

A Data-Driven Design Framework for Wind Turbines via Design-by-Morphing and Bayesian Optimization

Hamza Latif Mehr^{a,*,1}, Imran Akhtar^b and Haris Moazam Sheikh^{c,1}

^aGraduate School of Sustainable Engineering for Energy, Water and Environment (ENSE³), Grenoble INP, Université Grenoble Alpes, Grenoble, 38000, France

^bSchool of Interdisciplinary Engineering & Sciences (SINES), National University of Sciences & Technology (NUST), Islamabad, 44000, Pakistan

^cUniversity of Southampton, University Road, Southampton, SO17 1BJ, United Kingdom

ARTICLE INFO

Keywords:

Shape optimization
Design-by-morphing
Bayesian optimization
Vertical-axis wind turbines
Computational fluid dynamics

ABSTRACT

This study presents a data-driven design optimization framework that integrates Design-by-Morphing (DbM) with Bayesian optimization (BO) to facilitate the discovery of novel, high-performance aerodynamic geometries. Applied to vertical-axis wind turbines (VAWTs), the DbM approach enables the generation of novel shapes that go beyond the conventional constraints of lift-based (Darrieus) and drag-based (Savonius) geometries. Candidate designs were evaluated via two-dimensional Unsteady Reynolds-Averaged Navier-Stokes (URANS) simulations under representative urban wind conditions, with the design space systematically optimized using BO to maximize the power coefficient, C_p . From a design space of 5.80 million unique physical configurations, the best-performing configuration was first identified after 13 BO-guided evaluations following a 50 design initialization. The best design achieved a 21.26% higher maximum C_p than the best existing baseline design, while maintaining a stable periodic torque response across the investigated Reynolds number range. Ultimately, this work presents a scalable and computationally efficient DbM–BO framework for aerodynamic shape optimization. While demonstrated here for a two-bladed VAWT using primarily two-dimensional URANS simulations, the modular framework is not specific to this configuration and can be extended to other aerodynamic shape-design applications.

1. Introduction

Vertical-axis wind turbines (VAWTs) have attracted increasing interest for urban and decentralized applications because of their omnidirectional operation, ease of installation, and lower noise levels [1, 2]. Their wider adoption, however, has been limited by inherently lower aerodynamic efficiency: while horizontal-axis wind turbines (HAWTs) can reach efficiencies of 50–55%, VAWTs typically achieve only 30–40% [3]. This performance gap stems largely from the strong dependence of aerodynamic output on blade shape, turbine configuration, and tip-speed ratio (TSR). Historically, drag- and lift-based blade shapes for VAWTs were derived from engineering intuition built on decades of HAWT research, but VAWTs exhibit fundamentally different flow physics, so aerodynamically optimal design concepts do not directly translate. Therefore, systematic VAWT optimization that avoids existing designer biases is needed [4], and recent research has increasingly focused on developing high-performance VAWT rotors.

To overcome the low aerodynamic efficiency associated with conventional blade geometries, recent studies have increasingly employed intuitive hybrid designs that combine lift and drag-based turbine characteristics [5, 6]. The emergence of advanced additive manufacturing techniques has enabled the realization of geometries that were previously impractical to fabricate [7], enabling the exploration of more

complex and non-conventional turbine geometries. Several review studies have summarized the aerodynamic design space of VAWTs and highlighted the key parameters governing their performance. Bhutta et al. [8] provided an early comprehensive review of VAWT configurations and design methodologies, emphasizing the influence of rotor solidity, blade number, and airfoil selection on turbine efficiency. Recent work by Jassim et al. [9] further investigated a leaf-shaped hybrid Savonius–Darrieus VAWT and showed that turbine staging and blade number strongly affect the balance between torque, power output, and aerodynamic efficiency. Similarly, Shijer et al. [10] demonstrated through CFD and experimental analysis that wind-turbine performance is strongly sensitive to geometric and operating parameters, including blade angle, turbine diameter, rotational speed, blade number, and incoming wind velocity. Shen et al. [11] reviewed performance enhancement strategies for Darrieus turbines and reported that improvements are typically achieved through blade-geometry refinement, flow control techniques, and optimized operational parameters. These studies collectively demonstrate that VAWT performance depends on a complex interaction between geometric and operational variables, making systematic optimization approaches particularly valuable.

In addition, several recent works have investigated specific aerodynamic parameters affecting VAWT performance. Kuang et al. [12] performed a systematic analysis of rotor solidity over the range 0.12–0.60 and demonstrated that turbine efficiency strongly depends on the balance between blade loading and aerodynamic losses, with low-to-moderate solidity values below approximately 0.3 generally providing

*Corresponding author

✉ hamzamehr.hm@gmail.com (H.L. Mehr)

ORCID(s): 0009-0003-1889-2983 (H.L. Mehr)

¹Hamza Latif Mehr and Haris Moazam Sheikh contributed equally to this work.




Figure 1: Overview of the methodology combining DbM, CFD, and Bayesian optimization in a loop.

appreciable peak power performance. Experimental investigations by Peng et al. [13] further showed that turbulence intensity and inflow conditions can significantly alter the aerodynamic behavior of straight-bladed VAWTs, with turbulence intensities of 0.9%–14.6% producing reported efficiency improvements of 7.9%–38.0%, depending on the turbine configuration. From a numerical perspective, Fazlizan et al. [14] reviewed computational fluid dynamics (CFD) methodologies for VAWT analysis and emphasized the growing role of CFD-based optimization frameworks for exploring complex design spaces.

Other studies have investigated the aerodynamic optimization of VAWTs using various geometry modifications. Early efforts demonstrated that blade shape adaptation through morphing or deformation techniques combined with CFD can significantly improve power output across operating conditions with Baghdadi et al. [15] reporting C_P improvements of 18.6%–50%. Subsequent work extended shape optimization to lift-based Darrieus and hybrid VAWT configurations using evolutionary and heuristic optimization approaches, improving performance by tuning rotor-level parameters such as blade shape, radius, height, and blade number; Chen et al. [16] reported a 12.5% improvement in C_P at the optimal wind speed. More recent studies employed high-fidelity CFD to optimize multi-turbine arrangements, helical and hybrid blade designs, and coupled geometric and operational parameters, highlighting the importance of turbine spacing, relative positioning, rotation direction, and blade geometry for power enhancement [17]. Hassanpour [18] reported a 26.6% increase in power generation for an optimized pair of VAWTs relative to a single turbine. Sanaye and Farvizi increased the average power coefficient from 0.058 to 0.1845, corresponding to a 218% increase, while Abdel-razak et al. [19] increased the average C_P from 0.2475 to 0.3159, corresponding to a 27.65% enhancement for an optimized twin hybrid Darrieus–Savonius configuration. While these approaches demonstrated notable performance gains, many relied on intuition-based trial-and-error methods or conventional shape parameterization [20], which limits the potential for truly optimized solutions free from pre-existing biases. Several studies have been confined to traditional blade shapes that either favor drag-based (Savonius) or lift-based (Darrieus) designs, without fully exploring geometries outside these conventional categories. A major constraint that has limited the design scope of previous efforts has been the prohibitive cost of evaluating a large design space using high-fidelity CFD simulations or experiments [21].

Despite successful applications of DbM to other aerodynamic shape-optimization problems, its application to VAWT blade and rotor design remains largely unexplored. Furthermore, the integration of DbM with a sample-efficient global optimization strategy for CFD-based VAWT design has received little attention. In this paper, we address the limitations of previous efforts by introducing a systematic and computationally efficient design approach to VAWT optimization. The aim of this research is to explore non-intuitive combinations of drag-based and lift-based blade shapes using the parameterization method [22]–[23], Design-by-Morphing (DbM), and to optimize them through Bayesian optimization (BO), which is well-suited to costly black-box design spaces. Using CFD simulations validated on both lift-based and drag-based turbines, we evaluate the performance of hybrid blade configurations by computing the power coefficient C_P and assessing their efficiency improvements relative to traditional designs. The key objective of the presented work is to propose a scalable and computationally efficient data-driven framework for VAWT design using DbM and Bayesian optimization, while establishing a generalized methodology for discovering non-intuitive high-performance aerodynamic configurations across broader engineering applications. Given the large number of CFD evaluations required by the optimization framework, two-dimensional unsteady Reynolds-averaged Navier–Stokes (URANS) is employed as a computationally tractable design-screening and relative-ranking model. Accordingly, the reported performance improvements should be interpreted primarily as relative predictions within the common 2D numerical framework, since three-dimensional effects such as tip losses and spanwise flow are not captured.

2. Methodology

This study employs a three-component design methodology for aerodynamic shape optimization of VAWTs using the DbM framework. The main components are: (1) a morphing strategy for generating hybrid VAWT blade and rotor geometries through smooth transformations between conventional lift and drag-based turbines, (2) CFD simulations using the $k-\epsilon$ model with Enhanced Wall Treatment (EWT) to evaluate aerodynamic performance, and (3) optimization by utilizing multi-variable, single-objective Bayesian optimization, a specialized case of the MixMOBO algorithm by Sheikh et al. [24] to identify high-performing blade designs with minimal CFD evaluations. The integrated framework

couples DbM geometry generation with MixMOBO optimization. The optimizer proposes a set of discrete design variables for the VAWT design, which are used by DbM to generate a computer-aided design (CAD) valid blade geometry. This geometry is then evaluated using CFD to compute the objective function. The process is repeated iteratively within the prescribed evaluation budget to identify the best-performing configuration. Each component is detailed in subsequent subsections, and an overview of the methodology is provided in Figure 1.

2.1. Design Space Creation using DbM

Blade-shape parameterization is a central component of the DbM framework for VAWTs. Conventional shape parameterization techniques, such as Free-Form Deformation (FFD), PARSEC, Bézier curves, and B-splines provide flexibility for aerodynamic design but often impose geometric constraints, require a large number of control variables, and limit systematic global modifications [25, 26, 27]. As a result, these approaches tend to restrict the design space to incremental, intuition-driven modifications of existing geometries rather than enabling radical global transformations. This limitation is particularly critical in aerodynamic optimization problems, where performance is strongly dependent on the overall geometry. The DbM approach addresses these shortcomings by reducing high-dimensional geometry into a compact set of design parameters, enabling global, non-local transformations and extrapolation beyond existing designs [22].

DbM, originally introduced by Oh [22], addresses these limitations by reducing high-dimensional geometric representations into a compact set of design parameters that enable global, non-local shape transformations. Instead of directly manipulating control points or local surface parameters, DbM generates new geometries by morphing among a set of topologically compatible baseline shapes. Any two-dimensional shapes bounded by a single continuous surface are homeomorphic and can therefore be mapped onto one another through a set of discrete collocation points. This mapping establishes a one-to-one correspondence between geometries, allowing shapes to be combined smoothly within a continuous design space. A key advantage of DbM is that it allows both interpolation and extrapolation between baseline shapes. While interpolation generates intermediate geometries between existing designs, extrapolation through negative weights can produce entirely new shapes that do not directly resemble any of the original baselines. Although extrapolation may occasionally introduce nonphysical geometrical artifacts such as self-intersections, these can be detected and corrected through geometric conditioning procedures, allowing the design space to remain large without compromising geometric validity.

Initial proof-of-concept studies applied DbM to high-speed train nose optimization, achieving drag reductions of up to 13% compared to conventional designs [22]. Recently, Sheikh et al. [23] showed that DbM could reconstruct the

entire UIUC airfoil database with over 99.5% accuracy using a limited number of baseline airfoils. When combined with numerical simulation and optimization algorithms, the framework enabled the discovery of novel airfoil geometries that simultaneously improved lift-to-drag ratio and stall tolerance. These applications demonstrate the ability of DbM to explore unconventional aerodynamic geometries, motivating its extension to the strongly unsteady and rotating aerodynamic environment of VAWTs.

