Proceedings of CONF-MCEE 2026 Symposium: Geomaterials and Environmental Engineering
DOI: 10.54254/2755-2721/2026.MH31014

Machine Learning in Wind Turbine Blade Aerodynamics: Methods, Applications, and Future Directions

Chenleqian Man

College of Packaging Engineering, Jinan University, Guangzhou, China
Manchenleqian12@stu2023.jnu.edu.cn

Abstract. The aerodynamic design of wind turbine blades plays a decisive role in energy capture, structural loading and overall cost of wind energy, yet conventional design workflows remain heavily reliant on Blade Element Momentum (BEM) theory and high-fidelity computational fluid dynamics (CFD). As rotor sizes and design requirements become more demanding, these physics-based tools make high-dimensional, multi-objective optimization increasingly expensive. Over roughly the last five years, machine learning (ML) has emerged as a promising way to alleviate these limitations by providing fast, flexible surrogates and new forms of inverse and data-driven design. This review summarizes recent progress in machine-learning-based aerodynamic optimization of wind turbine blades, with a particular emphasis on wind-turbine-specific airfoils. We first recall the fundamentals of blade aerodynamics and the structure of conventional BEM/CFD-based optimization workflows. We then discuss the main classes of ML frameworks used in this context, and review their concrete applications to airfoil-level modeling and design, focusing on surrogate models for predicting lift and drag coefficients and surrogate-assisted optimization of airfoil shapes. These airfoil-level surrogates are highlighted as reducing the caculation cost and improving the optimization efficiency compared with the single-fidelity surrogate model. The review concludes by outlining key challenges, including data quality, generalization, physics consistency and multidisciplinary integration, and by identifying promising research directions towards physics-informed, uncertainty-aware and multi-fidelity ML frameworks for the aerodynamic design of next-generation wind turbine blades.

Keywords: Wind turbine blade aerodynamics, Machine learning, Surrogate modeling, Aerodynamic optimization.

1. Introduction

Wind energy has become one of the most rapidly growing renewable energy sources worldwide, and modern wind turbines are continuously increasing in size, especially in offshore applications. Serving as the primary medium for utilizing wind energy, the aerodynamic design of rotor blades plays a central role in determining power capture, structural loading, and overall cost of energy [1]. Achieving high aerodynamic efficiency over a wide range of operating conditions, while controlling loads and noise, remains a key challenge for next-generation turbines [2].

DOI: 10.54254/2755-2721/2026.MH31014

Traditionally, blade aerodynamic design has relied on a combination of low-order models such as blade element momentum (BEM) theory and high-fidelity computational fluid dynamics (CFD). Inverse design and optimization are usually carried out by coupling these aerodynamic solvers with gradient-based algorithms, adjoint methods, or population-based metaheuristics such as genetic algorithms and particle swarm optimization. Although these approaches have led to several successful commercial blade families, they are computationally demanding, especially when three-dimensional effects, unsteady inflow, and aeroelastic constraints are considered. As rotor diameters grow into the 10–20 MW class, straightforward CFD-driven optimization quickly becomes prohibitively expensive, limiting the size of the design space and the number of design iterations that can be explored in practice [3].

Recent advancement in aerodynamic optimization of wind turbine continue to improve the predictive optimization of the blade design through machine learning. Data-driven surrogate models, trained on CFD, BEM, or aeroelastic simulations, can emulate the mapping between blade design parameters and aerodynamic performance at a fraction of the original computational cost [4]. Deep learning architectures, including variational autoencoders and attention-enhanced networks, have been used for inverse design and performance prediction of wind-turbine airfoils and blade sections [5]. At the same time, reinforcement learning and other sequential decision-making methods are being explored to optimize control strategies that interact with blade aerodynamics in real time, such as pitch, torque, yaw, and even direct aerodynamic optimization of blade shapes [6].

