circulation development and slot-flow behavior through the relative positioning of the elements. Within the present design space, larger X_{CoR} and smaller g_{12} are consistently associated with high-performance configurations. The remaining parameters act as secondary variables: k_{y2} controls the transition between lift-dominated and drag-assisted configurations, while the influence of β increases toward downwind conditions. Consistent with the constraint analysis, feasible designs are primarily limited by separation and unsteady drag rather than by lift, indicating that the dominant variables govern both performance and feasibility.

The regime-dependent bifurcation observed across AWA cases highlights a limitation of single-objective aggregation for multi-point optimization. When the objective is defined over a broad AWA range (30°–150°, Case 1), the corresponding weighting ($w_D = 0$) removes

the influence of drag and recovers the same solution as pure lift maximization. As a result, the formulation favors lift-dominated configurations and does not promote designs that exploit drag-assisted propulsion. For applications spanning both upwind and downwind conditions, regime-specific or multi-objective formulations are therefore required to represent competing performance mechanisms. The small geometric difference between the two optimal configurations, primarily in k_{y2} , further suggests that transitions between regimes may be achieved through limited parameter variation, which may be relevant for adaptive or reconfigurable wingsail concepts.

6. Conclusions

A regime-dependent single-objective aerodynamic optimization of a two-element wingsail cross-section has been performed using a Bayesian optimization framework coupled with three-dimensional URANS simulations in a six-dimensional geometric design space. The use of verified URANS enables resolving the nonlinear and unsteady flow phenomena governing two-element performance, which could not be reliably captured by lower-fidelity or quasi-steady approaches. The observed sensitivity of thrust to small geometric variations, particularly




Fig. 13. Time- and span-averaged pressure coefficient distributions on the upper and lower surfaces of the optimal designs obtained from the lift-dominated AWA ranges (Best \overline{C}_L) and the drag-assisted AWA ranges (Best \overline{C}_L & \overline{C}_D).

in the gap size and hinge position, reflects the strong coupling between geometry and flow physics resolved by this approach.

A regime-dependent thrust formulation is introduced based on a weighted angular average of the thrust coefficient over prescribed AWA intervals, enabling consistent comparison of aerodynamic performance across upwind, beam-reach, and downwind conditions. By construction, the formulation captures the changing role of drag in propulsion, which penalizes performance for $\theta < 90^\circ$, vanishes at $\theta = 90^\circ$, and contributes positively for $\theta > 90^\circ$. Within this framework, a clear bifurcation in optimal designs is observed, governed by the sign of the drag weighting w_D . For AWA regimes with $w_D \geq 0$ (upwind and beam reach), the optimal configuration is lift-dominated and converges to a single geometry characterized by high \overline{C}_L and fully attached flow. In contrast, for regimes with $w_D < 0$ (downwind), a distinct configuration emerges in which a moderate increase in drag contributes positively to thrust. Despite this shift, \overline{C}_L remains nearly invariant across all regimes, indicating that performance differences are primarily driven by drag rather than lift. This transition is achieved through a localized adjustment of the flap thickness parameter k_{y2} , while the hinge position X_{CoR} and inter-element gap g_{12} remain the dominant variables across all cases. These results highlight a separation of roles in the design space, where global performance is governed by a consistent baseline configuration and regime-specific adaptation is achieved through targeted parameter variation. These findings also underscore the critical influence of flap deflection in regulating flow separation and its contribution to drag, despite separation being constrained in the current optimization strategy. Nevertheless, this brings about a penalty of strong flow fluctuations due to separation.

From a design perspective, the results provide guidance for the development of wingsail-assisted propulsion systems. The consistent dominance of X_{CoR} and g_{12} indicates that the relative positioning of the main element and flap defines the primary performance envelope and should be prioritized in early-stage design. At the same time, the regime-dependent role of k_{y2} and β suggests that performance across operating conditions can be enhanced through targeted geometric adaptation. The small geometric separation between the lift-dominated and drag-assisted configurations indicates that transitions between regimes may be achieved through limited parameter variation, which is relevant for adaptive or reconfigurable wingsail concepts. More broadly, the results show that high-performance designs are bounded by flow attachment constraints, emphasizing the importance of maintaining stable slot-flow behavior while increasing aerodynamic loading.

From a methodological perspective, the observed regime-dependent bifurcation in optimal configurations highlights a limitation of single-objective aggregation for multi-point optimization. When performance

is averaged over a broad AWA range, the resulting objective suppresses competing propulsion mechanisms and favors lift-dominated solutions, thereby overlooking configurations that exploit drag-assisted thrust under downwind conditions. This suggests that regime-specific or multi-objective formulations are required for wingsails operating across diverse conditions, particularly when different physical mechanisms contribute to performance.