In this study, DbM enables non-intuitive shape exploration by morphing between homeomorphic baseline shapes: the Darrieus NACA 0021 (lift-based) [28], Hybrid cusped Gottingen (goe) 395 airfoil [29], and Savonius (drag-based) [30] profiles as shown in Figure 2. These baselines are conventional geometries for VAWTs in the literature. Previous DbM applications have shown that this strategy can work well even with synthetic baseline shapes since DbM can negatively weight these shapes, but such shapes can increase the infeasible regions of the design space [31, 23]. In order to keep the dimensionality of the VAWT design space reasonable and maximize the number of candidate designs, we chose existing VAWT geometries as DbM baselines.

By morphing these conventional baseline shapes, DbM creates a flexible design space beyond conventional parameterization, enabling both interpolation and extrapolation of shapes while reducing user-imposed bias. Since simulations are primarily two-dimensional, the baseline shapes are adopted in two dimensions. These geometries must be homeomorphic, with point-to-point correspondence across all blades, ensured through a collocation procedure. Morphed designs are generated by applying weights in spherical coordinates, enabling interpolation and extrapolation of the baseline designs. The DbM design framework is based on an n -sphere formulation, as explained by Lee et al. [32]. The morphing equations for DbM in spherical coordinate system allow mapping from Cartesian to spherical space:

$$\begin{aligned} w_1 &= r \cos(\varphi_1), \\ w_2 &= r \sin(\varphi_1) \cos(\varphi_2), \\ &\vdots \\ w_n &= r \sin(\varphi_1) \cdots \sin(\varphi_{n-1}) \cos(\varphi_n), \\ w_{n+1} &= r \sin(\varphi_1) \cdots \sin(\varphi_{n-1}) \sin(\varphi_n). \end{aligned} \quad (1)$$

Here $\mathbf{w} \in \mathbb{R}^{n+1}$: $\|\mathbf{w}\| = R$, where R is the radius of the VAWT. For the three baseline shapes ($n = 2$), morphing weights (w_1, w_2, w_3) on the n -sphere are:

$$w_1 = \cos(\varphi_1), \quad w_2 = \cos(\varphi_2) \sin(\varphi_1), \quad w_3 = \sin(\varphi_2),$$

yielding the morphed blade through:

$$\mathbf{A}_{\text{morphed}} = \sum_{i=1}^3 w_i \mathbf{A}_i,$$

where $\mathbf{A}_i$ represents the collocated baseline shapes corresponding to the drag-based, goe-395 based, and lift-based profiles. Normalization of the turbine to a radius $R = 0.5\text{m}$ ensures that the L_2 -norm is equal to R , placing all weights




Figure 2: DbM framework for generating hybrid VAWT blade geometries. Collocated baseline blades (Savonius, goe 395, and NACA 0021) are morphed using weights (w_1 , w_2 , w_3) to produce new blade shapes, which are then reconstructed and positioned using geometric parameters dx , dy , blade rotation θ , and TSR in either direction for the CFD simulations

on the half-unit sphere and guaranteeing smooth transitions.

$$\|\mathbf{w}\|_2 = \sqrt{\sum_{i=1}^n w_i^2} = R \quad (2)$$

Collocation forms the basis of DbM and is applied to the VAWT baseline shapes to establish point-wise correspondence between the three blades of baseline geometries, enabling smooth morphing transformations in the DbM framework. The blade curve is parameterized by arc length and uniformly resampled. The upper part of the contour is then mirrored to form a continuous open curve for morphing. After applying the morphing weights, the new contour is mirrored back to close the curve and form the hybrid blade. Negative morphing weights may generate blade profile intersections. Rather than discarding these geometries, they are automatically corrected to preserve design diversity and maximize exploration of the design space.

Morphed blade geometries are generated by varying the spherical weight factors ϕ_1 and ϕ_2 within the range 1° to 179° , sampled using 13 discrete values, $\{1^\circ, 15^\circ : 15^\circ : 165^\circ, 179^\circ\}$. This results in a total of $13 \times 13 = 169$ distinct blade configurations. The endpoints were set to 1° and 179° instead of 0° and 180° to avoid reproducing the baseline shapes. The morphing loop spans drag-based and lift-based geometries, transitioning through hybrid profiles around $\phi_1 = 90^\circ$ and $\phi_2 = 90^\circ$, and producing a wide range of unconventional and intersected blade shapes, including non-intuitive configurations. Following blade morphing, turbine configurations were defined using three geometric parameters normalized by the unit diameter: inter-blade distance $d_x \in \mathcal{D}_x = \{0 : 0.045 : 0.45\}$, vertical placement $d_y \in \mathcal{D}_y = \{0 : 0.25 : 1\}$, and blade tilt/orientation

angle θ . Internally, the optimizer represents blade orientation using 13 angular positions over a 180° branch together with a binary 180° shift, giving 26 encoded states. However, the 180° orientation occurs in both branches, and 360° is physically equivalent to 0° due to angular periodicity. The physical design space therefore contains 24 unique blade orientations spanning $0^\circ \leq \theta < 360^\circ$ in 15° increments. The motivation for discretizing the design space is discussed in Subsection 2.3.

During the morphing process, intersections are categorized into two types: (1) self-intersections where a blade profile overlaps itself due to excessive curvature or distortion, and (2) inter-blade intersections where one blade intrudes into another in a two-bladed configuration. Discarding such cases would reduce design diversity. Hence, a set of Python-based algorithms and CAD correction techniques was developed to detect and resolve these intersections automatically.

Self-intersecting blade geometries generated during morphing were identified and corrected using geometric intersection detection, followed by reordering of coordinate points between intersection locations to restore a continuous contour. Local stiffening was applied by removing closely spaced points near intersections and interpolating between remaining coordinates to eliminate sharp curvature, with the process repeated until all self-intersections were resolved. For more complex cases involving multiple independent intersections, each blade segment was corrected sequentially before addressing cross-segment overlaps. In multi-bladed configurations, inter-blade overlaps were handled through geometric reconstruction in CAD that preserves the outer envelope of the intersecting blades, yielding a unified, physically feasible geometry suitable for meshing, CFD simulation, and subsequent analysis. The complete correction procedure is illustrated in Figure 3. The influence of




Figure 3: Intersection removal process in the DbM framework. (a) Detection of self-intersection regions in the morphed blade geometry (highlighted). (b) Local view of the intersection point showing coordinate directions. (c) Removal of the zero-area region by flipping the intersecting coordinate segment. (d) Stiffening and smoothing of the corrected contour to restore geometric continuity. (e) Transformation from intersecting blade geometry to final non-intersecting blade geometry after correction

the geometric correction procedure is quantified separately in Section 4.3, where the magnitude of the boundary and enclosed-area modifications and their occurrence among the evaluated designs are assessed.

Once the corrected blade geometry is obtained, the morphing parameters and turbine configuration variables define the complete optimization design space. The operational parameter TSR was evaluated in the range from 0.5 to 4.5 based on previous VAWT studies [28, 30], with 13 discrete values $\mathcal{T} = \{0.5 : 0.3333 : 4.5\}$. Another parameter of rotation direction was introduced in terms of turbine orientation of axis of rotation $\hat{a}$, as some shapes rotate in a clockwise direction and some in an anti-clockwise direction based on the resultant force on the blade shape. Together, these parameters enable full coverage of the design space for all morphed blade geometries, while allowing extrapolation beyond interpolation through negative morphing weights, leading to novel hybrid lift-drag VAWT configurations. The resulting blade transformations are illustrated in Figure 2. This yields a unique physical design space $\Phi_1 \times \Phi_2 \times D_x \times D_y \times \Theta \times \mathcal{T} \times \mathcal{A}$ comprising approximately 5.80×10^6 configurations. Table 1 outlines the discrete design levels that define the comprehensive design space. The complete unique physical design space therefore contains

$$\begin{aligned}
 N_{\text{unique}} &= N_{\varphi_1} N_{\varphi_2} N_{d_x} N_{d_y} N_{\theta} N_{\text{TSR}} N_{\hat{a}} \\
 &= 13 \times 13 \times 11 \times 5 \times 24 \times 13 \times 2 \\
 &= 5,800,080 \\
 &\approx 5.80 \times 10^6.
 \end{aligned} \tag{3}$$

Table 1

Summary of design levels for VAWT design variables.

Variable	Design Set
Morphing Parameter φ_1	$\Phi_1 = \{1, 15 : 15 : 165, 179\}$
Morphing Parameter φ_2	$\Phi_2 = \{1, 15 : 15 : 165, 179\}$
Inter-blade Distance d_x	$D_x = \{0 : 0.045 : 0.45\}$
Vertical Placement d_y	$D_y = \{0 : 0.25 : 1\}$
Blade Tilt Angle θ	$0^\circ \leq \theta < 360^\circ, \quad \Delta\theta = 15^\circ$
TSR	$\mathcal{T} = \{0.5 : 0.3333 : 4.5\}$
Rotation Direction $\hat{a}$	$\mathcal{A} = \{1, -1\}$

configurations.

2.2. CFD Solver

CFD simulations have become indispensable for assessing turbine performance during shape parameterization studies; however, they face inherent trade-offs between accuracy and computational cost. Conventional 2D models often overpredict efficiency by up to 32% due to neglected 3D vortex dynamics, while high-fidelity 3D simulations with large-eddy simulation (LES) provide more realistic predictions [33, 34]. Approaches such as RNG and realizable $k-\epsilon$ models [35] and the $k-\omega$ SST formulation [36] have been widely used in VAWT simulations, while hybrid RANS-LES strategies offer a compromise between accuracy and cost [37]. However, the EWT model [38] used with the $k-\epsilon$ model provides a more flexible alternative by smoothly bridging the viscous sublayer and log-law region through a damping function, ensuring accurate near-wall predictions without fine-tuning initialization parameters as in the $k-\omega$




Figure 4: Validation of Savonius (Drag-based) with Blackwell et al. [30] and Darrieus (Lift-based) turbines with Bachant and Wosnik [28].

SST model. This flexible formulation makes EWT particularly attractive for relatively unexplored VAWT applications. This section presents the numerical methodology used to simulate and validate VAWT performance. The objective is to develop an efficient CFD model using EWT that remains robust across a range of TSR values.