The application of machine learning to aerodynamic optimization in wind turbine remains a dynamic and emerging field,crtical to the ongoing success of wind energy. This review aims to delve deeply into these advancements, exploring their implications for the future of wind energy generation.This introduction set an in-depth exploration of several critical aspects of wind turbine technology. The remainder of this paper is organized into two main sections. Section 2 briefly reviews the fundamentals of wind turbine blade aerodynamics and outlines the conventional BEM- and CFD-based design workflow, highlighting why high-dimensional, multi-objective optimization remains computationally demanding. Section 3 indicates the core of the review and is devoted to machine-learning-based aerodynamic optimization. We first introduce the main classes of machine learning frameworks used in this context, including surrogate models, deep-learning-based inverse design and, where relevant, reinforcement-learning-inspired optimization strategies. Section 4 presents the core of the review and is devoted to machine-learning-based aerodynamic optimization of wind-turbine airfoils. We first summarize supervised and deep-learning surrogate models for predicting lift and drag coefficients and their ratio as functions of airfoil geometry and operating conditions, and then discuss surrogate-assisted shape optimization strategies, with particular emphasis on multi-fidelity convolutional neural network surrogates trained via transfer learning that combine low- and high-fidelity data. Finally, section 5 concludes with a synthesis of current capabilities and limitations, and an outlook on future research directions for machine-learning-enhanced aerodynamic design of next-generation wind turbine blades.

2. Basic principles and framework

This section gives an exploration of the fundamental aerodynamic principles that form the basis of the wind turbine design and efficiency. From an aerodynamic point of view, a modern rotor can be regarded as a set of spanwise blade elements, each operating approximately as a two-dimensional airfoil embedded in a three-dimensional, rotating flow.The wind turbine extract kinetic energy from the wind by converting it into torque on a rotating rotor. The overall aerodynamic performance of the turbine, usually expressed in terms of the power coefficient C_P and thrust coefficient C_T , is

DOI: 10.54254/2755-2721/2026.MH31014

primarily governed by the tip-speed ratio λ, defined as the ratio between the blade tip speed and the free-stream wind speed, as well as by the detailed geometry of the blades. Figure 1 shows the detailed schematics which foucses on the construction detail of the wind turbine blade.

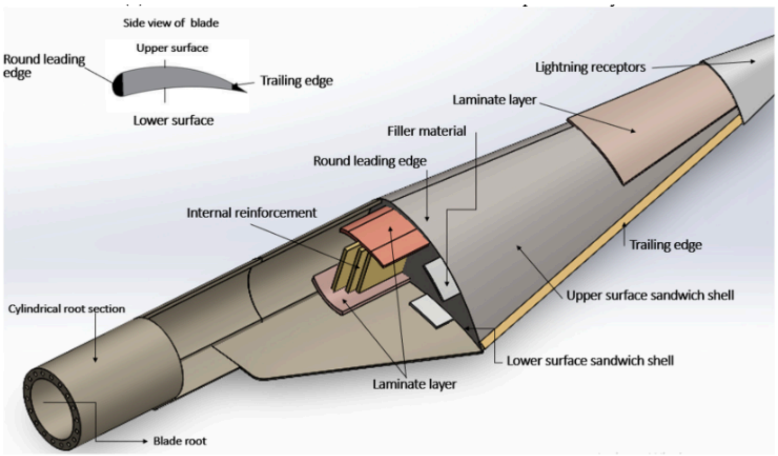

Figure 1. Detailed schematics of offshore wind turbine and blade construction

For a given operating condition, the relative flow velocity seen by a blade element is obtained by combining the incoming wind speed with the local tangential rotational speed of the blade. This defines a local inflow angle ϕ, from which the angle of attack α is computed as the difference between ϕ and the local pitch or twist angle β. The sectional lift and drag forces are then determined by the airfoil's aerodynamic characteristics, typically represented by lift and drag coefficients $C_L(\alpha, Re)$ and $C_D(\alpha, Re)$, which depend on angle of attack and Reynolds number. These sectional forces are integrated along the span to give the global rotor thrust and torque.