The present study is subject to several limitations. The analysis is conducted at a fixed angle of attack ($\alpha = 12^\circ$) and a representative Reynolds number ($Re_c \approx 5.7 \times 10^6$), such that variations in operating conditions may alter the balance between lift, drag, and separation, and consequently the optimal configurations. In addition, the study focuses on a short section and does not account for effects such as tip vortices or full-system integration. Future work should extend the framework to variable operating conditions, including α , Reynolds number, and in-flow characteristics, and incorporate experimental validation to assess the robustness of the identified regime-dependent designs.

Declaration of competing interest

On behalf of all authors, we hereby confirm that there are no known competing financial interests or personal relationships that could have appeared to influence the work reported in this manuscript.

All authors have reviewed and approved this declaration and agree that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Acknowledgments

The research is financed by the Swedish Transport Administration, Sweden in the project ‘‘Generic Multidisciplinary optimization for sail INstallation on wInd-assisted ships’’ (GEMINI, Grant No. TRV 2023/32107). The computations and data handling were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council through grant agreement no. 2022-06725.

References

- Celik, I.B., Ghia, U., Roache, P.J., Freitas, C.J., Coleman, H., Raad, P.E., 2008. Procedure for estimation and reporting of uncertainty due to discretization in CFD applications. *J. Fluids Eng.* 130 (7), 078001.
- Cerantola, D., Ferchichi, M., Guellouz, M.S., Ben Chiekh, M., 2025. Assessment of the sensitivity of high-lift airfoil performance to flap gap and overlap. *Results Eng.* 28, 108244. <http://dx.doi.org/10.1016/j.rineng.2025.108244>.