The flow is modeled using the incompressible URANS equations. Conservation of mass is expressed as

$$\frac{\partial \bar{u}_i}{\partial x_i} = 0, \quad (4)$$

while conservation of momentum is written as

$$\frac{\partial \bar{u}_i}{\partial t} + \bar{u}_j \frac{\partial \bar{u}_i}{\partial x_j} = -\frac{1}{\rho} \frac{\partial \bar{p}^*}{\partial x_i} + \frac{\partial}{\partial x_j} \left[(\nu + \nu_t) \left(\frac{\partial \bar{u}_i}{\partial x_j} + \frac{\partial \bar{u}_j}{\partial x_i} \right) \right], \quad (5)$$

where $\bar{u}_i$ is the Reynolds-averaged velocity, ν is the molecular kinematic viscosity, ν_t is the turbulent eddy viscosity, and $\bar{p}^*$ is the modified mean pressure defined as

$$\bar{p}^* = \bar{p} + \frac{2}{3} \rho k, \quad (6)$$

where $\bar{p}$ is the mean pressure and k is the turbulent kinetic energy. Turbulence closure is provided by the standard $k-\epsilon$ model. The transport equation for turbulent kinetic energy is

$$\frac{\partial k}{\partial t} + \bar{u}_j \frac{\partial k}{\partial x_j} = \frac{\partial}{\partial x_j} \left[\left(\nu + \frac{\nu_t}{\sigma_k} \right) \frac{\partial k}{\partial x_j} \right] + P_k - \epsilon, \quad (7)$$

where P_k denotes the production of turbulent kinetic energy, ϵ is its dissipation rate, and σ_k is the turbulent Prandtl number for k . Near-wall effects are treated using EWT [38], which combines a two-layer near-wall formulation with enhanced wall functions. The transition between the viscosity-affected and fully turbulent regions is governed by the wall-distance-based turbulent Reynolds number

$$Re_y = \frac{\rho y \sqrt{k}}{\mu}, \quad (8)$$

Table 2

Verification and Validation Error Summary

Verification / Validation Source	Error (%)
Turb. Model 2D Drag ($k-\epsilon$ EWT vs CFD [34])	2.20
Turb. Model 2D Drag ($k-\epsilon$ EWT vs $k-\omega$)	1.00
Turb. Model 3D Drag ($k-\epsilon$ EWT vs Exp.[30])	6.20
Turb. Model 3D Drag ($k-\epsilon$ EWT vs $k-\omega$)	0.33
Mesh Verification Drag	0.03
Mesh Verification Lift	0.91
Time-step verification	0.003
Validation 2D Lift vs CFD [33]	5.00
Validation 2D Lift vs Experiment [28] (Avg.)	16.70
Validation 2D Drag vs Experiment [30] (Avg.)	20.00
Validation 3D Drag vs Experiment [30] (Avg.)	6.50

where y is the distance from the cell center to the nearest wall. For sufficiently fine near-wall meshes, EWT resolves the viscous sublayer while smoothly transitioning to the outer turbulent formulation away from the wall.

The CFD verification and validation methodology was structured systematically, beginning with model verification in a simpler static configuration, followed by mesh and time-step independence studies for rotating cases, and concluding with validation against experimental data for both drag and lift-based designs across the complete range of TSR to ensure the accuracy and reliability of the simulations as summarized in Table 2; further details are provided in Appendix A. Overall, the CFD model using the $k-\epsilon$ EWT provided accuracy comparable to that of the $k-\omega$ SST model within 1% while reducing sensitivity and computational cost, as the optimization required a large number of unsteady CFD simulations. Then, mesh verification was performed using Richardson extrapolation for both drag- and lift-based turbines, resulting in a Grid Convergence Index (GCI) of less than 1%. Since the simulations were unsteady RANS, time-step independence was also evaluated, and a time step of 0.0025 s was selected as optimal, yielding less than 1% error across the full TSR range.

Validation against experimental studies on both lift-based and drag-based turbines confirmed the suitability of the CFD model for the subsequent optimization study, with an average (over TSR range) of 16.7% error for 2D lift-based turbines and an average of 20% error for drag-based turbines. Validation against experimental data of Bachant's NACA 0021 matched the C_p curve with $< 10\%$ error at mid-TSRs as shown in Figure 4. For drag-based turbines, simulations compared well with Blackwell's experiments, with C_p errors $< 15\%$ at most TSRs, though 2D overpredicted at very low TSR and high TSRs due to 2D simulation neglecting 3D turbulence effects. The overall trend of C_p with respect to TSR remains the same, making 2D CFD a tractable solver for optimization.

To identify the source of this error in the 2D analysis, the results were compared with 2D results reported in the literature, resulting in an average of 3% difference. However, validation of 3D CFD using the same CFD model as used in

2D before against experimental data resulted in only 6.5% error. The data are shown in Table 2. These results indicate that 2D simulations tend to overpredict separation effects, leading to relatively higher errors at lower TSR. It was observed that the 2D simulations overpredicted the C_Q , particularly at low TSR, due to the limitation of two-dimensional models in capturing turbulence and tip vortices; however, they still follow the torque trend closely. To balance computational cost and accuracy, 2D simulations were selected for the optimization process. The following section describes the multivariable optimization methodology adopted.

2.3. Bayesian Optimization

The design optimization problem in the present study to find the maximum coefficient of performance C_p , is formally posed as:

$$\mathbf{x}_{\text{opt}} = \arg \max_{\mathbf{x} \in \mathcal{X}} C_p(\mathbf{x}) \quad (9)$$

where $\mathbf{x} \in \mathcal{X}$ is the mixed-variable design vector, defined as $\mathbf{x} = (\varphi_1, \varphi_2, d_x, d_y, \theta, \text{TSR}, \hat{a})$. The corresponding discrete VAWT design space is $\mathcal{X} = \Phi_1 \times \Phi_2 \times D_x \times D_y \times \Theta \times \mathcal{T} \times \mathcal{A}$, with the individual design levels summarized in Table 1.

CFD simulations in VAWT optimization problems are considered to be expensive, hence a cost-effective optimization algorithm is critical. Gradient-based methods converge quickly but risk local optima, while heuristic strategies such as genetic algorithms or particle swarm optimization require many evaluations, especially for multivariable optimization [39, 40, 41, 42]. Bayesian optimization, particularly in multi-objective and mixed-variable algorithms like MixMOBO [24], offers a promising approach by using surrogate models to balance exploration and exploitation of the design space [43, 44, 45, 46]. This section presents the multivariable Bayesian optimization methodology adopted for VAWT blade design. MixMOBO is well-suited for such black-box functions and constructs a Gaussian-process (GP) surrogate model to guide the search using relatively few evaluations. The framework was developed in the previous work of Sheikh [24] and has subsequently been applied to several design and shape optimization problems [47, 48, 23, 32, 31].

The optimization follows an iterative loop: (i) initialization with a small set of CFD-evaluated points, (ii) surrogate model construction, (iii) acquisition function maximization, (iv) objective evaluation by CFD, (v) model update, and (vi) repetition until the prescribed evaluation budget is reached. For this study, we adopt the MixMOBO strategy developed by Sheikh et al. [24], which can optimize a mixed-variable design space using very few function calls, i.e. numerical simulations. In the present implementation, MixMOBO uses a GP surrogate with a radial basis function (RBF) covariance kernel. Candidate selection is performed using the Hedge strategy described by Sheikh [24], which adaptively selects from a portfolio of acquisition functions. The portfolio used in the present study comprises Expected Improvement (EI), Probability of Improvement (PI), GP Upper




Figure 5: Validation of MixMOBO on five standard benchmark functions. Mean normalized reward is plotted with shaded regions representing 0.5 $\times$ standard deviation [24].

Confidence Bound (GP-UCB), and Thompson Sampling (TS), which are commonly used acquisition strategies in Bayesian optimization [49]. GP hyperparameters are determined using leave-one-out cross-validation (LOO-CV), following the MixMOBO implementation [24].

The mixed-variable design space is handled using the mixed-distance covariance treatment described by Sheikh [24]. Prior to surrogate construction, all input variables are scaled to the interval [0, 1], while the observed objective values are standardized to zero mean and unit variance. The CFD objective is treated as deterministic, with no observational noise model introduced; however, a diagonal jitter term of 0.01 is added to the GP covariance matrix for numerical conditioning.

We note here that the optimization problem is not inherently mixed-variable and can be posed in continuous space. Most geometric and operating variables could in principle be treated continuously; however, the present study intentionally discretizes them, as probabilistic global optimizers such as Bayesian optimization balance exploration and exploitation, and may devote excessive evaluations to very small ‘neighborhoods’ when the objective is high-frequency and/or noisy [47]. This necessitates larger budgets to globally optimize an expensive design space. By discretizing the search space and treating variables as integers, this issue is somewhat mitigated, as it prevents the algorithm from spending too much time exploring near-identical data points [32]. Compared to conventional BO algorithms, MixMOBO is efficient with mixed variables, making the optimization more sample efficient.

Before applying the framework to the VAWT optimization problem, the MixMOBO optimization framework was validated on a set of benchmark optimization problems. Five standard test functions [24] were used (Rastrigin, Styblinski-Tang, Rosenbrock, Ackley, and Spherical), where each test function was evaluated over 10 independent runs. The design space dimensionality for these test functions was fixed to be similar to our design problem. For each benchmark run, 50 initial samples were drawn, followed by 180 sequential optimization evaluations, giving a total budget of 230 function evaluations. The same initialization size and

sequential optimization budget were subsequently adopted for the VAWT optimization campaign, i.e., 50 initialization evaluations followed by 180 BO-guided evaluations. The metric used for comparison is the normalized reward, which converges to 1 when the global optimum is reached. It is defined as:

$$\text{Norm. Reward} = \frac{f_{\text{current opt}} - f_{\text{random sample opt}}}{f_{\text{global opt}} - f_{\text{random sample opt}}}$$

Figure 5 shows the mean normalized reward achieved by MixMOBO across the benchmark functions, with shaded regions representing $0.5 \times$ standard deviation (for visualization). The results demonstrate that MixMOBO consistently approaches the global optimum within the prescribed budget, confirming its suitability for our expensive-to-compute black-box VAWT design problem.

2.4. Optimization Setup

The study targets small-scale turbines for urban winds, i.e., $V_\infty \approx 6$ m/s, $Re \approx 4.0 \times 10^5$ [50]. The optimization is conducted under several simplifying constraints. The rotor diameter is fixed at 1 m to provide a normalized reference scale for the design variables, as also explained in Section 2.1. A two-dimensional computational fluid dynamics model is used because it has been validated for relative performance comparisons while remaining computationally efficient for large optimization sweeps. In addition, the study is limited to two-bladed rotors to maintain geometric consistency and reduce the dimensionality of the design space. The optimization results are described in the following sections. The optimization study is performed using an integrated DbM-CFD-MixMOBO framework executed in an iterative loop with two manual checkpoints at mesh generation and post-processing. Figure 6 shows the computational mesh used for the rotating VAWT simulations, including the stationary outer domain and the circular rotating sliding-mesh region surrounding the rotor. Local mesh refinement is applied around the blade surfaces. The enlarged view highlights the refined mesh distribution near the blade surface and tip, where smaller cells are used to capture the local flow features more accurately.

An initial batch of 50 designs was sampled uniformly from the discrete design space and evaluated by CFD to initialize the optimization database and construct the initial GP surrogate model. No fixed random seed was specified for this initialization. Following initialization, one new BO-guided candidate was evaluated per iteration for 180 sequential iterations. The optimization was terminated after this prescribed computational budget, giving a total of $50+180 = 230$ CFD evaluations.

Among these initial evaluations, the DbM generated several non-intuitive blade geometries, highlighting the flexibility of the morphing-based parametrization. The best-performing initialization design corresponded to a hybrid blade geometry defined by $\phi_1 = 60^\circ$ and $\phi_2 = 150^\circ$, with configuration parameters $\theta = 165^\circ$, $dx = 0$, and $dy = 0.75$, achieving a power coefficient of $C_P = 0.1333$ at a TSR




Figure 6: Computational mesh used for the rotating VAWT simulations, showing the stationary outer domain, the circular rotating sliding-mesh region.

of 0.83. The configuration has the features of airfoil-like blade characteristics morphed into a Savonius type turbine, demonstrating the ability of DbM to uncover unconventional yet high-performing designs. This represented the highest-performing point in the initial 50 design database before sequential BO began. At each BO iteration, MixMOBO proposed one new candidate configuration, which was generated using DbM, evaluated by CFD, and returned to the optimizer through the resulting C_P value to update the surrogate model before the next iteration.