Blade Element Momentum (BEM) theory, introduced by Glauert [7], has long been the workhorse for rotor aerodynamic analysis, combining one-dimensional momentum theory with sectional blade-element aerodynamics. BEM provides a computationally efficient way to relate blade geometry and operating conditions to global performance metrics. Furthermore, its accuracy depends on empirical corrections for tip and root losses, dynamic stall, and yawed inflow. For more detailed analyses, especially when considering unsteady inflow, complex inflow shear, or interaction with the tower and wake, high-fidelity computational fluid dynamics (CFD) simulations based on Reynolds-averaged Navier–Stokes (RANS) or large-eddy simulation (LES) are employed. These models can resolve three-dimensional flow structures and provide more accurate pressure and load distributions, at the cost of a much higher computational expense.

3. Machine-learning frameworks for 3D wind-turbine blade aerodynamic optimization

Machine-learning methods have not replaced classical CFD-based design workflows, but they increasingly act as a "smart layer" on top: accelerating the evaluation of candidate blade geometries, enabling inverse design, and closing the loop between aero-loads and control in a 3D rotor context. In this section, we classify recent work into four representative categories: (i) surrogate-based optimization frameworks for full 3D blades, (ii) deep-learning–based inverse or generative design, (iii) reinforcement-learning approaches for aero-control of the rotor, and (iv) multi-fidelity and transfer-learning strategies used to reduce cost and improve robustness. Throughout, we emphasize how these frameworks interface with 3D blade aerodynamics rather than isolated 2D airfoils.

DOI: 10.54254/2755-2721/2026.MH31014

3.1. Surrogate-based aerodynamic optimization of 3D blades

A first family of approaches uses supervised learning surrogates to replace or augment expensive high-fidelity aerodynamic simulations in the optimization loop. The typical pipeline is: (1) generate a database of 3D blade designs and associated performance metrics using CFD and/or BEMT-based aeroelastic tools; (2) train a regression surrogate (e.g., radial basis functions, Kriging, Gaussian processes, or neural networks) to map geometry/design variables to aerodynamic or system-level quantities of interest; and (3) run a global optimizer (often a genetic algorithm or gradient-based method) on this surrogate, with occasional high-fidelity calls for correction and validation.

A recent innovation is the aerodynamic design of a 15 MW floating offshore rotor tailored to Mediterranean conditions which used the multi-fidelity surrogate-based framework [8]. high-fidelity CFD in the innovation is used to obtain accurate sectional aerodynamics, while a BEMT-based rotor model evaluates overall performance. The link between these models and the optimizer is provided by a stochastic radial basis function surrogate of the cost function, which is iteratively refined in regions of interest. This allows them to explore chord/twist distributions and thick root profiles for a 3D blade under realistic offshore conditions with a fraction of the CFD cost. For your review, this paper can serve as a representative case of how supervised surrogate models are embedded into a full 3D aerodynamic design loop, including how the training set is constructed, how the surrogate's accuracy is monitored, and how the optimizer navigates the surrogate space.

3.2. Deep-learning inverse and generative design for blade aerodynamics

A second line of research uses deep neural networks and generative models to perform inverse or generative design, often starting at the airfoil level but with clear implications for 3D blades. Instead of manually choosing shape parameters and running forward simulations, the idea is to learn a mapping from performance targets such as pressure distribution and lift–drag profile to geometry, or to learn a low-dimensional latent space that encodes physically realistic aerodynamic shapes.

Recent research showacases a two-step deep-learning inverse design framework for a wind-turbine airfoil that designed to illustrate this paradigm [5]. This innovation first train a variational autoencoder (VAE) on pressure distributions to learn a compact latent representation of feasible aerodynamic loads. Then, a multi-layer perceptron (MLP) is trained to map from this latent space to both quantities of interest and geometric parameters. During optimization, the design variables are latent coordinates rather than direct geometric parameters, which helps enforce realism and continuity of the pressure distribution. Active learning and transfer learning are further used to refine the surrogate as new high-fidelity samples are added.

3.3. Reinforcement learning for 3D rotor aero-control and performance optimization

Most shape optimization studies treat the blade geometry as fixed and focus on finding an optimal static shape. In contrast, reinforcement learning (RL) tackles aerodynamic optimization from a control perspective, actively adjusting the 3D rotor's operating conditions (pitch, yaw, and rotational speed) to maximize energy capture or reduce loads under turbulent inflow. Here the "design variables" are not geometric parameters but control policies that interact with the full 3D aerodynamics of the rotor.