- Chapin, V., Gourdain, N., Verdin, N., Fiumara, A., Senter, J., 2015. Aerodynamic study of a two-elements wingsail for high-performance multihull yachts. In: Proceedings of the 5th High Performance Yacht Design Conference. HPYD5, Auckland, New Zealand, 10–12 March 2015.
- Durbin, P.A., 1996. On the $k-\epsilon$ stagnation point anomaly. *Int. J. Heat Fluid Flow* 17 (1), 89–90. [http://dx.doi.org/10.1016/0142-727X\(95\)00073-Y](http://dx.doi.org/10.1016/0142-727X(95)00073-Y).
- Fiumara, A., Gourdain, N., Chapin, V.-G., Senter, J., 2017. Aerodynamic analysis around a C-class catamaran in gust conditions using LES and unsteady RANS approaches. In: Proceedings of the 4th International Conference on Innovation in High Performance Sailing Yachts. INNOV'sAIL 2017, Lorient, France, pp. 85–96. HAL: hal-01583529.
- Fiumara, A., Gourdain, N., Chapin, V., Senter, J., 2018. Aerodynamic analysis of a C-class-like Catamaran in simplified unsteady wind conditions using LES and URANS modeling. *J. Wind Eng. Ind. Aerodyn.* 180, 262–275. <http://dx.doi.org/10.1016/j.jweia.2018.08.008>.
- Fiumara, A., Gourdain, N., Chapin, V., Senter, J., Bury, Y., 2016. Numerical and experimental analysis of the flow around a two-element wingsail at Reynolds number 0.53×10^6 . *Int. J. Heat Fluid Flow* 62, 538–551. <http://dx.doi.org/10.1016/j.ijheatfluidflow.2016.08.005>.
- Head, T., Kumar, M., Nahrstaedt, H., Louppe, G., Shcherbatyi, I., 2020. Scikit-optimize/scikit-optimize: v0.8.1. <http://dx.doi.org/10.5281/zenodo.4014775>.
- Hillenbrand, A., Marimon Giovannetti, L., Dhomé, U., Kuttenukeuler, J., 2024. Wind tunnel tests of a two-element wingsail with focus on near-stall aerodynamics. *J. Sail. Technol.* 9 (1), 110–127. <http://dx.doi.org/10.5957/jst/2024.9.1.110>, Paper No. SNAME-JST-2024-07.
- Jones, D.R., Schonlau, M., Welch, W.J., 1998. Efficient global optimization of expensive black-box functions. *J. Global Optim.* 13 (4), 455–492. <http://dx.doi.org/10.1023/A:1008306431147>.
- Khan, L., Macklin, J.J.R., Peck, B.C.D., Morton, O., Soupez, J.-B.R.G., 2021. A review of wind-assisted ship propulsion for sustainable commercial shipping: Latest developments and future stakes. In: Proceedings of the Wind Propulsion Conference. London, UK, 15–16 September 2021. The Royal Institution of Naval Architects (RINA).
- Li, C., Wang, H., Sun, P., 2020. Numerical investigation of a two-element wingsail for ship auxiliary propulsion. *J. Mar. Sci. Eng.* 8 (5), 333. <http://dx.doi.org/10.3390/jmse8050333>.
- Li, C., Wang, H., Sun, P., 2023. Study on the influence of gradient wind on the aerodynamic characteristics of a two-element wingsail for ship-assisted propulsion. *J. Mar. Sci. Eng.* 11 (1), 134. <http://dx.doi.org/10.3390/jmse11010134>.
- Li, C., Wu, B., Wen, H., Yuan, J., 2025. Study on influence of configuration of bulges on stall characteristics of two-element Wingsails for ships. *J. Mar. Sci. Eng.* 13 (2), 337. <http://dx.doi.org/10.3390/jmse13020337>.
- Makram, T.E., Panagiotou, P., Mattheou, D., 2023. Wingsail layout design and shape optimization using a CFD-aided Taguchi approach: The Aegean Marathon case study. *Ocean Eng.* 276, 114055. <http://dx.doi.org/10.1016/j.oceaneng.2023.114055>.
- Mason, W.H., 2018. Configuration Aerodynamics. Virginia Tech, Blackburg VA.
- Menter, F.R., 1994. Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA J.* 32 (8), 1598–1605. <http://dx.doi.org/10.2514/3.12149>.
- Negoita, M.-F., Bogoi, A., Bunesco, I., Hothazie, M.-V., 2025. Deterministic optimization of single-slotted flaps using an automated CFD workflow. *INCAS Bull.* 17 (4), 103–121. <http://dx.doi.org/10.13111/2066-8201.2025.17.4.8>.
- Rasmussen, C.E., Williams, C.K.I., 2006. Gaussian Processes for Machine Learning. MIT Press, Cambridge, MA.
- Roache, P.J., 1994. Perspective: A method for uniform reporting of grid refinement studies. *J. Fluids Eng.* 116 (3), 405–413.
- Siemens Digital Industries Software, 2025. Simcenter STAR-CCM+ user guide. Plano, TX, USA.
- Smirnov, P.E., Menter, F.R., 2009. Sensitization of the SST turbulence model to rotation and curvature by applying the Spalart–Shur correction term. *J. Turbomach.* 131 (4), 041010. <http://dx.doi.org/10.1115/1.3070573>.
- Smith, A.M.O., 1975. High-lift aerodynamics. *J. Aircr.* 12 (6), 501–530. <http://dx.doi.org/10.2514/3.59830>.
- Spalart, P.R., Shur, M., 1997. On the sensitization of turbulence models to rotation and curvature. *Aerosp. Sci. Technol.* 1 (5), 297–302. [http://dx.doi.org/10.1016/S1270-9638\(97\)90051-1](http://dx.doi.org/10.1016/S1270-9638(97)90051-1).
- Srinivas, N., Krause, A., Kakade, S.M., Seeger, M.W., 2010. Gaussian process optimization in the bandit setting: No regret and experimental design. In: Proceedings of the 27th International Conference on Machine Learning. ICML, Haifa, Israel, pp. 1015–1022, Extended version: *IEEE Trans. Inform. Theory*, 58(5):3250–3265, 2012, doi:10.1109/TIT.2011.2182033.
- van Dam, C.P., 2002. The aerodynamic design of multi-element high-lift systems for transport airplanes. *Prog. Aerosp. Sci.* 38 (2), 101–144. [http://dx.doi.org/10.1016/S0376-0421\(02\)00002-7](http://dx.doi.org/10.1016/S0376-0421(02)00002-7).
- van Reen, S., Serbüent, B., Yao, H.-D., 2025. Machine learning-based multipoint optimisation for improving aerodynamics of symmetrically cambered wing sails in wind-assisted ship propulsion. *Ocean Eng.* 342, 122829. <http://dx.doi.org/10.1016/j.oceaneng.2025.122829>.
- von Klemperer, C.J., Horwitz, R.A., Malan, A.G., 2023. An articulating wingsail design for wind assisted ship propulsion (WASP) applications. *Sci. Afr.* 20, e01699. <http://dx.doi.org/10.1016/j.sciaf.2023.e01699>.
- Yao, H.-D., 2025. Hybrid parameterization of symmetrically cambered (crescent-shaped) airfoil profiles for rigid wingsail design in wind-assisted ship propulsion. Technical Report TR-2025-1, Chalmers University of Technology, Gothenburg, Sweden, <http://dx.doi.org/10.63959/m2.techreport/2025.1>.
- Zhu, H., Ringsberg, J.W., Ramne, B., Yao, H.-D., 2025. Aerodynamic analysis of full- and model-scale wingsails with cambered profiles. In: Proceedings of the XI International Conference on Computational Methods in Marine Engineering. MARINE 2025, Edinburgh, Scotland, 23–25 June 2025.
- Zhu, H., Yao, H.-D., Ringsberg, J.W., 2024. Unsteady RANS and IDDES studies on a telescopic crescent-shaped wingsail. *Ships Offshore Struct.* 19 (1), 134–147. <http://dx.doi.org/10.1080/17445302.2023.2256601>.
- Zhu, H., Yao, H.-D., Thies, F., Ringsberg, J.W., Ramne, B., 2023. Propulsive performance of a rigid wingsail with crescent-shaped profiles. *Ocean Eng.* 285, 115349. <http://dx.doi.org/10.1016/j.oceaneng.2023.115349>.