3. Results

3.1. Conventional VAWT Configuration Analysis

To evaluate the effectiveness of the proposed optimization framework, two conventional VAWT configurations were selected as benchmark cases: a drag-based Savonius rotor and a lift-based Darrieus rotor. Since both geometries are included as baseline shapes in the DbM framework, they were reconstructed directly through their corresponding morphing angles. The Savonius rotor was generated using $\phi_1 = 0^\circ$ and $\phi_2 = 0^\circ$, while the Darrieus rotor was generated using $\phi_1 = 90^\circ$ and $\phi_2 = 90^\circ$. This allowed the conventional shapes to be reproduced within the same DbM design constraints used for the optimized configurations.

The reconstructed Savonius benchmark adopted an overlap ratio of $s/d = 0.18$, which is 10.0% lower than the $s/d = 0.20$ configuration reported by Blackwell [30] and represents the closest realizable value in the discretized DbM geometry, for which the overlap was found to be aerodynamically beneficial. The Darrieus benchmark used a two-bladed rotor with chord length $c = 0.05$ m, which is 7.4% shorter than the 54 mm NACA 0021 configuration of Bachant et

al. [28]. The small differences in chord length for the lift-based benchmark and overlap ratio for the drag-based benchmark result from the discrete parameter increments used in the DbM design space. Consequently, the performance improvements reported here are calculated relative to the reconstructed internal benchmarks evaluated using the same numerical framework as the best-found design, rather than being interpreted as direct performance differences from the exact experimental geometries. The original literature geometries are considered separately in the CFD validation analysis. Both benchmarks were evaluated at $Re = 4 \times 10^5$ and $D = 1$ m using the same computational domain, mesh strategy, numerical schemes, and two-bladed constraint as the optimization cases.

A complete TSR sweep from 0.5 to 4.5 was performed for both benchmarks to establish a conventional performance standard for calculating the improvement of the optimized design. As shown in Figure 11, the Savonius rotor reached its maximum $C_P = 0.2709$ at TSR = 0.83, then decreased rapidly with increasing TSR, becoming negative beyond TSR = 1.83. This confirms its suitability for low-TSR operation. The Darrieus rotor showed the opposite behavior, producing very low C_P at small TSRs, with $C_P = 0.00498$ at TSR = 0.5 and $C_P = 0.00940$ at TSR = 0.83, but increasing steadily with TSR and reaching its peak $C_P = 0.1847$ at TSR = 3.5. At the Re under consideration, the Savonius baseline outperforms the Darrieus rotor, and the drag-based baseline is used to calculate the performance improvement of the optimized designs in the subsequent sections.

3.2. Optimization Results

Using this conventional VAWT design as the reference, the DbM-MixMOBO framework was then applied to determine how far the aerodynamic performance of a VAWT could be improved through non-intuitive hybrid shape generation and optimization. Across the complete campaign of 230 CFD evaluations, comprising 50 random initialization cases and 180 sequential BO-guided evaluations, MixMOBO identified several promising VAWT configurations. The best-found configuration reached a maximum $C_P = 0.3285$ at TSR = 1.17, representing a 21.26% increase relative to the maximum $C_P = 0.2709$ of the best existing baseline, which was Savonius turbine, as shown in Figure 11. For theoretical reference, the best-found value of $C_P = 0.3285$ corresponds to approximately 55% of the classical Betz limit, $C_{P,Betz} = 16/27 = 0.593$, for an ideal unblocked actuator disk. Because the URANS solution is deterministic for fixed numerical settings, the relevant variability is the periodic cycle-to-cycle variation. Each simulation was advanced through 15 rotor revolutions; the first 5 revolutions were discarded to remove transient development, and the reported power coefficient was obtained by averaging the following 10 revolutions,

$$\overline{C_P} = \frac{1}{N} \sum_{k=1}^N C_{P,k}, \quad N = 10, \quad (10)$$




Figure 7: Comparison of power coefficient (C_P) with epochs. Darrieus benchmark is scaled for visualization.

where $C_{P,k}$ denotes the revolution-averaged power coefficient of revolution k . For the best-found configuration, the retained cycles yielded $\overline{C_P} = 0.3285$, with a standard deviation of 3.07×10^{-3} and a coefficient of variation of 0.93%. After cycle 6, successive cycle-to-cycle differences remained below 0.5%, while the final two revolutions differed by only 0.031%, indicating satisfactory periodic convergence. The optimization history with respect to optimizer epochs is shown in Figure 7. Among the 180 BO-guided CFD evaluations, 129 configurations (71.7%) contained at least one negative morphing weight and therefore represented DbM extrapolation. Among the 28 configurations yielding positive C_P , 22 (78.6%) involved extrapolated morphologies, and all 10 of the highest-performing positive- C_P configurations contained at least one negative morphing weight. Further interpretation of representative negative-power configurations, including their phase-averaged torque behavior, is provided in Appendix A.5. The best-performing configuration identified was obtained after 13 BO-guided evaluations following the 50 point initialization, corresponding to the 63rd CFD evaluation of the complete optimization campaign.

The best-found design is an extrapolated DbM geometry resembling a high-solidity drag-based shape with contours similar to those of a lift-based shape, as shown in Figure 7. The DbM variables in this configuration are an anticlockwise rotation direction with a blade shape rotation angle in DbM as $\theta = 180^\circ$, blade phase angles $\phi_1 = 120^\circ$ and $\phi_2 = 179^\circ$, zero horizontal offset ($dx = 0$), and full vertical placement ($dy = 1$). Azimuthal torque for this best configuration in Figure 8 shows minimal torque near $\theta = 0^\circ$, a sharp rise to 10.01 at $\theta = 115^\circ$, followed by a decline with stall beyond $\theta \approx 135^\circ$, with mirrored behavior from 180° – 360° . The relatively smooth and periodic shape of the curve indicates stable aerodynamic behavior without abrupt changes in torque.




Figure 8: Torque versus azimuth and corresponding FFT for the best VAWT design

The vorticity contours in Figure 10 illustrate the flow mechanisms responsible for this torque variation. In frame (a), as the blade approaches $\theta \approx 90^\circ$, a strong shear layer develops along the outer blade surface and begins to roll up into a coherent vortex. In frame (b), the vortex detaches and is convected downstream, forming the beginning of a periodic vortex street in the wake. In frame (c), the shed vortices interact with the downstream blade position, creating alternating positive and negative vorticity structures that influence the pressure distribution around the rotor. Finally, in frame (d), the vortices move further downstream, and the flow begins to reorganize around the blade and initiates a new vortex-formation cycle. This periodic vortex formation and downstream convection contribute to the stable torque response observed in Figure 8. The hybrid profile delays the performance dip to a slightly higher TSR than Savonius, consistent with added lift-based features associated with a Darrieus type VAWT as shown in Figure 11.

To investigate the unsteady aerodynamic behavior of the optimized VAWT, a Fast Fourier Transform (FFT) was applied to the time-domain torque signal to convert it into the frequency domain as shown in Figure 8. This spectral analysis allows for the identification of dominant frequencies and potential aerodynamic instabilities that may not be evident in the time signal alone. The best-found configuration showed a dominant spectral peak at 4.44 Hz. For the operating condition $\lambda = 1.1667$, $U_\infty = 6.0$ m/s, and $R = 0.5$ m, the rotor frequency is

$$f_{\text{rotor}} = \frac{\omega}{2\pi} = \frac{\lambda U_\infty}{2\pi R} \approx 2.23 \text{ Hz.} \quad (11)$$

For the two-bladed rotor, the corresponding blade-passing frequency is

$$f_{\text{BPF}} = 2f_{\text{rotor}} \approx 4.46 \text{ Hz,} \quad (12)$$

which closely matches the dominant FFT peak at 4.44 Hz. This agreement shows that the principal torque fluctuation is associated with the periodic blade-passing aerodynamic loading.




Figure 9: Mesh sensitivity study for the best-found DbM configuration at $Re = 4 \times 10^5$ and $TSR = 1.1667$. The Richardson extrapolated value is also shown, together with the corresponding GCI values.

Table 3

Timestep sensitivity study for the best-found DbM configuration at $Re = 4 \times 10^5$ and $TSR = 1.1667$.

Δt (s)	($\approx \text{deg}/\text{timestep}$)	C_p	Difference(%)
0.00125	1	0.3298	0.00
0.0025	2	0.3285	0.39
0.005	4	0.2590	21.47

Because the best-found DbM blade differs topologically from the conventional benchmark turbines, an additional spot-check mesh and timestep sensitivity analysis was performed for the best-found configuration, $(\varphi_1, \varphi_2) = (120^\circ, 179^\circ)$, at $Re = 4 \times 10^5$ and $TSR = 1.1667$. As shown in Figure 9, three mesh levels with characteristic mesh sizes $h = 0.006$, 0.004 , and 0.00267 m yielded $C_p = 0.3101$, 0.3285 , and 0.3266 , respectively. The selected medium mesh differed from the finest mesh by only 0.58%, while Richardson extrapolation gave $C_p = 0.3264$, with $GCI_{\text{medium}} = 0.807\%$ and $GCI_{\text{fine}} = 0.084\%$, confirming acceptable spatial convergence for the selected mesh.

A timestep sensitivity study was also carried out using $\Delta t = 0.00125$, 0.0025 , and 0.005 s. At the present operating condition, these correspond to approximately 1° , 2° , and 4° of rotor rotation per timestep, respectively. The resulting values of C_p are summarized in Table 3. The selected timestep of $\Delta t = 0.0025$ s differed from the finer timestep by only 0.39%, whereas the coarsest timestep underpredicted C_p by 21.47%. Accordingly, the selected mesh and 0.0025 s timestep were considered sufficient for the best-found DbM configuration.

3.3. Comparison with Benchmarks

The best-found configuration was compared with the Savonius and Darrieus benchmarks using the same Reynolds number and numerical setup as shown in Figure 11. These benchmark configurations were reconstructed separately from the optimization candidate set using the exact baseline morphing angles and normalized diameter. Performance




Figure 10: Vorticity contours for the best-found shape at (a) $T/8$, (b) $T/4$, (c) $3T/8$, and (d) $T/2$. Contours are symmetric from $T/2$ to T .

improvements were calculated as

$$\Delta C_P = \frac{C_{P,DbM} - C_{P,ref}}{C_{P,ref}} \times 100. \quad (13)$$

Two comparisons are distinguished. First, comparison of the respective maximum C_P values gives a 21.26% improvement over the best benchmark case which was Savonius (0.3285 at $TSR = 1.17$ versus 0.2709 at $TSR = 0.83$) and a 77.8% improvement over Darrieus (0.3285 versus 0.1847 at $TSR = 3.5$). These comparisons therefore use the maximum performance of each turbine and do not correspond to the same TSR . Second, at the same TSR of 1.17, the Savonius turbine gives $C_P = 0.2040$, resulting in an increased improvement of 61.0% for the best-found design. Unlike the conventional baselines, where Savonius performs efficiently only at low TSR and Darrieus becomes effective primarily at high TSR , the best-found geometry exhibits combined drag- and lift-based aerodynamic characteristics and maintains its best performance within the low-to-moderate TSR range. Since the benchmark and best-found configurations were evaluated under the same numerical conditions, the results demonstrate the ability of the DbM–MixMOBO framework to efficiently explore a large, non-intuitive design space and identify configurations that outperform the conventional benchmark designs within the investigated conditions.