A clear example proposed by coupling a double deep Q-learning agent with a blade-element momentum (BEMT) model of a wind turbine [9]. The RL agent observes turbine states and learns control actions for rotor speed, blade pitch, and yaw angle, with the reward designed to maximize

DOI: 10.54254/2755-2721/2026.MH31014

energy generation under changing wind conditions. The trained controller is then tested on turbulent and real wind profiles and is shown to outperform both a traditional PID controller and a value-iteration-based RL controller in annual energy production and adaptability to gusty winds.

4. Machine-learning surrogate applications for wind turbine airfoils

In the context of wind turbine blades, the aerodynamic behavior of each spanwise section is often reduced to a 2D airfoil problem characterized by lift and drag coefficients over a range of angles of attack and Reynolds numbers. Machine-learning-based surrogates have therefore become a key enabler for accelerating both performance evaluation and shape optimization at the airfoil level. This section focuses on two main application lines: (i) surrogate models for predicting $\boldsymbol{C_L}$, $\boldsymbol{C_D}$ and related performance indicators, and (ii) surrogate-assisted optimization of airfoil shapes for wind turbine applications.

4.1. Surrogate models for CL/CD prediction

The most direct use of machine learning in airfoil aerodynamics is to learn a regression surrogate that predicts lift and drag behavior from geometric and flow descriptors. For wind turbine design, such surrogates are attractive because they can be queried orders of magnitude faster than CFD or XFOIL, yet still provide accurate estimates of C_L , C_D or their ratio C_L/C_D , which are needed when constructing airfoil polars for BEM-based rotor analyses.

A representative example of this approach which first generated a large dataset of airfoils, each annotated with aerodynamic performance metrics obtained from numerical simulations, with particular emphasis on predicting the lift-to-drag ratio across operating conditions [10]. Five supervised learning algorithms are systematically compared: Random Forest, Gradient Boosting Regression, Decision Tree Regressor, AdaBoost, and Linear Regression. The input features encode geometric information and flow parameters, while the output is typically C_L/C_D or related performance quantities at specified angles of attack and Reynolds numbers.

$$(\text{shape}, \text{operating condition}) \mapsto \left(C_L, C_D, \frac{C_L}{C_D}, \ldots\right) \tag{1}$$

Table 1 encompassed a multificated anysis that rigorously compared the different algorithms.The Coefficient of Determination (R^2) which serves as a barometer of predictive accuracy, and the MSE, a vital indicator of prediction precision, are key evaluation metrics in examination.The results show that ensemble methods such as Random Forest and Gradient Boosting deliver significantly lower prediction errors and more robust generalization than simple linear models, especially when the dataset is sufficiently large and diverse. The trained models are able to approximate the expensive high-fidelity solver with good accuracy, and can therefore be embedded in higher-level design workflows to rapidly evaluate candidate airfoils. For wind turbine applications, this means that instead of repeatedly calling CFD to populate polar tables for multiple candidate sections, a trained ML surrogate can be used to predict C_L and C_D curves on demand, enabling much faster exploration of families of airfoils tailored to different radial positions on the blade.

DOI: 10.54254/2755-2721/2026.MH31014

Table 1. Performance evaluation metrics of the ML algorithms

ML Algorithm	train/test ratio	R^2	MSE	Training time	Evaluation time
RF	0.2	0.944	0.0008	271.98s	0.23s
GBR	0.2	0.896	0.0014	197.30s	0.16s
DTR	0.2	0.857	0.0020	4.43s	0.10s
AdaBoost	0.7	0.747	0.0037	43.41s	17.35s
LR	0.3	0.832	0.0027	0.33s	0.01s

4.2. Surrogate-assisted optimization of airfoil shapes using multi-fidelity data

Given an accurate surrogate for C_L and C_D, the next logical step is to embed it in an optimization loop that searches directly over airfoil shape parameters. In practice, the main difficulty is that high-fidelity (HF) CFD simulations are expensive, while low-fidelity (LF) models are cheaper but less accurate. Recent work has therefore focused on multi-fidelity surrogate strategies, where LF and HF data are combined in a single model to achieve high accuracy at reduced cost.