The reported 21.26% improvement in maximum cycle averaged C_P should be interpreted within the common 2D URANS numerical framework. The average validation discrepancies of 16.7% and 20.0% for the lift- and drag-based reference turbines, respectively, primarily reflect model-form differences between the two-dimensional simulations and experiments and therefore do not represent statistical uncertainty bounds on the relative improvement. Since the best-found and benchmark configurations were evaluated using the same CFD model, computational domain, numerical schemes, and operating assumptions, the relative comparison is less sensitive to systematic numerical bias than the absolute C_P values. Accordingly, the 21.26% value is reported as a numerically predicted relative improvement rather than an experimentally demonstrated performance gain.

4. Sensitivity Analysis

The best-found design was compared with an intuitive turbine configuration to assess the impact of spacing, as conventional designs and intuition suggest that introducing a gap (s/d) can improve efficiency. To investigate this, the best-found configuration ($s/d = 0$) was compared with overlap ratios $s/d = 0, 0.09$ and 0.18 to understand how the gap affects performance across various $TSRs$. Additionally, the effect of Reynolds number was examined, as typical high-efficiency shapes tend to be lift-based with low solidity.




Figure 11: Comparison of the reconstructed drag-based Savonius benchmark ($s/d = 0.18$), lift-based Darrieus benchmark (scaled for visualization), and the best-found design.

However, the best found design is dominated by drag-based characteristics, with high solidity, making it important to assess its performance at different Reynolds numbers. Since the best found configuration performs best at a Reynolds number of 4×10^5 , this analysis helps to determine the reason for the high solidity of the best-found geometry.

4.1. Effect of Overlap (s/d) Ratio

Overlap ratios (s/d) of 0.09 and 0.18 were introduced in the best design to compare it with the intuitive turbine configuration. As shown in Figure 12, C_p is significantly influenced by the spacing ratio s/d across different TSRs. The best-found configuration with zero spacing ($s/d = 0$) consistently outperforms the spaced configurations across all TSR values. At $\text{TSR} = 1.1667$, C_p reaches its maximum value among the investigated cases, of 0.3285 in the zero-spacing case, compared to 0.1998 for $s/d = 0.09$ and 0.2224 for $s/d = 0.18$, demonstrating that eliminating the gap yields the highest energy capture in the moderate TSR range. For low TSR (0.5 and 0.8333), all configurations perform relatively well, but the best found configuration still maintains a noticeable lead. With increasing TSR, the spaced designs deteriorate rapidly; at $\text{TSR} = 1.8333$ both spaced cases show negative C_p , whereas the optimizer's best design remains positive (0.0431). Overall, spacing negatively impacts performance as rotational speed increases. Thus, comparing the respective best-performing cases, i.e., the maximum C_p attained by each configuration over the investigated TSR range, the zero-spacing best-found design ($C_p = 0.3285$ at $\text{TSR} = 1.1667$) performs 29.4% better than the $s/d = 0.09$ case ($C_p = 0.2538$ at $\text{TSR} = 0.8333$) and 38.8% better than the $s/d = 0.18$ case ($C_p = 0.2366$ at $\text{TSR} = 0.8333$), corroborating the efficiency of the DbM-MixMOBO approach in locating a superior configuration among millions of candidates. Having examined the influence of blade overlap, the sensitivity of the best-found geometry to Reynolds number was assessed next.

To further identify the aerodynamic mechanism responsible for the overlap-ratio sensitivity, the pressure fields were




Figure 12: Comparison of overlap ratios of 0, 0.09, and 0.18 for the best design blade shape.

compared for $s/d = 0, 0.09$, and 0.18 at the same Reynolds number and TSR and at matched rotor phases, as shown in Figure 13. Here, T denotes one complete rotor revolution, with the rows corresponding to $T/8$, $T/4$, $3T/8$, and $T/2$. A common pressure range is used for all configurations. For $s/d = 0$, the blades form a continuous inner contour that limits direct pressure communication through the rotor center and maintains a stronger pressure difference across the torque-producing blade surfaces. When overlap is introduced, an inter-blade passage is formed and a localized pressure depression develops within the gap, particularly around $T/8$ and $T/4$. This pressure equalization reduces the pressure difference available to generate useful aerodynamic torque, consistent with the lower phase-averaged and cycle-averaged C_p of the spaced configurations. Corresponding velocity-field comparisons are provided in Appendix A.4.

4.2. Effect of Reynolds Number

It was observed that the optimization design space leaned towards low TSR at this Reynolds number and a drag-based, i.e., high solidity shape. Therefore, the influence of Reynolds number on the aerodynamic performance of the VAWT was investigated. The impact of Reynolds number on turbine efficiency was first assessed by comparing the lift-based baseline at $Re = 4 \times 10^5$ with the same turbine at $Re = 1.3 \times 10^6$. A TSR sweep for the benchmark Darrieus at both Re shows substantial gains with increasing Re , confirming that lift-based designs benefit strongly from higher Re and higher TSRs. The increased Darrieus performance at higher Reynolds number suggests that re-optimization under higher- Re conditions may favor lower-solidity, more lift-dominant design configurations. This highlights the importance of Reynolds number, i.e., inlet air flow, on VAWT efficiency and optimum blade shape. The detailed performance evaluation of the Darrieus turbine is given in Appendix A. Therefore, the optimizer's best blade shape was investigated at different Reynolds numbers to analyze its effect.

The optimizer's best VAWT blade was evaluated across five Reynolds numbers (2×10^5 to 1.3×10^6) and TSR to assess its efficiency. As shown in Figure 14, the general




Figure 13: Pressure contours for the best-found blade geometry at overlap ratios (a) $s/d = 0$, (b) $s/d = 0.09$, and (c) $s/d = 0.18$, shown at matched rotor phases $T/8$, $T/4$, $3T/8$, and $T/2$ from top to bottom.

trend indicates that the C_p increases with Reynolds number, reaching a peak around $Re = 4 \times 10^5$, followed by a slight decline or fluctuation at higher Reynolds numbers. The maximum C_p of 0.3285 is observed at $Re = 4 \times 10^5$ and $TSR = 1.17$, confirming this as the highest-performing operating condition among the tested cases for this blade design.

Between $Re = 2 \times 10^5$ and $Re = 4 \times 10^5$, C_p improves by 46.7% at $TSR = 1.17$, 12.7% at $TSR = 0.5$, and 12.6% at $TSR = 0.83$. However, at higher Reynolds numbers, the performance either plateaus or declines; e.g., from $Re =$

4×10^5 to $Re = 8 \times 10^5$, C_p drops by 20.0% at $TSR = 1.17$ and 10.5% at $TSR = 1.5$. At $TSR = 1.1667$, the best-found design reaches $C_p = 0.3285$ at $Re = 4 \times 10^5$, compared with approximately $C_p = 0.2836$ at $Re = 1.3 \times 10^6$. The former is therefore approximately 15.83% higher than the latter. At $TSR = 1.83$, C_p becomes negative for several Reynolds numbers, reaching -0.2735 at $Re = 2 \times 10^5$ and -0.2691 at $Re = 8 \times 10^5$. Since $C_p = \lambda C_Q$ and the imposed TSR λ is positive, negative C_p corresponds to a negative cycle-averaged torque coefficient. The mean aerodynamic torque therefore opposes the imposed direction of rotation, such that




Figure 14: Comparison of C_p as a function of TSR for the optimizer's best VAWT blade across different Re .

external mechanical power would be required to maintain the prescribed rotational speed rather than net power being extracted from the flow.

The design remained power-producing up to approximately $TSR = 1.5$ across the investigated Reynolds numbers, while the peak C_p varied noticeably. The highest value, $C_p = 0.3285$, occurs at $Re = 4 \times 10^5$ and $TSR = 1.1667$. The phase averaged C_p analysis was used to identify changes in the magnitude and azimuthal distribution of power production, rather than to resolve the detailed underlying flow structures. The Reynolds-number dependence is non-monotonic within the present 2D URANS framework; resolving its flow-physics origin would require additional pressure, velocity, and vorticity-field analysis beyond the scope of this sensitivity study. These results suggest that the optimizer's best VAWT blade is best suited for moderate Reynolds number operation at $Re = 4 \times 10^5$, aligning with the characteristics of drag-dominant high-solidity designs. The non-monotonic trend beyond $Re = 4 \times 10^5$ should be interpreted as a trend of the present two-dimensional URANS results, which cannot clearly distinguish physical Reynolds-number effects from residual model-form effects. Three-dimensional simulations across multiple Reynolds numbers are therefore needed to confirm this behavior quantitatively.

Finally, Figure 15 compares the C_p response for the Reynolds number and overlap ratio studies of their respective best performing cases. At $TSR = 1.1667$, the cases at $Re = 4 \times 10^5$, 6×10^5 , and 1.3×10^6 show very similar peak locations, around $\psi \approx 120^\circ$ and 300° , with instantaneous C_p peaks of approximately 0.9–1.0. This indicates that Reynolds number mainly changes the magnitude of the aerodynamic loading, with little change in the phase of the dominant power-producing regions. For the overlap cases, the peak instantaneous C_p decreases to approximately 0.55–0.60, while the corresponding cycle-averaged C_p decreases from about 0.328 for the zero-overlap best case to 0.254 and 0.237 for $s/d = 0.09$ and 0.18, respectively. Although




Figure 15: Phase averaged C_p for the Reynolds number and overlap ratio sensitivity cases.

overlap can improve conventional Savonius performance by directing flow through the inter-bucket gap [30], this benefit is not evident for the present DbM geometry, whose shape differs substantially from a conventional Savonius rotor.

4.3. Influence of Geometry Correction Procedure

The influence of the geometric correction procedure was assessed by examining both its occurrence within the design space and the magnitude of the resulting geometric modification. Of the 169 morphologies defined by (φ_1, φ_2) , 64 (37.9%) required self-intersection correction. Among the 180 BO-guided CFD evaluations, 41 (22.8%) used corrected morphologies. Of the 28 evaluations yielding positive C_p , 8 (28.6%) involved correction, corresponding to seven unique morphologies. Only two of the ten highest-performing configurations required correction, while none of the three highest-performing configurations, including the best-found $(\varphi_1, \varphi_2) = (120^\circ, 179^\circ)$ design, required self-intersection correction.

The magnitude of the correction was quantified using three complementary geometric measures: a symmetric root-mean-square (RMS) nearest-boundary distance, a normalized symmetric-area difference, and the Hausdorff distance. Boundary-distance measures based on nearest-point comparisons are commonly used for geometric and shape comparison [51, 52], while the Hausdorff distance provides a measure of the maximum separation between two point sets [53].

Let Γ_1 and Γ_2 denote the uncorrected and corrected boundaries, respectively, with corresponding boundary lengths L_1 and L_2 . The symmetric RMS nearest-boundary distance is defined as

$$d_{\text{RMS}} = \sqrt{\frac{L_1 \langle d(\mathbf{x}, \Gamma_2)^2 \rangle_{\Gamma_1} + L_2 \langle d(\mathbf{x}, \Gamma_1)^2 \rangle_{\Gamma_2}}{L_1 + L_2}}, \quad (14)$$

where $d(\mathbf{x}, \Gamma)$ denotes the shortest Euclidean distance from point $\mathbf{x}$ to boundary Γ , and $\langle \cdot \rangle_{\Gamma}$ denotes averaging over uniformly sampled points along the corresponding boundary. The quantity is normalized by the rotor diameter D as

Table 4

Geometric modification metrics for corrected morphologies yielding positive C_P . Rank refers to their position among the 28 positive- C_P configurations.