A representative example is the multi-fidelity convolutional neural network with transfer learning (MFCNN-TL) proposed by Liao [11]. Figure 2 depicts the framework which the airfoil geometry is parameterized using B-spline control points, and aerodynamic performance is evaluated by two CFD solvers: a coarse-grid solver providing LF data and a fine-grid solver providing HF data. The key idea is to first train a convolutional neural network on a relatively large LF dataset, using airfoil "images" as input and aerodynamic coefficients (e.g. C_D) as output, and then to fine-tune the upper layers of this network using a much smaller HF dataset. From a machine-learning perspective, LF samples act as the source domain and HF samples as the target domain, so that the final MFCNN-TL surrogate leverages transfer learning to capture the nonlinear correlation between LF and HF responses.

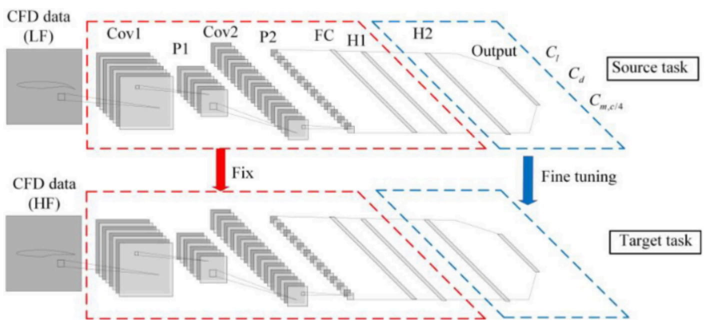

Figure 2. Architecture of the MFCNN-TL

This multi-fidelity surrogate is then used inside an iterative optimization loop. At each iteration, a particle swarm optimization (PSO) algorithm searches for improved airfoil shapes by querying the MFCNN-TL model instead of calling CFD directly. To keep the surrogate accurate in the regions explored by the optimizer, Liao et al. introduce a multi-fidelity infilling strategy: new candidate shapes selected by PSO are evaluated either with the LF or HF CFD solver, and the resulting samples are added to the training set to update the surrogate. In their study, the framework is applied to minimize the drag coefficient of the NACA0012 and RAE2822 airfoils, and the results show that

DOI: 10.54254/2755-2721/2026.MH31014

the multi-fidelity CNN–TL surrogate can reach comparable or better optima than a single-fidelity surrogate, while significantly reducing the number of HF simulations required.

For wind turbine airfoils, the same strategy is highly attractive. LF data can be generated, for example, from XFOIL or coarse-mesh RANS simulations across a grid of angles of attack and Reynolds numbers, while HF data come from fine-mesh CFD or detailed unsteady simulations. An MFCNN-TL-type surrogate can then be trained to map turbine-specific airfoil parameterizations to C_L , C_D , and derived performance metrics, and a global optimizer can operate on this surrogate. Because LF samples are cheap, they enrich the feature space and help the model learn general trends, whereas a limited number of HF samples correct systematic LF biases through transfer learning. In this way, multi-fidelity surrogate-assisted optimization enables designers to search richer shape spaces and account for more operating conditions than would be feasible with HF CFD alone.

The framework illustrates two important ideas that are directly relevant for wind-turbine aerodynamics: first, combining LF and HF information is an efficient way to train powerful surrogates for airfoil performance, and second, these surrogates can be tightly coupled with global optimizers and adaptive infilling strategies to build practical airfoil design tools that balance accuracy and computational cost. These concepts form a natural bridge to multi-fidelity 3D blade optimization, where the same principles can be applied to chord/twist distributions and spanwise airfoil families.