Rank	Morphology	RMS change (% D)	Sym. area diff. (%)	Hausdorff (% D)
7	45u150	0.0205	0.2435	0.2238
10	165u45	0.0517	0.7394	0.4058
13	30u179	0.0146	0.9347	0.1566
15, 21	179u60	0.0535	0.7092	0.4195
16	45u179	0.0251	0.7427	0.2369
22	179u45	0.0427	0.6031	0.3487
26	179u165	0.0136	0.5254	0.1390
Mean	–	0.0317	0.6426	0.2758
Maximum	–	0.0535	0.9347	0.4195

$$D_{\text{RMS}}(\%) = 100 \frac{d_{\text{RMS}}}{D}. \quad (15)$$

The change in enclosed geometry was quantified using the normalized symmetric-area difference. This measure corresponds to the area-based Jaccard distance between the two enclosed regions [52]. Because the uncorrected boundary may be self-intersecting, all bounded regions formed by the raw profile were first identified and their union was taken as the pre-correction region Ω_1 . The corrected enclosed region is denoted by Ω_2 . The metric is defined as

$$D_A(\%) = 100 \frac{A(\Omega_1 \triangle \Omega_2)}{A(\Omega_1 \cup \Omega_2)}, \quad (16)$$

where $\triangle$ denotes the symmetric difference. Thus, $D_A = 0$ corresponds to identical enclosed regions, while larger values indicate a greater area modified by the correction.

Finally, the maximum local boundary deviation was quantified using the Hausdorff distance [53],

$$d_H(\Gamma_1, \Gamma_2) = \max \left\{ \sup_{\mathbf{x} \in \Gamma_1} d(\mathbf{x}, \Gamma_2), \sup_{\mathbf{x} \in \Gamma_2} d(\mathbf{x}, \Gamma_1) \right\}, \quad (17)$$

which was normalized by the rotor diameter as

$$D_H(\%) = 100 \frac{d_H}{D}. \quad (18)$$

As summarized in Table 4, the normalized symmetric-area difference remained below 0.94% for all corrected positive- C_P morphologies, with a mean of 0.64%. The RMS boundary modification had a mean of 0.0317% of the rotor diameter D , while the maximum Hausdorff distance was 0.4195% of D . These values indicate that the correction produced localized geometric modifications. Furthermore, corrected geometries represented only two of the ten highest-performing configurations and none of the three highest-performing designs, indicating no systematic tendency of the optimizer toward corrected shape families.

5. Limitations Discussion

The present optimization uses 2D URANS as a computationally tractable design-screening and relative-ranking framework, since the complete campaign required 230 unsteady CFD evaluations. The approach does not capture three-dimensional effects such as tip losses and spanwise flow. In addition, the best-found DbM geometry has not yet been validated experimentally or using 3D CFD. Consequently, the reported 21.26% increase in maximum C_P should be interpreted as a relative numerical improvement within the common 2D URANS framework rather than as an experimentally validated performance gain. Three-dimensional CFD and experimental validation of the best-found geometry are therefore identified as future work. The present optimization is formulated as a single-objective aerodynamic design problem in which the cycle-averaged C_P is maximized. Consequently, the term ‘best-found design’ refers to the highest- C_P configuration identified within the prescribed design space and computational budget, rather than to an overall engineering optimum. Other design criteria, including starting torque, torque ripple, structural and fatigue loads, and acoustic noise, were not included in the objective function. Their incorporation within a multi-objective optimization framework represents an important extension of the present work.

The optimization results are also conditional on the prescribed design space and operating conditions. A single VAWT Bayesian optimization trajectory was performed; therefore, the best-found configuration should not be interpreted as a demonstrated global optimum of the complete design space. In addition, its aerodynamic performance is tuned to the target Reynolds number range, particularly around $Re = 4 \times 10^5$, and different operating conditions may favor different blade characteristics. Numerical sensitivity associated with mesh resolution, timestep, and rotor-cycle convergence is assessed through the verification analyses presented in this study, while remaining model-form uncertainty associated with the 2D approximation requires future three-dimensional CFD and experimental validation.

6. Conclusion

This research presents a data-driven design framework based on the DbM methodology coupled with Bayesian optimization (MixMOBO), applied in the present study to improve the aerodynamic performance of VAWTs. The framework explored a high dimensional design space composed of approximately 5.80 million unique physical configurations and the best-performing configuration was obtained after 13 BO-guided evaluations following the 50 design initialization, corresponding to the 63rd CFD evaluation of the 230-evaluation optimization campaign. Within the common 2D URANS framework, the best-found design achieves higher C_P than both the reconstructed drag-based (Savonius) and lift-based (Darrieus) benchmarks at $Re = 4 \times 10^5$. The benchmark sweep shows that C_P of the best-found design achieved a 21.26% improvement over the

C_p^{max} of Savonius and 77.85% improvement over the C_p^{max} of the Darrieus benchmark. The azimuthal torque of the optimizer's best shape shows hybrid aerodynamic behavior, combining strong drag and lift-based characteristics, with torque rising from -4.96 at $\theta \approx 27^\circ$ to 10.01 at $\theta \approx 115^\circ$. FFT analysis showed a dominant frequency of 4.44 Hz, closely matching the expected two-blade passing frequency of approximately 4.46 Hz, indicating a repeatable smooth periodic aerodynamic loading response.

Blade spacing was also analyzed, showing that the zero-spacing (optimizer's best) case ($s/d = 0$) achieved 29.4% better C_p^{max} than 0.09 s/d and 38.8% better than 0.18 s/d. The Reynolds number sweep further highlighted that the optimizer's best geometry is tuned for moderate Re (4×10^5), performing 15.83% better than at $Re = 1.3 \times 10^6$, where lift-based designs generally dominate. The operating condition $Re \approx 4 \times 10^5$ represents the target small-scale wind regime considered in the present optimization.

Overall, this study establishes a DbM–Bayesian optimization framework for the aerodynamic design of VAWTs, enabling systematic exploration of large, non-intuitive design spaces with a limited number of CFD evaluations. The framework demonstrates how DbM-based geometry generation can be coupled with sample-efficient optimization to identify high-performing configurations beyond conventional blade-shape families. Although demonstrated here for a two-bladed VAWT using primarily two-dimensional URANS simulations, the framework architecture is not specific to this configuration and can be extended to other aerodynamic shape-optimization problems subject to application-specific validation. While the present work uses two-dimensional CFD as a computationally tractable design-screening and relative ranking framework, the absolute performance of the best-found geometry requires future three-dimensional CFD. Three-dimensional simulations would allow additional flow physics, including tip losses and spanwise effects, to be resolved. Future work will also focus on experimental validation of the optimized design and the incorporation of system-level parameters, such as Reynolds number, to further assess performance under realistic operating conditions and broaden the applicability of the proposed approach.

CRedit Authorship Statement

Hamza Latif Mehr: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – Original Draft, Writing – Review & Editing, Visualization. **Imran Akhtar:** Conceptualization, Methodology, Investigation, Resources, Writing – Review & Editing, Supervision, Project administration, Funding acquisition. **Haris Moazam Sheikh:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing – Original Draft, Writing – Review & Editing, Visualization, Supervision, Project administration.

Data Availability Statement

The data supporting the findings of this study are presented within the manuscript. Any additional data related to this work will be made available by the corresponding author upon reasonable request.

References

- [1] A. Hijazi, A. ElCheikh, M. Elkhoury, Numerical investigation of the use of flexible blades for vertical axis wind turbines, *Energy Conversion and Management* 299 (2024) 117867. doi:10.1016/j.enconman.2023.117867.
- [2] V. Sivaram, J. O. Dabiri, D. M. Hart, The need for continued innovation in solar, wind, and energy storage, *Joule* 2 (2018) 1639–1642. doi:10.1016/j.joule.2018.07.025.
- [3] M. Abrams, Vertical-axis wind turbines work well together, 2021. ASME. Available at: <https://www.asme.org/topics-resources/content/vertical-axis-wind-turbines-work-well-together>.
- [4] J. R. R. A. Martins, A. Ning, *Engineering Design Optimization*, Cambridge University Press, 2021. Available at: <https://mdobook.github.io>.
- [5] M. Basumatary, A. Biswas, R. Misra, Cfd study of a combined lift and drag-based novel savonius vertical axis water turbine, *Journal of Marine Science and Technology* 28(1) (2023). doi:10.1007/s00773-022-00907-4.
- [6] B. P. Hand, G. Kelly, A. Cashman, Aerodynamic design and performance parameters of a lift-type vertical axis wind turbine: A comprehensive review, *Renewable and Sustainable Energy Reviews* 139 (2021) 110699. doi:10.1016/j.rser.2020.110699.
- [7] U. Singh, M. Lohumi, H. Kumar, Additive manufacturing in wind energy systems: A review, in: *Advances in Industrial and Production Engineering*, Springer, 2020, pp. 757–766. doi:10.1007/978-981-15-2647-3_71.
- [8] M. M. A. Bhutta, N. Hayat, A. U. Farooq, Z. Ali, S. R. Jamil, Z. Husain, Vertical axis wind turbine – a review of various configurations and design techniques, *Renewable and Sustainable Energy Reviews* 16 (2012) 1926–1939. doi:https://doi.org/10.1016/j.rser.2011.12.004.
- [9] L. Jassim, N. M. Mahdi, J. A. Hussein, M. A. Eleiwi, H. S. Majdi, Numerical and performance analysis of a hybrid savonius–darrieus vertical axis wind turbine with leaf-shaped blades, *International Journal of Heat and Technology* 44 (2026) 342–356. doi:10.18280/ijht.440130.
- [10] S. S. Shijer, H. J. Akeiber, M. S. Kassim, M. A. Eleiwi, H. S. Majdi, Energy consumption of wind turbines mounted on evaporative condenser for energy efficiency improvement, *Journal Européen des Systèmes Automatisés* 58 (2025) 467–478. doi:10.18280/jesa.580305.
- [11] Z. Shen, S. Gong, Z. Zuo, Y. Chen, W. Guo, Darrieus vertical-axis wind turbine performance enhancement approach and optimized design: A review, *Ocean Engineering* 311 (2024) 118965. doi:10.1016/j.oceaneng.2024.118965.
- [12] L. Kuang, R. Zhang, J. Su, Y. Shao, K. Zhang, Y. Chen, Z. Zhang, Y. Tu, D. Zhou, Z. Han, Y. Bao, Y. Cao, Systematic investigation of effect of rotor solidity on vertical-axis wind turbines: Power performance and aerodynamics analysis, *Journal of Wind Engineering and Industrial Aerodynamics* 233 (2023). doi:10.1016/j.jweia.2022.105284.
- [13] H. Peng, B. Zhong, G. Hu, H. Liu, Optimization analysis of straight-bladed vertical axis wind turbines in turbulent environments by wind tunnel testing, *Energy Conversion and Management* 257 (2022) 115411. doi:10.1016/j.enconman.2022.115411.
- [14] A. Fazlizan, W. K. Muzammil, N. A. Al-Khawlani, A review of computational fluid dynamics techniques and methodologies in vertical axis wind turbine development, *Computer Modeling in Engineering & Sciences* 144 (2025) 1371–1437. doi:10.32604/cmescs.2025.067854.
- [15] M. Baghdadi, S. Elkoush, B. Akle, M. Elkhoury, Dynamic shape optimization of a vertical-axis wind turbine via blade morphing

- technique, *Renewable Energy* 154 (2020) 239–251. doi:10.1016/j.renene.2020.03.015.
- [16] Y. Chen, J. Su, Z. Han, Y. Zhao, D. Zhou, H. Yang, Y. Bao, H. Lei, A shape optimization of ϕ -shape darrieus wind turbine under a given range of inlet wind speed, *Renewable Energy* 159 (2020) 286–299. doi:10.1016/j.renene.2020.05.038.
 - [17] S. Sanaye, A. Farvizi, Artificial neural network and genetic algorithm for optimization of helical-blade vertical axis wind turbine, *Cleaner Engineering and Technology* 29 (2025) 101088. doi:10.1016/j.clet.2025.101088.
 - [18] M. Hassanpour, L. N. Azadani, Aerodynamic optimization of the configuration of a pair of vertical axis wind turbines, *Energy Conversion and Management* 238 (2021) 114069. doi:10.1016/j.enconman.2021.114069.
 - [19] M. H. Abdel-razak, M. Emam, S. Ookawara, H. Hassan, Study the performance of a novel design of twin hybrid darrieus–savonius vertical-axis wind turbines integrated with building water storage tanks: 3d optimization study, *Renewable Energy* 256 (2026) 124024. doi:10.1016/j.renene.2025.124024.
 - [20] E. A. Attia, H. Saber, H. E. Gamal, Performance and dynamic characteristics of a multi stages vertical axis wind turbine, *Journal of Vibroengineering* 18 (2016) 4015–4032. doi:10.21595/jve.2016.17027.
 - [21] S. Rasekh, S. K. Aliabadi, Toward improving the performance of a variable pitch vertical axis wind turbine (vp-vawt), part 2: Multi-objective optimization using nsga-ii with cfd in the loop, *Ocean Engineering* 278 (2023) 114308. doi:10.1016/j.oceaneng.2023.114308.
 - [22] S. Oh, Finding the Optimal Shape of an Object Using Design-By-Morphing, Ph.D. thesis, University of California, Berkeley, 2016. Available at: <https://escholarship.org/uc/item/6rd0b8rr>. ProQuest ID: Oh_berkeley_0028E_16470; Merritt ID: ark:/13030/m5xw9gv4.
 - [23] H. M. Sheikh, S. Lee, J. Wang, P. S. Marcus, Airfoil optimization using design-by-morphing, *Journal of Computational Design and Engineering* 10 (2023) 1443–1459. doi:10.1093/jcde/qwad059.
 - [24] H. M. Sheikh, Mixed-Variable Multi-Objective Bayesian Optimization, Design-by-Morphing and their Applications, Ph.D. thesis, University of California, Berkeley, 2022. Available at: <https://escholarship.org/uc/item/66m2b3m9>.
 - [25] G. Anderson, M. Aftosmis, M. Nemec, Parametric deformation of discrete geometry for aerodynamic shape design, *AIAA Paper* 2012–0965 965 (2012). doi:10.2514/6.2012-965.
 - [26] X. Han, D. W. Zingg, An adaptive geometry parametrization for aerodynamic shape optimization, *Optimization and Engineering* 15 (2014) 69–91. doi:10.1007/s11081-013-9213-y.
 - [27] D. A. Masters, N. J. Taylor, T. Rendall, C. B. Allen, D. J. Poole, Review of aerofoil parameterisation methods for aerodynamic shape optimisation, in: 53rd AIAA Aerospace Sciences Meeting, Kissimmee, Florida, AIAA 2015-0761, 2015, pp. 1–13. doi:10.2514/6.2015-0761.
 - [28] P. Bachant, M. Wosnik, Effects of reynolds number on the energy conversion and near-wake dynamics of a high solidity vertical-axis cross-flow turbine, *Energies* 9 (2016) 73. doi:10.3390/en9020073.
 - [29] UIUC Applied Aerodynamics Group, Gottingen 395 airfoil (goe395.dat and goe395.gif), 2025. Department of Aerospace Engineering, University of Illinois at Urbana-Champaign. Available at: <https://m-selig.ae.illinois.edu/ads.html>, accessed April 2025.
 - [30] B. F. Blackwell, R. E. Sheldahl, L. V. Feltz, Wind Tunnel Performance Data for Two- and Three-Bucket Savonius Rotors, Technical Report, Sandia Laboratories, Albuquerque, NM, 1978. doi:10.2514/3.47966, issued for the United States Energy Research & Development Administration. Report No. SAN D76-0131.
 - [31] S. Lee, H. M. Sheikh, Airfoil optimization using design-by-morphing with minimized design-space dimensionality, *Journal of Computational Design and Engineering* 13 (2026) 108–124. URL: <https://doi.org/10.1093/jcde/qwaf124>. doi:10.1093/jcde/qwaf124.
 - [32] S. Lee, D. Lim, G. Gu, P. Marcus, Bayesian-optimized riblet surface design for turbulent drag reduction via design-by-morphing with large eddy simulation, *Journal of Mechanical Design* 146 (2024) 1–16. doi:10.1115/1.4064413.
 - [33] M. S. Siddique, N. Durrani, I. Akhtar, Quantification of the effects of geometric approximations on the performance of a vertical axis wind turbine, *Renewable Energy* 74 (2015) 661–670. doi:10.1016/j.renene.2014.08.068.
 - [34] E. Maican, S.-S. Biris, Comparative analysis of a wind turbine's performances by means of cfd simulations, *Tarım Makinaları Bilimi Dergisi (Journal of Agricultural Machinery Science)* 4 (2008) 247–252.
 - [35] S. A. Orszag, V. Yakhot, W. S. Flannery, F. Boysan, Renormalization group modeling and turbulence simulations, in: *Near-Wall Turbulent Flows*, 1993, pp. 1031–1046.
 - [36] F. R. Menter, Two-equation eddy-viscosity turbulence models for engineering applications, *AIAA Journal* 32 (1994). doi:10.2514/3.12149.
 - [37] O. L. J. Verhoeven, Trailing Edge Noise Simulations using IDDES in OpenFOAM, Master's thesis, Delft University of Technology, 2011. Available at: <http://resolver.tudelft.nl/uuid:f382d5f3-bd63-44fe-aaf5-792e410da95f>.
 - [38] ANSYS Inc., Enhanced wall treatment, 2009. Available at: <https://www.afs.enea.it/project/neptunius/docs/fluent/html/th/node101.htm>.
 - [39] J. Li, M. Zhang, Data-based approach for wing shape design optimization, *Aerospace Science and Technology* 112 (2021) 106639. doi:10.1016/j.ast.2021.106639.
 - [40] B. Bastl, M. Brandner, J. Egermaier, H. Horníková, K. Michálková, E. Turnerová, Gradient-free and gradient-based methods for shape optimization of water turbine blade, in: *Programs and Algorithms of Numerical Mathematics*, 2019, pp. 15–26. Available at: <http://eudml.org/doc/294929>.
 - [41] D. Knight, Design Optimization in Computational Fluid Dynamics, Springer US, Boston, MA, 2009, pp. 666–677. URL: https://doi.org/10.1007/978-0-387-74759-0_121. doi:10.1007/978-0-387-74759-0_121.
 - [42] D. E. Goldberg, Genetic Algorithms in Search, Optimization and Machine Learning, Addison-Wesley Longman Publishing Co., Inc., USA, 1989.
 - [43] Y. Morita, S. Rezaeiravesh, N. Tabatabaei, R. Vinuesa, K. Fukagata, P. Schlatter, Applying bayesian optimization with gaussian process regression to computational fluid dynamics problems, *Journal of Computational Physics* 449 (2022) 110788. doi:10.1016/j.jcp.2021.110788, available at: <https://www.sciencedirect.com/science/article/pii/S0021999121006835>.
 - [44] Z. Vangelatos, H. M. Sheikh, P. S. Marcus, C. P. Grigoropoulos, V. Z. Lopez, G. Flamourakis, M. Farsari, Strength through defects: A novel bayesian approach for the optimization of architected materials, *Science Advances* 7 (2021). doi:10.1126/sciadv.abk2218.
 - [45] S. Gopakumar, S. Gupta, S. Rana, V. Nguyen, S. Venkatesh, Algorithmic assurance: An active approach to algorithmic testing using bayesian optimisation, in: *Advances in Neural Information Processing Systems*, 2018, pp. 1–9. Available at: https://proceedings.neurips.cc/paper_files/paper/2018/file/cc70903297fe1e25537ae50aea186306-Paper.pdf.
 - [46] D. Golovin, B. Solnik, S. Moitra, G. Kochanski, J. Karro, D. Sculley, Google vizier: A service for black-box optimization, in: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2017, pp. 1487–1495. doi:10.1145/3097983.3098043.
 - [47] H. M. Sheikh, T. A. Callan, K. J. Hennessy, P. S. Marcus, Optimization of the shape of a hydrokinetic turbine's draft tube and hub assembly using design-by-morphing with bayesian optimization, *Computer Methods in Applied Mechanics and Engineering* 401 (2022) 115654. URL: <https://www.sciencedirect.com/science/article/pii/S0045782522006090>. doi:https://doi.org/10.1016/j.cma.2022.115654.
 - [48] H. M. Sheikh, T. Meier, B. Blankenship, Z. Vangelatos, N. Zhao, P. S. Marcus, C. P. Grigoropoulos, Systematic design of cauchy symmetric structures through bayesian optimization, *International Journal of Mechanical Sciences* 236 (2022) 107741. doi:10.1016/j.ijmecsci.2022.107741.

- [49] E. Brochu, V. M. Cora, N. de Freitas, A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning, arXiv preprint arXiv:1012.2599 (2010). arXiv:1012.2599.
- [50] S. V. Hudisteanu, C. G. Popovici, N. C. Chereches, Wind tunnel study of natural ventilation of building integrated photovoltaics double skin façade, in: E3S Web of Conferences, volume 32, 2018, p. 01020. doi:10.1051/e3sconf/20183201020.
- [51] M.-P. Dubuisson, A. K. Jain, A modified hausdorff distance for object matching, in: Proceedings of the 12th International Conference on Pattern Recognition, volume 1, 1994, pp. 566–568. doi:10.1109/ICPR.1994.576361.
- [52] A. A. Taha, A. Hanbury, Metrics for evaluating 3d medical image segmentation: Analysis, selection, and tool, BMC Medical Imaging 15 (2015) 29. doi:10.1186/s12880-015-0068-x.
- [53] D. P. Huttenlocher, G. A. Klanderman, W. J. Rucklidge, Comparing images using the Hausdorff distance, IEEE Transactions on Pattern Analysis and Machine Intelligence 15 (1993) 850–863. doi:10.1109/34.232073.
- [54] P. K. Talukdar, A. Sardar, V. Kulkarni, U. K. Saha, Parametric analysis of model savonius hydrokinetic turbines through experimental and computational investigations, Energy Conversion and Management 158 (2018) 36–49. doi:10.1016/j.enconman.2017.12.011.

Appendix

A. Computational Fluid Dynamics Simulation Setup

ANSYS Fluent was used for CFD analysis, where a sliding mesh approach was employed, with a rotating inner domain of $1.2D$ inside a stationary outer domain. Domains extended $10D$ downstream and $3D$ upstream, with $6D$ lateral spacing to minimize blockage as shown in Figure 16. The domain for 3D CFD analysis is shown in (a), high-fidelity Savonius cases included end plates of thickness $1.1D$ [54], 2D Savonius geometry at (b), 2D Darrieus geometry following Bachant et al [28] at (c), and finally 2D Darrieus geometry following Siddique et al. [33] at (d). The outer domain is the same, whereas the difference is only in the blade shape and the inner domain. Inlet velocity and pressure outlet boundaries were applied. Due to the rotational motion of the turbine and the continuously changing relative velocity around the blades, the local Reynolds number varies throughout the rotation. To maintain consistency and comparability across simulations, an average Reynolds number is used, calculated based on the inlet velocity and the turbine diameter.