5. Conclusion

This review examined how machine learning can support and accelerate the aerodynamic optimization of wind turbine blades, with a particular focus on wind-turbine airfoils as the basic building blocks of three-dimensional rotors. Starting from the fundamentals of blade aerodynamics and conventional design workflows, we recalled that current practice still depends heavily on BEM-based rotor analysis and high-fidelity CFD, coupled with gradient-based or metaheuristic optimization. While these tools are mature and reliable, their computational cost limits the size of the design space, the number of operating conditions, and the degree of multidisciplinary coupling that can be considered in routine design.

Against this background, we reviewed recent work in which supervised and deep-learning surrogate models are trained to predict lift and drag coefficients and their ratio as functions of airfoil geometry and operating conditions. These surrogates can reproduce high-fidelity polars at a fraction of the cost, making it feasible to explore large libraries of candidate airfoils and to populate sectional databases for blade design. Building on this, we discussed surrogate-assisted airfoil shape optimization, highlighting in particular multi-fidelity convolutional neural network surrogates trained via transfer learning that combine abundant low-fidelity data with limited high-fidelity simulations. Such models balance accuracy and efficiency and are directly suitable for integration into higher-level blade optimization frameworks.

Overall, the evidence indicates that machine learning already provides practical benefits when used as a surrogate and optimization engine within physics-based workflows. Future work should focus on improving data quality and coverage, enforcing physical consistency, quantifying uncertainty, and more tightly coupling aerodynamic surrogates with structural, aeroelastic and control models to enable robust, machine-learning-enhanced design of next-generation wind turbine blades.

DOI: 10.54254/2755-2721/2026.MH31014

References

[1] El Mouhsine, S., Oukassou, K., Ichenial, M. M., Kharbouch, B., & Hajraoui, A. (2018). Aerodynamics and structural analysis of wind turbine blade. Procedia Manufacturing, 22, 747-756.Bootorabi, F., Haapasalo, J., Smith, E., Haapasalo, H. and Parkkila, S. (2011) Carbonic Anhydrase VII—A Potential Prognostic Marker in Gliomas. Health, 3, 6-12.
[2] Yang, K. (2020). Geometry design optimization of a wind turbine blade considering effects on aerodynamic performance by linearization. Energies, 13(9), 2320.
[3] Jasa, J., Bortolotti, P., Zalkind, D., & Barter, G. (2021). Effectively using multifidelity optimization for wind turbine design. Wind Energy Science Discussions, 2021, 1-22.
[4] Baisthakur, S., & Fitzgerald, B. (2024). Physics-Informed Neural Network surrogate model for bypassing Blade Element Momentum theory in wind turbine aerodynamic load estimation. Renewable Energy, 224, 120122.
[5] Yang, S., Lee, S., & Yee, K. (2023). Inverse design optimization framework via a two-step deep learning approach: application to a wind turbine airfoil. Engineering with Computers, 39(3), 2239-2255.
[6] Sierra-García, J. E., & Santos, M. (2020). Exploring reward strategies for wind turbine pitch control by reinforcement learning. Applied Sciences, 10(21), 7462.
[7] Hansen, M. (2015). Aerodynamics of wind turbines. Routledge.
[8] Cardamone, R., Broglia, R., Papi, F., Rispoli, F., Corsini, A., Bianchini, A., & Castorrini, A. (2025). Aerodynamic design of wind turbine blades using multi-fidelity analysis and surrogate models. International Journal of Turbomachinery, Propulsion and Power, 10(3), 16.
[9] Soler, D., Mariño, O., Huergo, D., de Frutos, M., & Ferrer, E. (2024). Reinforcement learning to maximize wind turbine energy generation. Expert Systems with Applications, 249, 123502.
[10] Teimourian, A., Rohacs, D., Dimililer, K., Teimourian, H., Yildiz, M., & Kale, U. (2024). Airfoil aerodynamic performance prediction using machine learning and surrogate modeling. Heliyon, 10(8).
[11] Bravo-Mosquera, P. D., Ceron-Munoz, H. D., & Catalano, F. M. (2022). Design, aerodynamic analysis and optimization of a next-generation commercial airliner. Journal of the Brazilian Society of Mechanical Sciences and Engineering, 44(12), 609.