For all Savonius type cases, the Reynolds number is approximately 4.32×10^5 [30]. In the lift-based Darrieus validation case by Siddique et al [33], the Reynolds number is also 4×10^5 . However, for the cases based on Peter Bachant's configurations, the Reynolds number is significantly higher, around 1.1×10^6 [28]. Validation and verification cases included Savonius [30, 34] and Darrieus turbines [33, 28]. Table 5 summarizes all verification and validation cases performed for both drag-based (Savonius) and lift-based (Darrieus) VAWTs. Each case targets a specific aspect of model fidelity, including turbulence modeling, mesh independence, timestep sensitivity, and comparison with CFD and experimental data (validation). Their CAD domains are

shown in Figure 16 and their schematics are shown in Figure 17.

Definition of VAWT parameters. The overlap ratio s/d is defined as the overlap distance s normalized by the bucket (1 blade) diameter d . The rotor diameter is denoted separately by $D = 2R$. A value of $s/d = 0$ indicates no overlap, while $s/d = 0.2$ corresponds to a 20% normalized spacing between blades. The angle $\alpha = 5^\circ$ denotes a static blade angle of attack used only for model verification in non-rotating conditions. Static cases isolate aerodynamic loading without rotational effects, allowing turbulence model validation. TSR is defined as $\lambda = \omega R/V_\infty$, governing dynamic turbine operation.

Case 1 – 2D Static (Maican). A two-dimensional static simulation at $\alpha = 5^\circ$, $s/d = 0.2$, and $\text{TSR} = 0.1586$ was conducted for turbulence model verification against literature CFD. This case validates the EWT implementation under controlled, non-rotating conditions.

Case 2 – 3D Static (Blackwell, Run 12). A three-dimensional static validation at $\alpha = 5^\circ$ and $s/d = 0$ compares simulated torque with experimental data. This builds confidence in the 3D EWT model before dynamic simulations.

Case 3 – 2D Dynamic (Blackwell, Run 37). Dynamic simulations at multiple TSR values with $s/d = 0.2$ are used for timestep verification, mesh independence assessment, and experimental validation. This case ensures temporal and spatial convergence of the rotating solver.

Case 4 – 3D Dynamic (Blackwell, Run 37). A full 3D rotating simulation at multiple TSR values validates the dynamic torque against experiment, confirming improved wake resolution and reduced separation overprediction relative to 2D.

Case 5 – 2D Dynamic NACA 0022 (Siddique et al.). A lift-based validation at $\text{TSR} = 3.0$ compares blade and overall torque with CFD literature results. This verifies aerodynamic lift prediction for rotating airfoil-based turbines.

Case 6 – 2D Dynamic NACA 0021 (Bachant et al.). Validation across multiple TSR values compares predicted C_p with experimental data at high Reynolds number. This case includes mesh verification and confirms the solver's ability to capture lift-dominant turbine behavior.

Simulations are primarily two-dimensional, incompressible URANS, with selected three-dimensional cases for validation, solved using the Finite Volume Method. Turbulence is modeled using the standard $k-\epsilon$ with EWT, which resolves near-wall regions without strict y^+ constraints ($y^+ = 1$) under $y^+ < 5$. The two transport equations are supplemented by EWT blending functions [38]. For comparison, the $k-\omega$ SST model [36] was also applied, though it requires fine-tuning of initialization parameters and a higher mesh sensitivity for each TSR and Re . Numerical solutions used the SIMPLE algorithm with second-order schemes and convergence criteria of 10^{-5} .




Figure 16: Domain for (a) Case 2 & 4, (b) Case 1 & 3, (c) Case 6, and (d) Case 5.

Table 5

Verification and Validation cases for Savonius and Darrieus VAWTs

Type of VAWT	Case #	Validation case type	Validation parameter w.r.t. C_Q & C_P	Case parameter
Savonius Two-bucket (Drag Type) VAWT	1	2D Static, Maican [34]	$5^\circ \alpha$, $s/d = 0.2$, TSR = 0.1586	Model Verification
	2	3D Static, Blackwell (Run 12) [30]	$5^\circ \alpha$, $s/d = 0$, TSR = 0.1586	Model Verification
	3	2D Dynamic, Blackwell (Run 37) [30]	Multiple TSR, $s/d = 0.2$	(i) Timestep Verification (ii) Mesh Verification (iii) Validation with Experiment
	4	3D Dynamic, Blackwell (Run 37) [30]	Multiple TSR, $s/d = 0.2$	Validation with Experiment
Darrieus (Lift Type) VAWT	5	2D NACA 0022 Dynamic, Siddique et al. [33]	TSR = 3.0	Validation with CFD
	6	2D NACA 0021 Dynamic, Peter Bachant [28]	Multiple TSR	(i) Mesh Verification (ii) Validation with Experiment

A.1. Independence Studies

First model independence (Case #1) was tested using both $k-\omega$ SST and $k-\epsilon$ EWT. Simulations were compared with static (non-rotating) VAWT torque data. In 2D static validation (Maican [34]), C_Q showed $< 3\%$ error, confirming EWT suitability. A 3D static case (Case#2) yielded $< 7\%$ deviation from experiment, which was comparable with $k-\omega$ SST with only an absolute difference of 0.3%.

Mesh independence (Case #3 and 6) was achieved at $\sim 0.55M$ elements for both Savonius and Darrieus as shown in Figure 18. To rigorously assess mesh convergence, Richardson extrapolation was performed for the torque

coefficient (C_Q) at TSR = 0.1586 using three systematically refined meshes (0.41M, 0.55M, and 0.87M elements) for the Savonius Turbine. The extrapolated solution was found to be $C_Q^{\text{ext}} \approx 0.61004$, with a GCI of approximately 0.03% between the medium and fine meshes. These results confirm that the medium mesh (0.55 million elements) provides sufficient accuracy for the present simulations, offering a good balance between computational cost and solution fidelity, while timestep independence (Case #3) at 0.1586 TSR converged at $\Delta t = 0.005$ s for lower TSRs, and for higher TSRs it was found to be 0.0025 s. Hence, 0.0025 s was chosen for the timestep (Table 6). Following CFD model




Figure 17: Schematic for (a) Savonius and (b) Darrieus turbines




Figure 18: Mesh independence for Savonius (Drag-based) and Darrieus (Lift-based) turbines

Table 6
Timestep sensitivity analysis for C_Q at TSR = 0.1586

Timestep (s)	CFD (Present)	Percentage Error (%)
0.01	0.61005	—
0.005	0.6102	0.0246
0.0025	0.61022	0.0033

verification, CFD validation against experimental data was conducted.

A.2. CFD validation

Validation was performed for both Drag-based and Lift-based turbines across the range of TSRs. Two validation cases were performed for lift-based turbines. First, simulations were validated with CFD of Siddique et al. [33] (Case #5) within 5% torque error as shown in Figure 19. Then, validation against experimental data of Bachant's NACA 0021 (Case #6) that matched the C_P curve with < 10% error at mid-TSRs and for drag-based turbines, Savonius simulations that also compared well with Blackwell's experiments, with C_Q errors < 15% at most TSRs as shown in section 2.2. To prove robustness of the model and the identification of overprediction of torque, a 3D Savonius (Case #4) was analyzed with 36 million cells, which reduced error to below 9% (Table 7), confirming the robustness of the EWT approach and the limitations of 2D analysis.




Figure 19: Overall (top) and blade torque (bottom) comparison with respect to azimuth angle for the present CFD results (solid line) and CFD results from Siddique et al. [33] (dashed line).

Table 7
Comparison of C_Q from 3D simulations with Blackwell [30] at TSR = 0.1586 and 0.4173.

TSR	C_Q Exp.	C_Q CFD (Present)	Error (%)
0.1586	0.382	0.365	4.4
0.4173	0.373842	0.3414	8.6

A.3. Study on Darrieus turbine at variable Reynolds number

The impact of Reynolds number on the benchmark Darrieus turbine was examined by comparing performance at the baseline condition of $Re = 4 \times 10^5$ with a higher Reynolds number of $Re = 1.3 \times 10^6$. At $Re = 4 \times 10^5$, the turbine achieved a peak C_P of 0.1847 at TSR = 3.5, which is considerably lower than the performance reported by Bachant et al. [28] under higher Reynolds number conditions. To systematically assess this discrepancy, a TSR sweep was conducted at both Reynolds numbers. The results show a pronounced improvement in performance with increasing Re , particularly in the high TSR regime, where the peak C_P increases from 0.1847 to 0.3223, corresponding to a 74.5%




Figure 20: Comparison of Reynolds number for the benchmark Darrieus turbine at $Re = 4 \times 10^5$ (circles with solid line) and $Re = 1.3 \times 10^6$ (squares with dashed line).




Figure 21: Phase-averaged torque coefficient $C_Q(\theta)$ at representative operating points, obtained by averaging the final two complete rotor revolutions. The horizontal line indicates $C_Q = 0$; negative values correspond to aerodynamic torque opposing the imposed direction of rotation.

rise. This confirms the established trend that lift-based, low-solidity turbines benefit significantly from higher Reynolds numbers due to enhanced aerodynamic lift and delayed separation. In contrast, the DbM-optimized configuration, which exhibits high-solidity and drag-dominant characteristics, is tuned for moderate Reynolds number operation around $Re = 4 \times 10^5$, where such geometries demonstrate superior performance at lower TSRs.

A.4. Velocity-field comparison for overlap-ratio sensitivity

Figure 22 compares the corresponding velocity-magnitude fields for $s/d = 0, 0.09$, and 0.18 at the same Reynolds number, TSR, and rotor phases. The spaced configurations show localized acceleration through the inter-blade passage, particularly during the early phases of the cycle. This accelerated gap flow is consistent with the pressure-equalization mechanism identified from the pressure contours in Figure 13, indicating that part of the flow is redirected through

the gap rather than contributing to the pressure difference across the torque-producing blade surfaces.

A.5. Torque-coefficient distribution at negative-power operating points

To further examine the negative-power operating points encountered during the optimization and sensitivity analyses, the phase-averaged torque coefficient $C_Q(\theta)$ was evaluated over the final two complete revolutions for representative cases, as shown in Figure 21. Because the tip-speed ratio λ is positive and $C_P = \lambda C_Q$, the sign of the cycle-averaged C_P is determined directly by the cycle-averaged aerodynamic torque coefficient. Negative C_Q therefore represents an instantaneous aerodynamic torque acting opposite to the prescribed direction of rotation. Such behavior is not associated with a single operating parameter alone, but may result from the combined effects of TSR, candidate blade morphology, and rotation direction.

The torque distributions illustrate these different loading regimes. For the $TSR = 2.8333$ case, C_Q remains negative throughout essentially the entire rotor cycle, whereas at $TSR = 3.1667$ limited positive-torque intervals are outweighed by broader negative-torque regions. In contrast, the best-performing case at $TSR = 1.1667$ develops strong positive torque lobes that dominate the cycle and produce positive net power extraction. During Bayesian optimization, such negative-power cases are expected because Mix-MOBO explores combinations throughout the prescribed design space, including sub-optimal morphologies and rotation directions that may oppose the aerodynamically preferred rotation of a given blade geometry. Their evaluation is therefore part of the exploration–exploitation process used to distinguish unfavorable regions of the design space from high-performing configurations. The negative cycle-averaged C_P values should consequently be interpreted as the result of net aerodynamic torque opposing the imposed rotation for those particular geometry–operating-condition combinations, rather than being attributed to a single mechanism such as backflow.




Figure 22: Velocity-magnitude contours for the best-found blade geometry at overlap ratios (a) $s/d = 0$, (b) $s/d = 0.09$, and (c) $s/d = 0.18$, shown at matched rotor phases $T/8$, $T/4$, $3T/8$, and $T/2$ from top to bottom.